

Improving statistical sensitivity of X-ray searches for axion-like particles

Joseph P. Conlon¹★ and Markus Rummel^{2,3}★

¹*Rudolf Peierls Centre for Theoretical Physics, 1 Keble Road, Oxford OX1 3NP, UK*

²*Department of Physics and Astronomy, McMaster University, Hamilton, ON L8S 4M1, Canada*

³*Perimeter Institute for Theoretical Physics, Waterloo, ON N2L 2Y5, Canada*

Accepted 2019 January 17. Received 2019 January 17; in original form 2018 September 5

ABSTRACT

X-ray observations of bright AGNs in or behind galaxy clusters offer unique capabilities to constrain axion-like particles. Existing analysis technique relies on measurements of the global goodness of fit. We develop a new analysis methodology that improves the statistical sensitivity to ALP–photon oscillations by isolating the characteristic quasi-sinusoidal modulations induced by ALPs. This involves analysing residuals in wavelength space allowing the Fourier structure to be made manifest as well as a machine learning approach. For telescopes with microcalorimeter resolution, simulations suggest these methods give an additional factor of 2 in sensitivity to ALPs compared to previous approaches.

Key words: astroparticle physics – elementary particles.

1 INTRODUCTION

Axion-like particles are one of the simplest and best motivated extensions of the Standard Model. They arise as a generalization of the QCD axion, which was introduced as an appealing solution to the strong CP problem of QCD (Peccei & Quinn 1977; Weinberg 1978; Wilczek 1978) – a recent review of axion physics is Marsh (2016). While the original QCD axion requires a coupling to the strong force, it is also interesting to consider more general *axion-like particles* (ALPs) that couple only to electromagnetism. Such ALPs arise generally in string compactifications (for example, see Conlon 2006; Svrcek & Witten 2006; Arvanitaki et al. 2010; Cicoli, Goodsell & Ringwald 2012). ALPs have a rich phenomenology, for instance they are a candidate for the dark matter (Abbott & Sikivie 1983; Preskill, Wise & Wilczek 1983; Arias et al. 2012). An ALP a interacts with photons via the coupling:

$$g_{a\gamma\gamma} a \mathbf{E} \cdot \mathbf{B}, \quad (1)$$

where $g_{a\gamma\gamma}$ is a dimensionful coupling parametrizing the strength of the interaction and \mathbf{E} and \mathbf{B} are the electric and magnetic fields, respectively. While we refer in this paper to ALPs, the physics is also relevant for photophilic models of the QCD axion in which the mass is much smaller (or the photon coupling significantly enhanced) compared to naive expectations, such as those in Farina et al. (2017) and Agrawal et al. (2018).

As they attain masses only by non-perturbative effects, ALPs naturally have extremely small masses. The relevant physics is

described by the Lagrangian

$$\mathcal{L} = \frac{1}{2} \partial_\mu a \partial^\mu a + \frac{1}{2} m_a^2 a^2 + g_{a\gamma\gamma} a \mathbf{E} \cdot \mathbf{B}. \quad (2)$$

The ALP–photon coupling produced by the $a \mathbf{E} \cdot \mathbf{B}$ interaction implies that, within background magnetic fields, the ALP state a has a two-particle interaction with the photon γ . Under this mixing, the ‘mass’ eigenstates of the Hamiltonian are a mixture of the photon and ALP ‘flavour’ eigenstates, causing oscillation between the modes in a way analogous to neutrino oscillations (Sikivie 1983; Raffelt & Stodolsky 1988).

ALP–photon conversion is enhanced by large magnetic field coherence lengths. As it extends over megaparsec scales and contains coherence lengths up to tens of kiloparsecs, this makes the intracluster medium of galaxy clusters particularly efficient (Burrage, Davis & Shaw 2009; Horns et al. 2012; Conlon & Marsh 2013; Angus et al. 2014; Powell 2015; Conlon, Marsh & Powell 2016; Schliederer & Sigl 2016). This has allowed the use of X-ray point sources, located in or behind clusters, to produce leading bounds on ALP parameter space (Wouters & Brun 2013; Berg et al. 2017; Conlon et al. 2017; Marsh et al. 2017; Chen & Conlon 2018; Conlon et al. 2018). While X-rays are particularly good at probing ALPs with $m \lesssim 10^{-11}$ eV, more massive ALPs can be probed via gamma-ray astronomy – see Wouters & Brun (2012), Ajello et al. (2016), Payez, Cudell & Hutsemekers (2012), and Majumdar, Calore & Horns (2018) for some related work in different wavebands.

The aim of this work is to improve the theoretical tools available for the analysis of ALP–photon conversion, and in particular the statistical methodology used to analyse data. Existing approaches to bounding ALP–photon conversion involve, roughly, a comparison of the overall goodness of fit (as measured by the reduced χ^2) of models with ALPs compared to models without ALPs. For

* E-mail: joseph.conlon@physics.ox.ac.uk (JPC); rummel@mcmaster.ca (MR)

sufficiently large ALP–photon coupling, models with ALPs lead to a bad fit.

However, this approach makes little use of the actual structure of ALP–photon conversion. While the precise form of ALP–photon conversion depends on the (unknown) magnetic field along the line of sight, there is a quasi-sinusoidal oscillatory structure that is common to all realizations of the magnetic field. The aim of this paper is to develop analysis techniques to isolate this structure, and thereby allow greater sensitivity to constraining ALPs. This development of improved statistical methods is also important in view of the launch over the next decade of microcalorimeter-based X-ray telescopes such as *XARM* and *Athena*; the use of an overall χ^2 fit would then fail to take full advantage of the energy resolution provided by such instruments.

With an eye on the opportunities that will come from future satellites such as *XARM* and *Athena*, we focus on energies and parameter ranges that are applicable for the case of X-ray photons converting to ALPs within galaxy cluster magnetic fields. However, our methods are in principle not limited to this case and can also be used for other photon energies whenever oscillations are expected. A version of our code can be found online.¹

2 STATISTICS OF ALP–PHOTON CONVERSION

To develop the analysis formalism, we start with the theory of photon–ALP conversion for a single magnetic field domain. For an X-ray photon passing through a single magnetic field domain (Sikivie 1983; Raffelt & Stodolsky 1988), we ignore contributions from vacuum birefringence and photon–photon dispersion as these are not relevant for X-ray photons in a typical intracluster medium. For a similar reason, we also neglect Faraday rotation and treat the plasma as optically inactive – such effects are negligible for X-rays in the intracluster medium.

$$P(\gamma \rightarrow a) = \frac{\Theta^2}{1 + \Theta^2} \sin^2 \left(\Delta \sqrt{1 + \Theta^2} \right), \quad (3)$$

where

$$\Theta = \frac{2g_{a\gamma\gamma} B_{\perp} \omega}{(\omega_{pl}^2 - m_a^2)}, \quad (4)$$

$$\Delta = \frac{(m_a^2 - \omega_{pl}^2) L}{4\omega}. \quad (5)$$

The numerical values of Θ and Δ are

$$\Theta = 0.02 \left(\frac{B}{1 \mu\text{G}} \right) \left(\frac{\omega}{1 \text{keV}} \right) \left(\frac{10^{-3} \text{cm}^{-3}}{n_e} \right) \left(\frac{g_{a\gamma\gamma}}{10^{-12} \text{GeV}^{-1}} \right), \quad (6)$$

$$\Delta = 0.54 \left(\frac{n_e}{10^{-3} \text{cm}^{-3}} \right) \left(\frac{L}{10 \text{kpc}} \right) \left(\frac{1 \text{keV}}{\omega} \right). \quad (7)$$

As values of $g_{a\gamma\gamma}$ significantly greater than 10^{-12}GeV^{-1} are already excluded (Berg et al. 2017; Conlon et al. 2017; Marsh et al. 2017), this implies that for typical parameters at X-ray energies $\Theta^2 \ll 1$ [for the regions with large $B \sim 10\text{--}50 \mu\text{G}$ such as at the centre of cool core clusters such as Perseus, Hydra, or A2199 cluster, $n_e \gtrsim 10^{-2} \text{cm}^{-3}$ (Taylor et al. 2006; Churazov et al. 2003) or Table 8

of Govoni et al. (2017)]. In this limit, we can simplify the single-domain conversion probability to

$$P(\gamma \rightarrow a) = \Theta^2 \sin^2 \Delta + \mathcal{O}(\Theta^4), \quad (8)$$

which as a function of energy behaves as

$$P(\omega) \sim \epsilon \left(\frac{\omega}{\omega_0} \right)^2 \sin^2 \left[\Delta_0 \left(\frac{\omega_0}{\omega} \right) \right] + \mathcal{O}(\Theta^4), \quad (9)$$

where ϵ and Δ_0 are constants, and ω_0 is a reference energy which we take to be 1 keV.

This equation involves an averaging over polarization states. As current X-ray telescopes are not sensitive to polarization, and the orientation of the cluster magnetic fields is independent of any polarization properties of the initial photons and randomized over many different domains, this is physically justified. For any one domain, the precise value of ϵ will depend on B_{\perp} , but in all cases ϵ is small. For X-ray telescopes with polarization sensitivity, these assumptions would need to be revisited.

Along a sightline through a cluster, the magnetic field undergoes many reversals. We can model the cluster magnetic field as a sequential series of domains, each differing in terms of direction and coherence length, with an overall field strength set by a radial falloff moving away from the centre of the cluster. As the field directions vary between each domain, the conversion amplitudes add incoherently. This allows us the total survival probability to be obtained by multiplying the survival probability for each domain (although note that for the simulated data we numerically solve the full transfer matrix for the conversion amplitude).²

We note that this model consists of separate sequential magnetic field domains, each with a different coherence length and magnetic field. A more accurate model would involve, instead of distinct domains, many different scales simultaneously present in the magnetic field. However, simulations of photon–ALP conversion in such cases do not show major qualitative differences from the case where the different scales are represented by separate sequential domains (Angus et al. 2014). As it is conceptually and computationally much simpler, and does not appear to be too misleading, we shall therefore work with a magnetic field model of separate and sequential domains. We have tested that the simplified approach discussed above leads to very similar results as a full simulation of an initial state through a simulated turbulent magnetic field that does not make use of the single domain formula but recursively solves the linearized Schrödinger equation of Raffelt & Stodolsky (1988).

In this paper, we will always assume that the overall photon–axion conversion probability is very small (less than 10 per cent) and the equations used in this section are only valid in this limit. This is not a significant restriction – if it is not the case, then the more advanced techniques of this paper are unnecessary as a simple test such as an overall χ^2 fit will already show that the ‘standard’ astrophysical fit fails to describe the data.

The combination of the $\gtrsim \mathcal{O}(100)$ domains required to describe a Mpc-scale cluster and an overall conversion probability of less than 10 per cent implies that the conversion probability within any individual single domain is highly suppressed. Under these approximations, the overall photon survival probability can be

¹https://github.com/mrummphys/axion_statistics

²Although the axion propagation equations evolve amplitude not probability, it is the fact that different domains add incoherently that makes it sufficient to combine probabilities across domains.

simplified,

$$P(\gamma \rightarrow \gamma, \omega) = \prod_i \left\{ 1 - \epsilon_i \left(\frac{\omega}{\omega_0} \right)^2 \sin^2 \left[\Delta_i \left(\frac{\omega_0}{\omega} \right) \right] \right\} \\ \simeq 1 - \sum_i \epsilon_i \left(\frac{\omega}{\omega_0} \right)^2 \sin^2 \left[\Delta_i \left(\frac{\omega_0}{\omega} \right) \right], \quad (10)$$

as by definition of the small conversion regime, any $\mathcal{O}(\epsilon^2)$ $\gamma \rightarrow a \rightarrow \gamma$ re-conversion terms are negligible. We stress that when simulating data we use the full 3×3 mixing formalism of Raffelt & Stodolsky (1988) to determine conversion probabilities. What the above formula gives us is an analytic understanding, valid in the small conversion probability regime, which we use as motivation for our methods of data analysis.

This survival probability modulates the arriving photon spectrum. We shall assume that we know the correct functional form of the original source spectrum of photons, which we denote as $\mathcal{M}(\omega)$. For the purpose of this paper, we shall take this as a featureless power law plus an Fe K α line, although for more complicated spectra, this may also need to include more emission and absorption lines at (known) atomic energies.

The spectrum of (measured) arriving photons $\mathcal{A}(\omega)$ is then the source spectrum modulated by the $P(\gamma \rightarrow \gamma, \omega)$ survival probability,

$$\mathcal{A}(\omega) = P(\gamma \rightarrow \gamma, \omega) \mathcal{M}(\omega) \\ = \left\{ 1 - \sum_i \epsilon_i \left(\frac{\omega}{\omega_0} \right)^2 \sin^2 \left[\Delta_i \left(\frac{\omega_0}{\omega} \right) \right] \right\} \mathcal{M}(\omega). \quad (11)$$

This can be decomposed into global and oscillatory terms:

$$\mathcal{A}(\omega) = \left\{ 1 - \sum_i \frac{\epsilon_i}{2} \left(\frac{\omega}{\omega_0} \right)^2 + \sum_i \frac{\epsilon_i}{2} \left(\frac{\omega}{\omega_0} \right)^2 \cos \left[2\Delta_i \left(\frac{\omega_0}{\omega} \right) \right] \right\} \mathcal{M}(\omega). \quad (12)$$

The distinctive ALP-induced modulations lie in the sinusoidal terms, which average to zero. To isolate these we want to first remove the global term, and so the function we should therefore fit to the data is not $\mathcal{M}(\omega)$ itself, but instead

$$\mathcal{M}_a(\omega) = \left[1 - A \left(\frac{\omega}{\omega_0} \right)^2 \right] \mathcal{M}(\omega). \quad (13)$$

In the presence of axions, a fit to $\mathcal{M}_a(\omega)$ will produce oscillatory residuals equally distributed about zero. These residuals will then take the form

$$\mathcal{R}(\omega) = \mathcal{A}_{\text{true}}(\omega) - \mathcal{A}_{\text{fitted}}(\omega) \\ = \sum_i \frac{\epsilon_i}{2} \left(\frac{\omega}{\omega_0} \right)^2 \cos \left[2\Delta_i \left(\frac{\omega_0}{\omega} \right) \right] \mathcal{M}(\omega), \quad (14)$$

where the prefactor of $\mathcal{M}(\omega)$ is a small quantity. Taking the data/model ratio, we therefore obtain the fractional residual

$$\mathcal{F}(\omega) = \frac{\mathcal{A}(\omega) - \mathcal{M}_a(\omega)}{\mathcal{M}(\omega)} = \sum_i \frac{\epsilon_i}{2} \left(\frac{\omega}{\omega_0} \right)^2 \cos \left[2\Delta_i \left(\frac{\omega_0}{\omega} \right) \right]. \quad (15)$$

In this expression, we have self-consistently neglected terms that are second order in smallness – if such terms are important, the simple overall χ^2 fit will have already revealed the failure of the ‘standard’ astrophysical fit.

The above quantity is purely numerical and equally distributed about zero. As expressed above, it is a function of energy. However, the ω^{-1} terms imply this is not best analysed in frequency space.

Rather, by rewriting it in wavelength space and multiplying through by λ^2 , we obtain

$$\left(\frac{\lambda}{\lambda_0} \right)^2 \mathcal{F}(\lambda) = \sum_i \frac{\epsilon_i}{2} \cos \left[2\Delta_i \left(\frac{\lambda}{\lambda_0} \right) \right]. \quad (16)$$

We now have a quantity that can be analysed in a reasonably direct fashion as a Fourier series. The contributing Fourier frequencies depend on the coherence lengths and electron densities of the individual domains. So long as these remain broadly similar throughout the cluster, by determining $\lambda^2 \mathcal{F}(\lambda)$ from data, we can search for axion-induced structure by performing a Fourier decomposition of it.

3 ALP SEARCH STRATEGIES

The oscillatory structure of the residuals described in equation (16) can be exploited using any of the following three strategies:

- (i) An analysis of the Fourier transformed data,
- (ii) a sinusoidal fit to the data,
- (iii) a machine learning approach.

The Fourier transform of the data set can be directly calculated using fast Fourier transform methods. As opposed to a fit, this does not require optimization of any of the parameters involved. Fitting to a small number of sinusoidal functions can be an efficient way to capture any oscillatory structure present in the residuals. The machine learning approach uses the complete list of residual data as an input to ‘learn’ from a training set what value of g is associated. Contrary to the Fourier and fit method, the machine learning method does not rely on the data being of oscillatory nature. It can be thought of as a universal fit: once the network is trained it includes all the relevant information that influence the conversion of photons to axions along the line of sight. Hence, it can be used as a benchmark for the Fourier and fit method: once those methods perform as well as the somewhat brute force machine learning approach, we know that we are able to resolve – and understand – all the possible structure of photon to ALP conversion imprinted on to the residuals.

In the following, we present the three methods in more detail, allowing a comparison between each of the three.

3.1 Analysis of Fourier-transformed data

Given a data set analysed into the form of equation (16), we want to perform a Fourier analysis of the fractional residuals

$$(u_i, y_i) = (\omega_i^{-1}, \omega_i^{-2} \mathcal{F}(\omega_i)), \quad (17)$$

where ω_i is the energy of the i -th bin. In general, we do not expect the values u_i to be equispaced. We therefore require the tools of non-equispaced Fourier analysis and in particular the calculation of the non-equispaced discrete Fourier transform (NDFT). In this context, we make use of the tools developed in the context of non-equispaced fast Fourier transforms (NFFT).

The NDFT is defined³ by expanding a function $y(x)$ specified on a total of M data points with respect to N trigonometric polynomials as

$$y(x) = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \hat{y}_k e^{-2\pi i k x}. \quad (18)$$

³For example, see <https://www-user.tu-chemnitz.de/potts/nfft/>.

For a data set of values y_j and non-equispaced nodes x_j with $j = 0, \dots, M-1$ we can write

$$y_j \equiv y(x_j) = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \hat{y}_k e^{-2\pi i k x_j}, \quad (19)$$

which can be written as $\mathbf{y} = \mathbf{A} \hat{\mathbf{y}}$ with

$$\mathbf{y} = (y_j)_{j=0, \dots, M-1}, \quad (20)$$

$$\mathbf{A} = (e^{-2\pi i k x_j})_{k=-N/2, \dots, N/2-1; j=0, \dots, M-1}, \quad (21)$$

$$\hat{\mathbf{y}} = (\hat{y}_j)_{j=0, \dots, M-1}. \quad (22)$$

The Fourier transform $\hat{\mathbf{y}}$ can hence be calculated from the inverse of the $N \times M$ matrix \mathbf{A} .

For periodicity the nodes x_j are to be chosen in the interval $-\frac{1}{2} \leq x_j < \frac{1}{2}$. Hence, we linearly rescale the nodes of our data set to place them in this interval,

$$x_j = \frac{M-1}{M} \frac{u_j - u_0}{u_{M-1} - u_0} - \frac{1}{2}. \quad (23)$$

Note that for the case of equispaced nodes

$$u_j = u_0 + \frac{j}{M-1} (u_{M-1} - u_0). \quad (24)$$

Equation (19) reduces to the standard discrete Fourier transform

$$y_j = \sum_{k=-\frac{N}{2}}^{\frac{N}{2}-1} \hat{y}_k e^{-2\pi i k \left(\frac{j}{M} - \frac{1}{2}\right)}, \quad (25)$$

for $j = 0, \dots, M-1$.

If N is large enough, this of course gives a perfect fit. However, if photon-ALP conversion occurs then the data set should be well described by a restricted number of Fourier modes. By truncating this sum,

$$y^{(K)}(x) \equiv \text{Re} \left(\sum_{k=-K}^{K-1} \hat{y}_k e^{-2\pi i k x} \right), \quad (26)$$

we obtain an approximation of the data to mode order $K \leq M/2$ [in the case of $K = M/2$, by equation (19) $y^{(K)}(x)$ passes through all real data points (x_i, y_i)].

The essence of the axion search is that we expect the axions to introduce modulations that allow this data to be well approximated by only its low-frequency Fourier modes, i.e. $K \ll M/2$ (at least in the limit of a very large number of data points).

3.2 Sinusoidal fit

For this method, we simply fit the residuals (17) that are expected to behave sinusoidally in the presence of ALPs. As fitting functions, we use N_{fit} sine functions with amplitudes A_i and phases ϕ_i :

$$y_{\text{fit}}^{(N_{\text{fit}})}(u) = \sum_{i=1}^{N_{\text{fit}}} A_i \sin(f_i u + \phi_i). \quad (27)$$

In the spirit of a Fourier analysis, we fix the frequencies of the sine functions in (27) to be

$$f_i = \frac{2\pi i}{u_M - u_0}, \quad (28)$$

i.e. the lowest frequency mode $i = 1$ has one period in the interval $u_0 < u < u_M$ and the higher frequency modes have i periods. The fit to the data then involves minimization of a χ^2 function with $2N_{\text{fit}}$ fit parameters. In principle, we could also treat the frequencies f_i as fit parameters but in practice we find that the fit converges too slowly in this case and the amplitudes and phases allow enough flexibility of the fit function to capture the ALP-induced oscillations.

3.3 Machine learning

We set up a deep neural network (DNN) with Tensorflow⁴ that uses the list of residuals \mathcal{F}_i in its entirety as ‘feature columns’ while the value of g that has been used to simulate those residuals is used as the ‘label column’. The design parameters of the DNN are the number of layers and the number of nodes as well as the loss function that the network tries to optimize. For the loss function, we use the sum of squared errors between the actual and the DNN predicted values of the residuals \mathcal{F}_i . The data are shuffled and divided into a training (75 per cent of data) and a test set (25 per cent of data). The training data is used to teach the network which value of g is associated to a list of residuals by optimizing the parameters of each node while the test data is used to test how well the DNN performs and to estimate the prediction error. The data contain a sample of different realizations of the magnetic field model and Poisson errors. Hence, the DNN does not learn the residual structure for a particular realization of the magnetic field but its marginalized statistical properties. The number of layers and nodes are chosen such that the difference between the test error and training error is minimized, i.e. the network is built with enough flexibility to capture the data but avoids overfitting. We try one, two, and three layers and vary the number of nodes in steps of 10. Examining the training and test error, we find the optimal value to be at two layers and 200 (40) nodes for *Athena* (*XMM*). Three layers generally lead to overfitting even for a small number of nodes.

4 ALP DATA SIMULATION

We now aim to test this method on simulated data for *XMM-Newton* and *Athena*. Our focus here is on comparison of this method to a pure overall χ^2 fit (and hence whether it is possible to improve statistical sensitivity), rather than on applications to any particular real data set or to precise forecasts of future telescope sensitivity.

However, as we do want our simulated data set to be motivated by real physics, we take as a starting point for our simulations an absorbed power law together with an iron line at 6.4 keV, reflecting a typical AGN spectrum. We consider the simulated spectra of the model

$$wabs * ALP * (zpowerlw + zgauss). \quad (29)$$

This equation represents the initial AGN spectrum, modulated by the effects of photon-ALP conversion, and then absorbed by the Galactic column density of neutral Hydrogen. Here, *ALP* represents the effect of photon-ALP conversion within the cluster magnetic field.

As one of the most promising targets for ALP-photon conversion is the central Perseus AGN NGC 1275, we let *ALP* represent simulated survival probabilities for the Perseus cluster environment, i.e. for ALP-photon conversion within an electron density n_e and

⁴We use the DNNRegressor estimator. For more details, see the Tensorflow documentation.

magnetic field model representing that of the Perseus cluster. As the B-field properties can only ever be known statistically, and each explicit realization leads to different conversion probabilities even for statistically identical B-field parameters, we need to marginalize over many realizations of the magnetic field in Perseus.

The properties of the magnetic field that enter the simulations are those used for Perseus in Berg et al. (2017): we use a central magnetic field strength of 25 μG (Taylor et al. 2006) and assume that B decreases with radius as in the Coma cluster $B(r) \propto [n_e(r)]^{0.7}$ (Bonafede et al. 2010). The electron density as a function of radius in Perseus is (Churazov et al. 2003)

$$n_e(r) = \frac{3.9 \times 10^{-2}}{\left[1 + \left(\frac{r}{80 \text{ kpc}}\right)^2\right]^{1.8}} + \frac{4.05 \times 10^{-3}}{\left[1 + \left(\frac{r}{280 \text{ kpc}}\right)^2\right]^{0.87}} \text{ cm}^{-3}. \quad (30)$$

Finally, the domain length is a random parameter that is drawn from a power-law distribution with power 0.8 and with values between 3.5 and 10 kpc, motivated by Vacca et al. (2012). Each magnetic field simulation comprises a total of 300 domains yielding a total magnetic field region between 1 and 3 Mpc with typical values around 2 Mpc. In the simulation, the mixed axion-photon state propagates through each magnetic field realization with the above parameters. As we can see from equation (30), the magnetic field varies on scales of the order of 80–280 kpc which is much larger than the coherence scale. Hence, the magnetic field can be approximated as constant within each domain. This is typical for cluster magnetic field environments.

We simulate ALP to photon couplings in the range

$$10^{-14} \text{ GeV}^{-1} \leq g \leq 10^{-11} \text{ GeV}^{-1}. \quad (31)$$

This represents a range from values beyond observational reach to those that have previously been excluded.

Based on *XMM-Newton* pn and predicted *Athena* XIFU responses,⁵ we generate fake data using the XSPEC command `fakeit` with a simulated exposure time of 500 ks. `fakeit` simulates Poisson errors based on the exposure time. In terms of fitting, we work with an energy interval between 1.5 and 4.5 keV for *XMM-Newton* and from 0.7 to 1.4 keV for *Athena*. This allows the demonstration of the ability of *Athena* (or any other telescope with microcalorimeter level resolution) to resolve the rapid ALP-induced oscillations that would be present at lower energies. For a real data set, one would of course use the full energy range of the telescope.

We now describe the methods used to analyse the fake data sets:

(i) We first fit a zero model

$$wabs * ALPGlob * (zpowerlw + zgauss) \quad (32)$$

spectrum to the data. The *ALPGlob* component represents the global modulation due to ALPs given in (13) and is given as

$$ALPGlob = 1 - A \omega^2, \quad (33)$$

where $A = \omega_0^{-2} \sum \epsilon_i$.

In the absence of ALPs, one would expect the best-fitting value of A to vanish. In the presence of ALPs, while such a term removes the ‘global’ effects it will not account for the oscillations contained within the *ALP* component in (29). Hence, for a strong axion–photon coupling χ_0^2/dof would be significantly larger than one.

(ii) We perform a Fourier analysis (as described in Section 3.1) and a sinusoidal fit (as described in Section 3.2) of $(x_i, y_i) = (\omega_i^{-1}, \omega_i^{-2} \mathcal{F}(\omega_i))$ for a given K .⁶ If low mode oscillations are present in the data, then

$$\chi_1^2 = \sum_{i=1}^N \frac{[\mathcal{F}(\omega_i) - \omega_i^2 y(\omega_i)]^2}{\sigma_i^2} \quad (34)$$

per dof will improve significantly compared to χ_0^2 . The function $y(\omega)$ equals $y^{(K)}(\omega_i)$ for the Fourier method and $y^{(N_{\text{fit}})}(\omega_i)$ for the sinusoidal fit.

The Fourier function $y^{(K)}(\omega)$ generally overestimates the fluctuations in the simulated data $\omega_i^{-2} \mathcal{F}(\omega_i)$ at the edges of the energy intervals. This is particularly significant for *XMM* data where the residuals are often constant for large $\omega \lesssim 4.5$ keV while $y^{(K)}(\omega)$ can oscillate strongly. This would lead to a bad fit to the data which is why we use a smoothing function $S(\omega)$ that multiplies $y^{(K)}(\omega)$ to smooth it out at the edges of the energy interval. For the Fourier method only, χ_1^2 is then calculated as

$$\chi_{1,\text{Fourier}}^2 = \sum_{i=1}^N \frac{[\mathcal{F}(\omega_i) - \omega_i^2 S(\omega_i) y^{(K)}(\omega_i)]^2}{\sigma_i^2}, \quad (35)$$

where we choose $S(\omega) = 1 \text{ keV}^2/\omega^2$ as a smoothing function.⁷

(iii) The DNN is fitted by using 75 per cent of the residuals \mathcal{F}_i as training data. For the remaining 25 per cent of the residual data set the DNN predicts the value of g .

We show in Fig. 1 how low Fourier modes can capture residuals in the data, for *XMM-Newton* and *Athena*, respectively. The values of N_{fit} and K are chosen such that the low Fourier mode oscillations in the data can be accommodated. For this purpose, we examine plots such as Fig. 1 for different realizations of the magnetic field model and check how many Fourier modes we need to include to represent the oscillatory structure. Since *Athena* has a much better energy resolution than *XMM-Newton*, it can resolve a higher number of high-frequency modes.

For the Fourier and fit method, whatever the precise method used, the quantity we are ultimately interested in is

$$\Delta\chi^2(g) = \chi_0^2(g) - \chi_1^2(g). \quad (36)$$

where g is the ALP–photon coupling. This represents the ability to improve the fit by including Fourier modes ($\chi_0^2(g)$ is equivalent to the fit statistic with no Fourier modes included). The methods described in this paper are most relevant when $\chi_0^2(g)$ represents an acceptable fit – but a fit which can be significantly improved by adding Fourier modes. If $\chi_0^2(g)$ already gives an unacceptable fit, then the techniques of this paper are largely superfluous.

Photon to ALP modulations are very well represented by the leading Fourier modes and hence this analysis is very sensitive to the presence of ALP physics. An atomic line for instance would lead to virtually no improvement in $\Delta\chi^2(g)$ when using a Fourier analysis and hence would not be picked up by our method. The method is also very efficient as the calculation of Fourier coefficients can be performed very quickly and no spectral fit beyond the standard absorbed power law is necessary.

⁶For *Athena* data, we rebin two data points to one data point to speed up the calculation as half the energy resolution of *Athena* is easily sufficient to resolve ALP modulations.

⁷We have tried different powers of ω as smoothing functions and found that ω^{-2} has the desired effect as it smooths along the edges but does not suppress too much structure in the bulk of the given energy range.

⁵These are taken from <https://www.cosmos.esa.int/web/xmm-newton/epic-response-files> and www.cfa.harvard.edu/jzuhone/soxs/responses.html.




Figure 1. For particular B-field and Poisson error realizations, we show plots of residuals $y_i = \omega_i^{-2} \mathcal{F}(\omega_i)$ versus low mode sinusoidal fit function $y_{fit}^{(8)}$ for *XMM* (left side), and versus Fourier decomposition $y_{fit}^{(16)}$ for *Athena* (right side) as a function of inverse energy. The ALP to photon coupling is $g = 5 \times 10^{-12} \text{ GeV}^{-1}$.

Note that $\Delta\chi^2(g=0)$ represents improvements in the fit due simply to Poisson noise – even in the absence of ALPs. It is expected that some improvement can be attained by fitting the Poisson noise of the residuals with sinusoidal functions. Constraints will be obtained when $\Delta\chi^2(g) \gg \Delta\chi^2(g=0)$.

For every g , we generate 40 realizations of the magnetic field in Perseus and subsequently 50 fake data sets with different Poisson errors. Hence, for every g there are 2000 sets of simulated data.

In order to set bounds on g , one has to compare the simulated $\Delta\chi^2$ distributions to the $\Delta\chi^2$ of an *actual* observation $\Delta\chi_{obs}^2$, i.e. for the observed data of e.g. a point source there is just one realization of g , the magnetic field and Poisson errors. If $\Delta\chi_{obs}^2$ is lower than the 95th percentile of the $\Delta\chi^2(g=0)$ distribution we cannot exclude that the modulations are purely due to Poisson noise and we cannot set any bounds. Let us now discuss the case that the observed $\Delta\chi^2$ is larger. We take as the null hypothesis the statement that ALPs exist with a given coupling g . The null hypothesis is excluded at 95 percent, if 95 percent of the simulated values of $\Delta\chi^2(g)$ are larger than the measured $\Delta\chi_{obs}^2$. Since any percentile value of the $\Delta\chi^2(g \neq 0)$ distribution is expected to grow monotonically in g the bound g_{bound} is defined via

$$\text{5thpercentile} [\Delta\chi^2(g_{bound})] = \Delta\chi_{obs}^2. \quad (37)$$

If there large modulations are present in the data (i.e. $\Delta\chi_{obs}^2$ is large) we can only set a weak bound on g , i.e. g_{bound} is large. However, if $\Delta\chi_{obs}^2$ is small we can set a stronger bound on g and g_{bound} is small. These effects are illustrated in Fig. 2, where we see that a larger coupling leads (statistically) to a larger $\Delta\chi^2(g)$.

For the DNN, we can use similar criteria for exclusion when directly using percentages of the DNN predicted value of $g = G_{DNN}$ instead of $\Delta\chi^2$ as for the Fourier/fit methods. For real data, the set of measured residuals is fed into the DNN that has been trained with simulated data which results in a predicted value of $G_{obs} := G_{DNN}(\mathcal{F}_{i,obs})$. If the DNN predicted value of G_{obs} is lower than the 95th percentile of the values the DNN predicts for $g=0$ from the test set, then the residuals being purely due to Poisson noise cannot be excluded. Similarly to the previous discussion of setting bounds on g for the Fourier and fit method, the 95 percent exclusion bound in the DNN method is defined via

$$\text{5th percentile} [G_{DNN}(\mathcal{F}_{i,sim} | g_{bound})] = G_{obs}, \quad (38)$$




Figure 2. Histogram of the $\Delta\chi^2$ distributions for $g = 2 \times 10^{-13} \text{ GeV}^{-1}$ (red) and $g = 5 \times 10^{-13} \text{ GeV}^{-1}$ (blue). As g increases the distribution moves towards larger $\Delta\chi^2$. The red dotted line shows the 5th percentile of the $\Delta\chi^2$ distribution for $g = 2 \times 10^{-13} \text{ GeV}^{-1}$. The 5th percentile of the simulated data is used to infer a bound on g , see main text. This plot uses the sinusoidal fit method for *XMM-Newton* simulated data.

where $G_{DNN}(\mathcal{F}_{i,sim} | g)$ are the DNN predictions on the simulated data for a given g in the test set.

5 RESULTS

In the following, we compare the standard χ_0^2 calculation that was previously used to look for ALPs to the Fourier method explained in Section 3.1, the sinusoidal fit method explained in Section 3.2, and the machine learning approach discussed in Section 3.3. We apply these methods to the simulated data for *XMM* and *Athena*.

5.1 Athena

The standard power law χ_0^2 method, the weighted Fourier method, the sinusoidal fit method, and the DNN machine learning method are shown in Fig. 3. Every data set contains ~ 1800 points so for small $g < 10^{-13} \text{ GeV}^{-1}$ where the power law is already a decent fit to the

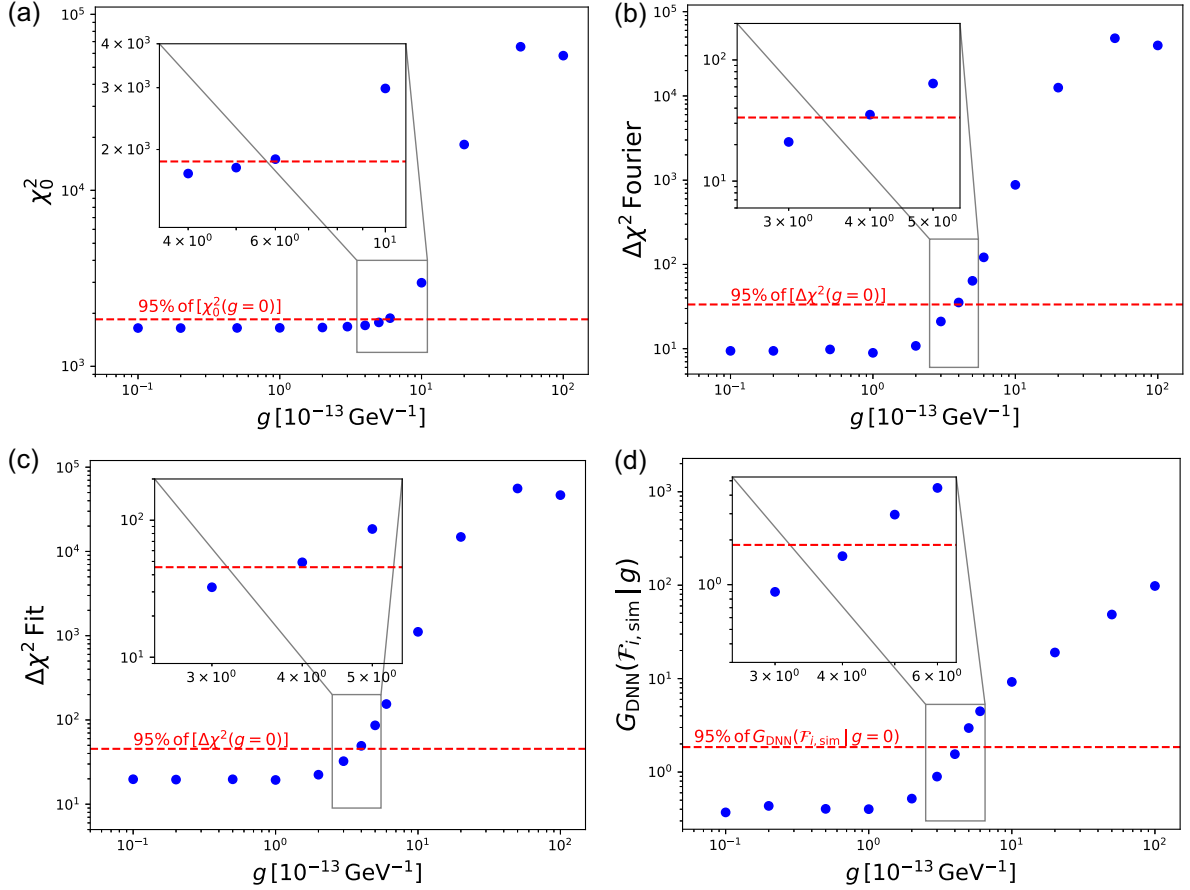


Figure 3. The four results for *Athena*. From left to right, top to bottom, these plots show (a) 5 percentile of the power-law fit $\chi_0^2(g)$ for *Athena*. There are 1747 degrees of freedom for this fit. The dashed red line is the 95th percentile of $\chi_0^2(g=0)$ and marks which values of g can be excluded: a point above it can be excluded at 95 per cent confidence. (b) 5 percentile of $\Delta\chi^2(g)$ for the Fourier method with ω^{-2} weighting (see the text) and $K = 16$ for *Athena*. (c) 5 percentile of $\Delta\chi^2(g)$ for a sinusoidal fit with $N_{\text{fit}} = 16$ for *Athena*. (d) 5th percentile of $G_{\text{DNN}}(\mathcal{F}_{i,\text{sim}} | g)$ for the DNN method for *Athena*. The dashed red line is the 95th percentile of $G_{\text{DNN}}(\mathcal{F}_{i,\text{sim}} | g=0)$ and marks which values of g can be excluded: a point above it can be excluded at 95 per cent confidence.

data, we expect a $\chi_0^2 \sim 1800$ which is what we find. Once we rebin two data points to one for the Fourier and fit method, the degrees of freedom are roughly divided by two so $\chi_0^2(\text{Fourier, Fit}) \sim 900$. χ_1^2 is calculated once the Fourier modes or the sinusoidal fit function are included. Even for very small g , we expect these methods to slightly improve the fit, hence $\chi_1^2 < \chi_0^2$, effectively fitting Poisson noise. We find $\chi_1^2 \lesssim \chi_0^2$ for very small $g < 10^{-13} \text{ GeV}^{-1}$ as expected. We examine the coverage of all four methods using our simulation data and find that while the χ_0^2 method generally gives the right coverage for all values of g the other three methods give undercoverage for small values of $g \lesssim 2 \times 10^{-13} \text{ GeV}^{-1}$ and overcoverage for $g \gtrsim 2 \times 10^{-12} \text{ GeV}^{-1}$. In the intermediate range, the method gives the right coverage. As large values of g are already excluded and small values of g are below the 95 per cent limit of the methods and hence below the discovery potential of the instrument, the intermediate range is the only one that is practically relevant.

As illustrated in these figures, the χ_0^2 method gives a bound of $g \lesssim 5.9 \times 10^{-13} \text{ GeV}^{-1}$. The weighted Fourier method and the fit method do significantly better with a bound of $g \lesssim 3.8 \times 10^{-13} \text{ GeV}^{-1}$. The DNN method also does better giving a bound of $g \lesssim 4.2 \times 10^{-13} \text{ GeV}^{-1}$. This represents an improvement by a factor of 1.5 in the coupling compared to the χ_0^2 method. As all actual

physical effects (interconversion of photons and axions) scale as $g_{a\gamma\gamma}^2$, the improvement in sensitivity is then slightly greater than a factor of 2.

5.2 XMM-Newton

The results for *XMM-Newton* are all shown in Fig. 4. In this case, the sinusoidal fit method can exclude g between $5 \times 10^{-13} \text{ GeV}^{-1}$ and $10^{-12} \text{ GeV}^{-1}$ and hence does slightly better than the χ_0^2 , weighted Fourier and DNN methods which can only exclude $g > 10^{-12} \text{ GeV}^{-1}$. Again, we check for coverage and find the same result as for *Athena*: the relevant intermediate range of g has the right coverage while there is undercoverage for small values and overcoverage for large values.

6 DISCUSSION

X-ray observations of AGNs offer one of the most sensitive probes of axion-like particles in the low-mass regime and it is therefore important to ensure that maximal statistical sensitivity is obtained in searches for them. In this paper, we have developed analysis methods aimed at exploiting the characteristic quasi-sinusoidal modulations that are induced by ALPs. Based on simulated data,

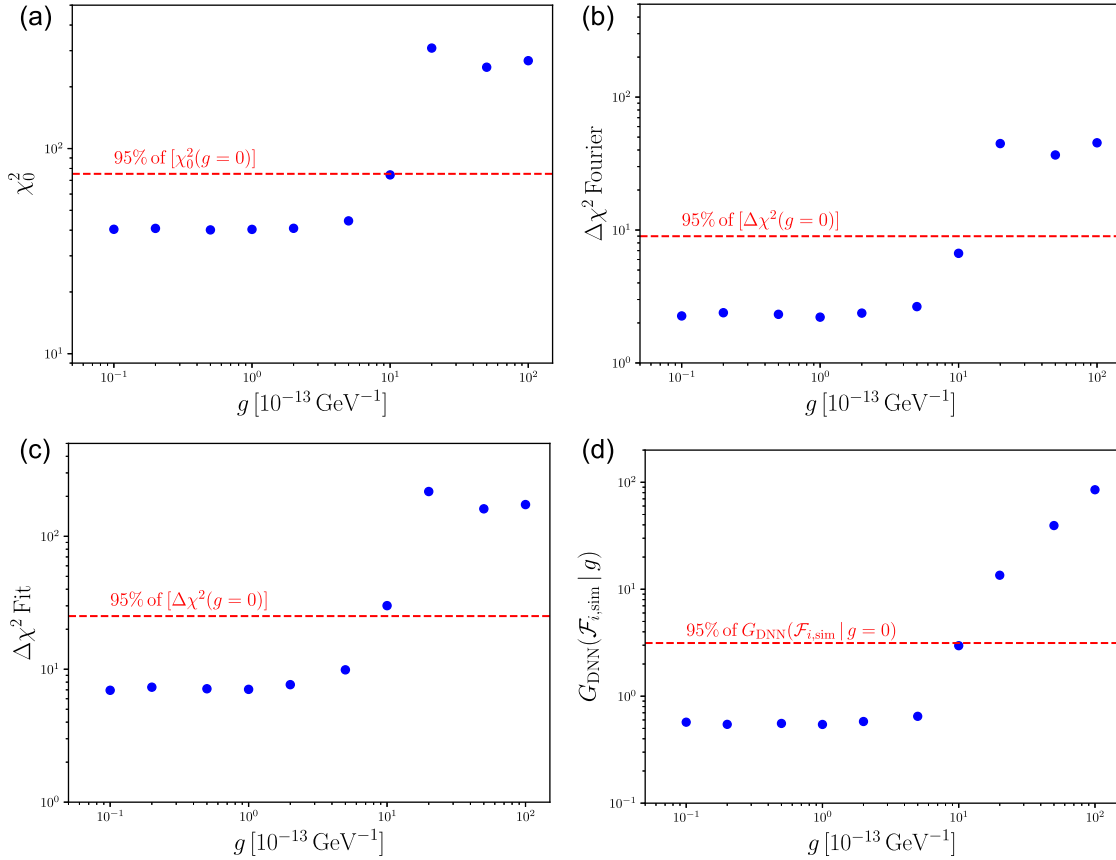


Figure 4. The four results for *XMM*. From left to right, top to bottom, these plots show (a) 5th percentile of the power-law fit $\chi_0^2(g)$ for *XMM*. There are 56 degrees of freedom for this fit. The dashed red line is the 95th percentile of $\chi_0^2(g=0)$ and marks which values of g can be excluded at 95 percent confidence. (b) 5th percentile of $\Delta\chi^2(g)$ for the Fourier method with ω^{-2} weighting (see the text) and $K=8$ for *XMM*. (c) 5th percentile of $\Delta\chi^2(g)$ for a sinusoidal fit with $N_{\text{fit}}=8$ for *XMM*. (d) 5th percentile of $G_{\text{DNN}}(\mathcal{F}_{i,\text{sim}} | g)$ for the DNN method for *XMM*. The dashed red line is the 95th percentile of $G_{\text{DNN}}(\mathcal{F}_{i,\text{sim}} | g=0)$ and marks which values of g can be excluded at 95 percent confidence.

these methods appear to improve sensitivity to photon–axion conversion by a factor of 2 (in conversion probability).

ACKNOWLEDGEMENTS

We thank Cliff Burgess, Francesca Day, Nicholas Jennings, Sven Krippendorf, David Marsh for discussions. This work was funded in part by ERC Starting Grant 307605.

REFERENCES

- Abbott L. F., Sikivie P., 1983, *Phys. Lett.*, B120, 133
Agrawal P., Fan J., Reece M., Wang L.-T., 2018, *J. High Energy Phys.*, 02, 006
Ajello M. et al., 2016, *Phys. Rev. Lett.*, 116, 161101
Angus S., Conlon J. P., Marsh M. C. D., Powell A. J., Witkowski L. T., 2014, *J. Cosmol. Astropart. Phys.*, 1409, 026
Arias P., Cadamuro D., Goodsell M., Jaeckel J., Redondo J., Ringwald A., 2012, *J. Cosmol. Astropart. Phys.*, 1206, 013
Arvanitaki A., Dimopoulos S., Dubovsky S., Kaloper N., March-Russell J., 2010, *Phys. Rev. D*, 81, 123530
Berg M., Conlon J. P., Day F., Jennings N., Krippendorf S., Powell A. J., Rummel M., 2017, *ApJ*, 847, 101
Bonafede A. et al., 2010, *A&A*, 513, A30
Burrage C., Davis A.-C., Shaw D. J., 2009, *Phys. Rev. Lett.*, 102, 201101
Chen L., Conlon J. P., 2018, *MNRAS*, 479, 2243
Churazov E., Forman W., Jones C., Bohringer H., 2003, *ApJ*, 590, 225
Cicoli M., Goodsell M., Ringwald A., 2012, *J. High Energy Phys.*, 10, 146
Conlon J. P., 2006, *J. High Energy Phys.*, 05, 078
Conlon J. P., Marsh M. C. D., 2013, *Phys. Rev. Lett.*, 111, 151301
Conlon J. P., Marsh M. C. D., Powell A. J., 2016, *Phys. Rev. D*, 93, 123526
Conlon J. P., Day F., Jennings N., Krippendorf S., Rummel M., 2017, *J. Cosmol. Astropart. Phys.*, 1707, 005
Conlon J. P., Day F., Jennings N., Krippendorf S., Muia F., 2018, *MNRAS*, 473, 4932
Farina M., Pappadopulo D., Rospignone F., Tesi A., 2017, *J. High Energy Phys.*, 01, 095
Govoni F. et al., 2017, *A&A*, 603, A122
Horns D., Maccione L., Meyer M., Mirizzi A., Montanino D., Roncadelli M., 2012, *Phys. Rev. D*, 86, 075024
Majumdar J., Calore F., Horns D., 2018, *J. Cosmol. Astropart. Phys.*, 1804, 048
Marsh M. C. D., Russell H. R., Fabian A. C., McNamara B. P., Nulsen P., Reynolds C. S., 2017, *J. Cosmol. Astropart. Phys.*, 1712, 036
Marsh D. J. E., 2016, *Phys. Rep.*, 643, 1
Payez A., Cudell J. R., Hutsemekers D., 2012, *J. Cosmol. Astropart. Phys.*, 1207, 041
Peccei R. D., Quinn H. R., 1977, *Phys. Rev. Lett.*, 38, 1440
Powell A. J., 2015, *J. Cosmol. Astropart. Phys.*, 1509, 017
Preskill J., Wise M. B., Wilczek F., 1983, *Phys. Lett.*, B120, 127
Raffelt G., Stodolsky L., 1988, *Phys. Rev. D*, 37, 1237

- Schleeder M., Sigl G., 2016, *J. Cosmol. Astropart. Phys.*, 1601, 038
- Sikivie P., 1983, *Phys. Rev. Lett.*, 51, 1415, [Erratum: *Phys. Rev. Lett.* 52, 695 (1984)]
- Svrcek P., Witten E., 2006, *J. High Energy Phys.*, 6
- Taylor G. B., Gugliucci N. E., Fabian A. C., Sanders J. S., Gentile G., Allen S. W., 2006, *MNRAS*, 368, 1500
- Vacca V., Murgia M., Govoni F., Feretti L., Giovannini G., Perley R. A., Taylor G. B., 2012, *A&A*, 540, A38
- Weinberg S., 1978, *Phys. Rev. Lett.*, 40, 223
- Wilczek F., 1978, *Phys. Rev. Lett.*, 40, 279
- Wouters D., Brun P., 2012, *Phys. Rev. D*, 86, 043005
- Wouters D., Brun P., 2013, *ApJ*, 772, 44

SUPPORTING INFORMATION

Supplementary data are available at https://github.com/mrummphys/axion_statistics online.

Please note: Oxford University Press is not responsible for the content or functionality of any supporting materials supplied by the authors. Any queries (other than missing material) should be directed to the corresponding author for the article.

This paper has been typeset from a \LaTeX file prepared by the author.