The Spherical Approximation in MOND:

Quantifying Geometric Biases in the Radial Acceleration Relation

Marijn van Dam

Student Number: 5813638

Supervisors: Dr. P. M. Visser
Dr. S. W. H. Eijt

July 3, 2026

BSc Applied Mathematics & BSc Applied Physics

Faculty of Applied Sciences & Delft Institute of Applied Mathematics

TU Delft Delft University of Technology

Abstract

The Radial Acceleration Relation (RAR) is one of the most striking empirical regularities in galactic dynamics: the observed centripetal acceleration in galaxies correlates tightly with the acceleration predicted from the baryonic mass distribution alone. Modified Newtonian Dynamics (MOND) naturally explains this relation, while the standard dark matter paradigm requires fine-tuning. However, nearly all observational tests of the RAR adopt a simplifying approximation: the baryonic acceleration at radius r is computed as $g_{\rm bar} = GM(< r)/r^2$, as if the galaxy's mass distribution were spherically symmetric. Real disk galaxies are highly flattened, and the validity of this approximation has never been directly quantified for the Radial Acceleration Relation.

This thesis addresses two questions: (1) How accurate is the spherical approximation for computing baryonic accelerations in disk galaxies? (2) Does its use bias previous tests of the MOND RAR? Using a numerical AQUAL solver applied to realistic three-dimensional stellar density reconstructions from Spitzer 3.6 μm imaging, we compare the spherical approximation to the true disk gravity for three galaxies spanning a range of morphologies: NGC 2841 (bulge-dominated), NGC 3198 (moderate bulge), and NGC 2976 (bulgeless).

We find that the accuracy of the spherical approximation for computing baryonic acceleration correlates strongly with bulge prominence. For bulge-dominated NGC 2841, the error is $\lesssim 2\%$; for moderate-bulge NGC 3198, it reaches $\sim 10\%$ in the outer disk; for bulgeless NGC 2976, the spherical approximation underestimates the true baryonic acceleration by $\sim 25\%$ on average, with localised peaks approaching a factor of two. Critically, for bulge-dominated and moderate-bulge galaxies, the spherical approximation does not bias MOND RAR tests: galaxies that follow the MOND prediction using the approximation also follow it using the correct disk gravity. However, for bulgeless galaxies, the approximation can produce a false violation of the RAR: the RAR of NGC 2976 appears to lie systematically above the curve predicted by MOND when the spherical approximation is used, but falls onto the prediction when the correct disk gravity is applied.

We conclude that existing RAR studies are valid for bulge-dominated and normal spiral galaxies, but may have systematically underestimated MOND's consistency with bulgeless, low-surface-brightness systems. Future large-sample tests should apply the full Poisson solver to galaxies without a significant bulge.

Keywords: Modified Newtonian Dynamics (MOND), Radial Acceleration Relation (RAR), galaxy dynamics, spherical approximation, disk galaxies, numerical Poisson solver, AQUAL, NGC 3198, NGC 2841, NGC 2976

Contents

1 Introduction and Motivation

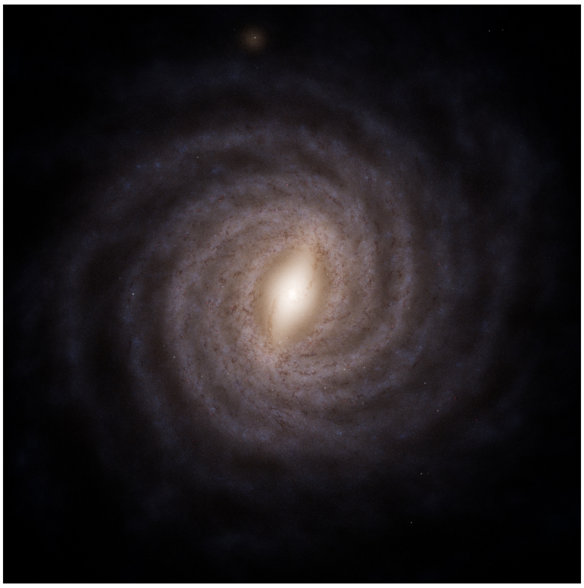

(a) Face-on view

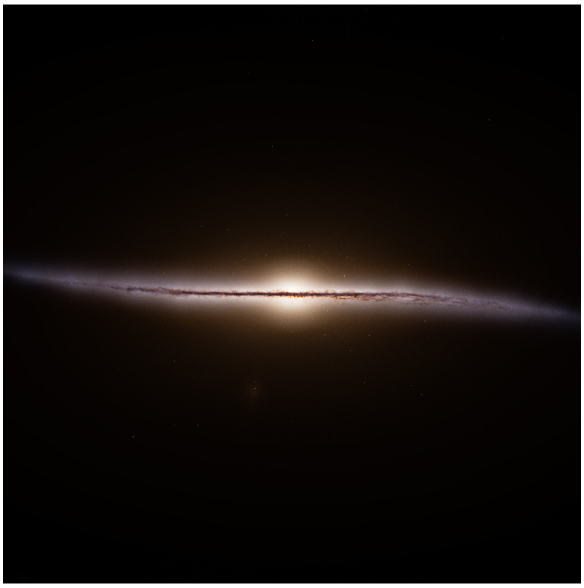

(b) Edge-on view

Figure 1: The flattened disk geometry of a typical spiral galaxy. The left panel shows a face-on view, revealing the exponential stellar disk, spiral arms, and central bulge. The right panel shows the same galaxy in edge-on projection, clearly illustrating the thin, flattened nature of the disk with a central bulge. This extreme flattening motivates the central question of this thesis: can the gravitational field of such a highly non-spherical system be accurately approximated as spherically symmetric?
Image credit: [European Space Agency (ESA) [1]].

1.1 The Galaxy Rotation Curve Problem: Dark Matter vs. Modified Gravity

One of the most enduring puzzles in modern astrophysics is the discrepancy between the luminous matter we observe in galaxies and the gravitational dynamics those galaxies exhibit. In a purely Newtonian framework, the circular velocity of a star orbiting at radius R in a galaxy should fall off as $v_c \propto R^{-1/2}$ once the bulk of the stellar mass has been enclosed, similar to the behaviour of planets in the outer Solar System. Instead, observations of spiral galaxies consistently show rotation curves that remain approximately flat at large radii, with $v_c \approx$ const extending far beyond the visible stellar disk [3, 4].

This discrepancy was first systematically characterised by Rubin et al. [3] and has since been confirmed for hundreds of galaxies spanning a wide range of morphologies, surface brightnesses, and sizes. The canonical interpretation within the standard cosmological model (ΛCDM) is that galaxies are embedded in extended halos of non-baryonic dark matter whose gravitational potential dominates at large radii. The dark matter hypothesis is well-motivated by independent evidence from the cosmic microwave background [5], large-scale structure, and gravitational lensing, and ΛCDM has achieved remarkable success as a cosmological framework.

Nevertheless, the dark matter interpretation of galaxy rotation curves faces persistent challenges on small scales. The predicted density profiles of dark matter halos (cuspy NFW profiles [6]), are in tension with the cored profiles inferred from observed dwarf galaxy rotation curves, a discrepancy known as the core-cusp problem [7]. More fundamentally, the observed kinematics of galaxies show a surprisingly tight empirical correlation with the *baryonic* mass distribution alone, a regularity that is difficult to understand in a framework where the dominant mass component is dark matter distributed independently of the baryons [8].

An alternative class of explanations modifies the law of gravity itself rather than invoking new matter. The most developed of these is Modified Newtonian Dynamics (MOND), introduced by Milgrom [9], which proposes that Newtonian dynamics breaks down at the very low accelerations characteristic of galaxy outskirts. As discussed in this thesis, MOND has achieved considerable empirical success in predicting the rotation curves of a diverse sample of galaxies from their baryonic mass distributions alone, without free parameters beyond a single universal acceleration scale a_0.

The question of whether the universe contains dark matter, or whether gravity is modified, remains open. This thesis does not attempt to resolve it. Instead, it addresses a more focused methodological question: are the standard tools used to test MOND, in particular, the Radial Acceleration Relation, subject to systematic biases introduced by geometric approximations in the analysis?

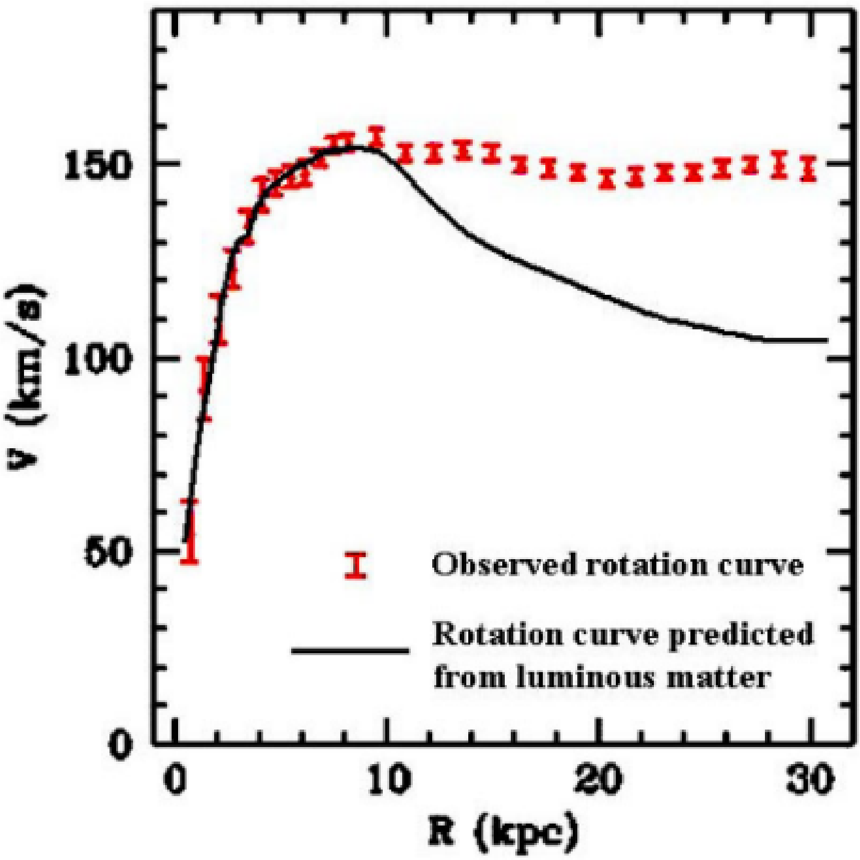

Figure 2: Schematic illustration of the galaxy rotation curve problem. The observed rotation curve of a typical spiral galaxy (red points with error bars) remains approximately flat at large radii, with $v_c \approx$ const extending far beyond the visible stellar disk. The Newtonian prediction from the stellar mass distribution alone (black line) rises to a peak and then declines as $v_c \propto R^{-1/2}$, failing to account for the observed velocities. Two classes of explanation exist: (i) the dark matter hypothesis, which posits an extended halo of non-baryonic matter whose gravitational potential dominates at large radii; and (ii) modified gravity, which alters the law of gravity itself in the low-acceleration regime. Modified Newtonian Dynamics (MOND) naturally produces asymptotically flat rotation curves without invoking dark matter, as the circular velocity in the deep-MOND regime becomes $v_c = (GMa_0)^{1/4}$, independent of radius.

Data points adapted from University of Glasgow [2]

1.2 Modified Newtonian Dynamics (MOND): Core Principles and Successes

MOND was proposed by Milgrom [9] as a phenomenological modification of Newtonian dynamics in the regime of very low accelerations. The central idea is that Newtonian mechanics holds when accelerations are large compared to a fundamental scale $a_0 \approx 1.2 \times 10^{-10}$ m s^{-2}, but transitions to a modified regime when accelerations fall below a_0. In the deep-MOND regime ($g \ll a_0$), the effective gravitational acceleration scales as $g \propto \sqrt{g_N a_0}$, where g_N is the Newtonian prediction. This modification naturally produces asymptotically flat rotation curves, since the circular velocity in the deep-MOND limit becomes $v_c = (GMa_0)^{1/4}$, independent of radius.

The theoretical basis for MOND was placed on firmer footing by Bekenstein and Milgrom [10], who developed the AQUAL (AQUAdratic Lagrangian) formalism. AQUAL is a fully self-consistent non-relativistic field theory that modifies the Poisson equation for the gravitational potential. An alternative formulation, QuMOND (Quasi-linear MOND), was introduced by Milgrom [11] and is often preferred for numerical work because it reduces to two sequential linear Poisson solves rather than a single non-linear one. Both formalisms reduce to the same algebraic approximation for spherically symmetric systems and yield nearly identical results for disk galaxies [12].

The empirical success of MOND has been convincing. The single parameter a_0 was originally calibrated from a handful of rotation curves, yet it predicts the kinematics of hundreds of galaxies with no additional free parameters [8, 13]. The tight correlation between observed centripetal accelerations and those predicted from baryonic mass alone, now quantified as the Radial Acceleration Relation (RAR; McGaugh et al. 8), is a natural consequence of MOND but requires careful fine-tuning in ΛCDM models. The value of a_0 is also intriguingly close to $cH_0/2\pi$, suggesting a possible connection between the MOND acceleration scale and the large-scale expansion of the universe, though this remains unexplained.

Despite these successes, MOND faces challenges of its own. It struggles to account for the dynamics of galaxy clusters without invoking additional unseen mass [14], and it has not yet been successfully extended into a fully relativistic framework, though several candidates exist [15, 16]. The question of whether MOND is the correct description of nature, or merely a phenomenological approximation to some deeper theory, remains unresolved.

1.3 Spherical Approximation in RAR and Galaxy Modelling

Observational tests of MOND typically proceed by comparing the observed centripetal acceleration $g_{\rm obs} = v_c^2(R)/R$ at each radius with the MOND prediction derived from the baryonic mass distribution. The latter requires knowledge of the Newtonian gravitational acceleration g_N produced by the stars and gas, which then enters the MOND interpolation function to yield the predicted $g_{\rm obs}$.

In practice, computing g_N for a realistic disk galaxy requires solving the three-dimensional Poisson equation for a flattened, extended mass distribution. For this reason, the vast majority of observational RAR studies, including the seminal work of McGaugh et al. [8] and the subsequent SPARC analysis of Lelli et al. [17], adopt a simplifying approximation: the baryonic acceleration at radius r is estimated as

$$g_{\rm bar} = \frac{GM(<r)}{r^2}, \tag{1}$$

where $M(<r)$ is the cumulative baryonic mass enclosed within a sphere of radius r, computed by integrating the azimuthally averaged surface density profile. This is the *spherical approximation*: it treats the inherently flattened disk mass distribution as if it were spherically symmetric.

The spherical approximation is attractive because it reduces the computation of g_N to a simple one-dimensional integral over a surface brightness profile, requiring no numerical solver and no knowledge of the three-dimensional structure of the disk. It is exact for spherically symmetric systems by virtue of Gauss's law, but disk galaxies are manifestly non-spherical: a typical spiral galaxy has a scale height-to-scale length ratio of order $h_z/R_d \sim 0.1$–0.2 (Figure 1), making it far more flattened than any sphere. It is therefore not obvious that the spherical approximation provides an accurate estimate of g_N for a disk galaxy, nor whether any systematic inaccuracy in this estimate biases the conclusions drawn from RAR studies.

1.4 Quantifying the Impact of Geometry in MONDian Galaxy Models

This thesis addresses the following two questions directly:

1. *How accurate is the spherical approximation $g_{\rm bar} = GM(<r)/r^2$ for computing the baryonic gravitational acceleration of disk galaxies?*

2. *Does any systematic inaccuracy in the spherical approximation bias empirical tests of the MOND Radial Acceleration Relation?*

To answer these questions, we implement a numerical AQUAL solver that computes the full three-dimensional gravitational field of a realistic disk galaxy, without assuming spherical symmetry. We apply this solver to three galaxies spanning a range of morphological types (Figure 3): NGC 3198 (a moderate-bulge spiral), NGC 2841 (a bulge-dominated spiral), and NGC 2976 (a bulgeless, low-surface-brightness dwarf). For this we use stellar mass density cubes constructed from

Spitzer 3.6 μm imaging from the S4G survey. By comparing the true disk gravity with the spherical approximation for each galaxy, and by computing the MOND RAR under both assumptions, we provide the first direct, quantitative test of whether the spherical approximation biases MOND analyses.

Previous work has studied whether the MOND algebraic approximation accurately reproduces the full AQUAL or QuMOND solution *given the correct Newtonian field* [12, 18]. Our study is complementary and distinct: we ask whether the standard method for computing the Newtonian field itself (the spherical enclosed-mass approximation) introduces systematic errors, and whether those errors matter for MOND tests. To our knowledge, this question has not previously been addressed in the literature.

The thesis is structured as follows. Section 2 introduces the theoretical foundations of MOND, including the AQUAL and QuMOND field equations, the Radial Acceleration Relation, and the role of geometry in non-linear modified gravity. Section 3 describes the numerical solver, including its validation against analytic test cases. Section 4 presents optimisations to the solver that substantially improve its accuracy for disk-shaped mass distributions. Section 5 describes the construction of stellar mass density cubes from infrared imaging. Section 6 presents the main results: the quantified error of the spherical approximation as a function of galaxy morphology, and its impact on the RAR. Section 7 discusses the implications, limitations, and directions for future work.

(a) Bulge-dominated galaxy (UGC 10043)

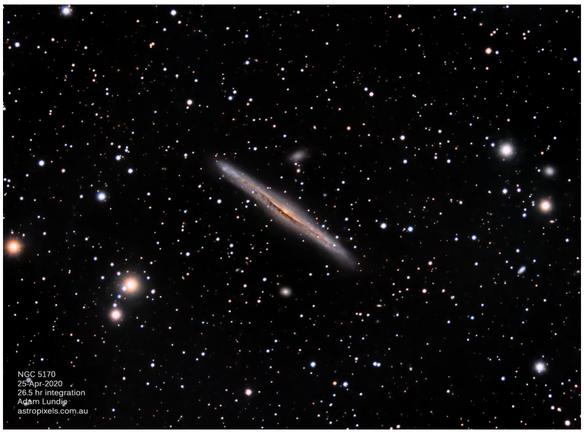

(b) Bulgeless galaxy (e.g., NGC 5170)

Figure 3: Contrasting galaxy morphologies: bulge-dominated versus bulgeless disks. The left panel shows a bulge-dominated spiral galaxy (UGC 10043), characterised by a prominent central concentration of stars that dominates the inner gravitational potential. The presence of a massive, approximately spherical bulge means that the mass distribution is closer to spherical symmetry, and the spherical approximation $g_{\rm bar} = GM(<r)/r^2$ is expected to be more accurate. The right panel shows a bulgeless galaxy (NGC 5170), which lacks a central bulge and is instead dominated by a thin, extended stellar disk. Such galaxies are highly flattened with no spherical component to "average over" the disk geometry, and the spherical approximation is expected to fail most severely. Understanding how the accuracy of the spherical approximation varies with bulge prominence is a central goal of this thesis.

Image credit: [ESA/Hubble and NASA [19], Lundie [20]].

2 Theoretical Foundations of Modified Newtonian Dynamics

2.1 The Radial Acceleration Relation (RAR)

Modified Newtonian Dynamics (MOND) was first proposed by Milgrom [9] as an alternative explanation for the observed mass discrepancies in galaxies without invoking dark matter. The theory introduces a fundamental acceleration scale a_0 that marks the transition between Newtonian and modified dynamics. Empirically, $a_0 \approx 1.2 \times 10^{-10}$ m s^{-2}, a value curiously close to $cH_0/2\pi$, where c is the speed of light and H_0 is the Hubble constant.

The cornerstone of MOND's empirical success is the Radial Acceleration Relation (RAR), first formulated by Milgrom [9]. This relation shows a tight correlation between the observed centripetal acceleration in galaxies, $g_{\rm obs} = v^2(r)/r$, and the gravitational acceleration predicted from the baryonic mass distribution alone, $g_{\rm bar}$, following:

$$g_{\rm obs} = \nu\left(\frac{g_{\rm bar}}{a_0}\right) g_{\rm bar}, \tag{2}$$

where $\nu(y)$ is a transition function between the Newtonian and deep-MOND regime that asymptotically satisfies:

$$\nu(y) \to 1 \qquad \text{for } y \gg 1 \quad \text{(Newtonian regime)}, \tag{3}$$

$$\nu(y) \to 1/\sqrt{y} \qquad \text{for } y \ll 1 \quad \text{(deep-MOND regime)}. \tag{4}$$

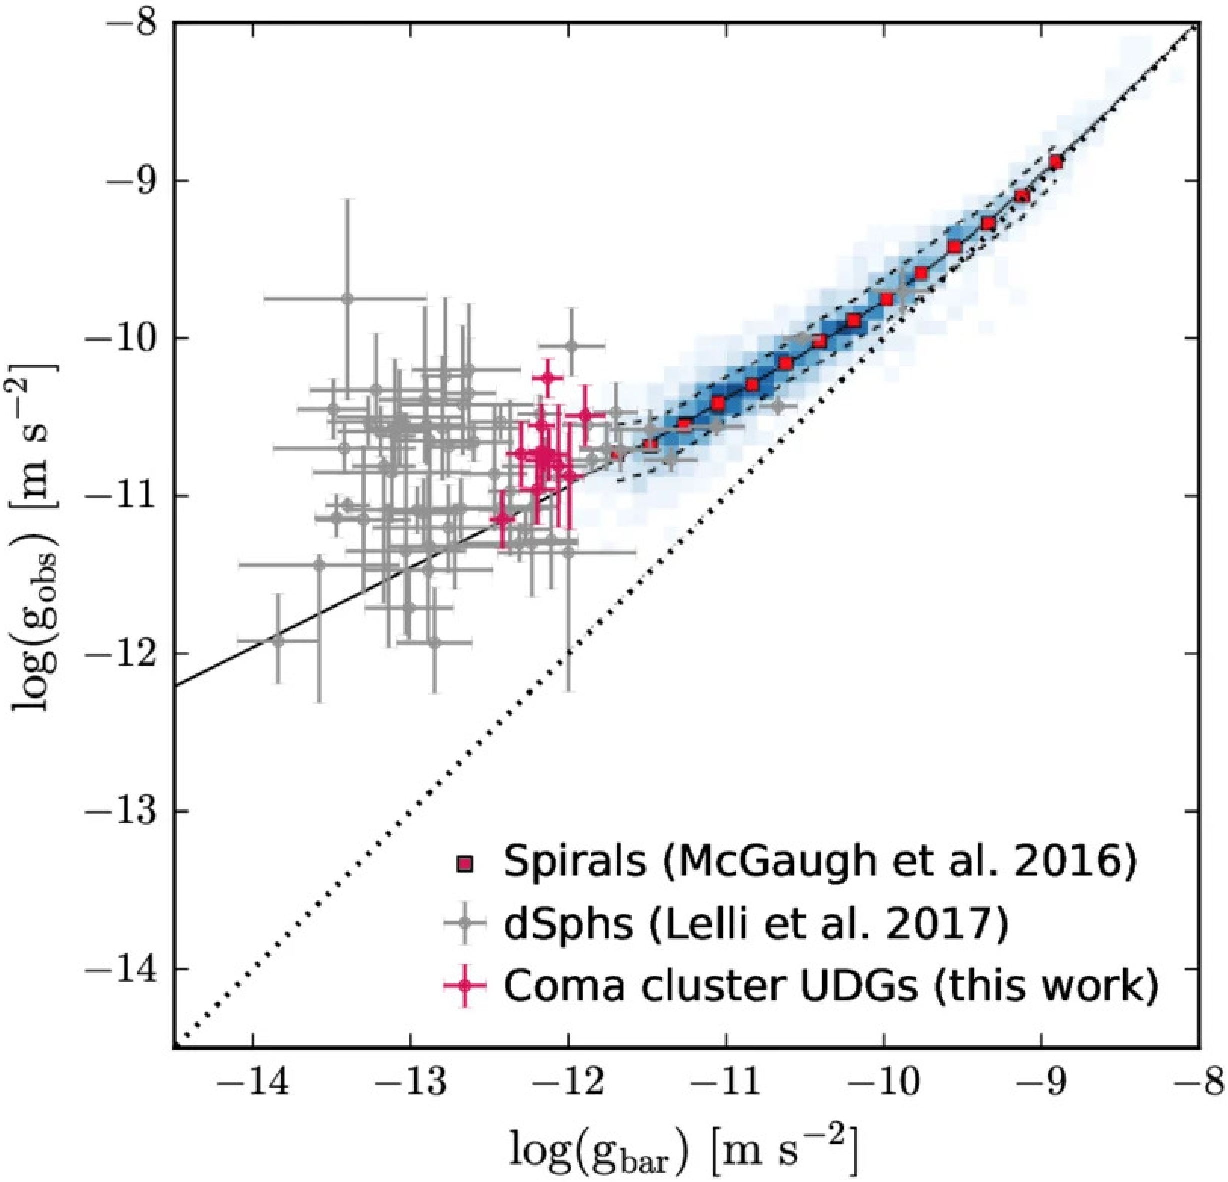

Figure 4: Radial acceleration relation (RAR) comparing the dynamics of different galaxy types. The magenta error bars show the measurements for ultra-diffuse galaxies (UDGs) in the Coma cluster. The background blue color scale and red squares represent the distribution and binned mean respectively, for the SPARC sample of local spiral galaxies (McGaugh et al. [8]). The gray error bars show the relation for dwarf spheroidal (dSph) satellites of the Local Group (Lelli et al. [17]). The solid black line is the empirical best-fit function derived from the SPARC sample, which encapsulates the phenomenological behavior of Milgromian dynamics (MOND). The diagonal dashed black line indicates the Newtonian expectation $g_{\text{obs}} = g_{\text{bar}}$ (i.e., no dark matter), while the additional dashed red line with a slope of approximately $1/2$ marks the deep-MOND regime where $g_{\text{obs}} \propto \sqrt{g_{\text{bar}}}$.

In figure 4 it is clear that for $g_{\text{bar}} \gg a_0$ the log-log plot of g_{obs} versus g_{bar} has gradient 1, as we would expect from regular Newtonian dynamics. However for $g_{\text{bar}} \ll a_0$, the gradient drops to $\frac{1}{2}$.

In the deep-MOND regime ($g_{\text{bar}} \ll a_0$): $\nu(\frac{g_{\text{bar}}}{a_0}) \approx (g_{\text{bar}}/a_0)^{-1/2}$, so $g_{\text{obs}} \approx \sqrt{g_{\text{bar}} a_0}$, giving a slope of $\frac{1}{2}$ on a log-log plot.

2.2 Non-Linear Poisson Formalism: The AQUAL Field Equation

The first complete, non-relativistic Lagrangian formulation of MOND was developed by Bekenstein and Milgrom [10], known as the AQUAL (AQUAdratic Lagrangian) formalism. The theory modifies the Poisson equation for the gravitational potential ϕ to:

$$\nabla \cdot \left[\mu \left(\frac{|\nabla \phi|}{a_0} \right) \nabla \phi \right] = 4\pi G \rho_b, \tag{5}$$

where ρ_b is the baryonic mass density and $\mu(x)$ is the interpolation function that encodes the transition between the two regimes, with $x = \frac{|\nabla \phi|}{a_0} = \frac{g_M}{a_0}$.

2.2.1 The Interpolation Function and its Relation to the RAR

The interpolation function $\mu(x)$ must satisfy the asymptotic conditions:

$$\mu(x) \to 1 \qquad \text{for } x \gg 1, \tag{6}$$

$$\mu(x) \to x \qquad \text{for } x \ll 1. \tag{7}$$

The connection between the AQUAL formalism and the observed Radial Acceleration Relation (RAR) emerges through the relation between the interpolation function $\mu(x)$ and the transition function $\nu(y)$ appearing in Eq. (2). For spherically symmetric systems, the AQUAL equation reduces to an algebraic relation between the Newtonian acceleration g_N and the MOND acceleration g_M:

$$\mu \left(\frac{g_M}{a_0} \right) g_M = g_N, \tag{8}$$

where $g_M = |\nabla \phi|$ and g_N is the Newtonian acceleration corresponding to the same baryonic mass distribution.
We can define $x = g_M/a_0$ and $y = g_N/a_0$. Then Eq. (8) becomes:

$$\mu(x)x = y. \tag{9}$$

The observed acceleration in the RAR is $g_{\text{obs}} = g_M$, while the baryonic prediction is $g_{\text{bar}} = g_N$. Therefore, the RAR transition function $\nu(y)$ satisfies:

$$g_M = \nu(y) g_N = \nu(y) a_0 y. \tag{10}$$

But since $g_M = a_0 x$, we have:

$$a_0 x = \nu(y) a_0 y \quad \Rightarrow \quad x = \nu(y) y. \tag{11}$$

Comparing this with Eq. (9), we see that:

$$\nu(y) y = x = \frac{y}{\mu(x)}. \tag{12}$$

Thus, the functions μ and ν are related through:

$$\nu(y) = \frac{1}{\mu(x)} \quad \text{and} \quad \mu(x) = \frac{1}{\nu(y)}, \tag{13}$$

where x and y are connected by $\mu(x)x = y$.
For the commonly used standard interpolating function:

$$\mu(x) = x/\sqrt{1 + x^2} \tag{14}$$

we can derive the corresponding $\nu(y)$ explicitly. (Section A, Appendix)

$$\nu(y)^2 = \frac{1 + \sqrt{1 + 4/y^2}}{2}. \tag{15}$$

Taking the square root gives the standard ν-function:

$$\nu(y) = \sqrt{\frac{1}{2} + \frac{1}{2}\sqrt{1 + \frac{4}{y^2}}}. \tag{16}$$

The asymptotic behavior is consistent: when $y \gg 1$ (Newtonian regime), $\nu(y) \to 1$, and when $y \ll 1$ (deep-MOND regime), $\nu(y) \to 1/\sqrt{y}$.

2.2.2 The Deep-MOND and Newtonian Limits

In the deep-MOND regime ($|\nabla\phi| \ll a_0$), where $\mu(x) \approx x$, Eq. (5) becomes:

$$\nabla \cdot \left(\frac{|\nabla\phi|}{a_0} \nabla\phi \right) = 4\pi G \rho_b, \tag{17}$$

which is non-linear but scale-invariant under $(t, \vec{r}) \longrightarrow (\lambda t, \lambda\vec{r})$. For an isolated point mass M at the origin, the solution of 17 yields the logarithmic potential:

$$\phi(r) = \sqrt{GMa_0} \ln \left(\frac{r}{r_0} \right), \tag{18}$$

giving the constant circular velocity $v_c = (GMa_0)^{1/4}$.
In the Newtonian limit ($|\nabla\phi| \gg a_0$), $\mu(x) \approx 1$ and Eq. (5) reduces to the standard Newtonian Poisson equation:

$$\nabla^2\phi = 4\pi G \rho_b. \tag{19}$$

2.3 Quasi-Linear MOND (QuMOND): An Alternative Formulation

Milgrom [11] introduced an alternative formulation called Quasi-linear MOND (QuMOND), which can be computationally more tractable. In this formulation, the MOND acceleration field $\vec{g}_\mathrm{M} = -\nabla\phi_\mathrm{M}$ is derived from the Newtonian potential. The field equation is:

$$\nabla^2\phi_\mathrm{M} = \nabla \cdot \left[\nu \left(\frac{|\nabla\phi_\mathrm{N}|}{a_0} \right) \nabla\phi_\mathrm{N} \right], \tag{20}$$

where ϕ_N satisfies the standard Poisson equation $\nabla^2\phi_\mathrm{N} = 4\pi G\rho_b$, and $\nu(y)$ is the QuMOND interpolation function. The resulting acceleration can be expressed as an algebraic relation:

$$\vec{g}_\mathrm{M} = \nu \left(\frac{|\vec{g}_\mathrm{N}|}{a_0} \right) \vec{g}_\mathrm{N} + \vec{H}, \tag{21}$$

where $\vec{g}_\mathrm{N} = -\nabla\phi_\mathrm{N}$. Here, $\vec{H}$ is the divergence-free (magnetic) field that is introduced to ensure that the modified gravity field remains curl-free, such that $\vec{g}_\mathrm{M} = -\nabla\phi_\mathrm{M}$.
QuMOND is often preferred for numerical calculations because it involves solving linear Poisson equations twice (for ϕ_N and ϕ_M), which is simpler than solving the non-linear AQUAL equation directly.

2.4 Non-Spherical Geometries

The non-linearity of the MOND field equations means that the superposition principle does not hold. This makes the geometry of the mass distribution critically important. While many analytical results and fitting functions (like the algebraic relation in Eq. (21)) are derived under the assumption of spherical symmetry, real galaxies are predominantly axisymmetric disks. When scientists study the RAR, they typically assume spherical symmetry to relate the observed circular velocity to the radial gravitational field ($g_r = -\partial\Phi/\partial r$). However, in flattened systems, the radial acceleration experienced by a particle on a circular orbit in the mid-plane is not simply the radial derivative of the potential; the vertical structure and the disk's geometry contribute non-trivially.

2.4.1 The Enclosed Mass Approximation and Spherical Symmetry

In a spherically symmetric mass distribution, the Newtonian gravitational field at radius r depends only on the mass enclosed within that radius. In this special case, Gauss's law implies that the radial gravitational acceleration can be written as

$$g_N(r) = \frac{GM(< r)}{r^2}, \tag{22}$$

where $M(< r)$ denotes the mass enclosed within a sphere of radius r. This relation is frequently employed in the construction of rotation curves and is often implicitly assumed in empirical analyses of the radial acceleration relation (RAR).
However, Eq. (22) is exact only for spherically symmetric systems. For flattened mass distributions such as disk galaxies, mass elements located at radii larger than r contribute non-negligibly to the radial force in the disk plane. As a result, the gravitational acceleration cannot, in general, be expressed solely in terms of an enclosed mass. Applying Eq. (22) to disk galaxies therefore constitutes an approximation that effectively enforces spherical symmetry.
In this work, this spherical approximation is treated explicitly as a modelling assumption rather than a fundamental property of MOND. By comparing MOND predictions obtained using the enclosed-mass approximation to those derived from a full solution of the QuMOND field equations for an axisymmetric disk, we quantify the impact of abandoning spherical symmetry on the predicted radial acceleration relation.

Jones-Smith et al. [12] studied whether the full QuMOND field equations reduce to the simple algebraic prescription $g_M = g_N \cdot \nu(g_N/a_0)$, assuming the correct disk-derived Newtonian field g_N. They concluded that for exponential disks the algebraic prescription provides an excellent approximation to the full QuMOND solution for thin disks, with deviations of at most a few percent in the inner regions and slightly larger discrepancies (up to $\sim 10\%$) in the outer low-acceleration regime. However, they found that these deviations are primarily confined to the vertical direction, while the radial acceleration is reproduced to within $\sim 1\%$ accuracy. Importantly, they attributed the residual differences to the breakdown of the curl-free assumption for the MOND field in regions where the Newtonian field is not spherically symmetric, but concluded that for most practical applications the algebraic prescription is sufficiently accurate. They did not examine the effect of computing g_N itself under the simplifying assumption of spherical symmetry ($g_N = GM(r)/r^2$) rather than solving the Poisson equation for a realistic disk geometry. That question is the subject of this study as it remains thus far unexplored in the literature.

2.5 Disk Galaxies and Numerical Implementation

To address the research question, we implement a numerical pipeline to compare MOND predictions under different assumptions for a disk galaxy.

2.5.1 QuMOND Poisson Solver: The General Case

The QuMOND field equation (Eq. (20)) is valid for any mass distribution. This equation must generally be solved in three dimensions ($\vec{r} = (x, y, z)$. However, for a *thin* disk galaxy, one can often make the approximation of solving it in three dimensions and only looking at the mass distribution in two dimensions (the mid-plane or $\vec{r} = (x, y, 0)$), effectively treating the vertical dimension as infinitely thin. This gives us the following density relation:

$$\rho(x, y, z) = \sigma(x, y) \cdot \delta(z) \tag{23}$$

In this relation $\delta(z)$ describes the Dirac-Delta function.

2.5.2 The Relation Between Algebraic and QuMOND Predictions

The algebraic relation (Eq. (21)), $\vec{g}_\mathrm{M} = \nu(|\vec{g}_\mathrm{N}|/a_0)\vec{g}_\mathrm{N}$, is exact only when the MOND acceleration field $\vec{g}_\mathrm{M}$ is curl-free. This occurs in two special cases:

1. For spherically symmetric systems, where $\vec{g}_\mathrm{N}$ is purely radial.
2. For infinite systems (infinite plane, infinite cylinder) with translational symmetry.

Even though $\vec{g}_\mathrm{N}$ is curl-free, in general $\vec{g}_\mathrm{N} \cdot \nu(|\vec{g}_\mathrm{N}|/a_0)$ is not curl-free. The QuMOND equation accounts for this through the additional divergence operation. The difference arises because the non-linear function ν distorts the field lines. As mentioned in the previous section, numerical studies have shown that for exponential disks, the curl term is surprisingly small [12]. This explains why in the model galaxy calculation, QuMOND and the algebraic approximation of MOND give nearly identical results.
The algebraic approximation applied to a disk's Newtonian acceleration,

$$g_\mathrm{alg}(R) = \nu\left(\frac{g_N^\mathrm{disk}(R)}{a_0}\right) g_N^\mathrm{disk}(R),$$

implicitly assumes that applying ν to the magnitude of $\vec{g}_\mathrm{N}$ and then multiplying by its direction yields the correct MOND acceleration.

2.5.3 Implementation of the Numerical Solver

The code implements a 3D QuMOND solver in the midplane that:

1. Create a 3D density array either from a either model (Chapter 3) or observations (Chapter 5).
2. Constructs an axisymmetric 3D Newtonian vector field $\vec{g}_N(x, y, 0)$ on a Cartesian grid using the Poisson Equation.
3. Computes the source term $S = \nabla \cdot [\nu(|\vec{g}_N|/a_0)\vec{g}_N]$.
4. Solves $\nabla^2\phi_\mathrm{M} = S$ via Fourier transforms.

The assumption of axisymmetry enters at step 2, where we construct $\vec{g}_N(x, y, 0)$ from a radial profile. This is valid for an axisymmetric disk but would need modification for barred or otherwise non-axisymmetric galaxies, like galaxies with spiral arms. The more significant assumption is the **thin-disk approximation**, assuming mass is only in the midplane rather than in full 3D.

2.5.4 Should These Approximations Be a Concern?

For the study of the RAR in typical spiral galaxies:

- The axisymmetry assumption is reasonable for most SPARC [21] galaxies, which are generally well-described by exponential disks.
- The thin-disk approximation may introduce small errors, particularly in the inner regions where the disk has finite thickness. However, studies suggest this effect is minor for rotation curve analysis [18]. Brada and Milgrom [18] showed that for exponential disks, the difference between thin-disk and finite-thickness QuMOND solutions is typically less than 5% for the rotation curve.
- The key comparison is between the full QuMOND solution and the spherical approximation.

Thus, while our implementation of the solver assumes axisymmetry, this is physically justified for the systems studied. The more important distinction is between the spherical approximation and the QuMOND solution.

3 Methodology and Tests: Comparing Spherical and Disk-Based MOND Solutions

3.1 Baseline Solver Description

The numerical calculations in this chapter and the optimizations in Chapter 4 are based on a Poisson solver originally developed by de Nijs [22]. This section briefly describes the baseline implementation to contextualize both the validation tests (Section 3.3) and the subsequent improvements.
The solver implements a fixed-point iteration scheme for the AQUAL field equation (Eq. 5) on a Cartesian grid of size 64^3 cells. Key features include FFT-based Poisson solving and spectral projection operators.
Although validated for 2D FLAT systems, several limitations motivate the optimizations described in Chapter 4: periodic boundary conditions distort the field at large radii, and finite-difference gradients introduce $O(\Delta x^2)$ errors.

3.2 The Spherical Approximation: Common Practice and Its Limitations

In most MOND analyses of galaxy rotation curves, including the seminal RAR study by McGaugh et al. [8], a key simplification is the *spherical approximation*. We define the surface brightness profile $\Sigma(R)$ as the projected surface density of baryonic matter as a function of the cylindrical radius R in the galactic plane for a disk galaxy. Then, one can compute the following.

$$M_{\rm bar}(r) = 2\pi \int_0^r \Sigma(R) R\, dR, \quad g_{\rm bar,sph}(r) = \frac{G M_{\rm bar}(< r)}{r^2}, \tag{24}$$

and subsequently applies the MOND relation (either algebraically via the algebraic relation Eq. (21) or by solving the full AQUAL model Eq. (5) assuming spherical symmetry) to obtain $g_{\rm obs,sph}(r)$. This procedure treats the intrinsically flattened disk mass distribution as if it were spherically symmetric.
The appeal of this approximation lies in computational simplicity: solving the modified Poisson equation (Eq. (5)) for realistic disk geometries requires non-linear numerical methods. However, because disk galaxies are manifestly non-spherical systems, the accuracy of this assumption warrants quantitative investigation.

3.3 Numerical Implementation and Validation

To assess the impact of the spherical approximation on the predicted radial acceleration relation, we implemented a numerical solver for the AQUAL field equation (Eq. (5)) using a fixed-point iteration scheme combined with Fast Fourier Transforms (FFTs). The solver operates on a three-dimensional Cartesian grid with $(2 \times 128)^3$ cells and adopts the interpolation functions defined previously in Eqs. (14) and (16).
The code is written in Python using CuPy for GPU acceleration, allowing efficient treatment of large computational grids.

3.4 Code units for validation.

For the validation tests in this section (point mass, exponential disk, and uniform sphere), we adopt dimensionless code units with $G = a_0 = 1$, box size $L = 4$, and masses treated as dimensionless scalars (e.g., $M = 0.01$). Lengths are expressed in grid units (256 cells span $L = 4$). These choices simplify the MOND equations and isolate numerical behavior from astrophysical scaling. Physical units (kpc, $M_\odot$, SI) are reintroduced in Section 5 when we construct a density cube from infrared imaging.

3.5 Numerical method

Starting from the Newtonian acceleration field $\vec{g}_N = -\nabla\phi_N$, obtained by solving the linear Poisson equation $\nabla^2\phi_N = 4\pi G\rho_b$ via an FFT-based solver, we seek the MOND acceleration $\vec{g}_M = -\nabla\phi_M$ satisfying Eq. (5). Following Brada and Milgrom [18], we introduce a residual field $\vec{H}$. Then the iteration process proceeds as

$$\vec{F} = \vec{g}_N + \vec{H}, \tag{25}$$

$$\vec{g}_M = \mathcal{P}_{\rm cf}\Big[\nu(|\vec{F}|/a_0)\,\vec{F}\Big], \tag{26}$$

$$\vec{H}_{\rm new} = \mathcal{P}_{\rm df}\Big[\mu(|\vec{g}_M|/a_0)\,\vec{g}_M - \vec{g}_N\Big], \tag{27}$$

where $\mathcal{P}_{\rm cf}$ and $\mathcal{P}_{\rm df}$ denote curl-free and divergence-free projection operators implemented in Fourier space.
In Fourier space, the curl-free projection of a vector field $\vec{A}$ corresponds to its longitudinal component,

$$\widehat{\mathcal{P}_{\rm cf}[\vec{A}]}(\vec{k}) = \begin{cases} \dfrac{\vec{k}(\vec{k}\cdot\hat{\vec{A}}(\vec{k}))}{k^2}, & \vec{k} \neq 0, \\ \hat{\vec{A}}(0), & \vec{k} = 0, \end{cases} \tag{28}$$

while the divergence-free projection extracts the transverse component,

$$\widehat{\mathcal{P}_{\rm df}[\vec{A}]}(\vec{k}) = \hat{\vec{A}}(\vec{k}) - \begin{cases} \dfrac{\vec{k}(\vec{k}\cdot\hat{\vec{A}}(\vec{k}))}{k^2}, & \vec{k} \neq 0, \\ 0, & \vec{k} = 0. \end{cases} \tag{29}$$

Here, $\vec{k} = (k_x, k_y, k_z)$ is the wave vector in Fourier space, with $k = |\vec{k}| = \sqrt{k_x^2 + k_y^2 + k_z^2}$ being the magnitude of the wave vector. The components k_x, k_y, and k_z correspond to the spatial frequencies in the x, y, and z directions, respectively. The special case $\vec{k} = 0$ represents the zero-frequency (DC) mode, which is handled separately to avoid division by zero in the projection operators. The Fourier transform convention used is $\hat{\vec{A}}(\vec{k}) = \int \vec{A}(\vec{x})e^{-i\vec{k}\cdot\vec{x}}\,d^3x$, with the inverse transform $\vec{A}(\vec{x}) = \int \hat{\vec{A}}(\vec{k})e^{i\vec{k}\cdot\vec{x}}\,d^3k/(2\pi)^3$.
All Fourier transforms use CuPy real-to-complex FFT routines, and the Poisson equation is solved using the kernel k^{-2}.
The iteration is initialized with $\vec{H} = 0$. A relaxation step,

$$\vec{H} \leftarrow \omega\,\vec{H}_{\rm new} + (1-\omega)\,\vec{H}, \tag{30}$$

is included for stability; here we simply adopt $\omega = 1$ ($\vec{H} \leftarrow \vec{H}_{\rm new}$). Convergence is reached when the maximum change in $\vec{g}_M$ falls below 10^{-6}, typically after 10–15 iterations.

3.5.1 Point mass test (spherical symmetry)

As a first validation, we model a point mass with mass M smoothed by a Gaussian density profile

$$\rho(r) = \frac{M}{(2\pi)^{3/2}\sigma^3}e^{-r^2/(2\sigma^2)}, \tag{31}$$

with $\sigma = 4$ cells. The analytic Newtonian acceleration is

$$g_N^{\rm ana}(r) = \frac{GM}{r^2}\,\mathrm{erf}\!\left(\frac{r}{\sqrt{2}\sigma}\right). \tag{32}$$

Because of spherical symmetry, the MOND solution must satisfy the algebraic relation $g_M^{\rm ana} = \nu(g_N^{\rm ana}/a_0)\,g_N^{\rm ana}$. Figure 5 compares numerical and analytic results. Outside the softened region ($r > 3\sigma$), the Newtonian field agrees with the analytic curve to better than 1%, and the MOND solution coincides with the algebraic prediction, confirming the spherical limit.

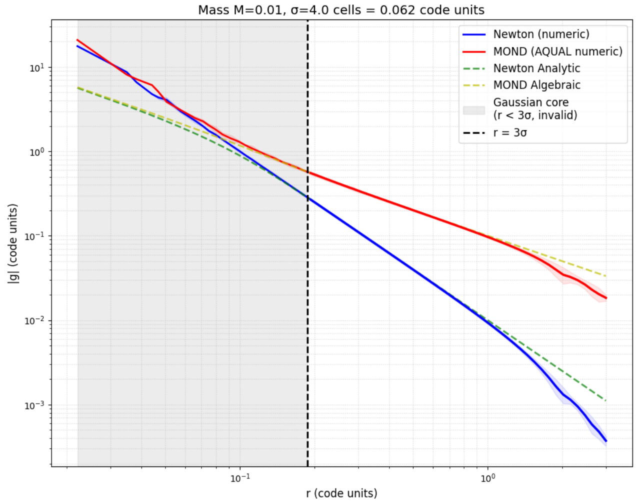

Figure 5: Validation of the AQUAL solver for a spherically symmetric Gaussian-broadened point mass. The mass distribution is given by $\rho(r) = M(2\pi\sigma^2)^{-3/2}\exp\left[-r^2/(2\sigma^2)\right]$ with total mass $M = 0.01$ in code units ($G = a_0 = 1$), softening length $\sigma = 0.062$ (4 cells), and box size $L = 4$ on a 256^3 grid. The numerical Newtonian acceleration g_N (blue solid line) and the full AQUAL solution g_M (red solid line) are compared against their analytic counterparts: $g_N^{\rm ana}$ (blue dashed) from Eq. (32) and the algebraic MOND prediction $g_M^{\rm alg} = \nu(g_N^{\rm ana}/a_0)g_N^{\rm ana}$ (red dashed) from Eq. (16). At $r < 3\sigma$, the Gaussian softening suppresses the divergence of the point-mass potential; this region is excluded from the analysis.

3.5.2 Exponential disk test (flattened geometry)

Next, we consider a thin exponential disk with scale length $R_d = 0.01$ and Gaussian vertical thickness $z_0 = 0.5$ cells. The density is

$$\rho(R, z) = \frac{M}{4\pi R_d^2 z_0} e^{-R/R_d} e^{-z^2/(2z_0^2)}, \tag{33}$$

with total mass $M = 0.01$.
For an infinitely thin exponential disk, the Newtonian midplane acceleration is given by the Freeman [23] solution:

$$g_N^{\rm Freeman}(R) = \frac{v_c^2(R)}{R}, \tag{34}$$

$$v_c^2(R) = 4\pi G\Sigma_0 R_d\, y^2\left[I_0(y)K_0(y) - I_1(y)K_1(y)\right], \tag{35}$$

where $y = R/(2R_d)$ and $\Sigma_0 = M/(2\pi R_d^2)$. Here, I_0 and I_1 are the modified Bessel functions of the first kind of order 0 and 1, respectively, while K_0 and K_1 are the modified Bessel functions of the second kind of order 0 and 1, respectively.
The numerical solution for the Newtonian acceleration closely follows the analytic Freeman profile. The full AQUAL solution and the algebraic MOND estimate are nearly indistinguishable, confirming earlier the findings that non-spherical corrections are small for typical disk galaxies [12, 18]. Error plots are displayed in Fig. 7.

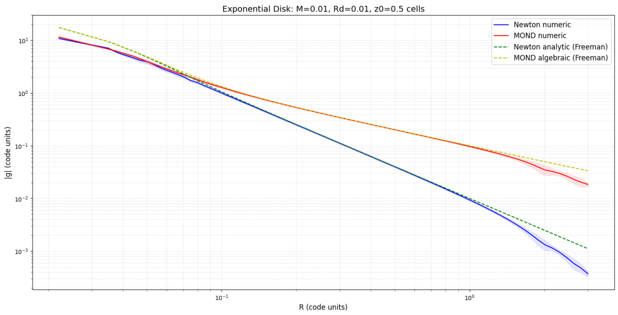

Figure 6: Validation of the AQUAL solver for a flattened exponential disk. The mass distribution follows $\rho(R, z) = (M/4\pi R_d^2 z_0)\exp(-R/R_d)\exp\left[-z^2/(2z_0^2)\right]$ with total mass $M = 0.01$, scale length $R_d = 0.01$, vertical thickness $z_0 = 0.5$ cells, and box size $L = 4$ on a 256^3 grid (code units: $G = a_0 = 1$). The numerical Newtonian acceleration g_N (blue solid) is compared against the analytic Freeman solution for an infinitesimally thin exponential disk (blue dashed, Eq. (34)), showing excellent agreement where the finite vertical thickness does not significantly affect the midplane gravity. The full AQUAL solution g_M (red solid) is nearly indistinguishable from the algebraic MOND approximation $g_M^{\rm alg} = \nu(g_N/a_0)g_N$ (yellow dashed. This confirms earlier findings [12, 18] that the curl-free algebraic approximation accurately reproduces the full QuMOND/AQUAL solution for typical disk geometries, despite the non-spherical mass distribution.

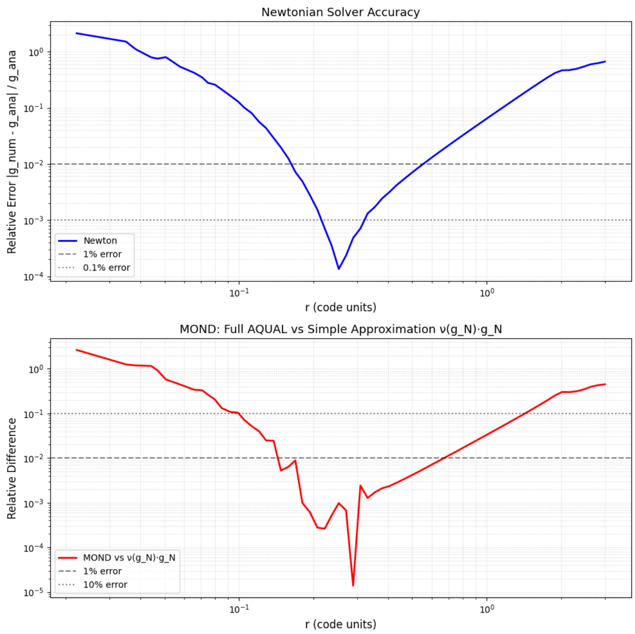

Figure 7: Relative errors in the numerical solutions for the exponential disk test (same setup as Figure 6: $M = 0.01$, $R_d = 0.01$, $z_0 = 0.5$ cells, 256^3 grid, code units $G = a_0 = 1$, box size $L = 4$). Top panel: Relative error of the numerical Newtonian solver, defined as $|g_N^{\rm num} - g_N^{\rm Freeman}|/g_N^{\rm Freeman}$, where $g_N^{\rm Freeman}$ is the analytic Freeman solution for an infinitesimally thin exponential disk (Eq. (34)). The error remains below 10% across the full radial range and is typically $< 1\%$ for $0.15 \lesssim r \lesssim 0.7$ (code units). At small radii ($r \lesssim 0.15$), the error increases due to limited grid resolution: the disk's central region is under-resolved, introducing discretisation errors in the density field and its corresponding potential. At large radii ($r \gtrsim 0.7$), the error rises again due to the finite box size: the FFT solver's periodic boundary conditions introduce spurious forces from image masses, and the mass distribution is no longer well-contained within the computational domain. Bottom panel: Relative difference between the full AQUAL solution and the algebraic MOND approximation, defined as $|g_M^{\rm AQUAL} - g_M^{\rm alg}|/g_M^{\rm alg}$, where $g_M^{\rm alg} = \nu(g_N^{\rm num}/a_0)g_N^{\rm num}$ is the algebraic prediction. The difference is typically $< 1\%$ across the disk, confirming that the algebraic approximation accurately reproduces the full AQUAL solution for exponential disk geometries, consistent with earlier findings [12, 18]. The slight increase at the largest radii reflects the growing influence of the curl term in the QuMOND equation, though even there the discrepancy remains well below 10%. The grey shaded region indicates radii where the Newtonian solver error exceeds 10% and the numerical solution is considered unreliable.

3.5.3 Convergence and limitations

All simulations use a cubic domain of side $L = 4$ with 256^3 cells, ensuring the mass distribution is well contained ($R_d < \frac{L}{2}$). The Gaussian softening ($\sigma = 0.125$) avoids aliasing while resolving the transition between Newtonian and deep-MOND regimes ($g_N/a_0 \sim 10$ down to 10^{-3}).
Rather than adopting an infinitely thin disk, a finite vertical thickness of 0.5 cells is used to maintain numerical smoothness. The excellent agreement with analytic expectations demonstrates that the solver reliably computes MOND fields for both spherical and flattened mass distributions.

3.6 Validation of the Spherical Approximation: Uniform Sphere Test

As a first controlled test of the spherical approximation, we consider a mass distribution for which the relation

$$g_N(r) = \frac{GM(<r)}{r^2} \tag{36}$$

is expected to hold exactly, namely a spherically symmetric system with uniform density.

3.6.1 Numerical setup

We construct a three–dimensional density field on a Cartesian grid representing a sphere of radius R and total mass M. The density is defined as

$$\rho(\mathbf{r}) = \begin{cases} \rho_0 & \text{if } r \leq R \\ 0 & \text{otherwise} \end{cases} \tag{37}$$

where the constant density ρ_0 is determined by normalisation to the total mass,

$$\rho_0 = \frac{M}{V_{\text{sphere}}}, V_{\text{sphere}} = \frac{4\pi R^3}{3}. \tag{38}$$

The Newtonian potential ϕ_N is obtained by solving the Poisson equation

$$\nabla^2 \phi_N = 4\pi G \rho \tag{39}$$

using a Fast Fourier Transform (FFT) method. The corresponding acceleration field is computed as

$$\mathbf{g}_N = -\nabla \phi_N. \tag{40}$$

3.6.2 Spherical averaging procedure

To compare the numerical solution with the enclosed-mass approximation, we compute the spherically averaged acceleration profile. The enclosed mass is obtained by binning the density field in radial shells and cumulatively summing:

$$M(<r_i) = \sum_{r<r_i} \rho(\mathbf{r}) \, \Delta V, \tag{41}$$

where $\Delta V = \sigma^3$ is the volume of a grid cell. The corresponding spherical approximation is then

$$g_{\text{sph}}(r) = \frac{GM(<r)}{r^2}. \tag{42}$$

In addition, the radial component of the numerical acceleration field is computed via projection onto the radial unit vector,

$$g_{N,\text{rad}}(\mathbf{r}) = -\mathbf{g}_N \cdot \hat{\mathbf{r}}. \tag{43}$$

This quantity is subsequently averaged in radial bins to obtain a one–dimensional profile $g_N(r)$.

3.6.3 Analytic expectation

For a uniform sphere, the Newtonian acceleration is known analytically:

$$g_{\text{analytic}}(r) = \begin{cases} \frac{GMr}{R^3} & r \leq R \\ \frac{GM}{r^2} & r > R. \end{cases} \tag{44}$$

This provides a direct benchmark for both the numerical solver and the spherical averaging procedure.

3.7 Results

As a first controlled test of the numerical solver, we consider a uniform density sphere of radius $R = 30\,\mathrm{kpc}$ and total mass $M = 10^{10}\,M_\odot$, for which the Newtonian acceleration profile is known analytically (Eq. (44)). This provides a clean benchmark before proceeding to more complex mass distributions.

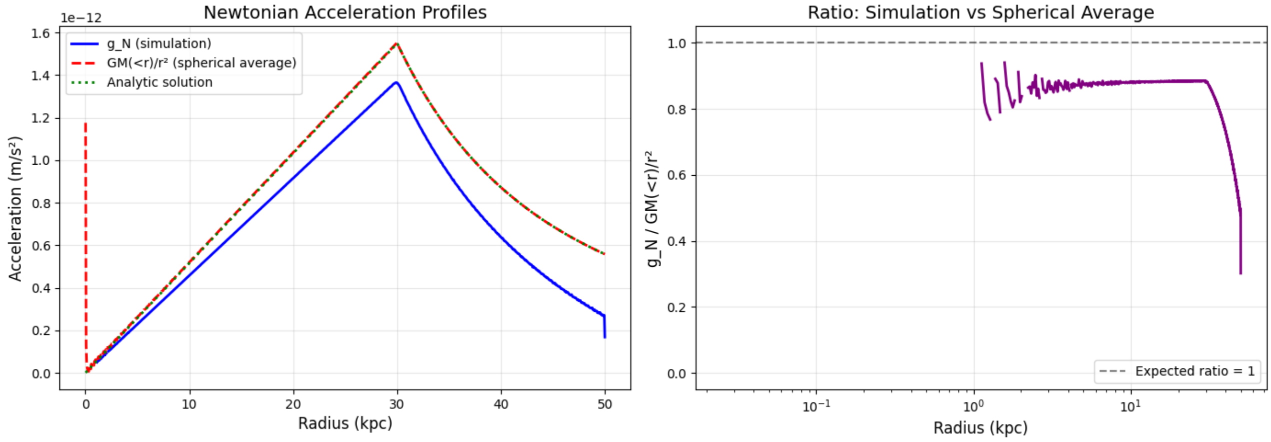

Figure 8: Validation of the Newtonian solver for a uniform density sphere, demonstrating the effect of periodic boundary conditions. The sphere has radius $R = 30\,\mathrm{kpc}$ and total mass $M = 10^{10}\,M_\odot$, discretised on a 256^3 grid with box size $L = 100\,\mathrm{kpc}$. Left panel: Radial acceleration profiles comparing the numerical FFT-based Poisson solver g_N (blue solid), the spherical enclosed-mass approximation $g_{\mathrm{sph}} = GM(< r)/r^2$ (red dashed), and the analytic solution (green dotted, Eq. (44)). The spherical approximation and analytic solution agree closely throughout, as expected for a sphere. However, the numerical solver systematically underestimates the true acceleration by 10–15% across all radii. Right panel: Ratio g_N/g_{sph} as a function of radius. The deviation from unity arises from the FFT solver's implicit periodic boundary conditions: the density field is tiled across the simulation volume, and the gravitational field includes spurious contributions from periodic image masses outside the box. These image masses pollute the solution at all radii, producing the $\sim$ 10–15% offset. The sharp drop at $r \gtrsim 30\,\mathrm{kpc}$ coincides with the sphere boundary and marks the onset of strong boundary contamination. The innermost region ($r \lesssim 1\,\mathrm{kpc}$) shows discretisation noise due to the small number of grid cells per radial shell; this is a purely numerical artifact independent of the solver accuracy. This test reveals that naive FFT-based Poisson solvers with periodic boundaries are inadequate for isolated mass distributions without additional mitigation strategies.

The numerical Newtonian acceleration $g_N(r)$, obtained by radially binning the projection of the FFT-based acceleration field onto the radial unit vector, is compared against the spherical approximation $g_{\mathrm{sph}}(r) = GM(<r)/r^2$ and the analytic solution. The results are shown in Figure 8.

Outside the innermost region, g_{sph} and the analytic solution agree to better than 1%, confirming that the enclosed-mass binning procedure correctly recovers the cumulative mass profile. However, the numerical solver g_N systematically underestimates the acceleration by approximately 10–15% across all radii, as visible in the ratio panel on the right of Figure 8.

Two limitations are identified. First, at small radii ($r \lesssim 1\,\mathrm{kpc}$), the spherical approximation shows noise and non-monotonic behaviour caused by discretisation: near the centre, each radial shell contains very few grid cells, and the irregular sampling of the cubic grid causes the enclosed mass to grow unevenly between bins. This is a purely numerical artifact that diminishes with higher resolution. Second, and more significantly, the FFT-based Poisson solver assumes *periodic* boundary conditions at 50 kpc, which means that the density field is implicitly tiled across the simulation volume. The gravitational field therefore includes contributions from periodic image masses outside the box, which pollute the solution at large radii and suppress the radial acceleration relative to the isolated-sphere expectation. This effect manifests as the sharp drop in g_N near $r \approx 50\,\mathrm{kpc}$ and the systematic $\sim$10% offset visible at intermediate radii.

4 Optimizing the Poisson Solver for Disk-shaped galaxy

The solver in chapter 2 seems to be a start when solving for MOND galaxy in the medium radial range. However, for the radial end range there is quite an error. Also, as mentioned before, a spherical assumptions has been made: constructing $\vec{g}_N$ from a radial profile.

In this chapter I will quantify how to improve a numerical Poisson equation solver and explore the removel of the radial profile.

In the remainder of this report, we restrict our analysis to the mid-plane of the galaxy, i.e., $z = 0$. This is the region where the disk mass distribution is most concentrated and where observational data for the Radial Acceleration Relation are typically reported. By focusing on the plane $z = 0$, we can directly compare our numerical results to both theoretical predictions and observational constraints. All subsequent scatter plots and analyses therefore consider only grid cells lying in this plane.

4.1 Zero-padding correction

The FFT-based Poisson solver implicitly assumes periodic boundary conditions: the density distribution is treated as repeating infinitely in all three spatial dimensions. Consequently, the gravitational potential at any point receives spurious contributions from the periodic images of the true mass distribution. For a finite isolated system such as a galaxy, these image masses produce an unphysical additional force that can significantly distort the solution, particularly near the boundaries of the computational domain.

One might consider two alternative strategies to mitigate this artifact. First, simply increasing the grid resolution N while keeping the physical box size fixed does not address the periodicity problem; it only refines the sampling of the same periodic volume. The image masses remain at the same distance, and their spurious contribution is unchanged. Second, one could manually extend the density grid by adding zeros in all directions without changing the solver strategy — i.e., embedding the galaxy in a larger volume of empty space. While this is conceptually similar to zero-padding, a naive implementation would require rewriting the density array generation and modifying the boundary logic. More importantly, without using the FFT's native ability to handle power-of-two grid sizes efficiently, such manual extension may lead to suboptimal performance or force the use of non-power-of-two transforms, which are slower.

Zero-padding offers a clean, systematic alternative. Instead of increasing N or manually appending zeros, the density grid is zero-padded to twice its linear size in each dimension before solving the Poisson equation. The padded grid of size $2N \times 2N \times 2N$ is passed to the FFT solver, and the resulting potential is cropped back to the original $N \times N \times N$ domain. This effectively doubles the separation between the physical mass distribution and its nearest periodic image, substantially reducing the spurious contribution to the gravitational field within the region of interest. The box size is also doubled. This approach leverages the efficiency of FFTs on power-of-two grids (since $2N$ remains a power of two if N is) while requiring no changes to the underlying density generation code other than embedding the existing grid in a larger array of zeros. All this requires significantly more memory — a factor of eight increase in grid cells — but by adjusting the code to clean up memory after using certain elements, it can be made to work within practical limits.

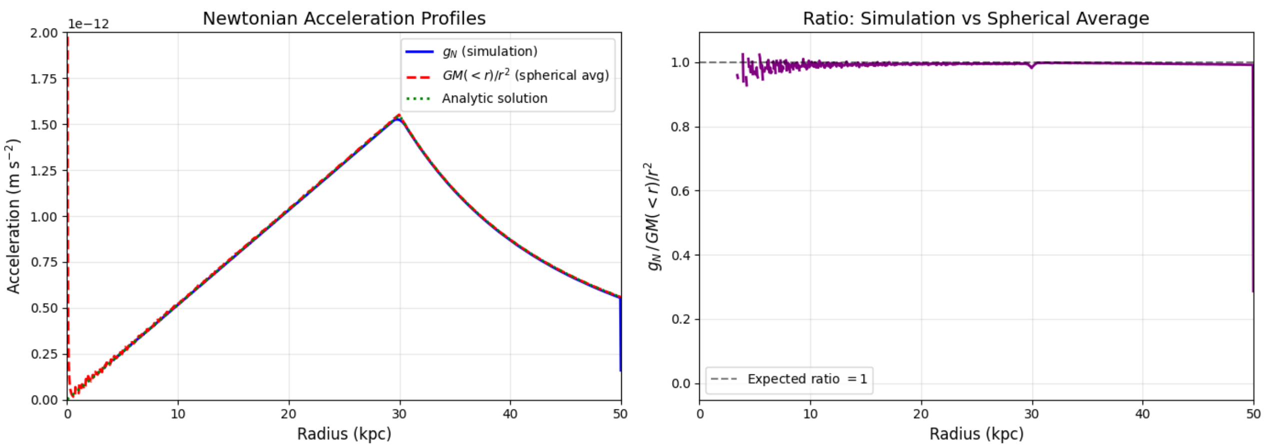

Figure 9: Validation of the Newtonian solver for a uniform density sphere after applying zero-padding to mitigate periodic boundary artefacts. The sphere has radius $R = 30\,\mathrm{kpc}$ and total mass $M = 10^{10}\,M_\odot$, embedded in a 256^3 grid that is zero-padded to 512^3 prior to the FFT Poisson solve, effectively doubling the box size to $L = 200\,\mathrm{kpc}$ and increasing the separation between the physical mass distribution and its nearest periodic images. Left panel: The numerical solver g_N (blue solid) is now nearly indistinguishable from both the spherical enclosed-mass approximation $g_{\rm sph} = GM(<r)/r^2$ (red dashed) and the analytic solution (green dotted, Eq. (44)) across the full domain. Right panel: The ratio $g_N/g_{\rm sph}$ lies within 1% of unity over the range $5\,\mathrm{kpc} \lesssim r \lesssim 49.7\,\mathrm{kpc}$, representing a substantial improvement over the unpadded case (Figure 8). Two residual deviations remain: (i) At small radii ($r \lesssim 5\,\mathrm{kpc}$), the ratio is noisy due to the limited number of grid cells per radial shell at finite resolution; this discretisation artefact is inherent to the cubic grid and independent of the solver accuracy. The MOND regime of interest occurs at low accelerations in the outer disk, so this inner-region noise does not affect our scientific conclusions. (ii) At large radii ($r \gtrsim 49.7\,\mathrm{kpc}$), a sharp drop reflects the residual influence of periodic image masses at the boundary of the padded domain. Results are therefore considered reliable within $5\,\mathrm{kpc} \lesssim r \lesssim 49.7\,\mathrm{kpc}$, which comfortably covers the physically relevant range of the mass distribution. The factor-of-eight increase in memory usage required for zero-padding is justified by the dramatic improvement in accuracy, confirming that zero-padding is an essential step for FFT-based Poisson solvers applied to isolated systems.

The result is shown in Figure 9. The ratio $g_N/g_{\rm sph}$ now sits within 1% of unity over the range $5\,\text{kpc} \lesssim r \lesssim 49.7\,\text{kpc}$, a substantial improvement over the unpadded case. Two residual deviations remain. At small radii ($r \lesssim 5\,\text{kpc}$), the ratio is noisy due to the limited number of grid cells per radial shell at this resolution; this is a discretisation artifact inherent to the cubic grid and is independent of the solver accuracy. Also, since the MOND-regime we are interested in is always at the edge of a galaxy it is not inherently relevant for our study. At large radii ($r \gtrsim 49.7\,\text{kpc}$), a sharp drop reflects the residual influence of periodic image masses at the boundary of the padded domain. Results are therefore considered reliable within $5\,\text{kpc} \lesssim r \lesssim 49.7\,\text{kpc}$, which comfortably covers the physically relevant range of the mass distribution. This level of agreement is sufficient to validate the Newtonian solver before proceeding to the full MOND solution.

4.2 From Binned to Scattered Validation: The Need for Unbinned Comparisons

The preceding validation test used a spherical mass distribution and binned the radial acceleration data. This approach is valid for spherically symmetric systems because, at a given radius r, every point in the spherical shell experiences the same gravitational acceleration. Binning and averaging over r therefore preserves the full information and does not conceal any underlying errors.

However, this situation changes dramatically when we move beyond spherical symmetry. For a non-spherical mass distribution — such as an axisymmetric disk or a triaxial halo — the gravitational acceleration at a fixed radial distance r is no longer unique. Points at the same r but different angular positions experience different values of g_N due to the geometry of the mass distribution. In such cases, binning by radius and averaging would wash out this angular variation, potentially hiding systematic errors or biases in the solver. A simple average over r might accidentally cancel positive and negative deviations, giving a false impression of accuracy.

Therefore, for non-spherical geometries, we must switch to an unbinned validation approach: comparing the numerical solver's output to the analytic (or reference) solution on a point-by-point basis, presented as a scatter plot. Each point in the computational domain contributes one pair of values $(g_{\rm sph}, g_{\rm num})$, and any systematic offset or scatter becomes immediately visible. This method does not assume any symmetry and reveals the full error distribution, including outliers and angular dependencies.

Figure 10 shows the result of applying this unbinned validation to the same spherical test case. While the binned result in Figure 9 suggested near-perfect agreement across most radii, the scatter plot reveals two important features that were previously hidden.

First, a noticeable "bump" appears at $r \approx 30$ kpc in both figures, the edge of the spherical mass distribution. At this radius, the density drops sharply to zero, creating a steep gradient in the potential. The FFT solver, which assumes smooth periodic functions, struggles to represent this discontinuity perfectly, leading to a small but systematic deviation visible as a clustering of points offset from the line of perfect agreement.

Second, at small radii ($r \lesssim 10$ kpc, the scatter increases somewhat. This is purely a numerical discretization artifact: the inner radial shells contain very few grid cells. For a Cartesian grid, the number of cells at a given radius scales roughly as r^2, so at small r, each spherical shell may contain only a handful of points. With such limited sampling, the computed acceleration varies significantly from cell to cell due to the discrete representation of the density field, producing the observed spread. This scatter does not indicate a fundamental problem with the solver; it simply reflects the finite resolution of the grid.

Importantly, neither of these features compromises our scientific goals. The scatter at small radii affects the innermost region of the galaxy, while our MOND analysis focuses on the outer, low-acceleration regime. The bump at $r \approx 30$ kpc is small in magnitude and caused by a sudden overlap in density which is not physical. Furthermore, at the largest radii shown ($r \approx 50\,\text{kpc}$), the error remains only about 1%, consistent with the binned result. The solver is therefore sufficiently accurate for our purposes.

4.3 Asymmetric Test

Before applying our solver to realistic disk geometries, we must verify that it performs correctly for non-spherical, asymmetric mass distributions. The spherical test case validated the solver for a symmetric configuration where the analytic solution is known. However, real galaxies are neither spherical nor centred at the coordinate origin relative to the computational domain. To test the solver's behaviour under more general conditions, we construct a simple asymmetric system: two Gaussian spheres of different widths and masses, placed at distinct positions offset from the origin.

4.3.1 Analytic Reference Solution

For a spherically symmetric Gaussian density profile

$$\rho(r) = A\exp\left(-\frac{r^2}{2\sigma^2}\right), A = \frac{M}{(2\pi)^{\frac{3}{2}}\sigma^3} \tag{45}$$

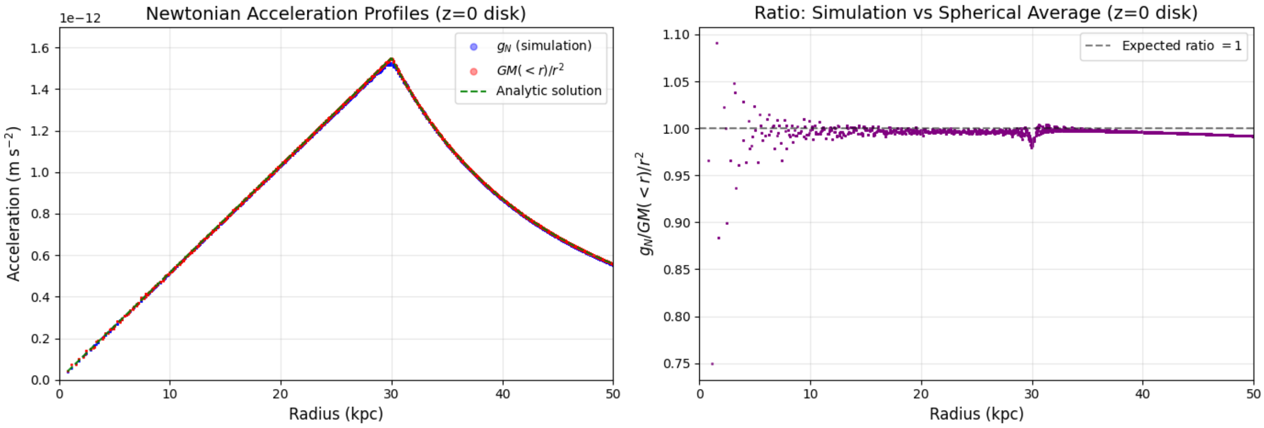

Figure 10: Unbinned (scatter) validation of the Newtonian solver for a uniform density sphere after applying zero-padding, revealing features hidden by radial binning. The sphere has radius $R = 30\,\mathrm{kpc}$ and total mass $M = 10^{10}\,M_\odot$, with the density field zero-padded to 512^3 prior to the FFT Poisson solve (same setup as Figure 9). Each point in the scatter plot represents a single grid cell in the computational domain, plotting the numerical radial acceleration g_N against the analytic spherical solution g_{analytic} (Eq. (44)). The solid red line indicates perfect agreement ($g_N = g_{\mathrm{analytic}}$), while the dashed green lines show 1% and 10% deviation bounds. Unlike the binned validation in Figure 9, which suggested near-perfect agreement across most radii, the scatter plot reveals the full distribution of errors. Two features become apparent: (i) A noticeable "bump" or clustering of points offset from the line of perfect agreement appears near $r \approx 30\,\mathrm{kpc}$, corresponding to the edge of the spherical mass distribution. At this radius, the density drops sharply to zero, creating a steep gradient in the potential. The FFT solver, which assumes smooth periodic functions, struggles to represent this discontinuity perfectly, leading to a small but systematic deviation visible as a clustering of points offset from the line of perfect agreement. (ii) At small radii ($r \lesssim 10\,\mathrm{kpc}$), the scatter increases significantly. This is a numerical discretisation artifact caused by the small number of grid cells per radial shell at these inner radii; for a Cartesian grid, the number of cells at a given radius scales roughly as r^2, so at small r each spherical shell may contain only a handful of points. With such limited sampling, the computed acceleration varies significantly from cell to cell due to the discrete representation of the density field, producing the observed spread. This scatter does not indicate a fundamental problem with the solver; it simply reflects the finite resolution of the grid. Importantly, neither of these features compromises our scientific goals: the scatter at small radii affects the innermost region of the galaxy, while our MOND analysis focuses on the outer, low-acceleration regime; and the bump at $r \approx 30\,\mathrm{kpc}$ is small in magnitude and caused by a sharp density discontinuity that is not present in realistic galaxy models. At the largest radii shown ($r \approx 50\,\mathrm{kpc}$), the error remains only about 1%, consistent with the binned result. The scatter plot demonstrates that while binning can hide angular variations and discretisation artefacts, the solver is sufficiently accurate for our purposes when the full error distribution is examined. This unbinned validation approach is essential for non-spherical geometries, where points at the same radius experience different accelerations, and binning would wash out systematic errors.

where A is determined by the normalisation condition $\int \rho\, \mathrm{d}^3 r = M$, the enclosed mass within radius r is

$$M(<r) = M\left[\mathrm{erf}\left(\frac{r}{\sqrt{2}\,\sigma}\right) - \sqrt{\frac{2}{\pi}}\,\frac{r}{\sigma}\,\exp\left(-\frac{r^2}{2\sigma^2}\right)\right]. \tag{46}$$

The radial gravitational acceleration follows from Newton's shell theorem:

$$g(r) = \frac{G\,M(<r)}{r^2}, \tag{47}$$

and the full acceleration vector, directed inward toward the sphere's centre $\mathbf{r}_\mathrm{c}$, is

$$\mathbf{g}(\mathbf{r}) = -g(|\mathbf{r}-\mathbf{r}_\mathrm{c}|)\,\frac{\mathbf{r}-\mathbf{r}_\mathrm{c}}{|\mathbf{r}-\mathbf{r}_\mathrm{c}|}. \tag{48}$$

Because the Poisson equation is linear, the total acceleration for a system of non-overlapping spheres is the vector superposition of the individual contributions:

$$\mathbf{g}_{\mathrm{total}}(\mathbf{r}) = \mathbf{g}_1(\mathbf{r}) + \mathbf{g}_2(\mathbf{r}) + \cdots \tag{49}$$

This provides a straightforward analytic reference against which we compare our numerical solver.
A Gaussian profile is preferred over a uniform (top-hat) sphere for this validation. A hard-edged density distribution introduces a discontinuity that is represented in Fourier space by a slowly decaying series, producing Gibbs ringing in the

reconstructed potential and its gradient. The Gaussian profile, being analytic and rapidly decaying, has a Fourier transform that it is also a Gaussian and decays exponentially, so the spectral solver can represent it to near machine precision at any resolved scale. The smooth profile therefore provides a much cleaner test of the solver's accuracy.

4.3.2 Numerical Setup

We adopt a two-sphere configuration with parameters chosen to produce a distinctly asymmetric mass distribution. Sphere 1 (the smaller and less massive sphere) is placed to the left of the origin with a 1σ radius $\sigma_1 = 10.0$ kpc, total mass $M_1 = 5 \times 10^8\, M_\odot$, and centre at $x_1 = -25.0$ kpc. Sphere 2 (the larger and more massive sphere) is placed to the right of the origin with $\sigma_2 = 8.0$ kpc, $M_2 = 1 \times 10^{10}\, M_\odot$, and centre at $x_2 = +25.0$ kpc.
The computational domain is a cube of side length $L_{\rm box} = 200$ kpc, discretised on a grid of $N = 256$ cells per dimension, giving a cell spacing of $\Delta x = L_{\rm box}/N \approx 0.781$ kpc. The density field is constructed by evaluating the Gaussian profile of each sphere at every grid-cell centre and summing the two contributions. Each profile is normalised so that the discrete sum over all cells exactly recovers the input mass M, regardless of resolution:

$$\rho_{\rm num} = \frac{M}{N_{\rm cells}\, \Delta x^3}, \tag{50}$$

where $N_{\rm cells}$ is the effective number of cells weighted by the Gaussian envelope.

4.4 Spectral Differentiation: From Finite Differences to Fourier Gradients

In an initial implementation, the gravitational acceleration was obtained from the potential Φ using a second-order centred finite-difference scheme:

$$g_x(x, y, z) = -\frac{\Phi(x + \Delta x, y, z) - \Phi(x - \Delta x, y, z)}{2\Delta x}, \tag{51}$$

with analogous expressions for g_y and g_z. While straightforward, this approach has several drawbacks. It is a local approximation, accurate only to order $O(\Delta x^2)$, and therefore susceptible to discretisation errors near regions where the potential has significant curvature. It also requires a separate post-processing step after the potential has been computed.
A more accurate approach is to compute the acceleration directly in Fourier space. Since the Poisson equation in Fourier space reads

$$\hat{\Phi}(\mathbf{k}) = -\frac{4\pi G\, \hat{\rho}(\mathbf{k})}{k^2}, \tag{52}$$

and differentiation in real space corresponds to multiplication by $i\mathbf{k}$ in Fourier space, the acceleration components are obtained without any additional approximation:

$$\hat{g}_x = -ik_x\hat{\Phi}, \qquad \hat{g}_y = -ik_y\hat{\Phi}, \qquad \hat{g}_z = -ik_z\hat{\Phi}. \tag{53}$$

Each component is then recovered by a single inverse FFT:

$$g_x(\mathbf{x}) = \mathcal{F}^{-1}\left[-ik_x\hat{\Phi}(\mathbf{k})\right], \quad g_y(\mathbf{x}) = \mathcal{F}^{-1}\left[-ik_y\hat{\Phi}(\mathbf{k})\right], \quad g_z(\mathbf{x}) = \mathcal{F}^{-1}\left[-ik_z\hat{\Phi}(\mathbf{k})\right]. \tag{54}$$

This spectral differentiation is exact on the discrete grid — to within floating-point precision — and requires no additional post-processing beyond the inverse transforms already needed for the potential.

4.4.1 Comparison

Figure 11 shows the result when the finite-difference gradient is used. The ratio $g_N/g_{\rm analytic}$ deviates from unity by up to $\pm 20\%$ across the domain, with the largest errors occurring at intermediate radii where the contributions of the two off-centre spheres produce the steepest gradients in the potential. These errors are a direct consequence of the $O(\Delta x^2)$ truncation error of the finite-difference stencil. Note the difference in scale on the y-axis. It may look as though the spread of figure 11 is less but this is simply because there are larger outliers.

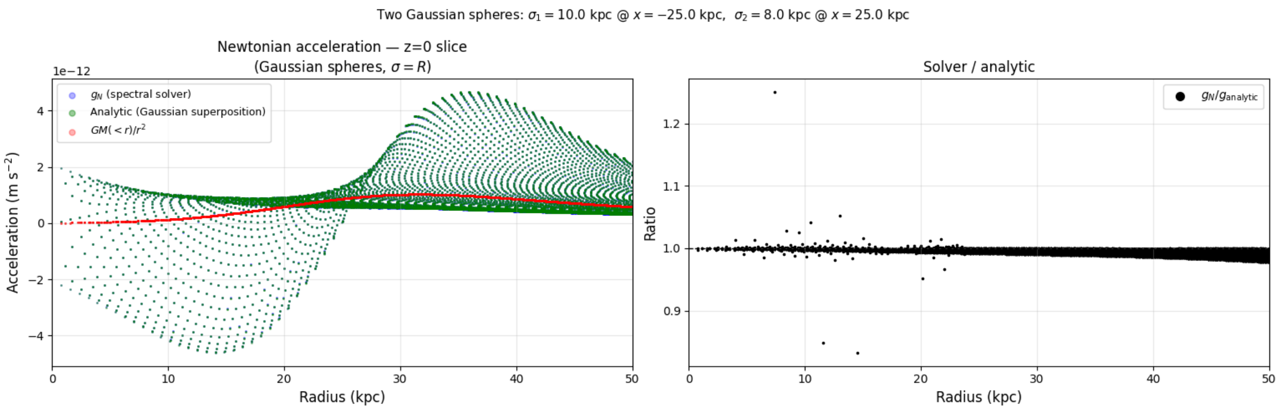

Figure 11: Validation of the Newtonian solver for an asymmetric two-Gaussian-sphere configuration using a second-order centred finite-difference gradient, revealing significant truncation errors. The mass distribution consists of two off-centre Gaussian spheres: Sphere 1 (smaller, less massive: $\sigma_1 = 10.0$ kpc, $M_1 = 5 \times 10^8\,M_\odot$, centred at $x_1 = -25.0$ kpc) and Sphere 2 (larger, more massive: $\sigma_2 = 8.0$ kpc, $M_2 = 1 \times 10^{10}\,M_\odot$, centred at $x_2 = +25.0$ kpc). The computational domain is a cube of side $L = 200$ kpc, discretised on a 256^3 grid ($\Delta x \approx 0.781$ kpc). The Newtonian potential is obtained via the FFT-based Poisson solver with zero-padding, but the acceleration is computed using a second-order centred finite-difference stencil: $g_x = -[\Phi(x + \Delta x) - \Phi(x - \Delta x)]/(2\Delta x)$ (Eq. (51)). Left panel: Radial acceleration profiles in the $z = 0$ plane for the numerical solver (blue), the analytic vector superposition of the two Gaussian spheres (green, Eq. (49)), and the spherical enclosed-mass estimate $GM(< r)/r^2$ (red). The numerical solution deviates significantly from the analytic reference across the domain. Right panel: Ratio $g_N/g_{\rm analytic}$ as a function of radius. Deviations of up to $\pm 20\%$ are visible, with the largest errors occurring at intermediate radii where the steepest gradients in the potential arise from the superposition of the two off-centre Gaussian profiles. These errors are a direct consequence of the $O(\Delta x^2)$ truncation error of the finite-difference stencil: the potential has significant curvature due to the asymmetric mass distribution, and the local finite-difference approximation fails to capture the gradient accurately. The strong scatter and systematic offsets confirm that finite-difference gradients are unsuitable for use with FFT-based Poisson solvers when high accuracy is required, particularly for non-spherical mass distributions with steep potential gradients. The vertical dashed lines indicate the centres of the two Gaussian spheres; the point $x = 0$ lies between them, where the potential gradient is strongest and the finite-difference error is largest.

Figure 12 shows the result after replacing the finite-difference gradient with spectral differentiation. The ratio $g_N/g_{\rm analytic}$ now remains within $\pm 1.5\%$ across the entire domain, with the residual error arising from the finite resolution of the Gaussian profiles rather than from any truncation error in the gradient calculation. This confirms that the spectral gradient is the correct choice for use alongside the FFT-based Poisson solver. This results in both less outliers and a more accurate result.

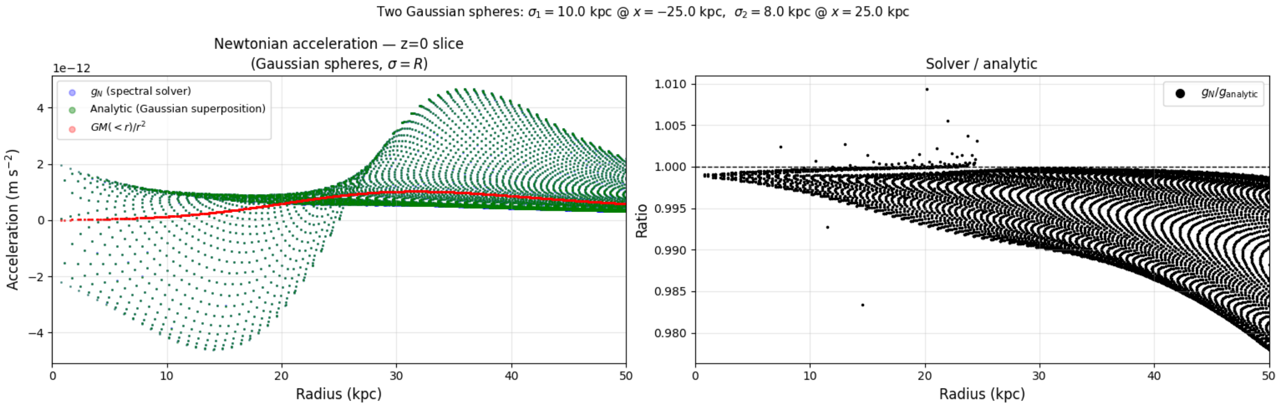

Figure 12: Validation of the Newtonian solver for an asymmetric two-Gaussian-sphere configuration using spectral differentiation, achieving near-machine-precision accuracy. The mass distribution is identical to Figure 8: two off-centre Gaussian spheres with Sphere 1 ($\sigma_1 = 10.0$ kpc, $M_1 = 5 \times 10^8\, M_\odot$, $x_1 = -25.0$ kpc) and Sphere 2 ($\sigma_2 = 8.0$ kpc, $M_2 = 1 \times 10^{10}\, M_\odot$, $x_2 = +25.0$ kpc), on a 256^3 grid with box size $L = 200$ kpc and zero-padding applied. Unlike Figure 8, the acceleration is now computed directly in Fourier space via spectral differentiation: $\hat{g}_x = -ik_x\hat{\Phi}$, $\hat{g}_y = -ik_y\hat{\Phi}$, $\hat{g}_z = -ik_z\hat{\Phi}$ (Eqs. (53)–(54)), which is exact on the discrete grid to within floating-point precision. Left panel: The numerical solver g_N (blue) now agrees almost perfectly with the analytic vector superposition (green), while the spherical enclosed-mass estimate $GM(<r)/r^2$ (red) deviates significantly due to the non-spherical mass distribution. Right panel: The ratio $g_N/g_{\mathrm{analytic}}$ now remains within ±1.5% across the entire domain, a dramatic improvement over the ±20% errors in Figure 8. The residual error arises from two sources: (i) the finite grid resolution ($\Delta x \approx 0.781$ kpc) limits the accuracy with which the Gaussian profiles are represented on the discrete grid, and (ii) the periodic boundary conditions of the FFT Poisson solver introduce small spurious contributions from image masses, even with zero-padding. Critically, no error is attributable to the gradient calculation itself, because spectral differentiation is exact on the discrete grid. The $O(\Delta x^2)$ truncation error of the finite-difference stencil is eliminated, resulting in a solver that is accurate to better than 2% for this asymmetric test case. This confirms that spectral differentiation is the correct choice for use alongside FFT-based Poisson solvers, particularly for non-spherical mass distributions where finite-difference gradients introduce unacceptably large systematic errors. The vertical dashed lines indicate the centres of the two Gaussian spheres.

5 Construction of a Stellar Mass Density Cube from Infrared Imaging

To study the gravitational field of a galaxy, a numerical pipeline was developed that converts a two-dimensional infrared image into a three-dimensional stellar mass density distribution. The input data consist of a FITS (Flexible Image Transport System) image at 3.6 μm from the Spitzer Space Telescope, taken from the Spitzer Survey of Stellar Structure in Galaxies (S4G; NASA/IPAC Infrared Science Archive 21). Emission at this wavelength traces the old stellar population and therefore provides a good approximation to the stellar mass distribution of the galaxy.

The pipeline performs several steps (described in detail below): image cleaning, geometric deprojection with correct pixel area conservation, conversion from infrared flux to surface mass density using the S4G standard conversion factor, optional bulge component addition, and construction of a three-dimensional density cube with a realistic vertical profile. Diagnostic plots are generated throughout the procedure to verify that the reconstructed density distribution behaves physically and conserves mass.

5.1 Loading and Cleaning the Image

The FITS image is read using the `astropy` library and converted into a two-dimensional numerical array representing the infrared surface brightness of the galaxy in units of MJy sr^{-1}.

Crucially, no background subtraction is performed. The S4G images are already calibrated to absolute surface brightness units (MJy sr^{-1}) with the background properly subtracted by the survey pipeline. Subtracting additional background would artificially reduce the stellar mass and distort the gravitational field.

If an external mask file is available (provided by S4G to flag foreground stars, artefacts, and cosmic rays), masked pixels are set to zero. Negative values produced by noise are clipped to zero, and extremely high values (cosmic rays) are clipped to the 99.9th percentile of the positive pixels. This produces a cleaned image in which the dominant signal originates from the stellar disk of the galaxy itself.

5.2 Geometric Deprojection

Observed galaxies are generally inclined relative to the line of sight. As a result, intrinsically circular stellar disks appear elliptical in projection. To reconstruct the approximate face-on stellar distribution, the image is geometrically deprojected. The deprojection is performed in two stages:

1. The image is rotated so that the major axis becomes horizontal, using the observed position angle of the galaxy.
2. The image is stretched along the minor axis by a factor $1/\cos i$, where i is the inclination angle of the disk.

An important correction is applied to the pixel area. When stretching the image, the physical area represented by each pixel changes. The correct pixel area after deprojection is

$$A_{\text{pixel}} = \frac{A_{\text{orig}}}{\cos i} = \frac{p_{\text{orig}}^2}{\cos i}, \tag{55}$$

where p_{orig} is the original pixel scale in metres. This ensures that the total flux (and therefore total stellar mass) is conserved during deprojection. The effective pixel scale after deprojection is $\sqrt{A_{\text{pixel}}}$.

5.3 Centre Finding

The galaxy centre is identified as the peak of the image after applying a heavy Gaussian smoothing ($\sigma = 20$ pixels). This smooths over small-scale structure (spiral arms, individual star-forming regions) to locate the overall photometric centre of the stellar disk.

5.4 Conversion to Stellar Surface Density

The Spitzer 3.6 μm images are calibrated in units of MJy sr^{-1}. Following standard S4G practice, we convert the 3.6 μm surface brightness into stellar surface density using

$$1 \text{ MJy sr}^{-1} = 20.9\, M_\odot \text{ pc}^{-2}. \tag{56}$$

For a general mass-to-light ratio (M/L), the stellar surface mass density is therefore

$$\Sigma(x, y) = I(x, y) \times 20.9 \times (M/L) \quad [M_\odot \text{ pc}^{-2}], \tag{57}$$

where $I(x, y)$ is the infrared surface brightness in MJy sr^{-1}. For use in the gravitational solver, this is converted to SI units (kg m^{-2}) using $1\, M_\odot \text{ pc}^{-2} = M_\odot/\text{pc}^2 \approx 2.09 \times 10^{-3}$ kg m^{-2}.

5.5 Radial Surface Density Profile

Although the full two-dimensional structure of the galaxy is preserved for the final density cube, an azimuthally averaged radial profile is also constructed as a diagnostic quantity.
For each cylindrical radius R, the mean stellar surface density is computed,

$$\Sigma(R) = \langle \Sigma(x, y) \rangle_{\text{ring}} = \frac{1}{2\pi} \int_0^{2\pi} \Sigma(R\cos\theta, R\sin\theta)\, d\theta. \tag{58}$$

This radial profile provides a simplified description of how the stellar mass is distributed across the galactic disk. The total stellar mass of the disk is then estimated by integrating the surface density profile,

$$M_\star = \int 2\pi R\, \Sigma(R)\, dR. \tag{59}$$

This quantity provides a useful consistency check against values reported in the literature. The half-mass radius R_{half}, defined by $M(< R_{\text{half}}) = M_\star/2$, is also computed for diagnostic purposes.

5.6 Bulge Component (Optional)

Many spiral galaxies, including NGC 3198, contain a significant central bulge component that is not well traced by the 3.6 μm disk emission alone. The pipeline therefore includes an optional bulge component modelled as a Plummer sphere (e.g., Plummer 24), with density profile

$$\rho_{\text{bulge}}(x, y, z) = \frac{3M_{\text{bulge}}}{4\pi a^3} \left(1 + \frac{x^2 + y^2 + z^2}{a^2}\right)^{-5/2}, \tag{60}$$

where M_{bulge} is the total bulge mass and a is the Plummer scale radius, these are fitted using the rotation curve of a galaxy. The bulge is added directly to the density cube in three dimensions, centred on the galaxy centre $(x, y, z) = (0, 0, 0)$.

5.7 Construction of the Three-Dimensional Density Cube

The final three-dimensional density distribution is constructed directly from the full two-dimensional surface density map rather than from the azimuthally averaged radial profile. This preserves non-axisymmetric structures such as spiral arms, bars, and rings.

The stellar disk is assumed to have a separable vertical structure. Rather than the simple Gaussian profile used in earlier work, we adopt a more realistic sech^2 vertical profile (van der Kruit and Searle 25), which better represents the vertical distribution of stars in disk galaxies:

$$\rho(x, y, z) = \frac{\Sigma(x, y)}{2z_0} \operatorname{sech}^2\left(\frac{z}{z_0}\right), \tag{61}$$

where z_0 is the vertical scale height of the stellar disk. The factor $1/(2z_0)$ ensures that $\int \rho(x, y, z)\, dz = \Sigma(x, y)$.

The pipeline also supports an alternative two-component *thin+thick disk* profile,

$$\rho(z) = \frac{f}{2z_{\mathrm{thin}}} \operatorname{sech}^2\left(\frac{z}{z_{\mathrm{thin}}}\right) + \frac{1-f}{2z_{\mathrm{thick}}} \operatorname{sech}^2\left(\frac{z}{z_{\mathrm{thick}}}\right), \tag{62}$$

where f is the fraction of mass in the thin disk ($f \approx 0.75$ typically), $z_{\mathrm{thin}} \approx 250$ pc, and $z_{\mathrm{thick}} \approx 900$ pc.

The two-dimensional surface density map is interpolated onto the (x, y) grid of the cube using bilinear interpolation. The resulting density distribution is discretized onto a cubic grid of size $N \times N \times N$ (typically $N = 256$), where each grid cell stores the local stellar mass density in units of $\mathrm{kg\,m^{-3}}$.

Mass conservation is then explicitly verified. The total mass in the final cube,

$$M_{\mathrm{cube}} = \sum_{i,j,k} \rho(x_i, y_j, z_k)\, \Delta x^3, \tag{63}$$

is compared to the total mass in the deprojected image,

$$M_{\mathrm{image}} = \sum_{\mathrm{pixels}} \Sigma(x, y)\, A_{\mathrm{pixel}}. \tag{64}$$

The ratio $M_{\mathrm{cube}}/M_{\mathrm{image}}$ should be close to unity (typically within a few percent), confirming that the interpolation and vertical integration have not introduced significant numerical errors.

The resulting three-dimensional cube serves as the source term for numerical gravitational solvers.

5.8 Diagnostic Plots

To verify the correctness of the reconstruction procedure, several diagnostic plots are generated.

- **Raw FITS image:** Shows the original 3.6 μm image before processing.
- **Cleaned and deprojected image:** Displays the galaxy after masking, pixel area correction, and geometric deprojection. A successful deprojection should produce an approximately face-on disk.
- **Surface density map:** The logarithm of the stellar surface density, $\log_{10} \Sigma\ [M_\odot\,\mathrm{pc}^{-2}]$, visualizes the reconstructed stellar mass distribution across the galaxy.
- **Radial surface density profile:** Shows the azimuthally averaged stellar surface density $\Sigma(R)$ as a function of radius, typically compared to an exponential fit.
- **Enclosed mass profile:** Shows the cumulative stellar mass $M(< R)$ as a function of radius, used to compute the half-mass radius.
- **Midplane density slice:** A two-dimensional slice through the midplane of the reconstructed density cube ($z = 0$) illustrates the final stellar density distribution used in the gravitational solver.

Together, these diagnostics provide visual verification that the image processing, physical unit conversion, and density reconstruction behave consistently with expectations for spiral galaxies.

The final output of the pipeline is a three-dimensional density cube,

$$\rho(x, y, z), \tag{65}$$

in SI units of $\mathrm{kg\,m^{-3}}$, suitable for use in numerical simulations of galactic gravitational fields.

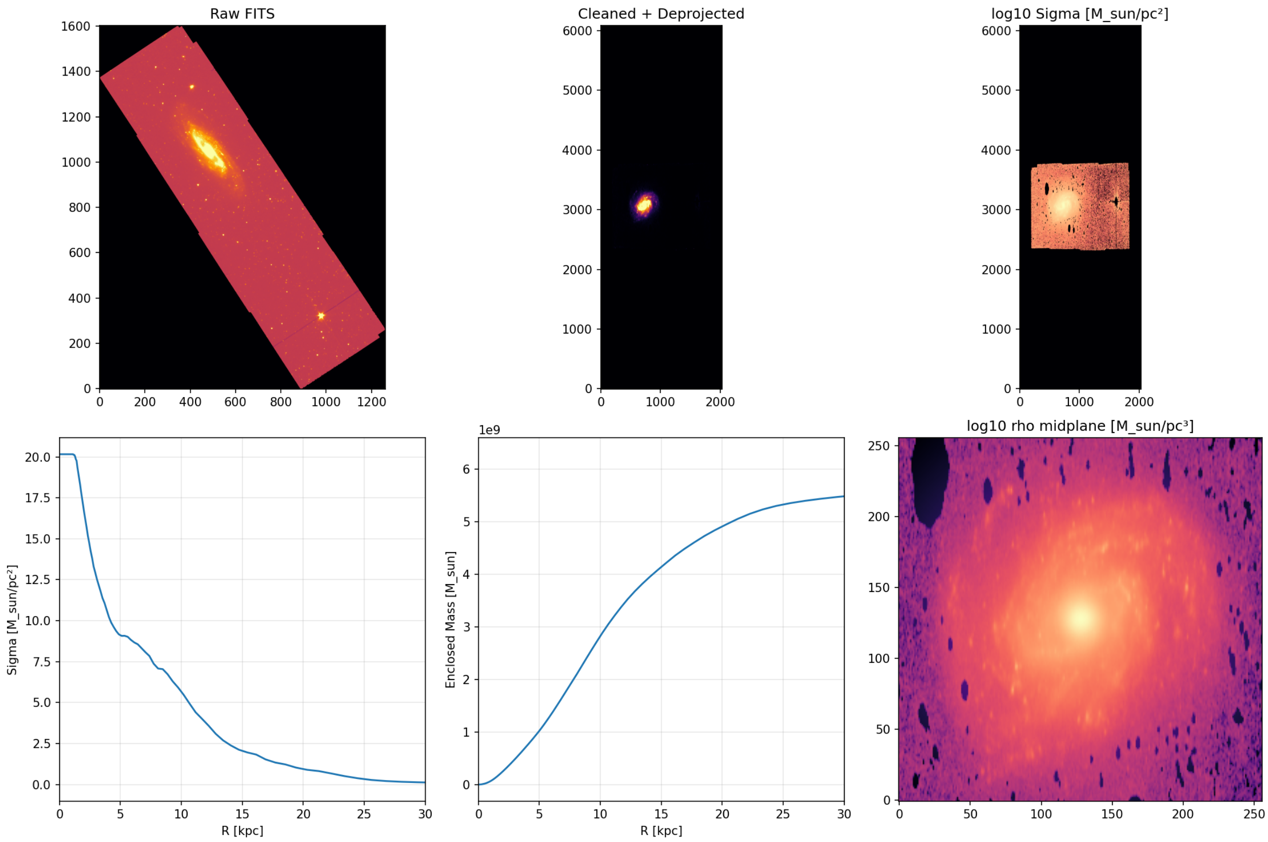

Figure 13: Diagnostic plots produced by the stellar density reconstruction pipeline for NGC 3198. **Top left:** raw 3.6 μm FITS image showing the observed infrared emission of the galaxy. **Top centre:** cleaned and geometrically deprojected image after masking and inclination correction. **Top right:** logarithmic stellar surface density map $\log_{10}\Sigma$ [$M_\odot$ pc^{-2}], obtained from the infrared luminosity using the S4G conversion factor with $M/L = 0.5$. **Bottom left:** azimuthally averaged stellar surface density profile $\Sigma(R)$. **Bottom centre:** enclosed stellar mass profile $M(< R)$. **Bottom right:** midplane slice of the reconstructed three-dimensional stellar density cube used as input for the gravitational solver.

5.9 Calibration of Model Parameters

The stellar mass distribution contains several free parameters that were calibrated by matching the predicted MOND rotation curve to the observed HI rotation curve of NGC 3198 [26]. The underlying assumption is that if MOND is correct, the circular velocity computed from the MOND acceleration, $v = \sqrt{g_{\mathrm{MOND}} \cdot r}$, should closely follow the observed velocities.

Three geometric parameters are fixed to literature values: distance (14.0 Mpc), inclination (71.5°), and position angle (213.0°). The remaining free parameters are the stellar mass-to-light ratio M/L, the vertical scale height z_0, the choice of vertical profile (single sech^2 or thin+thick disk), and the bulge mass M_{bulge} and Plummer scale radius a. The thin+thick disk sub-parameters ($z_{\mathrm{thin}} = 250$ pc, $z_{\mathrm{thick}} = 900$ pc, $f_{\mathrm{thin}} = 0.75$) are held fixed at typical values from the literature.

After systematic exploration, the following values provided the best agreement with the observed rotation curve:

$$M/L = 0.5, \quad z_0 = 400\ \mathrm{pc}, \quad \mathrm{profile} = \mathrm{thin{+}thick}, \quad M_{\mathrm{bulge}} = 3 \times 10^9\, M_\odot, \quad a = 3\ \mathrm{kpc}. \tag{66}$$

Figure 14 presents this comparison, showing the Newtonian and MOND-predicted circular velocities alongside the observational data of Venkataramani & Newell (2021). The fitted rotation curves for NGC 2841 and NGC 2976 are displayed in the Appendix B.

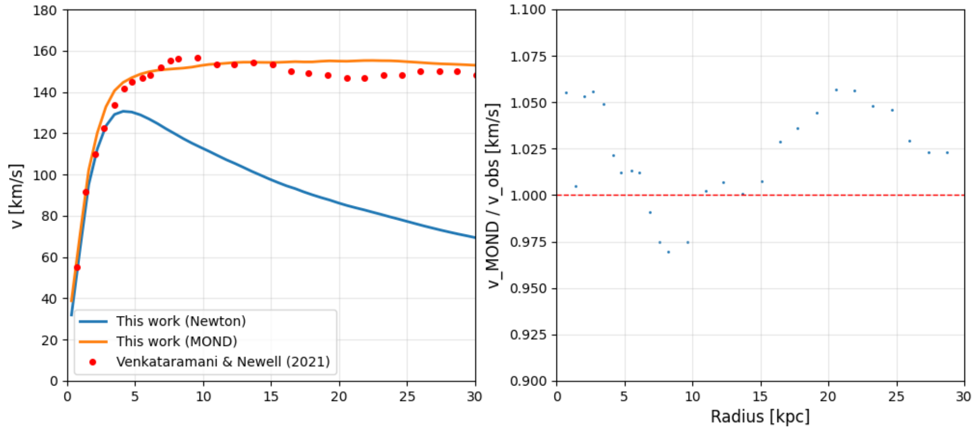

Figure 14: Comparison of the reconstructed rotation curve for NGC 3198 with observational HI data, demonstrating that the full AQUAL MOND solution successfully reproduces the flat outer rotation curve. The stellar mass distribution was constructed from Spitzer 3.6 μm imaging using the pipeline described in Section 5, with parameters calibrated to match the observed rotation: distance $D = 14.0$ Mpc, inclination $i = 71.5°$, position angle PA $= 213.0°$, mass-to-light ratio $M/L = 0.5$, vertical scale height $z_0 = 400$ pc with a thin+thick disk profile ($z_{\rm thin} = 250$ pc, $z_{\rm thick} = 900$ pc, $f_{\rm thin} = 0.75$), and a Plummer bulge component with $M_{\rm bulge} = 3 \times 10^9 M_\odot$ and scale radius $a = 3$ kpc. Left panel: The Newtonian baryonic rotation curve $v_{\rm bar} = \sqrt{g_N^{\rm radial} \cdot r}$ (blue solid) rises to approximately 130 km s^{-1} and then declines, failing to account for the flat outer rotation curve. The full AQUAL MOND prediction $v_{\rm MOND} = \sqrt{g_M \cdot r}$ (orange solid) with $a_0 = 1.2 \times 10^{-10}$ m s^{-2} successfully reproduces the flat rotation curve out to 30 kpc, closely following the digitised HI rotation curve measurements from Venkataramani and Newell [26] (red points). Right panel: Ratio of the MOND-predicted velocity to the observed velocity, $v_{\rm MOND}/v_{\rm obs}$. The scatter around unity (indicated by the dashed red line) is within $\sim 5\%$ across the disk, demonstrating good agreement between the model and observations. This calibration confirms that the reconstructed stellar mass distribution, combined with MOND, provides an accurate description of the galaxy's kinematics, and validates the use of the density cube for subsequent RAR analysis.

6 Results

This section addresses the central questions of this work: *(1) How accurate is the spherical approximation for computing baryonic accelerations in RAR studies, and (2) does using the true disk Newtonian field instead of $GM(<r)/r^2$ shift galaxies off the MOND RAR prediction?*
To answer these questions, we compare two different definitions of the baryonic acceleration $g_{\rm bar}$ for three galaxies that span a range of morphological types: NGC 3198 (a typical spiral with a prominent bulge), NGC 2841 (a massive, bulge-dominated spiral), and NGC 2976 (a bulgeless, low-surface-brightness galaxy). The used notation and definitions are:

1. **The spherical approximation:** $g_{\rm sph} = GM(<r)/r^2$, where $M(<r)$ is the enclosed stellar mass within a sphere of radius r. This is the definition commonly used in observational RAR studies (e.g., McGaugh et al., 2016) because it can be computed directly from the azimuthally averaged surface density profile.
2. **The true disk gravity:** $g_N^{\rm radial}$, the radial component of the Newtonian acceleration obtained by solving the full three-dimensional Poisson equation for the flattened stellar disk. This is the physically correct $g_{\rm bar}$ for a disk galaxy.

We then compute the MOND-predicted acceleration $g_{\rm MOND}$ via the full AQUAL solver for both definitions of $g_{\rm bar}$ and compare the resulting positions on the Radial Acceleration Relation.

6.1 Quantifying the Error of the Spherical Approximation

Figure 15 shows the direct comparison between $g_{\rm sph}$ and the true disk gravity $g_N^{\rm radial}$ for all three galaxies. The top row plots both quantities as functions of radius; the bottom row shows their ratio.
The key result is that the accuracy of the spherical approximation correlates strongly with the prominence of the bulge: more bulge-dominated galaxies yield more accurate results, while bulgeless galaxies show the largest discrepancies.
NGC 2841 (Fig. 15a), Bulge-dominated: This galaxy exhibits the smallest discrepancy. The ratio $g_N^{\rm radial}/g_{\rm sph}$ remains close to unity ($\sim 0.98 - 1.02$) across most of the disk, deviating only slightly at the outermost radii. The spherical approximation works well because the mass distribution is dominated by the central bulge, which is approximately spherical. The disk's flattening has a negligible effect on the radial acceleration when a massive bulge is present.

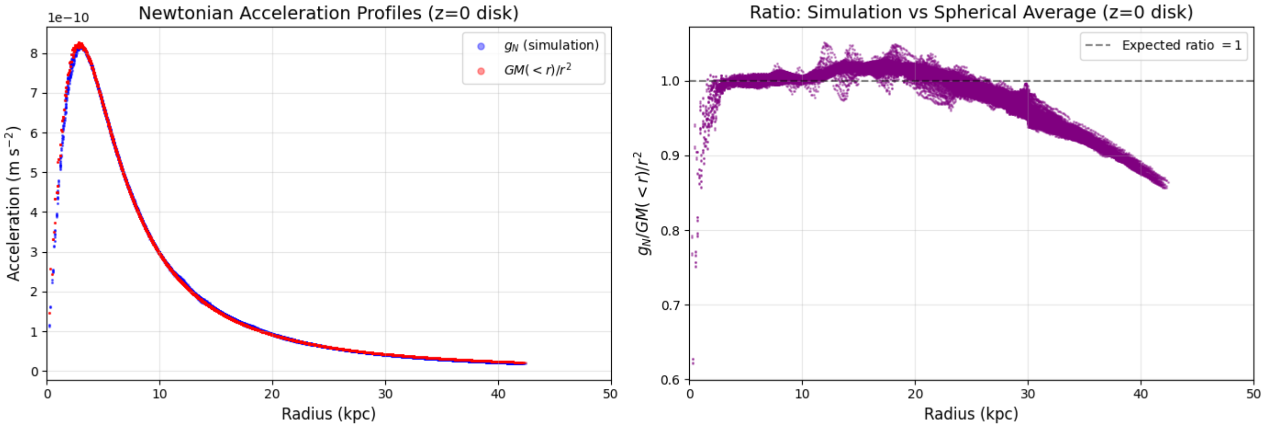

(a) NGC 2841 (bulge-dominated)

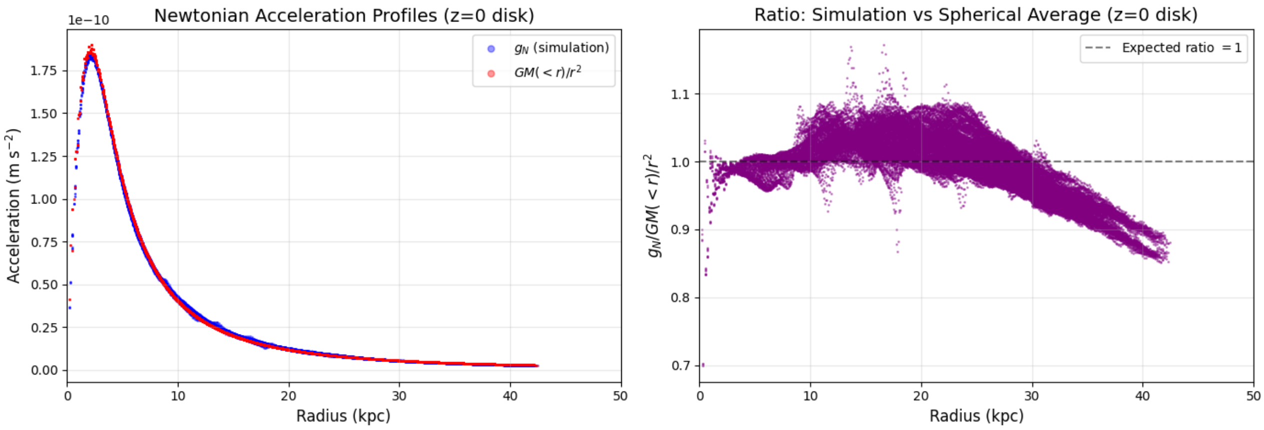

(b) NGC 3198 (moderate bulge)

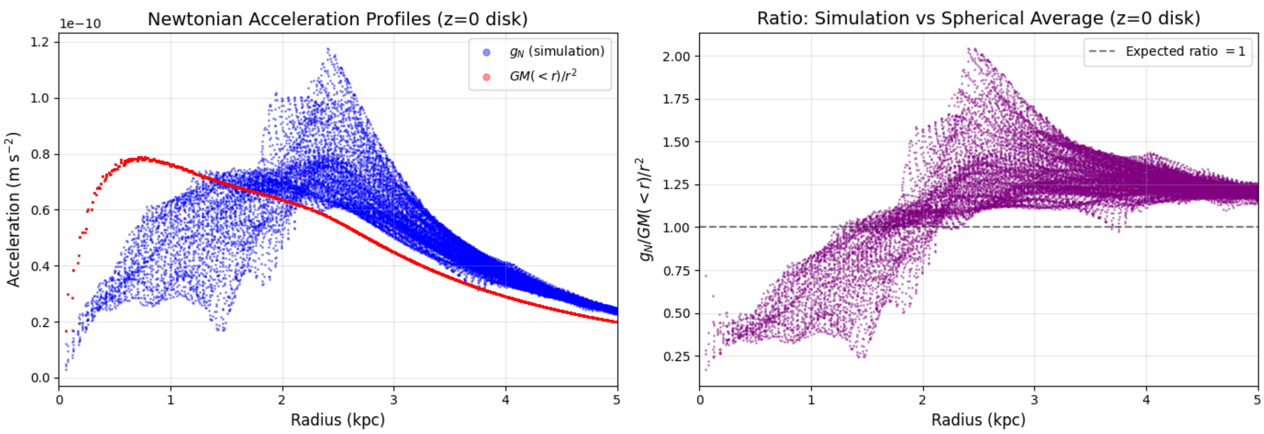

(c) NGC 2976 (bulgeless)

Figure 15: Comparison between the spherical enclosed-mass approximation $g_{\rm sph} = GM(<r)/r^2$ and the true disk Newtonian gravity $g_N^{\rm radial}$ for three galaxies. For each galaxy, left panels show radial acceleration profiles for $g_{\rm sph}$ (red) and $g_N^{\rm radial}$ (blue); right panels show $g_N^{\rm radial}/g_{\rm sph}$ (dashed line at unity). The accuracy correlates with bulge prominence: $\sim 2\%$ for NGC 2841, 5–10% for NGC 3198, and $\sim 25\%$ average for NGC 2976 with localised peaks approaching a factor of two. The spherical approximation neglects in-plane gravitational contributions from mass at larger radii, enhancing the radial force in flattened systems. Scatter at small radii is a discretisation artefact.

NGC 3198 (Fig. 15b), Moderate bulge: This galaxy shows intermediate behaviour. The ratio deviates from unity across the disk, reaching a maximum of ~ 1.1 in the range $10-25$ kpc, meaning the spherical approximation underestimates the true baryonic acceleration by approximately $5-10\%$ in the outer disk. At large radii ($35-40$ kpc), the ratio falls below unity to approximately 0.9, indicating that the spherical approximation slightly overestimates the true gravity at the outermost regions. The different radial values reflect the transition from bulge-dominated to a more disk-dominated character.

NGC 2976 (Fig. 15c) — Bulgeless: This galaxy exhibits the largest discrepancy. The ratio $g_N^{\rm radial}/g_{\rm sph}$ averages ~ 1.25 across most of the disk, with localised peaks reaching nearly 2.0 (i.e., a 100% error) at $r \approx 2.5$ kpc. It is also visible that $g_N^{\rm radial}/g_{\rm sph}$ is much below 1 for $r \lesssim 1.5$ kpc. Unlike NGC 3198, the ratio never drops below unity for large radii; the spherical approximation systematically underestimates the true baryonic acceleration at these radii.

The physical interpretation is clear: the spherical approximation neglects the gravitational contribution of mass at larger radii that lies in the same plane, which enhances the radial force in the midplane of a flattened system. This effect is weakest when the mass distribution is already nearly spherical (bulge-dominated NGC 2841) and strongest when the mass distribution is purely disk-like (bulgeless NGC 2976). Crucially, the error is not constant — it varies with both galaxy type and radius, and can be as large as 25% for bulgeless systems.

6.2 The Impact on the Radial Acceleration Relation

Having established that the spherical approximation can be significantly wrong for some galaxies, we now ask: *Does this error matter for MOND tests?* If we use the wrong $g_{\rm bar}$, do galaxies still fall on the MOND RAR, or are they systematically shifted off the predicted relation?

Figure 16 answers this question by comparing the MOND RAR for all three galaxies using both definitions of $g_{\rm bar}$. For each galaxy, the left subpanel uses $g_{\rm bar} = g_{\rm sph}$ (the spherical approximation), and the right subpanel uses $g_{\rm bar} = g_N^{\rm radial}$ (the true disk gravity). The blue points show $g_{\rm MOND}$ from the full AQUAL solution (the same in both subpanels for each galaxy). The red curve shows the algebraic MOND prediction $g_{\rm algebraic} = \nu(g_{\rm bar}/a_0)g_{\rm bar}$.

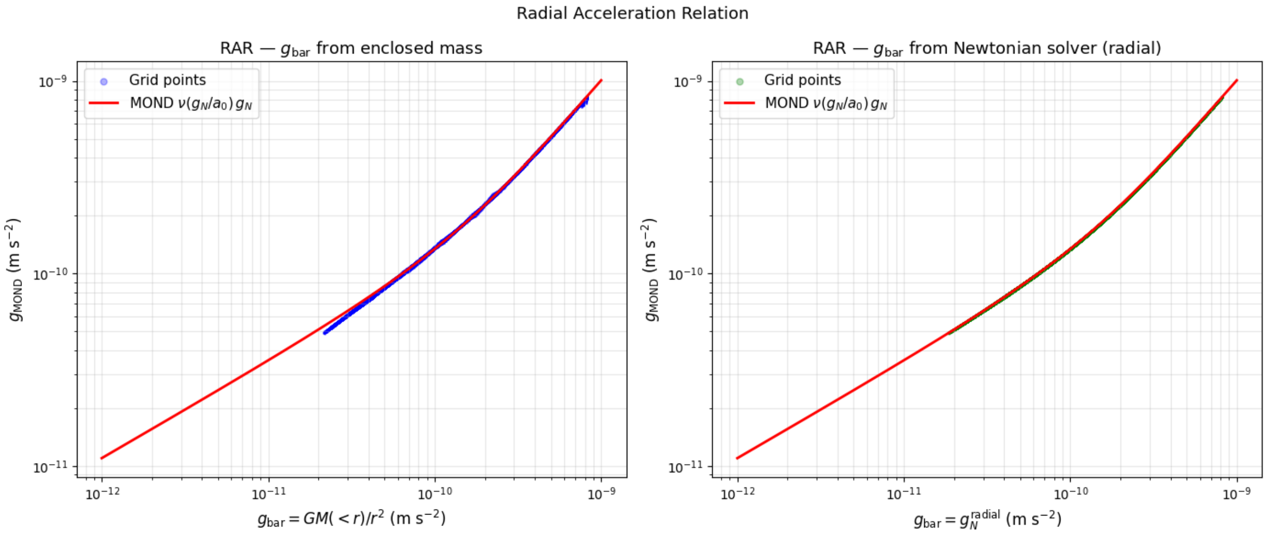

(a) NGC 2841 (bulge-dominated)

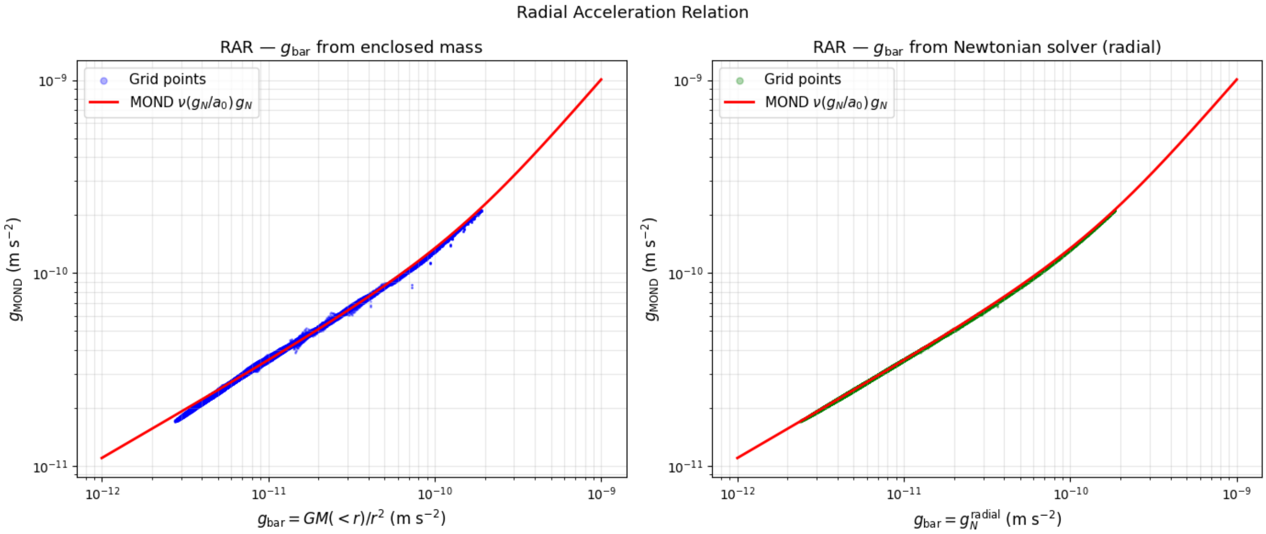

(b) NGC 3198 (moderate bulge)

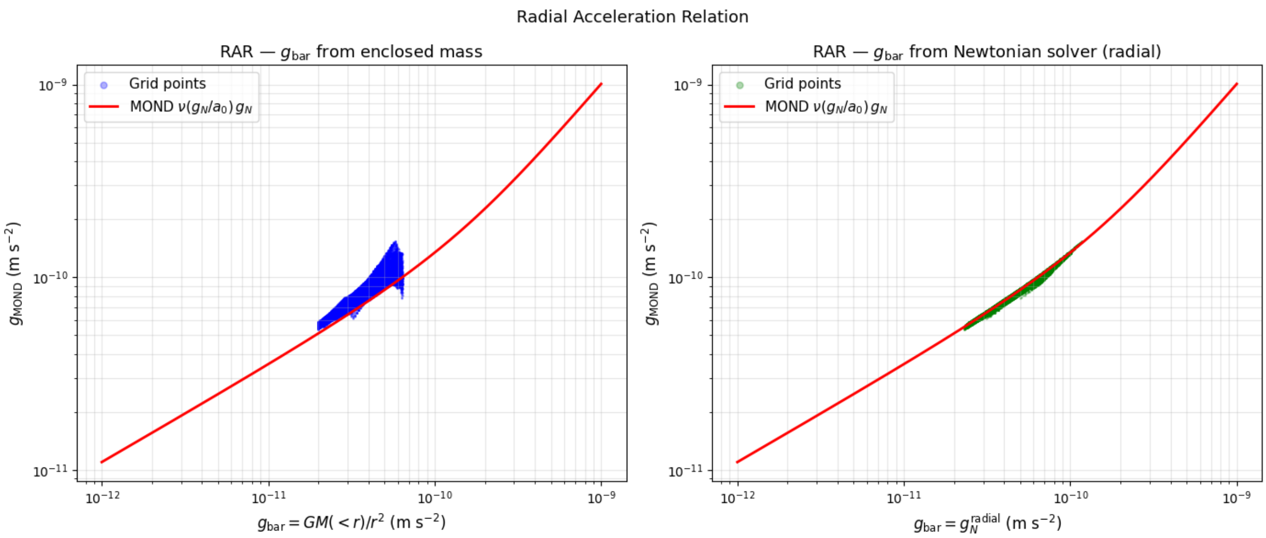

(c) NGC 2976 (bulgeless)

Figure 16: Impact of the spherical approximation on the MOND Radial Acceleration Relation. Green points use $g_{\rm bar} = g_{\rm sph} = GM(< r)/r^2$; blue points use $g_{\rm bar} = g_N^{\rm radial}$ from the full Poisson solver. Solid red line: algebraic MOND $\nu(g_{\rm bar}/a_0)g_{\rm bar}$ with $\nu(y) = \sqrt{\frac{1}{2} + \frac{1}{2}\sqrt{1 + 4/y^2}}$ (Eq. 16) and $a_0 = 1.2 \times 10^{-10}\,{\rm m\,s^{-2}}$. Points are restricted to reliable radii ($r \gtrsim 2\,{\rm kpc}$, $r \lesssim 49.7\,{\rm kpc}$).

The results are again galaxy and bulge-dependent:
NGC 2841 (Fig. 16a) — Bulge-dominated: Both the spherical approximation (blue points) and the true disk gravity (green points) lie almost exactly on top of each other and on the MOND RAR curve, except for the blue points beloning to the lowest accelerations. The excellent agreement reflects the fact that the spherical approximation is accurate for this galaxy (as shown in Fig. 15a), so switching to the true disk gravity makes only small visible difference.
NGC 3198 (Fig. 16b) — Moderate bulge: The blue points are slightly separated from the red line but both remain close to the MOND RAR curve, even through the intermediate MOND regime.
NGC 2976 (Fig. 16c) — Bulgeless: This galaxy tells a dramatically different story. Using the spherical approximation (blue points), the galaxy appears to lie systematically *above* the MOND RAR curve across most radii, the points fall significantly off the red line. However, when the true disk gravity is used (green points), the galaxy shifts almost perfectly onto the algebraic MOND prediction. The horizontal shift is substantial (given the $\sim 25\%$ error in $g_{\rm bar}$ shown in Fig. 15c), and it is sufficient to bring the galaxy out of agreement with MOND.
This is a critical finding: *For bulgeless galaxies, the spherical approximation can make a MOND-consistent galaxy appear to violate the RAR.* The effect arises because bulgeless galaxies have no central spherical component to "average over" the disk flattening; the discrepancy between $g_{\rm sph}$ and $g_N^{\rm radial}$ manifests directly in the RAR plot. In contrast, bulge-dominated galaxies (NGC 2841) and even moderate spirals with bulges (NGC 3198) are more forgiving. The spherical approximation, though quantitatively wrong, preserves the qualitative alignment with MOND. Also note that regardless of galaxy, the green points lie almost perfectly on the algebraic MOND line. This confirms earlier findings that suggest that the predicted intrinsic scatter between QuMOND (20) and algebraic MOND (21) is small compared to the observational uncertainties [12].

6.3 Interpretation: Does the Spherical Approximation Matter?

The answer depends sensitively on the galaxy type and the question being asked:
(1) For testing whether galaxies follow the MOND RAR:

- For bulge-dominated and normal spiral galaxies (NGC 2841, NGC 3198), the spherical approximation is sufficient. The systematic error in $g_{\rm bar}$ does not shift the galaxy off the MOND RAR more than observational erros.
- For bulgeless galaxies (NGC 2976), the spherical approximation can lead to a false negative, a galaxy that satisfies MOND when using the correct $g_{\rm bar}$ appears to violate it when using $GM(<r)/r^2$.

This could suggests that existing RAR studies that rely on the spherical approximation may have systematically underestimated MOND's consistency with bulgeless, low-surface-brightness galaxies.
(2) For determining the correct $g_{\rm bar}$ value of a galaxy: The spherical approximation is wrong by up to 25% for NGC 2976, $5-10\%$ for NGC 3198, and only $\sim 2\%$ for NGC 2841. If one cares about the absolute value of $g_{\rm bar}$ at a given radius, for example, when comparing to observed rotation curves, then using $GM(<r)/r^2$ introduces a significant, radius-dependent, galaxy-dependent systematic error.
(3) For determining whether we should solve the full Poisson equation for RAR studies: The answer is *yes for bulgeless galaxies*, and *no for bulge-dominated ones*. Galaxies with a significant spherical component (bulge or large central concentration) are safely analysed with the spherical approximation for MOND tests. Galaxies without a bulge require the full Poisson solution to avoid biasing the RAR test against MOND.

7 Discussion and Conclusions

7.1 Summary of Findings

This work set out to answer two questions: *(1) How accurate is the spherical approximation $g_{\rm bar} = GM(<r)/r^2$ for disk galaxies*, and *(2) does its use bias tests of the MOND Radial Acceleration Relation?*
Using a numerical AQUAL solver applied to realistic three-dimensional stellar density reconstructions of three galaxies (NGC 3198, NGC 2841, and NGC 2976) derived from Spitzer 3.6 μm imaging, we compared two definitions of the baryonic acceleration: (1) the spherical enclosed-mass approximation commonly used in observational RAR studies, and (2) the true radial component of the Newtonian gravitational field obtained by solving the full Poisson equation for the flattened disk. The MOND-predicted acceleration $g_{\rm obs}$ was then computed via the full AQUAL iteration for both definitions.
The key quantitative findings are:

- The accuracy of the spherical approximation depends strongly on galaxy morphology. For the bulge-dominated galaxy NGC 2841, the ratio $g_N^{\rm radial}/g_{\rm sph}$ remains within $\sim 2\%$ of unity. For the moderate-bulge NGC 3198, the discrepancy reaches $\sim +10\%$ in the outer disk (10–25 kpc) and $\sim -10\%$ at the largest radii (35–40 kpc). For the bulgeless galaxy NGC 2976, the spherical approximation underestimates the true baryonic acceleration by $\sim 25\%$ on average, with localised peaks reaching a factor of two and overestimates for low R.

- The impact of this error on MOND tests is equally shape-dependent. For bulge-dominated and moderate-bulge galaxies (NGC 2841 and NGC 3198), the MOND RAR is essentially unchanged: galaxies that lie on the algebraic relation $g_{\rm obs} = \nu(g_{\rm bar}/a_0)\, g_{\rm bar}$ using the spherical approximation also lie on it within observational error using the correct disk gravity. For bulgeless galaxies (NGC 2976), however, the spherical approximation can make a MOND-consistent galaxy appear to systematically violate the RAR. When the correct disk gravity is used, NGC 2976 shifts back onto the MOND prediction.
- The full AQUAL solution for a realistic disk galaxy follows the algebraic MOND prediction with negligible scatter, confirming that the curl-free condition required for the algebraic approximation is approximately satisfied for axisymmetric systems.

7.2 Comparison with Previous Work

Our finding that the algebraic MOND approximation works well for disk galaxies is consistent with previous studies. Brada and Milgrom [18] showed that for exponential disks, the full AQUAL field equations yield rotation curves nearly identical to the simple algebraic relation. Jones-Smith et al. [12] extended this result to the QuMOND formalism, finding that the curl term in the QuMOND equation is surprisingly small for realistic disk geometries.

However, previous work focused on the accuracy of the MOND approximation *given the correct Newtonian field*. Our study addresses a distinct question: the accuracy of the *spherical approximation* for computing the Newtonian field itself. To our knowledge, this is the first direct quantification of how $GM(<r)/r^2$ compares to the true disk gravity across a range of galaxy morphologies, and the first test of whether this approximation biases the RAR.

Our result that the spherical approximation is accurate to within $\sim 10\%$ for NGC 3198 (a moderate-bulge galaxy) is perhaps surprising given how flattened disk galaxies are. The relatively modest error can be understood from the radial surface density profile: NGC 3198 has a long, extended disk with no sharp truncation, so the mass distribution remains approximately spherical in an integrated sense even though the instantaneous density is highly flattened. The much larger errors found for the bulgeless NGC 2976 confirm the physical expectation that the spherical approximation breaks down most severely when there is no central spherical component to dominate the potential.

7.3 Implications for MOND Tests

Our results have two important implications for empirical tests of MOND using the RAR.

First, the spherical approximation is sufficient for testing whether *bulge-dominated and normal (moderate-bulge) spiral* galaxies follow the MOND relation. An observer who computes $g_{\rm bar}$ via $GM(<r)/r^2$ from an azimuthally averaged surface density profile will introduce a systematic error of $\lesssim 10\%$ for such galaxies, which does not shift them off the RAR curve beyond observational uncertainties. The existing RAR literature [8] is therefore not systematically biased by the spherical approximation for this class of galaxy.

Second, and more critically, the spherical approximation is *inadequate* for bulgeless, low-surface-brightness galaxies. For NGC 2976, the $\sim 25\%$ underestimate of $g_{\rm bar}$ for a large part of the galatic extension causes the galaxy to appear to violate the MOND RAR when it in fact satisfies it. This implies that existing RAR studies may have systematically underestimated MOND's consistency with bulgeless galaxies: apparent outliers above or below the RAR in samples such as SPARC may be artefacts of the spherical approximation rather than genuine failures of MOND.

Third, the spherical approximation is inadequate for precise rotation curve modelling of individual galaxies regardless of type. If one aims to infer the mass-to-light ratio Υ_*, constrain the vertical scale height, or test MOND against detailed HI kinematics, using $GM(<r)/r^2$ introduces a radius-dependent systematic error that is comparable to or larger than the reported scatter of the RAR. For such applications, solving the full Poisson equation for the true disk gravity is necessary.

7.4 Limitations and Future Work

Several limitations of this study should be acknowledged.

Small galaxy sample. Our analysis covers three galaxies that span a range of morphological types, but a statistically robust conclusion requires a larger sample. Future work should extend this analysis to the full SPARC catalogue [17], spanning a wide range of Hubble types, surface brightnesses, and bulge-to-disk ratios, to determine how the error of the spherical approximation varies systematically across the galaxy population.

Axisymmetry. Our density reconstructions assumed axisymmetry after deprojection. Real galaxies contain spiral arms, bars, and other non-axisymmetric structures. These features may introduce additional scatter in the RAR that is not captured by our axisymmetric model. However, the empirical RAR shows remarkably little scatter (~ 0.1 dex), suggesting that non-axisymmetric features do not strongly affect the radial acceleration on average.

Numerical resolution. At small radii ($r \lesssim 2$ kpc), discretisation artefacts due to the finite Cartesian grid produce scatter in the ratio plots. This region is not the primary focus of MOND tests, the deep-MOND regime of interest occurs at low accelerations in the outer disk, but improved resolution would cleanly separate numerical noise from physical effects near the centre, which is particularly relevant for the bulgeless NGC 2976 where the largest discrepancies occur at small radii.

Baryonic mass model. Our stellar mass distributions were constructed from 3.6 μm imaging with assumed mass-to-light ratios calibrated to match the observed rotation curves. The gas component (atomic and molecular hydrogen) was not included because the S4G images trace only the old stellar population. For galaxies like NGC 3198, the gas mass becomes important in the outermost disk and its inclusion would slightly modify $g_{\rm bar}$ at large radii.
Other MOND formulations. Our calculations used the AQUAL formulation with the standard interpolation function $\mu(x) = x/\sqrt{1+x^2}$. The QuMOND formulation [11] might produce slightly different results, particularly regarding the curl-free condition. While Jones-Smith et al. [12] found QuMOND and the algebraic approximation to agree closely, a direct comparison with our AQUAL implementation for the same density cubes would be valuable.

7.5 Conclusions

The central conclusions of this work are as follows:

1. The accuracy of the spherical approximation $g_{\rm bar} = GM(<r)/r^2$ depends strongly on galaxy morphology. For bulge-dominated galaxies (NGC 2841), the error is $\lesssim 2\%$. For moderate-bulge spirals (NGC 3198), it reaches $\sim 10\%$ in the outer disk. For bulgeless galaxies (NGC 2976), the spherical approximation underestimates the true baryonic acceleration by $\sim 25\%$ on average, with localised errors approaching a factor of two.

2. For bulge-dominated and moderate-bulge galaxies, the spherical approximation does *not* bias tests of the MOND RAR. Galaxies that follow the algebraic relation $g_{\rm obs} = \nu(g_{\rm bar}/a_0)\, g_{\rm bar}$ using the spherical approximation also follow it using the correct disk gravity. Existing RAR studies employing $GM(<r)/r^2$ are therefore valid for this class of galaxy.

3. For bulgeless galaxies, the spherical approximation *does* bias MOND tests. NGC 2976 appears to systematically violate the RAR when the spherical approximation is used, but is consistent with the MOND prediction when the correct disk gravity is applied. Existing RAR studies may therefore have systematically underestimated MOND's consistency with bulgeless, low-surface-brightness galaxies.

4. For applications requiring precise rotation curve predictions, such as inferring Υ_* or testing MOND against detailed kinematics, the spherical approximation is inadequate for all galaxy types studied, and the full Poisson equation for the true disk gravity should be solved.

5. The full AQUAL solution for a realistic disk galaxy follows the algebraic MOND prediction with negligible scatter, confirming that the curl-free condition is approximately satisfied for axisymmetric systems.

7.6 Outlook

The results presented here suggest a clear path forward for future RAR studies. For large-sample analyses aimed at testing the existence of the RAR and its consistency with MOND in bulge-dominated and normal spiral galaxies, the spherical approximation remains a practical and sufficiently accurate tool.
However, for bulgeless and low-surface-brightness galaxies the spherical approximation introduces biases large enough to produce false apparent violations of the RAR. Future large-sample studies should therefore apply the full three-dimensional Poisson solver to bulgeless galaxies. The computational cost is modest (a single galaxy can be solved in minutes on a GPU), and the increased accuracy is essential for an unbiased test of MOND across the full range of galaxy morphologies.
Finally, extending this analysis to a diverse sample of galaxies would provide a quantitative calibration of the spherical approximation as a function of galaxy properties such as bulge fraction, disk scale length, and surface brightness. Such a calibration would enable observers to identify which galaxies require the full Poisson treatment and to correct for the spherical approximation bias when needed, while retaining the simplicity of $GM(<r)/r^2$ for the majority of galaxies in statistical samples.

References

[1] European Space Agency (ESA). The best milky way map, by gaia (edge-on). ESA Multimedia, January 2025. URL https://www.esa.int/ESA_Multimedia/Images/2025/01/The_best_Milky_Way_map_by_Gaia_edge-on. Artist's impression of the Milky Way based on Gaia data.

[2] University of Glasgow. Astronomy 1y: Introduction to cosmology. Lecture notes, University of Glasgow, 2024. URL https://radio.astro.gla.ac.uk/a1cosmology/pw/Part_I_lecture_cosmology.pdf.

[3] V. C. Rubin, W. K. J. Ford, and N. Thonnard. Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 (R=4kpc) to UGC 2885 (R=122kpc). *Astrophysical Journal*, 238:471, 1980. doi: 10.1086/158003.

[4] T. S. van Albada, J. N. Bahcall, K. Begeman, and R. Sancisi. Distribution of dark matter in the spiral galaxy NGC 3198. *Astrophysical Journal*, 295:305, 1985. doi: 10.1086/163375.

[5] Planck Collaboration, N. Aghanim, Y. Akrami, and et al. Planck 2018 results. VI. cosmological parameters. *Astronomy & Astrophysics*, 641:A6, 2020. doi: 10.1051/0004-6361/201833910.

[6] J. F. Navarro, C. S. Frenk, and S. D. M. White. A universal density profile from hierarchical clustering. *Astrophysical Journal*, 490:493, 1997. doi: 10.1086/304888.

[7] W. J. G. de Blok. The core-cusp problem. *Advances in Astronomy*, 2010:789293, 2010. doi: 10.1155/2010/789293.

[8] S. S. McGaugh, F. Lelli, and J. M. Schombert. Radial acceleration relation in rotationally supported galaxies. *Physical Review Letters*, 117(20):201101, 2016. doi: 10.1103/PhysRevLett.117.201101.

[9] M. Milgrom. A modification of the Newtonian dynamics as a possible alternative to the hidden mass hypothesis. *Astrophysical Journal*, 270:365, 1983. doi: 10.1086/161130.

[10] J. Bekenstein and M. Milgrom. Does the missing mass problem signal the breakdown of Newtonian gravity? *Astrophysical Journal*, 286:7, 1984. doi: 10.1086/162570.

[11] M. Milgrom. Quasi-linear formulation of MOND. *Monthly Notices of the Royal Astronomical Society*, 403:886, 2010. doi: 10.1111/j.1365-2966.2009.16184.x.

[12] Katherine Jones-Smith, Harsh Mathur, Luca V. Delacrétaz, and Pranjal Sharma. The radial acceleration relation and a magnetostatic analogy in quasilinear MOND. *New Journal of Physics*, 20:083045, 2018. doi: 10.1088/1367-2630/aad5e6.

[13] R. H. Sanders and S. S. McGaugh. Modified Newtonian dynamics as an alternative to dark matter. *Annual Review of Astronomy & Astrophysics*, 40:263, 2002. doi: 10.1146/annurev.astro.40.060401.093923.

[14] D. Clowe, M. Bradač, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, and D. Zaritsky. A direct empirical proof of the existence of dark matter. *Astrophysical Journal Letters*, 648:L109, 2006. doi: 10.1086/508162.

[15] J. D. Bekenstein. Relativistic gravitation theory for the modified Newtonian dynamics paradigm. *Physical Review D*, 70:083509, 2004. doi: 10.1103/PhysRevD.70.083509.

[16] C. Skordis and T. Złosnik. New relativistic theory for modified Newtonian dynamics. *Physical Review Letters*, 127: 161302, 2021. doi: 10.1103/PhysRevLett.127.161302.

[17] F. Lelli, S. S. McGaugh, J. M. Schombert, and M. S. Pawlowski. One law to rule them all: The radial acceleration relation of galaxies. *Astrophysical Journal*, 836:152, 2017. doi: 10.3847/1538-4357/836/2/152.

[18] Radim Brada and Mordehai Milgrom. Exact solutions and approximations of MOND fields of disk galaxies. *Astrophysical Journal*, 512(2):L83–L86, 1999. doi: 10.1086/311866.

[19] ESA/Hubble and NASA. Hubble captures an edge-on spiral with curve appeal (ugc 10043). NASA Science, November 2024. URL https://science.nasa.gov/missions/hubble/hubble-captures-an-edge-on-spiral-with-curve-appeal/. Image of edge-on spiral galaxy UGC 10043. Composite of data from 2000 and 2023.

[20] Adam Lundie. Edge on spiral galaxy ngc 5170. Eatons Hill Observatory (Astropixels), October 2020. URL https://astropixels.com.au/astrophotography/Edge_On_Spiral_Galaxy_NGC_5170/. Amateur astrophotography image of NGC 5170.

[21] NASA/IPAC Infrared Science Archive. Spitzer survey of stellar structure in galaxies (S4G). Online data service, 2026. URL https://irsa.ipac.caltech.edu/data/SPITZER/S4G/. Accessed: 2026-05-29.

[22] J. V. de Nijs. Fast particle-mesh code for Milgromian dynamics. Bachelor's thesis, Delft University of Technology, 2023. Supervised by P. M. Visser and S. W. H. Eijt.

[23] K. C. Freeman. On the disks of spiral and S0 galaxies. *Astrophysical Journal*, 160:811, 1970. doi: 10.1086/150474.

[24] H. C. Plummer. On the problem of distribution in globular star clusters. *Monthly Notices of the Royal Astronomical Society*, 71:460, 1911. doi: 10.1093/mnras/71.5.460.

[25] P. C. van der Kruit and L. Searle. Surface photometry of edge-on spiral galaxies. *Astronomy & Astrophysics*, 95: 105, 1988.

[26] Shankar C. Venkataramani and Alan C. Newell. Pattern dark matter and galaxy scaling relations. *European Physical Journal Special Topics*, 230(9):2139–2165, 2021. doi: 10.1140/epjs/s11734-021-00196-9.

[27] D. Zobnina and A. Zasov. Galaxies with declining rotation curves. *Astronomy Reports*, 64:295, 2020. doi: 10.1134/S1063772920050054.

[28] John Marr. Galaxy rotation curves with log-normal density distribution. *Monthly Notices of the Royal Astronomical Society*, 448:482, 2015. doi: 10.1093/mnras/stv216.

A Derivation of Interpolation Function $\nu(y)$

From $\mu(x)x = y$, we have:

$$\frac{x^2}{\sqrt{1+x^2}} = y. \tag{67}$$

Squaring both sides:

$$\frac{x^4}{1+x^2} = y^2. \tag{68}$$

Rearranging:

$$x^4 = y^2(1+x^2) \quad \Rightarrow \quad x^4 - y^2x^2 - y^2 = 0. \tag{69}$$

This is a quadratic in x^2:

$$(x^2)^2 - y^2(x^2) - y^2 = 0. \tag{70}$$

Solving for x^2:

$$x^2 = \frac{y^2 \pm \sqrt{y^4 + 4y^2}}{2} = \frac{y^2 \pm |y|\sqrt{y^2+4}}{2}. \tag{71}$$

Taking the positive root (since $x^2 > 0$):

$$x^2 = \frac{y^2 + |y|\sqrt{y^2+4}}{2}. \tag{72}$$

For $y > 0$, this simplifies to:

$$x^2 = \frac{y^2 + y\sqrt{y^2+4}}{2}. \tag{73}$$

Since $x = \nu(y)y$, we have:

$$\nu(y)^2 y^2 = \frac{y^2 + y\sqrt{y^2+4}}{2}. \tag{74}$$

Dividing by y^2:

$$\nu(y)^2 = \frac{1 + \sqrt{1 + 4/y^2}}{2}. \tag{75}$$

Taking the square root gives the standard ν-function:

$$\nu(y) = \sqrt{\frac{1}{2} + \frac{1}{2}\sqrt{1 + \frac{4}{y^2}}}. \tag{76}$$

B Rotation Curves of NGC 2841 and NGC 2976

For completeness, Figures 17 and 18 present the observed rotation curves and the MOND-predicted rotation curves (using the true disk gravity) for NGC 2841 and NGC 2976. These complement the analysis presented in the main text.

C Diagnostic Plots for NGC 2841 and NGC 2976

For completeness, Figures 19 and 20 present the same diagnostic plots shown for NGC 3198 in the main text (Figure 13), now for the two additional galaxies analysed in this work. Each figure illustrates the key steps in the stellar density reconstruction pipeline.

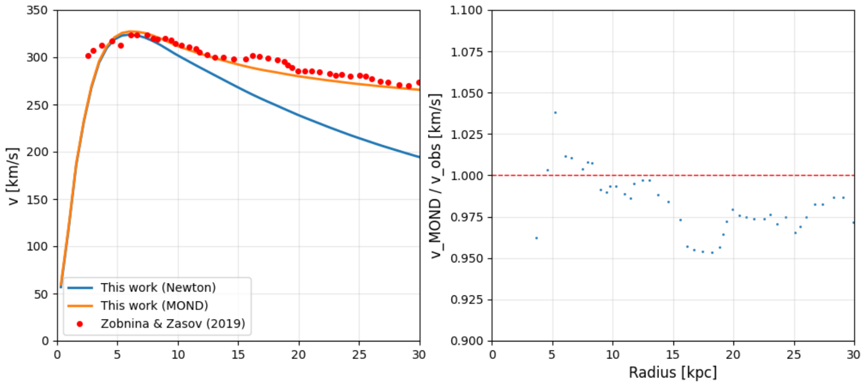

Figure 17: Comparison of the reconstructed rotation curve for the bulge-dominated galaxy NGC 2841 with observational data. The stellar mass distribution was constructed from Spitzer 3.6 μm imaging using the pipeline described in Section 5, with parameters calibrated to match the observed rotation. Left panel: The Newtonian baryonic rotation curve $v_{\text{bar}} = \sqrt{g_N^{\text{radial}} \cdot r}$ (blue solid) rises steeply in the inner regions due to the prominent bulge and then declines, failing to account for the flat outer rotation curve. The full AQUAL MOND prediction $v_{\text{MOND}} = \sqrt{g_M \cdot r}$ (orange solid) with $a_0 = 1.2 \times 10^{-10}\,\text{m}\,\text{s}^{-2}$ successfully reproduces the observed rotation curve, closely following the digitised measurements from Zobnina and Zasov [27] (red points). Right panel: Ratio of the MOND-predicted velocity to the observed velocity, $v_{\text{MOND}}/v_{\text{obs}}$. The scatter around unity (indicated by the dashed red line) is within $\sim 5\%$ across the disk, demonstrating good agreement between the model and observations. The excellent agreement reflects the fact that NGC 2841 is bulge-dominated; the mass distribution is approximately spherical, and MOND's predictions are relatively insensitive to the detailed disk geometry. This calibration confirms that the reconstructed stellar mass distribution, combined with MOND, provides an accurate description of the galaxy's kinematics.

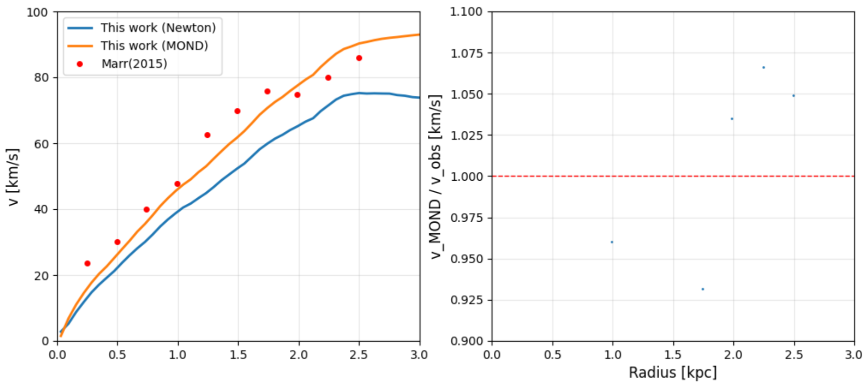

Figure 18: Comparison of the reconstructed rotation curve for the bulgeless galaxy NGC 2976 with observational data. The stellar mass distribution was constructed from Spitzer 3.6 μm imaging using the pipeline described in Section 5, with parameters calibrated to match the observed rotation. Left panel: The Newtonian baryonic rotation curve $v_{\text{bar}} = \sqrt{g_N^{\text{radial}} \cdot r}$ (blue solid) rises slowly and then declines, failing to account for the observed kinematics. The full AQUAL MOND prediction $v_{\text{MOND}} = \sqrt{g_M \cdot r}$ (orange solid) with $a_0 = 1.2 \times 10^{-10}\,\text{m}\,\text{s}^{-2}$ successfully reproduces the observed rotation curve, following the digitised measurements from Marr [28] (red points). Right panel: Ratio of the MOND-predicted velocity to the observed velocity, $v_{\text{MOND}}/v_{\text{obs}}$. The scatter around unity (indicated by the dashed red line) is within $\sim 10\%$ across most of the disk, with slightly larger deviations in the outermost regions where the observational uncertainties are also largest. This calibration confirms that MOND provides an accurate description of this bulgeless galaxy when the true disk gravity is used, consistent with the RAR analysis in Figure 13c which showed that NGC 2976 falls onto the MOND prediction when the spherical approximation is not employed. The agreement is particularly notable given that NGC 2976 has no central bulge to "average over" the disk flattening, demonstrating that MOND works for purely disk-like systems when the correct Newtonian field is computed.

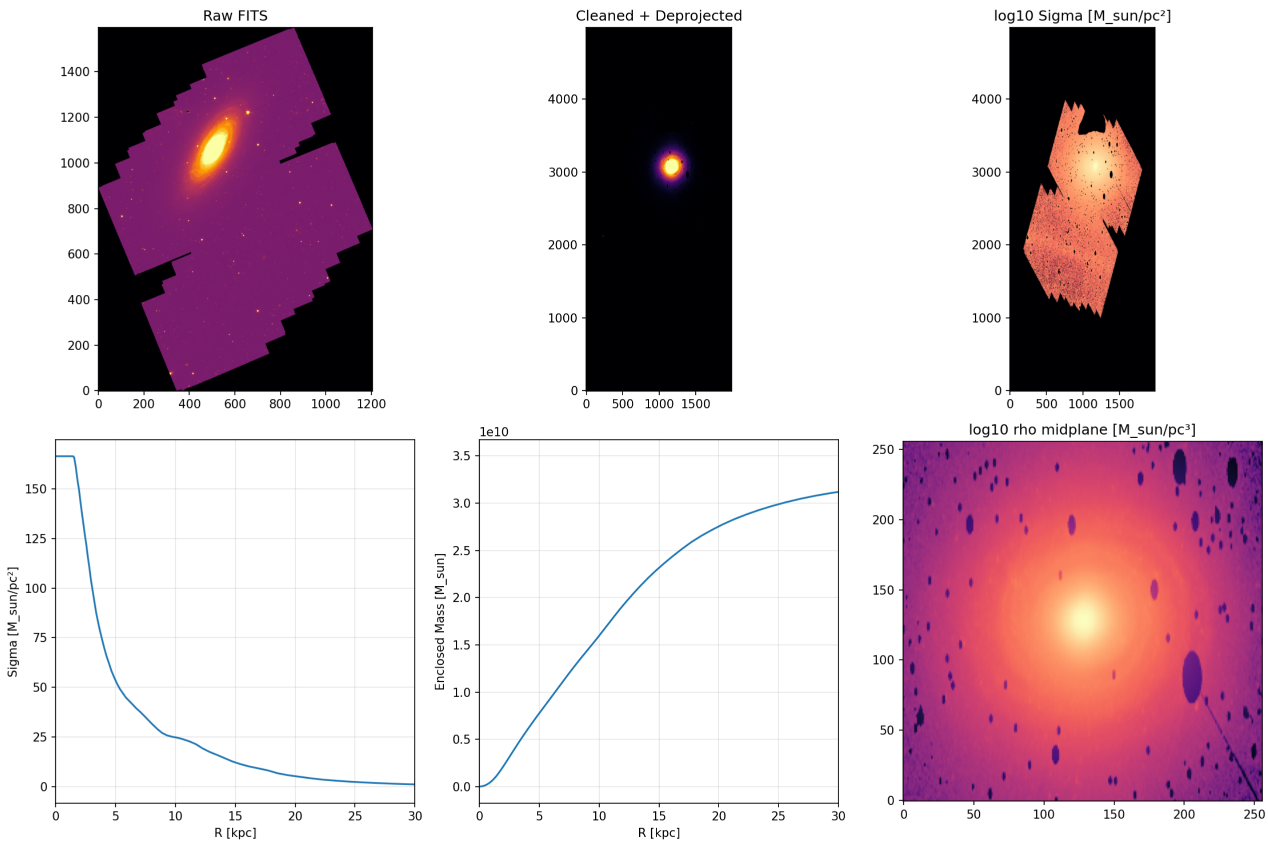

Figure 19: Diagnostic plots produced by the stellar density reconstruction pipeline for NGC 2841. **Top left:** raw 3.6 μm FITS image showing the observed infrared emission of the galaxy. **Top centre:** cleaned and geometrically deprojected image after masking and inclination correction. **Top right:** logarithmic stellar surface density map $\log_{10} \Sigma[M_\odot \, \mathrm{pc}^{-2}]$, obtained from the infrared luminosity using the S4G conversion factor with $M/L = 0.5$. **Bottom left:** azimuthally averaged stellar surface density profile $\Sigma(R)$. **Bottom centre:** enclosed stellar mass profile $M(< R)$. **Bottom right:** midplane slice of the reconstructed three-dimensional stellar density cube used as input for the gravitational solver.

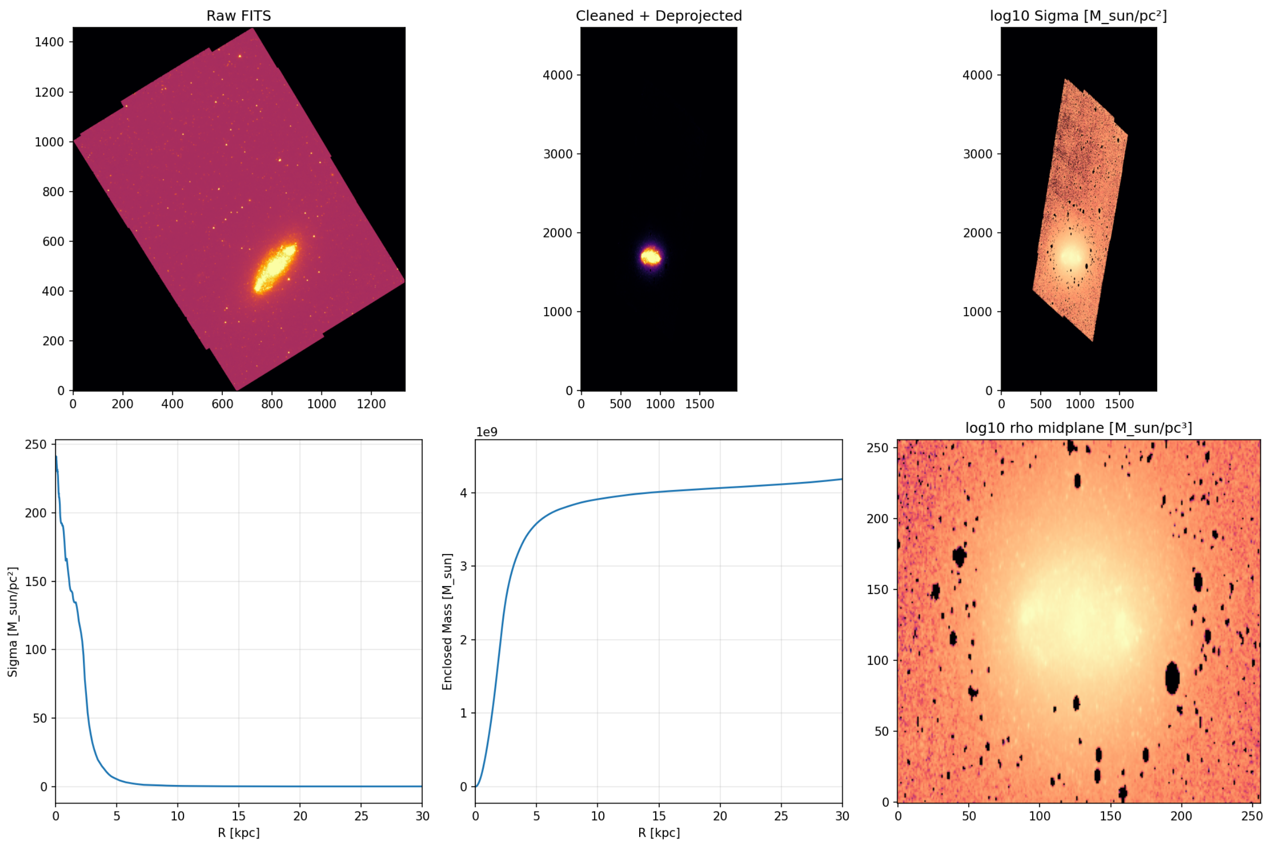

Figure 20: Diagnostic plots produced by the stellar density reconstruction pipeline for NGC 2976. **Top left:** raw 3.6 μm FITS image showing the observed infrared emission of the galaxy. **Top centre:** cleaned and geometrically deprojected image after masking and inclination correction. **Top right:** logarithmic stellar surface density map $\log_{10} \Sigma[M_\odot\,\mathrm{pc}^{-2}]$, obtained from the infrared luminosity using the S4G conversion factor with $M/L = 0.5$. **Bottom left:** azimuthally averaged stellar surface density profile $\Sigma(R)$. **Bottom centre:** enclosed stellar mass profile $M(< R)$. **Bottom right:** midplane slice of the reconstructed three-dimensional stellar density cube used as input for the gravitational solver.