



Astrophysical Limits on Very Light Axion-like Particles from *Chandra* Grating Spectroscopy of NGC 1275

Christopher S. Reynolds¹ , M. C. David Marsh² , Helen R. Russell^{1,3}, Andrew C. Fabian¹ , Robyn Smith⁴ ,
Francesco Tombesi^{4,5,6,7} , and Sylvain Veilleux^{1,4,8}

¹ Institute of Astronomy, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK; csr12@ast.cam.ac.uk

² The Oskar Klein Centre, Department of Physics, Stockholm University, Stockholm SE-106 91, Sweden

³ School of Physics and Astronomy, University of Nottingham, Nottingham NG7 2RD, UK

⁴ Dept. of Astronomy, University of Maryland, College Park, MD 20742, USA

⁵ Dept/of Physics, University of Rome “Tor Vergata,” Via della Ricerca Scientifica 1, I-00133 Rome, Italy

⁶ NASA/Goddard Space Flight Center, Code 662, Greenbelt, MD 20771, USA

⁷ INAF Astronomical Observatory of Rome, Via Frascati 33, I-00078 Monteporzio Catone (Rome), Italy

⁸ Kavli Institute for Cosmology, University of Cambridge, Madingley Road, Cambridge CB3 0HA, UK

Received 2019 July 5; revised 2019 December 19; accepted 2020 January 8; published 2020 February 12

Abstract

Axions/axion-like particles (ALPs) are a well-motivated extension of the Standard Model and are generic within String Theory. The X-ray transparency of the intracluster medium (ICM) in galaxy clusters is a powerful probe of light ALPs (with mass $< 10^{-11}$ eV); as X-ray photons from an embedded or background source propagate through the magnetized ICM, they may undergo energy-dependent quantum mechanical conversion into ALPs (and vice versa), imprinting distortions on the X-ray spectrum. We present *Chandra* data for the active galactic nucleus NGC 1275 at the center of the Perseus cluster. Employing a 490 ks High Energy Transmission Gratings exposure, we obtain a high-quality 1–9 keV spectrum free from photon pileup and ICM contamination. Apart from iron-band features, the spectrum is described by a power-law continuum, with any spectral distortions at the $< 3\%$ level. We compute photon survival probabilities as a function of ALP mass m_a and ALP-photon coupling constant $g_{a\gamma}$ for an ensemble of ICM magnetic field models, and then use the NGC 1275 spectrum to constrain the $(m_a, g_{a\gamma})$ -plane. Marginalizing over magnetic field realizations, the 99.7% credible region limits the ALP-photon coupling to $g_{a\gamma} < 6 - 8 \times 10^{-13} \text{ GeV}^{-1}$ (depending upon magnetic field model) for masses $m_a < 1 \times 10^{-12}$ eV. These are the most stringent limit to date on $g_{a\gamma}$ for these light ALPs, and have already reached the sensitivity limits of next-generation helioscopes and light-shining-through-wall experiments. We highlight the potential of these studies with the next-generation X-ray observatories *Athena* and *Lynx*, but note the critical importance of advances in relative calibration of these future X-ray spectrometers.

Unified Astronomy Thesaurus concepts: Particle astrophysics (96); Intracluster medium (858); X-ray active galactic nuclei (2035); Extragalactic magnetic fields (507); Dark matter (353)

1. Introduction

Astrophysical observations have great potential to uncover new physics beyond the Standard Model (SM). Indeed, the clearest experimental indications that new physics must be manifest at low energies is the astrophysical observation that SM particles and fields only account for 4% of the energy density of our universe (Planck Collaboration et al. 2014).

Especially interesting, and the focus of this paper, are axion-like particles (Graham et al. 2015; Irastorza & Redondo 2018). The axion is a consequence of a well-motivated extension to the SM, namely the Peccei–Quinn (PQ) mechanism that protects the strong interaction from CP (i.e., time-reversal) violating effects (Peccei & Quinn 1977; Weinberg 1978; Wilczek 1978). This quantum chromodynamics axion couples to two photons with a coupling strength $g_{a\gamma}$ that is proportional to their mass (and inversely proportional to the energy scale of PQ symmetry breaking). Furthermore, many extensions of the SM, including in particular string theory, commonly feature very light axion-like particles (ALPs) that do not couple to the strong interactions, but can interact with photons with a strength that is independent of their masses. The coupling constant $g_{a\gamma}$ and ALP mass m_a can, in a natural way, be small compared with other SM mass scales. Axions and ALPs may be produced in the early universe via non-thermal mechanisms,

vacuum realignment and the decay of topological defects, and hence would be produced with very small velocity dispersion. Thus, despite being very light, they can still be a viable candidate for cold dark matter (Abbott & Sikivie 1983; Dine & Fischler 1983; Preskill et al. 1983).

There is a rich literature on searches for and constraints on axions/ALPs using astrophysical and cosmological observations; for an up-to-date review see Tanabashi et al. (2018). The historically most important limit on sufficiently light ALPs comes from SN1987A (Brockway et al. 1996; Grifols et al. 1996; Payez et al. 2015). Other searches include those based on the structure and luminosity of stars (Vysotskii et al. 1978), the polarization of the cosmic microwave background (Tiwari 2012; Mukherjee et al. 2019), the timing of radio pulsars and fast radio bursts (Caputo et al. 2019), and the possible discrepancies between the excess of Cosmic IR background radiation at $1 \mu\text{m}$ and the TeV opacity of the universe (Kohri & Kodama 2017).

The transparency of astrophysical systems is a particularly simple and effective way to search for the effects of ALPs; as astrophysical photons traverse through cosmic magnetic fields, they are susceptible to conversion to an ALP via the two-

photon interaction described by the Lagrangian term

$$\mathcal{L}_a = -g_{a\gamma} a \mathbf{E} \cdot \mathbf{B}, \quad (1)$$

where a is the ALP field, and \mathbf{E} and \mathbf{B} are the electric and magnetic fields. The transparency of astrophysical systems can be used to set upper limits on $g_{a\gamma}$ and may be a route to the eventual detection of ALPs.

X-ray observations of active galactic nuclei (AGN) in rich clusters of galaxies are particularly suited to ALP searches. Faraday rotation measure (RM) studies demonstrate that the hot intracluster medium (ICM) in such clusters is magnetized, with a ratio of thermal-to-magnetic pressure of $\beta \sim 100$ (Taylor et al. 2006). Furthermore, we expect significant regions of this field to be coherent on scales 1–10 kpc (Vacca et al. 2012). The result is that, if ALPs exist with a sufficiently high coupling to photons, clusters will be efficient converters of X-ray photons into ALPs. For typical parameters relevant to the ICM, the conversion probability will be energy dependent thereby imprinting distortions into the observed spectrum of any embedded (or background) object, with a precise form that depends upon the magnetic field structure as well as the ALP properties. We note that these transparency studies probe the physics of ALPs independently of whether they actually constitute a significant component of the non-baryonic dark matter.

Wouters & Brun (2013) used *Chandra* imaging spectroscopy of the central AGN Hydra-A in the Hydra cluster of galaxies to set an upper limit on spectral distortions from ALPs, showing that any ALP with mass $m_a < 7 \times 10^{-12}$ eV must have $g_{a\gamma} < 8.3 \times 10^{-12}$ GeV $^{-1}$ (95% confidence level). Subsequently, Berg et al. (2017) examined the AGN NGC 1275 at the center of the Perseus cluster, the target of our current study, using *Chandra* and *XMM-Newton* imaging spectroscopy. NGC 1275 is almost 100 \times brighter than Hydra-A in the 2–10 keV band and, furthermore, does not have the heavy intrinsic absorption of Hydra-A. While this dramatically reduces the statistical errors on the spectrum of the NGC 1275, the high source flux creates systematic issues. For the large body of *Chandra* imaging data (almost 1Ms) for which NGC 1275 is close to the optimal aim-point, the subarcsecond focusing leads to very severe photon pileup; thus there are very strong spectral distortions that are entirely instrumental in origin. For this reason, Berg et al. (2017) focus their attention on the smaller quantities of *Chandra* data for which NGC 1275 is off-axis and on data from *XMM-Newton*. Since the poorer focus for these data sets reduces but does not eliminate the effect of pileup, they employ the pileup model of Davis (2001) and conclude that any spectral distortions must be below the 10% (once smoothed to the spectral resolution of the CCD detectors in *Chandra* and *XMM-Newton*). The resulting ALP constraint is $g_{a\gamma} < 4 \times 10^{-12}$ GeV $^{-1}$ for massless ALPs. Chen & Conlon (2018) extended this analysis to massive ALPs, and verified that this limit applies to $m_a \lesssim 10^{-12}$ eV.⁹

To circumvent the photon pileup issue, Marsh et al. (2017) used 370 ks of short (0.4 s) frame time *Chandra* imaging spectroscopy of the core of M87 in the Virgo cluster. Combined with the fact that this AGN is almost an order of magnitude fainter than NGC1275, the short frame time exposures led to a high-quality spectrum with negligible

pileup, and ALP constraints of $g_{a\gamma} < 2.6 \times 10^{-12}$ GeV $^{-1}$ for $m_a < 10^{-12}$ eV (95% confidence level).

In this paper, we present new *Chandra* observations of NGC 1275 that set the tightest limits to date on light ALP conversion. By employing the High Energy Transmission Gratings (HETG) and investing almost 500 ks of on-source exposure, we obtain a high-quality separation of the AGN emission from the ICM with no discernible photon pileup. Outside of the astrophysically rich iron band (6–7 keV), we find that the resulting 1–9 keV spectrum of the AGN is well described by a power-law continuum (modified only by the effects of absorption by cold gas in our Galaxy) with any remaining spectral distortions below the 3% level.

We proceed to marginalize over a set of representative realizations of the magnetic field to determine the posterior distribution on the $(m_a, g_{a\gamma})$ -plane. The resulting 99.7% credible region limits the ALP-photon coupling to $g_{a\gamma} < 6 - 8 \times 10^{-13}$ GeV $^{-1}$ (depending upon the magnetic field model) for most masses $m_a < 1 \times 10^{-12}$ eV. This is the most stringent limit to date on the ALP-photon coupling.

For one of our magnetic field models, the posterior peaks at a non-vanishing value of $g_{a\gamma}$, and the 95% credible region pick out a preferred non-zero value of the couple constant, $g_{a\gamma} \approx 4 - 8 \times 10^{-13}$ GeV $^{-1}$, for a range of masses $m_a < 1 \times 10^{-12}$. We attribute this “detection” to residual, low-level, errors in the *Chandra*/HETG calibration.

This paper is organized as follows. Section 2 presents the new observations and describes some subtleties encountered during the data reduction. After discussing the modeling of the ICM magnetic field and the associated ALP distortions in Section 3, the new constraints on ALPs are given in Section 4. We put these results into context and draw our conclusions in Section 5. When necessary, we assume a standard *Planck* cosmology with $H_0 = 68$ km s $^{-1}$ Mpc $^{-1}$ (Planck Collaboration et al. 2014), which, at a redshift of $z = 0.0173$ (Hitomi Collaboration et al. 2018), places NGC 1275 at a distance of 76 Mpc.

2. Observations, Data Reduction and Initial Spectral Fitting

Chandra observed NGC 1275 in 15 separate segments (ObsIDs) between 2017 October 24 and 2017 December 5 using the HETG readout on the Advanced CCD Imaging Spectrometer (ACIS) S array. The total on-source exposure time was 490 ks. To ameliorate the risk of modest photon pileup in the event that the source was brighter than expected, we turned off the two outlying ACIS-S chips and used 1/2 sub-arrays on the remaining four chips, resulting in a reduction of the frame readout time to 2.4 s with no loss of observing efficiency. While loss of the two outlying chips in principle affects our ability to observe the softest regions of the spectrum, the contaminant that has built up on the ACIS optical blocking filters unavoidably removes those soft photons anyway.

Figure 1 shows the image of the ACIS-S array for one of our ObsIDs (20449). This is dominated by the zeroth order image of the core of the Perseus cluster with its famous cavity system (Fabian et al. 2000). The two two-sided dispersed spectra of the bright central AGN emission, one from the High Energy Grating (HEG) and one from the Medium Energy Grating (MEG), are clearly visible. Although the gratings are slitless and so the cluster light is also dispersed, it is clear that the AGN

⁹ For ALP limits from *Chandra* observations of seven sources (with more poorly constrained magnetic fields), see also Conlon et al. (2017).




Figure 1. Full-band image of the ACIS-S array for ObsID 20449 (exposure 45 ks). For purposes of display only, the raw pixel data have been binned by 4×4 , and the color bar shows total photon count per (new) binned pixel for this ObsID. The extraction region for the dispersion spectrum is shown in white.

light is distinct and can be well isolated from the bulk of the ICM emission. Order-sorting, whereby only photons with CCD-detected energies compatible with their spatial position along the dispersion spectrum are accepted, enables further isolation of the AGN spectrum.

Much but not all of our data reduction is standard. All data were reprocessed with CIAO-4.10 and CALDBv4.8.1. The extraction of the AGN spectra for each ObsIDs then follows the standard CXC science threads¹⁰ with two exceptions. First, we halve the width of the extraction regions (`width_factor_hetg` = 18) in order to reduce MEG/HEG overlap at the center of the array and hence access the higher energy band in the HEG. Second, the standard algorithm for automatically locating the zeroth order image and hence setting the energy scale of the spectrum failed for most of the ObsIDs, presumably due to the surrounding high-surface brightness and structured ICM. Instead, we force the zeroth order point to be at the known coordinates of NGC1275, and visually confirm for each ObsID that this correctly locates the zeroth order image of the point-like AGN (i.e., that astrometry errors are within one pixel or $0''.5$). The result is four spectra and associated background spectra, response matrices and effective area files for each ObsID, namely the +1 and -1 order spectra for each of the MEG and HEG. In a final step, we combine the spectra to produce a single HEG and a single MEG spectrum (with associated background spectra, response matrices, and effective area files), summing the $+/-1$ orders from all ObsIDs.

All spectral fitting presented in this paper uses the 1 – 7keV band for the MEG, the 1.5 – 9keV band for the HEG, and employs Cash (C) statistic minimization to fit the unbinned spectrum, modified to allow for the subtraction of a background spectrum.¹¹ Fitting is performed with the XSPECv12.10.1 code (Arnaud 1996).

An initial fit of the spectra with a power-law continuum modified by Galactic absorption ($N_H = 1.32 \times 10^{21} \text{ cm}^{-2}$; Kalberla et al. 2005) finds 10%–15% excesses in three energy




Figure 2. Illustration of the background under subtraction issue when using the default-normalized background spectra. Shown here is the best-fitting power-law model to the combined HEG spectrum (blue) that has been background subtracted using the background spectrum obtained from the standard pipeline (green). The significant deviations from the power-law at ~ 2.5 keV clearly mirror a feature in the background spectrum.

bands; below 1.3 keV (MEG only), 2.2 – 2.5keV (HEG and MEG), and 6 – 7keV (principally in the HEG). The 6 – 7keV structure corresponds to the well-established iron fluorescent line from cold gas in the vicinity of the AGN (Hitomi Collaboration et al. 2018) and will be the subject of another paper (C. S. Reynolds et al. 2020, in preparation). As illustrated in Figure 2 for the HEG, the other residuals closely mirror structure in the background spectra (which is actually dominated by the core ICM emission) and suggests that the standard grating extraction algorithms have underestimated the background normalization. Broadly, this is not a surprise. The background spectrum is determined from strips that flank the

¹⁰ http://cxc.cfa.harvard.edu/ciao/threads/spectra_hetgacis/

¹¹ See <https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/node304.html#AppendixStatistics>.

Table 1

Power-law Fits (Modified by Galactic Absorption) to the Combined First-order HEG and MEG Spectra

Parameter	HEG Value	MEG Value
Galactic absorption, N_H	$1.32 \times 10^{21} \text{ cm}^{-2}$ (fixed)	$1.32 \times 10^{21} \text{ cm}^{-2}$ (fixed)
Photon index, Γ	1.852 ± 0.017	1.882 ± 0.011
Normalization, A	$(7.74 \pm 0.14) \times 10^{-3}$	$(8.34 \pm 0.08) \times 10^{-3}$
Flux (1–9 keV)	$3.06 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$	$3.18 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$
C/DOF	2835/2750	2045/2122

source extraction region with algorithms that are designed and optimized for a spatially uniform background around a point source. In our case, the ICM emission that forms our background is centrally concentrated around the AGN, so we expect that the spectrum extracted from the background/flanking regions will be normalized too low. It is not possible, however, to simply estimate the size of this effect from the surface brightness profile—the dispersion of the ICM emission in this slitless grating system together with the order-sorting algorithm (where the intrinsic energy resolution of the ACIS is used to reject all photons that definitely lie outside of the expected map of dispersed position to energy) makes the background normalization a non-trivial function of the ICM spatial and spectral structure.

Instead, we follow an empirical approach. By scanning through a range of possible renormalization values, we find that the C -statistic of the power-law fit is minimized if the HEG and MEG backgrounds are scaled up by factors of 2.32 and 1.92, respectively. This reduces the deviations from the power-law to below the 3%–5% level. We validate this process with simulations. We use the MARX package to simulate an HETG observation of a Perseus-like cluster with the following components (i) a point-like AGN with a power-law spectrum (photon index $\Gamma = 1.9$), (ii) an ICM core described by a β -profile with core radius $2'$ and an optically thin thermal plasma spectrum (temperature $kT = 4$ keV), and (iii) a model for the ICM cavity/shell structure consisting of two annular rings offset so that they just overlap at the AGN with inner radius $0.6'$, outer radius $1'$, and an optically thin thermal plasma spectrum (temperature $kT = 2$ keV). We then pass the simulated events files through the standard extraction pipeline used for the real data. This confirms that, when attempting to analyze the spectrum of the AGN, the spatial structure of the cluster leads to a systematic underestimate in the normalization of the background spectrum by approximately a factor of two.

With these adjusted background normalizations, the AGN photons comprise 80% and 82% of the HEG and MEG spectra, respectively. A joint HEG/MEG fit gives a power-law index of $\Gamma = 1.890 \pm 0.009$ and normalization at 1 keV of $A = (8.28 \pm 0.06) \times 10^{-3} \text{ phs}^{-1} \text{ keV}^{-1} \text{ cm}^{-2}$, with $C = 4956$ for 4874 degrees of freedom. This corresponds to a 1–9 keV band flux and luminosity of $F_{1-9\text{keV}} = 3.8 \times 10^{-11} \text{ erg cm}^{-2} \text{ s}^{-1}$ and $L_{1-9\text{keV}} = 2.2 \times 10^{43} \text{ erg s}^{-1}$, respectively. If we allow the power-law parameters to float freely between the MEG and HEG spectra, the fit improves significantly ($\Delta C = -76$) and we find marginally significant slope and highly significant normalization differences between the two gratings (Table 1). While highly significant, the absolute flux difference is below 10% and hence within the realm of what can reasonably be

attributed to an instrumental cross-calibration uncertainty. Hence, for all subsequent fitting in this paper, we will permit normalization and slope offsets between the HEG and MEG spectra.

Figure 3 shows this free fit of the absorbed power-law to the HEG and MEG spectra, heavily binned for plotting purposes (but unbinned for fitting purposes). The residuals about the power-law are not entirely random, with a slight broad dip around 1.2 keV, a very subtle broad hump between 3 and 4 keV, and an obvious feature in the iron band (6.4–7 keV). Still, outside of the iron band, the remaining residuals are less than 5% and mostly less than 3% and so are entirely consistent with the expected level of residual effective area calibration errors (see Figure 7 of Marshall 2012). Having probably reached the level dominated by systematic calibration uncertainties, this is the highest quality 1–9 keV band spectrum of this AGN obtained to date and, as we shall see, permits the most sensitive search yet for light ALPs.

3. Modeling the ALP signatures

Our modeling of ALP spectral distortions follows that of Marsh et al. (2017). We solve the linearized Schrödinger-like equation that describes the quantum mechanical oscillations between photons and ALPs as they traverse through the cluster magnetic field toward the observer. The survival probability of photons emitted from the nucleus that exit the cluster as photons, rather than ALPs, depends on the plasma density and the cluster magnetic field. For the electron density, we use the simple analytic approximation derived from *XMM-Newton* observations of the Perseus cluster observations by Churazov et al. (2003):

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

For the magnetic field, we use two stochastic models.

Model-A: We first consider a slight modification to the magnetic fields used in Berg et al. (2017). This model is motivated by VLBA observations of the nucleus of NGC 1275 by Taylor et al. (2006), which found Faraday RM of the order of $6500\text{--}7500 \text{ rad m}^{-2}$ across the tip of the southern active jet of 3C 84, the radio source associated with NGC 1275. Interpreting these RMs as arising from a narrow (~ 2 kpc) Faraday screen in the high-density ($n_e \approx 0.3 \text{ cm}^{-3}$) central region, Taylor et al. (2006) estimated the central magnetic field strength as $B_0 \approx 25 \mu\text{G}$. Taking this value for the central magnetic field, Berg et al. (2017) modeled the magnetic field along the line of sight as a series of domains. Within each domain, the field is taken to be randomly oriented with a constant magnitude given by $B(r_c) = B_0[n_e(r_c)/n_e(0)]^{0.7}$ (where r_c is the radial coordinate at the center of the domain). Motivated by the detailed RM study of A2199 by Vacca et al. (2012), the size of each domain L is drawn from a random distribution with probability density proportional to $L^{-1.2}$ (corresponding to the index of the 3D RM power-spectrum found by Vacca et al. 2012), between 3.5 and 10 kpc (motivated by scaling length scales in A2199 to Perseus).

Our Model-A modifies this description in two ways. First, we note that the large value of B_0 combined with the moderate central plasma density of Equation (2) leads to non-negligible ALP-photon oscillations from the inner-most region of the

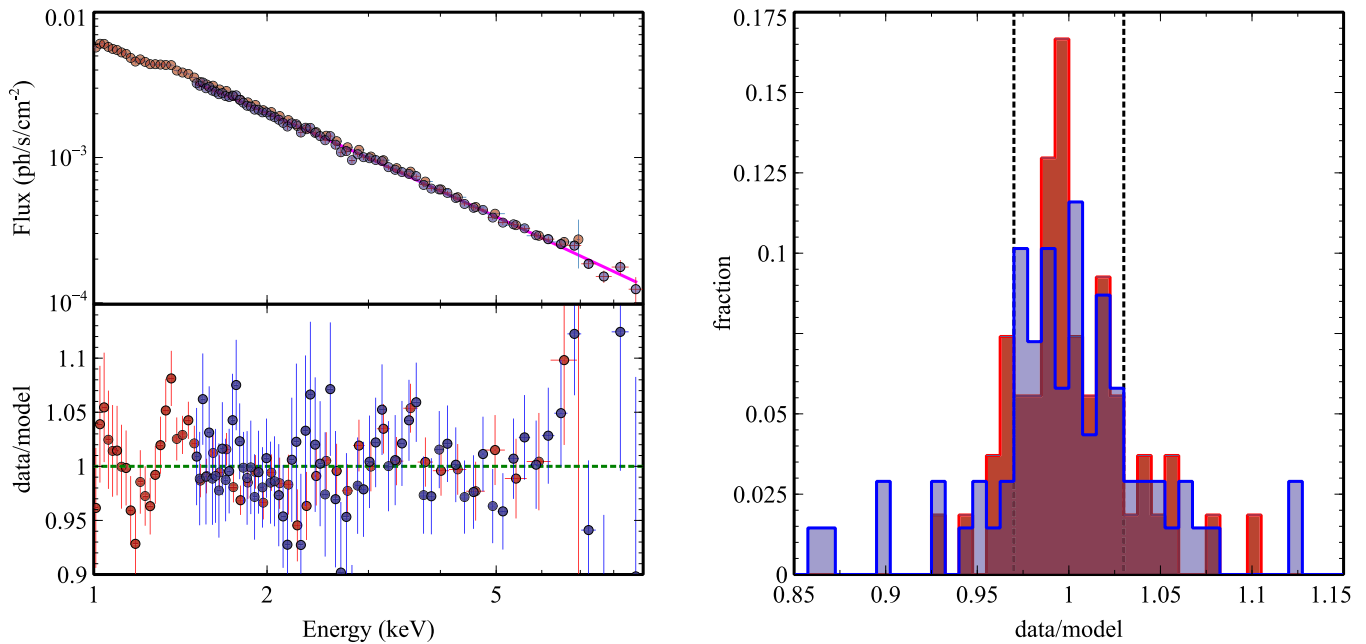


Figure 3. Left panel: best-fitting power-law model to the combined HEG (blue) and MEG (red) first-order spectra (top) with corresponding ratios of the data to the best-fitting model. The data have been heavily binned for the purposes of plotting, but spectral fitting is performed on unbinned data. Right panel: distribution of the ratios of the data to the best-fitting model for the HEG (blue) and MEG (red). The vertical dotted lines denote the $\pm 3\%$ levels.

cluster. However, the simple analytic model of Churazov et al. (2003) does not apply to small radii, $r < 10$ kpc, where it underestimates the electron density, and where the spherically symmetric approximation is not justified. Applying Equation (2) to this region leads to an overestimate of the ALP-photon conversion probability. In our work, we conservatively exclude the central region, and simulate the ALP-photon oscillations from 10 kpc out to the virial radius, $R_{\text{vir}} = 1.8$ Mpc. Second, we note that the bulk of the ICM also acts as a Faraday screen and sources RMs in addition to those arising from the central region. Since this model attributes the observed RMs to the Faraday screen close to the center of the cluster, we consistently select only those magnetic field configurations in which the cluster contribution is subleading: $\text{RM}_{\text{cluster}} \leq 2000 \text{ rad m}^{-2}$. However, we have found that this restriction has no statistically significant impact on the typical conversion probabilities.

Model-B: We furthermore consider a model in which the ratio of the thermal-to-magnetic pressure is fixed to $\beta = 100$ throughout the cluster. We use the Perseus pressure profile of Fabian et al. (2006) to derive a magnetic field strength of $B_{25} \approx 7.5 \mu\text{G}$ at $r = 25$ kpc. Approximating the cluster as isothermal, the magnetic field decreases with radius as $\sim \sqrt{n_e(r)}$, where we again use Equation (2) for the electron density. With a central field that is suppressed with respect to Model-A, the ALP-photon conversion from the central region is negligible and so we can use this model from $r = 0$ to the virial radius. The coherence lengths of the magnetic field can be expected to grow with distance from the center. We model this effect by drawing the coherence lengths randomly from $(1+r/50 \text{ kpc}) \times 3.5 \text{ kpc}$ to $(1+r/50 \text{ kpc}) \times 10 \text{ kpc}$, with a power-law fall-off as $\sim L^{-1.2}$. This model produces Faraday RMs of the same order as those observed by Taylor et al. (2006), with the cluster as the Faraday screen. While this choice of domain-size structure is somewhat arbitrary, it is

designed to allow comparison and connection with the results of Model-A as well as previous studies.

For each of our two field models, we generate 500 RM-acceptable magnetic field configurations and solve the Schrödinger-like equation in order to calculate photon survival probabilities across a grid of m_a and $g_{a\gamma}$. Our models sample the $(m_a, g_{a\gamma})$ -plane, spanning the range $\log_{10}(m_a/\text{eV}) \in [-13.6, -11.1]$ and $\log_{10}(g_{a\gamma}/\text{GeV}^{-1}) \in [-13, -10.7]$. The result is a library of approximately 260,000 energy-dependent photon survival probability curves for each of our two magnetic field models (Model-A and Model-B), $p_\gamma(E; m_a, g_{a\gamma}, i_{A/B})$, where $i_{A/B}$ indexes the 500 RM-acceptable magnetic field realizations for that given magnetic field model. Some representative photon survival probability curves, and their functional dependence on m_a and $g_{a\gamma}$, are shown in Figure 4.

In Figure 5, we compare photon survival probability curves from two representative realizations for each of our two magnetic field models at an illustrative point in the ALP parameter space, $m_a = 10^{-12} \text{ eV}$, $g_{a\gamma} = 10^{-12} \text{ GeV}^{-1}$. Below 4 keV, the two field models give spectral distortions of similar magnitude, although the distortions produced by Model-A are typically narrower. Above 4 keV, both field models show a transition to more periodic energy structures, with Model-A showing a marked increase in the magnitude of the distortions.

The domain models that we consider are simple enough to make the extensive calculations required below feasible, but complex enough to agree qualitatively with several of the features of more elaborate stochastic models in which the cluster magnetic field is taken to be a divergence-free function derived from Gaussian random fields (as in Angus et al. 2014, see in particular Section 5.2.1). The discontinuity of the magnetic field at the boundaries of the domain does not, of course, lead to discontinuities in the conversion probability as a function of the radius. Neither the Gaussian random field model nor the discrete cell model correspond to magnetic fields that




Figure 4. Some example photon survival probability curves using one representative realization of Model-B for the magnetic field structure, $p_\gamma(E)$. Left panel: curves for fixed mass ($\log_{10}(m_a/\text{eV}) = -12.7$) and magnetic field configuration, but various values of the coupling constant $\log_{10}(g_{a\gamma}/\text{GeV}^{-1}) = -11.9$ (black), -12.1 (blue), -12.3 (green), -12.5 (magenta), -12.7 (red). Right panel: curves for fixed coupling constant $\log_{10}(g_{a\gamma}/\text{GeV}^{-1}) = -12.1$ and magnetic field configuration, but various ALP masses $\log_{10}(m_a/\text{eV}) = -12.3$ (black), -12.5 (blue), -12.7 (green), -12.9 (red).




Figure 5. Example photon survival probability curves for one choice of ALP parameters ($m_a = 10^{-12} \text{ eV}$, $g_{a\gamma} = 10^{-12} \text{ GeV}^{-1}$) and two representative magnetic field realizations from each of Model-A (black) and Model-B (red).

are actually realized in nature, but are simple models that capture some aspects of the underlying physical magnetic field. For photon-ALP oscillations, the relevant aspects are the (non-radial) strength of the magnetic field and its (radial) coherence length. Testing the detailed properties of photon-to-ALP conversion in more realistic turbulent magnetic fields such as those derived from MHD simulations is an interesting exercise but beyond the scope of the current paper.

4. Constraints on ALP Parameters

Equipped with the library of photon survival probability curves, we can now use our HETG spectra of NGC 1275 to determine the allowed regions of the $(m_a, g_{a\gamma})$ -plane.

In order to determine the probability of the parameters given this data, appropriately marginalized over the unknown cluster magnetic field configuration, we follow the Bayesian procedure of Marsh et al. (2017). We assume flat priors on $\ln m_a$ and $\ln g_{a\gamma}$ in the range $\log_{10}(m_a/\text{eV}) \in [-30, -11.1]$ and $\log_{10}(g_{a\gamma}/\text{GeV}^{-1}) \in [-19, -10.7]$. We will find that our results are insensitive to the minimum allowed mass, and the particular choice considered here corresponds to cosmologically large Compton wavelengths of the ALPs. The minimum allowed coupling constant corresponds to the inverse Planck mass, below which quantum gravitational corrections are expected to become important. We also assume flat priors on the randomly generated magnetic field configurations, labeled by $i_{A/B}$. Motivated by the initial fitting presented in Section 2, our baseline spectral model for NGC 1275 consists of a power-law continuum modified by the effects of Galactic absorption ($N_H = 1.32 \times 10^{21} \text{ cm}^{-2}$; Kalberla et al. 2005).

For a given magnetic field model (Model-A and Model-B), we take each of our photon survival probability curves (indexed by m_a , $g_{a\gamma}$ and the magnetic field realization $i_{A/B}$), multiply by the power-law spectrum (modified by Galactic absorption), and then fit to the unbinned HEG/MEG spectra, minimizing the C -statistic over the HEG/MEG photon indices and HEG/MEG normalizations. The lowest masses in our model library ($\log_{10}(m_a/\text{eV}) = -13.6$) yield fits that are indistinguishable from the massless case, and hence these model fits are used as proxies for the very low-mass region of parameter space. Similarly, the smallest coupling constant in our model library ($\log_{10}(g_{a\gamma}/\text{GeV}^{-1}) = -13$) are indistinguishable from the zero coupling case and hence these model fits are used as proxies for the very small coupling region of parameter space.

From the resulting set of $C(m_a, g_{a\gamma}, i_{A/B})$, we form posterior probabilities $\mathcal{P}_{A/B}(m_a, g_{a\gamma}, i) \propto \exp(-C^2/2)$ normalized such that

$$\sum_{m_a, g_{a\gamma}, i_{A/B}} \mathcal{P}_{A/B}(m_a, g_{a\gamma}, i_{A/B}) = 1. \quad (3)$$

To obtain the posterior over the ALP parameters alone (and account for the “look-elsewhere effect” associated with the unknown magnetic field), we then marginalize over the magnetic field configurations,

$$\mathcal{P}_{A/B}(m_a, g_{a\gamma}) = \sum_{i_{A/B}} \mathcal{P}_{A/B}(m_a, g_{a\gamma}, i_{A/B}). \quad (4)$$

The maximum value of $\mathcal{P}_{A/B}(m_a, g_{a\gamma})$ gives the best-fit values of the ALP parameters. Contours of equal $\mathcal{P}_{A/B}(m_a, g_{a\gamma})$ provide the boundaries of the credible regions of the ALP parameters. The $x\%$ credible region includes the points with largest $\mathcal{P}_{A/B}(m_a, g_{a\gamma})$ so that their sum accounts for $x\%$ of the total probability. For the marginalized probabilities of this work, this method presents no ambiguities.

Figure 6 shows the main result of this paper, our new constraints on the $(m_a, g_{a\gamma})$ -plane. For reference, Figure 6 also shows the previously tightest constraints on ALPs in this mass range from SN1987A (Payez et al. 2015), and the lack of spectral distortions in the X-ray spectrum of M87 (Marsh et al. 2017) and the *Fermi*/ γ -ray spectrum of NGC 1275 (Ajello et al. 2016). The posterior probability from our new analysis over the ALP parameters has a complex structure. For $m_a < 1 \times 10^{-12} \text{eV}$, we set a strong upper limit of $g_{a\gamma} < 6.3 \times 10^{-13} \text{GeV}^{-1}$ (Model-A) and $g_{a\gamma} < 7.9 \times 10^{-13} \text{GeV}^{-1}$ (Model-B) at the 99.7% level. This stems directly from the fact that the HEG/MEG spectra display no significant spectral distortions exceeding $\pm 3\%$.

At the 95% level and for Model-B, the marginalization process picks out a preferred value of $g_{a\gamma} \approx 4 - 8 \times 10^{-13} \text{GeV}^{-1}$.¹² An examination of a sample of the photon survival probability curves for the most probable models shows that they all share a modest dip at 1.2 keV and a broad hump at 3–4 keV, structure that is visually apparent in the spectra (Figure 3). At these levels, however, such spectral structures could easily result from remaining calibration errors in the HETG energy-dependence effective area (Marshall 2012) and thus we cannot claim even a tentative detection of ALPs on the basis of the enclosed 95% confidence contour.

We end this section with a brief discussion of the Bayesian evidence for ALP models compared with a no-ALP hypothesis. For magnetic field Model-A, we obtain a Bayes factor of $K = 1.5$, “barely worth mentioning” on the Jeffrey’s scale. For Model-B, we obtain $K = 22.8$ which constitutes “strong evidence” for ALPs on the Jeffrey’s scale. This is a restatement of the fact that the data does, indeed, possess structure that can be fitted by ALP distortions.

¹² We note that the strength of the preference for a non-vanishing value of the coupling $g_{a\gamma}$ is prior-dependent: should we have restricted the prior range to only the sampled region of $\log_{10}(g_{a\gamma}/\text{GeV}^{-1}) \in [-13, -10.7]$, the 99.7% confidence regions obtained from either of the magnetic field models would exclude a vanishing ALP-photon coupling. In contrast, the upper limit on $g_{a\gamma}$ is highly prior-independent, and the large prior range considered in this paper is therefore conservative.

5. Discussion and Conclusions

A deep exposure of NGC 1275 at the center of the Perseus cluster of galaxies with the *Chandra*/HETG has allowed us to obtain the highest quality spectrum of this AGN free from strong contamination by the surrounding cluster emission and free from the effects of photon pileup. Apart from subtle structure in the iron band (6–7 keV), we find that the 1–9 keV spectrum of NGC 1275 is accurately described by a power-law continuum form modified by the effects of modest Galactic absorption; deviations from the power-law are at the $\pm 3\%$ level. Taking this to be the intrinsic spectrum of the AGN, we proceed to use these data to constrain models for ALP-photon oscillations in the magnetic field of the Perseus cluster. We have obtained the most stringent limit yet on the ALP-photon coupling constant of very light ALPs with masses $m_a < 1 \times 10^{-12} \text{eV}$:

$$g_{a\gamma} < 7.9 \times 10^{-13} \text{GeV}^{-1}, \quad (99.7\% \text{ confidence}), \quad (5)$$

for magnetic field Model-B, with an even more stringent limit of $g_{a\gamma} < 6.3 \times 10^{-13} \text{GeV}^{-1}$ for Model-A. Even at this much higher level of confidence, our limits are 3–4 \times stronger than those obtained by Marsh et al. (2017) for M87/Virgo, and over 5 \times stronger than the those found by Berg et al. (2017) with non-dispersive spectroscopy of NGC 1275/Perseus.

For one of our two magnetic field models, the ALP conversion models pick out some remaining structure in the HEG/MEG spectra and hence, at the 95% level (but not the 99.7% level) lock onto a preferred non-zero coupling constant of $g_{a\gamma} \approx 4 - 8 \times 10^{-13} \text{GeV}^{-1}$. However, we acknowledge that $\pm 3\%$ is very plausibly at the level of residual calibration uncertainties in the HETG energy-dependent effective area and hence this result must be viewed with extreme caution.

For these very light ALPs ($m_a < 1 \times 10^{-11} \text{eV}$) our astrophysical limits are now tighter than those obtained by the fact that SN1987A did not generate a gamma-ray burst via the ALP-mediated escape of gamma-rays from the collapsing stellar core (Payez et al. 2015). Our limits have also exceeded the projected sensitivity of the next-generation helioscope, the International Axion Observatory (IAXO; Armengaud et al. 2014), as well as the next-generation light-shining-through-walls experiments such as ALPS-II (Bähre et al. 2013).

While the focus of this paper are the constraints on ALPs, we note that our data allow a strong test of the Conlon et al. (2017) fluorescent dark matter model. These authors note that there are hints of an absorption feature at 3.5 keV in the *XMM-Newton*/EPIC-pn of NGC 1275. In an attempt to reconcile the claim of a 3.5 keV dark matter emission line in the *XMM-Newton* spectrum of the Perseus ICM (Boyarsky et al. 2014; Bulbul et al. 2014) with the non-detection of any such feature in the *Hitomi* spectrum of the the system (Hitomi Collaboration et al. 2017), Conlon et al. (2017) proceed to formulate a two-level dark matter model in which a dark matter absorption line in the AGN spectrum offsets the dark matter emission line from the cluster in system-integrated spectra such as that produced by *Hitomi*. This hypothesis requires that the dark matter imprints an absorption line close to 3.5 keV with an equivalent width of 15 eV. Our HETG current spectrum sets an upper limit of 4eV with 99% confidence, allowing us to rule out this version of the fluorescent dark matter model.

Following on from Wouters & Brun (2013), Berg et al. (2017) and Marsh et al. (2017), our work is just the latest to




Figure 6. Constraints on the $(m_a, g_{a\gamma})$ -plane from this study for magnetic field Model-A (top) and Model-B (bottom). We show 99.7% confidence limit (heavy black line bounding the excluded, light blue shaded region), the 95% confidence region (allowed region with yellow shading), and best-fit parameter values (red dots). Also shown are the previous most stringent 95% limit from the Marsh et al. (2017) study of M87 (blue line), the limit from the absence of a γ -ray burst associated with SN1987 (Payez et al. 2015) (black line), the limit from FERMI (gray line bounding blue excluded region at the 95% confidence level) as well as the projected sensitivity from the next-generation helioscope IAXO. For other, weaker limits obtained using X-ray astronomy in the same region of parameter space, see also Wouters & Brun (2013), Berg et al. (2017), Conlon et al. (2017).

highlight the power of X-ray transparency studies of galaxy clusters to probe the physics of ALPs. Equipped with high-resolution, non-dispersion, high count rate microcalorimeter arrays at the focus of X-ray telescopes with $A > 1\text{m}^2$, the next generation of X-ray flagship missions *Athena* and *Lynx* will dramatically advance the quality of the spectrum possible for an embedded AGN such as NGC 1275.

But the current study also points to the challenges that need to be overcome if we are to extend this technique further and dig deeper into the parameter space. Even if photon pileup/deadtime and contamination by ICM emission is rendered negligible, the relative calibration of the energy-dependent effective area of the X-ray spectrometer will set a floor on the sensitivity of any studies based on spectral distortions

irrespective of the photon statistics. For example, for a representative Perseus magnetic field model, a very light ALP with $g_{a\gamma} = 3 \times 10^{-13} \text{GeV}^{-1}$ produces distortions at the 1.5% (3%) level in the $E \sim 2 - 4 \text{keV}$ ($5 - 8 \text{keV}$) band. For $g_{a\gamma} = 2 \times 10^{-13} \text{GeV}^{-1}$, the corresponding distortion is 0.75% (1.5%). Thus, to significantly further this technique for searching for ALPs, future X-ray spectrometers must achieve relative effective area calibrations of 1% or better. It is clear that in order to push the ALP constraints significantly further than we have done here (or actually obtain a robust detection of very light ALPs) will require advances in broad-band, on-orbit, relative calibration of future X-ray spectrometers.

Should future studies actually detect spectral distortions due to ALPs, the recovered ALP parameters $(m_a, g_{a\gamma})$ will clearly

depend upon the magnetic field model assumed in the analysis. As an illustration of this issue, we have created simulated 490 ks HEG/MEG spectra of NGC 1275 modified by ALPs with $\log(m_a/\text{eV}) = -12.5$ and $\log(g_{a\gamma}/\text{GeV}^{-1}) = -12$ employing one realization of magnetic field Model-A. We then ran the simulated data through our full analysis pipeline using magnetic field Model-B. The ALP signal is detected at the 99.7% level, with $g_{a\gamma}$ constrained to 0.2 dex albeit biased low compared with the true value by ~ 0.1 dex. The 68% constraints on the mass are $\log(m_a/\text{eV}) \in [-12.3, -12.1]$, slightly higher than the injected signal, but the mass is unconstrained at the 99.7% level. This simple exercise highlights the importance of a high-quality magnetic field model, and hence the need for more detailed RM mapping of target clusters.

We thank the *Chandra* Science Center, and especially Hermann Marshall, for advice and guidance in the execution of these observations. D.M. thanks Jens Jache for discussions. We also thank the anonymous referee for their suggestions and insights that improved the manuscript. C.S.R. thanks the UK Science and Technology Facilities Council (STFC) for support under the New Applicant grant ST/R000867/1, and the European Research Council (ERC) for support under the European Union’s Horizon 2020 research and innovation programme (grant 834203). D.M. acknowledges support from the Swedish Science Agency grant 2018-03641 and from European Research Council Grant 742104. H.R.R. acknowledges support from an STFC Rutherford Fellowship and an Anne McLaren Fellowship. A.C.F. acknowledges support from ERC Advanced Grant 340442. F.T. acknowledges support by the Programma per Giovani Ricercatori—anno 2014 “Rita Levi Montalcini.” S.V. acknowledges support from a Raymond and Beverley Sackler Distinguished Visitor Fellowship and thanks the host Institute, the Institute of Astronomy, where this work was concluded. S.V. also acknowledges support by the Science and Technology Facilities Council (STFC) and by the Kavli Institute for Cosmology, Cambridge. S.V. also thanks support from NASA under the *Chandra* Guest Observer program. This work used, in part, the Cambridge COSMOS SMP system, part of the STFC DiRAC HPC Facility supported by BIS NeI capital grant ST/J005673/1 and STFC grants ST/H008586/1, ST/K00333X/1.

ORCID iDs

Christopher S. Reynolds  <https://orcid.org/0000-0002-1510-4860>

M. C. David Marsh  <https://orcid.org/0000-0001-7271-4115>

Andrew C. Fabian  <https://orcid.org/0000-0002-9378-4072>

Robyn Smith  <https://orcid.org/0000-0001-5626-5209>

Francesco Tombesi  <https://orcid.org/0000-0002-6562-8654>

Sylvain Veilleux  <https://orcid.org/0000-0002-3158-6820>

References

- Abbott, L. F., & Sikivie, P. 1983, *PhLB*, **120**, 133
- Ajello, M., Albert, A., Anderson, B., et al. 2016, *PhRvL*, **116**, 161101
- Angus, S., Conlon, J. P., Marsh, M. C. D., Powell, A. J., & Witkowski, L. T. 2014, *JCAP*, **2014**, 026
- Armengaud, E., Avignone, F. T., Betz, M., et al. 2014, *JInst*, **9**, T05002
- Arnaud, K. A. 1996, in ASP Conf. Ser. 101, *Astronomical Data Analysis Software and Systems V*, ed. G. H. Jacoby & J. Barnes (San Francisco, CA: ASP), 17
- Bähre, R., Döbrich, B., Dreyling-Eschweiler, J., et al. 2013, *JInst*, **8**, T09001
- Berg, M., Conlon, J. P., Day, F., et al. 2017, *ApJ*, **847**, 101
- Boyarsky, A., Ruchayskiy, O., Iakubovskiy, D., & Franse, J. 2014, *PhRvL*, **113**, 251301
- Brockway, J. W., Carlson, E. D., & Raffelt, G. G. 1996, *PhLB*, **383**, 439
- Bulbul, E., Markevitch, M., Foster, A., et al. 2014, *ApJ*, **789**, 13
- Caputo, A., Sberna, L., Frías, M., et al. 2019, *PhRvD*, **100**, 063515
- Chen, L., & Conlon, J. P. 2018, *MNRAS*, **479**, 2243
- Churazov, E., Forman, W., Jones, C., & Bohringer, H. 2003, *ApJ*, **590**, 225
- Conlon, J. P., Day, F., Jennings, N., Krippendorf, S., & Rummel, M. 2017, *PhRvD*, **96**, 123009
- Conlon, J. P., Day, F., Jennings, N., Krippendorf, S., & Rummel, M. 2017, *JCAP*, **1707**, 005
- Davis, J. E. 2001, *ApJ*, **562**, 575
- Dine, M., & Fischler, W. 1983, *PhLB*, **120**, 137
- Fabian, A. C., Sanders, J. S., Ettori, S., et al. 2000, *MNRAS*, **318**, L65
- Fabian, A. C., Sanders, J. S., Taylor, G. B., et al. 2006, *MNRAS*, **366**, 417
- Graham, P. W., Irastorza, I. G., Lamoreaux, S. K., Lindner, A., & van Bibber, K. A. 2015, *ARNPS*, **65**, 485
- Grifols, J. A., Massó, E., & Toldrà, R. 1996, *PhRvL*, **77**, 2372
- Hitomi Collaboration, Aharonian, F., Akamatsu, H., et al. 2018, *PASJ*, **70**, 13
- Hitomi Collaboration, Aharonian, F. A., Akamatsu, H., et al. 2017, *ApJL*, **837**, L15
- Irastorza, I. G., & Redondo, J. 2018, *PrPNP*, **102**, 89
- Kalberla, P. M. W., Burton, W. B., Hartmann, D., et al. 2005, *A&A*, **440**, 775
- Kohri, K., & Kodama, H. 2017, *PhRvD*, **96**, 051701
- Marsh, M. C. D., Russell, H. R., Fabian, A. C., et al. 2017, *JCAP*, **12**, 036
- Marshall, H. L. 2012, *Proc. SPIE*, **8443**, 844348
- Mukherjee, S., Spergel, D. N., Khatri, R., & Wandelt, B. D. 2019, arXiv:1908.07534
- Payez, A., Evoli, C., Fischer, T., et al. 2015, *JCAP*, **2**, 006
- Peccei, R. D., & Quinn, H. R. 1977, *PhRvL*, **38**, 1440
- Planck Collaboration, Ade, P. A. R., Aghanim, N., et al. 2014, *A&A*, **571**, A16
- Preskill, J., Wise, M. B., & Wilczek, F. 1983, *PhLB*, **B120**, 127
- Tanabashi, M., Hagiwara, K., Hikasa, K., et al. 2018, *PhRvD*, **98**, 030001
- Taylor, G. B., Gugliucci, N. E., Fabian, A. C., et al. 2006, *MNRAS*, **368**, 1500
- Tiwari, P. 2012, *PhRvD*, **86**, 115025
- Vacca, V., Murgia, M., Govoni, F., et al. 2012, *A&A*, **540**, A38
- Vysotskii, M. I., Zel’dovich, Y. B., Khlopov, M. Y., & Chechetkin, V. M. 1978, *JETPL*, **27**, 502
- Weinberg, S. 1978, *PhRvL*, **40**, 223
- Wilczek, F. 1978, *PhRvL*, **40**, 279
- Wouters, D., & Brun, P. 2013, *ApJ*, **772**, 44