Topics *related* to simulations of axions

1. Idealised model for interference in DM filaments
2. Emulator for the non-linear matter power-spectrum in mixed DM models

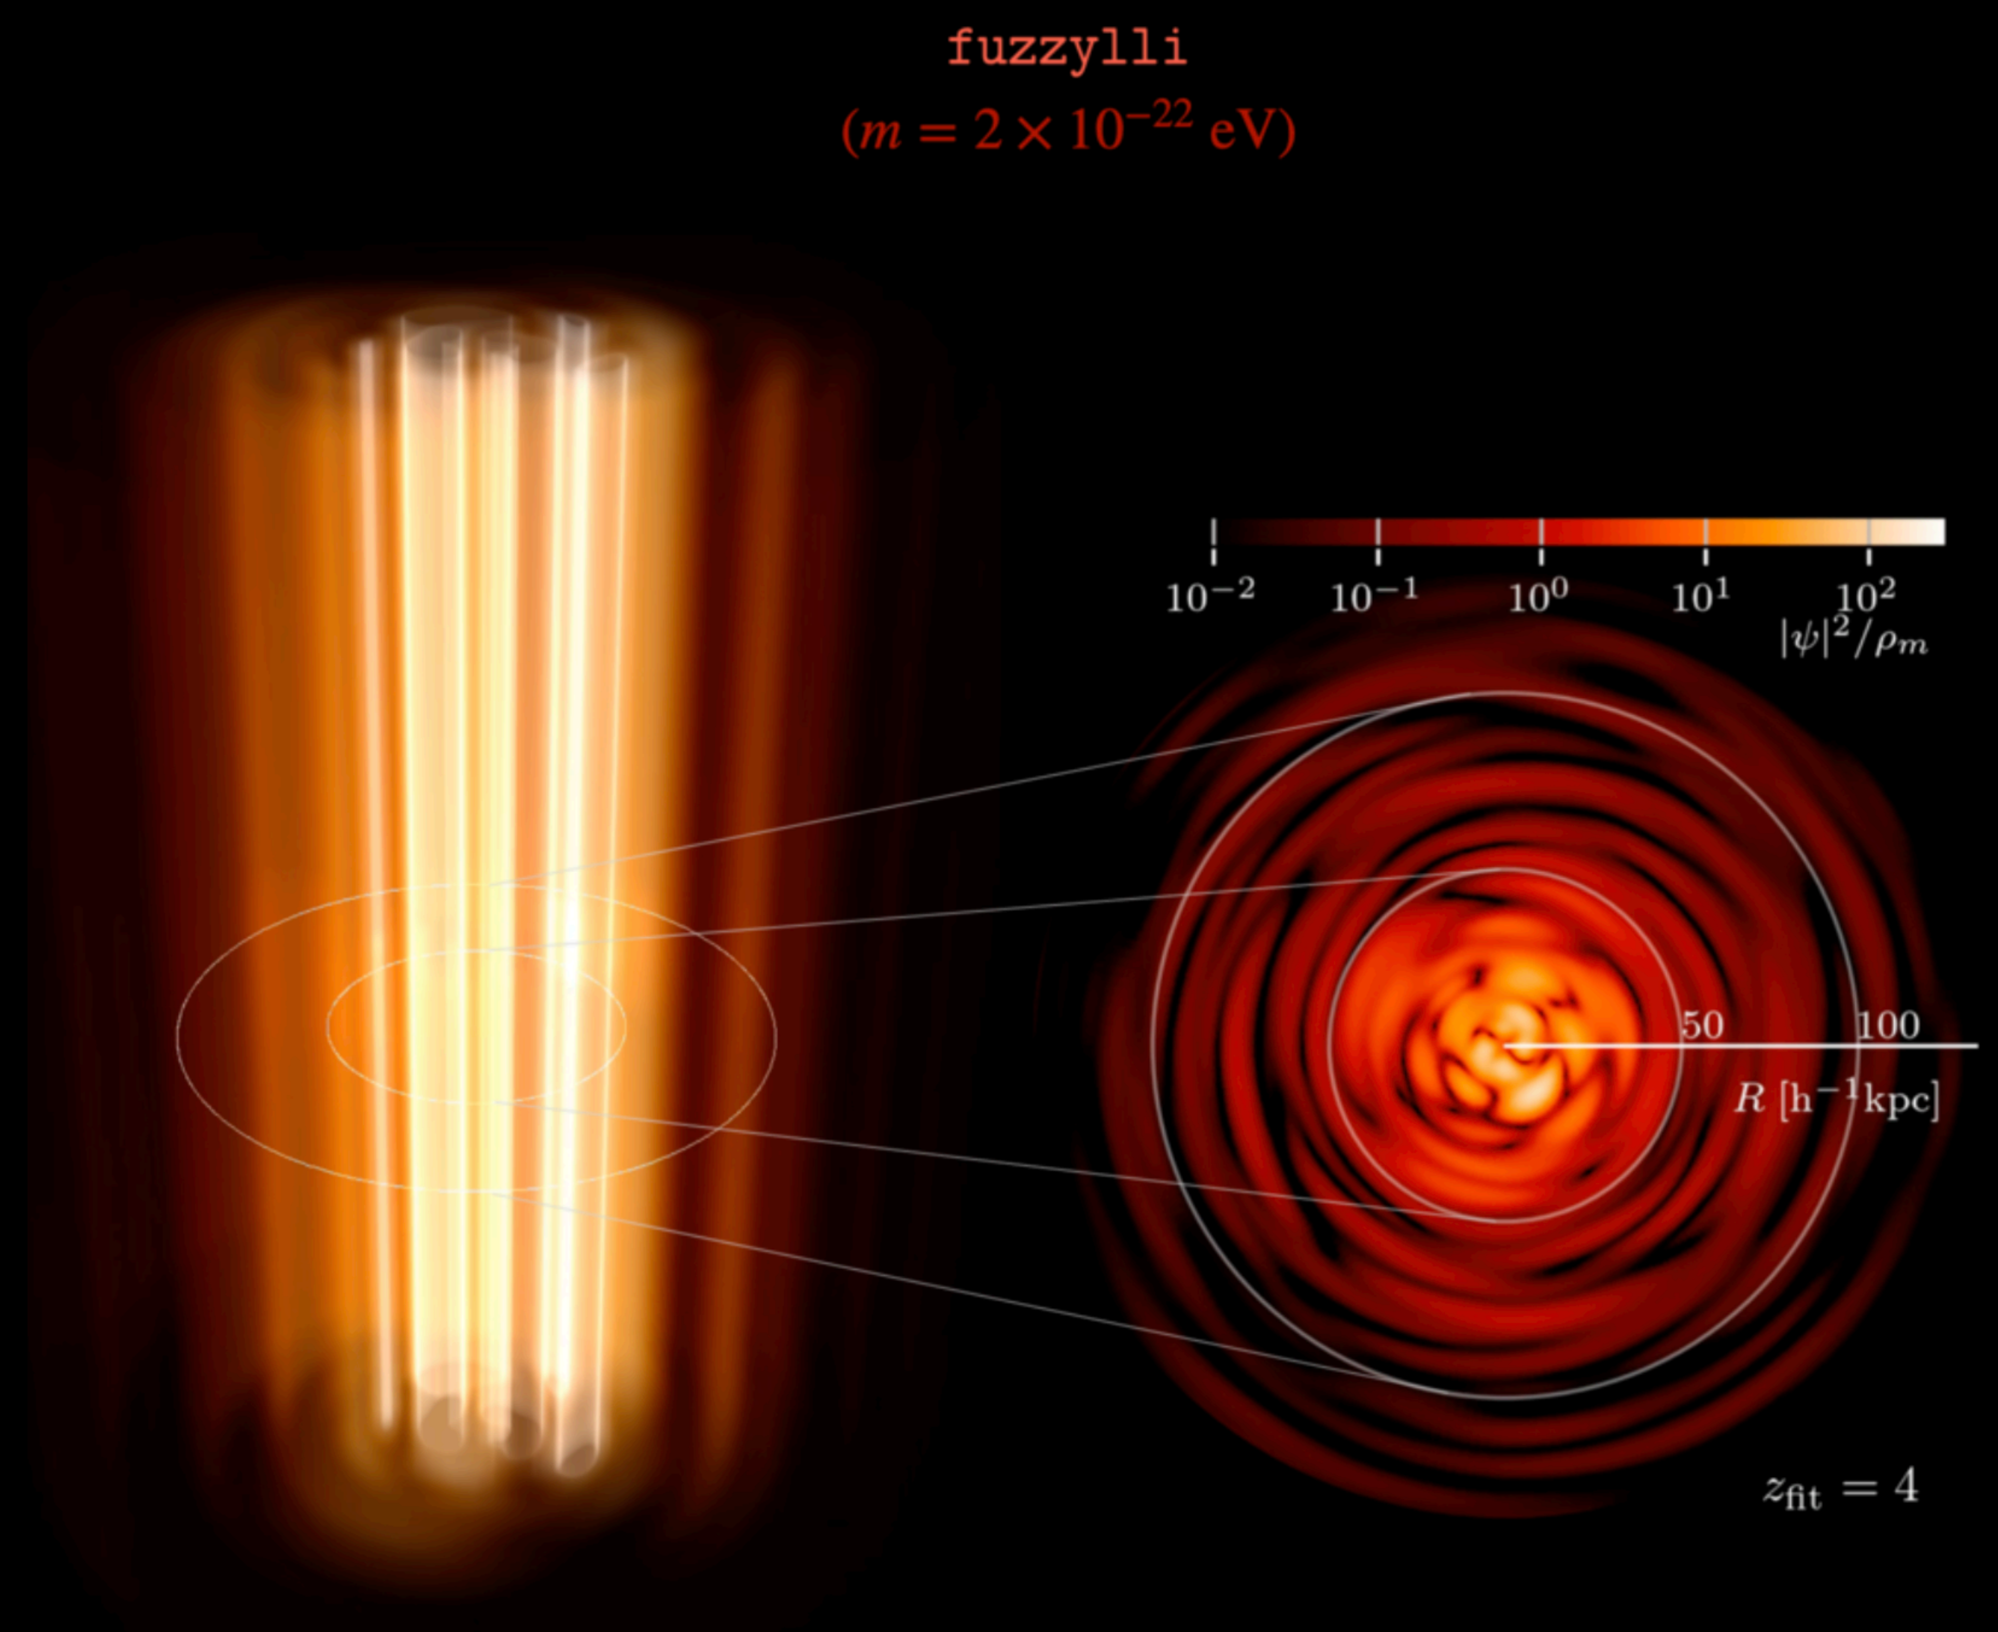

News from the Dark 10

Hans Winther, ITA, Oslo 10 September 2025

Introduction

$$i\hbar\partial_t\psi = \left[-\frac{\hbar^2}{2ma^2(t)}\triangle + mV(\boldsymbol{x},t)\right]\psi$$

$$\triangle V = \frac{4\pi G}{a(t)}\left(|\psi|^2 - 1\right) .$$

- Ultra light axions is a dark matter candidate (a free scalar field) whose mass is, as the name suggests, very light: <~1e-22 eV and has many intriguing signatures in the cosmic web: cores inside halos, interference patters in the DM distribution and many more.

- Many of these signatures require simulations to study, but these are very expensive to perform. Main problem in that the Schrödinger equation determining its evolution is first order in time and second order in space so we need to satisfy the stability constraint dt < dx^2.

- The characteristic size we need to resolve is the so-called de-Broglie scale which is very small (~kpc) for typical axion masses studied. Together this means we require very small time steps and often very small simulation boxes.

- Some methods are faster than others and can do larger volumes (e.g. SPH, but it can't do interference). There have been many developments in hybrid method that can solve some of these issues.

- Simulations of axions are great… and even though the title said I'd speak about it I will today instead talk about some works that are *related* to it. Still numerically though, but just some things we can do without having to perform full-fledged simulations explicitly to study these models.

Paper I :

Interference in Fuzzy Dark Matter Filaments

Idealised Models and Statistics

T. Zimmermann[1], D. J.E. Marsh[2], K. K. Rogers[3,4], H. A. Winther[1], and S. Shen[1]

- Extended interference fringes in cosmic filaments is one key signature of ultra light axions / fuzzy dark matter (FDM)

- A detailed understanding of this may strengthen existing limits on the boson mass but also break the degeneracy with warm dark matter, and provide a unique fingerprint of interference in cosmology.

- Aim of this work is to we build a theoretically motivated steady-state approximation for filaments and express the equilibrium dynamics of such in an expansion of FDM eigenstates

- Allows us to study filaments and interference without having to run expensive simulations

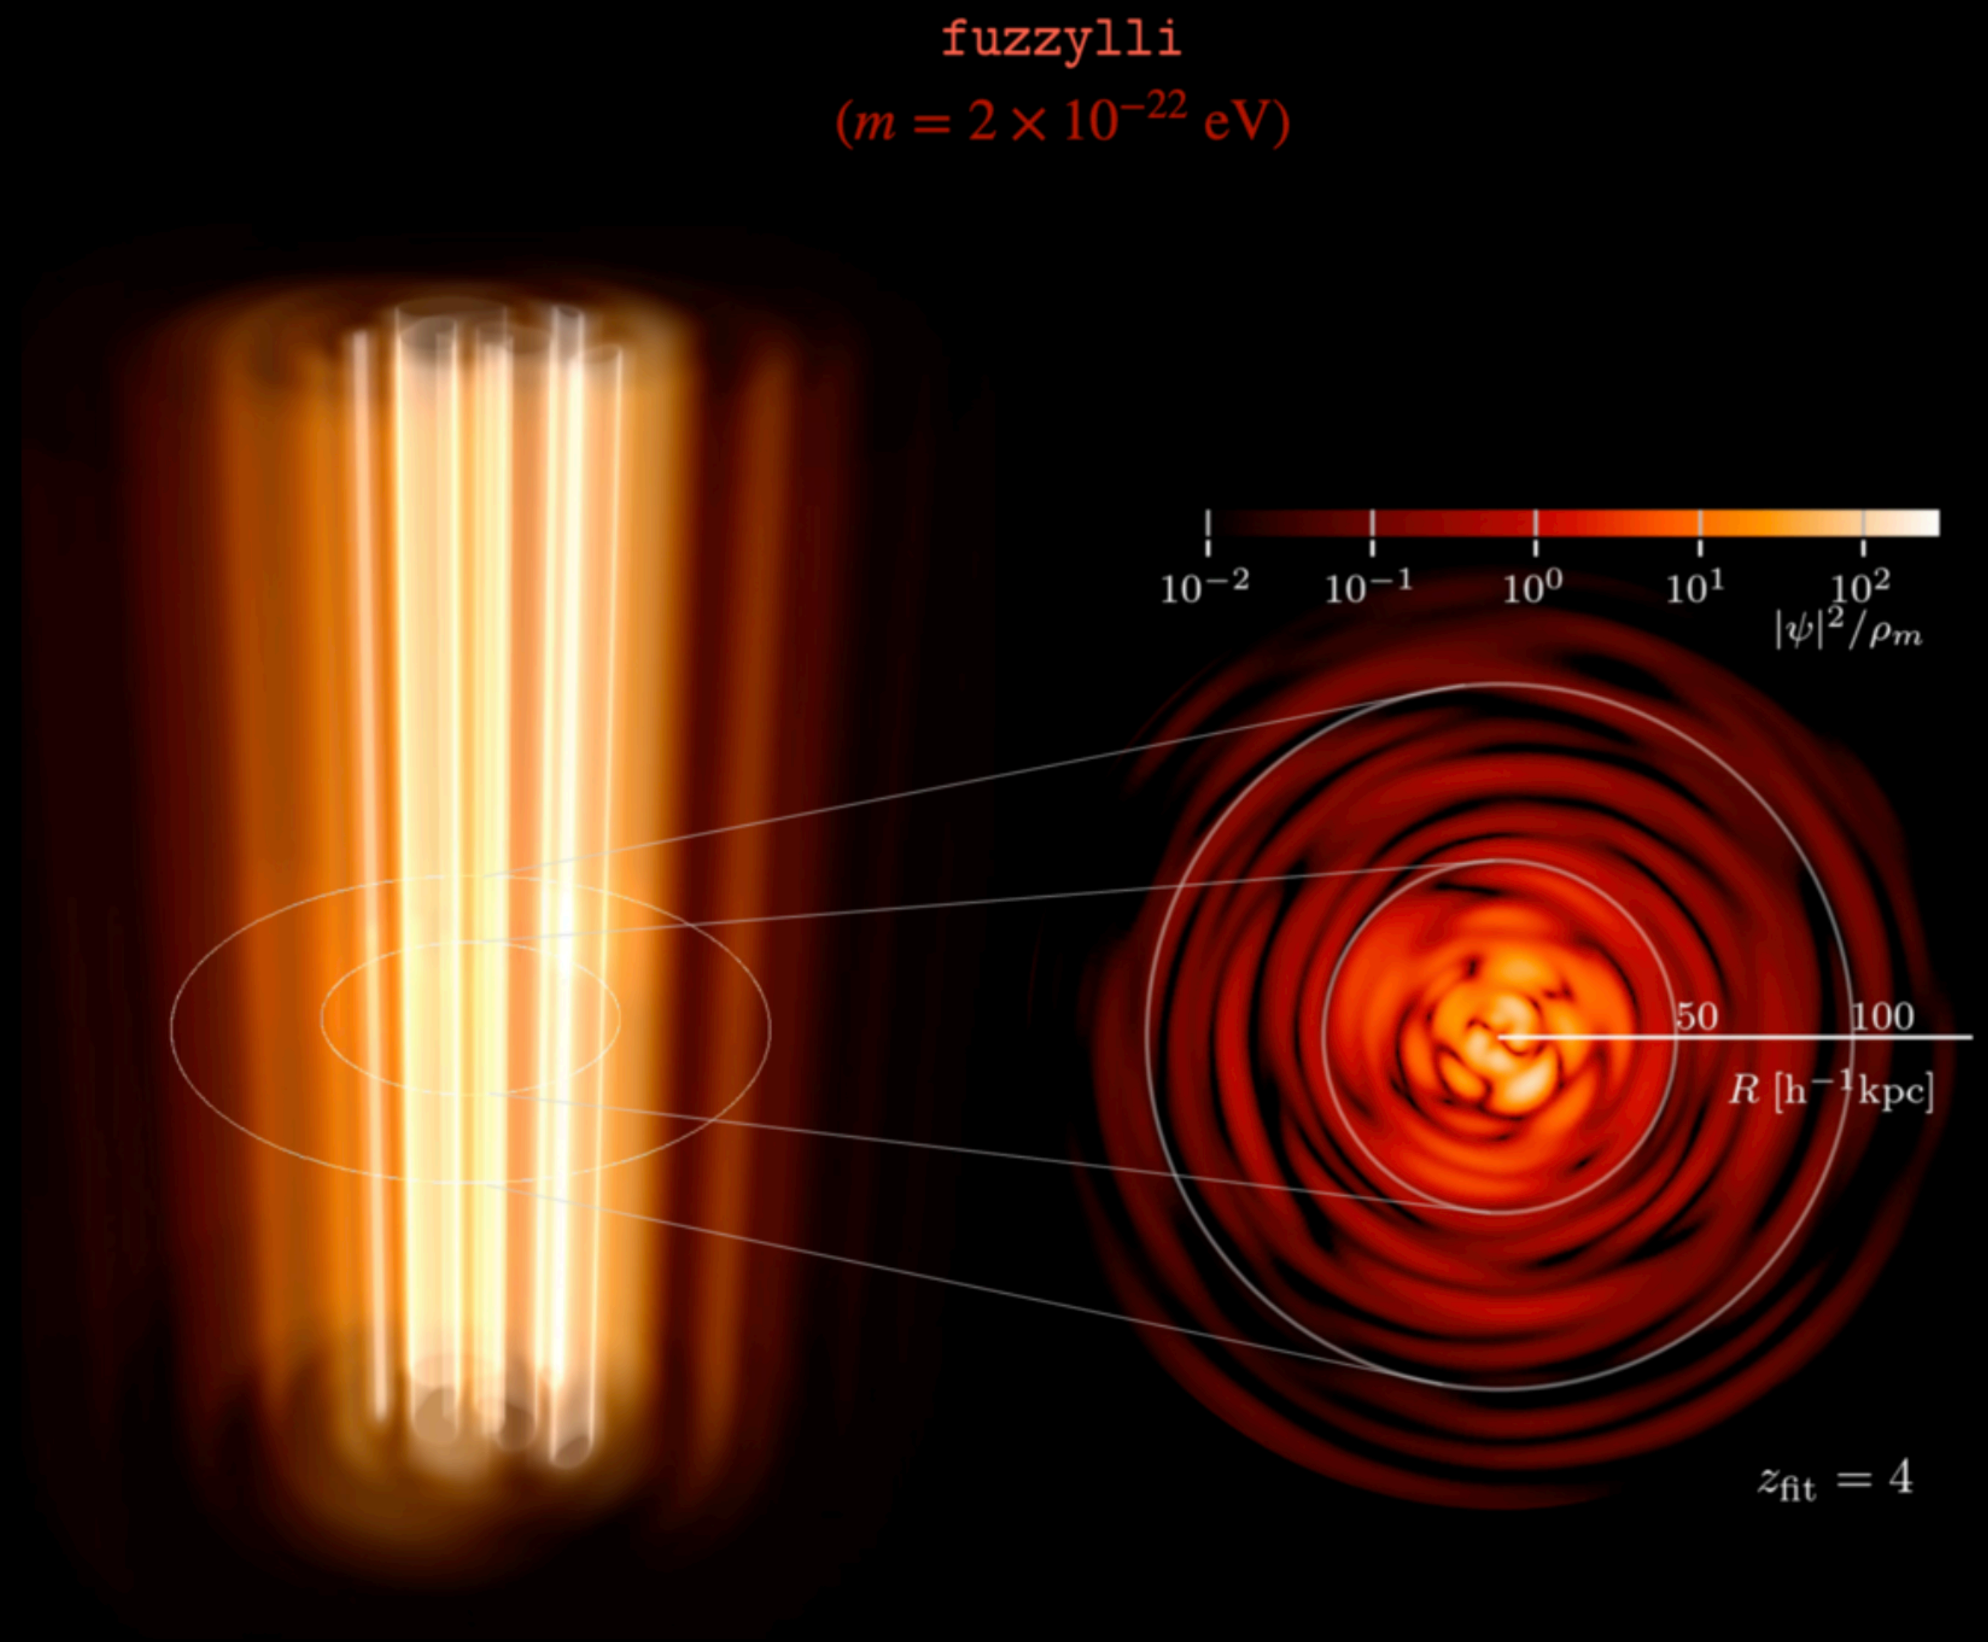

Self-consistent FDM Interference in Steady-State Systems

- FDM is governed by the Schrödinger-Poisson equation:

$$i\hbar\partial_t\psi = \left[-\frac{\hbar^2}{2ma^2(t)}\Delta + mV(\boldsymbol{x},t)\right]\psi$$
$$\Delta V = \frac{4\pi G}{a(t)}\left(|\psi|^2 - 1\right) .$$

- We assume the detailed dynamics of $\psi(x, t)$ take place in a smooth gravitational potential that is effectively static in time

$$i\hbar\partial_t\psi = \left[-\frac{\hbar^2}{2ma^2}\Delta + mV(\boldsymbol{x})\right]\psi$$
$$\Delta V = \frac{4\pi G}{a}\rho_{\mathrm{BG}}(\boldsymbol{x}) ,$$

- Integrating in time then reduces to diagonalising the Hamiltonian and expanding $\psi(x, t)$ in the eigenbasis $\psi_j(x)$ with energy eigenvalues E_j

$$\psi(\boldsymbol{x},t) = \sum_j a_j\psi_j(\boldsymbol{x})e^{iE_jt/\hbar}$$

- Interference emerges as the cross-terms in the *time-dependent* density! Important for later: even if each eigenstate obeys the same symmetries as the steady-state background (cylindrical for our filaments) $|\psi(x, t)|$, the interference term on the right does not.

$$|\psi(\boldsymbol{x},t)|^2 = \sum_j |a_j|^2|\psi_j(\boldsymbol{x})|^2 + \sum_{j\neq k} a_ja_k^*\psi_j(\boldsymbol{x})\psi_k^*(\boldsymbol{x})e^{i(E_k-E_j)t/\hbar}$$

- Given a static density profile $\rho_{BG}(x)$ we need to find complex coefficients a_j such that the time-independent term recovers this static background density $\rho_{BG}(x)$.

$$\rho_{\mathrm{BG}}(\boldsymbol{x}) \stackrel{!}{=} \sum_j |a_j|^2|\psi_j(\boldsymbol{x})|^2 = \left\langle|\psi|^2\right\rangle$$

Idealised model for filaments

- As always in physics, we start with choosing the simplest approach: filaments are mathematically modelled as infinite long, isolated and isothermal cylinders.
- Quasi-virialized: in it each cross section we assume a steady-state is attained, virial equilibrium and dynamics along the longitudinal direction are suppressed
- These kinds of filaments are consistent with some observations of filaments and also a good model (for certain things related to filaments) in CDM simulation.
- Importantly: it is also consistent with what is found in FDM simulations!

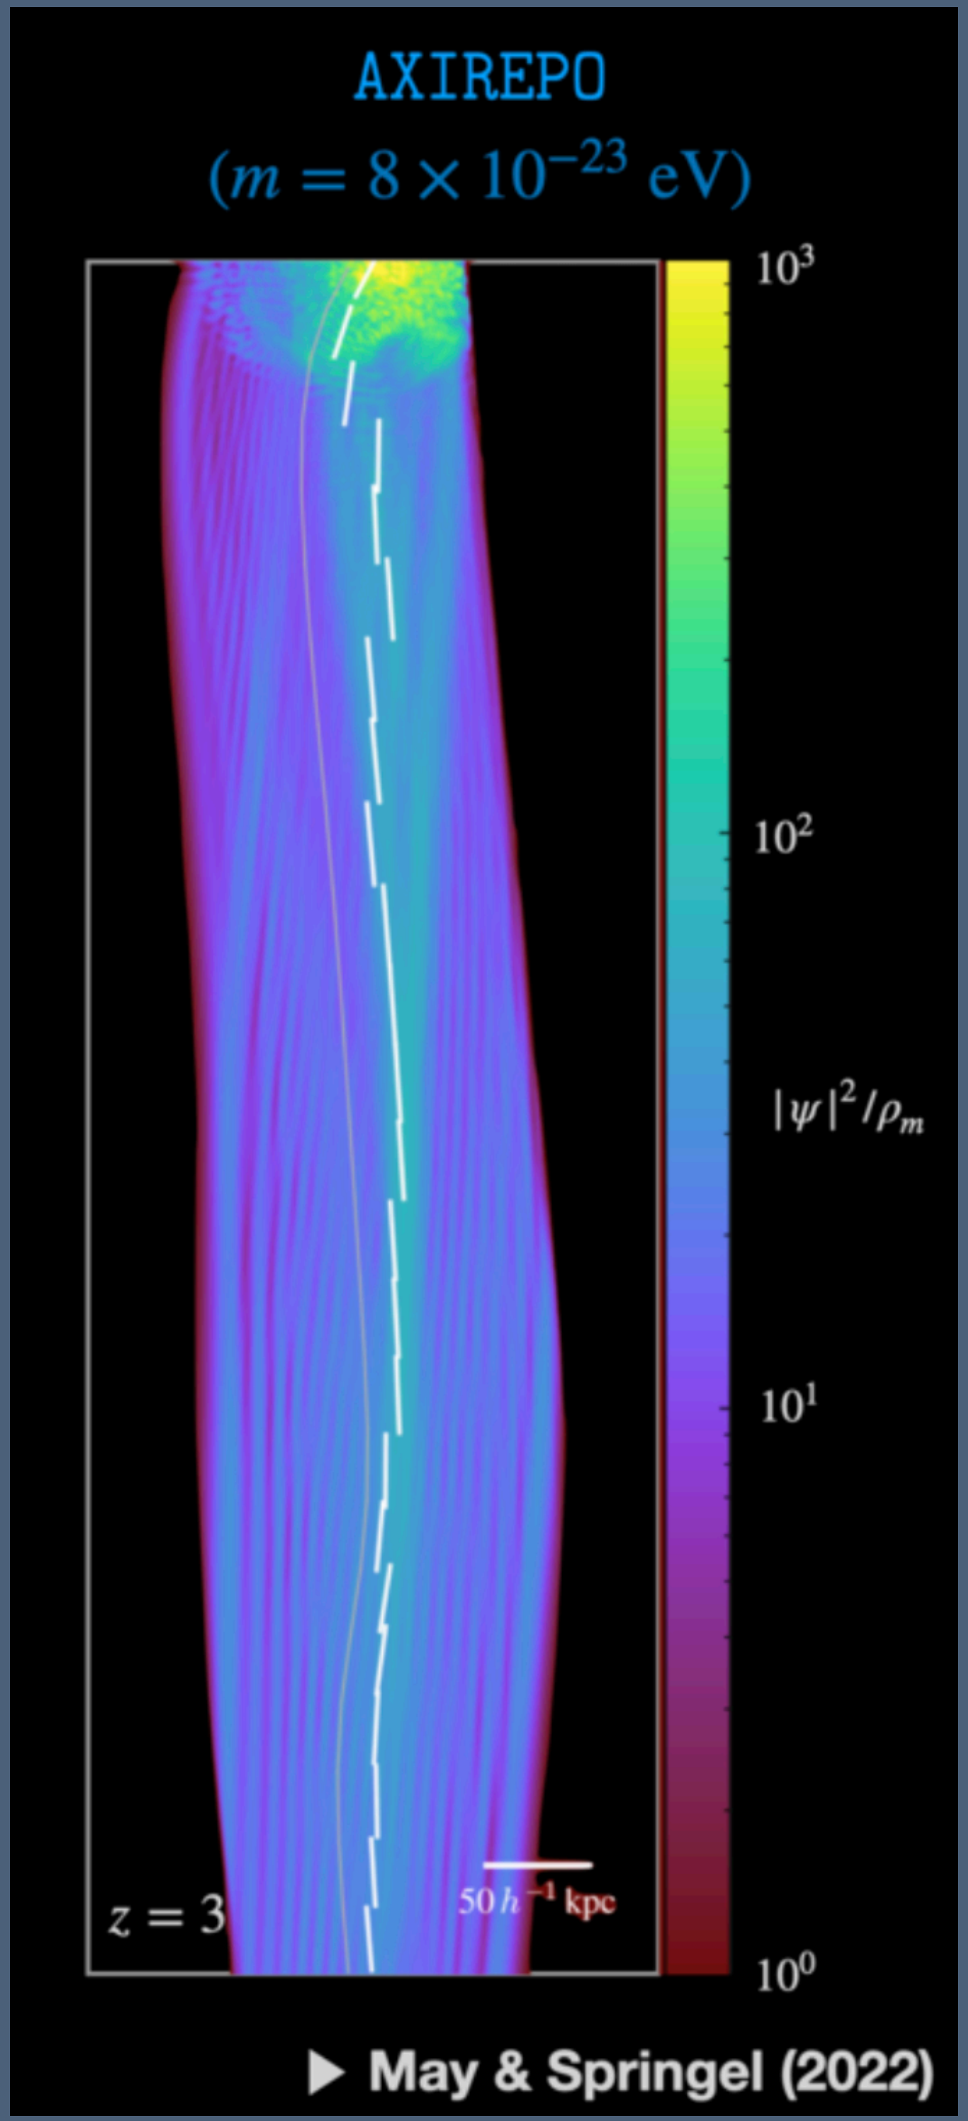

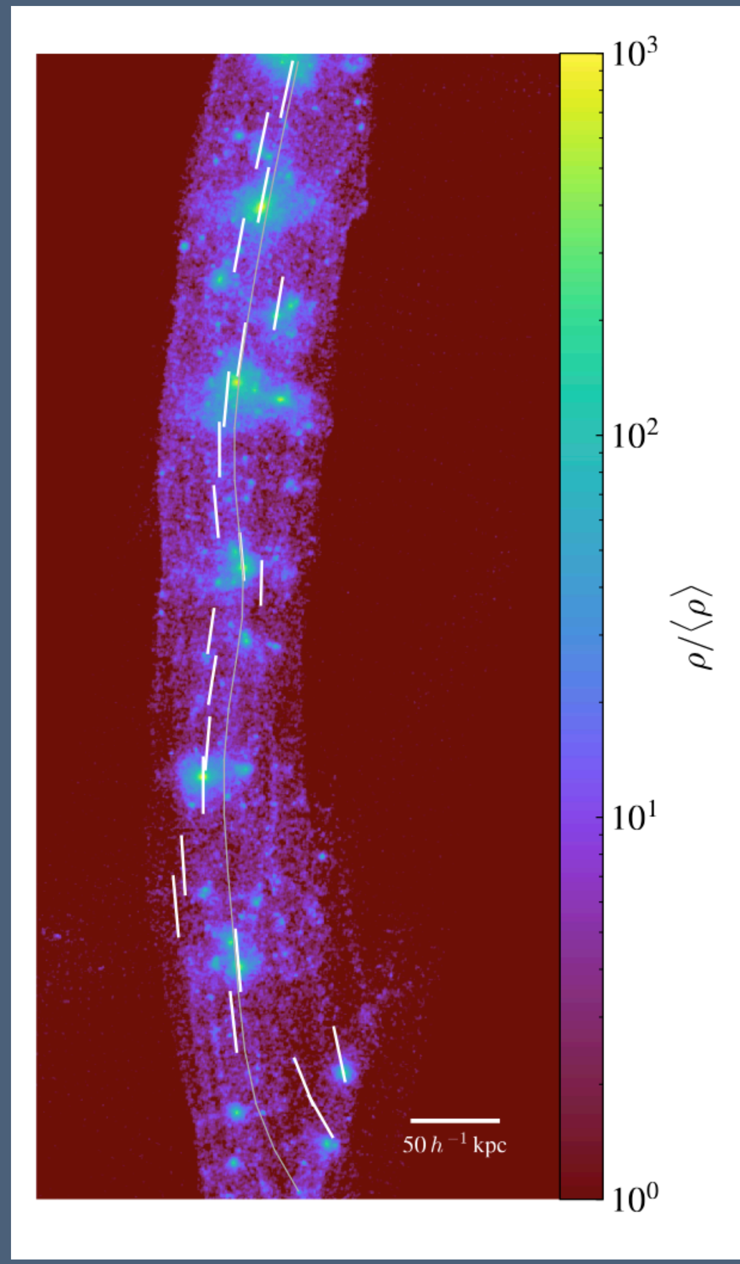

Real space density

- On scales larger than λdB the Schrödinger-Vlassov correspondence allows us to treat FDM as collisionless particles. We let the behaviour of these particles be described by a velocity dispersion tensor of the form

$$\sigma^2 = \mathrm{Diag}\left(\langle v_r^2\rangle, \langle v_\phi^2\rangle, 0\right)$$

- Isothermal assumption: the radial moment, ⟨vr^2⟩, the azimuthal moment, ⟨vϕ^2⟩, and the anisotropy parameter β are constant throughout the filament

$$\beta \equiv 1 - \frac{\langle v_\phi^2\rangle}{\langle v_r^2\rangle}$$

- From the cylindrical Jeans equation we can then derive a closed form, analytic solutions for the density

$$\rho(r \mid r_0, \sigma^2, \beta) = \frac{(2-\beta)^2\sigma^2}{2\pi G r_0^2}\frac{y^{-\beta}}{(y^{2-\beta}+1)^2}\,, \quad y \equiv \frac{r}{r_0}$$

- Consistent with simulations. May & Springer 2022 find filaments have cuspy profiles. This can be achieved by assuming a radially biased velocity anisotropy β > 0.

Our density profile

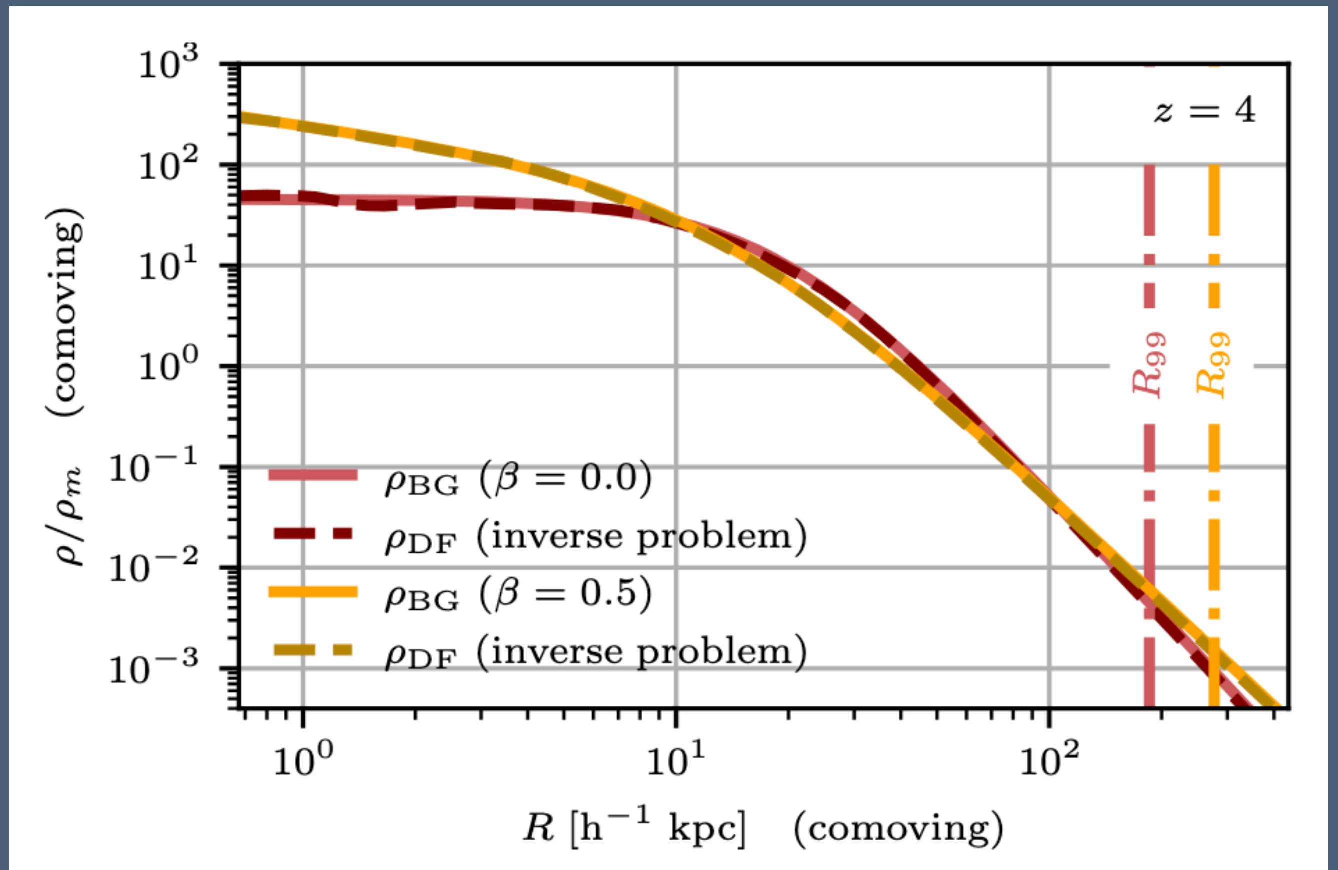

FDM simulations

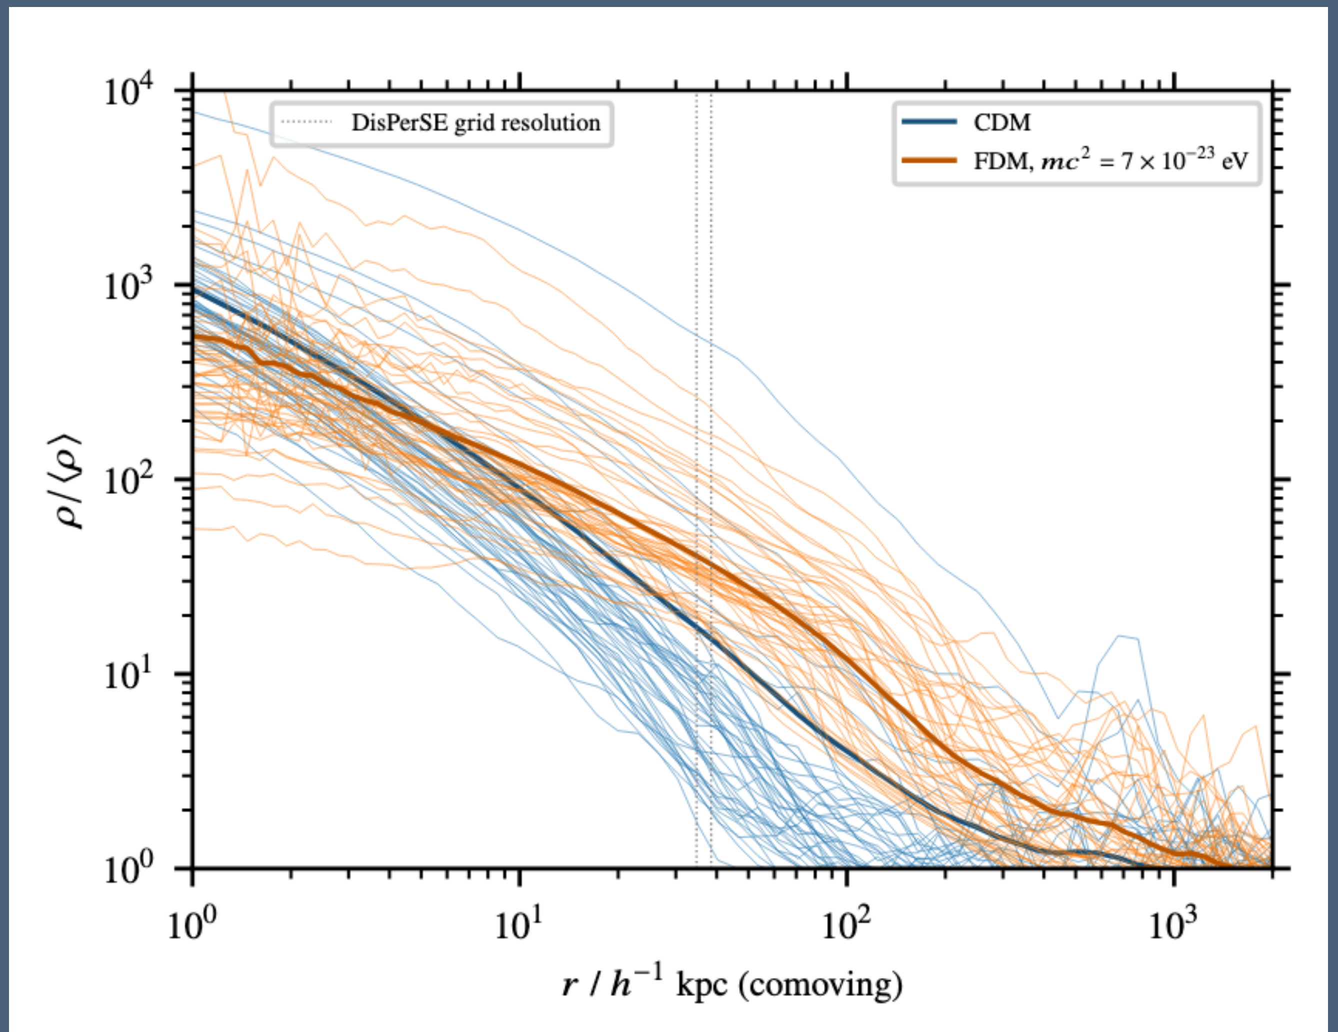

Phase space distribution

- According to Jeans' theorem, we may assume that a steady-state DF is a function of three isolating integrals of motion only (the specific energy contribution of the longitudinal, Ez, and transversal motion, E, and the specific angular momentum L)
- Jeans theorem then suggests the ansatz
- The real-space density is then given by (assuming a constant anisotropy):
- We know V, and the density profile and can then invert this to obtain the distribution function f

$$E = \frac{1}{2a^2}(u_R^2 + u_\Phi^2) + V(R)\,, \quad E_Z = \frac{u_Z^2}{2a^2}\,, \quad L = Ru_\Phi\,.$$

$$\tilde{f}(E, E_Z, L) = \mathcal{N} f(E, |L|)\delta_D\left(\frac{u_Z^2}{2a^2} - E_Z\right) \quad \text{with}$$

$$\mathcal{N} \equiv \mu\left(\int \mathrm{d}R\,\mathrm{d}\Phi\,\mathrm{d}^3\boldsymbol{u}\,R\tilde{f}(E, E_Z, |L|)\right)^{-1},$$

$$\rho_{\rm BG}(R) = 2a^2\sqrt{\pi}\frac{\Gamma\left(\frac{1-\beta}{2}\right)}{\Gamma\left(\frac{2-\beta}{2}\right)}R^{-\beta}\int_{V(R)}^{\infty}\mathrm{d}E\frac{g(E)}{[2a^2(E - V(R))]^{\beta/2}}$$

$$f(E, |L|) = |L|^{-\beta}g(E)$$

Eigenstate library

- We may realise an infinitely long, cylindrical geometry by solving the Schrodinger-Poisson system in cylindrical symmetry. Factorizing it in eigenmode we need to solve:
- Solving this numerically results in a dense matrix representation of the Hamiltonian
- More tricky than the spherical case, but doable. To fully characterise the wave-function we need 100-1000 eigenstates so its quite a bit of computation.

$$\psi_j(\boldsymbol{x}) = \psi_{nl}(R,\Phi) = \sqrt{\frac{\mu}{2\pi}} R_{nl}(R) e^{il\Phi}\,, \quad \int_0^\infty \mathrm{d}R\, R\, R_{nl}^2(R) = 1,$$

$$E_{nl} u_{nl} = -\frac{\hbar^2}{2m^2a^2}\partial_R^2 u_{nl} + V_{\mathrm{eff},\psi} u_{nl}\,,$$

$$V_{\mathrm{eff},\psi}(R) = V(R) + \frac{\hbar^2}{2m^2a^2}\frac{l^2 - 1/4}{R^2}\,,$$

$$\frac{1}{R}\frac{\partial}{\partial R}\left(R\frac{\partial}{\partial R}V\right) = \frac{4\pi G}{a}\rho_{\mathrm{BG}}(R)\,,$$

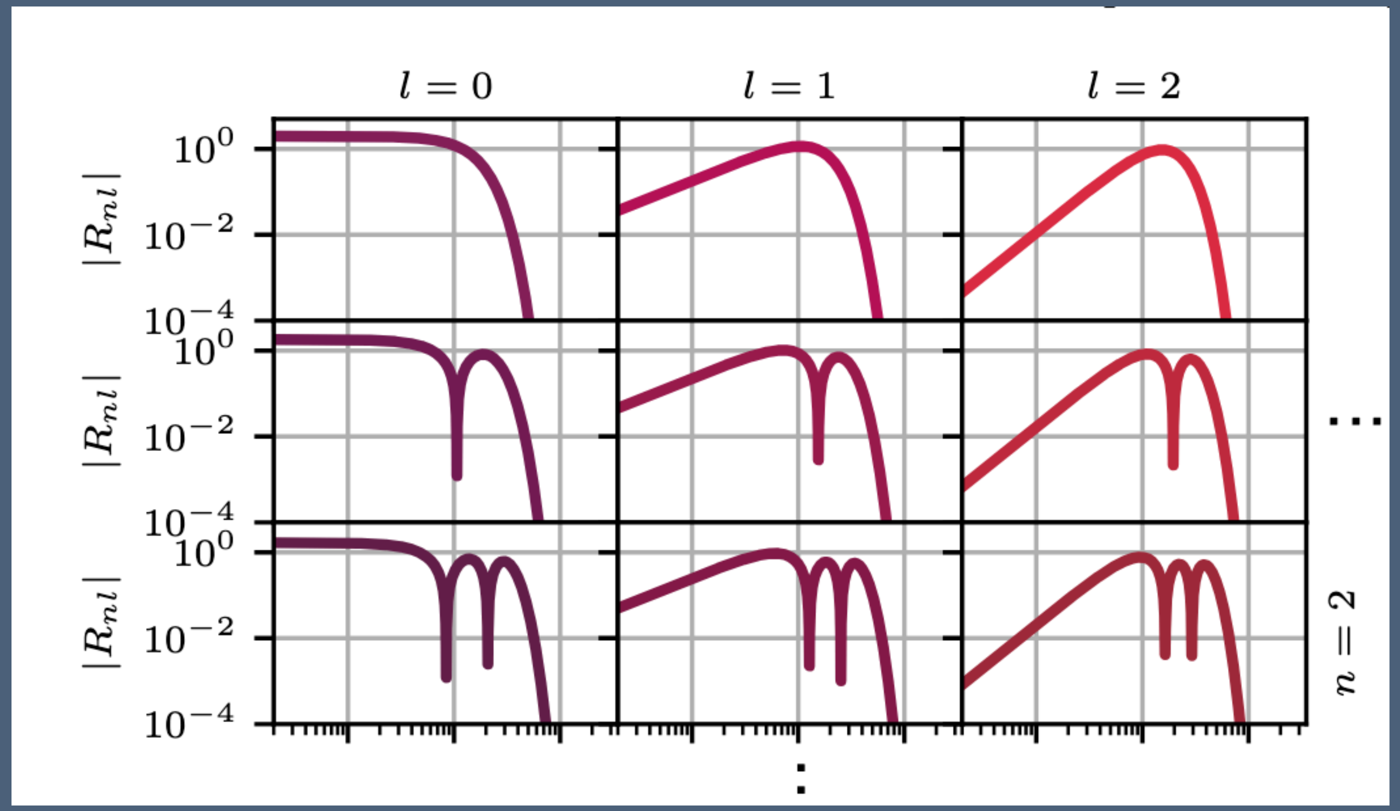

Wave-function reconstruction

- Now that we have the density profile, the distribution function and a large set of eigenmodes we can finally perform a wave-function reconstruction from

$$\rho_{BG}(R) = \frac{\mu}{2\pi} \sum_{n,l\geq 0} N_l |a_{nl}|^2 |R_{nl}(R)|^2 \ , \ N_l = \begin{cases} 1, & l = 0 \\ 2, & l > 0 \end{cases}$$

- There exist multiple approaches for computing the coefficients anl (e.g. WKB, numerical fitting, …). E.g. |anl|^2 may be interpreted as the probability of finding state |ψnl|^2. The value of the distribution function represents this classical probability so we should expect (and WKB indeed gives us this):

$$|a_{nl}| \propto f(E_{nl}, L_{nl}) \Delta E \Delta L$$

- How we did all this in practice is technical so I will skip it due to time-constraints, but Tim found a way of significantly reduce the amount of modes needed and still get accurate results.

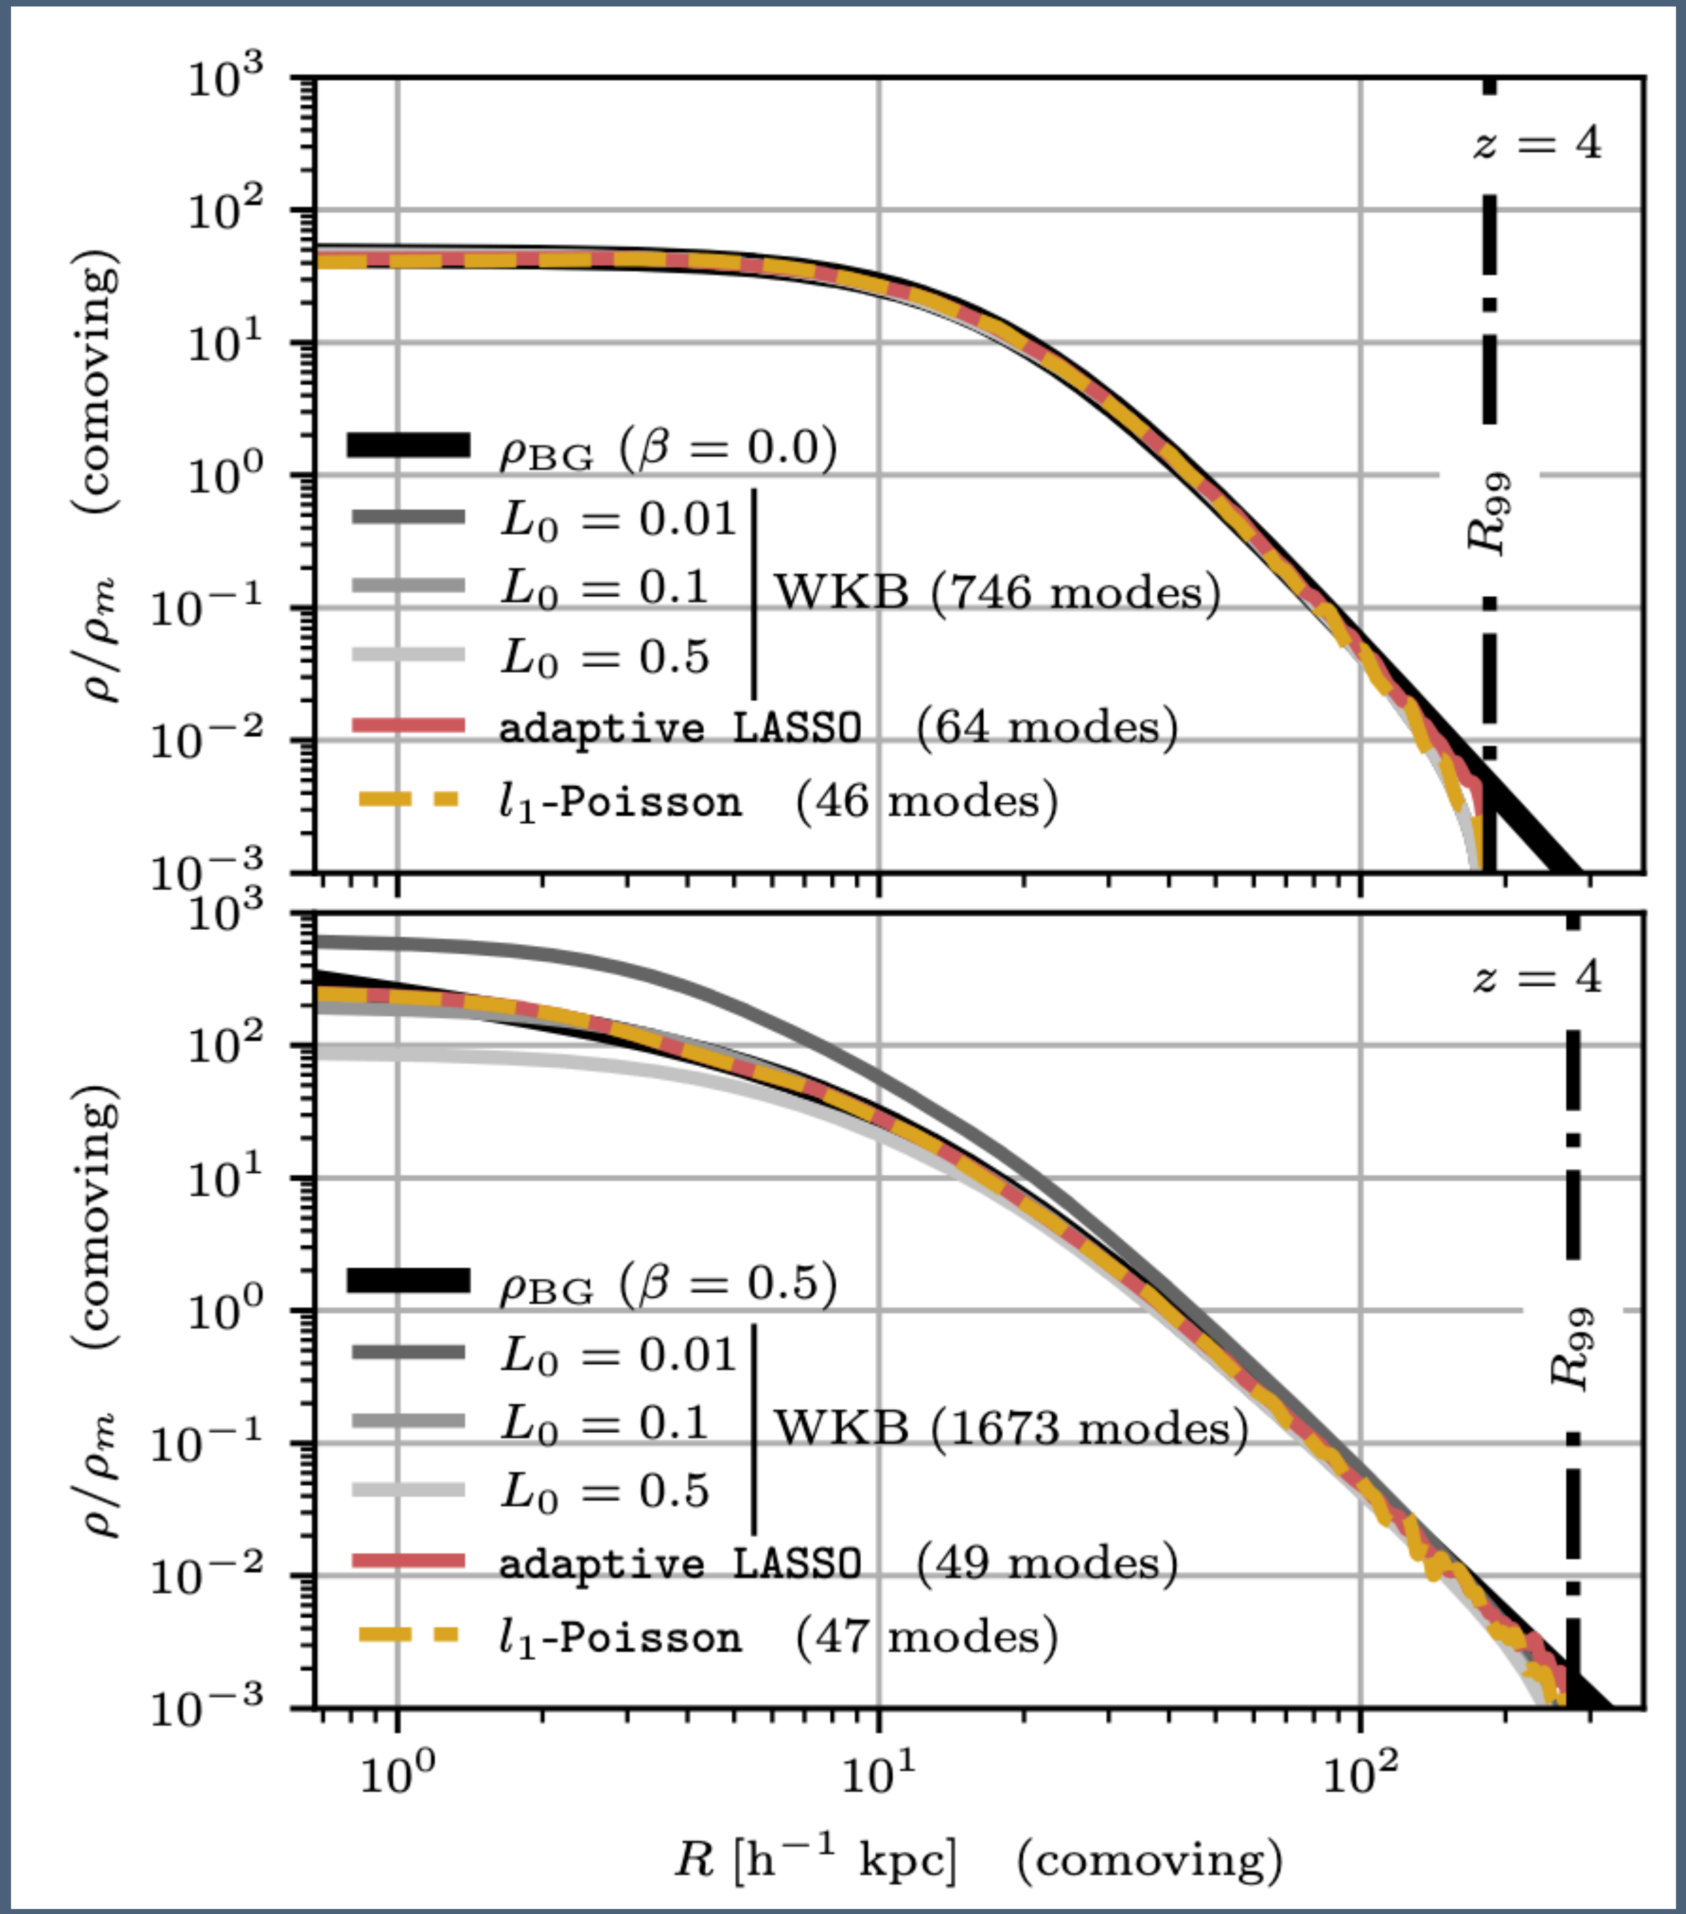

Results

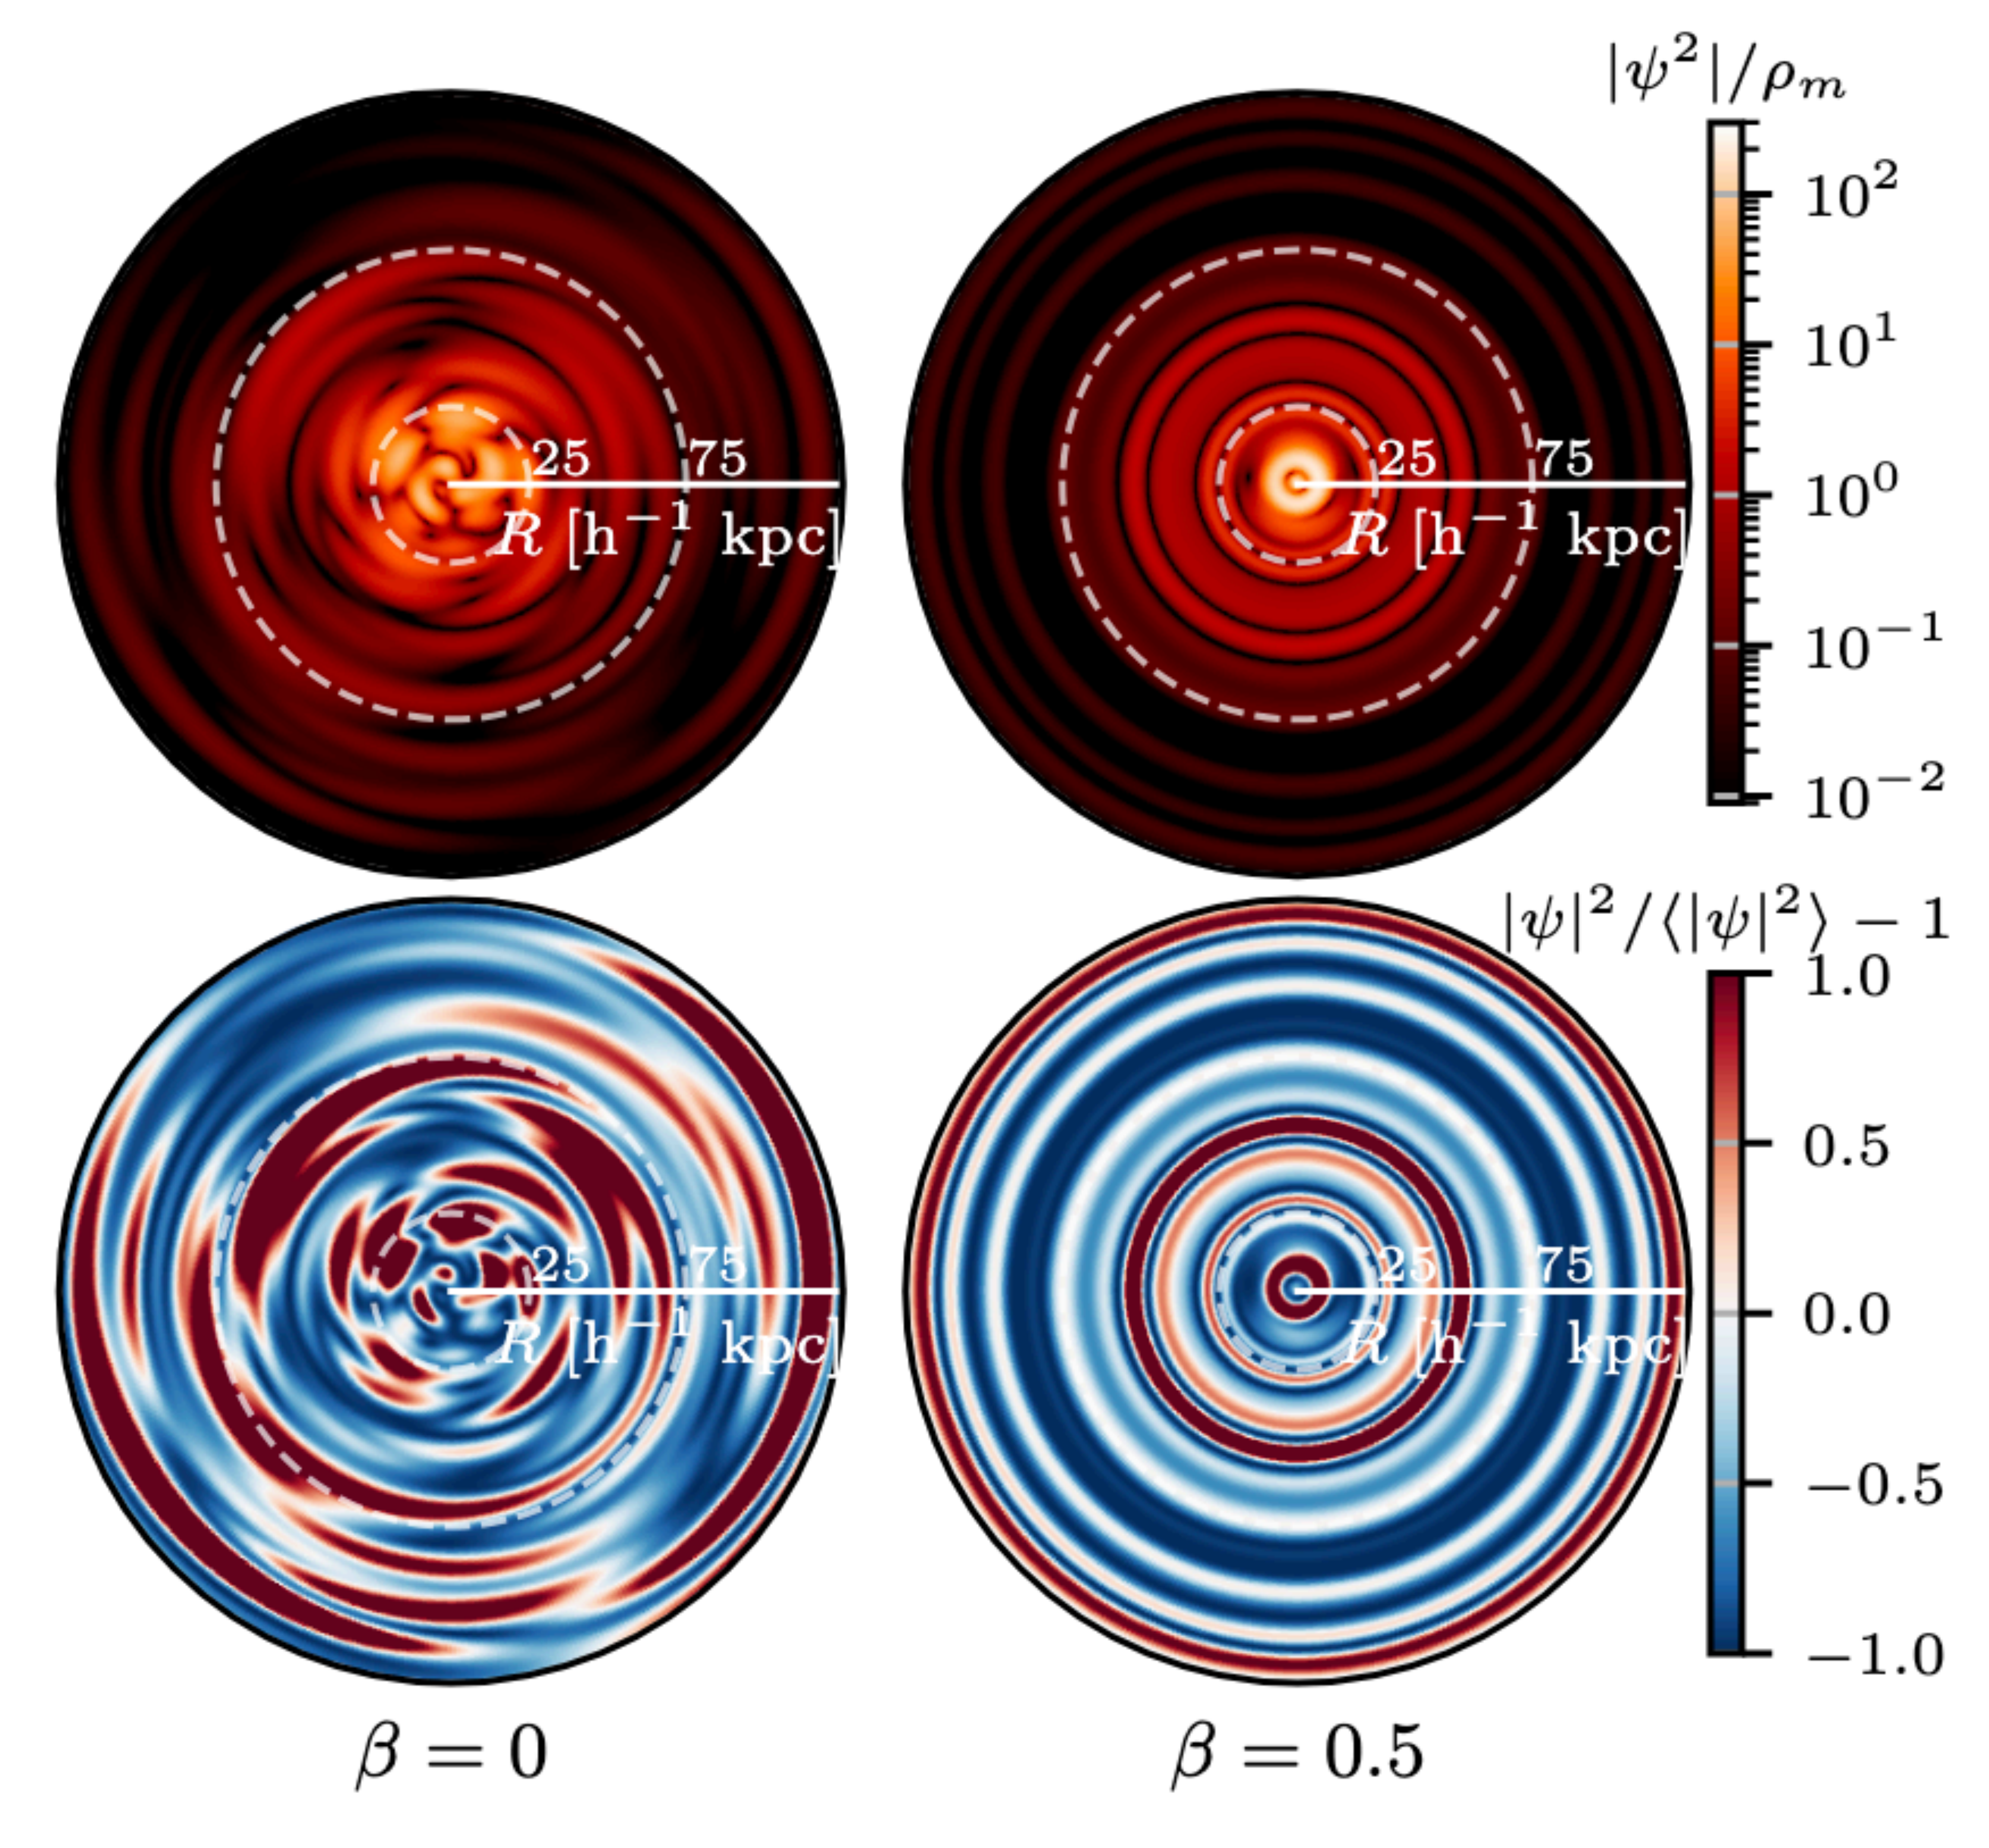

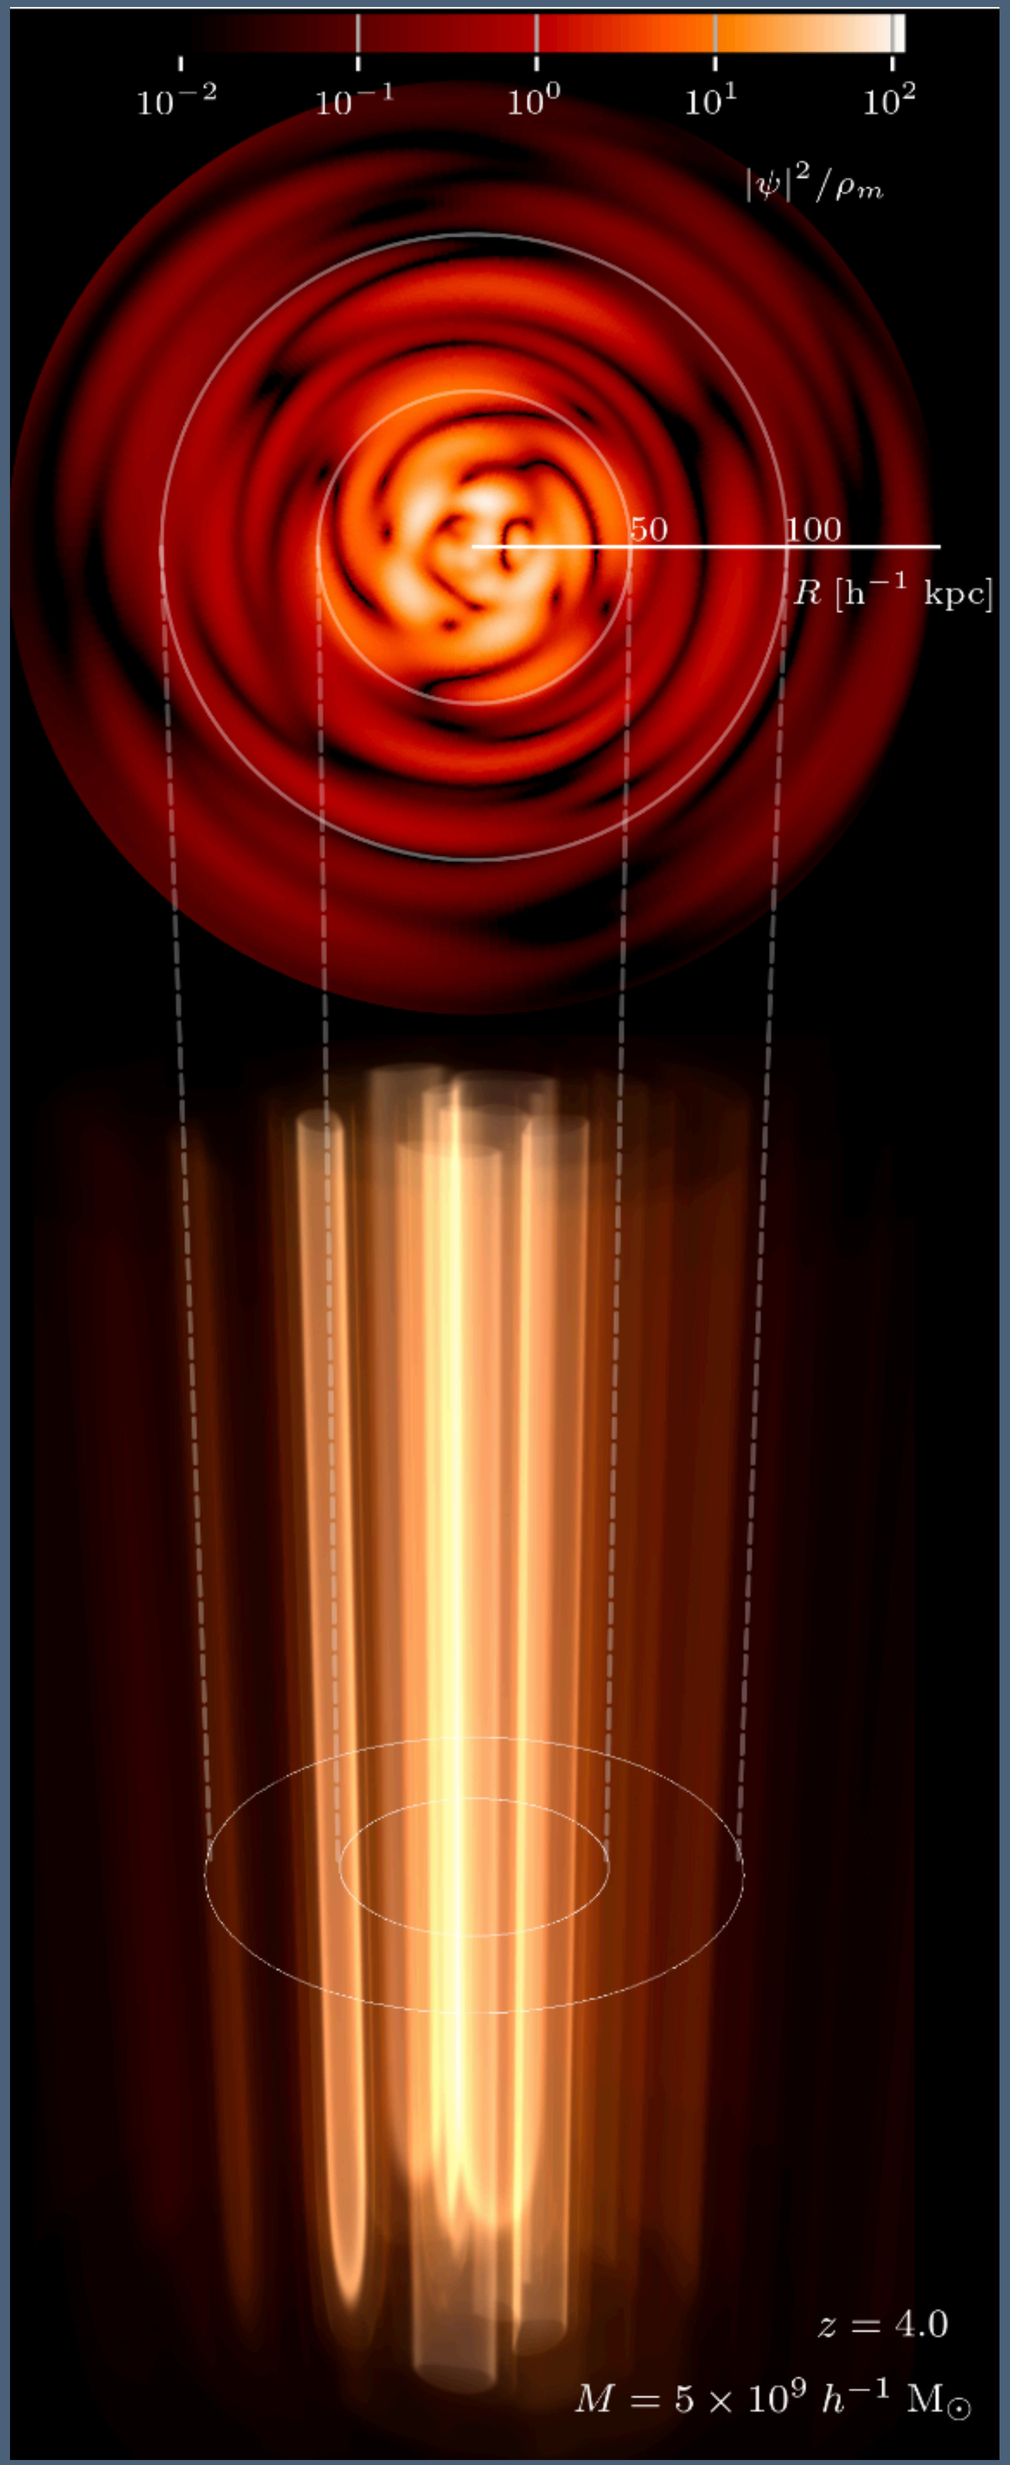

Interference modulates the density comparable to the magnitude of the background density leading to a concentric ring pattern for the radially biased case (right) and a more mixed interference fringe configuration under isotropic conditions (left).

Results

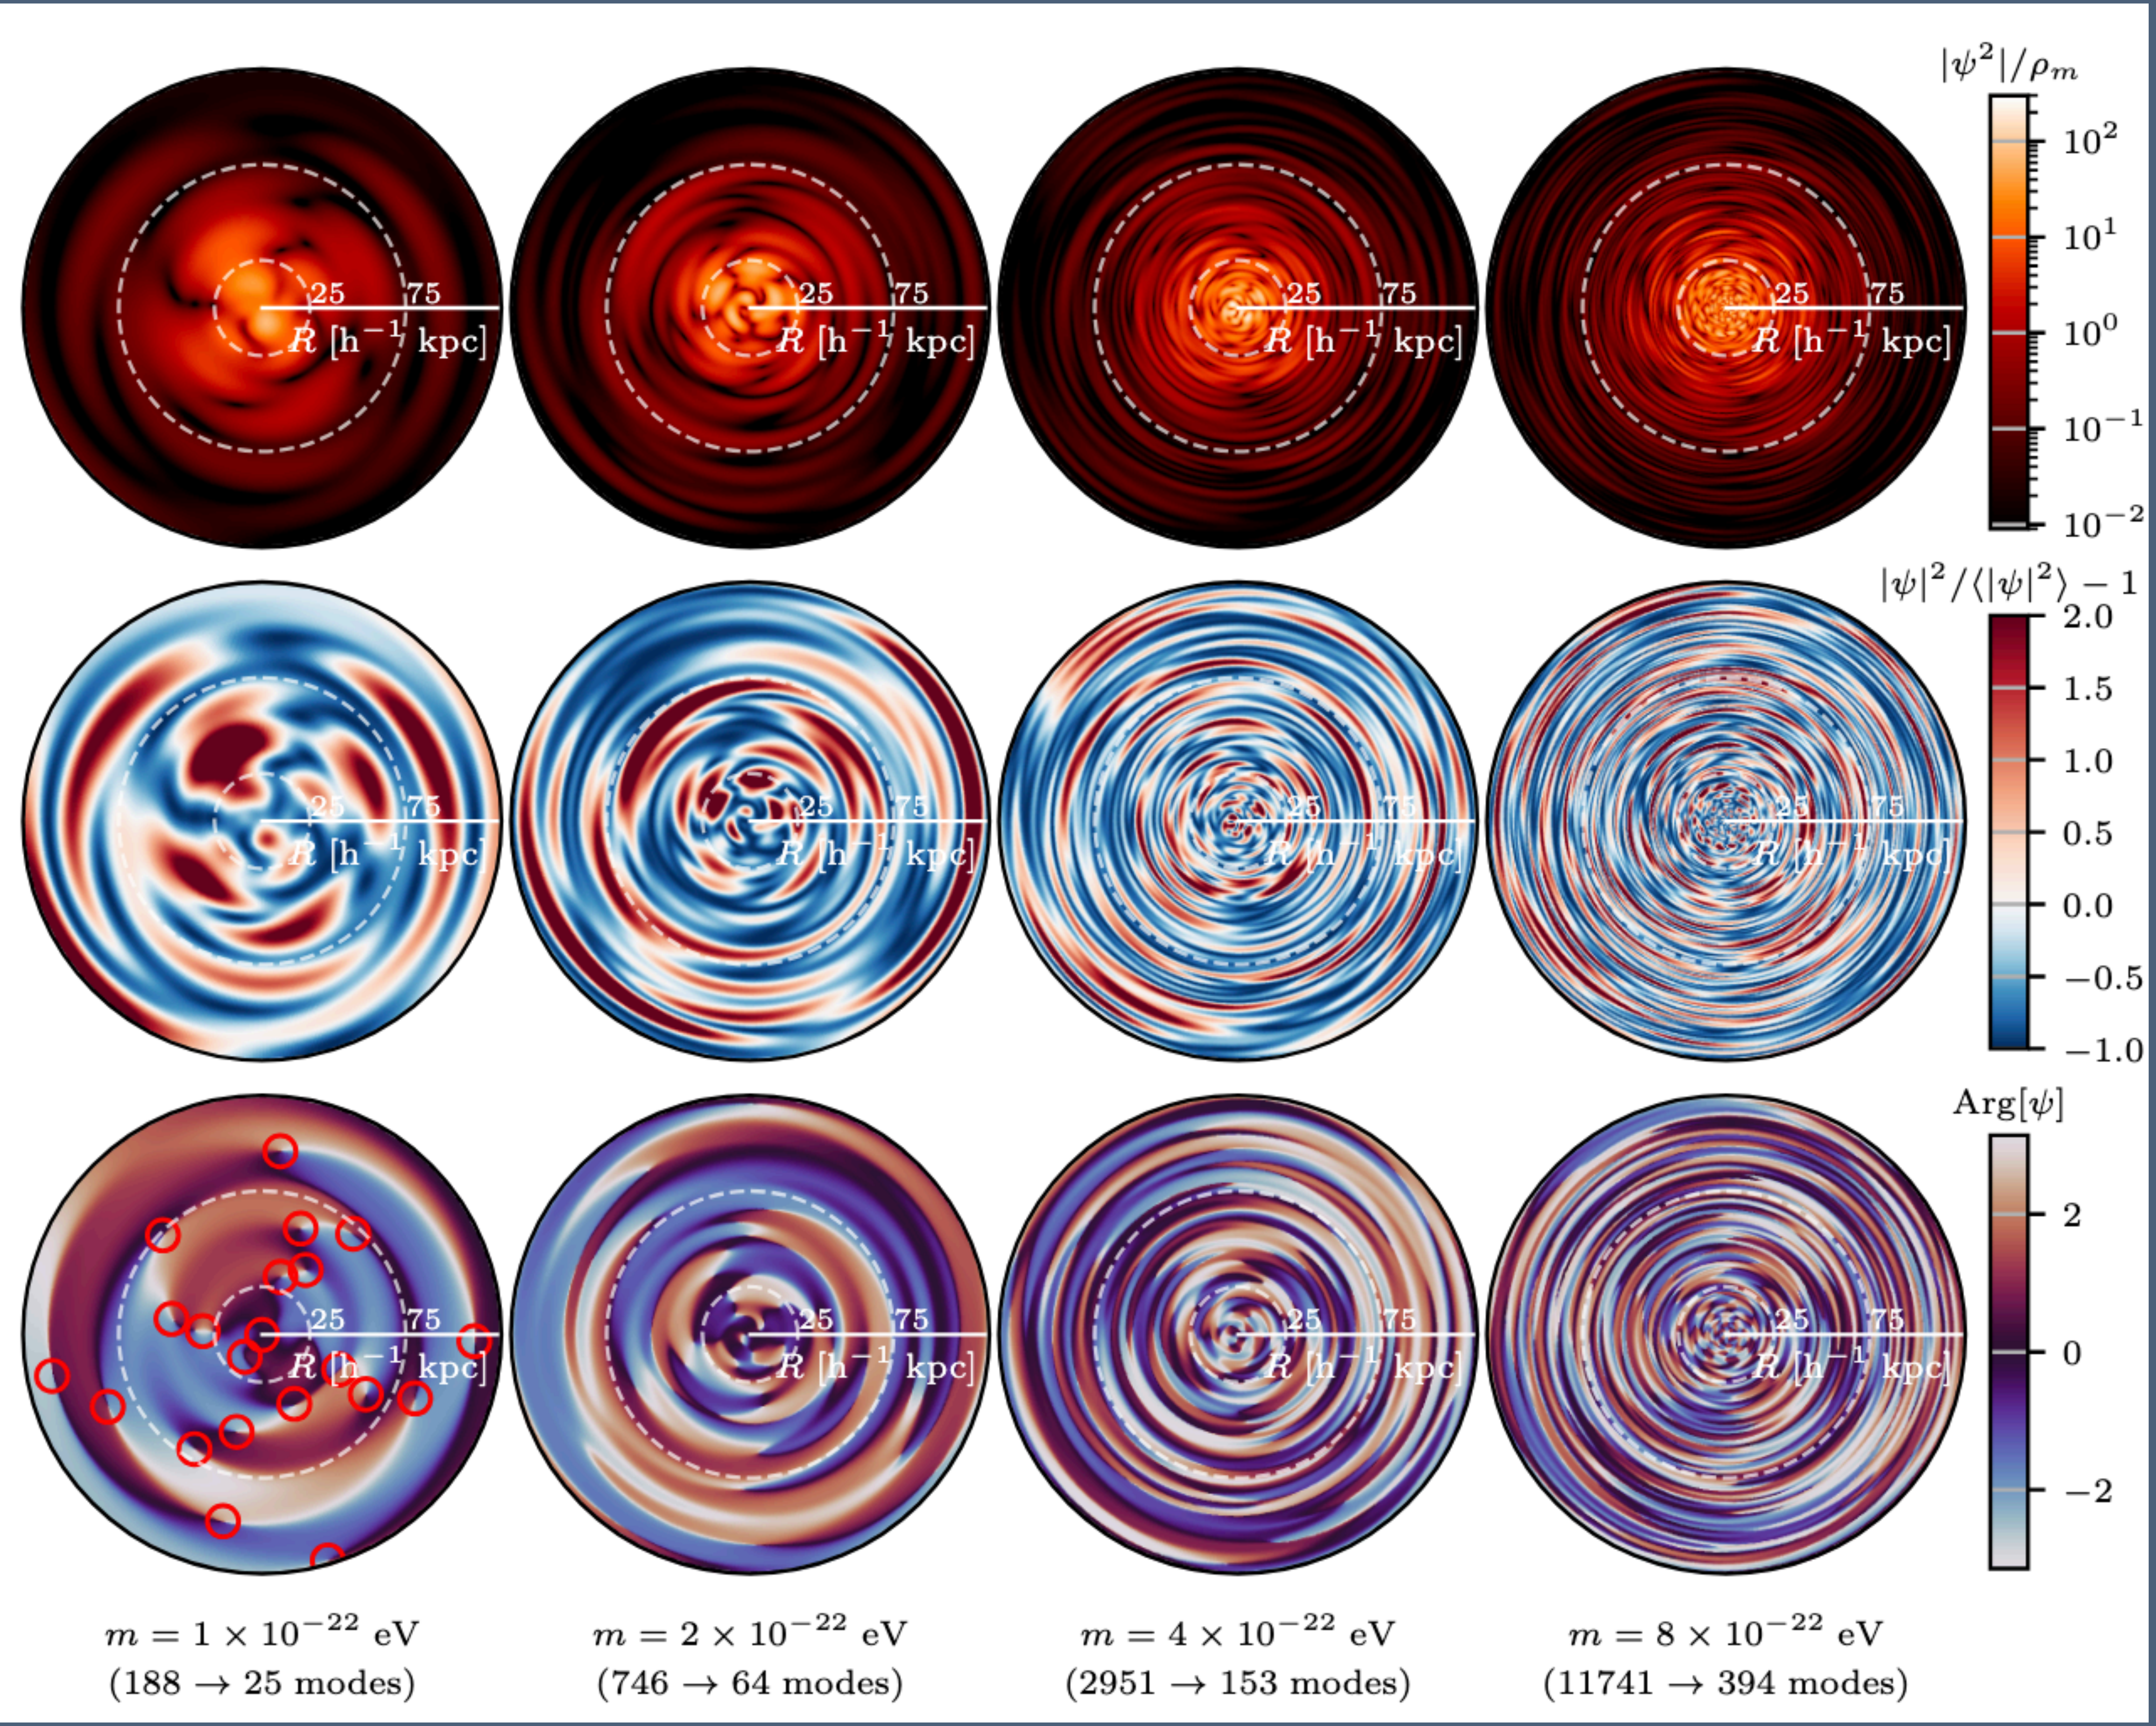

Total density

Interference term (overdensity)

Phase of the wave-function

Statistical distribution

- The final step is to go from our effectively two-dimensional, single cylinder to a three dimensional filament population.
- For this we need to know the distribution of filaments with mass/size.
- We use **ellipsoidal collapse.** A filament is identified as *an object for which two out of the three principal axes are frozen*
- Allows us to construct the filament mass function which we can then sample from to generate a simulation box of realistic filaments

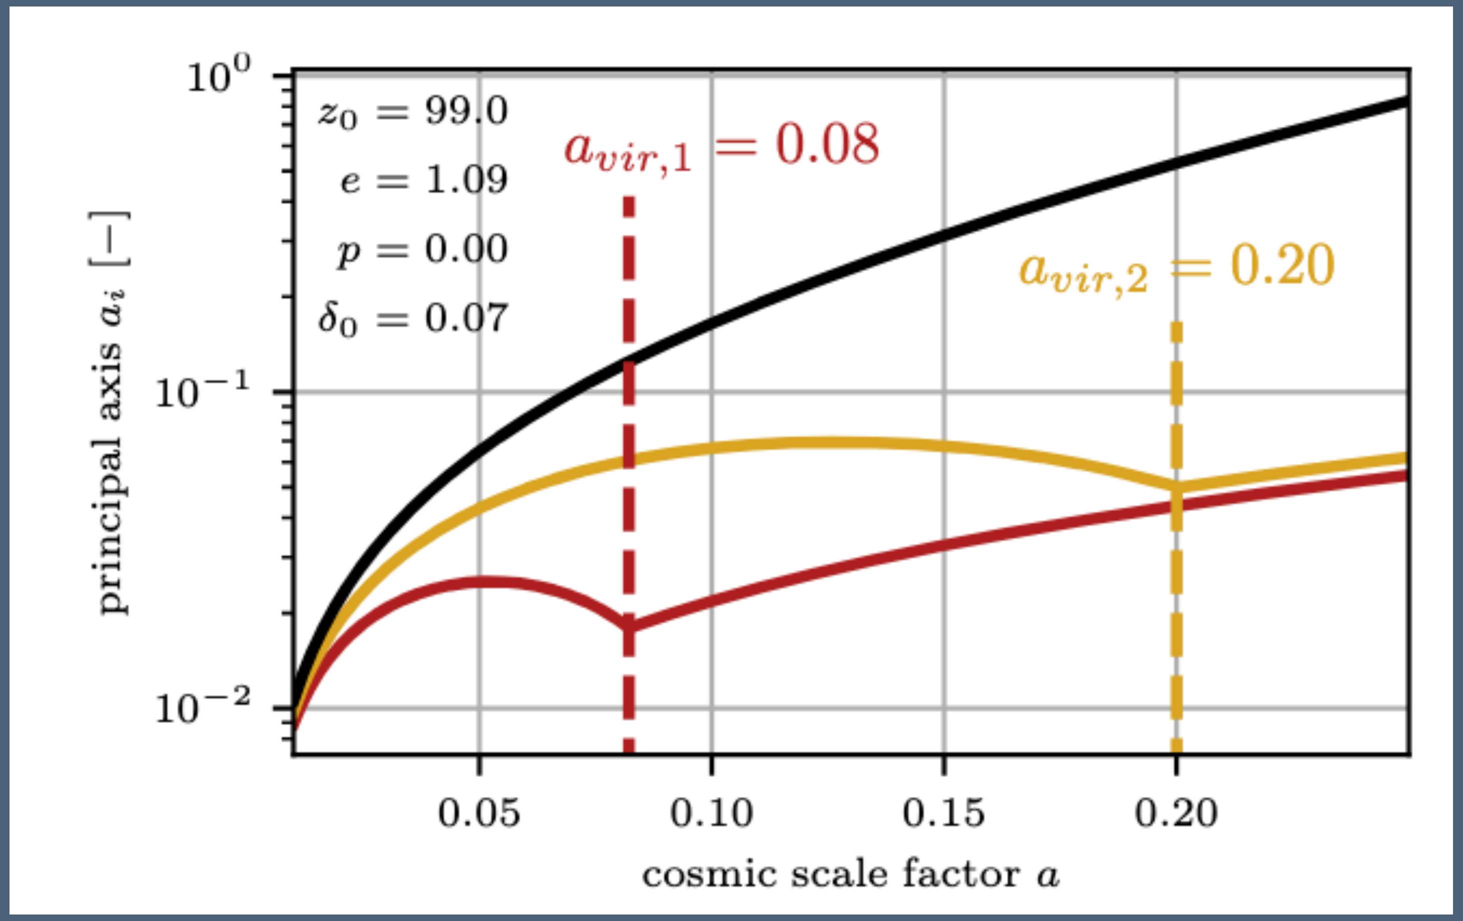

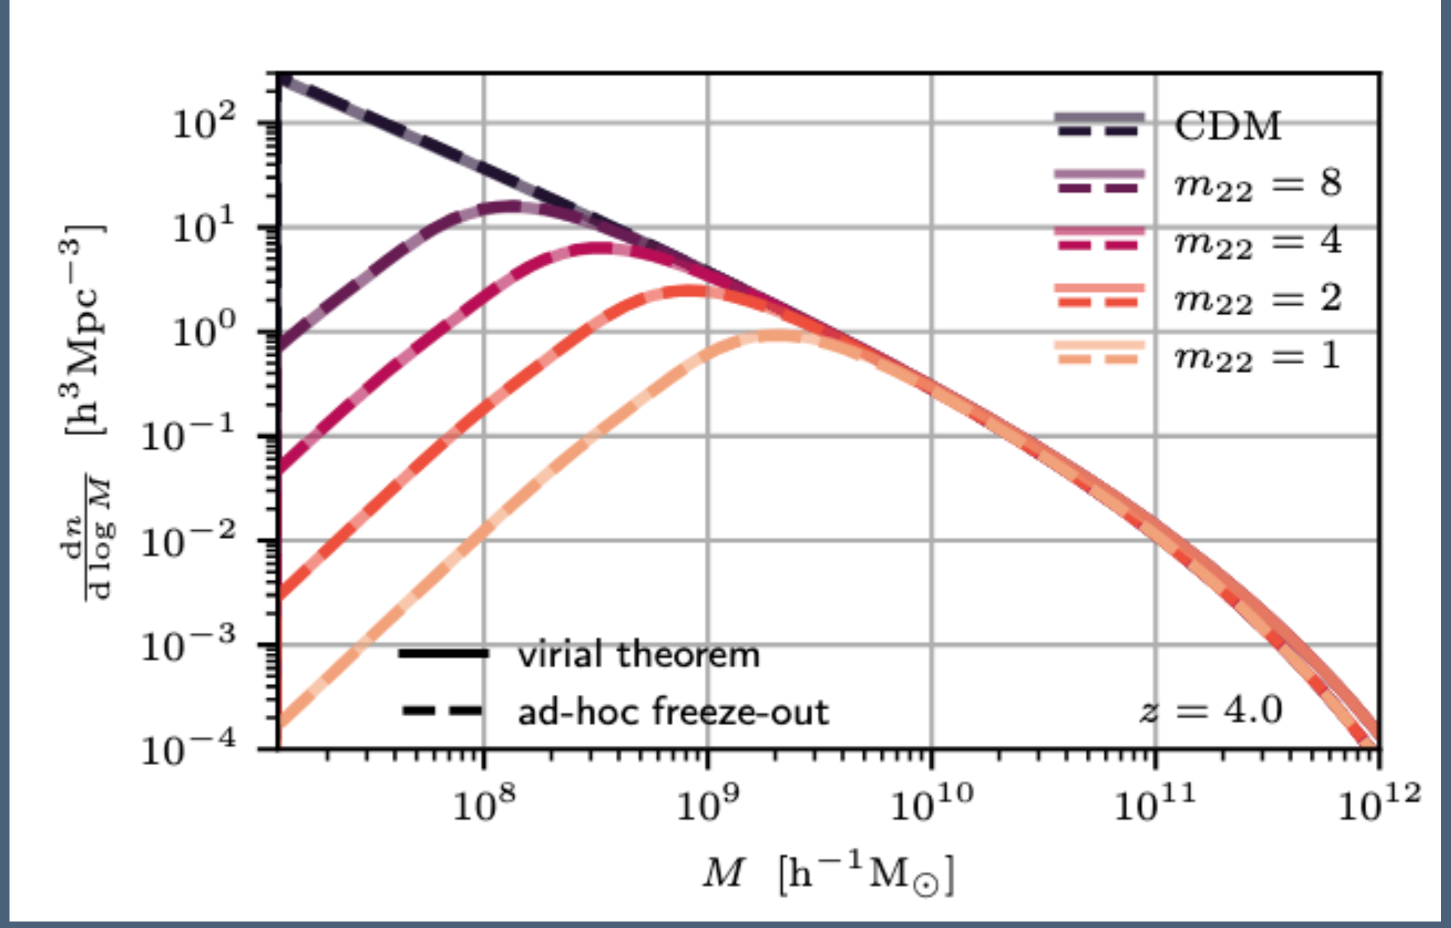

3D distribution of the filaments

- An example realisation of our filament population assuming cylinder locations and random orientations.
- Samples are accepted if no cylinders overlap (otherwise violates our assumption of all filaments being gravitationally isolated)
- **Currently shortcoming:** ignores any large scale filament-filament cross-correlation (which, like for the halo model, at leading order should be given by the linear power spectrum). More refined placement techniques (e.g. peak-patch description by Bond & Myers 1996) may be used to alleviate this shortcoming.

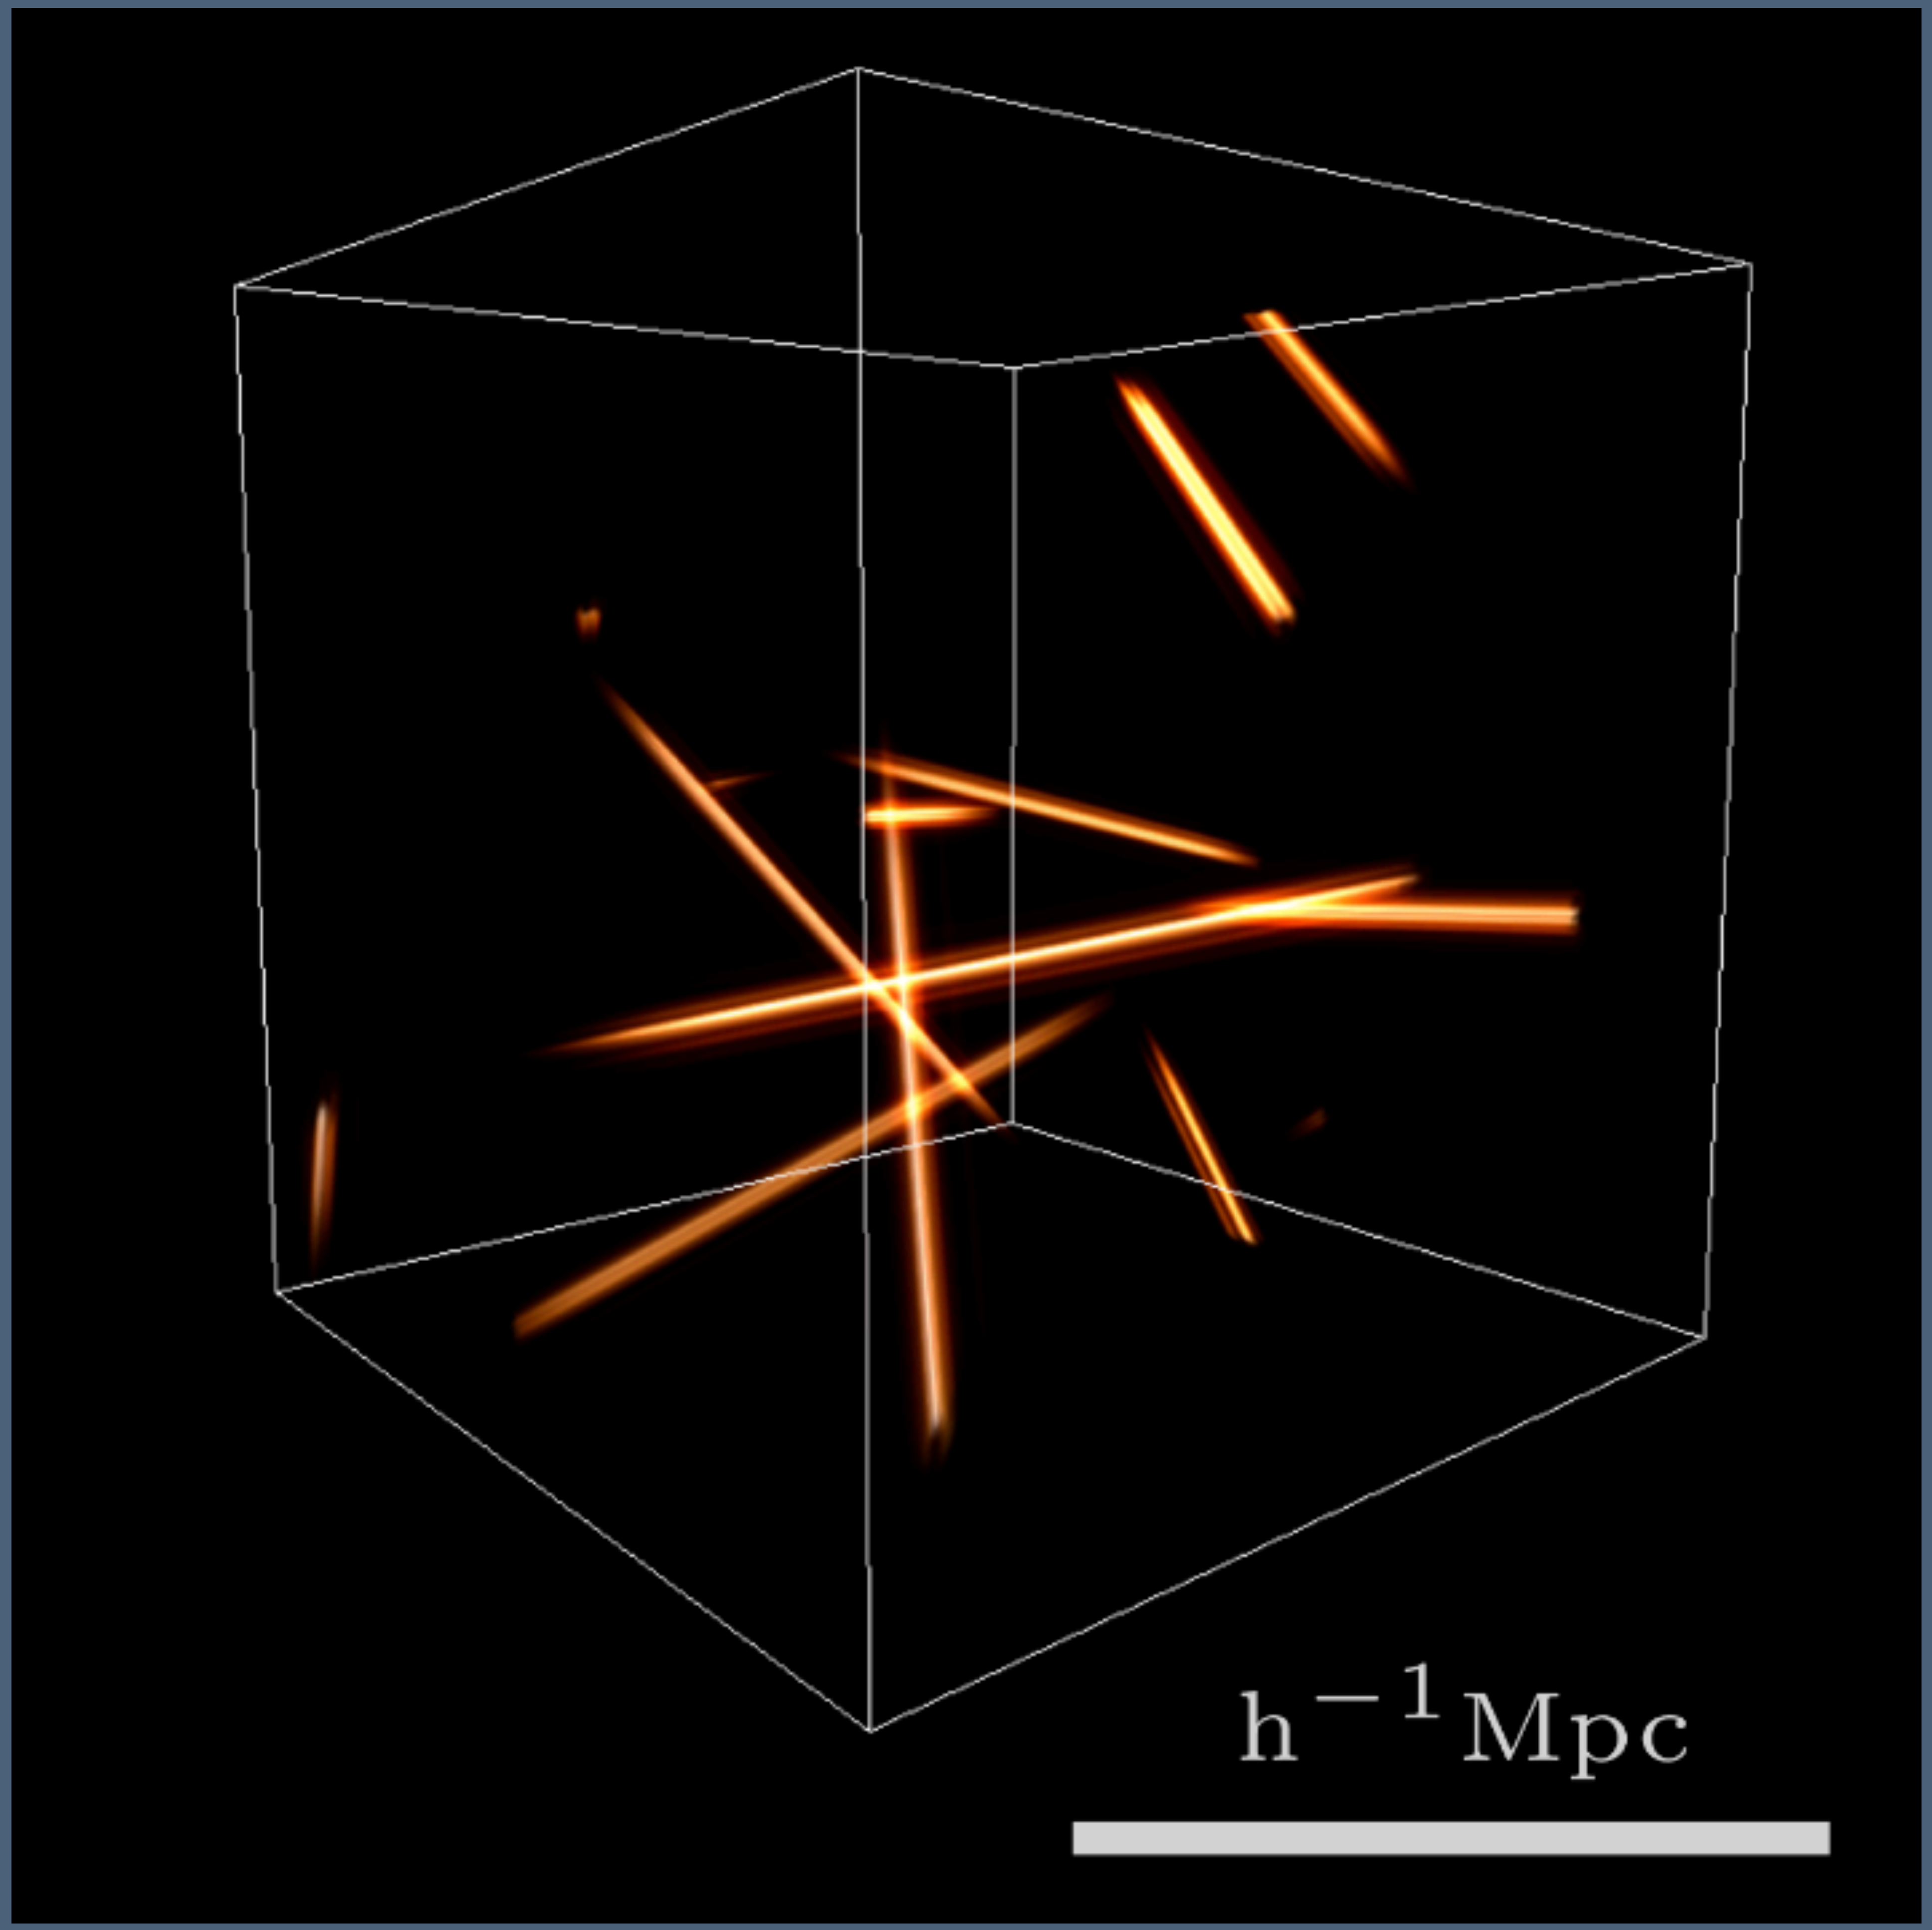

Matter power-spectrum

- Motivated by the halo-model we compute the 1-term filament power-spectrum

$$P^{1f}(k) = \frac{1}{\bar{\rho}_m^2} \int \mathrm{d}M\, M^2 n(m) P_{\mathrm{single}}(k \mid M), \quad \bar{\rho}_m = \int \mathrm{d}M M n(M)$$

- The FDM spectrum detaches from the CDM at some wavenumber k-detach where we start to probe scales interior to the filament.

- We have a cut-off scale k-cut which we interpret as a non-linear extension of the linear Jeans suppression scale set by the uncertainty principle.

- Full cosmological simulations of FDM have found an excess correlation in the FDM matter power spectrum compared to its CDM counterpart for highly non-linear k > O(100) h/Mpc (Veltmaat & Niemeyer 2016; Mocz et al. 2020; May & Springel 2021, 2022; Laguë et al. 2024). Conjectured to originate from interference fringes. Consistent with what we see in our model!

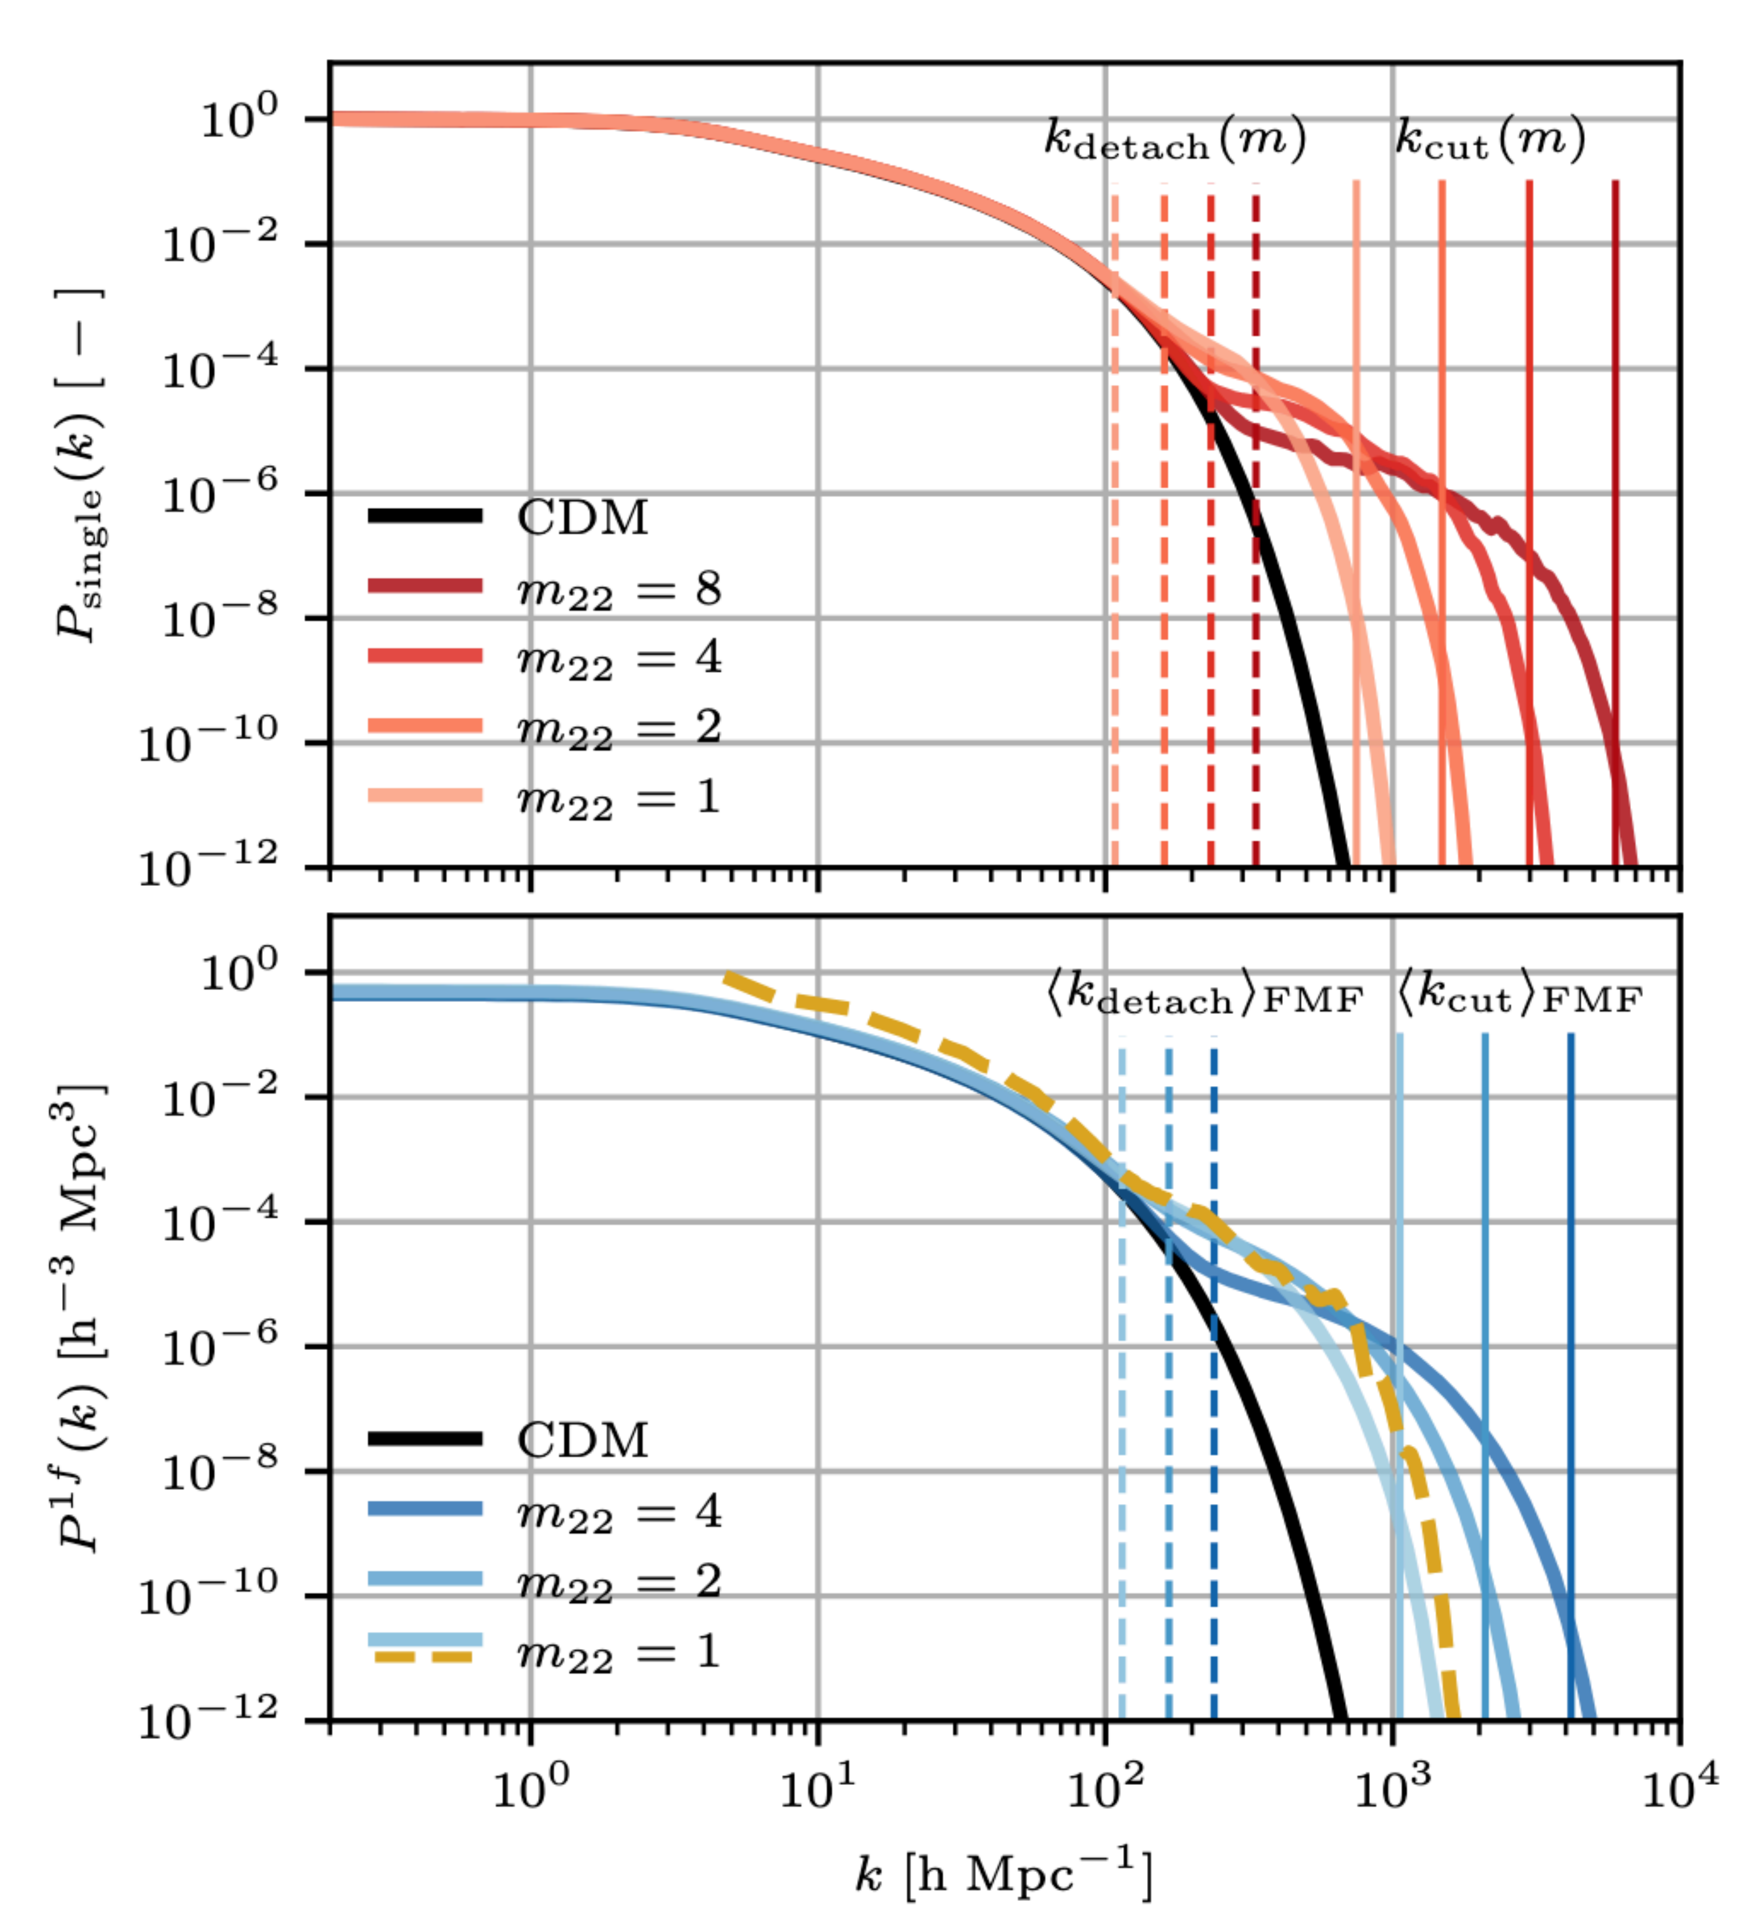

Summary

- We have created an idealised model for cosmic filaments in model with an ultra light axion / fuzzy dark matter.
- Even though its simplicity, the model has most of the features we expect to see like interference and quantum vortices and looks consistent with simulations.
- Just a first model. Many ways it can be extended or improved.
- The model can be used to study filaments in FDM, but also as a post-processing tool for CDM simulations of steady-state objects (i.e. as a way of injecting FDM into CDM).

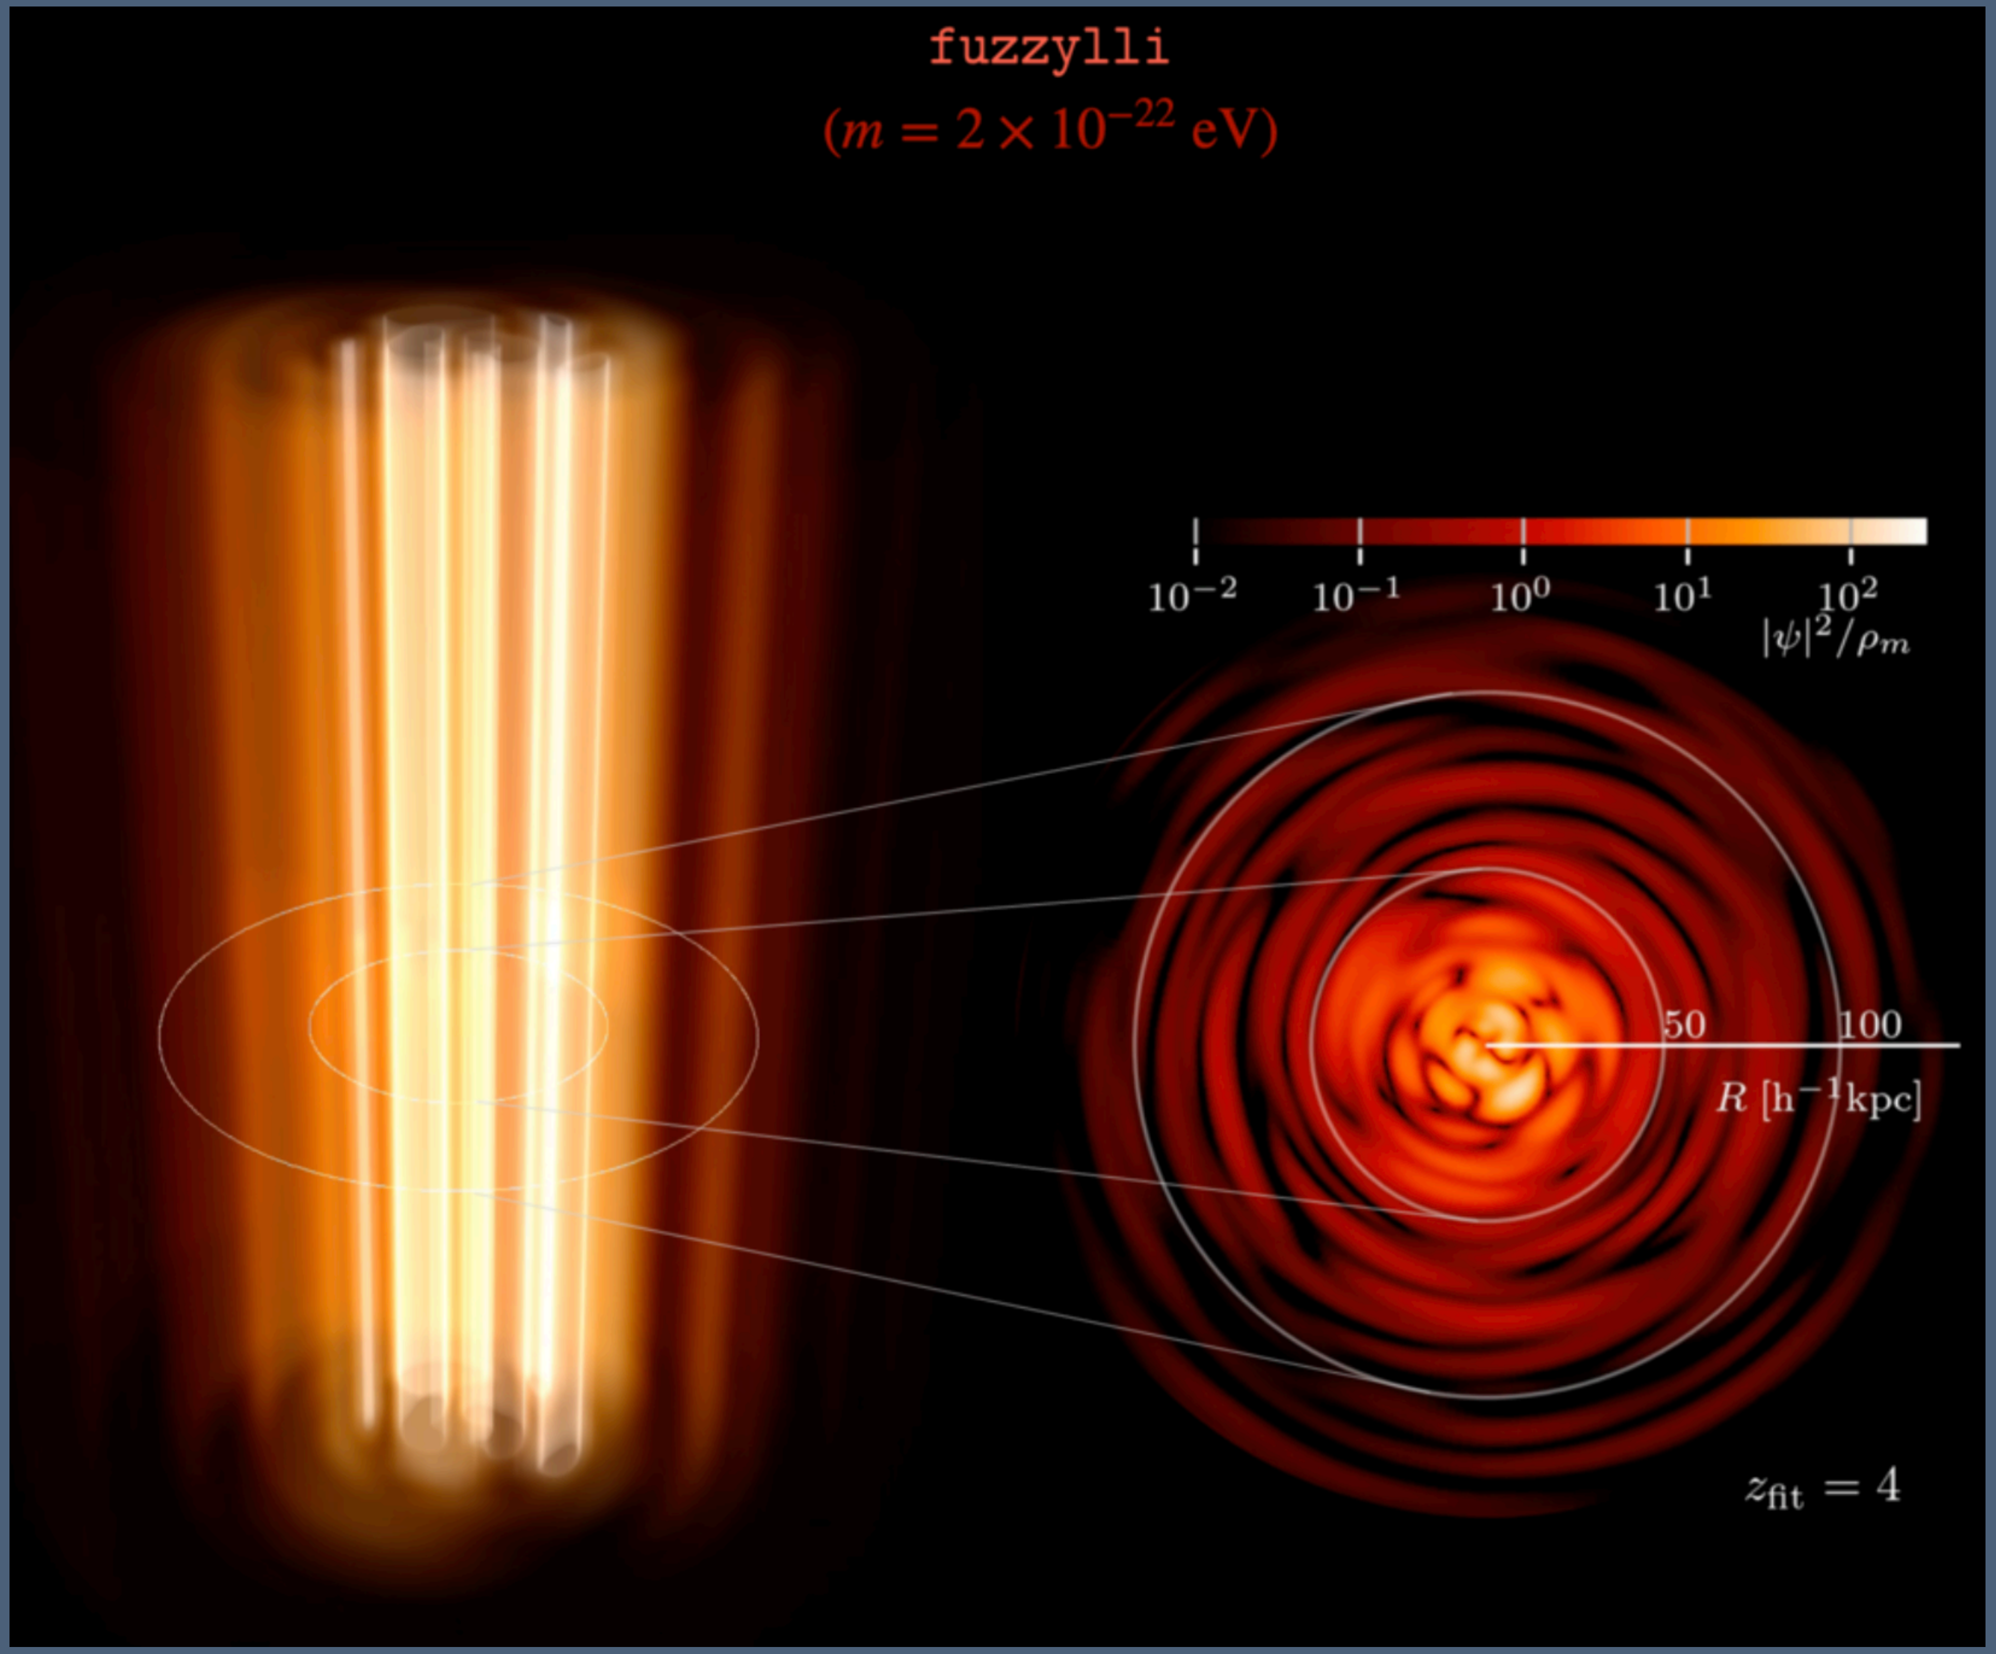

GitHub: The code used to generate the results in this work can be found at timzimm/fuzzylli and timzimm/fdm_filaments.

Paper II : Emulating the Non-Linear Matter Power-Spectrum in Mixed Axion Dark Matter Models

Dennis Fremstad and Hans A. Winther

- The axion mass and axion fraction (in mixed CDM + axion model) can be constrained by combining many cosmological datasets (CMB, LSS, SN, ...)
- For LSS, current and future weak-lensing surveys will probe the non-linear regime and help narrow down the allowed regime.
- For the theoretical predictions we need to know the non-linear matter power-spectrum down to k ~ O(1-5) h/Mpc.
- For LCDM many high quality emulators for this quantity have been created (also including effects of baryons).
- For axions the main tools are axionCAMB (linear - Hlozek et al. 2015) and axionHMcode (a semi-analytical method for non-linear clustering - Vogt et al. 2022, Dome et al. 2024).
- Our goal was to make a simulation based emulator for axion-like models where the axion has a mass m_axion and contributes a fraction f_axion of the DM budget (the rest is CDM)

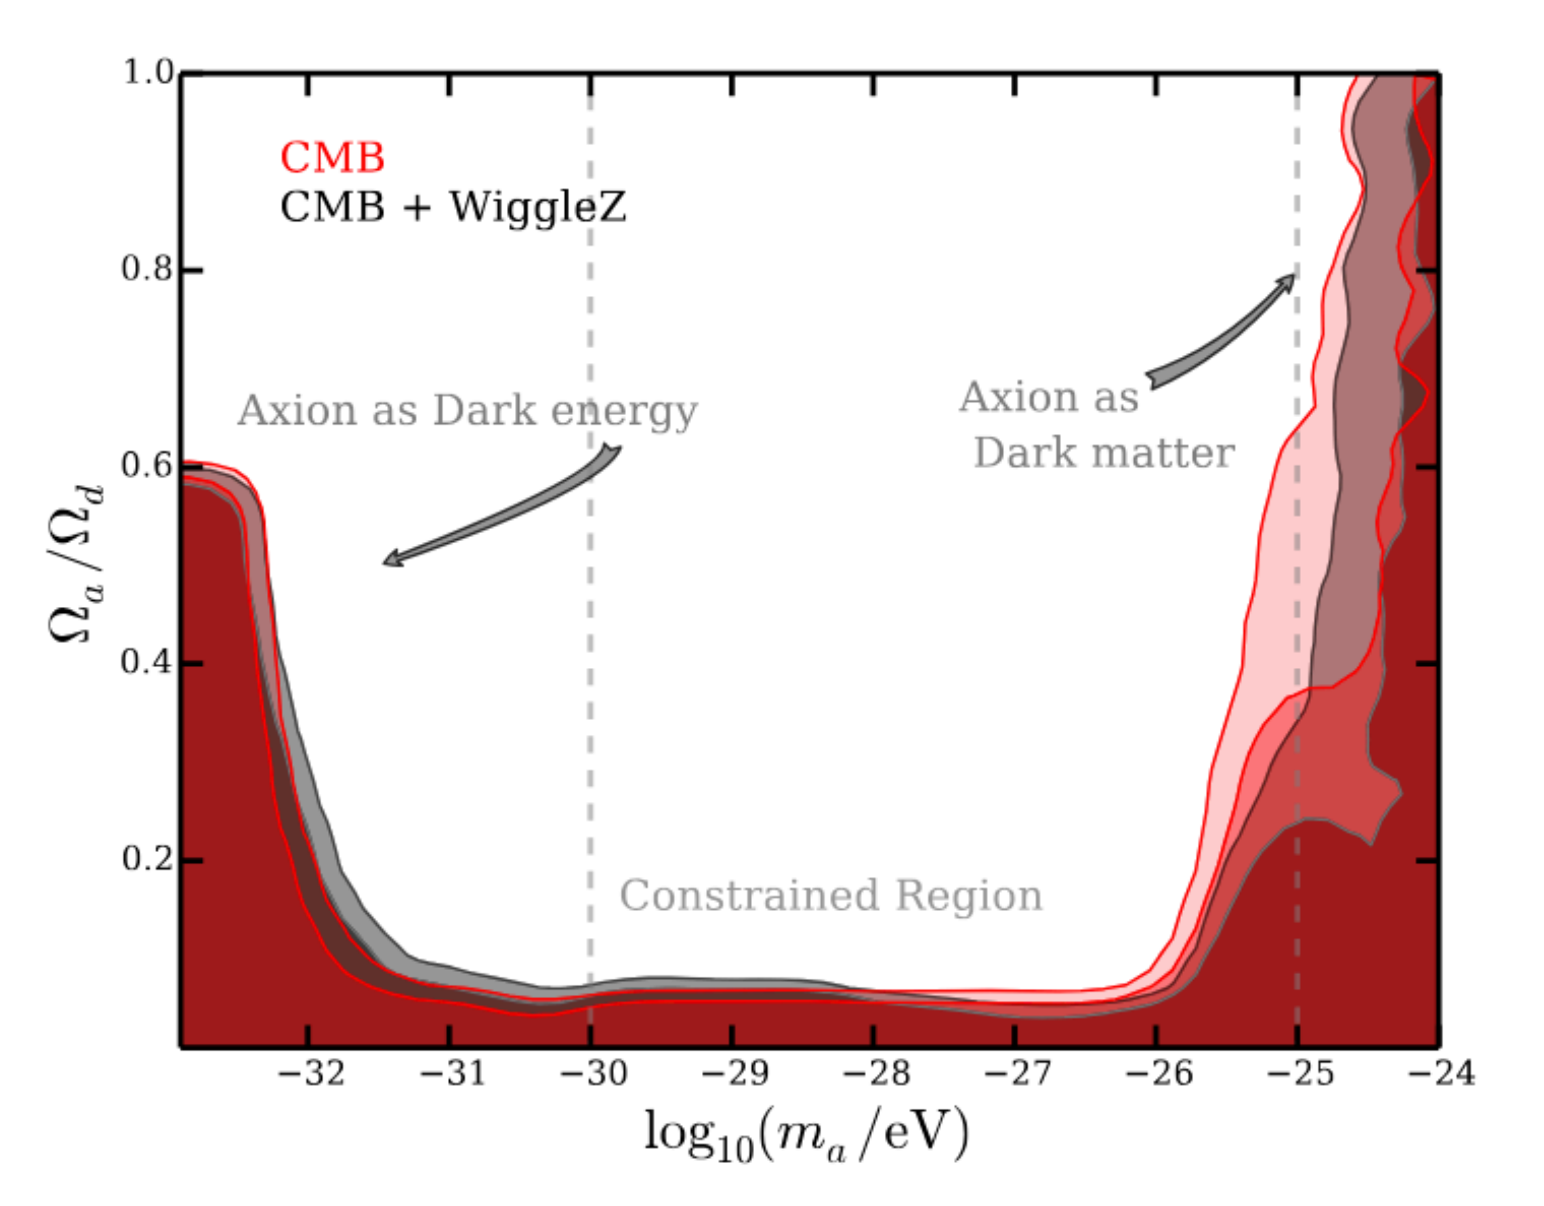

Hlozek et al. 2015

Effect of axions on P(k)

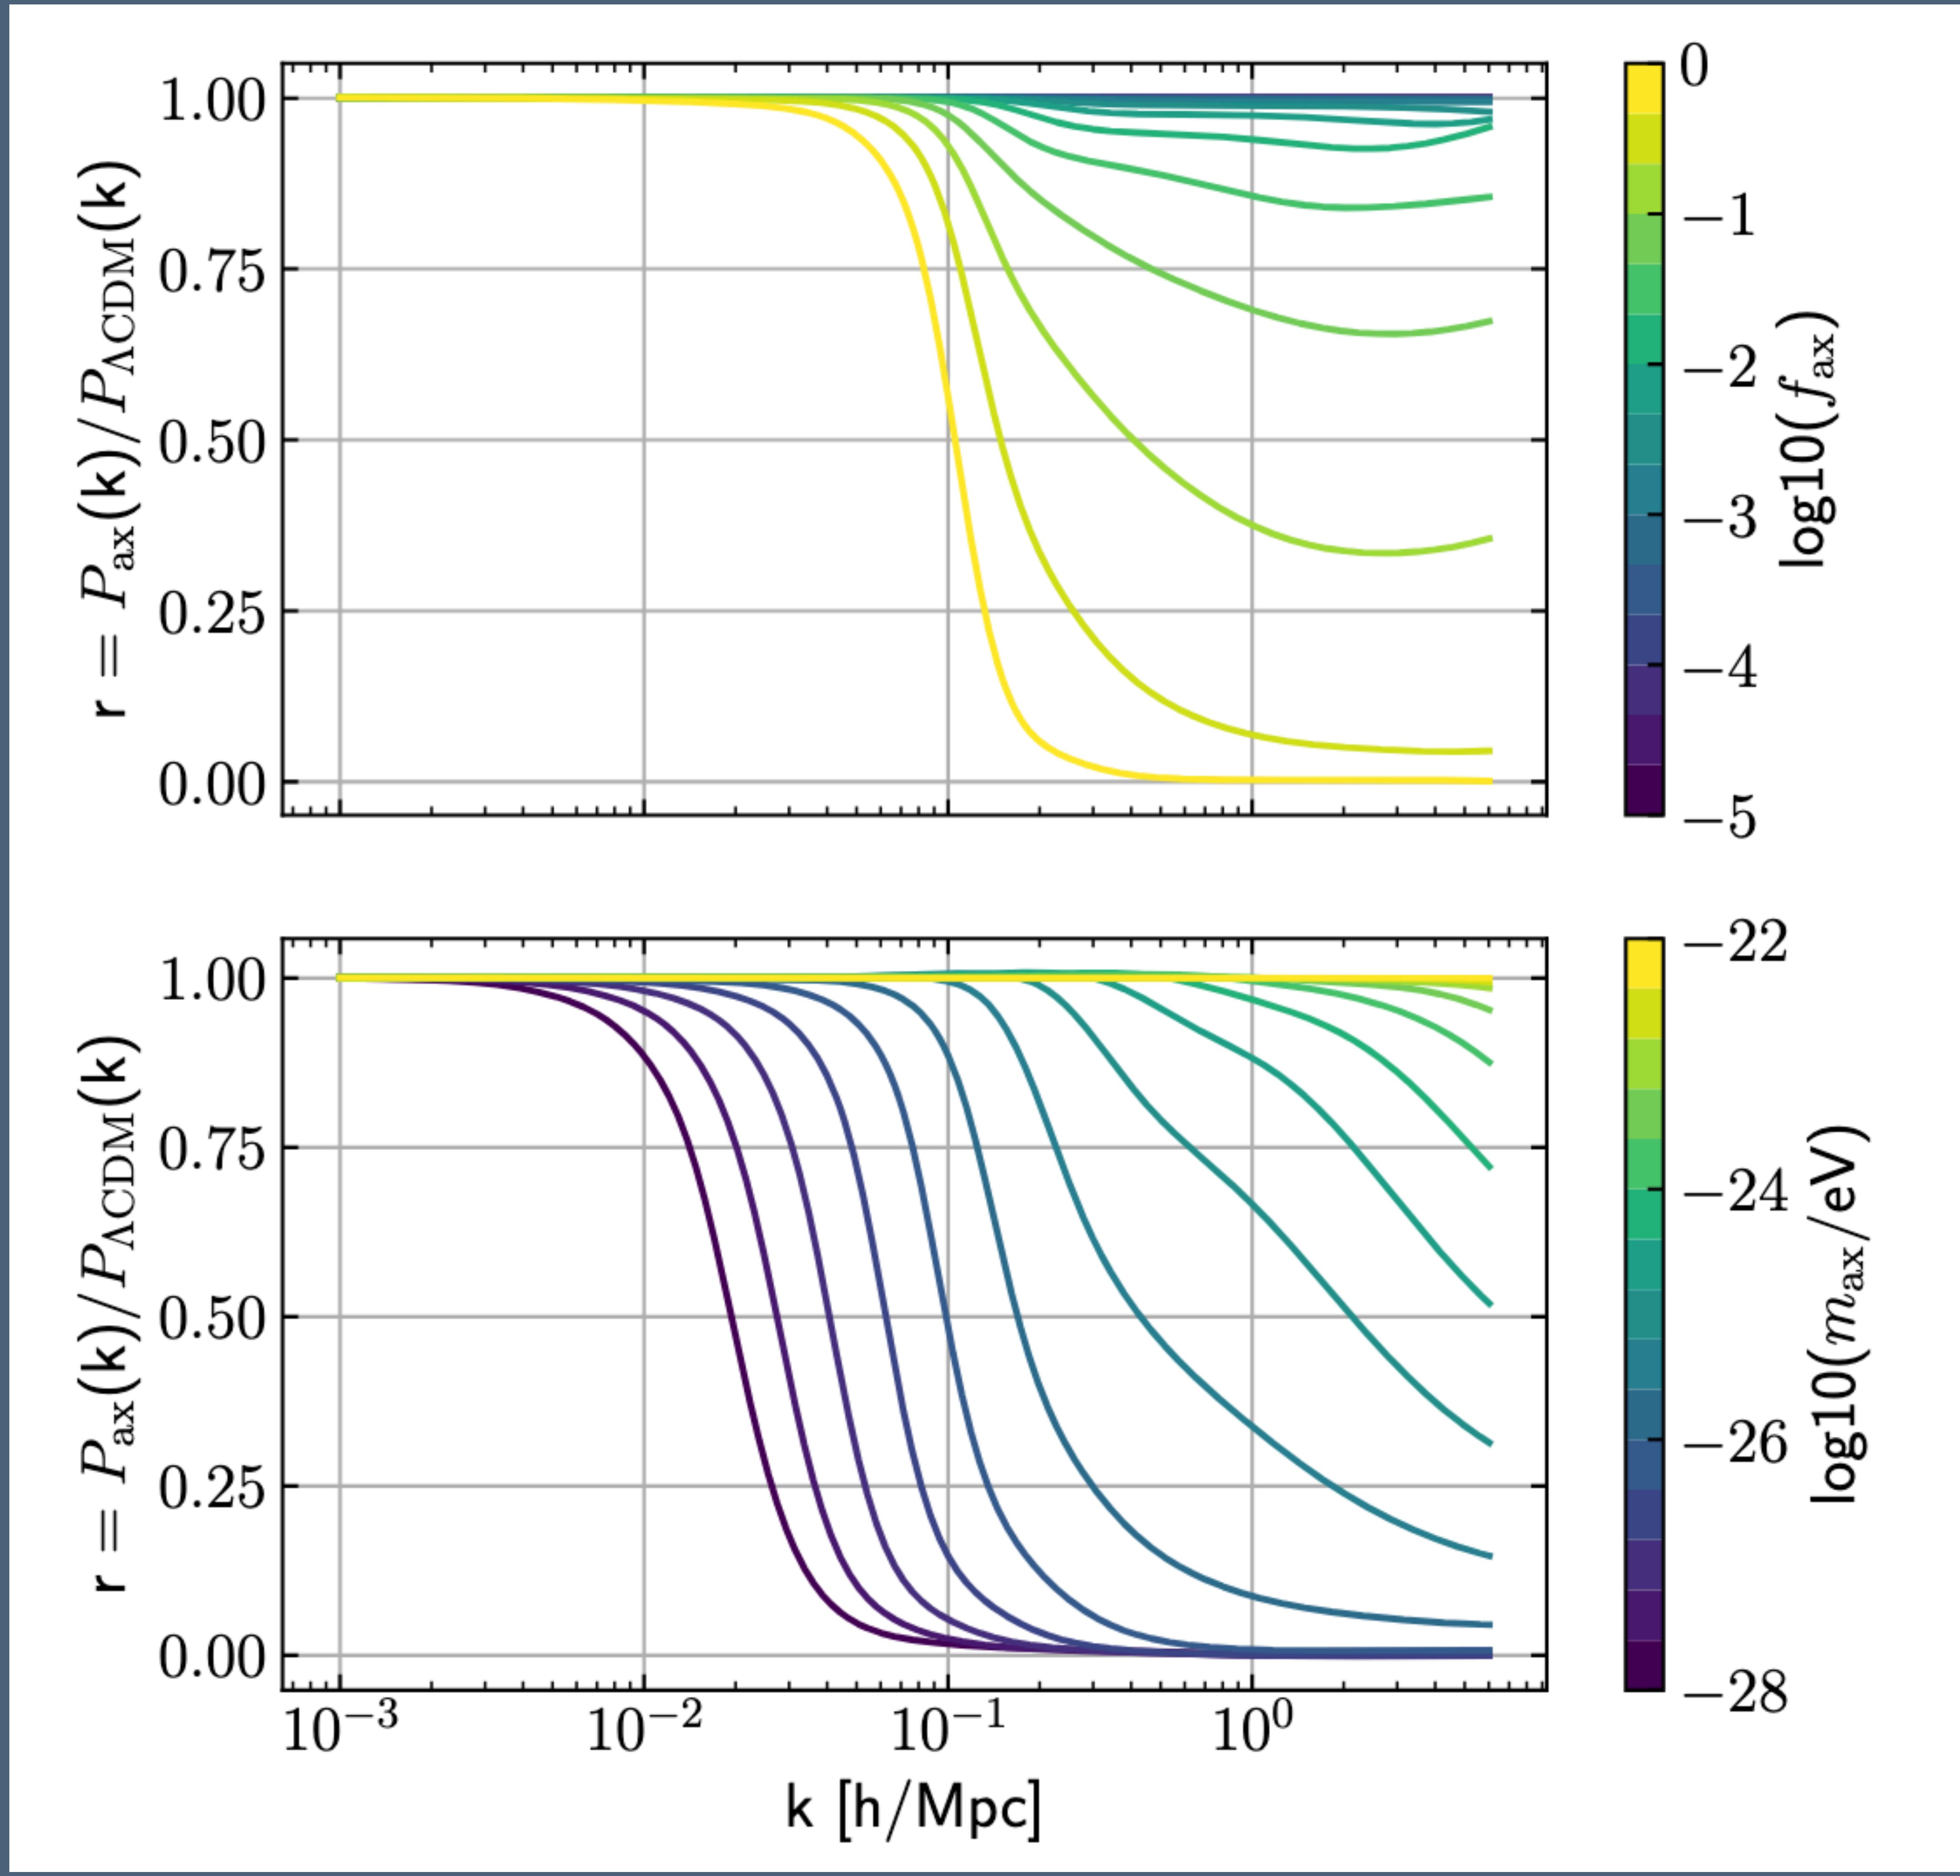

Approach

- **Problem:** simulations are too expensive. We do the simplest thing and only keep the axion physics in the IC. Big approximation, but allows us to run very fast simulations (PM).

- Results from some simulations suggest that this is an “ok” approximation for global clustering statistics atleast at higher redshifts (but still remains to be checked in more detail for a wider range of masses / axion-fractions)

- We choose to emulate the ratio **P(k, z | params) / PLCDM(k, z | params)**. Factors out some of the non-linearity. Allows us to piggy-back on the work done for high quality LCDM emulators.

- **Steps:**
  - Figure out what parameters we are sensitive to
  - Generate samples of the parameters
  - Run simulations for each set of parameters (plus corresponding LCDM)
  - Estimate the power-spectrum boost P(k, z | params) / PLCDM(k, z | params)
  - Run the machine learning and ensure the result is accurate

- All these steps are included in the package SESAME (Mauland-Hus and HAW 2024)

May & Springel 2022

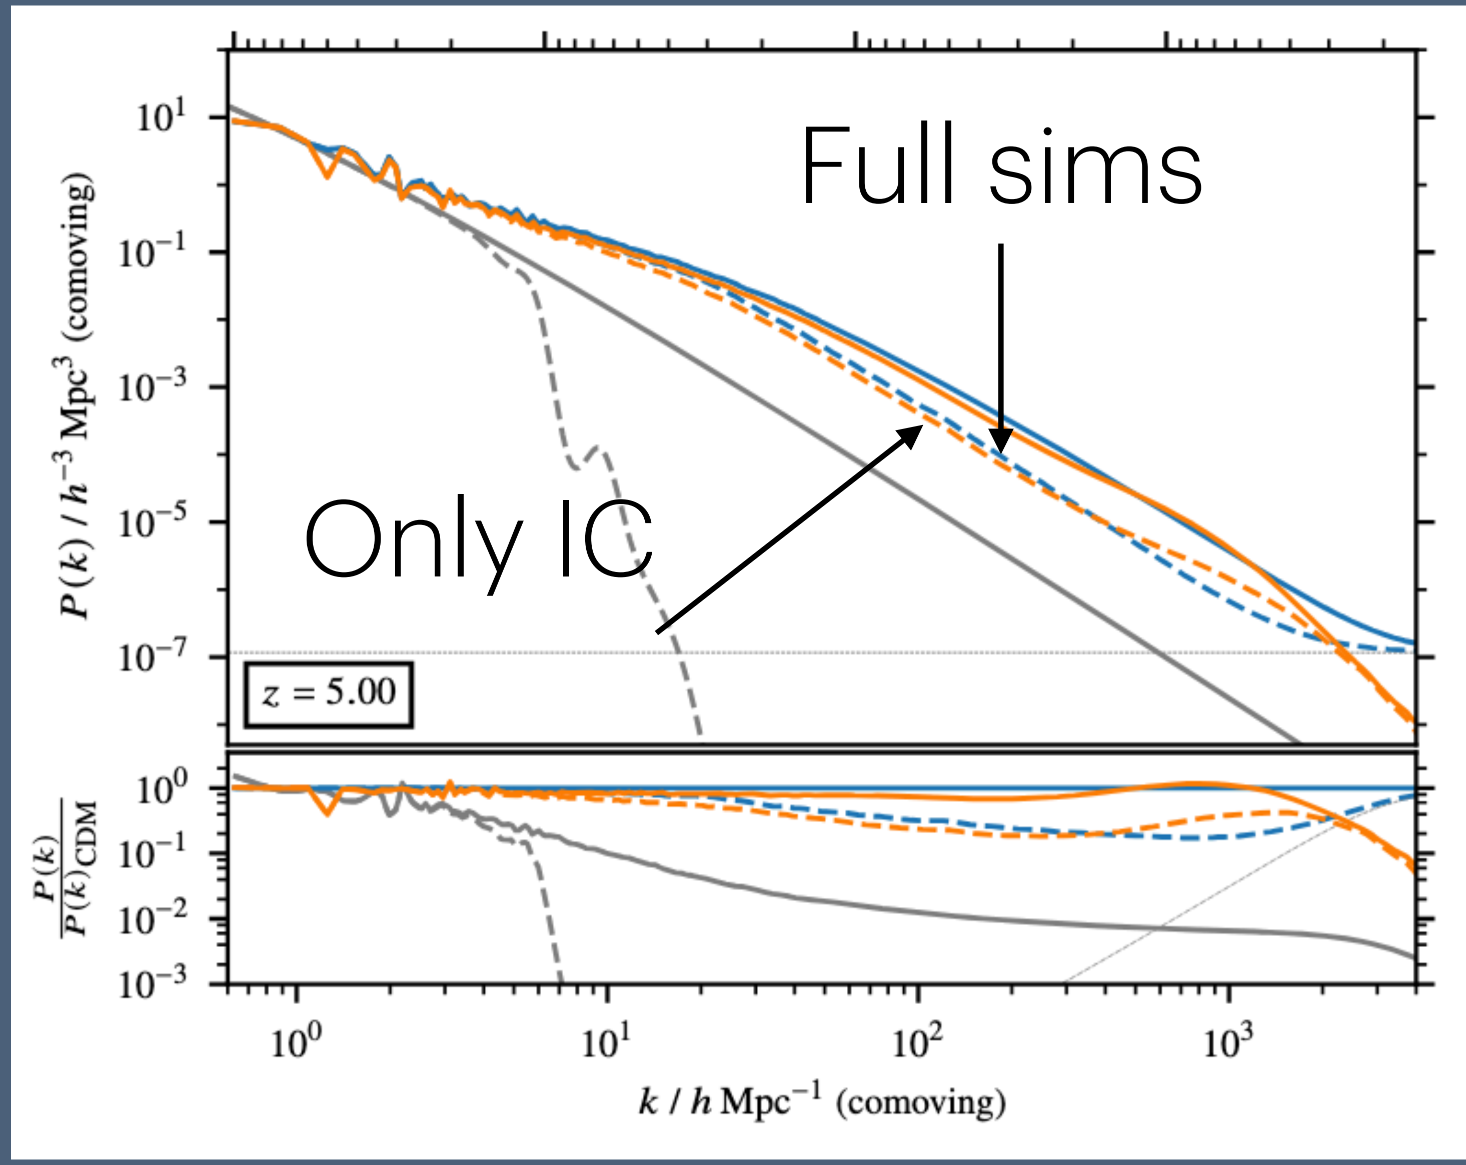

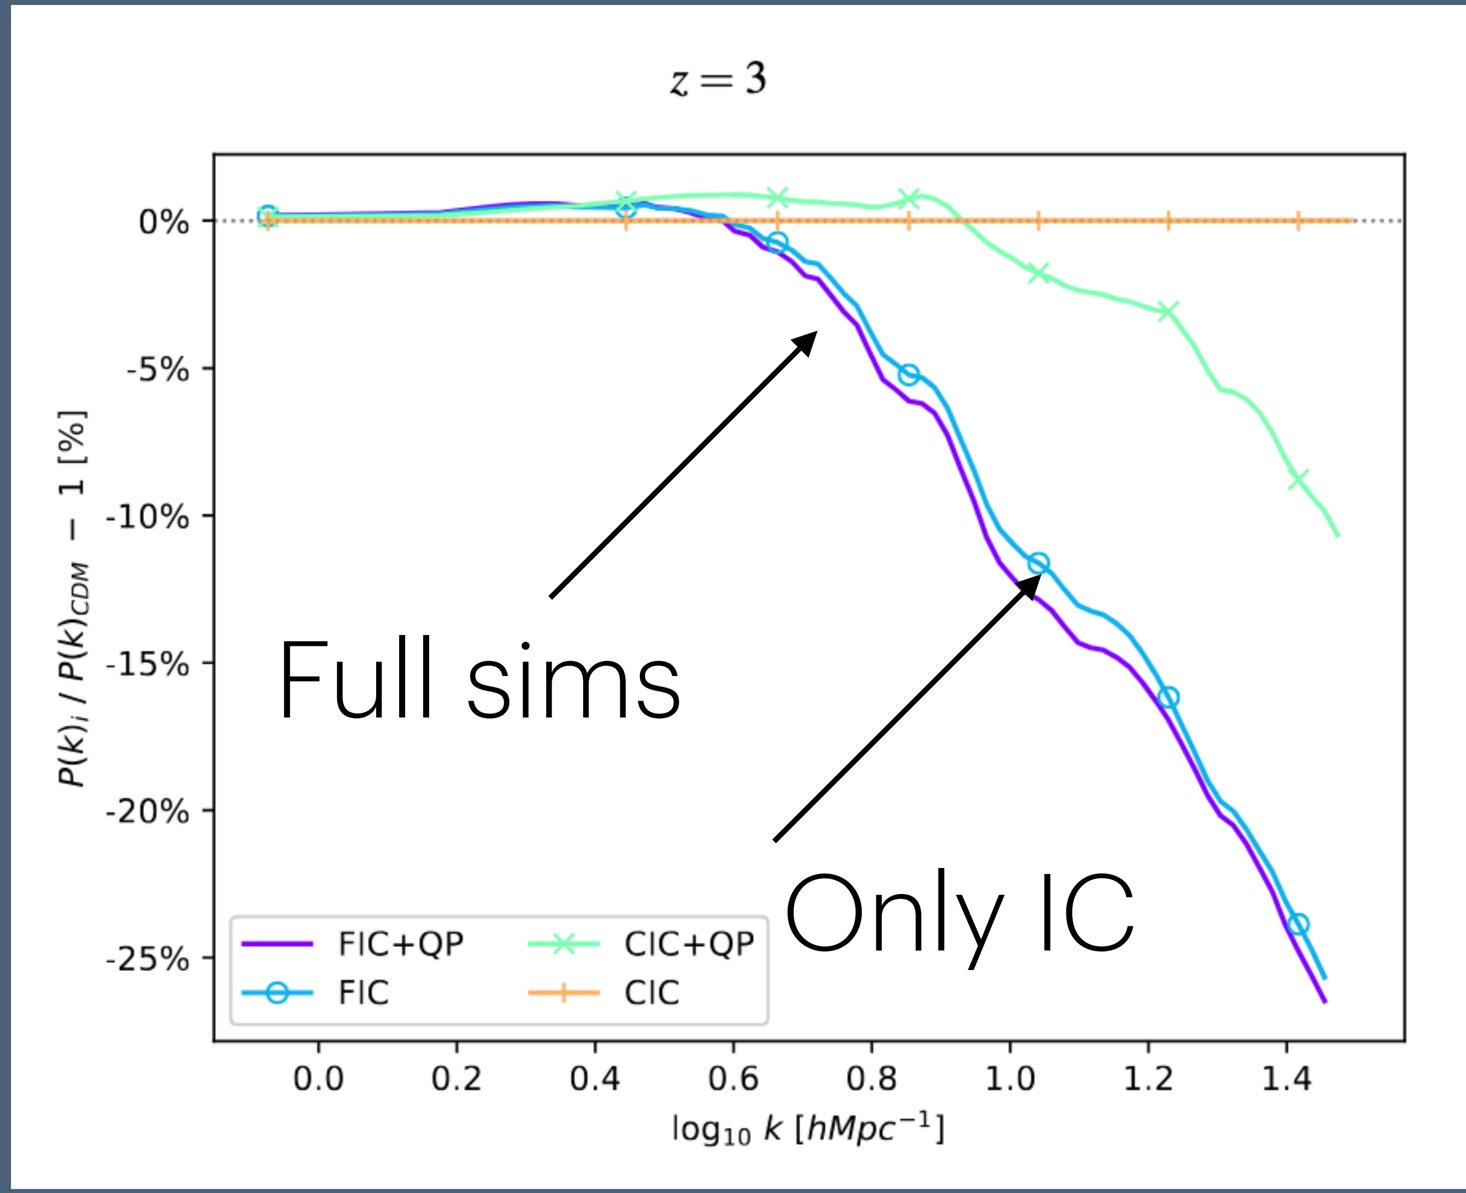

Nori & Balid 2019

Cosmological parameter dependence

- The power-spectrum ratio P/PLCDM is only sensitive to the axion parameters and the main parameters determining the clustering in LCDM (ΩCDM and As)
- We therefore choose to only sample the 4 parameters m_axion, f_axion, ΩCDM and As and sample them inside the given range and run PM simulations for each choice of parameters

$$
\begin{aligned}
\log_{10} m_{\rm ax}/{\rm eV} &\in [-28, -22],\\
f_{\rm ax} &\in [0.001, 1],\\
\log_{10} A_s &\in [-9, -8.52],\\
\Omega_{\rm DM} &\in [0.2, 0.4],\\
z &\in [0, 6.75],
\end{aligned}
$$

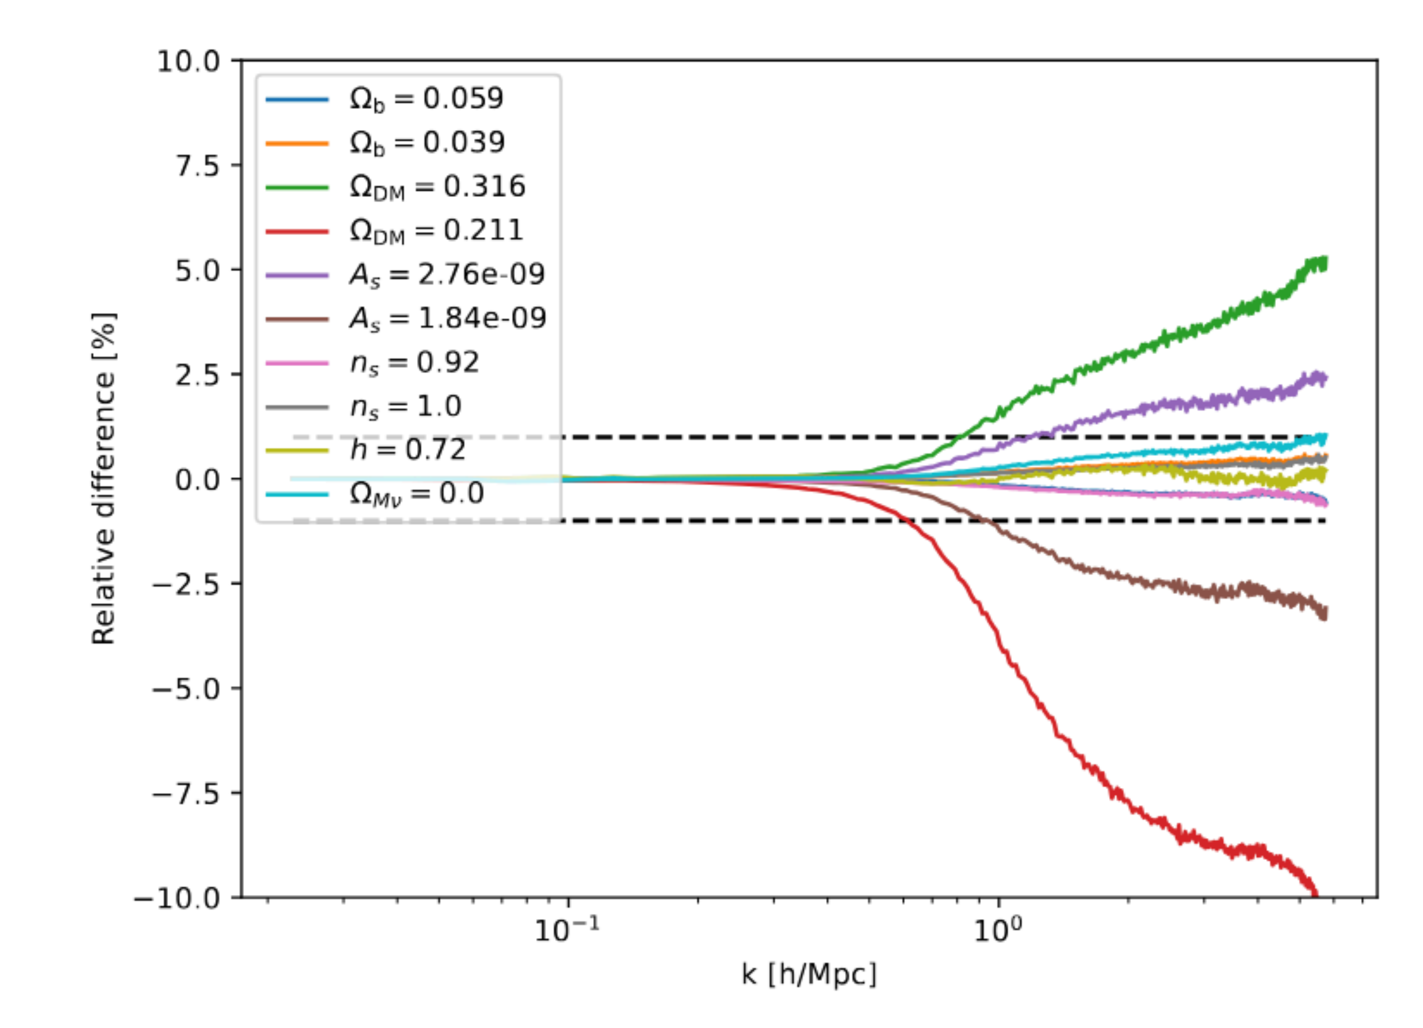

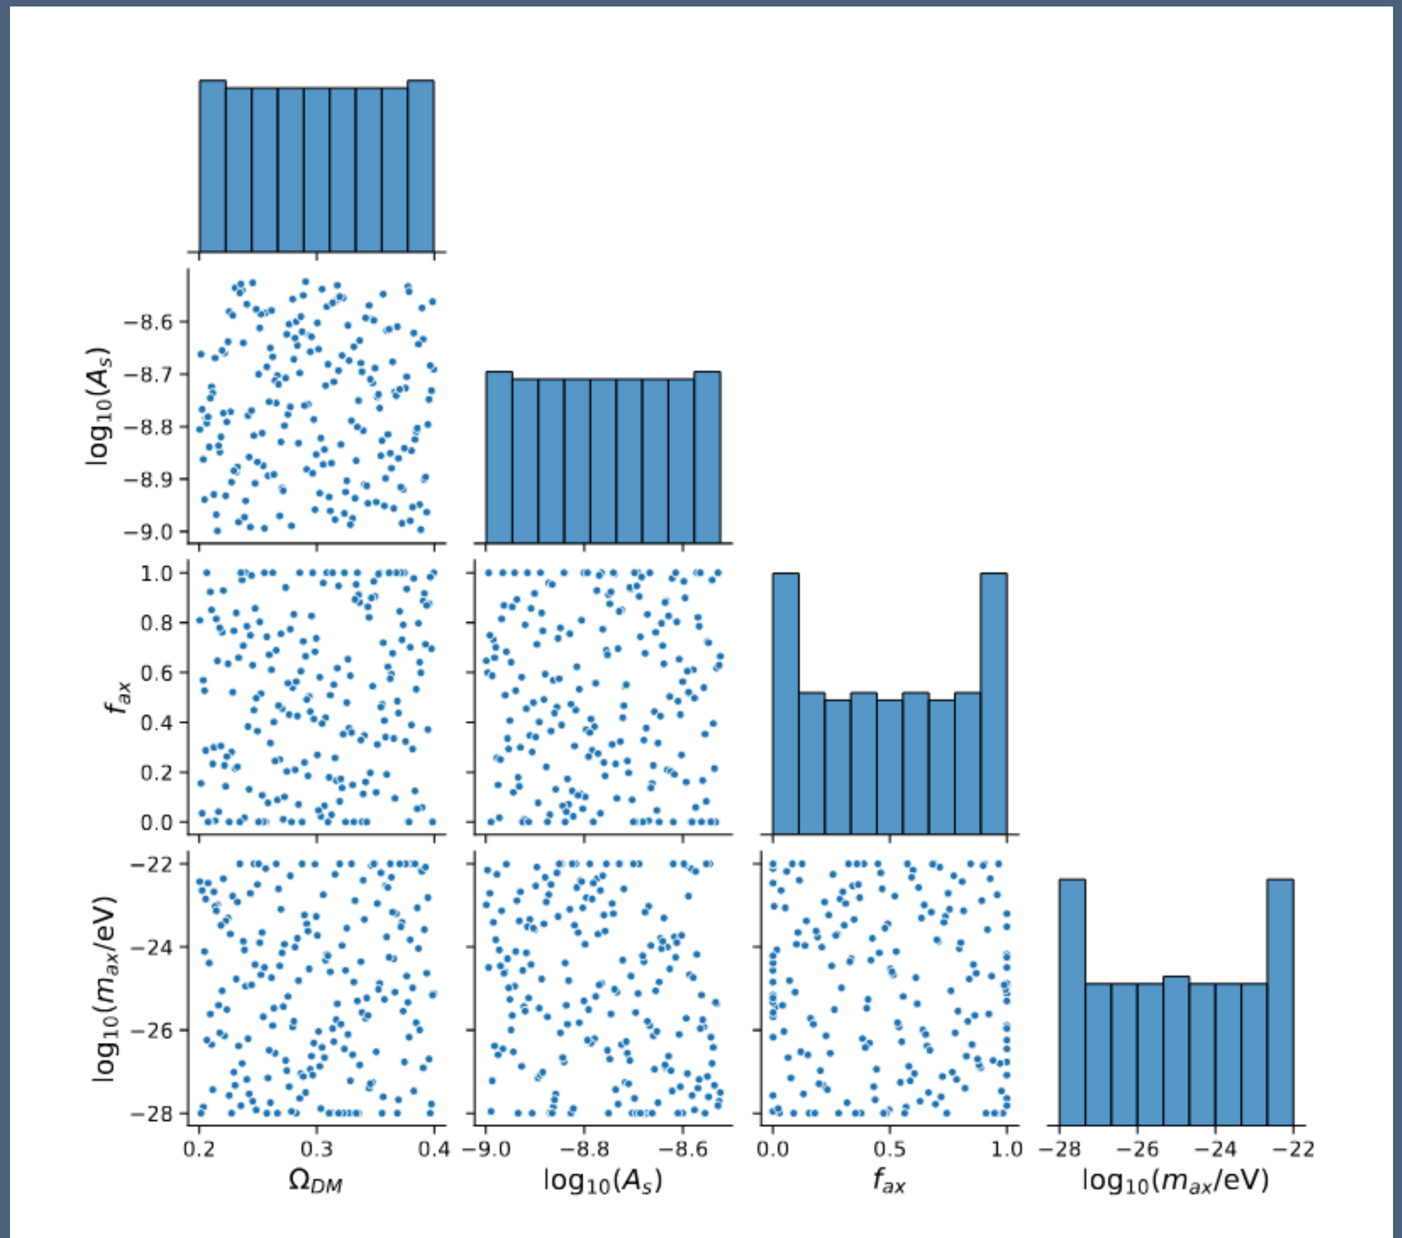

Machine learning

- We employ a simple Feed Forward Neural Network architecture with 2 hidden layers.
- The input layer accepts 6 parameters (ΩCDM, log10 As, f_ax, log10 m_ax, z and k) and outputs a single value: the prediction for the power-spectrum ratio for the given input values.
- The training may take between ~15 CPU minutes and 3 CPU hours to train. Very fast and achieves high accuracy

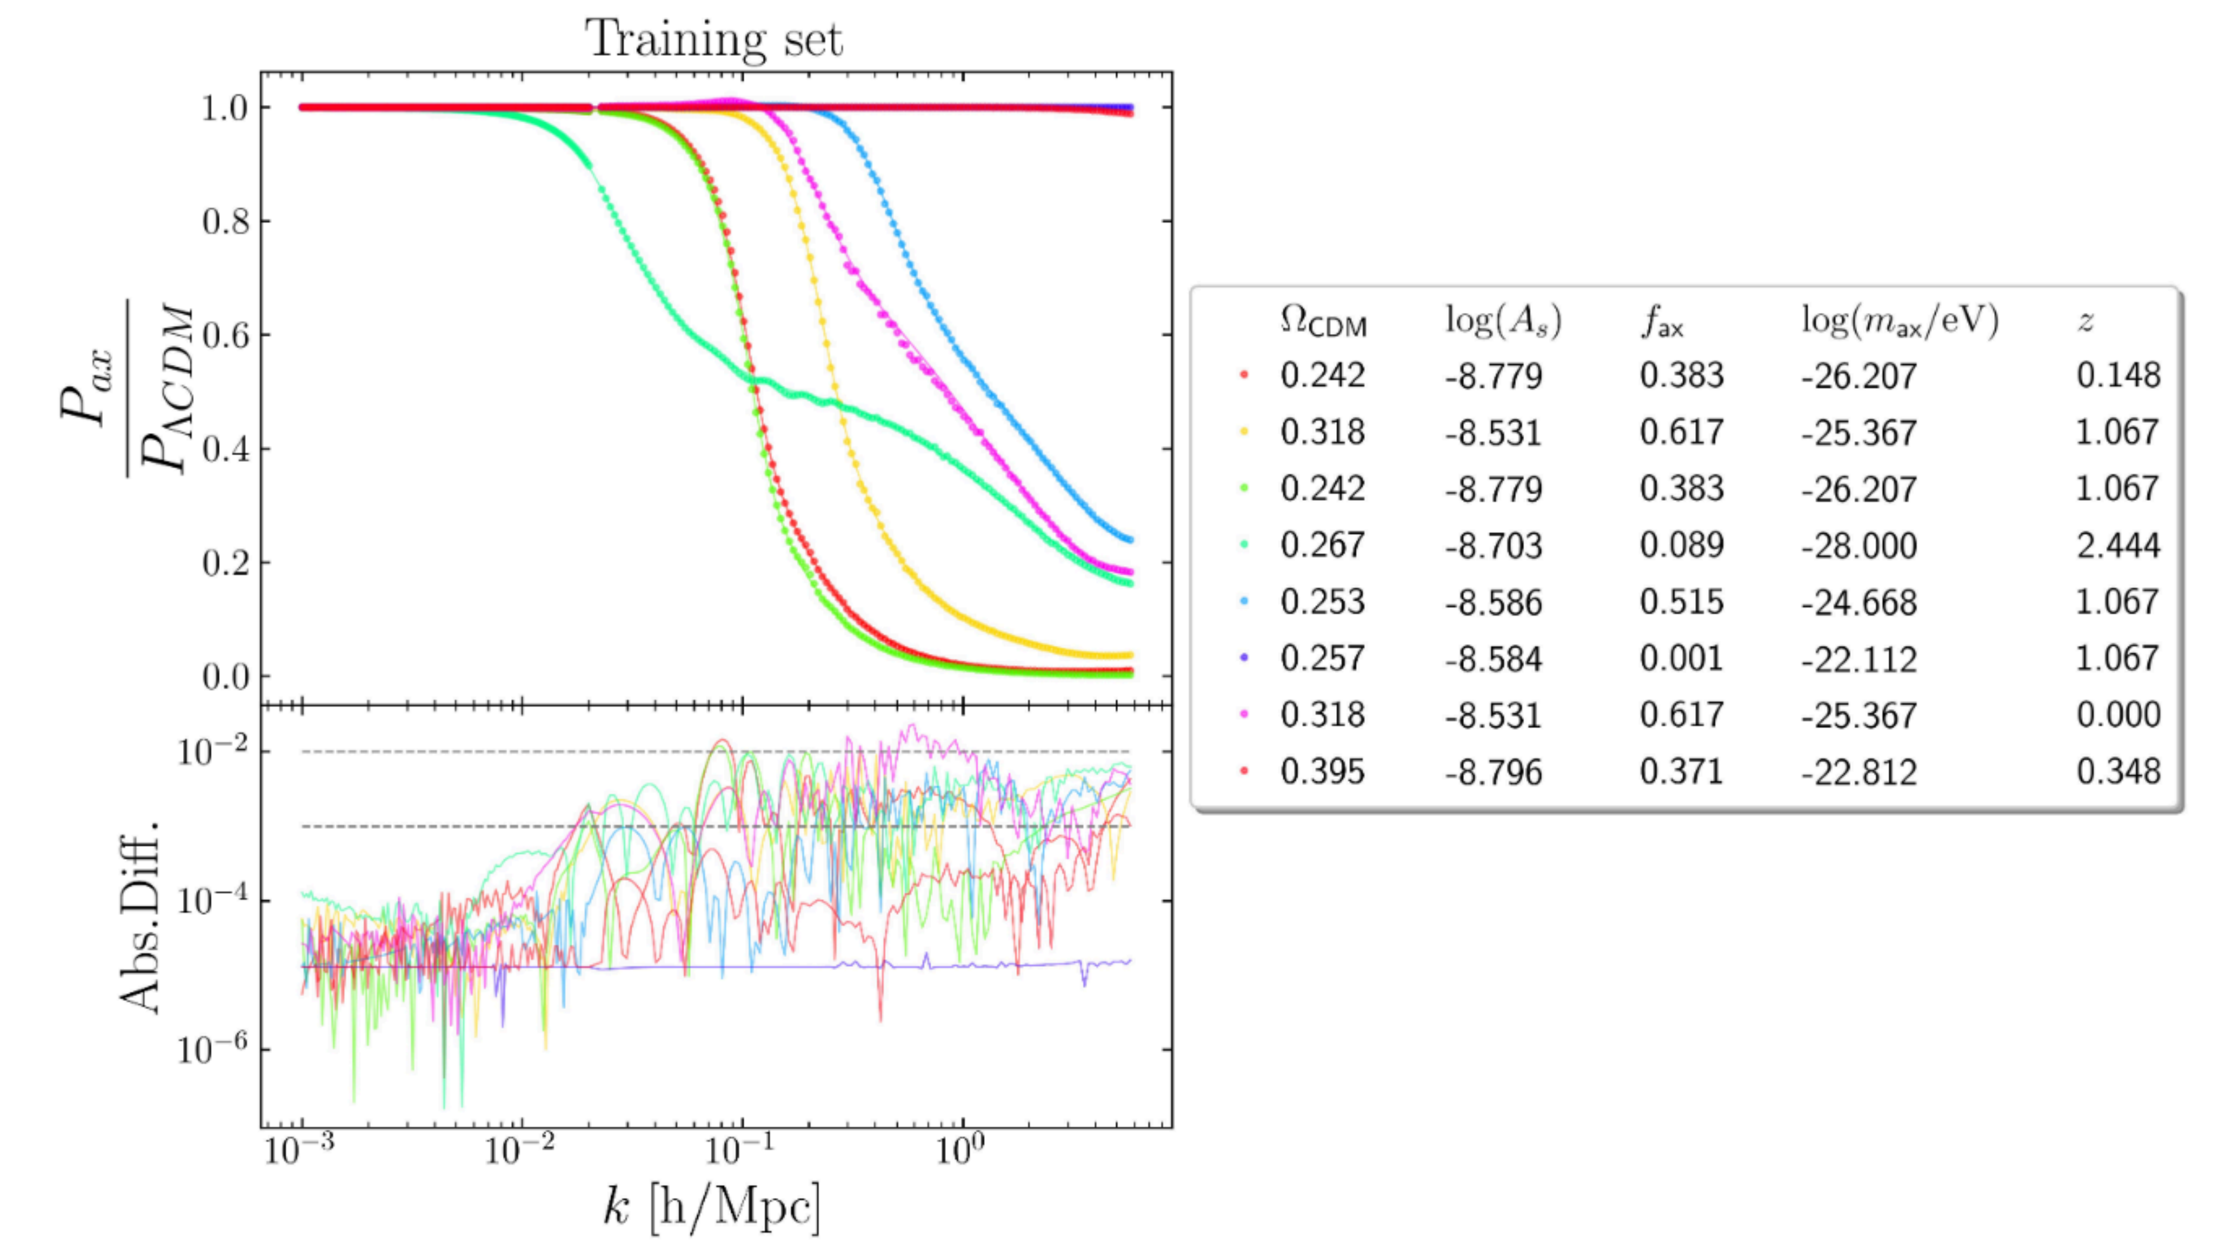

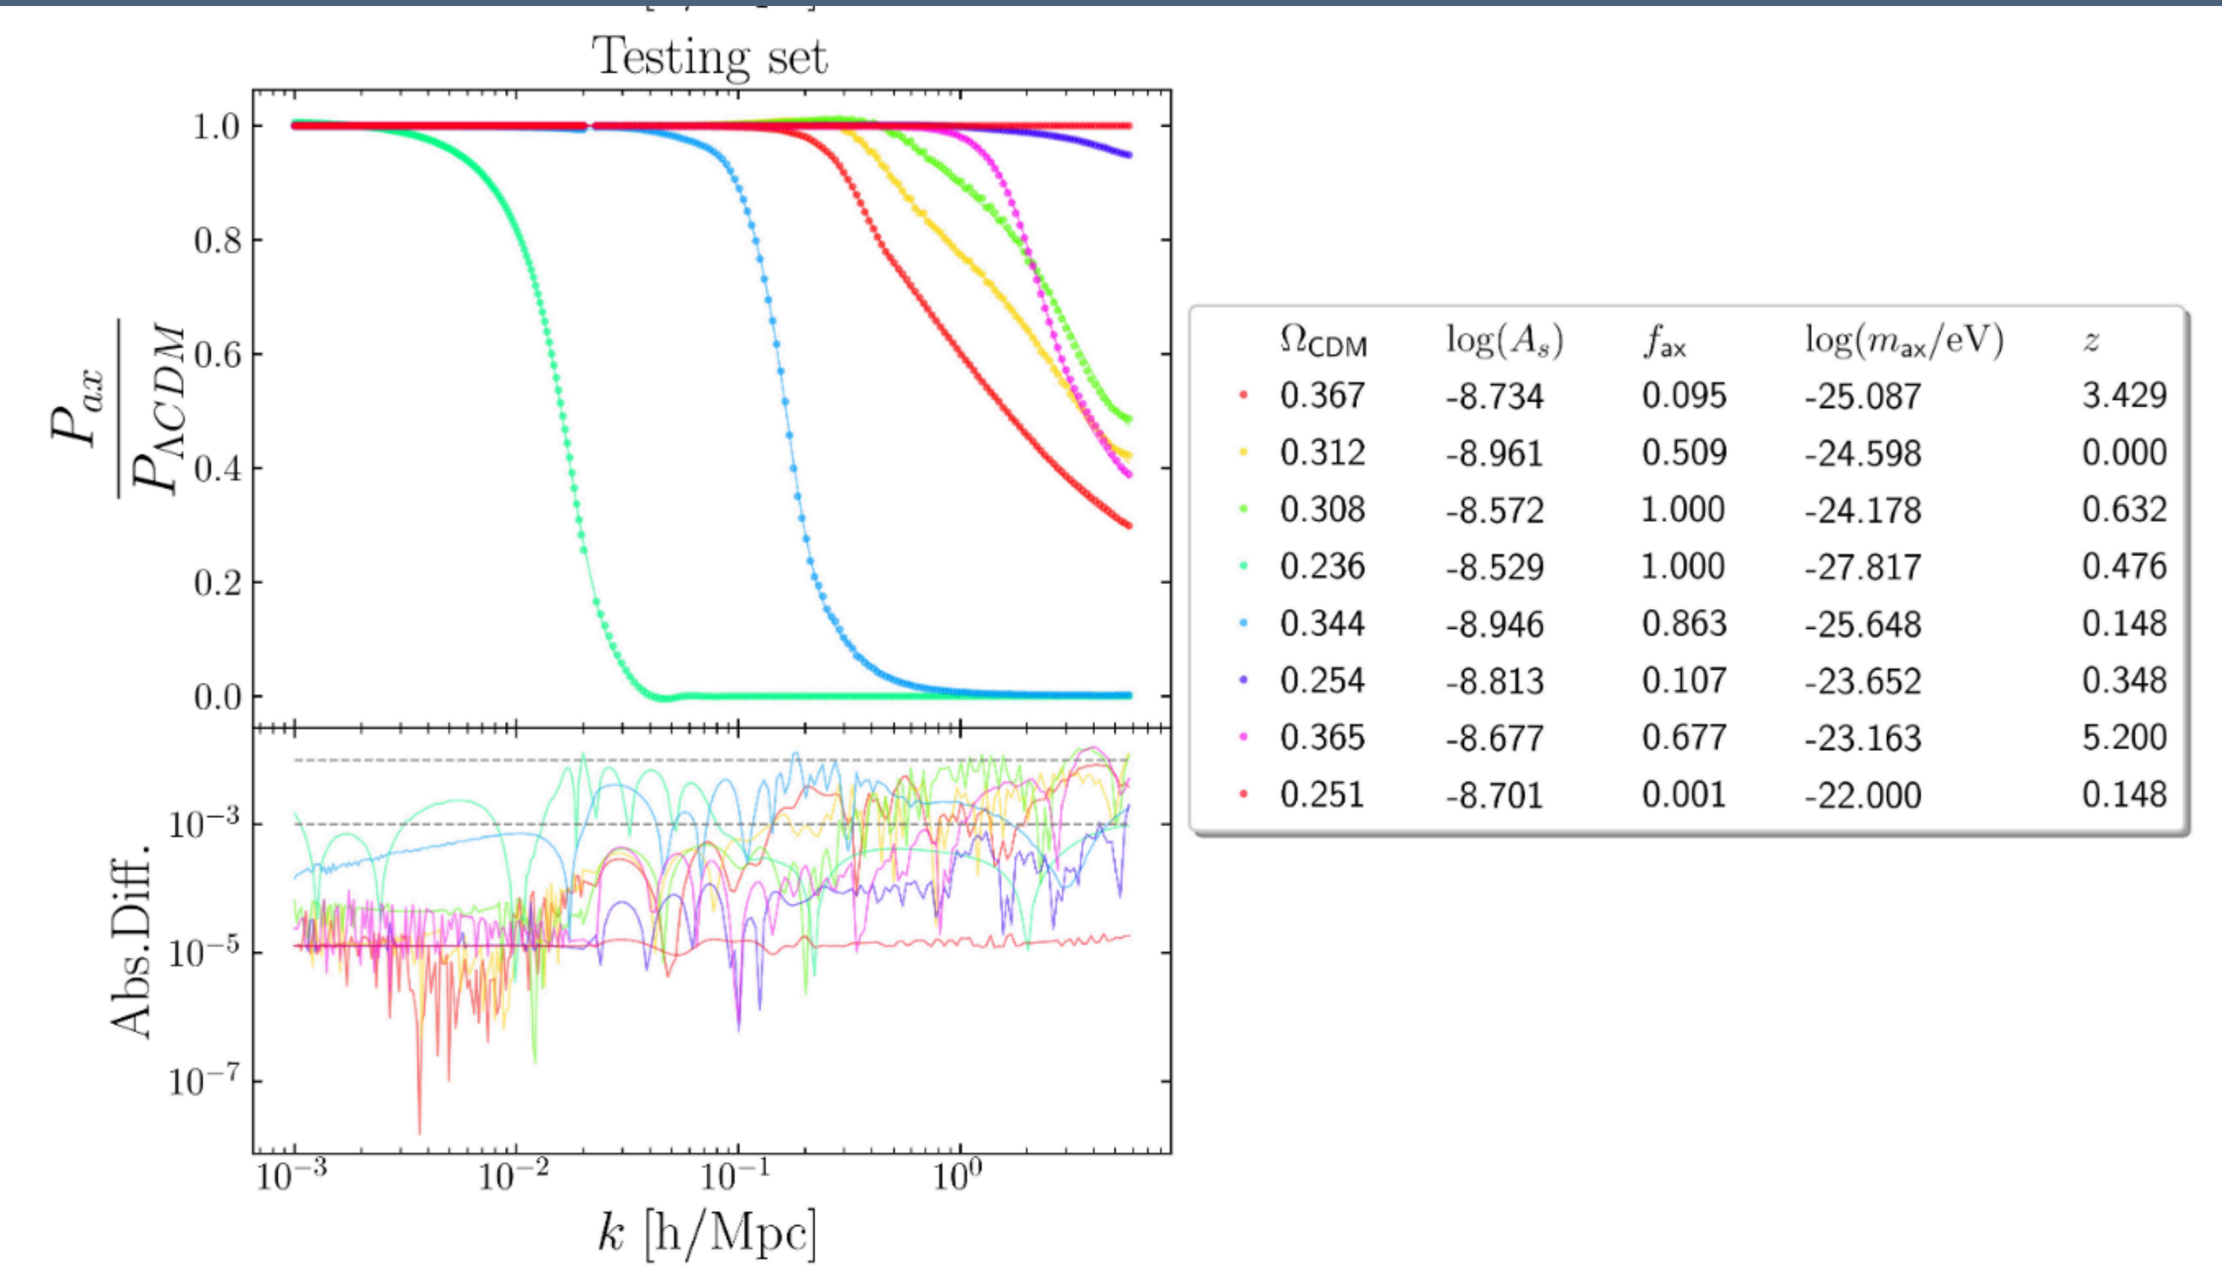

Results

- Comparison to other approaches:

  - axionCAMB (linear only)

  - HMcode ("fitting formula" to go from linear -> non-linear for LCDM)

  - axionHMcode (HMcode with axion physics included and tuned to simulations)

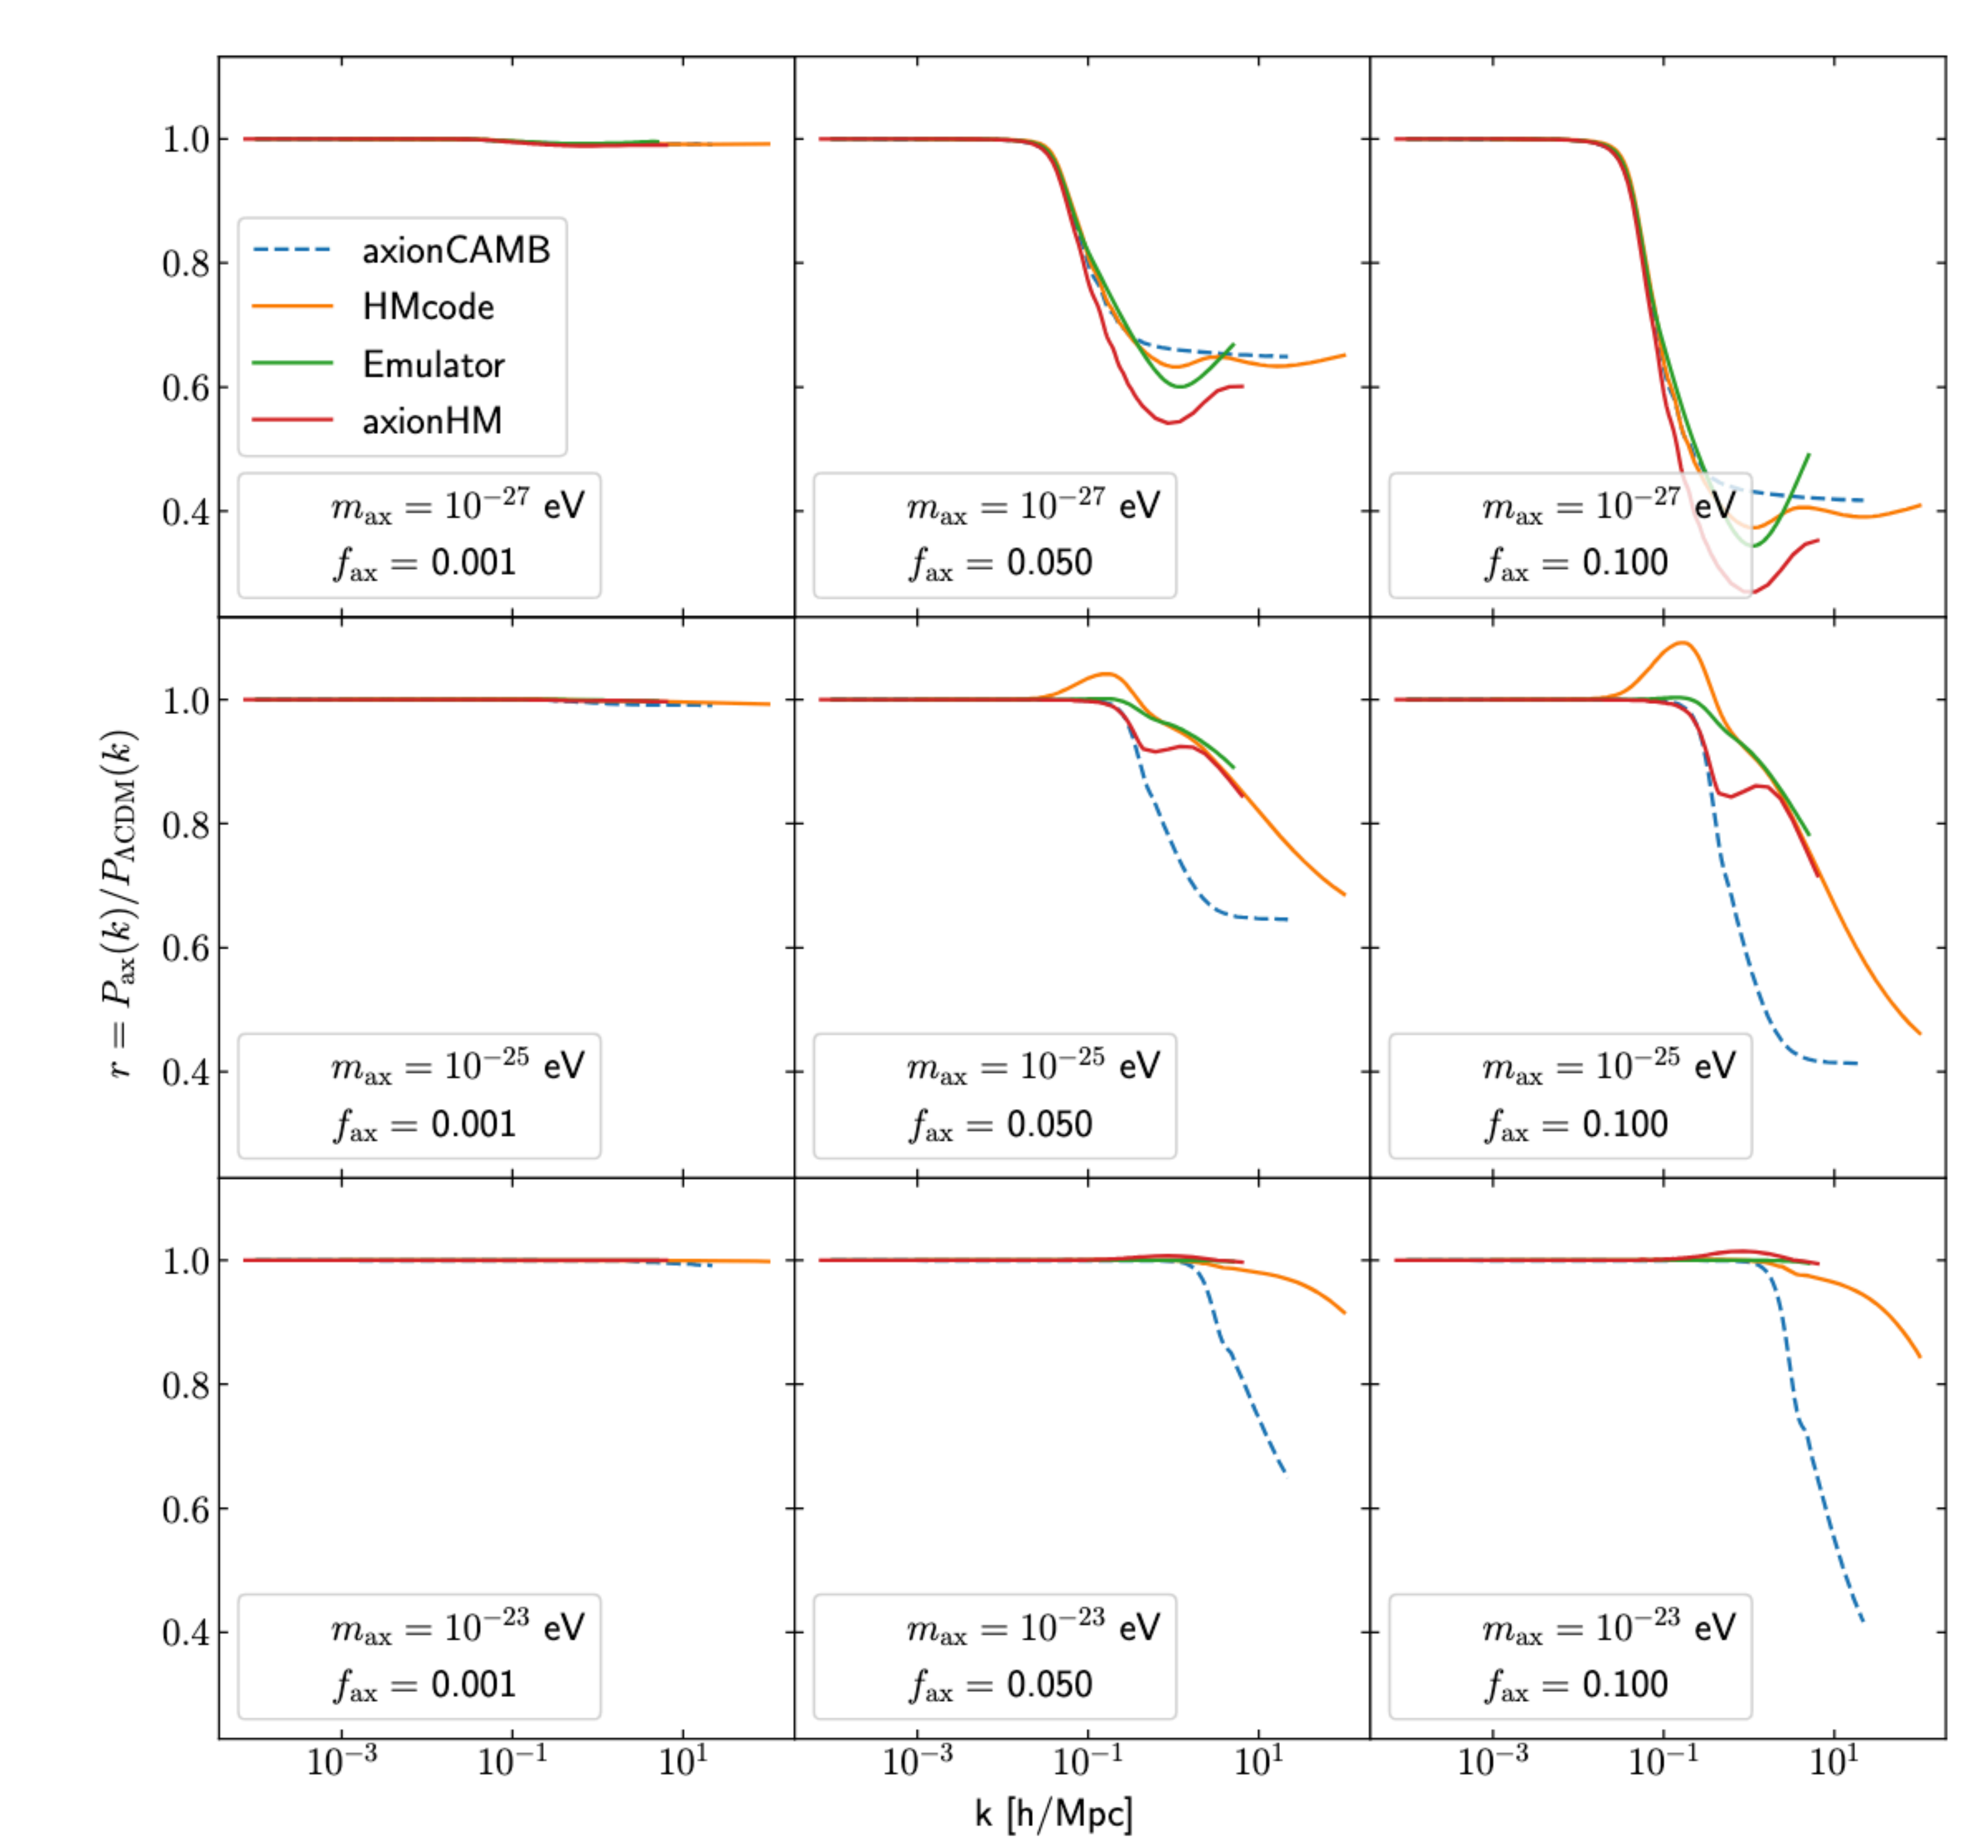

- Tests with full simulations *should* be performed and would be the ultimate test to see if this method is accurate enough, but we have not been able to do this test yet.

https://github.com/frdennis/Axion-Emulator.

Summary

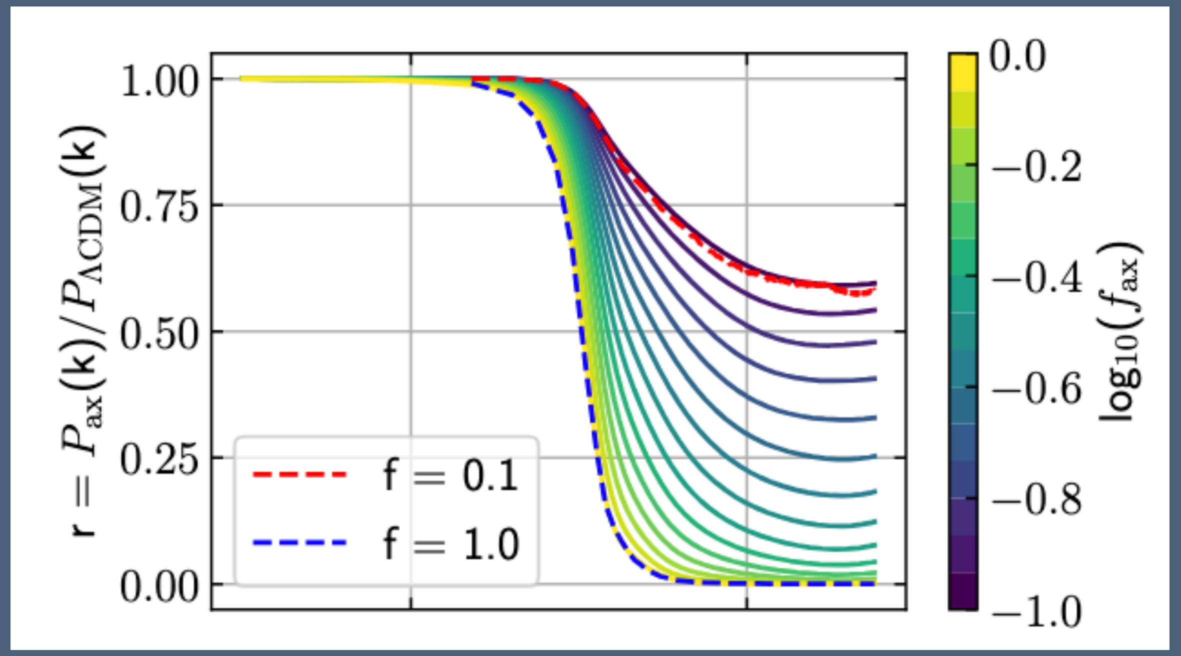

- We have created a neural network emulator for mixed axion models using fast approximate simulations (PM) where the axion physics is only injected only in the initial conditions (i.e. provides the suppression of power of small scales).

- **Advantages:**
  - Very cheap: a total cost of ~5000 CPU hours. Can be done on any local cluster. Compare this to the cost of one high-resolution axion simulation in a O(Mpc) box which is often O(100k - 1M) CPU hours.
  - Gives us an alternative non-linear prescription to axionHMcode.

- **Caveats:**
  - We don't have much of the axion physics (quantum pressure), everything is in the initial conditions, so we have to compare to simulations to get an estimate for this "theoretical error".
  - This test will have to be performed in the future.