Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Mapping the Milky Way's Rotation Curve from Archival 21-cm HI Observations

Shibam Ghosh

technoshibam@gmail.com

Univeristy of Petroleum and Energy Studies https://orcid.org/0009-0003-4780-5324

Joshal Bhokare

Univeristy of Petroleum and Energy Studies

Research Article

Keywords: Milky Way, neutral hydrogen, 21-cm line, HI4PI, Galactic rotation, rotation curve

Posted Date: November 4th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11312879/v2

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: The authors declare no competing interests.

Mapping the Milky Way's Rotation Curve from Archival 21-cm H I Observations

Shibam Ghosh Joshal Bhokare

Computational analysis of HI4PI Galactic neutral-hydrogen spectra using the tangent-point method

Abstract

The 21-cm hyperfine line of neutral atomic hydrogen (H I) provides a direct spectroscopic probe of Galactic gas because radio emission penetrates interstellar dust and Doppler shifts encode line-of-sight motion. This work uses archival HI4PI data to construct a Galactic longitude–velocity (l–v) diagram, inspect representative spectra, identify terminal-velocity edges, and transform those measurements into an inner-Milky-Way rotation curve with the tangent-point method. In the supplied run, the displayed l–v product spans approximately $10^\circ \leq l \leq 85^\circ$ and averages a broader latitude strip $|b| \leq 2^\circ$. Terminal velocities are measured with an absolute brightness-temperature threshold of 20 K and a persistence requirement across adjacent channels. The real diagnostic figures show that this configuration is not uniformly successful: representative sightlines near $l = 69.9^\circ$ and 84.9° return no valid terminal velocity, and the intended high-longitude calibration interval begins beyond the available accepted terminal-velocity sequence. Consequently the displayed edge-bias correction is 0.0 $\mathrm{km\,s^{-1}}$ rather than a measured calibration. The derived curve extends from roughly 1.5 to 7.4 kpc and rises from about 95 $\mathrm{km\,s^{-1}}$ in the innermost sampled region to approximately 230–255 $\mathrm{km\,s^{-1}}$ over much of the outer part of the measured interval. The exercise therefore demonstrates both the astrophysical power of 21-cm kinematics and the importance of quality control: a plausible-looking rotation curve must be interpreted together with threshold diagnostics, missing detections, geometric validity, and calibration coverage. The analysis is presented as a reproducible undergraduate investigation rather than a precision Galactic mass model.

Keywords: Milky Way; neutral hydrogen; 21-cm line; HI4PI; Galactic rotation; rotation curve.

1 Introduction

1.1 Motivation: mapping a Galaxy from inside it

The Milky Way cannot be photographed externally, so its three-dimensional structure and gravitational field must be reconstructed from angular positions, distances, spectra, and motions. Radio spectroscopy is especially valuable because centimetre-wavelength radiation penetrates dust that obscures the optical Galactic plane. Neutral atomic hydrogen (H I) is widespread through the disk and its 21-cm emission therefore provides a nearly all-sky kinematic tracer. Modern surveys such as HI4PI combine wide angular coverage with fine velocity sampling, making it possible to revisit classical Galactic rotation experiments using archival data rather than a dedicated telescope campaign (HI4PI Collaboration et al., 2016).

1.2 Scientific questions addressed

This project asks four linked questions. First, what structure is visible in a longitude–velocity diagram of Galactic H I? Second, how can an operational terminal velocity be extracted from a finite-width, multi-component spectrum? Third, how does tangent-point geometry convert that terminal velocity into Galactocentric radius and circular speed? Fourth, what do the resulting rotation and enclosed-mass curves imply, and where do the assumptions of the method become limiting? These questions connect atomic physics, spectroscopy, coordinate geometry, numerical analysis, and Galactic dynamics in one reproducible workflow.

1.3 Scope of the present run

The paper documents the supplied parameter configuration rather than replacing it with an idealized result. The displayed data use a broad latitude average, $|b| \leq 2^\circ$, and a 20 K absolute edge threshold. These settings produce a useful inner-Galaxy sequence but fail to return terminal velocities for some high-longitude spectra. That failure is retained and analysed because a scientifically useful pipeline must expose its limitations as clearly as its successful detections. The Milky Way is observed from within its disk, so its global structure and dynamics must be inferred indirectly. Optical tracers become difficult to use through heavily extinguished regions of the Galactic plane, whereas the 21-cm line of neutral atomic hydrogen can be detected through dust. It therefore supplies two complementary observables: the distribution of atomic gas on the sky and the Doppler velocity of that gas along the line of sight.

The H I line arises from the hyperfine splitting of the hydrogen ground state. Its rest frequency is approximately 1420.4058 MHz, corresponding to a wavelength near 21.1 cm. Although the spontaneous transition probability of an individual atom is very small, the enormous abundance of interstellar H I makes the integrated Galactic signal strong. Modern surveys consequently provide densely sampled spectral cubes in Galactic longitude, latitude and velocity. Differential rotation imprints a characteristic envelope on these cubes. For sightlines through the inner Galaxy, the extremal radial velocity can be associated with a tangent point under the assumptions of approximately circular motion and axisymmetry. Repeating the measure-

ment across longitude converts the spectral envelope into a circular-speed curve. The present project follows that chain explicitly, with special attention to the fact that the terminal velocity is an operationally defined edge rather than a directly tabulated observable.

2 Scientific Background

2.1 The 21-cm hyperfine transition

The H I line originates from the hyperfine splitting of the electronic ground state caused by the interaction between the magnetic moments of the proton and electron. The transition frequency is approximately $\nu_0 = 1420.4058$ MHz, corresponding to $\lambda \simeq 21.106$ cm. The spontaneous transition is intrinsically weak for an individual atom, but interstellar columns contain so many neutral hydrogen atoms that Galactic emission is readily detected. The observable in a survey cube is commonly expressed as brightness temperature T_B, a radio-astronomical representation of specific intensity through the Rayleigh–Jeans relation.

2.2 Doppler velocity and the Local Standard of Rest

For non-relativistic Galactic gas, a small spectral displacement is interpreted as a radial velocity. In first-order form,

$$v_{\rm rad} \simeq c\,\frac{\nu_0 - \nu}{\nu_0}, \tag{1}$$

with the sign convention depending on the adopted spectral definition. Survey products are normally provided in a standard velocity frame such as the Local Standard of Rest (LSR), removing a conventional component of the Sun's peculiar motion. The resulting $v_{\rm LSR}$ is not the full three-dimensional gas velocity: it is only the line-of-sight projection relative to the adopted reference frame.

2.3 Galactic differential rotation

Gas on approximately circular orbits has an angular speed that depends on Galactocentric radius. Because the Sun also orbits the Galactic centre, a cloud's observed radial velocity is a difference of projected orbital motions. In the inner Galaxy, a single sightline passes through a range of Galactocentric radii. The smallest radius reached along that sightline is the tangent point, where the line of sight is tangent to a circular orbit. Under axisymmetry and circular motion, the most extreme allowed radial velocity is associated with gas near this point.

2.4 Historical context for terminal-velocity rotation curves

Terminal-velocity methods have a long history in H I and CO studies. CO tangent-point measurements were used to derive detailed inner-disk rotation curves (Clemens, 1985), while later H I work used dense terminal-velocity sampling to quantify the rotation curve and streaming motions (McClure-Griffiths & Dickey, 2007). The present analysis follows the same physical logic but emphasizes transparent diagnostics and reproducible archival processing. A radio spectral cube can be represented as brightness temperature $T_B(l, b, v)$, where l and b are Galactic longitude and latitude and v is velocity with respect to the Local Standard of Rest (LSR). Averaging over a narrow latitude interval produces the two-dimensional diagnostic $T_B(l, v)$. Bright emission near $v \simeq 0$ often contains local gas, while extended wings trace gas with substantial projected rotational motion. The HI4PI survey combines EBHIS and GASS into a full-sky H I data set. The published survey has an angular resolution of about $16.2'$ and an rms sensitivity near 43 mK. Those characteristics are more than sufficient for the high-brightness edge criteria used in this student analysis. The threshold is therefore not set by thermal receiver noise alone; it is chiefly a methodological choice about where a finite-width astrophysical emission profile is declared to end.

3 Data Source and Processing

3.1 HI4PI survey characteristics

HI4PI is an all-sky Galactic H I survey formed from EBHIS in the northern sky and GASS in the southern sky. The survey publication reports an angular resolution of about $16.2'$ and an rms sensitivity of about 43 mK, improving substantially on the earlier LAB all-sky compilation (HI4PI Collaboration et al., 2016; Kalberla et al., 2005). The public CDS/VizieR products make the survey particularly suitable for reproducible student projects because the same FITS cubes can be downloaded and re-analysed independently (CDS/VizieR, 2016).

3.2 FITS cubes and coordinate metadata

Each spectral cube contains brightness temperature sampled along two celestial axes and one spectral axis. The pipeline reads the FITS header, reconstructs the spectral coordinate from the reference pixel, reference value and channel increment, and uses Astropy WCS for sky coordinates. Explicit unit checks are important because a velocity axis expressed in $\mathrm{m\,s^{-1}}$ but interpreted as $\mathrm{km\,s^{-1}}$ would introduce a factor-of-one-thousand error.

3.3 Spatial selection and latitude averaging

The requested longitude interval is extracted from the relevant overlapping cubes. Pixels satisfying $|b| \leq 2^\circ$ are averaged to produce one spectrum per longitude. Latitude averaging increases the effective amount of emission contributing to each sightline and can suppress pixel-scale fluctuations, but it also mixes gas at different heights and can alter the terminal wing. Consequently the latitude width is an analysis parameter, not merely a plotting choice.

3.4 Cube stitching and duplicate removal

Neighbouring survey cubes overlap on the sky. The program concatenates the extracted pieces, sorts them in longitude, and removes duplicate sightlines within a small coordinate tolerance. When spectral grids differ, spectra are interpolated onto a master velocity grid before concatenation. These bookkeeping operations are scientifically important because a duplicated or misregistered sightline can create artificial discontinuities in an l–v map.

3.5 Reproducible data provenance

The pipeline records downloaded file sizes and SHA-256 checksums, stores derived arrays and CSV tables, and writes a run record containing adopted constants and configuration values. This makes the analysis auditable: a later rerun can establish whether a difference comes from changed code, changed parameters, or changed input data rather than from an undocumented manual step. The computational workflow uses Galactic Cartesian FITS cubes from the CDS archive associated with HI4PI. The program reconstructs velocity coordinates from FITS header metadata, obtains sky coordinates through Astropy WCS, selects the requested longitude and latitude region, and averages the retained latitude pixels. Overlapping cube sightlines are sorted and de-duplicated before construction of the final l–v array. The present figures correspond to a different parameter configuration from the earlier $|b| \leq 0.5^\circ$, 10 K run. Here the displayed map averages $|b| \leq 2^\circ$, and the terminal-edge diagnostic uses 20 K. This distinction is scientifically important: broadening the latitude strip mixes emission from a larger vertical region of the Galactic disk, while increasing the edge threshold tends to move the detected boundary inward toward brighter emission. The figures in this paper are therefore interpreted as the actual outcome of this run, not as interchangeable illustrations from another configuration.

4 Longitude–Velocity Structure

4.1 Construction of the l–v diagram

After latitude averaging, the cube is reduced to $T_B(l, v)$. A longitude–velocity diagram is therefore not a conventional sky image: horizontal position is Galactic longitude, vertical position is radial velocity, and pixel intensity is H I brightness temperature (Figure 1). This representation compresses a large fraction of the kinematic information in the inner disk into a single diagnostic plane.

4.2 Interpreting coherent emission ridges

Bright ridges arise because many sightlines intersect substantial columns of gas at related velocities. Spiral-arm streaming, local gas, shells, and other non-circular structures can produce substructure. The terminal envelope is the high-absolute-velocity boundary of significant emission, not necessarily the location of maximum brightness. This distinction motivates direct inspection of individual spectra in addition to the two-dimensional map.

4.3 Effect of the broad latitude strip

A $|b| \leq 2^\circ$ average includes more off-plane emission than a narrower extraction. It can make extended components more visible and may reduce stochastic variation, but a faint high-velocity wing confined close to the plane can be diluted. Thus the same absolute 20 K threshold can behave differently when the latitude averaging width changes.

The emission structure in the diagram changes systematically with longitude. At intermediate longitudes the diagram contains bright structured emission extending to velocities of order 100 $\mathrm{km\,s^{-1}}$ or more; toward the high-longitude end, the dominant emission contracts toward lower velocity. A negative-velocity component is also visible below the main ridge and reflects additional Galactic gas along the selected sightlines. A broad latitude average increases the amount of gas contributing to each longitude bin. That can improve visual continuity but also changes the spectral shape used for edge detection. For this reason, the l–v diagram should always be inspected before accepting an automated terminal-velocity table. The correct diagnostic question is not whether the extracted points form a smooth curve, but whether they actually trace the outer boundary of real emission in the underlying data.

The diagram can also be understood as a compressed summary of hundreds of individual spectra. A vertical cut at one longitude gives a single brightness-temperature spectrum; moving horizontally compares neighbouring sightlines. The fact that the high-positive-velocity boundary varies coherently over many adjacent longitudes is therefore a powerful internal check that it is Galactic structure rather than isolated noise. An extraction algorithm should preserve that coherence while still allowing local departures produced by real gas dynamics. The broad $|b| \leq 2^\circ$ average used here deserves explicit attention. Near the Galactic plane, different latitude rows can contain somewhat different mixtures of clouds and velocity components. Averaging improves the signal representation of widespread H I but can broaden or reshape an edge. Thus the latitude width is part of the measurement definition, just like the brightness threshold. A direct comparison with a narrower $|b| \leq 0.5^\circ$ extraction would quantify how much of the terminal-velocity change comes from vertical averaging rather than from the edge threshold itself.

Another useful check is the empty velocity space above the envelope. The plot extends to velocities far larger than the detected positive wing, so the middle-longitude terminal edges are not being clipped by the displayed velocity ceiling. This matters because an insufficient search window can create a false edge exactly at its upper boundary. In the present figure the main limitation is instead the brightness criterion at high longitude, where the outer wing becomes too weak to satisfy 20 K.

5 Spectral Morphology

The representative spectra in Figure 2 show that a Galactic H I profile is rarely a single Gaussian line. Several components can overlap, and the location of the highest peak can differ greatly from the most extreme velocity at which significant emission persists. For example, the $l \approx 25^\circ$ and 40° profiles contain substantial positive-velocity wings, whereas the high-longitude spectra concentrate much closer to zero velocity. This morphology motivates an edge-based method. A fixed absolute threshold has the advantage that the physical brightness level is the same from sightline to sightline. It also has a limitation: if the high-velocity wing never reaches that level, the algorithm returns no detection even though weaker emission may still be present. The supplied 20 K run demonstrates exactly this issue at high longitude.

The six profiles in Figure 4 also illustrate the changing relationship between peak brightness and terminal velocity. A very bright narrow component near zero velocity can dom-

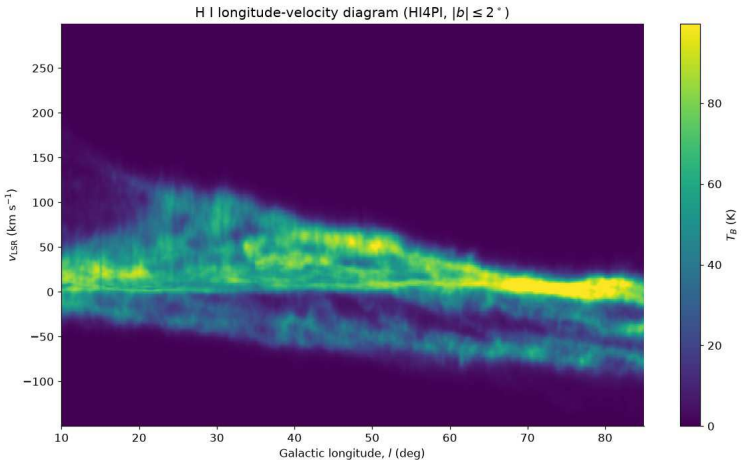


Figure 1. H I longitude–velocity diagram from the supplied HI4PI run, averaged over $|b| \leq 2^\circ$. The bright emission ridge and positive-velocity envelope provide the observational basis for the first-quadrant terminal-velocity analysis.

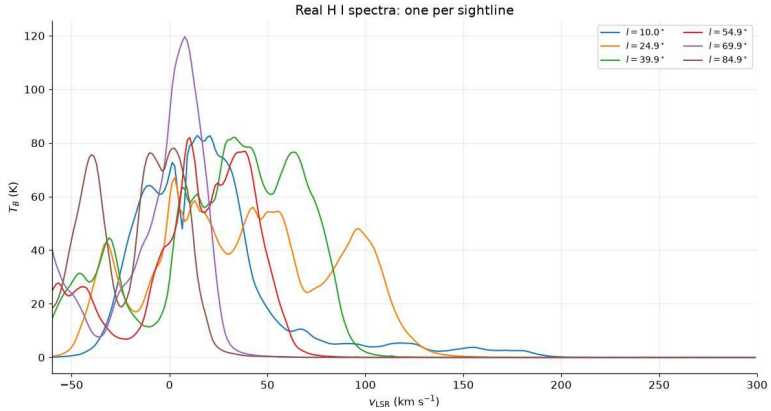


Figure 2. Representative H I spectra for six sightlines from $l = 10.0^\circ$ to 84.9°. Multiple peaks, shoulders and broad wings show why the terminal edge, rather than the brightest spectral peak, is used as the kinematic estimator.

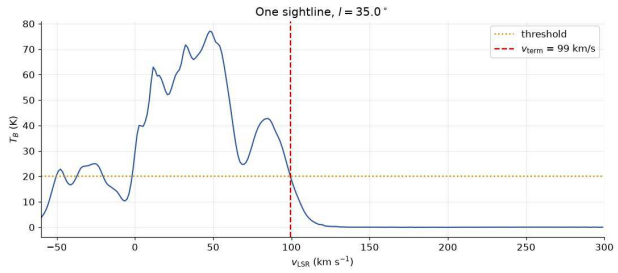


Figure 3. Single-sightline edge measurement at $l = 35.0^\circ$. The dotted line marks the 20 K threshold and the dashed line marks the measured terminal velocity, approximately 99 km s^{-1}.

inate the visual appearance of a spectrum without carrying the extremal rotational information. At another longitude, a moderate-brightness wing can extend far beyond the main peak and determine $v_{\rm term}$. This is why an algorithm based on the velocity of maximum T_B would not reproduce the tangent-point method.

Spectral complexity has a physical origin. A single line of sight intersects gas at many distances, and those regions need not share exactly the same circular speed or peculiar velocity. Spiral-arm streaming, cloud-to-cloud motions, turbulence and expanding structures can create shoulders or subsidiary peaks. Latitude averaging adds still more components because neighbouring sightlines are combined. The resulting spectrum is therefore a superposition of Galactic structures rather than a single laboratory line profile. For an undergraduate analysis, these profiles are valuable because they make the measurement procedure auditable. One can inspect the spectrum, identify the chosen threshold, locate the last persistent crossing and compare the automated answer with the visible wing. This is much stronger than reporting only a final table of terminal velocities. It also provides a direct way to diagnose why a sightline becomes NaN: the reader can see whether the 20 K line never intersects a suitable outer wing or whether the profile becomes ambiguous.

6 Terminal-Velocity Criterion

For each spectrum the analysis searches the appropriate positive-velocity side for contiguous emission above the selected threshold. In the example at $l = 35^\circ$ (Figure 3), the spectrum remains above 20 K until approximately 99 km s^{-1}, after which it falls rapidly toward the baseline. The terminal velocity is therefore associated with the threshold crossing on the outer wing rather than with the principal emission maximum near 50 km s^{-1}. The distinction is phys-

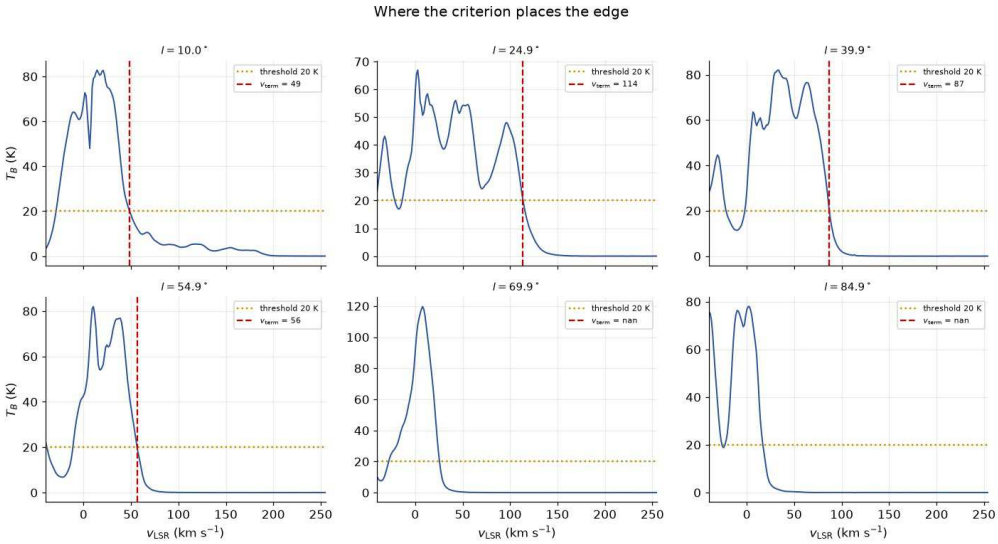


Figure 4. Threshold diagnostic for six representative sightlines. Valid 20 K terminal velocities are found at $l = 10.0°$, $24.9°$, $39.9°$ and $54.9°$, while the displayed $l = 69.9°$ and $84.9°$ cases return no valid value (NaN).

ically essential. A peak identifies where a large amount of gas emits at one velocity; the terminal edge identifies the largest projected rotational velocity represented by sufficiently bright gas along the sightline. In the tangent-point interpretation, it is the latter quantity that is related to the circular speed at the minimum Galactocentric radius reached by that line of sight.

6.1 Persistence and false detections

A useful edge algorithm should reject isolated channels that cross the threshold because of narrow features or noise. Requiring a minimum contiguous run provides a simple safeguard. The procedure remains sensitive, however, to real astrophysical profile complexity. Shells, streaming motions, spiral-arm features and blended clouds can all produce emission at velocities that do not perfectly follow an axisymmetric circular model.

7 Failure Modes Revealed by the Real Run

For the $l = 69.9°$ and $84.9°$ sightlines in Figure 4, the plotted legend reports $v_{\rm term} = \rm nan$. In other words, the selected 20 K rule does not yield a qualifying terminal edge for those examples. This is not a plotting defect; it is a result of the chosen criterion applied to the actual spectra. A scientifically transparent workflow retains such failures rather than silently filling them by interpolation. Missing detections can arise when the outer wing lies below the absolute threshold, when the search window and emission geometry do not overlap appropriately, or when the persistence rule rejects a short feature. Their presence argues for a parameter-sensitivity study, for example comparing 20 K with 10 K or 5 K while independently checking the l–v envelope. The low-longitude $l = 10°$ example also deserves caution. The tangent-point method becomes increasingly vulnerable to the Galactic bar and other non-circular motions toward the Galactic centre. Including very low longitudes can extend the nominal radial range inward, but the resulting circular-speed interpretation becomes less secure.

8 Quality Control on the Emission Envelope

The overlay in Figure 5 provides the most direct check on whether the algorithm measures what it claims to measure. Accepted points should hug the boundary between detectable emission and the dark background. Across much of the middle-longitude range the white trace follows the high-positive-velocity envelope convincingly. A discontinuity near the low-20-degree region and the termination of accepted points well before 90° are equally informative. These features show why a smooth final rotation curve cannot be the only quality criterion. Binning and interpolation can visually suppress missing or problematic sightlines. The underlying envelope plot retains the connection to the observed data and should therefore accompany any derived kinematic curve in a research presentation. A useful way to read this diagram is to follow a vertical slice at fixed longitude. Each slice is one H I spectrum, and every bright patch in the image corresponds to a velocity channel that contributes to that spectrum. The coherent upper boundary across neighbouring longitudes is therefore not a fitted curve imposed on the data; it emerges from many independent spectral profiles. This continuity is what makes terminal-velocity work possible. If an apparent edge jumps abruptly from one longitude to the next while the underlying emission does not, the extraction rule is probably following a secondary feature rather than the true envelope.

The map also shows why the velocity range must be wider than the expected terminal velocities. Baseline channels well beyond the Galactic emission provide a direct visual reference for what “no emission” looks like. Conversely, restricting the velocity axis too tightly can clip the wing and create an artificial terminal edge at the search boundary. The displayed range leaves substantial empty velocity space above the positive envelope, so the principal middle-longitude edges are not simply imposed by the plotting limit. Finally, the brightness scale should not be confused with mass density. T_B is the observed radio brightness temperature in a velocity channel. Converting it to H I column density requires integrating over velocity and usually in-

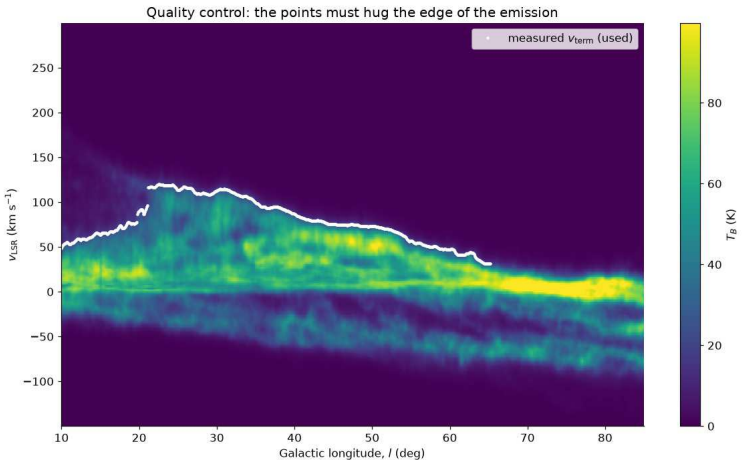


Figure 5. Quality-control overlay of accepted terminal velocities on the l–v diagram. The white trace follows much of the outer positive-velocity envelope, while gaps and changes in behaviour identify regions requiring additional scrutiny.

vokes an optically thin approximation. The present project does not need that conversion for the rotation measurement: it uses the location of the spectral boundary in velocity space. This keeps the kinematic inference conceptually separate from a gas-mass estimate.

9 Tangent-Point Geometry

9.1 Geometry inside the Solar circle

For first-quadrant longitudes $0^\circ < l < 90^\circ$, the tangent point lies at

$$R_t = R_0 \sin l, \tag{2}$$

where R_0 is the Sun's Galactocentric radius. The distance from the Sun to the tangent point is $d_t = R_0 \cos l$ in the idealized plane geometry. The important advantage is that R_t follows from geometry alone; no independent distance to a particular H I cloud is required.

9.2 Conversion from terminal velocity to circular speed

For circular motion in the first quadrant, the terminal velocity can be written

$$v_{\rm term} = \Theta(R_t) - \Theta_0 \sin l, \tag{3}$$

so that

$$\Theta(R_t) = v_{\rm term} + \Theta_0 \sin l. \tag{4}$$

The adopted constants $R_0 = 8.15 \pm 0.15$ kpc and $\Theta_0 = 236 \pm 7$ km s^{-1} are consistent with modern maser-parallax modelling (Reid et al., 2019). They set the radial scale and velocity zero point of the derived curve.

9.3 Domain of validity

The method does not apply uniformly to all Galactic longitudes. Between 90° and 270°, inner-Solar-circle tangent points are not sampled in the same way, and outer-Galaxy rotation requires different distance information. At very small longitudes, bar-driven and other non-circular motions become increasingly important. The useful domain is therefore deliberately restricted. For a sightline in the first Galactic

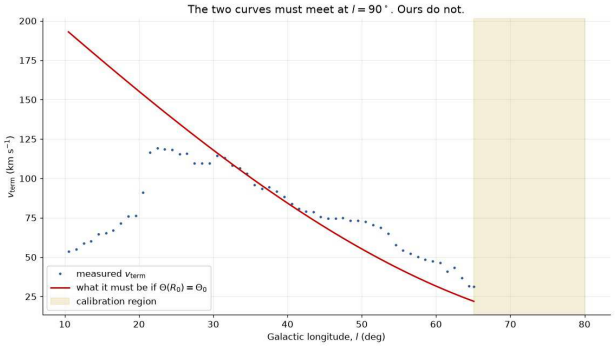


Figure 6. Comparison of measured terminal velocities with the ideal high-longitude behaviour. The shaded calibration region begins near 65°, but the accepted measured sequence ends at about the same longitude, so this run does not provide the intended high-longitude calibration sample.

quadrant and an idealized circular velocity field, the tangent point lies at

$$R = R_0 \sin l, \tag{5}$$

where R_0 is the Sun's Galactocentric radius. The circular speed is obtained from the terminal velocity through

$$\Theta(R) = v_{\rm term,corr} + \Theta_0 \sin l, \tag{6}$$

with Θ_0 the adopted circular speed at the Solar radius. The analysis uses $R_0 = 8.15$ kpc and $\Theta_0 = 236$ km s^{-1} as external calibration parameters. They are not measured by the H I extraction itself.

The geometry is strongest for inner-Galaxy sightlines that avoid the central non-axisymmetric region and approach, but do not require, the exact $l = 90^\circ$ limit. At $l = 90^\circ$, $R = R_0$ and the ideal terminal velocity tends to zero. This boundary condition can be used to diagnose systematic edge offsets only when valid high-longitude measurements actually reach the calibration interval.

10 Edge-Bias Calibration: Why It Fails Here

An earlier configuration could estimate an additive edge bias from accepted sightlines near the Solar-circle limit.

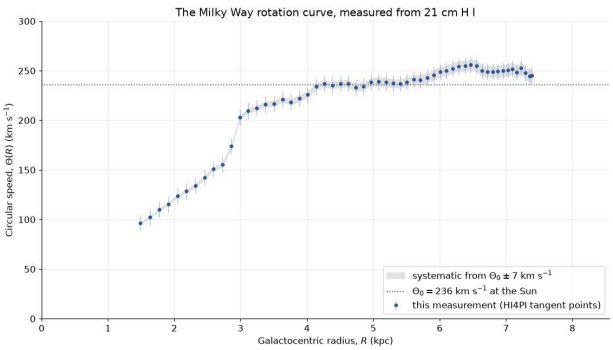


Figure 7. Rotation curve obtained from the accepted tangent-point measurements in the supplied 20 K, $|b| \leq 2^\circ$ run. The curve rises strongly at small radius and reaches approximately 230–255 $\mathrm{km\,s^{-1}}$ over much of the outer measured interval.

The present real run is different. Figure 6 shows the calibration region beginning around $l = 65^\circ$, while the measured terminal-velocity sequence ends at approximately that longitude. Thus there are not enough accepted high-longitude points inside the calibration interval to determine a robust median excess. This explains the later diagnostic label "edge bias of 0.0 km/s removed." Zero here should not be interpreted as evidence that the 20 K edge is unbiased. It indicates that no empirical correction was successfully established for this configuration. This distinction materially affects the scientific interpretation of the derived rotation curve, particularly near its outer radial limit. A better follow-up run would either lower the absolute threshold, adopt a carefully validated fractional criterion, or extend the detection strategy so that valid terminal edges are recovered at sufficiently high longitude. Any revised criterion should then be checked against the l–v envelope rather than selected solely because it makes the final curve approach Θ_0. The strong rise at small radius and the flattening of the curve near Θ_0 (Figure 7) therefore reflect the absence of a usable calibration estimate, not proof that the raw edge is unbiased.

11 Rotation-Curve Results

The displayed rotation curve spans roughly 1.5–7.4 kpc, shorter than the Solar radius because valid terminal velocities do not extend to $l \simeq 90^\circ$. The innermost points rise from about 95 $\mathrm{km\,s^{-1}}$ at $R \simeq 1.5$ kpc to roughly 200 $\mathrm{km\,s^{-1}}$ by $R \simeq 3$ kpc. Beyond about 4 kpc, many points lie in the approximate 230–255 $\mathrm{km\,s^{-1}}$ range. The rapid inner rise should not automatically be interpreted as a precision measurement of circular speed. It is produced by sightlines at small longitude, precisely where the bar and non-circular gas flows are most important. A conservative scientific use of this run would emphasize the intermediate-radius behaviour and treat the smallest-radius points as method-sensitive. The horizontal $\Theta_0 = 236$ $\mathrm{km\,s^{-1}}$ reference provides context but does not validate the curve. Because the high-longitude edge calibration failed, agreement or disagreement with that reference near the largest measured radius is not a closed consistency test.

12 Systematic Interpretation of the Zero Correction

In Figure 8, a genuine additive correction would leave the two point sets with the same shape but differ by a constant vertical offset. Here they are identical. The legend explicitly reports a removed edge bias of 0.0 $\mathrm{km\,s^{-1}}$. Combined with Figure 6, the correct interpretation is that the calibration could not be established from the available high-longitude accepted points. This is preferable to inventing a correction. A reproducible analysis should return "no calibration" when its calibration sample is inadequate. For publication-quality work, the run record should distinguish that state from a genuinely measured bias consistent with zero, because the two cases have different uncertainty implications.

13 Enclosed-Mass Diagnostic

For a spherical-equivalent interpretation, the circular speed can be converted into an enclosed dynamical mass,

$$M(<R) = \frac{\Theta^2(R)\,R}{G}, \tag{7}$$

where the computational pipeline uses $G = 4.30091 \times 10^{-6}$ $\mathrm{kpc\,(km\,s^{-1})^2\,M_\odot^{-1}}$. This relation is a useful dynamical diagnostic, although the Milky Way is a disk galaxy and the resulting $M(<R)$ should not be treated as a detailed decomposition of bulge, disk, gas and halo.

14 Mass Growth and Non-Keplerian Rotation

The enclosed-mass sequence (Figure 9) rises from a few 10^9 $M_\odot$ in the innermost plotted region to approximately 10^{11} $M_\odot$ near 7 kpc. Over the intermediate and outer measured range, the continuing mass increase is qualitatively consistent with a rotation curve that remains high rather than falling as $R^{-1/2}$. The comparison line $M \propto R$ represents the scaling expected for an exactly flat rotation curve under the same spherical-equivalent approximation. The measured points do not lie exactly on this line because the inferred circular speed is not exactly constant. In particular, the steep rise in the inner rotation curve produces corresponding curvature in the mass estimate.

15 Comparison with a Point-Mass Model

If essentially all relevant mass were concentrated inside the innermost anchor radius, circular speed outside that radius would decrease as $\Theta \propto R^{-1/2}$. Figure 10 shows that the supplied H I-derived points instead rise and then remain near a few hundred $\mathrm{km\,s^{-1}}$. The shaded separation is therefore a visual representation of the additional gravitating mass required relative to that deliberately simplified point-mass model. The comparison is pedagogical rather than a complete dark-matter measurement. The innermost anchor is itself located in a region where non-circular motions are important, and the Galaxy's baryonic mass is spatially extended rather than point-like. Nevertheless, the contrast demonstrates why Galactic rotation curves motivate extended mass models.

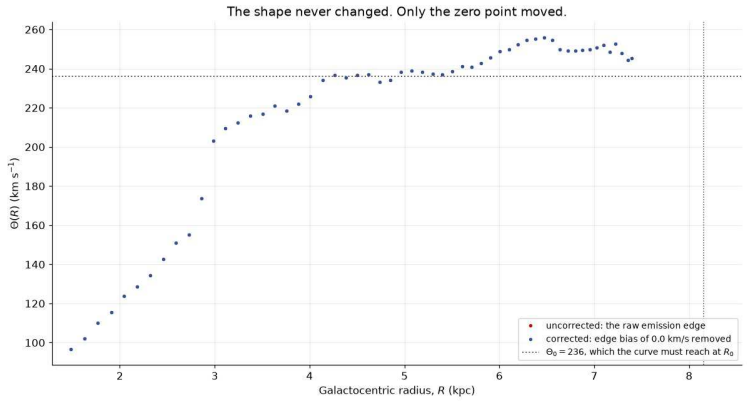


Figure 8. Before-and-after edge-bias diagnostic. The two point sets coincide because the run applied a correction of 0.0 $\mathrm{km\,s^{-1}}$. This reflects the absence of a usable calibration estimate, not proof that the raw edge is unbiased.

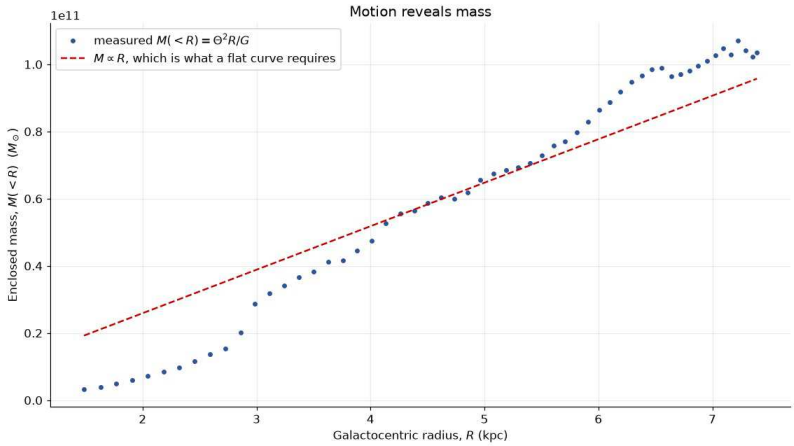


Figure 9. Spherical-equivalent enclosed mass derived from the displayed rotation curve. The inferred mass increases with Galactocentric radius and reaches of order $10^{11}\ M_{\odot}$ by the outer end of the measured interval.

16 Connection to Published Outer-Galaxy Work

The tangent-point method is intrinsically limited to the inner Galaxy. To show the broader context, Figure 11 overlays the linear circular-speed fit reported by Eilers et al. (2019), based on stellar kinematics from roughly 5 to 25 kpc. Their fit declines only gently with radius, whereas a Keplerian curve normalized at the Solar radius falls much more rapidly. The overlap between the H I points and the external fit should be interpreted cautiously in this run because the H I edge calibration is incomplete. The external curve is not used to repair the missing calibration; it is shown as an independent literature comparison. This separation avoids circular reasoning in which a threshold or correction is tuned merely to force agreement with a published answer.

17 Uncertainties and Limitations

17.1 Instrumental noise versus analysis systematics

The HI4PI rms noise is far below the 20 K threshold, so thermal noise is not the dominant uncertainty in the displayed edge. The larger uncertainties arise from the astrophysical width of the terminal wing, the adopted threshold, latitude averaging, finite velocity channels, and non-circular gas motions. This is a useful lesson: high signal-to-noise data can still yield a systematically uncertain derived quantity.

17.2 Uncertainty in R_0 and Θ_0

An uncertainty in R_0 changes every tangent-point radius coherently, while uncertainty in Θ_0 changes the circular-speed zero point with a longitude-dependent projection. These are correlated systematic effects rather than independent random errors on each plotted point. They should therefore be distinguished from local scatter caused by gas kinematics.

17.3 Failure of high-longitude self-calibration

The intended calibration compares measured high-longitude terminal velocities with the geometric requirement that the terminal velocity approach zero as l approaches 90°. In the supplied run, the 20 K threshold produces missing detections before a robust calibration sample is established. The displayed 0.0 $\mathrm{km\,s^{-1}}$ correction must therefore be read as "no correction measured," not as evidence that the threshold is unbiased.

17.4 Axisymmetry and the Galactic bar

The tangent-point equations assume circular orbits in an axisymmetric potential. The Milky Way contains a bar, spiral

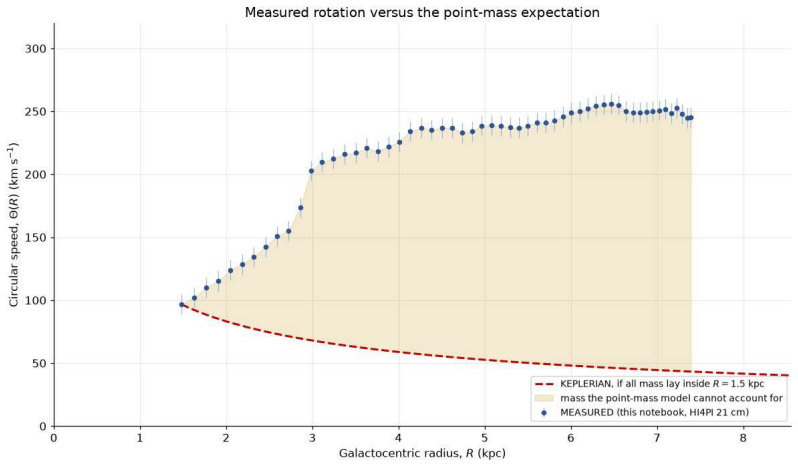


Figure 10. Measured circular speed compared with a Keplerian point-mass expectation anchored at the innermost plotted radius. The growing separation illustrates how strongly the measured intermediate-radius motion differs from a model in which all gravitating mass were already enclosed at the anchor radius.

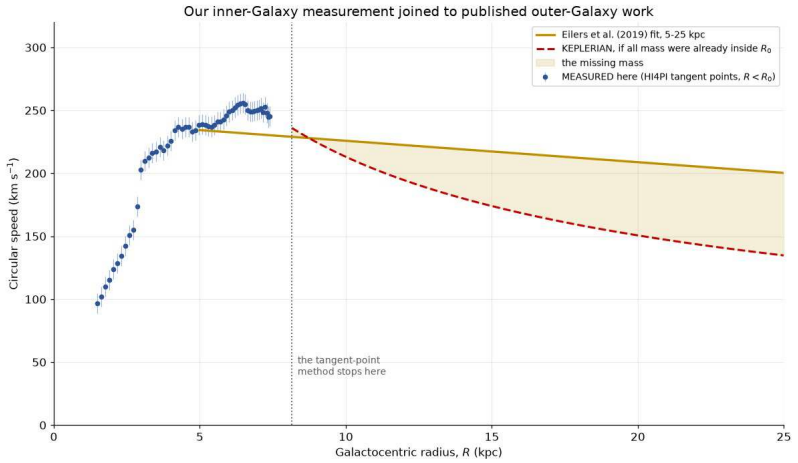


Figure 11. Inner-Galaxy H I result joined to the published Eilers et al. (2019) outer-Galaxy fit. The vertical line marks R_0 and the domain beyond it is external published context; the tangent-point method used in this project does not directly measure the outer Galaxy.

structure and local perturbations. These can shift terminal velocities away from the ideal circular value. The effect is especially important toward the inner few kiloparsecs, where the orbit family is strongly influenced by the bar.

17.5 Mass-model limitations

The expression $M(<R) = \Theta^2 R/G$ is exact only for spherical symmetry. Applied to a disk galaxy, it is best regarded as a spherical-equivalent diagnostic showing the mass scale required by the measured speed. A rigorous decomposition would separately model the bulge, stellar disk, gas disk and dark halo and would use the gravitational potential of flattened components rather than treating all enclosed matter as spherical. The largest limitations are methodological rather than instrumental. First, the tangent-point conversion assumes approximately circular, axisymmetric motion. The Galactic bar, spiral arms, shells, cloud motions and streaming violate that idealization. Second, an absolute 20 K threshold is much larger than the HI4PI thermal noise but can still miss faint high-velocity wings. The NaN results at high longitude directly demonstrate this sensitivity. Third, the $|b| \leq 2^\circ$ latitude average combines a wider range of Galactic heights than a narrow-plane extraction. This changes both the brightness profile and the mixture of kinematic components. Fourth, the calibration region is not populated by enough accepted points, so the edge-bias correction is unavailable. Finally, the spherical-equivalent mass formula is a diagnostic approximation and not a geometric description of the disk. These limitations do not make the run useless. They make it scientifically instructive: the figures expose exactly where an automated pipeline succeeds and where a parameter choice must be reconsidered.

18 Recommended Parameter Validation

A robust follow-up analysis should compare several edge criteria while keeping the data extraction fixed. Absolute thresholds of 20, 10, 5 and 2 K can be tested alongside fractional thresholds, but every result should be judged first by the envelope overlay and the fraction of valid high-longitude detections. The preferred criterion should recover a continuous physical edge without chasing isolated weak features. The low-longitude cutoff should also be tested. Excluding the smallest longitudes reduces the radial range but im-

proves the plausibility of the circular-orbit approximation. Likewise, the effect of the latitude width can be quantified by comparing $|b| \leq 0.5°$, $1°$ and $2°$ extractions. These are controlled scientific comparisons because only one analysis choice changes at a time.

19 Reproducibility

The Python workflow is designed to preserve reproducibility through explicit constants, parameter settings, downloaded-file checksums and saved derived tables. FITS input data are obtained from the HI4PI CDS archive, while NumPy, pandas, Matplotlib and Astropy handle numerical arrays, tables, plotting and coordinate metadata. A complete run should archive the terminal-velocity table together with flags for failed sightlines, the rotation-curve table and the configuration that generated them. The present LaTeX paper uses the supplied real figures as primary evidence. Numerical descriptions in the Results sections are therefore intentionally limited to values visible in those figures rather than invented intermediate statistics that were not supplied separately.

20 Discussion

20.1 What is robust in this experiment

The strongest result is methodological: the archival data clearly contain coherent longitude-dependent H I kinematics, and the tangent-point transformation provides a direct bridge from those spectra to an inner-Galaxy rotation curve. The diagnostic sequence — l–v map, individual spectra, threshold crossings, envelope overlay, calibration test, rotation curve, and mass comparison — makes each inference visually auditable.

20.2 What is configuration-dependent

The exact terminal velocities are not immutable properties of the data independent of analysis. They depend on the operational edge definition, the latitude average, persistence requirement, search interval, and any bias correction. The difference between a 10 K narrow-latitude run and the present 20 K broad-latitude run is therefore scientifically meaningful rather than cosmetic.

20.3 Relation to published outer-Galaxy work

Eilers et al. (2019) measured the Milky Way circular-velocity curve from 5 to 25 kpc using stellar kinematics and found only a gentle decline across that range. Their method and tracer population are independent of the H I tangent-point analysis. Joining the curves in a contextual figure is useful for illustrating the broader Galactic rotation problem, but the published outer curve must not be presented as if it were measured by this project.

20.4 Implications for extended Galactic mass

A point mass produces $\Theta \propto R^{-1/2}$ outside the mass concentration. The observed Galactic curve remains much higher than such a decline over large radial intervals. Historically, this behaviour is one of the key lines of evidence that the gravitating mass distribution is extended. Within the inner Galaxy, baryonic bulge and disk components contribute substantially, so the H I curve alone is not a clean dark-matter decomposition. The outer-Galaxy persistence of high circular speed provides stronger leverage on the halo problem.

20.5 Educational and computational significance

This project reproduces a genuine observational-astronomy workflow with open data: download calibrated survey cubes, inspect metadata, construct derived data products, define an estimator, validate it graphically, propagate geometry, compare with literature, and preserve provenance. That chain is more valuable than producing a single polished curve because it teaches how scientific conclusions depend on both physical assumptions and analysis choices. This parameter configuration demonstrates a useful distinction between obtaining a visually plausible scientific plot and obtaining a fully validated measurement. The l–v diagram and middle-longitude terminal edges are physically coherent, and the resulting intermediate-radius rotation curve displays the familiar non-Keplerian scale of Galactic rotation. At the same time, the 20 K threshold fails for representative high-longitude spectra and prevents the intended Solar-circle edge calibration. The correct response is not to hide those failures. They identify the next experiment: modify the edge criterion, rerun the same archival data, and test whether the accepted points continue to hug the observed emission boundary while extending far enough in longitude to establish the calibration. This iterative process is a central part of computational observational astronomy. The mass and Keplerian-comparison figures remain useful as demonstrations of the dynamical meaning of a high, slowly varying circular speed. However, precision statements about dark matter would require a validated terminal-velocity extraction, a treatment of non-circular motions, and a mass model for the baryonic disk, bulge and gas.

20.6 Interpretive value of the diagnostic sequence

An important strength of the workflow is that the final rotation curve is not treated as an isolated product. The longitude–velocity map, representative spectra, threshold crossings, accepted-edge overlay, calibration test, and mass diagnostics form a linked chain of evidence. This makes it possible to distinguish genuine Galactic structure from artefacts introduced by an edge criterion or by incomplete longitude coverage. In particular, the absence of accepted terminal velocities at some high-longitude sightlines is retained as information about the measurement procedure rather than being hidden by interpolation. Such traceability is essential when archival survey products are transformed into derived dynamical quantities.

20.7 Prospects for refinement

A natural extension is a controlled sensitivity study in which the latitude extraction is held fixed while several absolute and fractional edge criteria are compared. The preferred estimator should be selected by its ability to follow the observed emission envelope and recover a physically continuous high-longitude sequence, not merely by how closely

the resulting rotation curve approaches an expected reference value. A second refinement is to test narrower latitude strips, which may preserve faint terminal wings that are diluted by broader averaging. Together, these checks would strengthen the connection between the measured spectral boundary and the inferred circular speed.

21 Conclusion

Archival HI4PI 21-cm data provide a practical route from real radio spectra to the dynamics of the Milky Way. In the supplied run, averaging over $|b| \leq 2^\circ$ produces a clear longitude–velocity diagram, and a 20 K edge criterion successfully identifies terminal velocities over a substantial part of the first quadrant. The tangent-point transformation then yields a rotation curve rising from the inner region and reaching roughly 230–255 $\mathrm{km\,s^{-1}}$ across much of the intermediate-radius disk. Equally important, the real diagnostic figures expose the limitations of this configuration. Two representative high-longitude spectra return no valid terminal velocity, the accepted sequence does not populate the intended calibration region, and the resulting edge correction remains 0.0 $\mathrm{km\,s^{-1}}$ because no robust bias estimate is available. The project therefore illustrates good scientific practice: results are interpreted together with the diagnostics that determine their reliability. With a revised threshold and a validated high-longitude calibration, the same pipeline can provide a stronger quantitative inner-Galaxy rotation measurement. Even in its present form, the analysis demonstrates the physics of 21-cm Doppler mapping, tangent-point geometry, rotation curves and dynamical mass inference using openly accessible astronomical data.

References

CDS/VizieR. (2016). HI4PI spectra and column density maps (Catalogue J/A+A/594/A116). CDS/VizieR.

Clemens, D. P. (1985). Massachusetts-Stony Brook Galactic Plane CO Survey: The Galactic disk rotation curve. *The Astrophysical Journal*, 295, 422–436. doi:10.1086/163386

Eilers, A.-C., Hogg, D. W., Rix, H.-W., & Ness, M. K. (2019). The circular velocity curve of the Milky Way from 5 to 25 kpc. *The Astrophysical Journal*, 871, 120. doi:10.3847/1538-4357/aaf648

HI4PI Collaboration, Ben Bekhti, N., Flöer, L., Keller, R., Kerp, J., Lenz, D., Winkel, B., Bailin, J., Calabretta, M. R., Dedes, L., Ford, H. A., Gibson, B. K., Haud, U., Janowiecki, S., Kalberla, P. M. W., Lockman, F. J., McClure-Griffiths, N. M., Murphy, T., Nakanishi, H., Pisano, D. J., & Staveley-Smith, L. (2016). HI4PI: A full-sky H I survey based on EBHIS and GASS. *Astronomy & Astrophysics*, 594, A116. doi:10.1051/0004-6361/201629178

Kalberla, P. M. W., Burton, W. B., Hartmann, D., Arnal, E. M., Bajaja, E., Morras, R., & Pöppel, W. G. L. (2005). The Leiden/Argentine/Bonn (LAB) survey of Galactic H I. *Astronomy & Astrophysics*, 440, 775–782. doi:10.1051/0004-6361:20041864

McClure-Griffiths, N. M., & Dickey, J. M. (2007). Milky Way kinematics. I. Measurements at the subcentral point of the fourth quadrant. *The Astrophysical Journal*, 671, 427–438. doi:10.1086/522297

Reid, M. J., Menten, K. M., Brunthaler, A., Zheng, X. W., Dame, T. M., Xu, Y., Wu, Y., Zhang, B., Sanna, A., Sato, M., Hachisuka, K., Choi, Y. K., Immer, K., Moscadelli, L., Rygl, K. L. J., & Bartkiewicz, A. (2014). Trigonometric parallaxes of high mass star forming regions: The structure and kinematics of the Milky Way. *The Astrophysical Journal*, 783, 130. doi:10.1088/0004-637X/783/2/130

Reid, M. J., Menten, K. M., Brunthaler, A., Zheng, X. W., Dame, T. M., Xu, Y., Li, J., Sakai, N., Wu, Y., Immer, K., Zhang, B., Sanna, A., Moscadelli, L., Rygl, K. L. J., Bartkiewicz, A., Hu, B., Quiroga-Nuñez, L. H., & van Langevelde, H. J. (2019). Trigonometric parallaxes of high-mass star-forming regions: Our view of the Milky Way. *The Astrophysical Journal*, 885, 131. doi:10.3847/1538-4357/ab4a11