Searches for Axion-Like Particles with X-ray Astronomy

Nicholas Jennings
Somerville College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity 2018

Acknowledgements

Firstly I would like to thank my supervisor Joseph Conlon for his teaching and guidance over the course of my time at Oxford. I would like to thank all my collaborators from whom I have learned so much: Marcus Berg, Francesca Day, Sven Krippendorf, Andrew Powell, Francesco Muia and Markus Rummel. I would also like to thank the wider Theoretical Physics community at Oxford for many stimulating seminars, lectures and conversations. I have been funded by the Science and Technology Facilities Council and my supervisor's European Research Council grant.

I would like to thank Anna, Jeannette, Ian and Laura for the time and effort they took to help with proof-reading this thesis, and their general support and inspiration.

Statement of Originality

This thesis is entirely written by myself, except as indicated in the text. It contains work that I have conducted individually and in collaboration with my supervisor and his group. Here I detail my contributions to work done in collaboration. Chapter 1 contains an introductory review that is entirely my own work. Chapter 2 sets out the Aims and Objectives of the thesis, and is entirely my own work. Chapters 3, 4 and 5 are based collectively on the papers [1], [2] and [3]. The work in these papers was done collaboratively with Joseph Conlon, Francesca Day and Sven Krippendorf as well as: with Marcus Berg, Andrew Powell and Markus Rummel [1]; with Markus Rummel [2]; and with Francesco Muia [3].

My main contributions to these papers were: processing and analysing the *Chandra* satellite images of NGC 1275; and modelling the contamination in the *Chandra* NGC 1275 images using the `jdpileup` analysis tool, and writing up the relevant sections in [1]; processing and analysing *Chandra* satellite images of other point sources in [2]; running the `SIXTE` software to simulate the *Athena* satellite; deriving estimated ALP bounds from *Athena*; and writing up the paper [3]. In Chapter 5 I briefly cover work done in [4], which was done collaboratively with Joseph Conlon, Francesca Day, Sven Krippendorf and Markus Rummel. Chapter 6 is my conclusion which is entirely my own work. No part of this thesis has been submitted for any other qualification.

Abstract

Axion-Like Particles (ALPs) are a very well-motivated class of Beyond the Standard Model particles. They can be implemented as a minimal extension of the Standard Model that could solve many mysteries (the Strong CP Problem, Dark Matter, origin of Inflation), and also commonly arise in String Theory compactifications. This motivates the commissioning of dedicated experiments looking for these particles, as well as the utilisation of telescope data to search for their astrophysical effects. Through their coupling to electromagnetism, they will interconvert with photons in the presence of a magnetic field. A promising environment to look for this effect is the intracluster medium of Galaxy Clusters, which can have weak but very extensive magnetic fields. At X-ray energies, the conversion probability of photons to ALPs is periodic in energy, imprinting a quasi-sinusoidal oscillation on the energy spectrum of an object shining through the cluster. This thesis describes an analysis of data taken by the *Chandra* and *XMM-Newton* satellites of point sources (such as active galactic nuclei) shining through galaxy clusters. The absences of modulations in the spectra of these objects lead to constraints on the ALP-photon coupling. With this analysis, a previously unexplored region of parameter space is ruled out. This thesis also details simulations that were performed of the capabilities of the *Athena* X-ray Observatory, due to launch in 2028, and predicts its ability to place further constraints on ALPs.

Contents

List of Figures

Chapter 1

Introduction to Axions

1.1 The Strong CP Problem and the QCD axion

1.1.1 Quantum Chromodynamics and the Strong CP Problem

The Standard Model (SM) remains our best description of particle physics to date, and has an excellent track record of explaining data from experiments ranging from high-energy supercolliders to precision tests of its fundamental parameters. The SM organises the particles of nature into the gauge symmetry group SU(3)×SU(2)×U(1), where SU(N) is the special unitary group of dimension N. SU(2)×U(1) incorporates the electromagnetic and weak nuclear forces of nature, with photons and the W,Z bosons being the mediators. $\mathrm{SU}(3)_C$ relates to the strong nuclear force, with the C referring to the color charge of particles that determines how they interact via the strong force. The theory underpinning our understanding of the strong force is known as Quantum Chromodynamics (QCD), with the following Lagrangian density:

$$\mathcal{L} = -\frac{1}{g^2} G^a_{\mu\nu} G^{a\,\mu\nu} + \sum_f \bar{\psi}_f (i\mathrm{D}_\mu \gamma^\mu - m_f)\psi_f \tag{1.1}$$

As the masses for the up and down quarks are much smaller than Λ_{QCD} (the scale below which quarks are confined into hadrons), we can consider the limit where these masses are zero. In this case, there is a $\mathrm{SU}(2) \otimes \mathrm{SU}(2)$ chiral symmetry of the QCD Lagrangian. The diagonal subgroup is isospin, while the remainder is a Nambu-Goldstone symmetry with the pions as the massless Nambu-Goldstone bosons. This symmetry is dynamically broken by quark condensates $\langle u\bar{u}\rangle$, $\langle d\bar{d}\rangle$, leading to small masses for the pions. However, there is an additional U(1) symmetry of the Lagrangian when the up and down quark masses are zero:

$$\psi_f \rightarrow \mathrm{e}^{-i\alpha\gamma_5}\psi_f, f = 1, 2. \tag{1.2}$$

where ψ_1 and ψ_2 are the up and down quarks. This symmetry is broken by the same quark condensates as break the chiral $\mathrm{SU}(2)\otimes\mathrm{SU}(2)$ symmetry, which indicates there should be a fourth light meson. However, no such particle exists.

This was dubbed "The $\mathrm{U}(1)_{\mathrm{A}}$ problem" by S. Weinberg [5] because there appeared to be no Lagrangian term that could break this $\mathrm{U}(1)_{\mathrm{A}}$ symmetry. The chiral anomaly [6, 7, 8] does provide a contribution to the action:

$$\delta W = \alpha \frac{g_s^2 N}{32\pi^2} \int \mathrm{d}^4x \, G^a_{\mu\nu} \tilde{G}^{a\mu\nu} \, , \tag{1.3}$$

however the term inside the integral is a total divergence:

$$G^a_{\mu\nu}\tilde{G}^{a\mu\nu} = \partial_\mu K^\mu = \partial_\mu \epsilon^{\mu\alpha\beta\gamma} A_{a\alpha} \left(G_{a\beta\gamma} - \frac{g_s}{3} f_{abc} A_{b\beta} A_{c\gamma} \right) . \tag{1.4}$$

This means that the contribution to the action is a surface integral. If one assumes the boundary condition $A^\mu_a = 0$ at spatial infinity, the contribution disappears and there is no $\mathrm{U}(1)_A$-breaking term. However, it was shown in [9, 10] that the boundary condition is $A^\mu_a = 0$ *or a gauge transformation thereof.* These gauge transformations Ω_n are characterised by their behaviour at spatial infinity:

$$\Omega_n \to \mathrm{e}^{2\pi i n}, \tag{1.5}$$

where n is an integer characterising the topological winding number. Field configurations of different winding numbers cannot be continuously deformed into each other. Taking the superposition of these n-vacua gives the true vacuum:

$$|\theta\rangle = \sum_n \mathrm{e}^{in\theta} |n\rangle. \tag{1.6}$$

The existence of this complicated vacuum structure in QCD leads to an additional term in the effective action:

$$S_{eff}[A] = S_{\mathrm{QCD}}[A] + \theta \frac{g_s^2}{32\pi^2} \int \mathrm{d}^4x G^a_{\mu\nu} \tilde{G}^{a\mu\nu} \tag{1.7}$$

We now wish to calculate the experimental effects of this extra Lagrangian term. To do so, we must incorporate our above discussion of SU(3) gauge theory into a theory with quarks and weak interactions. The quark masses are described by the Cabibbo–Kobayashi–Maskawa (CKM) matrix, which is measured to be complex. Diagonalising to a physical basis requires a chiral transformation, which changes the theta vacuum [11]. Thus θ receives a correction:

$$\bar{\theta} = \theta + \arg\,\det\, M \tag{1.8}$$

for a mass matrix M_{ij}.

The extra term in Equation 1.7 breaks parity and time reversal invariance, but conserves charge conjugation invariance in the strong force. Thus the term breaks CP invariance. The main effect that can be experimentally detected is an induced neutron electric dipole moment (EDM) that takes the following form:

$$d_n \simeq \frac{e\,\bar{\theta}\,m_q}{m_N^2}\,. \tag{1.9}$$

The most sensitive constraint on the neutron EDM is $-3.8 < d_n < 3.4 \times 10^{-26}\,e\,\mathrm{cm}$ at 95% confidence limit (CL) [12]. This implies that $\bar{\theta} \lesssim 10^{-10}$. The smallness of this number presents a puzzle, known as the "Strong CP Problem". Both components of $\bar{\theta}$ are dimensionless quantities that are typically large. Given that they result from different interactions, there is nothing in the SM to indicate that they ought to cancel so exactly. If one makes an argument that $\theta = 0$, one must also explain why arg det M is so small. This places a severe constraint on models that can successfully solve the Strong CP Problem. I provide an incomplete review of proposed solutions in the literature below.

1.1.2 Solutions to the Strong CP Problem

Spontaneous breaking of CP

We can set $\theta = 0$ if CP is a symmetry of the Standard Model. Upon spontaneous breaking of this symmetry, radiative corrections from loop diagrams will generate a correction to θ. We require the corrections to be smaller than 10^{-10}, which in general constrains the models to require the contributions of 1-loop diagrams to cancel as well. The following papers contain models that achieve small θ [13, 14, 15, 16], although some of them are not compatible with the latest experimental constraints. Some of these rely on the high energy, CP-conserving theory being broken to the SM with additional particles that have hitherto escaped detection. It has also been challenging to include them in a Grand Unified Theory (GUT). An elegant class of GUTs that are broken to the SM and feature small $\bar{\theta}$ were discovered in [17, 18, 19]. These solutions to the Strong CP Problem are known collectively as the Nelson-Barr Mechanism.

An overview of the challenges facing this model, including hierarchy and fine-tuning problems, can be found in [20].

Massless up quark

An additional chiral symmetry would solve the Strong CP Problem. This could be achieved if arg det $M = 0$, which could occur if the lightest quark (the up quark) was massless [21]. However, lattice calculations indicate that the up quark mass $m_u \sim 1\,\mathrm{MeV}$, which strongly disfavours this explanation [22].

1.1.3 Axions as a solution to the Strong CP Problem

A compelling alternative to a massless up quark is to explicitly add a new global U(1) chiral symmetry to the Lagrangian [23, 24], known as the Peccei-Quinn symmetry $\mathrm{U}(1)_{PQ}$. This symmetry is spontaneously broken by the scalar field potential:

$$V(\varphi) = \lambda(|\varphi|^2 - f_a^2/2)^2, \tag{1.10}$$

where f_a is the symmetry breaking scale. This induces, after symmetry breaking, an expectation value for the field:

$$\langle\varphi\rangle = \left(\frac{f_a}{\sqrt{2}}\right)e^{ia(x)/f_a}, \tag{1.11}$$

which introduces a Nambu-Goldstone boson into the SM: the axion $a(x)$ [25, 26], which transforms under the chiral symmetry as

$$a(x) \to a(x) + \alpha f_a. \tag{1.12}$$

The existence of the axion adds the following interaction terms to the SM, in addition to the last term from Equation 1.7:

$$\mathcal{L} = \mathcal{L}_{SM} + \bar{\theta}\frac{g_s^2}{32\pi^2}G^a_{\mu\nu}\tilde{G}^{a\mu\nu} - \frac{1}{2}\partial_\mu a\partial^\mu a + \xi\frac{a}{f_a}\frac{g_s^2}{32\pi^2}G^a_{\mu\nu}\tilde{G}^{a\mu\nu} + \mathcal{L}_{int}, \tag{1.13}$$

where $\mathcal{L}_{int}$ represents interaction terms connecting the SM to derivative couplings of the axion. The penultimate term results from the chiral anomaly of the $\mathrm{U}(1)_{PQ}$ current.

The transformation shown in Equation 1.12 is broken by non-perturbative QCD effects. These come in at a scale $\Lambda_{QCD} \ll f_a$, and break the transformation to:

$$a(x) \to a(x) + 2\pi f_a. \tag{1.14}$$

This generates an effective potential for the axion which has its minimum at $\langle a\rangle = -\bar{\theta} f_a/\xi$, as described by the Vafa-Witten Theorem [27]. Therefore, the axion dynamically relaxes to a minimum where the $\bar{\theta}$ term is cancelled out, and the Strong CP Problem is solved. This creates a periodic effective potential for the axion, of the form (redefining $a \rightarrow a_{phys} = a - \langle a\rangle$ so that $V(0) = 0$):

$$V(a) = \Lambda_a^4 \left(1 - \cos\left(\frac{a_{phys}}{f_a}\right)\right). \tag{1.15}$$

(For brevity I will refer to a_{phys} as a from now on). This generates an effective mass for the axion with its second derivative of the form:

$$m_a^2 = \left\langle \frac{\partial^2 V}{\partial a^2} \right\rangle \Bigg|_{\langle a\rangle=0} = \frac{\Lambda_a^4}{f_a^2}. \tag{1.16}$$

Thus the axion is a pseudo-Nambu Goldstone boson.

The values of Λ_a and f_a are model dependent. They can be calculated using current algebra techniques or effective Lagrangian descriptions. For the QCD axion the mass takes the form:

$$m_a = \mathcal{A} \frac{m_\pi f_\pi}{f_a} \frac{\sqrt{m_u m_d}}{m_u + m_d}, \tag{1.17}$$

where $\mathcal{A}$ is a constant that represents the specific couplings of the axion to other SM particles for a particular theory (I will go into more detail below).

1.1.4 “Invisible” Axions

In the original PQWW (Peccei-Quinn-Weinberg-Wilczek) model [25, 26, 23], the $\mathrm{U}(1)_{\mathrm{PQ}}$ symmetry is achieved through two Higgs doublets: one giving mass to up-type quarks, and the other giving mass to down-type quarks and leptons. The Yukawa interactions are:

$$\mathcal{L} = \Gamma_{ij}^u \bar{Q}_{Li} \Phi_1 u_{Rj} + \Gamma_{ij}^d \bar{Q}_{Li} \Phi_2 d_{Rj} + \Gamma_{ij}^\ell \bar{L}_{Li} \Phi_2 \ell_{Rj} + \text{h.c.}, \tag{1.18}$$

with the axion being the common phase field of Φ_1 and Φ_2:

$$\Phi_1 = \frac{v_1}{\sqrt{2}} \mathrm{e}^{iax/v_F} \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad \Phi_2 = \frac{v_2}{\sqrt{2}} \mathrm{e}^{ia/xv_F} \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \tag{1.19}$$

where $x = v_2/v_1$ and $v_F = \sqrt{v_1^2 + v_2^2}$. The ALP-photon-photon interaction comes from the Lagrangian term:

$$\mathcal{L}_{a\gamma\gamma} = \frac{\alpha}{4\pi} K_{a\gamma\gamma} \frac{a}{f_a} F_{a\mu\nu} \widetilde{F_a^{\mu\nu}}, \tag{1.20}$$

where the coupling $K_{a\gamma\gamma}$ receives contributions from the axion mixing with the π_0 and η mesons:

$$K_{a\gamma\gamma} = \frac{N}{2}\left(x + \frac{1}{x}\right)\frac{m_u}{m_u + m_d}, \tag{1.21}$$

where N is the number of quark flavours. The mass of the axion is given by Equation 1.17 with:

$$\mathcal{A} = \frac{N}{2}\left(x + \frac{1}{x}\right). \tag{1.22}$$

In this model f_a is at the scale of electroweak symmetry breaking $v_F \simeq 250\,\mathrm{GeV}$. This has long since been ruled out by beam dump experiments [28] and colliders [29]. However, models that put f_a at a much higher scale, known as 'invisible axion' models [30], have not been ruled out across all parameter space.

KSVZ axion

The Kim-Shifman-Vainshtein-Zakharov (KSVZ) axion [31, 32] features a scalar field σ and a heavy quark Q that both carry PQ charge. They interact via the Yukawa term:

$$\mathcal{L}_Y = -\lambda_Q \sigma \bar{Q}_L Q_R + \text{h.c.} \tag{1.23}$$

The KSVZ axion interacts with light quarks via the color and EM anomalies:

$$\mathcal{L}^{KSVZ} \supset \frac{a}{f_a}\left(\frac{g_s^2}{32\pi^2} G_{a\mu\nu}\widetilde{G_a^{\mu\nu}} + 3e_Q^2 \frac{\alpha}{4\pi} F_{a\mu\nu}\widetilde{F_a^{\mu\nu}}\right), \tag{1.24}$$

where e_Q is the electromagnetic charge of the heavy quark. Note that there are no tree-level couplings to leptons for the KSVZ axion. The resulting axion-photon coupling is:

$$K_{a\gamma\gamma} = 3e_Q^2 - \frac{4m_d + m_u}{3(m_d + m_u)}. \tag{1.25}$$

Canonically one takes the Q field to be neutral, so $e_Q = 0$. For the KSVZ axion $\mathcal{A} = 1$ as defined in Equation 1.17, giving a mass:

$$m_a = 6.3\,\mathrm{eV}\left(\frac{10^6\,\mathrm{GeV}}{f_a}\right). \tag{1.26}$$

DFSZ axion

The Dine-Fischler-Srednicki-Zhitnisky (DFSZ) axion [33, 34] requires two Higgs doublets H_u and H_d like the PQWW axion, plus an additional scalar field ϕ. The Higgs doublets and the scalar interact via the potential:

$$V = \lambda_H \phi^2 H_u H_d. \tag{1.27}$$

This has a mass corresponding to Equation 1.17 where $\mathcal{A} = \mathcal{C} = 6$ is the color anomaly (also known as the Domain Wall Number i.e. the number of vacua in the range $a \in [0, 2\pi f_a]$). The axion-photon coupling is:

$$K_{a\gamma\gamma} = \frac{4}{3} - \frac{4m_d + m_u}{3(m_d + m_u)}. \tag{1.28}$$

1.2 Axion-Like Particles in String Theory

We have seen how the axion solves the Strong CP Problem, and how it can be included in a minimal extension of the Standard Model. We can search for the axion either through its coupling to the strong force ($G\tilde{G}$) or electromagnetism ($F\tilde{F}$). One could also consider searching for particles that have couplings in these forms, independently of whether they solve the Strong CP Problem. This removes the dependence of the axion mass on the order parameter as given by Equation 1.17, meaning that one could consider searching for such particles anywhere in the m_a–f_a plane. I refer to particles that do not exist on the line defined in Equation 1.17 as Axion-Like Particles (ALPs), to distinguish them from axions that solve the Strong CP Problem[1].

The question is, whether there are any compelling motivations for such particles to exist. The answer is emphatically yes, when one looks at String Theory models that compactify a 10-dimensional superstring theory (or 11-dimensional M-Theory) down to the 4 dimensions we are familiar with.

I present a very brief outline of the appearance of axions in heterotic string theory. For reviews, see [35, 36]. All superstring theories contain an antisymmetric field tensor B_{MN}, where $M, N = 0, ..., 9$ which ensures anomaly cancellation. Its gauge invariant field strength tensor is:

$$H = \mathrm{d}B - \omega_{3\mathrm{Y}} + \omega_{3\mathrm{L}}, \tag{1.29}$$

$$\omega_{3\mathrm{Y}} = \mathrm{tr}(AF - \frac{1}{3}A^3), \;\; \omega_{3\mathrm{L}} = \mathrm{tr}(\omega R - \frac{1}{3}\omega^3), \tag{1.30}$$

where A is the gauge field, F is the Yang-Mills field strength, ω the spin connection and R the Riemann tensor. The Bianchi identity for H is:

$$\mathrm{d}H = \frac{1}{16\pi^2}(\mathrm{tr}R \wedge R - \mathrm{tr}F \wedge F), \tag{1.31}$$

where the wedge products are equivalent to $F\tilde{F}$ terms written in differential form notation. The relevant terms in the low-energy 10-dimensional supergravity Lagrangian are:

$$\mathcal{L} \supset \frac{1}{2\kappa_{10}^2}\sqrt{-g}R - \frac{1}{4\kappa_{10}^2}H \wedge \star H - \frac{\alpha'}{8\kappa_{10}^2}\mathrm{tr}F \wedge \star F, \tag{1.32}$$

where R is the Ricci scalar and F is the curvature of SO(32), and $\kappa_{10}^2/g_{10}^2 = \alpha'/4$. Compactifying on a six-dimensional manifold Z with volume V_Z gives a 4-dimensional effective action with a term:

$$S \supset -\frac{2\pi V_Z}{g_s^2 l_s^4}\int \left(\frac{1}{2}H \wedge \star H\right). \tag{1.33}$$

Focussing on modes of the B-field tangent to the 4D manifold, and constant on Z, we introduce a Lagrange multiplier for the Bianchi identity of Equation 1.31:

$$S(a) = \frac{g_s^2 l_s^4}{2\pi V_Z}\int \mathrm{d}^4x\left(-\frac{1}{2}\partial_\mu a \partial^\mu a\right) + \int a\frac{1}{16\pi^2}(\mathrm{tr}F \wedge F - \mathrm{tr}R \wedge R). \tag{1.34}$$

Thus we have an axion with decay constant [37]:

---

[1] This is standard nomenclature in the literature. An alternative convention refers to these particles generally as axions, and specifies the QCD axion as the one that solves the Strong CP Problem.

$$f_a = \frac{g_s^2 l_s^2}{\sqrt{2\pi V_Z}}. \tag{1.35}$$

This is the so-called model-independent axion, as its properties are somewhat independent of the manifold Z. Such an axion would typically have a decay constant $\mathcal{O}(10^{16}\,\text{GeV})$. In addition, there are model-dependent heterotic string axions that correspond to zero-modes of the B-field on Z. After compactification we have a 4-dimensional effective action for the modes b_i, which only features derivative couplings. Axionic couplings arise due to the Green-Schwarz anomaly cancellation mechanism, which introduces one-loop couplings. The decay constant for a linear combination of b_i (for a near-isotropic Z) is:

$$f_a \sim \frac{V_Z^{1/3}}{2\pi g_s^2 l_s^4}, \tag{1.36}$$

giving typical values $f_b \sim 10^{17}\,\text{GeV}$.

As before, non-perturbative effects generate an effective potential for the axion, which generates a mass. The contribution to the superpotential from these effects is:

$$W = M^3 e^{-S_{inst}+ia}, \tag{1.37}$$

where M is the scale of the instanton effects. Plugging this into the effective scalar potential:

$$V = e^{K/M_P^2}(K^{i\bar{j}} D_i W D_{\bar{j}} W - 3|W|^2/M_P^2), \tag{1.38}$$

where K is the Kähler potential, we find an axion potential of the form shown in Equation 1.15, with:

$$\Lambda_a^4 = m_S^2 M_{pl}^2 e^{-S_{inst}}, \tag{1.39}$$

where m_S is the scale of SUSY breaking. In the case of model-dependent axions $S_{inst} \propto V_C$, where V_C is the volume of the cycle. This produces axions with mass:

$$m_a^2 = \left(\frac{m_S^2 M_P^2}{f_a^2}\right) e^{-A\,V_C}, \tag{1.40}$$

for a model-dependent constant A. It is therefore possible to have a large ($\mathcal{O}(100)$) number of axions with a similar number per decade of mass. This motivates experimental searches for ALPs at many different masses.

Axion-Like Particles can also appear in $E_8 \times E_8$ heterotic String Theory [38], and Type I, IIA and IIB String Theories [39], as well as M-Theory [40]. The higher dimensional gauge invariance of the field guarantees that the axions generated will be massless in all orders of perturbation theory. Therefore, with the only source of mass coming from instanton effects, there could be many light Axion-Like Particles [41]. This plenitude of ALPs is referred to as the string axiverse [42]. For more details on ALPs in String Theory see [43, 44, 45, 46, 47].

1.3 Dark Matter

At least 80% of the matter in the Universe is Dark Matter (DM) that interacts very weakly with the Standard Model. This conclusion comes from galaxy rotation curves and the CMB power spectrum, among other observations. This indicates the presence of at least one new BSM particle. The DM must be cold, stable and weakly interacting.

Axions are certainly weakly-interacting, and for low masses are very stable. Axions produced thermally in the early Universe that then freeze out once the interaction rate is smaller than the Hubble rate would be a hot component of Dark Matter. There are severe experimental restrictions on the proportion of Dark Matter that can be hot. Cold axions could be produced in the early Universe through another process: the 'misalignment mechanism'. The axion field, after U(1) symmetry breaking and instanton effects, evolves in a FLRW universe in the following way:

$$\ddot{a} + 3H\dot{a} + m_a(T)^2 a = 0, \tag{1.41}$$

where $H = \dot{R}/R$ is the Hubble parameter for scale parameter R and the axion mass $m_a(T)$ depends on the temperature of the Universe T, being practically zero for T greater than the scale of non-perturbative physics Λ_a. The density of the background axion field is:

$$\bar{\rho} = \frac{1}{2}\dot{a}^2 + m_a(T)^2 a^2. \tag{1.42}$$

At early times $H > m_a$, the axion field is overdamped and frozen at its initial value θ_a. Later when $H < m_a$ the axion begins to oscillate around its minimum. Here the density scales as $\rho_a \propto R(t)^{-3}$, meaning that it behaves as matter (unlike for example neutrinos which behave as radiation). The cosmic mass fraction of cold axions is set by the initial misalignment angle θ_a and f_a:

$$\Omega_a h^2 \approx 0.71 \left(\frac{f_a}{10^{12}\,\mathrm{GeV}} \right)^{7/6} \left(\frac{\theta_a}{\pi} \right)^2, \tag{1.43}$$

provided that the reheating temperature is below f_a and that there is negligible decay to BSM particles. The cosmic mass fraction is inferred to be [48]:

$$\Omega_{DM} h^2 = 0.1186 \pm 0.0020, \tag{1.44}$$

meaning that cold axions could comprise a fraction of cold DM up to 100%. To prevent overclosure of the universe, a bound of $f_a \lesssim 3 \times 10^{11}\,\mathrm{GeV}$ is inferred for an average misalignment angle of $\theta_a^2 = \pi^2/3$. This bound could be extended upwards if θ_a happened to be small.

A potential 'drawback' of the theory of axion Dark Matter is that there is no prediction for the initial misalignment in our Universe, and thus the abundance of DM in the Universe. It seems somewhat arbitrary that the DM abundance is such that the Universe is long-lived and supports complex structure, unless one invokes anthropic arguments. In addition, there is no guarantee that axions comprise 100% of the Dark Matter, meaning that any experiments attempting to detect it are sensitive not only to the coupling to the SM but also to the fraction of ALP DM to total DM. This makes such experiments highly model-dependent.

Axion-Like Particles can also be viable Dark Matter candidates. In fact, Cold Dark Matter has a number of small-scale 'problems' where predictions of the model are in tension with astrophysical observations [49]. These include the 'missing satellite' problem [50, 51] (too few low-mass satellites observed in the Milky Way), the too-big-to-fail problem [52] (massive satellites in the Milky Way are too dark), and the core-cusp problem (observations of many low-mass systems show a flat central density profile rather than a cusp expected from collisionless CDM). These 'problems' may be a result of incorrectly simulating the effects of baryonic feedback on structure formation [53], or could point to DM having a component that is not composed of cold, collisionless particles [54, 55].

There is a possibility that these problems could be solved if the Dark Matter consists of ultralight ALPs ($m_a \lesssim 10^{-22}\,\mathrm{eV}$). The density of ALPs produced by the misalignment mechanism depends on when the ALP oscillations begin in cosmic history. A good approximation, for the case that it happens during the radiation-dominated epoch, is [56]:

$$\Omega_a \approx \frac{1}{6}(9\Omega_r)^{3/4}\left(\frac{m_a}{H_0}\right)^{1/2}\left(\frac{f_a}{M_{pl}}\right)^2 \theta_a^2 \tag{1.45}$$

where Ω_r is the energy density of radiation. When simulating the cosmic evolution of the DM density, wavelike effects must be taken into account on scales of order the de Broglie wavelength of the ALP (for $m_a = 10^{-22}\,\mathrm{eV}$ this scale is around 1 kpc). N-body simulations of ultralight ALP DM indicate that smoother cores would result compared to CDM [57]. Other wavelike effects such as interference bands might also occur. In short, ultralight ALP DM could help explain current tensions with CDM models, while also providing unique phenomenology to search for [58, 59]. Other consequences of ALP DM could include axion miniclusters [60, 61, 62, 63] and caustic rings [64, 65].

1.4 Axion-photon conversion in magnetic fields

In summary, Axion-Like Particles are extremely versatile when it comes to solving extant problems in cosmology. Such models require relatively little fine-tuning as there are few parameters in ALP models to play with, and often a large range of masses and couplings are allowed, based on current experimental constraints. This motivates attempts to observe the existence of ALPs directly via their couplings to the SM. In particular, their coupling to Electromagnetism $g_{a\gamma\gamma}$ could allow detection of ALPs with telescopes or laser experiments. An ALP will decay to two photons with the rate:

$$\Gamma_{a\to\gamma\gamma} = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi}. \tag{1.46}$$

Therefore using telescopes to search for Extragalactic Background Light (EBL) allows one to constrain high mass, high coupling ALPs [66]. In addition these photons would cause distortion of the Cosmic Microwave Background Radiation blackbody spectrum, and influence the abundance of deuterium depending on the epoch of axion decay [67]. This constrains the QCD axion to have a mass $m_a \lesssim 10\,\text{eV}$.

If one wishes to look at lower masses and couplings, the conversion rate to photons needs to be enhanced. It was shown in [68] that this can be achieved with macroscopic magnetic fields. In this section, I review the derivation of the conversion probabilities for ALPs to photons.

1.4.1 Dispersion relations for photons in a magnetic field

The propagation of a photon through a time- and space-independent external magnetic field, for sufficiently small photon energy ω and magnetic field strength B (I will define 'small' shortly), is well described by the fourth-order expansion of the Euler-Heisenberg effective Lagrangian [69][2]:

---

[2] For all equations in this thesis I use natural units $\hbar = c = 1$.

$$
\begin{aligned}
\mathcal{L}_{\gamma\gamma} &= -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{\alpha^2}{360m_e^4}[4(F_{\mu\nu}F^{\mu\nu})^2 + 7(F_{\mu\nu}\widetilde{F}^{\mu\nu})^2] \\
&= \frac{1}{2}(\mathbf{E}^2 - \mathbf{B}^2) + \frac{2\alpha^2}{45m_e^4}[(\mathbf{E}^2 - \mathbf{B}^2)^2 + 7(\mathbf{E}\cdot\mathbf{B})^2], \qquad (1.47)
\end{aligned}
$$

where $\mathbf{E} = \mathbf{E}^{\mathbf{wave}}$ is the electric field of the photon, $\mathbf{B} = \mathbf{B}^{\mathbf{wave}} + \bar{\mathbf{B}}$ is the sum of the photon magnetic field and the external magnetic field, $\alpha = e^2/4\pi$, and m_e and e are the electron mass and charge respectively. This Lagrangian is a good approximation provided $B \ll B_{\text{crit}} = m_e^2/e \approx 4.41 \times 10^{13}$ Gauss (G) and $\omega \ll 2m_e$, i.e. the threshold for photopair production. The second term in each line of Equation 1.47 causes a polarisation of the vacuum, with a different index of refraction perpendicular to the magnetic field ($n_\perp$) to that parallel to the magnetic field ($n_\parallel$). We show this by calculating the electric displacement $\mathbf{D}^{\mathbf{wave}} \equiv \partial\mathcal{L}/\partial\mathbf{E}^{\mathbf{wave}}$ and magnetic intensity $\mathbf{H}^{\mathbf{wave}} \equiv -\partial\mathcal{L}/\partial\mathbf{B}^{\mathbf{wave}}$, keeping only terms linear in E or B:

$$
D_i^{\text{wave}} = \epsilon_{ij}E_j^{\text{wave}} = \left[\delta_{ij}(1 - \frac{2\alpha^2\bar{B}^2}{45\pi m_e^4}) + \frac{7\alpha^2\bar{B}^2}{45\pi m_e^4}\hat{b_i}\hat{b_j}\right]E_j^{\text{wave}} \qquad (1.48)
$$

$$
H_i^{\text{wave}} = \mu_{ij}^{-1}B_j^{\text{wave}} = \left[\delta_{ij}(1 - \frac{2\alpha^2\bar{B}^2}{45\pi m_e^4}) - \frac{4\alpha^2\bar{B}^2}{45\pi m_e^4}\hat{b_i}\hat{b_j}\right]B_j^{\text{wave}}. \qquad (1.49)
$$

where $\bar{B}$ and $\hat{b_i}$ are the external magnetic field magnitude and direction respectively. We then plug these quantities into the source-free Maxwell equations:

$$
\begin{aligned}
&\nabla\cdot\mathbf{D}^{\mathbf{wave}} = 0 && \nabla\cdot\mathbf{B}^{\mathbf{wave}} = 0 \\
&\nabla\times\mathbf{E}^{\mathbf{wave}} = -\partial\mathbf{B}^{\mathbf{wave}}/\partial t && \nabla\times\mathbf{H}^{\mathbf{wave}} = \partial\mathbf{D}^{\mathbf{wave}}/\partial t. \qquad (1.50)
\end{aligned}
$$

The result of this calculation is that there are two propagation eigenmodes where the electric field of the photon is parallel/perpendicular to the external magnetic field. In these two cases, the indices of refraction are:

$$
n_\perp = 1 + \frac{4}{2}\xi\sin^2\Theta, \quad n_\parallel = 1 + \frac{7}{2}\xi\sin^2\Theta, \qquad (1.51)
$$

where $\xi = (\alpha/45\pi)(\bar{B}/B_{crit})^2$ and $\cos\Theta = \widehat{\mathbf{b}}\cdot\widehat{\mathbf{k}}$. We note that the indices of refraction are small perturbations from unity. We can write out the stationary wave equation, for a photon travelling along the z-direction, with the plane wave ansatz $\mathbf{A}(t,z) = \mathbf{A}\exp i(\omega t - k_z z)$ in the following way:

$$\left(\omega^2 + \partial_z^2 + \begin{pmatrix} Q_\perp & Q_F \\ Q_F & Q_\parallel \end{pmatrix}\right)\begin{pmatrix} |A_\perp\rangle \\ |A_\parallel\rangle \end{pmatrix} = 0, \tag{1.52}$$

where $Q_{\perp/\parallel} = 2\omega^2(n_{\perp/\parallel} - 1)$, and Q_F is a term responsible for the Faraday Rotation effect. While the Faraday Rotation effect is negligible at the photon energies we will be interested in for this thesis, it is a useful diagnostic for the strength of an extragalactic magnetic field, and we will explore it in more detail in Chapter 3. If the magnetic field is unvarying on scales of the photon wavelength, we can expand $\omega^2 + \partial_z^2 = (\omega + i\partial_z)(\omega - i\partial_z) = (\omega + k)(\omega - i\partial_z) \approx 2\omega(\omega - i\partial_z)$ as $k = n_{\perp/\parallel}\omega \approx \omega$. We therefore write the linearised wave equation:

$$\left(\omega - i\partial_z + \begin{pmatrix} \Delta_\perp & \Delta_F \\ \Delta_F & \Delta_\parallel \end{pmatrix}\right)\begin{pmatrix} |A_\perp\rangle \\ |A_\parallel\rangle \end{pmatrix} = 0, \tag{1.53}$$

where $\Delta_{\perp/\parallel/F} = Q_{\perp/\parallel/F}/2\omega$. The above equation accounts for the vacuum effects on the propagation of a photon in a magnetic field. Additional effects can come from the medium through which the propagation occurs. In a plasma featuring electrons and heavy ions, the diagonal terms will be corrected by a term proportional to the plasma frequency:

$$\omega_{\mathrm{pl}} = \sqrt{4\pi\alpha n_e/m_e} = \left(\frac{n_e}{10^{-3}\mathrm{cm}^{-3}}\right)^{0.5} 1.2\times 10^{-12}\,\mathrm{eV}, \tag{1.54}$$

where n_e is the number density of electrons. The wave equation is then:

$$\left(\omega - i\partial_z + \begin{pmatrix} \Delta_{pl} + \Delta_\perp & \Delta_F \\ \Delta_F & \Delta_{pl} + \Delta_\parallel \end{pmatrix}\right)\begin{pmatrix} |A_\perp\rangle \\ |A_\parallel\rangle \end{pmatrix} = 0, \tag{1.55}$$

where $\Delta_{pl} = -\omega_{pl}^2/2\omega$.

1.4.2 Interconversion with ALPs

In this section, I follow the working of [68, 70, 71]. With the inclusion of ALPs, the Lagrangian shown in Equation 1.47 is supplemented by the terms:

$$\mathcal{L}_{a\gamma\gamma} = \frac{1}{2}(\partial_\mu a \partial^\mu a - m_a^2 a^2) + \frac{1}{4M}\, a\, F_{\mu\nu}\tilde{F}^{\mu\nu}\,, \tag{1.56}$$

where $M = 1/g_{a\gamma\gamma}$ is the scale of the PQ symmetry breaking. The Euler-Lagrange equation of motion for an ALP travelling in the z-direction (again using the plane wave ansatz) is:

$$(\omega^2 + \partial_z^2 - m_a^2)|a\rangle + \frac{\omega \bar{B}_T}{M}|A_\parallel\rangle = 0, \tag{1.57}$$

where $\bar{B}_T$ is the component of the magnetic field perpendicular to the z-axis. One can see that mixing occurs between the ALP state and a photon state in a manner analogous to neutrino mixing. An initial quantum state of an ALP will evolve in the following way:

$$|\psi_{init}\rangle = |a(\omega)\rangle \longrightarrow |\psi_{final}\rangle = \alpha|\gamma(\omega)\rangle + \beta|a(\omega)\rangle\ ,$$

where $|\alpha|^2 + |\beta|^2 = 1$, and $\gamma(\omega)$ $[a(\omega)]$ denotes a photon [ALP] with energy ω. Thus ALPs interconvert with photons as they propagate through the magnetic field. One can also see that only the parallel eigenmode of the photon interacts with the ALP. This is because the perpendicular eigenmode is even under CP, while the parallel eigenmode and the ALP are odd, and so it decouples from the ALP modes.

The equation of motion for the parallel eigenmode of the photon receives an identical correction from the presence of ALPs, allowing us to write the linearised equations of motion:

$$\left(\omega + \begin{pmatrix} \Delta_{pl} + \Delta_\perp & \Delta_F & 0 \\ \Delta_F & \Delta_{pl} + \Delta_\parallel & \Delta_{\gamma a} \\ 0 & \Delta_{\gamma a} & \Delta_a \end{pmatrix} - i\partial_z\right)\begin{pmatrix} |A_\perp\rangle \\ |A_\parallel\rangle \\ |a\rangle \end{pmatrix} = 0, \tag{1.58}$$

where $\Delta_{\gamma a} = B_\perp/2M$ and $\Delta_a = -m_a^2/2\omega$.

In the event that the term Δ_F corresponding to Faraday rotation is negligible compared to the other terms (which is the case for high energies), the perpendicular eigenmode decouples completely from the others. We can linearise the remaining 2×2 matrix by rotating to a new basis:

$$\begin{pmatrix} |A'_\parallel\rangle \\ |a'\rangle \end{pmatrix} = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}\begin{pmatrix} |A_\parallel\rangle \\ |a\rangle \end{pmatrix}, \tag{1.59}$$

where the rotation angle is:

$$\theta = \frac{1}{2}\arctan\left(\frac{2\Delta_{\gamma a}}{\Delta_{pl} + \Delta_{\|} - \Delta_a}\right). \tag{1.60}$$

For the extragalactic magnetic fields that I will be analysing in this thesis, it is generally the case that $\Delta_{pl} \gg \Delta_{\|}$, so we can neglect the latter term. Finally we are in a position to calculate the probability that an initial state photon, after propagating through a constant magnetic field for a distance L, has converted to an ALP:

$$|\langle\gamma(0)|a(z)\rangle|^2 = P_{a\to\gamma} = \sin^2(2\theta)\,\sin^2\left(\frac{\Delta}{\cos(2\theta)}\right), \tag{1.61}$$

where

$$\Delta = \frac{m_{eff}^2 L}{4\omega}, \qquad m_{eff}^2 = |m_a^2 - \omega_{pl}^2|\,. \tag{1.62}$$

Henceforth we assume $m_a \ll \omega_{pl}$ and set it to zero. After plugging in constants, $\tan(2\theta)$ and Δ evaluate to:

$$\tan(2\theta) = 4.9\times10^{-2}\left(\frac{10^{-4}\,\mathrm{cm}^{-3}}{n_e}\right)\left(\frac{B_\perp}{1\,\mu\mathrm{G}}\right)\left(\frac{\omega}{3.5\,\mathrm{keV}}\right)\left(\frac{10^{13}\,\mathrm{GeV}}{M}\right), \tag{1.63}$$

$$\Delta = 1.5\times10^{-2}\left(\frac{n_e}{10^{-4}\,\mathrm{cm}^{-3}}\right)\left(\frac{3.5\,\mathrm{keV}}{\omega}\right)\left(\frac{L}{10\,\mathrm{kpc}}\right). \tag{1.64}$$

For $\Delta \ll 1$ and $\theta \ll 1$ the conversion probability simplifies to:

$$P(a\to\gamma) = 2.3\times10^{-8}\left(\frac{B_\perp}{1\,\mu\mathrm{G}}\frac{L}{1\,\mathrm{kpc}}\frac{10^{13}\,\mathrm{GeV}}{M}\right)^2. \tag{1.65}$$

The study of photon-ALP interconversion in astrophysical magnetic fields is well-developed; I provide an incomplete list of articles on the subject here [72, 73, 74, 75, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 92].

1.5 Review of constraints

We have seen that there is strong motivation to search for ALPs across almost all of the $g_a - m_a$ plane, where g_a refers generally to the couplings to SM particles ($g_{a\gamma\gamma}, g_{aee}$ etc.). In this Section I briefly mention the strongest experimental constraints on ALP couplings to SM particles, as well as those from astrophysical observations. Broadly speaking, the experiments can be divided into two kinds. The first kind assumes that ALPs constitute a significant proportion of Dark Matter. As such, these bounds should be expressed in terms of the fraction $Z = \rho_{ADM}/\rho_{DM}$: the ratio of ALP DM density to the total DM density. The second kind is more model-independent, depending only on ALP couplings and mass (plus models of astrophysical systems such as stars, galaxy clusters where appropriate). For a review of experiments, see [93].

1.5.1 Light Shining through Walls

In this experimental set-up, a laser generates photons that pass through two superconducting dipole magnets separated by an optical barrier. Photons that convert to ALPs can pass through the "wall", and can be reconverted to photons on the other side. The advantage of this method is that one has an excellent understanding of the ALP beam for a given $g_{a\gamma\gamma}$ coupling. The disadvantage is that the probability $P_{\gamma\to a\to\gamma} = |P_{\gamma\to a}|^2 \propto g_{a\gamma\gamma}^4$ is highly dependent on the ALP coupling. The best limit is from the OSQAR experiment [94], which used two 9 T LHC magnets and an 18.5 W laser. They found a bound of $g_{a\gamma\gamma} < 3.5 \times 10^{-8}\,\mathrm{GeV}^{-1}$ at 95% confidence for $m_a \lesssim 0.3\,\mathrm{meV}$ [95].

1.5.2 Stellar energy-loss

Low-mass ALPs produced in the centre of stars would efficiently transport energy out of stars due to their weak interactions with SM particles. Stars would therefore be shorter-lived, allowing us to constrain the coupling by looking at populations of stars. An analysis comparing the number of stars on the horizontal branch (HB) to those on the Red-giant branch (RGB) in 39 globular clusters (GCs) found an upper bound at 95% confidence of $g_{a\gamma\gamma} < 6.6 \times 10^{-11}\,\mathrm{GeV}^{-1}$ for $m_a \lesssim 10\,\mathrm{keV}$ [96]. While such studies depend on what stellar models are used, the bound is three orders of magnitude better than LSW experiments, and represents the best bound across a large mass range. Weaker constraints can be derived from the neutrino flux from the Sun and helioseismology. In addition, the coupling to electrons g_{aee} can be constrained from analyses of RGB stars [97] and White Dwarf cooling [98].

Another stringent bound on the ALP-photon coupling can be derived using similar arguments of constraining energy transport in supernovae. In particular, the non-observation of a gamma-ray flux coincident with SN 1987A by the SMM satellite provides a constraint of $g_{a\gamma\gamma} \lesssim 5.3 \times 10^{-12}\,\mathrm{GeV}^{-1}$ for $m_a \lesssim 4.4 \times 10^{-10}\,\mathrm{eV}$ [99]. This provides the best broadband ALP-photon coupling constraint, albeit in a lower range than the constraint from GCs.

1.5.3 Helioscopes

Helioscopes aim to convert the ALPs produced by the Sun, as mentioned in the previous paragraph, into photons using a large bore dipole magnet. The strongest helioscope constraint comes from the CERN Axion Solar Telescope (CAST), which uses a former LHC magnet on a mount which tracks the position of the Sun. The experiment runs while the Sun is set, in order to use the Earth as a photon barrier. For $m_a \lesssim 0.02\,\mathrm{eV}$, a bound of $g_{a\gamma\gamma} < 6.6 \times 10^{-11}\,\mathrm{GeV}^{-1}$ was found at 95% confidence [100]. By filling the magnet with He^4, bounds were derived up to $m_a \sim 0.4\,\mathrm{eV}$ [101], while using He^3 produced bounds up to $m_a \sim 1.17\,\mathrm{eV}$ [102].

1.5.4 Resonant Cavities

If cold ALPs constitute a significant proportion of DM, then we can attempt to detect the local density with microwave cavities. An electromagnetic cavity is permeated by a strong magnetic field, and when the frequency of a cavity mode equals the ALP mass, resonant conversion to microwaves will occur[3]. In order to scan over a range of ALP masses, the cavity must be tunable. The Axion Dark Matter eXperiment (ADMX) has an 8 T magnetic field filling a volume of around 200 litres. A scan over the mass range 1.9–3.3 μeV has allowed them to constrain ALPs down to the KSVZ axion (around $10^{-15}\,\mathrm{GeV}^{-1}$) [103, 104]. Other microwave cavity experiments RBF and UF [105], HAYSTAC [106, 107] and ORGAN [108] have constrained parts of parameter space above the KSVZ axion for $\mathcal{O}(1 - 100\,\mu\mathrm{eV})$.

1.5.5 Black Hole superradiance

Axion-Like Particles could affect the spin evolution of black holes. If the ALPs have a de Broglie wavelength of order the black hole diameter, they form gravitational bound states around a rapidly spinning black hole. As ALPs are bosonic, the occupation number grows exponentially, provided the self-interaction is small. This process extracts angular momentum from the black hole and leads to the formation of a ‘cloud’ which superradiates gravitational waves. Note that the rate of occupation number growth, provided the ALP couplings are small, only depends on the ALP

---

[3] Technically, resonant conversion occurs when the frequency equals the total ALP energy (mass plus kinetic energy), but as the ALPs must be cold the kinetic contribution is negligible.

mass. Thus constraints derived from black hole superradiance extend to arbitrarily small couplings in a given mass range.

This phenomenon would lead to gaps in the mass-spin plot of black hole populations [109]. Precision studies of stellar mass black holes exclude ALPs at 95% confidence in a mass range $6 \times 10^{-13}\,\mathrm{eV} < \mathrm{m_a} < 2 \times 10^{-11}\,\mathrm{eV}$ [110]. The gravitational waves emitted by ALPs around a black hole could be detected by Advanced LIGO [110]. Advanced LIGO could also be used to look at gravitational wave signals from black hole binary mergers, which would be affected by the presence of ALPs [111].

1.5.6 Searches for ALPs with satellites

As searching for ALPs-photon conversion in astrophysical magnetic fields is the topic of this thesis, I do not go into detail here. However, I will briefly mention the main previous results. X-ray observations of the Hydra A cluster produce a bound of $g_{a\gamma\gamma} \lesssim 8.3 \times 10^{-12}\,\mathrm{GeV}^{-1}$ for masses $m_a \lesssim 7 \times 10^{-12}\,\mathrm{eV}$ [112]. Gamma-ray observations of NGC1275 with Fermi-Lat produce bounds of $g_{a\gamma\gamma} \lesssim 5 \times 10^{-12}\,\mathrm{GeV}^{-1}$ for $0.5 \lesssim m_a \lesssim 5\,\mathrm{neV}$ [113]. The H.E.S.S. collaboration observed the AGN PKS 2155-304, and produced bounds of $g_{a\gamma\gamma} \lesssim 2.1 \times 10^{-11}\,\mathrm{GeV}^{-1}$ for $1.5 \times 10^{-8}\,\mathrm{eV} \lesssim \mathrm{m_a} \lesssim 6.0 \times 10^{-8}\,\mathrm{eV}$ [114].

1.5.7 Summary

I summarise the current bounds on the ALP-photon coupling in Figure 1.1 (taken from the Particle Data Group review 2016 [48]). We see that, while a large part of parameter space has already been excluded, a significant unprobed region remains. Finding new ways to probe this parameter space is an important endeavour.

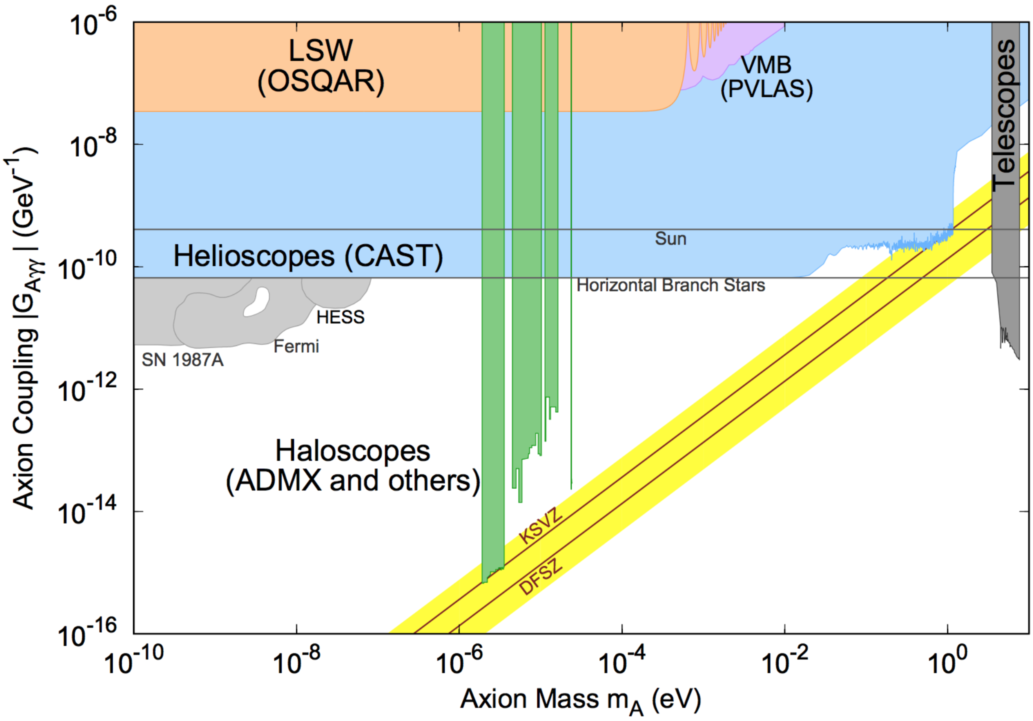

Figure 1.1: Overview of exclusion limits on the ALP-photon coupling, plotted against mass. Included in the graph are the mass-coupling relations for the KSVZ and DFSZ axions (within the yellow band). Full references can be found in the Particle Data Group review on *Axions and other similar particles* [48].

Chapter 2

Aims and Objectives

We have seen that Axion-Like Particles are very well-motivated particles that extend the Standard Model. The QCD axion solves the Strong CP Problem, and could be a solid candidate for being Dark Matter. The coupling of the axions to photons has been constrained through a number of experiments and astrophysical observations. At lower masses, the coupling of the QCD axion to electromagnetism becomes very weak, meaning that, for the time being at least, this region of parameter space is beyond our experimental reach. One might therefore conclude that experiments that do not reach the QCD axion 'line' are of no benefit other than as an incremental advance of technology that might eventually reach down to these exceptionally weak couplings. However, there is a strong motivation to analyse this data carefully to look for Axion-Like Particles that could exist in this parameter space. String theory provides many constructions that contain multiple ALPs that can exist in these areas of parameter space. In addition, simple model-building involving multiple ALPs can push the QCD axion to regions of parameter space where the coupling to electromagnetism is enhanced. We therefore have compelling reasons to find new ways to reach this area of parameter space experimentally.

We have seen that ALPs convert with photons in the presence of a magnetic field, acting as the third polarisation state of the effectively massive photon. The dependence of the conversion probability on certain key parameters can point the way for us to begin our search. The probability increases with the product of the magnetic field strength and the length of the domain squared. One can therefore look to high strength magnetic fields (with the corresponding budgetary limitation of only being able to produce one across a short distance) or look to weak magnetic fields that compensate by extending across great distances. Clearly we cannot generate such fields on Earth; however, we know that such fields exist in some of the largest structures in the universe: galaxy clusters and the giant lobes of radio galaxies.

These appear therefore to be good targets to examine the ALP-photon conversion probability in greater detail.

We will see in Chapter 3 that the conversion probability in galaxy cluster magnetic fields increases with the square of the energy, up to a point of saturation, implying that we should observe galaxy clusters at high energies to benefit from this. The most interesting energy region, between minimal conversion at low energies and saturation at high energies, has a probability that oscillates, which would modulate the spectra of astrophysical objects.

Galaxy clusters at X-ray energies sit at a sweet spot for photon-ALP physics. This is due to two key results. First, galaxy clusters are *particularly efficient* environments for photon-ALP interconversion. The electron densities are relatively low. Clusters have magnetic fields that are not significantly smaller than in galaxies, but in which the **B**-field extends over megaparsec scales, far greater than the tens of kiloparsecs applicable for galactic magnetic fields. The magnetic field coherence lengths in clusters are also larger than in galaxies, comfortably reaching tens of kiloparsecs. For massless ALPs, this feature singles out galaxy clusters as providing the most suitable environment in the universe for ALP-photon interconversion.[1]

The second key result is that, for the electron densities and magnetic field structures present within galaxy clusters, the photon-ALP conversion probability is energy-dependent, with a quasi-sinusoidal oscillatory structure at X-ray energies. This provides distinctive spectral features to search for.

The overall goal of this thesis is therefore to evaluate how X-ray observations of structures containing large magnetic fields can constrain the axion-photon coupling. With this in mind we have the following objectives:

1. Identify structures (galaxy clusters and radio galaxies) that have measurements of an extensive magnetic field that could produce significant conversion.

2. Model these magnetic fields using the best-fit parameters that have been measured so far, and thereby calculate the conversion probability of photons to ALPs in the X-ray regime for a given coupling.

3. Identify data sets from the publicly available archives for the latest generation of X-ray satellites.

---

[1] Although they appear appealing, magnetars and related objects do not provide efficient environments for ALP-photon conversion [71].

4. Reprocess and fit the data to models for the sources without ALPs. Use the goodness-of-fit to constrain the ALP-photon coupling.
5. Analyse any deviations from the best fit and discuss their possible systematic, astrophysical and new physics explanations.
6. Assess the experimental reach of future X-ray satellites to constrain this coupling further.

Chapter 3

Analysis of Galaxy Clusters

Galaxy clusters are the largest gravitationally bound structures in the Universe. Only about 1% of the total mass is contained in the component galaxies. A further 9% is accounted for by the intracluster medium (ICM), a magnetised plasma trapped within its gravitational well. The remaining 90% is dark matter. Galaxy clusters typically are hundreds of kiloparsecs in size. The ICM in a galaxy cluster can contain a magnetic field that also extends across similar distances. Typical temperatures are $\mathcal{O}(1 - 10\,\mathrm{keV})$, with X-ray emission occurring via thermal bremsstrahlung. Observations of atomic lines in this X-ray emission show that the ICMs have subsolar abundances of heavy ions. A subset of galaxy clusters contain a core $\mathcal{O}(20\,\mathrm{kpc})$ of dense gas with a cooling time less than the age of the universe. For a review of the properties of the ICM, see [115].

In this section, we outline the procedure we use to derive a bound on the ALP-photon coupling from X-ray data of point sources in galaxy clusters. This method is broadly similar to that described and used in [84].

3.1 Magnetic Fields in the Intracluster Medium

Below we discuss methods that constrain the properties of intracluster magnetic fields. The Zeeman effect can be used to measure directly the strength and orientation of magnetic fields:

$$\Delta\nu_0 = \frac{eB}{4\pi m_e c} = 2.8\,\frac{B}{\mu\mathrm{G}}\,\mathrm{Hz}, \tag{3.1}$$

meaning that an $\mathcal{O}(1\,\mu\mathrm{G})$ magnetic field causes only an $\mathcal{O}(1\,\mathrm{Hz})$ splitting in the frequency, far below the resolution of detectors.

Individual methods have degeneracies (for example between the relativistic electron density and the magnetic field strength), so often a combination of methods is required to evaluate the parameters.

3.1.1 Faraday Rotation Measures

We have already touched on the dispersion of a photon through a magnetic field in Subsection 1.4.1. The difference between the refractive indices perpendicular and parallel to the magnetic field induces rotation in the electromagnetic wave, which can be measured by polarimeters. Once the effects of the plasma are taken into account, the refractive indices are:

$$n_{L,R} \approx 1 - \frac{1}{2}\frac{\omega_p^2}{\omega(\omega \pm eB_\parallel/m_e)}, \tag{3.2}$$

for $\omega_p, eB_\parallel/m_e \ll \omega$ [116]. For two opposite-handed waves travelling a path length dl, this creates a phase difference:

$$d\phi = \omega dt = \omega \Delta n \, dl \approx \frac{4\pi e^3}{\omega^2 m_e^2} n_e \, B_\parallel \, dl, \tag{3.3}$$

where $B_\parallel$ is the magnetic field along the line of sight, and we have used $\omega_{pl}^2 = 4\pi\alpha n_e/m_e$. The intrinsic polarisation angle χ of the light after travelling a path length dl will change by an angle $d\chi = d\phi/2$. Writing in terms of the Faraday Rotation Measure:

$$RM = \frac{e^3}{2\pi m_e^2} \int_0^L n_e(l) B_\parallel(l) \, dl \ \ \mathrm{rad\,m^{-2}}, \tag{3.4}$$

the intrinsic polarisation of a source shining through a cluster of length L is:

$$\chi = \chi_0 + RM\,\lambda^2, \tag{3.5}$$

where λ is the wavelength. This allows one to probe the strength of the magnetic field parallel to the line of sight. The dependence on λ means that radio telescopes are best placed to look for the effects of Faraday rotation.

3.1.2 Synchrotron Radiation and Equipartition

Relativistic electrons propagating through a magnetic field will emit synchrotron radiation into a cone of half power width $\propto 1/\gamma$. The spectrum of synchrotron radiation from a single electron is sharply peaked around the critical frequency:

$$\nu_e = \frac{3}{4\pi}\frac{eB\sin\theta}{m_e c}\gamma^2\beta^2, \tag{3.6}$$

where θ is the angle between the electron's direction of motion and the magnetic field, and $\beta = v/c$ for electron velocity v. The emitted synchrotron power is:

$$P_{\text{syn}} = \frac{2e^4}{3m_e^2c^3}(B\sin\theta)^2\gamma^2\beta^2, \tag{3.7}$$

One can see that extracting the magnetic field strength B from an observed synchrotron emission energy spectrum requires knowledge of the energy spectrum of relativistic electrons in the intracluster medium. *Often* this is not well constrained theoretically or experimentally. The challenge is to separate the synchrotron emission from thermal emission. It also involves detecting hard X-rays, which will mean smaller data sets than radio waves.

We can estimate the strength of the magnetic field by assuming that the combined energy of relativistic particles (electrons and protons) and the magnetic field is minimised:

$$U_{tot} = U_B + U_{el} + U_{pr}, \tag{3.8}$$

where $U_B = B^2V/8\pi$ for a magnetic field filling a volume V and

$$U_{el} = V\int_{\epsilon_1}^{\epsilon_2} N(\epsilon)\,\epsilon\,\mathrm{d}\epsilon \tag{3.9}$$

is the energy density in relativistic electrons of energies between ϵ_1 and ϵ_2.

Assuming $U_{pr} = kU_{el}$ for some constant k, we can estimate the equipartition magnetic field strength. The parameter k depends on assumptions made about the creation of the relativistic protons and electrons, and in the standard analysis is taken to be 1.

3.1.3 Inverse Compton Scattering

An independent way to constrain B using synchrotron radiation is to compare it to Inverse Compton scattering. Highly energetic electrons can lose their energy by transferring it to photons through Inverse Compton Scattering, emitting photons with the frequency:

$$\nu_{\text{out}} = \frac{4}{3}\gamma^2\nu_{\text{in}}. \tag{3.10}$$

For electrons in the intracluster medium scattering off the Cosmic Microwave Background, which has a temperature of 3.1 K, will lead to hard X-rays being produced. Calculating the flux density $S_{IC}(\nu_x)$ and comparing it to the flux density of synchrotron radiation $S_{syn}(\nu_r)$ allows us to derive the magnetic field strength [117]:

$$B \propto \left(\frac{S_{syn}(\nu_r)}{S_{IC}(\nu_x)}\right)^{\frac{2}{g+1}} \left(\frac{\nu_r}{\nu_x}\right)^{\frac{g-1}{g+1}}, \tag{3.11}$$

where g is the power law index of the relativistic electron spectrum. The limitation of this method is separating the non-thermal IC emission from thermal emission in the hard X-ray regime.

3.1.4 Summary of galaxy cluster magnetic field measurements

For some of the clusters we will examine, information about the electron densities and coherence lengths will be experimentally inferred, while others will have to be inferred from similar clusters. In all cases, we will require there to be a confirmed measurement of an intrinsic magnetic field strength.

Coma Cluster

The Coma cluster has a well-measured set of magnetic field parameters. From Faraday rotation measures a central magnetic field value of $B_0 = 4.7^{+0.8}_{-0.7}\,\mu$G is inferred [118]. Its radial profile is $B(r) \propto n_e(r)^\eta$ with $\eta = 0.5^{+0.2}_{-0.1}$. The magnetic field power spectrum is well described by a Kolmogorov ($\eta = 5/3$) power spectrum with minimum/maximum scale 2 kpc/34 kpc. Its properties will be used to estimate those of other galaxies for which no observations appropriate to determine them have been performed. However this does not mean that the Coma cluster is the best target for ALP-photon conversion. Other clusters benefit from a higher inferred central magnetic field value, or a brighter point source shining through the cluster towards us. In all cases where the Coma cluster is used to infer the properties of other clusters, the value used will be that is most conservative for ALP-photon conversion while still remaining within the quoted error range.

Perseus Cluster

The Perseus galaxy cluster (A426) is the brightest X-ray cluster in the sky. It is a cool-core cluster at redshift $z = 0.0176$, centred around the Seyfert galaxy NGC 1275 and its Active Galactic Nucleus (AGN). Due to its proximity and brightness, the Perseus cluster has been a standard target for all X-ray satellites. The X-ray spectrum and emission specifically from the AGN are described in [119, 120, 121, 122].

The central magnetic field strength of the Perseus Cluster has been inferred using the Very Long Baseline Array [123]. This found a Rotation measure between 6500 and $7500\,\mathrm{rad\,m^{-2}}$. The authors assumed an average electron number density of $0.3\,\mathrm{cm^{-3}}$ and a path length of $2\,\mathrm{kpc}$ to derive a central magnetic field value of $25\,\mu\mathrm{G}$.

3.2 Modelling properties of the intracluster medium

In order to form a detailed description of the effects of ALP interactions with photons in galaxy clusters, we require theoretical models of the structure of the intracluster magnetic fields and electron densities. As the three-dimensional structure of these magnetic fields is not known, the precise configuration in each cluster, and hence the exact form of the energy-dependent ALP-photon conversion probability, can only be modelled statistically. We infer the properties of these statistical models from satellite data of these clusters. We have seen that the key parameters that determine ALP-photon conversion are: the magnetic field strength B, the photon energy ω, the coherence length of the magnetic field L and the electron number density n_e. The parameter most relevant to ALP-photon conversion is the central magnetic field strength B_0, with other parameters and the precise configuration being subdominant effects.

We account for our ignorance of the exact structure of the magnetic field by randomly generating many instances of the field from a given power spectrum. The magnetic field is divided into domains, within which the magnetic field strength and electron number density are constant (the values of $B(r)$ and $n_e(r)$ are evaluated at the centre of the domain). The orientation of the magnetic field is independently drawn from a flat distribution. This is done for two reasons. First, this allows for efficient computation of the conversion probability, which is vital for the thousands of different magnetic field configurations being modelled. The errors introduced compared to a continuous magnetic field model obeying Gauss' Law are small compared to the experimental uncertainties on the magnetic field strength [124]. Secondly, the assumption of completely random magnetic field directions is conservative for producing ALP-induced oscillations, as they will destructively interfere between domains. As we are seeking to constrain the ALP-photon coupling, we must ensure our assumptions are conservative to produce robust bounds.

Where there are direct estimates of the parameters for the galaxy cluster being studied, we use the best fit values. For the Perseus cluster, we separately calculate the bound for smaller magnetic field strength and coherence length, as an illustration of the effect of these parameters on the bounds. For galaxy clusters that do not have a direct estimate of a particular parameter, that value is inferred from a similar galaxy cluster, with the most conservative value taken from the quoted error range.

We assume the electron number density is spherically symmetric, depending only on the distance to the centre of the cluster r. In addition we assume that the magnetic field strength is proportional to the electron number density to some exponent $B(r) \propto n_e(r)^\eta$. We now illustrate our approach for the Perseus Cluster. The details for each galaxy cluster are shown in Chapter 4.

Direct measurements of magnetic fields of galaxy clusters are not always possible, as the necessary radio sources may not exist. The recent paper on the A194 magnetic field [125] summarises extant measurements of cluster magnetic fields (see its Table 8).

This thesis involves both sources that are at the centre of a cluster, and also sources that are significantly offset from the centre. ALP-photon constraints depend on the magnetic field along the line of sight from the source. For sources whose projected position is at a significant offset from the centre of the cluster, the field along the line of sight depends on *both* the overall central magnetic field of the cluster, denoted B_0, *and* the rate at which the field decreases away from the centre.

This is parametrised by assuming the cluster magnetic field to be radially symmetric,

$$B(r) \sim B_0 \left(\frac{n_e(r)}{n_0} \right)^\eta .$$

η is expected to lie somewhere between 0.5 and 1. The former represents an equipartition between magnetic field energy and themal energy ($\langle B^2 \rangle \sim n_e k_B T$), while a value $\eta = 2/3$ corresponds to a magnetic field that is frozen into the gas. A value of $\eta = 1$ has been found for the cool-core cluster Hydra A [126] (with a best-fit central magnetic field $B_0 = 36\ \mu$G). For the same value of B_0, a higher value of η will mean a more rapid drop-off of magnetic field strength with radius. We use an intermediate value of $\eta = 0.7$. A β-model for the electron density takes the form

$$n_e(r) = n_0 \left(1 + \frac{r^2}{r_c^2} \right)^{-3\beta/2} ,$$

where r_c is the core radius. Although not perfect, the β model captures the gross behaviour of the electron density in a cluster.

Perseus Cluster magnetic field

Based on results for the Coma cluster, we assume that B decreases with radius as $B \propto n_e^{0.7}$ [127]. The electron density n_e has the radial distribution found in [119],

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

We simulate each field realisation with 600 domains. The length l of each domain is between 3.5 and 10 kpc, randomly drawn from a power law distribution with minimum length 3.5 kpc and power 0.8. We therefore have:

$$P(l = x) = N \begin{cases} 0 & \text{for}\, x > 10\,\mathrm{kpc} \ , \\ x^{-1.2} & \text{for}\, 3.5\,\mathrm{kpc} < x < 10\,\mathrm{kpc} \ , \\ 0 & \text{for}\, x < 3.5\,\mathrm{kpc} \ , \end{cases} \tag{3.12}$$

with normalisation constant N.

The coherence length and power spectrum of the magnetic field in the centre of Perseus is not observationally determined. Instead, these parameters are motivated by those found for the cool core cluster A2199 [128], taking a conservative value for the magnetic field radial scaling.

3.3 Calculating ALP-photon conversion probability in clusters

We compute the energy-dependent conversion probability through our magnetic field model as described above. We choose an energy range that encompasses that of the detector (*Chandra's* ACIS, *XMM-Newton's* EPIC and *Athena's* X-IFU). For the majority of images, this is 0.1–10 keV. In cases where the data is contaminated by pile-up and we need to discard data at higher energies, we reduce the energy range of the simulation. From this energy range we select enough energies to ensure that the sampling is better than the energy resolution of the detector.

For the electron densities and magnetic field structures present within galaxy clusters, the photon-ALP conversion probability is energy-dependent, with a quasi-sinusoidal oscillatory structure at X-ray energies. We can see this by plugging in the relevant numbers into Equation 1.61. The amplitude of the oscillations, and their period, grow as the square of the energy, as is shown in Figure 3.1. The inefficiency of conversion at energies $E \lesssim 0.2$keV implies that effects of photon-ALP conversion are not visible in the optical (and below) range.

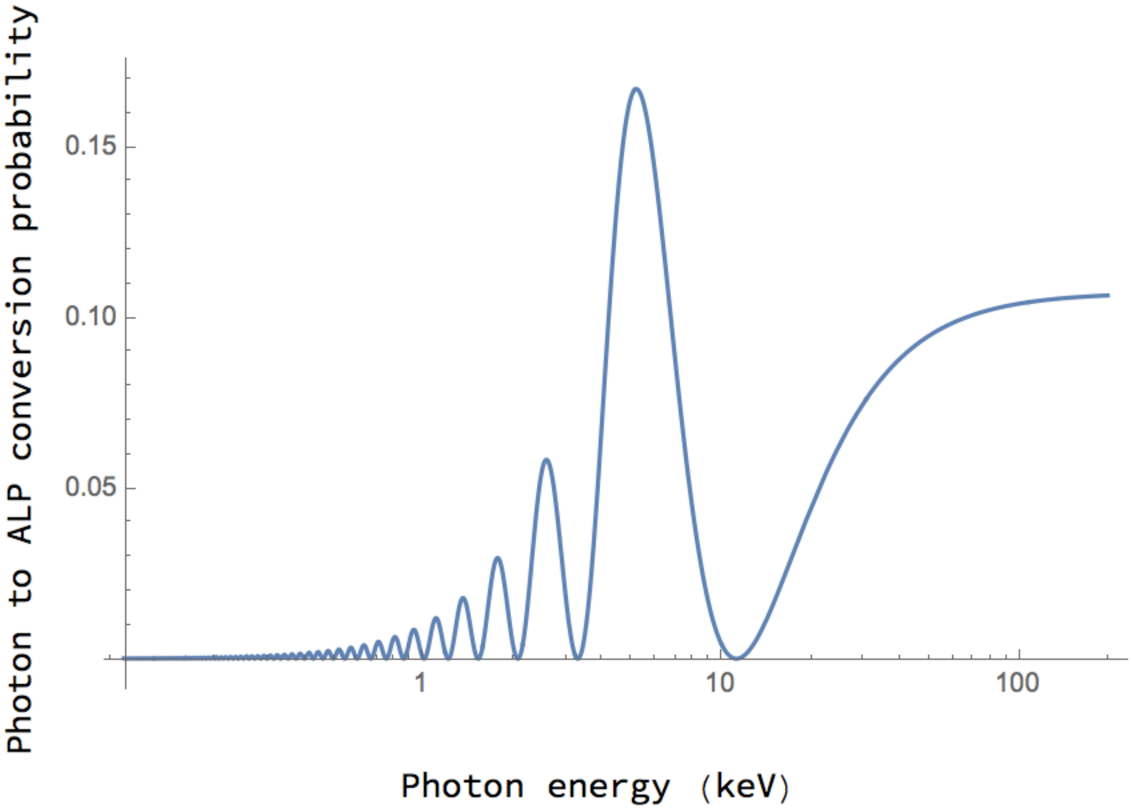

Figure 3.1: The photon survival conversion through a single-domain of a magnetic field due to ALP interconversion. The magnetic field strength is $1\,\mu$G, with an average electron number density of $10^{-3}\,\mathrm{cm}^{-3}$. The ALP-photon coupling is $g_{a\gamma\gamma} = 5 \times 10^{-12}\,\mathrm{GeV}^{-1}$, equal to the upper limit from SN1987A. One can see the oscillations growing with energy, being particularly pronounced in the $1 - 10\,$keV range, before the conversion probability saturates at higher energies. The domain length of 350 kpc is deliberately exaggerated over a typical domain length in a cluster magnetic field, in order to clearly show the oscillatory peaks. A more realistic magnetic field configuration is shown in Figure 3.2.

To calculate the conversion probability across the galaxy cluster magnetic field we use multiple domains, each of which has a conversion probability amplitude calculated. These amplitudes are then multiplied across all domains, and finally squared to produce a total conversion probability. We provide an example of the distinctive pattern of the features in Figure 3.2, where we plot a photon survival along a single line of sight modelled on that through the centre of the Perseus Cluster. However the location of the oscillations in the spectrum depends on the actual magnetic field structure along the line of sight. There is also an overall reduction in luminosity, but this can be absorbed into the overall normalisation of the spectrum. For unpolarised light the $\gamma \to a$ conversion probability cannot exceed fifty per cent, and in the limit of strong coupling saturates at an average value of $\langle P(\gamma \to a) \rangle = 1/3$ (for example, see [129]). It therefore follows that, expressed as a ratio of data to model, the maximal allowed range of ALP-induced modulations is approximately ±30%.[1]

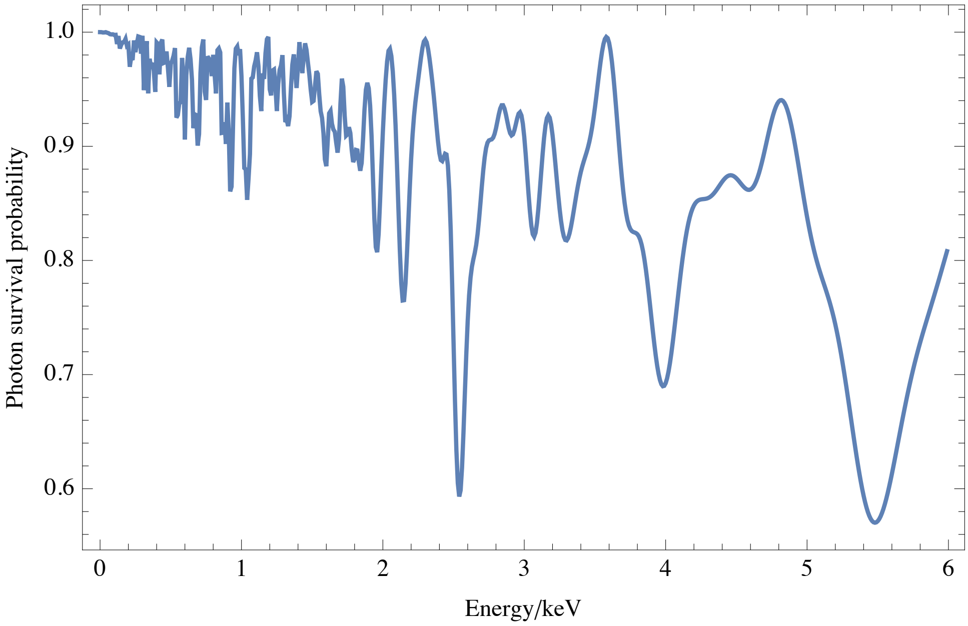

Figure 3.2: The photon survival probability through a randomly generated magnetic field due to ALP interconversion. A central magnetic field of $B_0 = 25\,\mu\text{G}$ was assumed, with a radial scaling of $\langle B(r) \rangle \sim n_e(r)^{0.7}$. There were 200 domains, with lengths randomly sampled from a Pareto distribution with range 3.5 – 10 kpc. The total propagation length was 1200 kpc. The ALP-photon coupling is $g_{a\gamma\gamma} = 1.5 \times 10^{-12}\,\text{GeV}^{-1}$ (roughly a factor of three beyond the current upper limit $g_{a\gamma\gamma} < 5 \times 10^{-12}\,\text{GeV}^{-1}$ from SN1987A). This quasi-sinusoidal structure arises generically in random field configurations.

3.4 Candidate target sources

Bright point-like sources either behind or embedded in a galaxy cluster are particularly attractive for searching for ALP-induced modulations. The galaxy cluster pro-

---

[1] We re-emphasise here that photon-ALP conversion involves quantum oscillations between states rather than absorption. Therefore for passage from $A \to C$ the survival probability $P(A \to C)$ does *not* equal $P(A \to B) \times P(B \to C)$.

vides a good environment for ALP-photon conversion; the bright point source ensures there are many photons, all passing along the same line of sight.

These factors make quasar or active galactic nucleus (AGN) spectra attractive for searching for ALPs. Emission from an AGN arises from matter accreting onto the central black hole. As evidenced by the rapid time variability of AGN luminosities, the physical region sourcing the X-ray AGN emission is tiny – of order a few Schwarschild radii of the central black hole. As cluster magnetic fields are ordered on kiloparsec scales, this implies that for all practical purposes every photon arising from the AGN has experienced an identical magnetic field structure during its passage to us.

To first approximation, at X-ray energies an AGN spectrum can be described as an absorbed power law. The effect of ALPs is then to imprint a quasi-sinusoidal modulation on this power law, of relative amplitude at most $\mathcal{O}(30\%)$ and with a modulation period of order a few hundred eV. As the fractional Poisson error on N counts is $\frac{1}{\sqrt{N}}$, and CCD detectors such as those on *Chandra* and *XMM-Newton* have intrinsic energy resolutions of around $\mathcal{O}(100\,\mathrm{eV})$, it therefore requires large numbers of counts to be able to distinguish any ALP-induced modulations from normal statistical fluctuations.

NGC 1275 (Perseus Cluster)

All the above facts make the AGN of the Seyfert galaxy NGC 1275 an excellent candidate for searching for ALP-photon interconversion. NGC 1275 is the central galaxy of the Perseus cluster, which as a cool core cluster should have a high central magnetic field (estimated as 25 μG in [123]) – implying the sightline from NGC 1275 to us should be efficient at ALP-photon conversion.

The nucleus is intrinsically bright and unobscured, with a spectrum that is well characterised by a power law and narrow Fe Kα line, absorbed by the galactic n_H column density [119]. Furthermore, there is enormous *Chandra* observation time on NGC 1275, encompassing 1.5 Ms in total. This results in over half a million photon counts originating from the central AGN, although for the on-axis observations quite a number are contaminated by pile-up. This is a level two orders of magnitude larger than the study in [84] involving Hydra A.

3.5 Data processing of X-ray satellites

X-ray satellites have a good chance to detect modulations in energy spectra resulting from ALP-photon conversion. Figure 3.3 shows that for galaxy cluster magnetic fields significant ALP-induced modulation occurs between 1–10 keV, an energy range that instruments on *Chandra*, *XMM-Newton* and *Athena* cover. Three features of these satellites are crucial to detecting these modulations:

1. **Energy resolution** The finite energy resolution of the detector will mean that rapid oscillations, as found at low energies, will be lost. This is shown in Figure 3.3, where the survival probability is convolved with a Gaussian distribution of 150 eV.

2. **Angular resolution** This must be good enough to resolve a single line of sight. If photons arriving on the same pixel have passed through different magnetic field configurations to each other, the oscillations will undergo destructive interference. In addition, for galaxy clusters that have significant X-ray emission from the intracluster medium, it is useful to be able to distinguish the AGN spectrum from the background.

3. **Effective Area** In order to distinguish the oscillations from Poisson fluctuations, we require sufficient photon statistics in the data set. A large effective area of the telescope will provide more photon counts for the same exposure time.

In this section we discuss the X-ray satellites currently in operation that are best suited to ALP searches. We review their limitations and describe the plans for future satellites that will have increased sensitivity to ALP induced oscillations.

3.5.1 *Chandra*

The Chandra X-ray Observatory was launched by *NASA* in 1999. It combines the High Resolution Mirror Array (HMRA), capable of excellent angular resolution, with the Advanced CCD Imaging Spectrometer (ACIS). The former consists of a nested set of four paraboloid-hyperboloid (Wolter-1) grazing-incidence X-ray mirror pairs, with the largest having a diameter of 1.2 m (taken from the Chandra Proposer's Guide[2]). The effective area is 600 cm^2 at 1.5 keV. The point-spread function (PSF) and the encircled energy fraction for a given radius depend upon off-axis angle and energy. The HRMA PSF has been simulated with numerical raytrace calculations: the increase in image size with off-axis angle is greatest for the inner shell, and hence is larger for higher X-ray energies.

---

[2] `http://cxc.harvard.edu/proposer/POG/html/chap6.html`

ACIS consists of two charged coupled device (CCD) arrays: a 2×2 chip array, ACIS-I; and a 1×6 chip array, ACIS-S. One ACIS-S chip provides an 8 arcminute by 8 arcminute field of view, while ACIS-I provides a 16 arcminute by 16 arcminute field of view. The CCD is a device composed primarily of silicon, with pixels separated by gates. Upon absorption of an X-ray photon a proportional number of electrons are liberated, which are then confined by electric fields. The array is exposed for a time that depends on the mode being used (in full frame this is $\sim 3.2\,\mathrm{s}$). The charge is then passed to a serial read-out at the end of a row. The energy resolution of the detector depends on the fraction of the charge collected, the charge transfer inefficiency from pixel to pixel in the read-out phase, the resolution of the read-out amplifiers, read noise and the off-chip analog processing electronics. Note that a single photon may create electron-hole pairs in multiple pixels. Pixels with a charge deposit above the threshold, plus the surrounding pixels, are "graded" based on the likelihood that the event was a photon absorption (instead of, say, a cosmic ray).

Figure 3.3 compares the photon survival probability shown in Figure 3.2 with the same probability convolved with a Gaussian with full width at half maximum (FWHM) of 150 eV, representing the approximate energy resolution of the CCD detectors present on *Chandra* and *XMM-Newton* satellites (the precise figure of 150eV is taken from the in-orbit performance of the ACIS-I detectors on *Chandra*, see table 6.4 of the *Chandra* proposer's guide[3]). While at lowest energies the oscillations are too rapid to be resolved by CCD detectors, and would require micro-calorimeters such as those that were present on *Hitomi*[4], in general it is fortuitous that the scales of the oscillations match those of the X-ray telescopes extensively used to observe galaxy clusters.

We use `CIAO 4.7` [130], `Sherpa` [131] and `HEASOFT 6.17` for the *Chandra* data analysis.[5] After the data is reprocessed using `CALDB 4.6.9`, it is cleaned from time periods that are polluted by flares using the program `chips`. We analyse the images and count rates of each observation using the image software `ds9`.

The boundaries of the ellipse are set by the location where the image of the point source ceases to dominate the background cluster emission. The background is taken from an elliptical annulus around the point source, with the inner ellipse slightly larger than for the source extraction region. Spectra and responses are created using `specextract` for each observation, and then stacked using `combine_spectra`. In

---

[3] `http://cxc.harvard.edu/proposer/POG/html/chap6.html`

[4] The energy resolution from *Hitomi* was around ~ 5 eV.

[5] The update to `CIAO 4.8` affects data taken in Continuous Clocking mode, which does not apply to these observations.

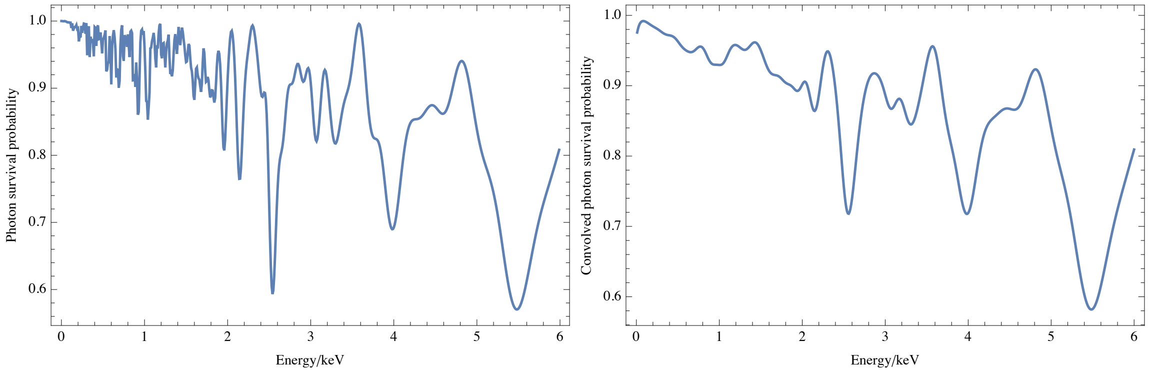

Figure 3.3: Left—The photon survival probability after passing through a randomly generated magnetic field. A central magnetic field of $B_0 = 25\,\mu$G was assumed, with a radial scaling of $\langle B(r) \rangle \sim n_e(r)^{0.7}$. There were 200 domains, with lengths randomly sampled from a Pareto distribution with range $3.5 - 10\,$kpc. The total propagation length was 1200 kpc. The ALP-photon coupling is $g_{a\gamma\gamma} = 1.5 \times 10^{-12}\,\text{GeV}^{-1}$ (roughly a factor of three beyond the current upper limit $g_{a\gamma\gamma} < 5 \times 10^{-12}\,\text{GeV}^{-1}$ from SN1987A). This quasi-sinusoidal structure arises generically in random field configurations. Right—The photon survival probability for the same magnetic field convolved with a Gaussian with FWHM of 150 eV.

general we bin so that there are approximately one hundred bins in total. If there are too few counts per bin, then the fit is insensitive to localised modulations as can be produced by ALPs, as the goodness-of-fit is insensitive to the sign of the residuals. We use `Sherpa`'s Levenberg-Marquardt fitting method with Poisson errors derived from the value of the data in each bin.

Pile-Up

Bright sources whose light shines on very few pixels of a CCD detector will have their inferred spectra significantly contaminated by pile-up. The energy recorded on the individual ACIS pixels (each approximately 0.5 arcseconds square) is read out approximately every three seconds (one frame time) when *Chandra* is operating in "full frame mode". Based on groupings of 3×3 pixels, events are graded. "Bad grades" are assigned to signals that are not likely to be the result of a photon interaction (an example of a non-photon interaction would be a cosmic ray), while "good grades" indicate the signal is likely to be a photon.

Pile-up refers to the arrival of more than one photon in this grouping within the same readout frame. This can lead to the energy of the two (or more) incident photons being summed, and either treated incorrectly as a single photon event of higher energy, or assigned a bad grade (grade migration). For an on-axis bright source (as in the ACIS-S observations of NGC 1275), the level of pile-up can be high, and the resulting spectrum contains events with two, three (and more) photons. As pile-up is a statistical feature of the number of arriving photons, some level of pile-up is inevitable in any observation. The question is always whether the magnitude of pile-up is sufficient to corrupt the science analysis being undertaken.

In terms of the measured photon distribution, the general effect of pile-up is to cause a hardening of the spectrum: two or more lower-energy photons are misidentified as a single higher-energy photon. This implies that for a fit of a single power law to a photon distribution, as pile-up increases the best-fit power law index will decrease. In a spectrum contaminated by pile-up, this makes it harder to determine the correct original power law index.

There are two basic methods we can employ to reduce the effects of pile-up on the analysis. The first is to reduce the amount of pile-up by using an annular extraction region and extracting the spectrum only from the wings of the point spread function. The point spread function of the telescope causes the arriving photons to be spread out on the detector, with the degree of spread increasing as one moves progressively off-axis. Furthermore, the point spread function is mildly energy-dependent: photons with higher energy are spread out further than photons with smaller energy. While fewer photons arrive in the wings of the point spread function, those that do suffer less from pile-up than those arriving on the central pixels. The second is to model the pile-up, and the distortion it induces on the spectrum, explicitly. We shall utilise both approaches below.

We set the `ds9` binning such that the `ds9` pixels are essentially the same size as the physical *Chandra* pixels (half an arc-second across). We then create an extraction region manually excluding all `ds9` pixels which are neighbours (either side-by-side or diagonally) to any pixel with total counts greater than 1% of the overall exposure time. Using the `pileup_map` tool we can estimate the pile-up fraction of each pixel.

The `jdpileup` model is described in [132]. This model is specific to modelling a source that is on-axis. The model assumes Poisson statistics to calculate the probability of different numbers of photons hitting an event-detection cell (a $3{\times}3$ pixel region in Chandra) within the read-out time adjusted to these respective observations. The model then convolves this probability with the probability of such events being assigned a ‘good’ grade, and the conversion from photon energy to pulse height by the detector. Obviously this is a non-linear process: the pile-up in any particular bin depends on the energy spectrum for all energies below the bin. The parameters of the model therefore need to be determined together with those for the spectrum under consideration. This can lead to degeneracies in parameter space, in particular for a simple power law.

As per [132] the two parameters of the pile-up model that we allow to vary are α and f, where α^{p-1} is the probability that p piled photons will be assigned a 'good' grade, and f is the fraction of events to which pile-up will be applied. The other parameter that will prove important is n, the number of regions to which `jdpileup` will be applied independently. For a point source this should be set to 1, as was done for the ACIS-S observations. For an extended source it should roughly correspond to the number of 3×3 pixel islands in the region. The reason for this is that `jdpileup` assumes spatial uniformity across the extraction region. For the ACIS-I midway observations, where the AGN is smeared out across several pixel islands, the value of n proved difficult to determine for large extraction regions, with the value corresponding to the best fit being unrelated to the number of pixel islands. We therefore constrained ourselves to a smaller central region, with little variation in count rate between pixels, to give us more control over the pile-up model.

It is also worth noting that the `jdpileup` model is set to zero for energies less than 0.5 keV and energies greater than 10 keV. As there are many events above the 10 keV range for the ACIS-S observations, we only model the spectrum up to 10 keV, as extending the fit beyond that would result in the model parameters being sent to unphysical values.

The overall balance here is between enhanced photon statistics – but with more pile-up and so with worse data quality – and fewer photons but better data. In particle physics language, this is the trade-off between efficiency and purity.

A more involved procedure to model the effects of pile-up is to perform MARX simulations of the data. Significantly piled-up data is known to produce complex effects, in particular with a bright background (see for instance [133]). We first simulate data without any axions present.

1. The thermal emission was simulated as a spatially extended Gaussian with a width of 30 arcminutes, centred on the AGN. It was modelled as an `apec` model with a temperature $T = 3.5\text{keV}$ and an abundance $Z = 0.48$. The normalisation of the model was adjusted so that it had the same fitted strength as for real data (as determined above) when extracted within the region used for the ACIS-I Edge data.

2. The AGN emission was simulated as a point source power law. The normalisation of the power law was adjusted so that - after it had been combined

with the thermal emission and the pileup processing applied - it had the same normalisation as the real data when extracted over the same extraction region.

3. The two individual simulations were then combined using `marxcat` and the MARX pileup processing applied to the combined simulation.

4. Counts were grouped to 500 and the combined spectrum was fitted with the sum of a power law plus thermal emission. The previous two steps were iterated until the fitted strength and index of the power law matched that of the real data.

We now repeat the MARX simulations including the effects of axions in the data. We multiply the AGN power law with a $P(\gamma \to \gamma)$ survival probability coming from photon-ALP conversion. As this leads to a net reduction in the number of photons present, we adjust the intrinsic normalisation upwards so that the fitted normalisation matches the actual data.

For a fixed ALP coupling, we generate 50 fake data samples. To compare these results to those described above using Sherpa's `fake_pha` command, we simulate three separate ALP couplings: $g_{a\gamma\gamma} = 1,\ 1.5,\ 2 \times 10^{-12}\text{GeV}^{-1}$. In all cases we assume the 'optimistic' magnetic field model.

The above analysis uses MARX's simulation of *Chandra*'s optics. The most advanced simulation of *Chandra*'s optics is through the *ChaRT* simulator ([134]), rather than MARX. As *ChaRT* requires the original source spectrum to be manually uploaded to a website, it is not possible to automate this process to produce axion bounds (as every different ALP conversion template represents a different source spectrum).

Errors in the Fitted Parameters

For fits of spectra to models, the paper includes the *fitted statistical errors.* However, significant caution should be applied when interpreting such a quantity (for example, a power law index) as an error on the intrinsic spectrum of the source. Systematic error on overall power law indices are likely to be much larger. While hard to quantify, there are several clear sources of such systematic effects:

1. Although the cleaned spectra have significantly reduced levels of pile-up, no spectrum can be entirely free of pile-up. Pile-up automatically leads to a hardening of spectral indices, as it moves photons from low to high energies.

2. The point spread function of *Chandra* is a function of energy, and more energetic photons tend to be spread out more. For spectra such as those used here, where

a central and highly piled-up core is excluded, this will bias the analysed photons to higher energy. While the *Chandra* analysis software aims to take this into account, this correction will not be perfect, particularly for off-axis sources where the form of the optical image is more complex (as for the ACIS-I Midway data, where the AGN image takes on a 'Maltese cross' form).

3. Generally, there are systematic errors on the overall power law index that will arise because the optical conditions are significantly different on-axis and off-axis. These are hard to quantify, but are certainly much larger than the purely statistical errors on fitted parameters.

4. The AGN luminosity is time-variable and the power law index will be time-variable as well. For observations taken at different times, the intrinsic power law of the source may be different.

3.5.2 *XMM-Newton*

The X-ray Multi-Mirror Mission (XMM-Newton) was launched by ESA in 1999. The optics are three X-ray telescopes, each consisting of 58 Wolter I grazing-incidence gold-coated mirrors. They have an effective area of 1500 cm^2 at 2 keV. There is also an effective area edge just above 2 keV due to the gold coating.

On board XMM Newton is the European Photon Imaging Camera (EPIC) consisting of three CCD cameras. Two of these are part of the MOS camera, comprising Metallic Oxide Semiconductor cameras. The third uses pn CCDs, and so is referred to as the pn camera. MOS receives photons passing through reflection gratings, while pn has an unobstructed view. The detector can be set up in full frame mode, with a frame time of 2.6 s for MOS and 73.4 ms for pn. In addition pn can be run in extended frame mode with a frame time of 199.1 ms.

The pixel size of pn is 4.1 arcseconds and 1.1 arcseconds for MOS. This means that pn in extended full frame mode is more susceptible to pile-up than MOS. In general, pile-up is an issue for all *XMM-Newton* observations of NGC 1275 as we will discuss in the *XMM-Newton* analysis.

A significant difference between *XMM-Newton* and *Chandra* is the angular resolution. For both MOS and pn, the radius of the disk containing 50 % of the photons collected in the focal plane (Half Energy Width) is around 8.5 arcseconds at 1.5 keV, while for *Chandra* it is much smaller at ~ 0.5 arcseconds (FWHM). The FWHM for MOS-1 in orbit is 4.3 arcseconds, for MOS-2 is 4.4 arcseconds, and for pn is 12.5 arcseconds due to the large pixel size. As the central region of the Perseus cluster is also intrinsically bright, this makes it harder to separate AGN and cluster emission for the *XMM-Newton* observations.

The effective area at 1 keV is $922\,\mathrm{cm}^2$ for MOS and $1227\,\mathrm{cm}^2$ for pn compared to 340 cm^2 for *Chandra*. This allows MOS and pn combined to collect roughly 7 times more photons in a given observation time than *Chandra* (although this also increases the amount of pile-up).

Analysis

We use `SAS` version 15.0.0. As data has been taken in the past, we reprocess it with up to date calibrations via `cifbuild`. To remove flares, we apply the standard filters of counts/s < 0.35 for MOS and counts/s < 0.4 for pn. After selecting the extraction region for the point source we check if pile-up is an issue using the `SAS` tool `epatplot`.

After applying `eveselect` to generate the spectra, we use `rmfgen` and `arfgen` to generate the response files. We use `epicspeccombine` to combine different observations. We compare this to using the ftools routines `mathpha`, `addrmf` and `addarf` (as `epicspeccombine` does not produce an arf file). We complete our analysis by fitting the spectra using the programs `Xspec` [135] and `Sherpa` [131].

3.5.3 The future satellite *Athena*

Searches for these oscillations can be used to constrain ALP parameter space. Current constraints on ALPs derived in this fashion [84, 1, 136, 2] are based on data taken with CCD detectors, which have an energy resolution of $\mathcal{O}(100\text{ eV})$. A large improvement with sensitivity will be achieved once data becomes available from microcalorimeters with $\mathcal{O}(\text{a few eV})$ energy resolution. Such microcalorimeters will be on board the Advanced Telescope for High ENergy Astrophysics (ATHENA), currently scheduled to launch in 2028. Its X-IFU instrument will have a large effective area, good imaging and an energy resolution of ~ 2.5eV, greatly enhancing the discovery potential for ALPs.

We estimate the experimental sensitivity of *Athena* to ALPs. We do so using simulated data for a mock observation of NGC 1275, hosting the radio source 3C 84, which contains the central AGN of the Perseus cluster. This allows us to compare directly the capabilities of *Athena* to the present-day satellites discussed above.

One major limiting constraint on existing data is the energy resolution of the detectors. If they exist, ALPs provide oscillatory structure all the way down to the lowest energies. However, as illustrated in Figure 3.4, detectors with energy

	Athena (X-IFU)	*Chandra* (ACIS-I)
Energy range	0.2–12 keV	0.3–10 keV
Energy resolution at 6 keV	2.5 eV	150 eV
Spatial resolution	5 arcsec	0.5 arcsec
Time resolution	10 μs	0.2 s (2.8 ms single row)
Effective area	2 m^2 @ 1 keV	600 cm^2 @ 1.5 keV

Table 3.1: Parameters taken from the *Athena* Mission Proposal and the *Chandra* Proposer's Guide.

resolutions of $\mathcal{O}(100 \text{ eV})$ cannot resolve this structure at lower energies – but this does become accessible once a resolution of $\mathcal{O}(2.5 \text{ eV})$ is achieved. We now discuss the future *Athena* X-ray observatory, whose greatly enhanced technical capabilities offer improved sensitivity to ALP-photon interconversion.

Athena is an ESA mission to explore the Hot and Energetic Universe, due to launch in 2028 [137]. The mirror will have a 2 m^2 effective area and a 5 arcsec angular resolution. There are two instruments: the X-ray Integral Field Unit (X-IFU) and the Wide Field Imager (WFI). Here we focus on the former, which will consist of an array of TiAu Transition Edge Sensor (TES) thermistors. An efficient photon absorber is attached to a film of TiAu superconductor. The array is cooled to a temperature of 50 mK, keeping the TiAu in equilibrium between the normal and superconducting state. The transition between the states is very sharp, meaning that a small deposit of energy can effect an excitation. The increase in resistance caused by an increase in heat is then measured. The energy resolution of a single micro-calorimeter pixel is predicted to be 2 eV.

Once noise from read-out, fluctuations in the environment and gain miscalibration is taken into account, X-IFU is planned to achieve an energy resolution of 2.5 eV below 7 keV [138], meaning it will be able to resolve narrow spectral oscillations. X-IFU will be sensitive to the energy range 0.2–12 keV [139]. A readout time of ~ 10 μs will ensure pileup contamination is minimised. Table 3.1 contains a summary of its properties, taken from the *Athena* Mission Proposal[6], compared to properties of the *Chandra* ACIS-I detector, taken from the *Chandra* Proposer's Guide[7].

The combination of larger effective area, greatly improved energy resolution and

---

[6] http://www.the-athena-x-ray-observatory.eu/images/AthenaPapers/The_Athena_Mission_Proposal.pdf

[7] http://cxc.harvard.edu/proposer/POG/html/chap6.html

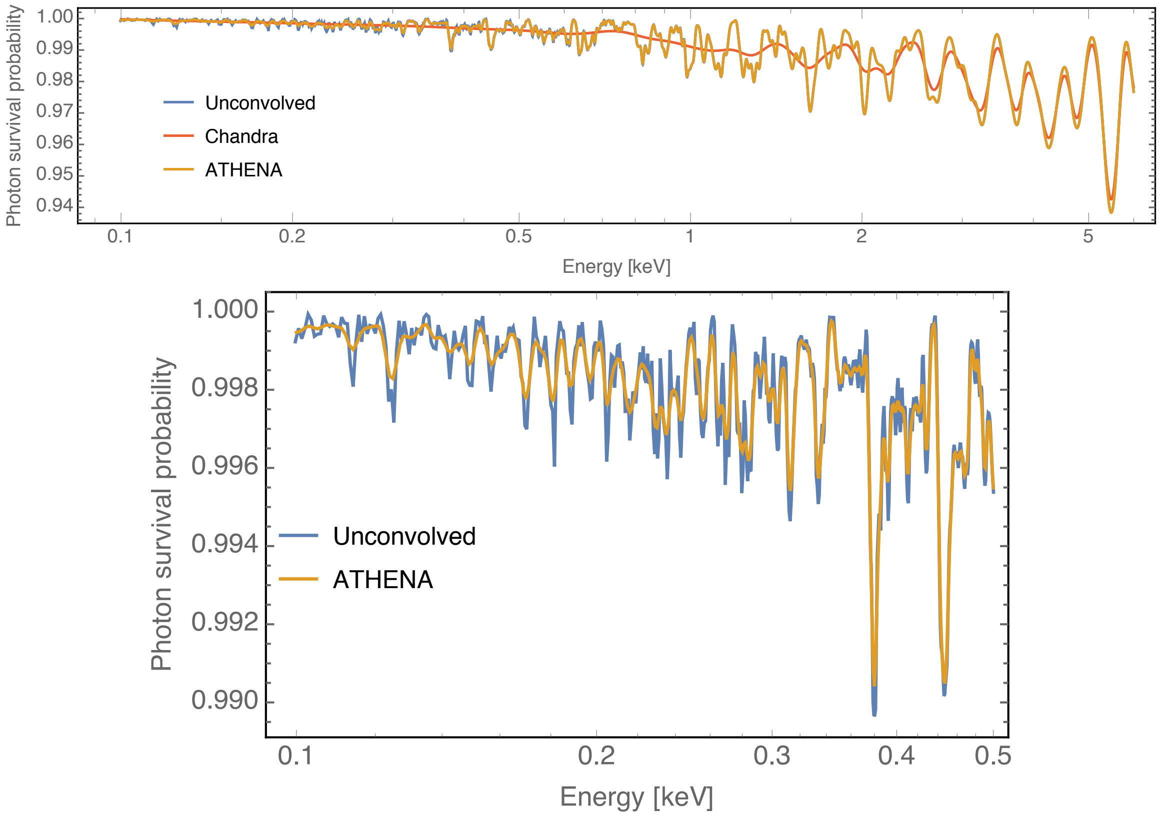

Figure 3.4: Above—A randomly generated photon survival probability along the line of sight from NGC 1275 to us: unconvolved (blue), convolved with a Gaussian with FWHM 150 eV (a typical energy resolution of *Chandra*'s ACIS-I detector (red)) and 2.5 eV for *Athena*'s X-IFU detector (orange). A central magnetic field of $B_0 = 25\,\mu$G was used, with a radial scaling of $B \sim n_e^{0.7}$, further details in Section 3.2. The ALP-photon coupling is $g_{a\gamma\gamma} = 5 \times 10^{-13}\,\text{GeV}^{-1}$. Small, rapid oscillations at low energies, and larger oscillations at high energies, are generic features of these survival probabilities. At energies $< 2\,$keV *Chandra* is unable to resolve oscillations while *Athena* performs much better. Left—The same photon survival probabilities, showing the sensitivity of X-IFU to oscillations at low energies.

reduced pileup contamination means *Athena* has far more potential to detect ALP-induced oscillations than the best current satellites. The aim of this paper is to make the first quantitative estimate of the extent to which *Athena* will be able to improve constraints on $g_{a\gamma\gamma}$.

3.6 Procedure to constrain the ALP-photon coupling

For convenience we will restrict our analysis to massless ALPs. As the efficiency of ALP-photon conversion is proportional to $1/|m_a^2 - \omega_{pl}^2|^{-2}$, the results will hold for all $m_a^2 \ll \omega_{pl}^2$. This is the case for all galaxy clusters we study when the ALP mass $m_a \lesssim 10^{-12}$eV.

To obtain approximate bounds on $g_{a\gamma\gamma}$, we compare two models for the flux $F(E)$ observed from NGC 1275:

- Model 0: An absorbed power law $F_0(E) = AE^{-\gamma} \times e^{-n_H\sigma(E)}$, as described in Equation (4.1).
- Model 1: An absorbed power law multiplied by the photon survival probability assuming the existence of ALPs with coupling $g_{a\gamma\gamma}$. In this case the predicted flux also depends upon the magnetic field $\mathbf{B}$ along the line of sight. We have $F_1(E, \mathbf{B}) = AE^{-\gamma} \times e^{-n_H\sigma(E)} \times P_{\gamma\to\gamma}(E(1+z), \mathbf{B}, g_{a\gamma\gamma})$.

We compute 95% confidence limits on $g_{a\gamma\gamma}$ by generating fake data from Model 1 and assessing how well it is fit by Model 0 i.e. how well the oscillations due to ALP-photon conversion can hide in the Poisson noise. As it has minimal pile-up, we use the clean ACIS-I edge observations for this analysis. We fit the spectrum between 1 and 4 keV (a region unaffected by pile-up) and bin such that there are 1000 counts in each energy bin. Our procedure to determine whether ALPs with coupling $g_{a\gamma\gamma}$ are excluded at the 95% confidence level is as follows:

1. Fit Model 0 to the real data and find the corresponding reduced χ^2, χ^2_{data}.
2. Randomly generate 50 different magnetic field realisations $\mathbf{B}_i$ for the line of sight to the point source.
3. For each $\mathbf{B}_i$, compute $P_{\gamma\to\gamma}(E, \mathbf{B}_i, g_{a\gamma\gamma})$ by numerically propagating photons at different energies through $\mathbf{B}_i$, as described for example in [87]. We take 300 photon energies equally spaced between 1 and 4 keV.

4. For each $\mathbf{B}_i$, generate 10 fake data sets from Model 1, using `Sherpa`'s `fake_pha` method.

5. Fit Model 0 to each of the 500 fake data sets and find the corresponding reduced χ^2, χ_i^2 for each.

6. If fewer than 5% of the χ_i^2 are lower than χ^2_{data}, $g_{a\gamma\gamma}$ is excluded at the 95% confidence level.

We scan over $g_{a\gamma\gamma}$ in steps of $10^{-13}\,\text{GeV}^{-1}$. For the $g_{a\gamma\gamma}$ value excluded, we also check that the three values above it in our grid are also excluded. For the magnetic field parameters described above, we find $g_{a\gamma\gamma} \lesssim 1.4 \times 10^{-12}\,\text{GeV}^{-1}$ which is shown in Figure 4.18.

If we consider a more pessimistic scenario with $B_0 = 15\,\mu\text{G}$ and a minimum coherence length of 0.7 kpc, we instead find $g_{a\gamma\gamma} \lesssim 2.7 \times 10^{-12}\,\text{GeV}^{-1}$. If we take an even more pessimistic scenario in which the central field is $B_0 = 10\,\mu\text{G}$ and the minimum coherence length is 0.7 kpc, the bound increases further to $g_{a\gamma\gamma} \lesssim 4.0 \times 10^{-12}\,\text{GeV}^{-1}$.

3.7 Simulating bounds for future satellites

We use the following procedure to determine whether a particular value of $g_{a\gamma\gamma}$ is excluded: we varied the ALP-photon coupling $g_{a\gamma\gamma}$ from $g_{a\gamma\gamma} = 5 \times 10^{-13}\,\text{GeV}^{-1}$ to $g_{a\gamma\gamma} = 1 \times 10^{-13}\,\text{GeV}^{-1}$, with stepsize $0.5 \times 10^{-13}\,\text{GeV}^{-1}$. As the bound is dependent on uncertainties in the magnetic field strength of a factor of 2, and we are only using simulated data, we do not consider step sizes smaller than this. For each $g_{a\gamma\gamma}$:

1. Generate 50 configurations of the magnetic field B_i.

2. Use the B_i to calculate the survival probability $P_{\gamma\to\gamma}$ along the line of sight for different photon energies (as done in [87]). We calculate for 8000 equally spaced photon energies in the range 0.01–10 keV.

3. Combine each $P_{\gamma\to\gamma}$ with the AGN spectrum.

4. Generate 10 fake PHAs for each spectrum, providing 500 fake data samples in total.

5. Fit the fake data to Model 0, and calculate the reduced chi-squareds χ_1^2.

6. Generate 100 fake PHAs based on Model 0, and compute the average of their reduced chi-squareds χ_0^2. Assuming the absence of ALPs, this represents the expected quality of the fit to the single real data set. If the actual data is a poor fit for some reason, then this will weaken the level of the resulting bounds that we can produce.

7. Determine the percentage of fake data sets that have a reduced chi-squared $\chi_1^2 < \max(\langle\chi_0^2\rangle, 1)$. If this is true for fewer than 5 per cent of the data sets, the value of $g_{a\gamma\gamma}$ is excluded at 95 per cent confidence.

Chapter 4

Results

4.1 Point sources in/behind Galaxy Clusters

We searched for appropriate sources by examining the *Chandra* archive, using a combination of manual inspection of images and SIMBAD [140] to locate bright point sources shining through galaxy clusters. We used the *Chandra* archive as its superior angular resolution makes it the best telescope to distinguish point sources from cluster emission. When candidate sources were found in the *Chandra* archive, we checked whether the *XMM-Newton* archive contained observations of the same source.

We focused on nearby clusters as they are more likely to have a distant source behind them. This also allows, for the cases of sources embedded in the cluster, a greater discrimination of the point source from the contiguous diffuse emission.

4.1.1 Summary of Point Sources Selected

We list the point sources included within this paper. We exclude previously studied sources (Hydra A, NGC 1275 and M87). We assume a cosmology with $H_0 = 73\,\mathrm{km\,s^{-1}}$, and source redshifts have been identified using either SIMBAD [140] or NED[1]. The point sources are:

1. The Sy1.5 galaxy NGC 1275 at the centre of the Perseus Cluster has been observed for 1 Ms by *Chandra* ACIS-S, with observations taken in 2002 and 2004, and also for a further 0.5 Ms by *Chandra* ACIS-I in 2009. There are also 180 ks of observation time with *XMM-Newton* taken in 2001 and 2006.

2. The bright Sy1 galaxy 2E3140 at z=0.05893 within A1795 located at (RA,DEC) = (13:48:35, +26:31:09) and offset from the centre of the cluster by 417$''$ (a

---

[1] http://ned.ipac.caltech.edu/

projected distance of 456 kpc). In total, this is the subject of 660 ks of *Chandra* time.

3. The quasar CXOU J134905.8+263752 at $z = 1.30$ behind A1795, at (RA, DEC) = (13:49:06, 26:37:48) and offset from the cluster centre by $177''$ at a projected distance of 194 kpc, and the subject of 985 ks of *Chandra* time.
4. The central AGN UGC9799 of the cluster A2052 (redshift z=0.0348) at (RA, DEC) = (15:16:45, +07:01:18), with 654 ks of *Chandra* time.
5. The quasar B1256+281 at redshift z=0.38 at (RA,DEC) = (12:59:17,+27:53:46) shining through the Coma cluster (z=0.023). This is offset by $521''$ from the cluster centre, at a projected distance of 232 kpc, and there is 493 ks of *Chandra* time containing it.
6. The quasar SDSS J130001.48+275120.6 at redshift $z = 0.975$ at (RA,DEC) = (13:00:01, +27:51:20) shining through the Coma cluster (z=0.023), offset by $484''$ from the cluster centre at a projected distance of 215 kpc, and with 574 ks of *Chandra* time on it.
7. The bright AGN, NGC3862, within the cluster A1367 (z=0.0216) and at an offset of $443''$ from the cluster centre (a projected distance of 186 kpc), at (RA,DEC) = (11:45:05, +19:36:23), and the subject of 75 ks of public *Chandra* time.
8. The central AGN of the Hydra A galaxy cluster (z=0.052) at (RA, DEC) = (09:18:05, -12:05:44). This has been analysed in [112] and we include it here for completeness. It has 240 ks of *Chandra* observation time.
9. The central AGN IC4374 of the cluster A3581 (z=0.023), at (RA, DEC) = (14:07:29, -27:01:04), with 85 ks of *Chandra* time.
10. The central LINER-type AGN of the Virgo cluster, at the heart of the galaxy M87, with redshift=0.004 and coordinates (RA, DEC) = (12:30:49, +12:23:28). This shines through the Virgo cluster towards us.

A summary of the sources used can be found in Table 4.1.

Source	(RA,DEC)	Cluster (redshift)	Source redshift	Offset arcsec|kpc	Chandra exposure (ks)
NGC 1275	(03:19:48.1, +41:30:42)	Perseus 0.0176	0.0176	0 0	1500
B1256+281	(12:59:17, +27:53:46)	Coma 0.023	0.38	521 232	493
SDSS J130001.48+275120.6	(13:00:01, +27:51:20)	Coma 0.023	0.975	484 215	574
NGC 3862	(11:45:05, +19:36:23)	A1367 0.0216	0.0216	443 186	75
IC 4374	(14:07:29, -27:01:04)	A3581 0.023	0.022	0 0	85
2E3140	(13:48:35, +26:31:09)	A1795 0.062	0.059	417 456	660
CXOU J134905.8+263752	(13:49:06, 26:37:48)	A1795 0.062	1.30	177 194	985
UGC9799	(15:16:45, +07:01:18)	A2052 0.0348	0.0345	0 0	654

Table 4.1: An enumeration of the sources used.

4.2 NGC 1275

4.2.1 Chandra observations

The deep Chandra observations involving NGC 1275 can be divided into three main groups. The first involves 200 ks of ACIS-S observations taken in 2002 together with 800 ks of ACIS-S observations taken in 2004. In these observations, NGC 1275 is close to the aimpoint. The second group involves 300ks of ACIS-I observations carried out in 2009, where NGC 1275 is approximately midway between the edge of the chips and the aimpoint. The third group involves 200 ks of ACIS-I observations also taken in 2009, in which NGC 1275 is close to the edge of one of the chips, around 8 arcminutes from the aimpoint. Finally, there are also some brief pre-2002 observations that we do not include.

This separation of the observations is necessitated by the growth of the point spread function away from on-axis. In the first group, the photons from the AGN suffer little dispersion and are highly concentrated on a few pixels. In the third group, the arriving photons are scattered over many pixels, whereas the second group is intermediate. This is illustrated in Figure 4.1, which shows images of NGC 1275

Obs ID	Exposure [ks]	Year	Instrument	Location of NGC 1275
3209	95.77	2002	ACIS-S	Central
3404	5.31	2002	ACIS-S	Central
4289	95.41	2002	ACIS-S	Central
4946	23.66	2004	ACIS-S	Central
4947	29.79	2004	ACIS-S	Central
6139	56.43	2004	ACIS-S	Central
6145	85	2004	ACIS-S	Central
4948	118.61	2004	ACIS-S	Central
4949	29.38	2004	ACIS-S	Central
6146	47.13	2004	ACIS-S	Central
4950	96.92	2004	ACIS-S	Central
4951	96.12	2004	ACIS-S	Central
4952	164.24	2004	ACIS-S	Central
4953	30.08	2004	ACIS-S	Central
11715	73.36	2009	ACIS-I	Midway
11716	39.64	2009	ACIS-I	Midway
12037	84.63	2009	ACIS-I	Midway
11714	91.99	2009	ACIS-I	Midway
11713	112.24	2009	ACIS-I	Edge
12025	17.93	2009	ACIS-I	Edge
12033	18.89	2009	ACIS-I	Edge
12036	47.92	2009	ACIS-I	Edge

Table 4.2: The Chandra observations used. The last column shows the location of NGC 1275 in the respective observation relative to the focal point.

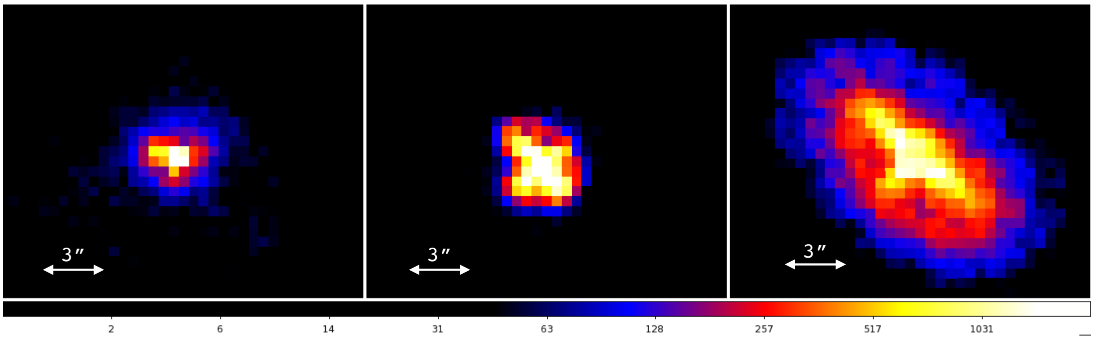

Figure 4.1: NGC 1275 in three types of observation, from left to right: centre of the chip in 2004 (*Chandra* ObsID 4952), midway between the edge and centre of the chip in 2009 (ObsID 11714), and the edge of the chip in 2009 (ObsID 11713). The colour coding is adjusted to account for the different observation times such that each colour corresponds to the same count-rate across images.

for each of the different observation types.

The observations were taken in "full frame" mode, with a readout time of 3.1 s.

We see from these observations that the brightness of the AGN varies substantially with time. As described in [122], the NGC 1275 AGN was brightest from 1970 to 1990, before rapidly declining by an order of magnitude until around 2000. Since then, it appears to have brightened significantly over the decade from 2003 onwards, although it has not yet returned to the luminosities it had pre-1990.

4.2.2 Chandra Analysis

After the data is reprocessed, it is cleaned from time periods that are polluted by flares using the program `chips`. We find that only Obs ID 4950 is affected by flares and the cleaning reduces the observation time slightly from 96.12 ks to 89.23 ks for this observation.

Spectral Analysis

We first extract the spectrum for NGC 1275 for the ACIS-I edge observations (11713, 12025, 12033, 12036), where the AGN is around seven arcminutes off the optical axis, using the full extraction region without any exclusion of the central core. This is only valid for the edge observations where pile-up is relatively low throughout the whole image. Spectra and responses were created using `specextract` for each observation, and then stacked using `combine_spectra`. An ellipse around NGC 1275 of radii 11.6 and 7.2 arcseconds was used for the extraction region. The boundaries of the ellipse were set by the location where the image of the AGN ceased to dominate the background cluster emission. The background was taken from an elliptical annulus around NGC 1275, with the outer radii being 19 and 13 arcseconds and the inner radii 13.3 and 9.3 arcseconds (for the two short observations 12025 and 12033 this region goes beyond the edge of the chip, and a rectangular box was used instead for the background).

The resulting stacked spectrum contains around 266000 counts, reducing to 230000 after background subtraction, giving a ratio of 6.5:1 for the AGN against the cluster emission. The resulting spectrum was binned to ensure a minimum of 2000 counts a bin, and fitted between 0.8 and 5 keV with an absorbed power law `xswabs` × `powlaw1d`,

$$AE^{-\gamma} \times e^{-n_H \sigma(E)} \,. \tag{4.1}$$

Here A denotes the normalisation of the power law, γ the power law index, and n_H the effective Hydrogen column density.

The resulting fit is shown in Figure 4.2, together with the fractional ratio of data to model. The best-fit value of n_H is $2.3 \times 10^{21}\,\mathrm{cm}^{-2}$ and the power law index is $\gamma = 1.83 \pm 0.01$. While the absorbed power law is a reasonable characterisation of the data, there are two large localised residuals: one positive between 2–2.2 keV and one negative around 3.4–3.6 keV. There is an upward trend at 5 keV. As the effective area of *Chandra* begins to fall off rapidly here, and there are also intrinsically fewer photons expected, pile-up plays a proportionately more important role. This rising trend continues beyond 5 keV and we attribute this to the effects of pile-up.

For the 300 ks of ACIS-I Midway observations with the AGN around 3 arcminutes from the aimpoint, a comparable full extraction results in a spectrum that is still highly piled up (the best-fit index is an unphysically hard $\gamma = 1.30$). This is even more the case for the ACIS-S observations, where a similar full extraction produces an index of $\gamma = 0.66$. For these cases, it is necessary to account for pile-up to extract physical results.

We next plot the complete spectra for the ACIS-I observations in which NGC 1275 is midway on the chip (11714, 11715, 11716, 12037). As for each observation NGC 1275 is in different locations relative to the aimpoint, the optical distortions differ for each case and customised extraction regions were used. For 11714, this was a circle of radius 3.5 arcseconds. For 11715, this was an ellipse of radii 4.6 and 6.4 arcseconds. For 11716, an ellipse of radii 4.4 and 5.7 arcseconds was used, while for 12037 an ellipse of radii 4.1 and 6.4 arcseconds was used. In all cases the background was taken from a circular annulus around NGC 1275 of radii 8.5 and 20 arcseconds.

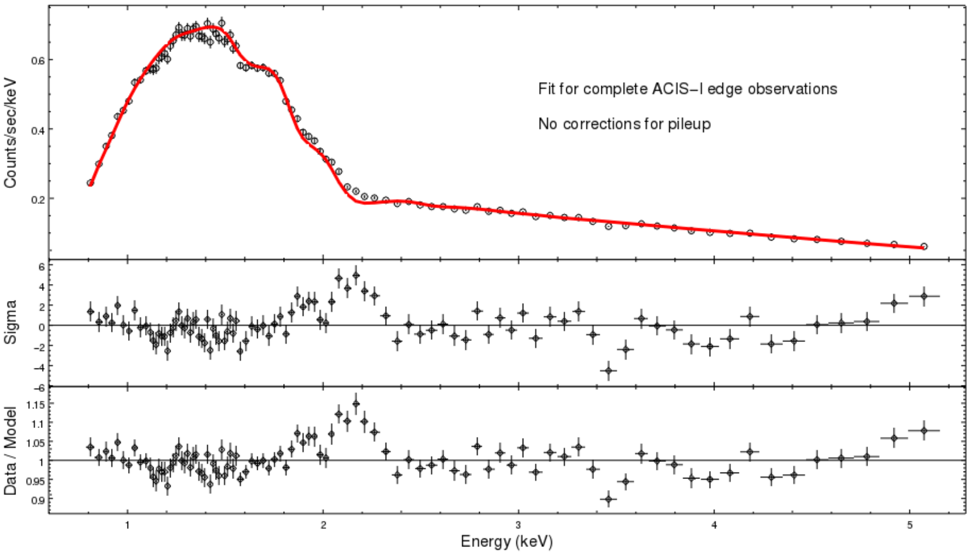

Figure 4.2: The stacked spectrum of the ACIS-I edge observations from the entirety of the extraction region, involving 230 000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 2.3 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index of $\gamma = 1.83$. σ refers to the standard deviation from the model expectation for a Poissonian count rate.

The resulting spectrum contains around 259000 counts, reducing to 242000 after background subtraction, giving an AGN to cluster ratio of 14.2:1. After background subtraction, the fraction of counts in the 7-10 keV band and 10-15 keV bands are 3.4% and 0.6% respectively. The resulting spectrum was binned to ensure a minimum of 2400 counts per bin, and fitted between 0.8 and 5 keV with an absorbed power law.

The resulting fit and the ratio of data to model are shown in Figure 4.3. The best-fit parameters are $n_H = 1.0 \times 10^{21}\,\mathrm{cm}^{-2}$ and spectral index $\gamma = 1.30$. The overall bad fit precludes extracting formal 1-sigma confidence intervals for the fit parameters, and we do not quote these for either this fit or the full ACIS-S data. There is a large overall modulation in the data to model ratio that arises from pile-up contamination, but superimposed on it there is a clearly visible large localised positive residual just below 2.2 keV. The data to model ratio for this residual is entirely consistent with the same feature in the ACIS-I edge data.

Finally, we extract spectra for the ACIS-S observations in which NGC 1275 is centrally located. A circle of radius 3.4 arcseconds is used for the extraction region, and the background is taken from an annulus of inner and outer radii 5 and 12 arcseconds. The resulting stacked spectrum contains around 233000 counts, reducing to 183000 after background subtraction, giving a ratio of 3.7:1 of AGN to cluster

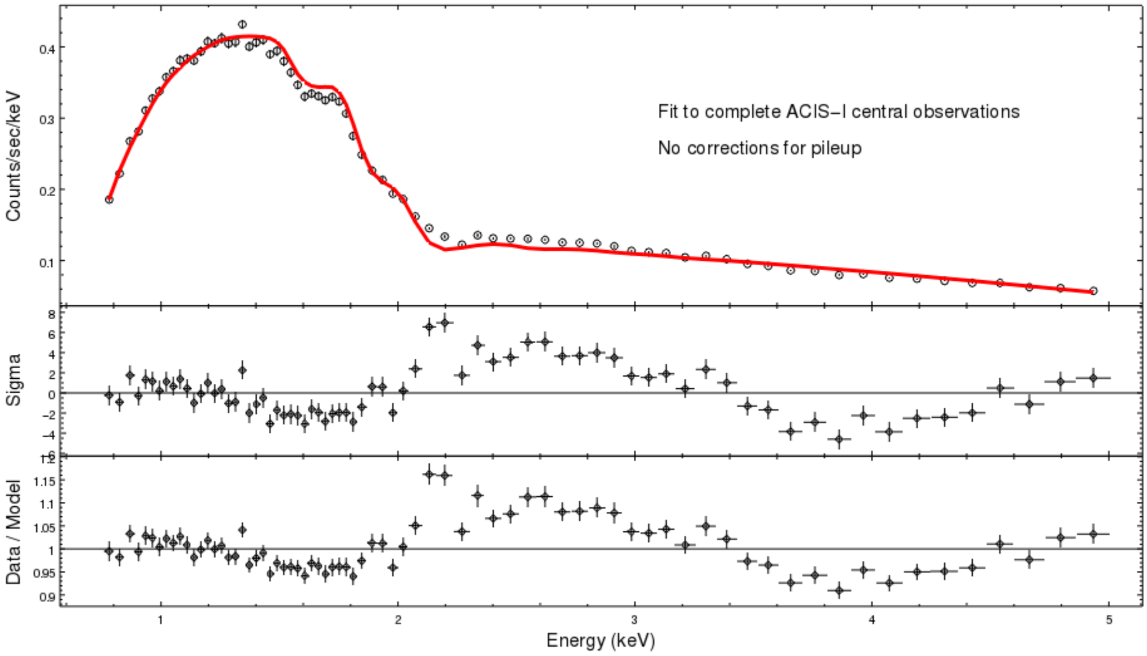

Figure 4.3: The stacked spectrum of the ACIS-I midway observations from the full extraction region, involving 242 000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 1.0 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index of $\gamma = 1.30$.

emission. After background subtraction, the fraction of counts in the 7-10 keV band and 10-15 keV bands are 10.4 % and 6.7% respectively (evidence of the very heavy pile-up present for these on-axis observations). The resulting spectrum was binned to ensure a minimum of 1800 counts per bin, and fitted between 1.5 and 5 keV with an absorbed power law.

The resulting fit is shown in Figure 4.4 together with the data to model ratio. The best-fit parameters are $n_H = 1.9 \times 10^{21}\,\mathrm{cm}^{-2}$ and an extremely hard, pile-up induced, spectral index of $\gamma = 0.66$. As with the ACIS-I midway observations, there is nonetheless a sharp localised excess in the 2–2.2 keV region superimposed on the pile-up.

4.2.3 Pileup

The level of pile-up differs greatly between the three observational groups. In the ACIS-S observations, the central pixels are very heavily affected by pile-up. As an illustration, for an extraction of the ACIS-S spectrum with no exclusion of a central region, over 5% of photon counts have energies in the 10-15 keV range. As the effective area of the telescope is zero at these energies, these counts all arise from piled-up

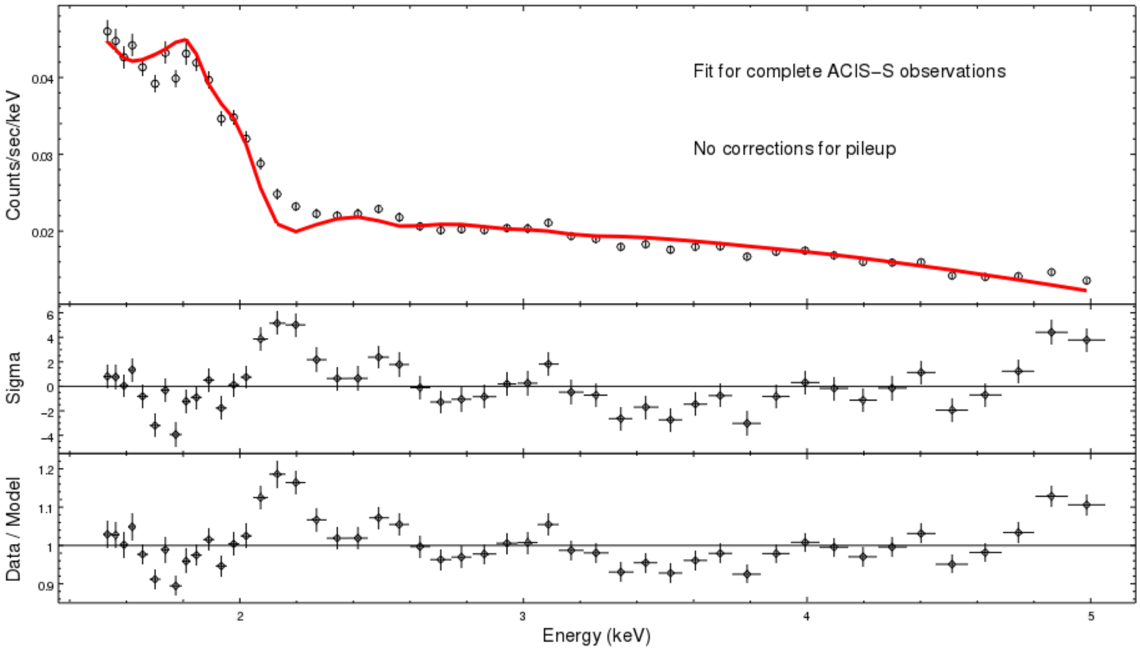

Figure 4.4: The complete stacked spectrum of the ACIS-S observations, involving 183 000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 1.9 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index of $\gamma = 0.66$.

events in which several photons have arrived in the same readout frame. In contrast, pile-up is far more moderate for the ACIS-I edge observations (for which a similar extraction across the entirety of the image results in only 0.1% of counts lying in the 10-15 keV range).

Given the data, there are two basic methods we can employ to reduce the effects of pile-up on the analysis. The first is to reduce the amount of pile-up by using an annular extraction region and extracting the spectrum only from the wings of the point spread function. The point spread function of the telescope causes the arriving photons to be spread out on the detector, with the degree of spread increasing as one moves progressively off-axis. We note that the point spread function is mildly energy dependent: photons with higher energy are spread out further than photons with smaller energy. While fewer photons arrive in the wings of the point spread function, those that do suffer less from pile-up than those arriving on the central pixels. The second is to model the pile-up, and the distortion it induces on the spectrum, explicitly. We shall utilise both approaches below.

Removing central pixels

We analyse the images and count rates of each observation using the image software `ds9`. We set the `ds9` binning such that the `ds9` pixels are the same size as the physical *Chandra* pixels (half an arc-second across). We then create an extraction region manually excluding all `ds9` pixels which are neighbours (either side-by-side or diagonally) to any pixel with total counts greater than 1% of the overall exposure time. Using the `pileup_map` tool, this results in almost all retained pixels having a pile-up fraction lower than 5%.[2] Using `arfcorr`, the responses reflect the presence of a central exclusion in the extraction region.

We first show in Figure 4.5 the ACIS-I edge observations with a central exclusion according to the above method. As these observations are relatively clean to begin with, the fractional change in the number of counts is relatively small. There are now 187000 counts before background subtraction and 153000 counts afterwards. The fraction of counts in the 7–10 keV and 10–15 keV regions are 1.2% and 0.04%. The best-fit spectral index is $\gamma = 1.85 \pm 0.015$ (the errors quoted here and in the remainder of this paper are statistical 1-σ errors, see Appendix 3.5.1 for more details) with $n_H = 2.2 \times 10^{21}\,\mathrm{cm}^{-2}$. While the statistical significance of the features at 2–2.2 keV and 3.4–3.5 keV reduces (consistent with the reduced photon counts), the magnitudes of the data to model fluctuations remain the same.

We now perform a similar cleaning of the ACIS-I observations with the source midway on the chip. As the central pixels are heavily piled-up, in this case the cleaning procedures significantly increases the quality of the fit. This is at the cost of a significant reduction in photon statistics: there are now only 88000 counts before background subtraction, and 74000 after background subtraction. We group counts so that there are 700 counts per bin, leading to a spectral index of $\gamma = 1.64 \pm 0.02$ with $n_H = 1.3 \times 10^{21}\,\mathrm{cm}^{-2}$. With a Q-value of 0.18 and a reduced χ^2 of 1.14 for 86 degrees of freedoms, this is now an overall good fit to the data (see Figure 4.6).[3]

---

[2] For these observations, the telemetry limit was 15 keV and any (piled-up) photons recorded with energy larger than this were not telemetered to the ground. For there to be a significant number of such events, there would also exist a tail of photons with energies close to, but below, the telemetry limit. From the $< 0.2\%$ fraction of photons measured in the 10-15 keV band for the cleaned spectra, we can infer that these contained very few un-telemetered events. The same analysis shows that there are very few un-telemetered events for the ACIS-I edge observations even without excising the central part of the image, but the central high pile-up regions of the ACIS-I Midway and ACIS-S observations did have significant numbers of un-telemetered events.

[3] Fitting the different observations individually produces results for γ that are consistent with the stacked result, both for the Midway and Edge observations.

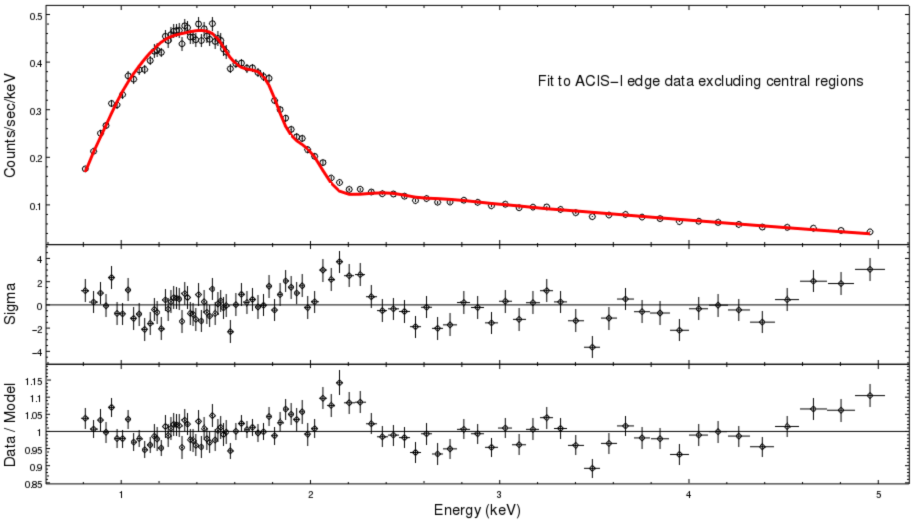

Figure 4.5: The stacked spectrum of the ACIS-I edge observations, with central pixels excluded according to the prescription in the text. There are 153 000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 2.2 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index of $\gamma = 1.85 \pm 0.015$.

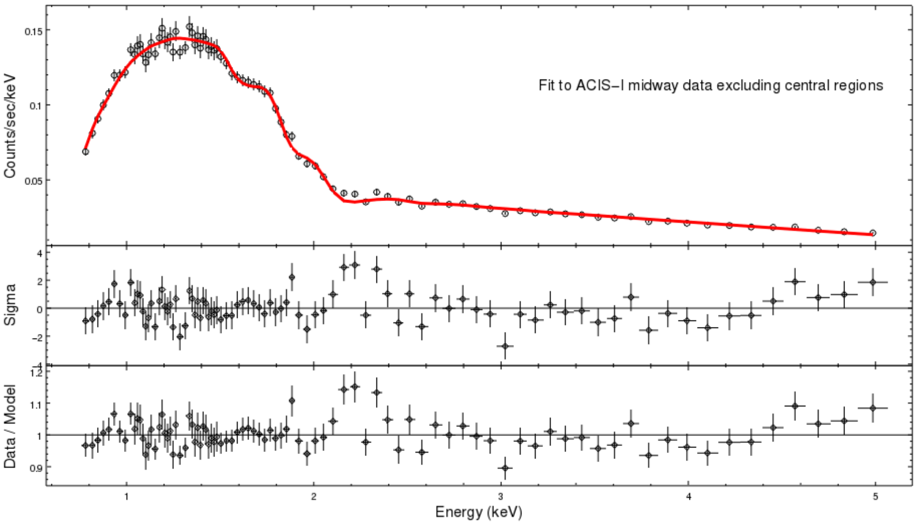

Figure 4.6: The stacked spectrum of the ACIS-I midway observations, with central pixels excluded according to the prescription in the text. There are 74 000 counts after background subtraction. The fit is to an absorbed power law, and results in $n_H = 1.3 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index of $\gamma = 1.64 \pm 0.02$.

We note that the spectral index γ is clearly different from that found with the ACIS-I edge data. However, this spectrum was obtained by excluding the central part of the image and extracting only from the wings of the point spread function. By comparison with the brightness of the AGN measured in the contemporaneous ACIS-I edge observations, we can infer that this spectrum retains only around one quarter of the total number of photons from the AGN. As the ACIS point spread function increases with energy, this biases the analysed photons towards a harder spectrum (and a lower γ). In principle, this is taken into account in the analysis procedure, as the aperture correction of `arfcorr` modifies the response to reflect the exclusion of the central part of the image. However, it is also the case that the image has a complex form with neither a spherical nor an elliptical shape. Instead, the image of the AGN – traced by the excluded central region – resembles a 'Maltese cross' shape, as can be seen in Figure 4.1.

The objective of this work is to constrain large, localised modulations in the spectrum of the AGN that could arise from ALP-photon conversion. For this purpose we therefore do not regard these different values of γ as problematic, particularly given the large difference in observational conditions between the two sets of observations and the large fraction of photons that are excluded in the ACIS-I Midway observations.

We perform a similar procedure for the ACIS-S observations. In this case we reduce pile-up by removing a central square of 16 pixels (2 arcseconds square) from an extraction region of a circle of radius 3.4 arcseconds. The resulting cleaned spectrum has 117000 counts before background subtraction and 74000 counts after background subtraction. After background subtraction there are now 0.1% counts in the 10–15 keV band and 1.1% of counts in the 7–10 keV band, indicating that this spectrum is now substantially cleaner. We group counts so that there are at least 700 counts per bin. In this case, an absorbed power law is not sufficient for a good fit and we supplement this by a soft thermal component using `xsapec` (the presence of a thermal component for NGC 1275 was also reported in [120]). The presence of a soft thermal component substantially improves the fit, resulting in an acceptable Q value of 10^{-2}. The resulting fit (see Figure 4.7) has $n_H = 1.6 \times 10^{21}\,\mathrm{cm}^{-2}$ and a power law index of $\gamma = 1.84 \pm 0.03$. The thermal component has a temperature $T = 0.92\,\mathrm{keV}$. At this temperature, the amplitude and abundance of the thermal component are largely degenerate in the fit. Fixing the abundance at solar abundance, the relative amplitude of thermal component to the power law is 0.15.

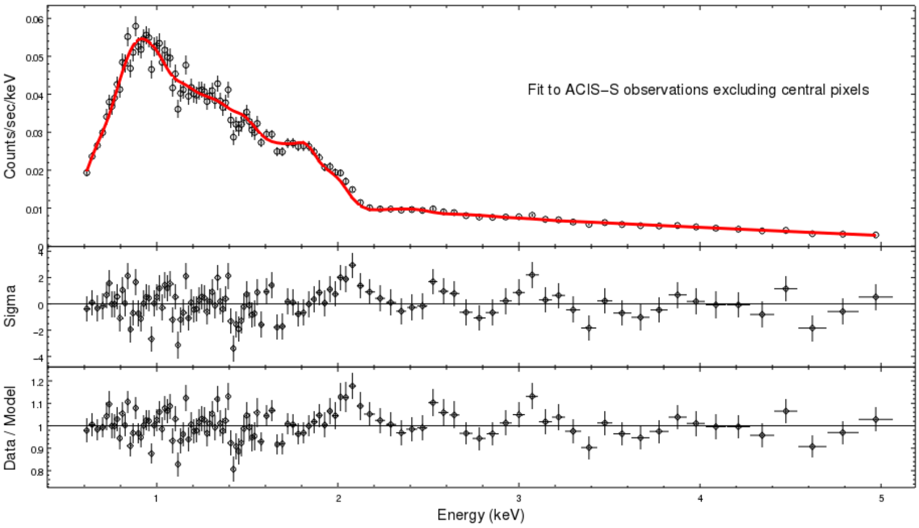

Figure 4.7: The stacked spectrum of the ACIS-S observations, with a central 16-pixel square excluded. The fit is to an absorbed power law and soft thermal component.

Modelling pileup

Here we aim to model the pile-up directly. While this modelling does not provide a complete picture of pile-up, it ameliorates its effect. A full treatment would require customised tools going beyond the scope of this paper (the pile-up tool `jdpileup` [132] provided with Sherpa is optimised for on-axis sources, while NGC 1275 is off-axis in all ACIS-I observations).

We first describe our pile-up modelling for the ACIS-S observations, where the image of NGC 1275 is on-axis. This results in a high degree of pile-up – the central pixels have large numbers of counts above 10 keV, all of which arise from multi-photon pile-up events.

We use the `jdpileup` model as described in [132]. As this model assumes the source is on-axis, it is directly appropriate here. The model assumes Poisson statistics to calculate the probability of different numbers of photons hitting an event-detection cell (a 3×3 pixel region in Chandra) within the read-out time adjusted to these respective observations. For the ACIS-S and ACIS-I midway observations, the read-out time was 3.1 seconds, and for the ACIS-I edge observations it was 3.2 seconds. The model then convolves this probability with the probability of such events being assigned a 'good' grade, and the conversion from photon energy to pulse height by the detector. We note that the non-linearity of pile-up (the dependence on the energy spectrum for all energies below the bin) will cause degeneracies in the inferred parameters of the `jdpileup` model (α and f) with those of the energy spectrum [132].

The parameter n, the number of regions to which `jdpileup` will be applied independently, was set to 1 for the ACIS-S observations (as the target was a point source). For an extended source it should roughly correspond to the number of 3×3 pixel islands in the region, as `jdpileup` assumes spatial uniformity across the extraction region. For the ACIS-I midway observations, where the AGN is smeared out across several pixel islands, the value of n proved difficult to determine for large extraction regions, with the value corresponding to the best fit being unrelated to the number of pixel islands. We therefore constrained ourselves to a smaller central region, with little variation in count rate between pixels, to give us more control over the pile-up model.

It is also worth noting that the `jdpileup` model is set to zero for energies less than 0.5 keV and energies greater than 10 keV. As there are many events above the 10 keV range for the ACIS-S observations, we only model the spectrum up to 10 keV, as extending the fit beyond that would result in the model parameters being sent to unphysical values.

For the ACIS-S pile-up analysis, we used a circular region around NGC 1275 of radius 3.5 arcseconds, with the background taken from an annular region of inner radius 4 arcseconds and outer radius 7 arcseconds. Counts were grouped to 1500 per bin. The spectrum was modelled using an absorbed power law and thermal emission with temperature $T = 0.85\,\mathrm{keV}$, and fitted with the `jdpileup` model for energy values between 1 and 10 keV. The data was fit using the `moncar` Monte-Carlo method, and run several times to ensure the global minimum had been found.

The resulting fit and the ratio of data to model are shown in Figure 4.8. The best-fit parameters involve $n_H = 2.6 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index $\gamma = 1.81$.[4] The best-fit α and f parameters of `jdpileup` were $\alpha = 0.660$ and $f = 0.943$ respectively. We can clearly see that this model provides a reasonable description of the data all the way up to 10 keV, and produces physically sensible values for n_H and γ, despite the model estimating a pile-up fraction of 82%. The reduced statistic of the fit is 1.75, with a Q-value of 10^{-5}. While we re-emphasise that this will not represent a perfect account of pile-up, it does capture the relevant physics, producing a sensible fit with physical parameters.

For the ACIS-I midway observations we considered a central 6×6 pixel extraction region (this is almost exactly the complement of the clean ACIS-I midway spectrum used in the previous subsection). The resulting spectrum contains around 136000

---

[4] As the uncertainties on the pile-up modelling are hard to quantify, we only quote best-fit parameters and do not include errors.

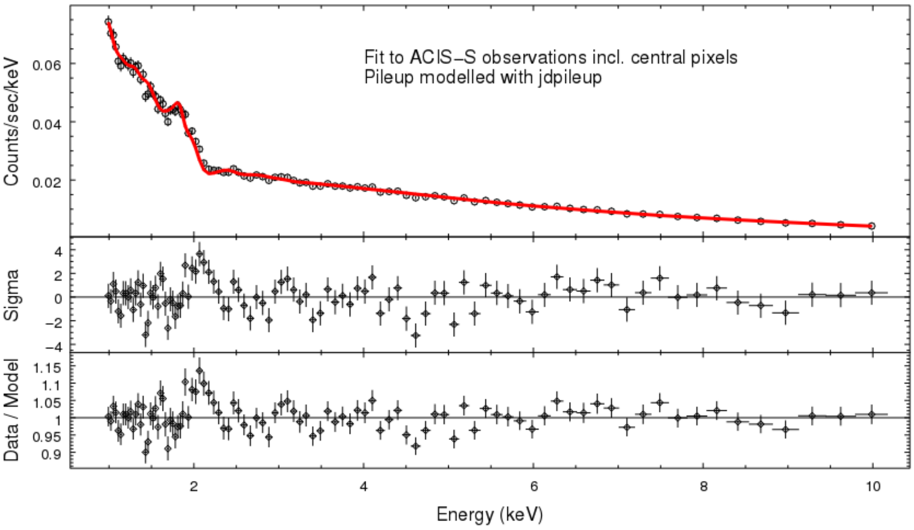

Figure 4.8: The ACIS-S observations, involving 177 000 counts after background subtraction. The fit is to an absorbed power law with a thermal component of $T = 0.85\,\mathrm{keV}$, and pile-up is modelled with `jdpileup`. The ratio of data to model is shown at the bottom of this figure.

counts, reducing to 134000 after background subtraction, giving a very high AGN to cluster contrast of 67:1. After background subtraction, the fraction of counts in the 7-10 keV band and 10-15 keV bands are 5.3% and 1.2% respectively. Counts were grouped to 1000 per bin. The spectrum was modelled using an absorbed power law and fitted with the `jdpileup` model, this time for energy values between 1 and 9 keV, to ensure no counts with energy greater than 10 keV were included in the final bin. The `jdpileup` parameter n was set to 4, the number of 3×3 pixel islands in the extraction region.

The resulting fit and ratio of data to model are show in Figure 4.9. The best-fit parameters involve $n_H = 2.5 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index $\gamma = 1.93$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.324$ and $f = 0.975$ respectively, and the estimated pile-up fraction was 35%. While the fit is not perfect, it does give a reasonable characterisation of the data.

We finally consider the case of the ACIS-I edge observations, where pile-up is relatively weak. Although a reasonable fit can be made with no pile-up modelling, we now aim at modelling the pile-up also in these observations for completeness. There is a rule of thumb[5] that 0.007 counts/second per 3×3 pixel island is about 1%

---

[5] `cxc.harvard.edu/csc/memos/files/Davis_pileup.pdf`

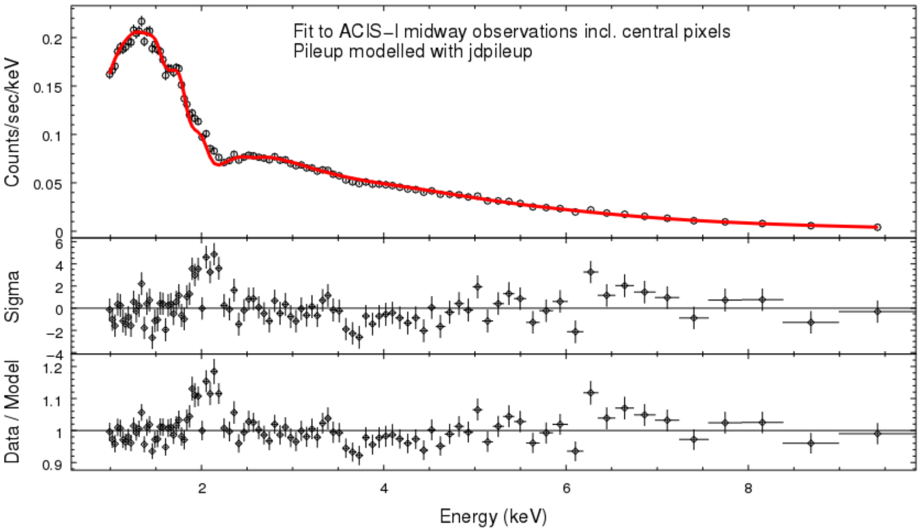

Figure 4.9: The ACIS-I midway observations from a central 6×6 pixel region, involving 134 000 counts after background subtraction. The fit is to an absorbed power law and pile-up is modelled with `jdpileup`.

pile-up, and 0.07 counts/second is about 10%. We will assume that the ACIS-I edge observations with mild cleaning have pile-up fractions of 10% or lower.

For the ACIS-I edge observations, we used as source an elliptical annulus of radii 3 and 5 arcseconds, with the 12 brightest pixels removed by a contour. The background was extracted from an elliptical annulus of outer radii 18 and 23 arcseconds, with inner radii 8 and 12 arcseconds. The resulting spectrum contains around 128000 counts, reducing to 111000 after background subtraction, giving an AGN to cluster ratio of 6:1. After background subtraction, the fraction of counts in the 7–10 keV band and 10–15 keV bands are 1.1% and 0.1% respectively. We grouped counts to 1000 per bin. Here there are fewer counts out at higher energies than in the previous two sets of observations, so we fit only out to 5 keV. The `jdpileup` parameter n was set to 16, the number of 3×3 pixel islands in the extraction region. This should only be thought of as a rough estimate, as the `jdpileup` model is stated to be accurate for on-axis point sources, and the edge observations are also not spatially uniform in terms of counts. The fit is to an absorbed power law, and results in $n_H = 2.7 \times 10^{21}\,\mathrm{cm}^{-2}$ and a spectral index of $\gamma = 1.89$. The best-fit α and f parameters of `jdpileup` were $\alpha = 0.52$ and $f = 0.87$ respectively. The pile-up fraction was 3%, consistent with estimates and very mild as expected.

Given the high flux levels of the source, a cleaner spectrum requires reducing the effects of pile-up in the spectrum. The first method involves excluding a central region

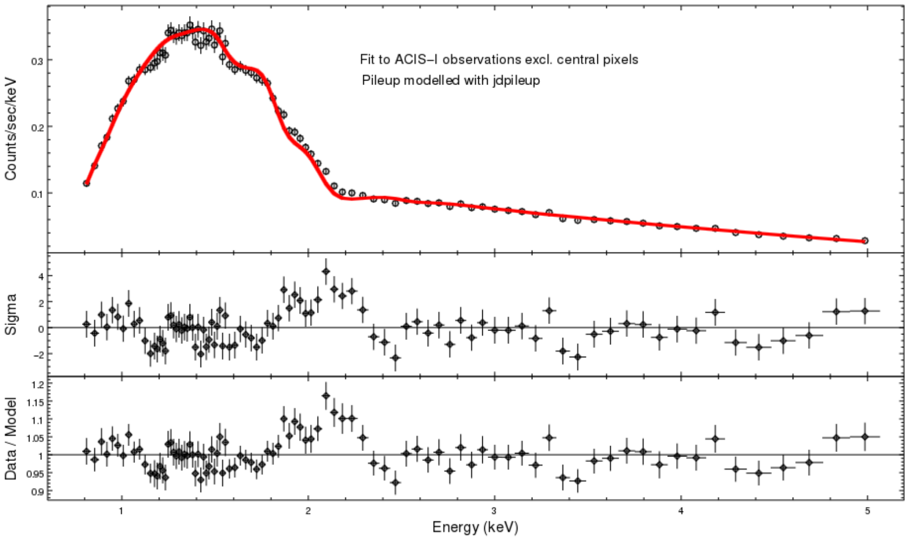

Figure 4.10: The ACIS-I edge observations, involving 110000 counts after background subtraction. The fit is to an absorbed power law and pile-up is modelled with `jdpileup`.

of highest pile-up from the analysis. While this reduces the photon counts – and so the statistical significance of any features in the spectrum – it also allows produces purer spectra of higher quality. The second method is to keep the full spectrum, but model the pile-up explicitly. For the purpose of placing bounds on ALPs and constraining the existence of ALP-induced spectral modulations, the most useful spectra are those where photon events occur at their true energies.

Finally, we also considered *Chandra*'s grating observations of NGC 1275 (ObsIDs 333 and 428 with HETG) as a means to reduce pile-up, but this reduced counts to the point where it did not seem worthwhile.

4.2.4 Analysis of residuals in *Chandra* NGC 1275 data

No residuals in the spectra exceed 10%. At the 10% level, the *Chandra* spectra presented in the main body of the paper show two main features departing from a power law: an excess at 2–2.2 keV, present in all *Chandra* observations with overwhelming statistical significance, and a deficit 3.4–3.5 keV that is not as strong but is still present at almost five (local) sigma in the ACIS-I edge data. These features ensure that an absorbed power law (plus thermal component) is not a good fit to the data.

We perform a statistical analysis of these features at $E \sim 2\,\mathrm{keV}$ and $E \sim 3.5\,\mathrm{keV}$. We first consider the $E \sim 2\,\mathrm{keV}$ feature. Considering the clean ACIS-I Edge data, we include a positive Gaussian ($xswabs * (xspowerlw + xsgaussian)$) and analyse its effect on the fit. For a zero-width Gaussian, this leads to an improvement in the fit by $\Delta\chi^2 = 33.7$ with a best-fit energy of 2.13 keV. For a finite-width Gaussian, the improvement is $\Delta\chi^2 = 48.0$, with a best-fit energy of 2.02 keV and a best-fit width of 0.2 keV.

For the ACIS-I midway data, a similar additional Gaussian improves the fit by $\Delta\chi^2 = 18.6$ with a best-fit energy of 2.17 keV (the fitted width is much smaller than the detector resolution and so results are identical for zero-width and finite-width Gaussians). For the ACIS-S data, the additional Gaussian improves the fit by $\Delta\chi^2 = 20.3$ with a best-fit energy of 2.06 keV (again, fitting the width gives a result smaller than the detector resolution and so does not affect the result).

For the 3.5 keV feature, we consider the cleaned datasets for the three sets of observations and include a negative Gaussian ($xswabs * (xspowerlw - xsgaussian)$). We formally treat the Gaussian as zero-width, but any finite width much narrower than the ACIS energy resolution gives an identical result. For the three data sets (ACIS-I Edge, ACIS-I Midway, and ACIS-S) we plot below (Figures 4.11-4.13) the improvement in the χ^2 that can come from adding an additional negative Gaussian $\Delta\chi^2_{\rm NG}$. The ACIS-I Edge data show a strong preference for an additional negative Gaussian ($\Delta\chi^2 \sim 17$) at 3.5 keV. For both the ACIS-I Midway and ACIS-S data, a negative Gaussian at 3.5 keV mildly improves the fit ($\Delta\chi^2 \sim 1.5$ and $\Delta\chi^2 \sim 0.4$) but is not required (as the ACIS-I Edge dataset is both larger and cleaner than the other two datasets, these results are consistent – in particular, the inferred *strength* of the dip is consistent within 2 σ).

Across all three plots, the largest feature is clearly seen to be the residual in the ACIS-I Edge data at 3.5 keV. The next strongest residuals are two features at $E \sim 1.2$ keV (ACIS-I Edge) and $E \sim 1.4$ keV (ACIS-S) with $\Delta\chi^2 \sim 10$. In both cases, there are very strong atomic lines visible in the background spectra at precisely these energies (Fe XXI, Fe XXII, Fe XXIII, Fe XXIV, Mg XII). These features can be reliably associated with these atomic lines, coming from a small mis-subtraction of the background (deep in the core of the Perseus cluster, the physical conditions of the gas in the signal region and background region will not be precisely identical).

In Figure 4.14, we produce a combined plot of the significance of the 3.5 keV deficit across all *Chandra* observations. In order to compare observations taken at different times and with different intrinsic AGN luminosities, to measure the strength of the deficit we have used the induced fractional deficit of photons in the 3.3 –

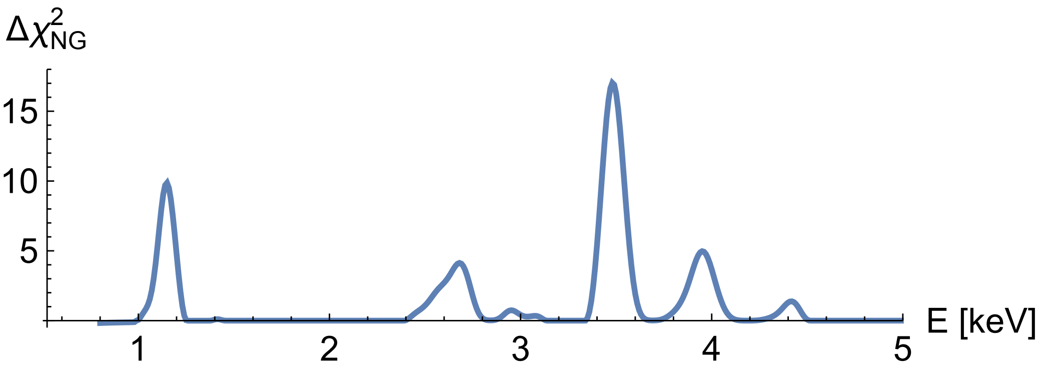

Figure 4.11: The improvement in χ^2 attainable by adding a negative Gaussian at the specific energy for ACIS-I Edge observations.

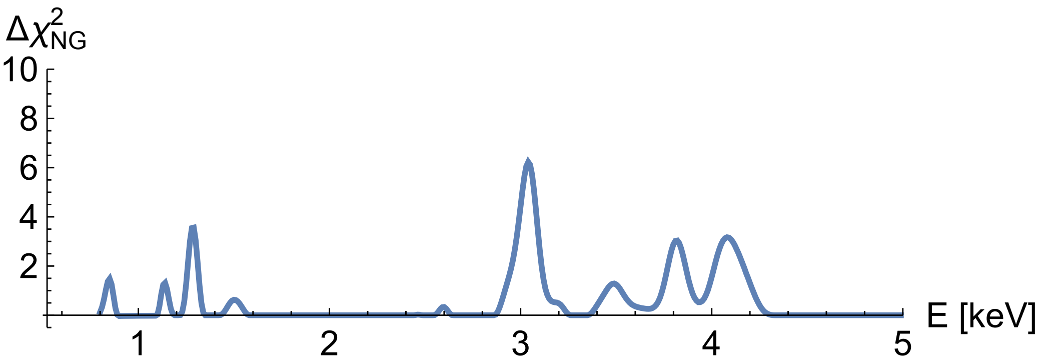

Figure 4.12: The improvement in χ^2 attainable by adding a negative Gaussian at the specific energy for ACIS-I Midway observations.

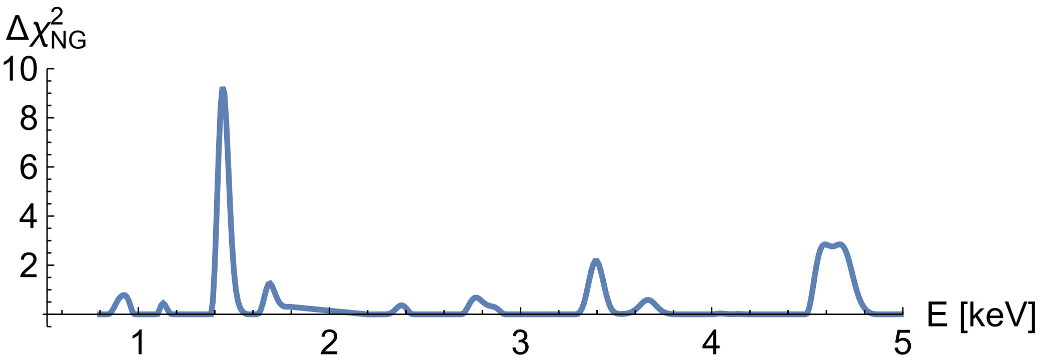

Figure 4.13: The improvement in χ^2 attainable by adding a negative Gaussian at the specific energy for ACIS-S observations.

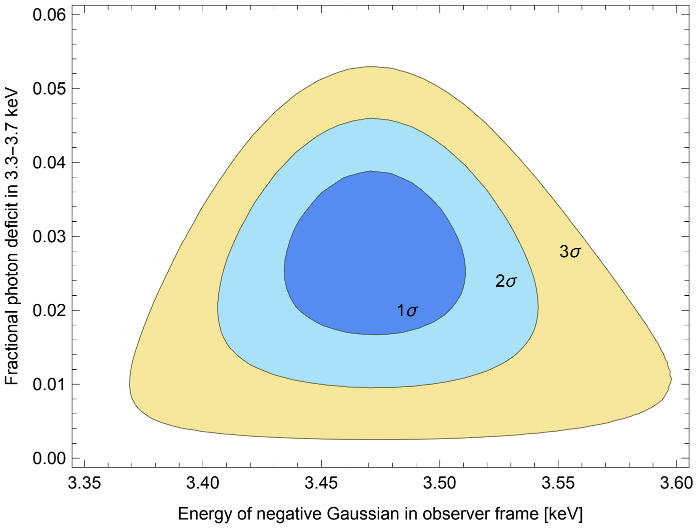

Figure 4.14: The overall significance and location for the 3.5 keV deficit, combined across all *Chandra* observations (cleaned datasets only). We have used the induced fractional deficit in photons in the 3.3 - 3.7 keV region to facilitate a uniform comparison for the strength of the dip, as the AGN luminosity varies between observations.

3.7 keV regime. The feature remains over 3σ significant globally, and is also located at precisely the same energy as the diffuse cluster excess observed in [141, 142].

ALP Interpretation of Features

The purpose of this paper was to use the extraordinary dataset of counts from the NGC 1275 AGN to search for spectral irregularities, with the intent of constraining ALP parameters. This search has resulted in two features being present in the data at high statistical significance, the former of which is consistent with arising from pile-up around the Iridium edge.

It is not possible for us to constrain ALP couplings beyond a level where they would produce residuals comparable to these features. We can estimate approximately that to produce such residuals would require an ALP-photon coupling of magnitude $g_{a\gamma\gamma} \sim 1 - 5 \times 10^{-12}\,\mathrm{GeV}^{-1}$. For example, Figure 4.15 shows a fit to the clean ACIS-I edge observations with an absorbed power law multiplied by the photon survival probability $P_{\gamma\to\gamma}$. In this case, $P_{\gamma\to\gamma}$ was calculated assuming the existence of ALPs with $g_{a\gamma\gamma} = 1 \times 10^{-12}\,\mathrm{GeV}^{-1}$ and a central magnetic field $B_0 = 25\,\mu\mathrm{G}$. We see that the anomalies at 2.2 keV and 3.5 keV have been alleviated by the presence of ALPs (although, for this magnetic field, at the expense of creating similarly sized anomalies at higher energies). Figure 4.16 shows the corresponding photon survival probability spectrum.

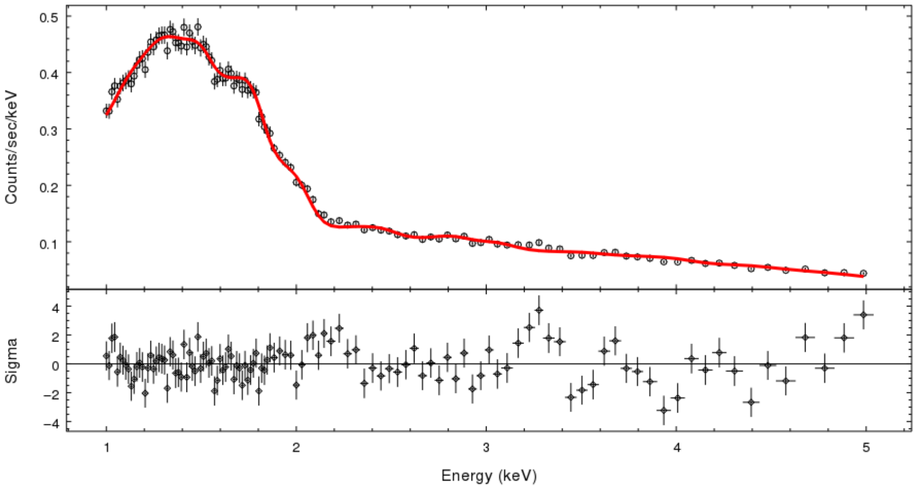

Figure 4.15: A fit to the clean ACIS-I edge observations with an absorbed power law multiplied by the photon survival probability for a specific example of ALP-photon conversion. We include ALPs with $g_{a\gamma\gamma} = 1 \times 10^{-12}\,\mathrm{GeV}^{-1}$ and assume a central field of $B_0 = 25\,\mu\mathrm{G}$. The reduced χ^2 is 1.51, compared to 1.65 for a fit to an absorbed power law without ALPs.

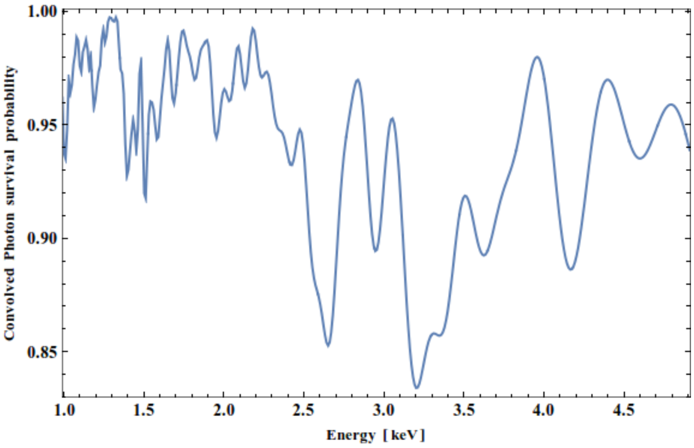

Figure 4.16: The photon survival probability for the fit shown in Figure 4.15.

4.2.5 The Presence of the 6.4 keV Iron line

As the cleaned spectra in the main body of the paper are only extracted up to 5 keV, it is a useful consistency check to ensure that the reflected iron Fe Kα line at 6.4 keV (6.3 keV after redshifting) is present in our spectra. For this purpose we use the cleaned spectra for each of the three sets of *Chandra* observations, and fit the spectrum only between 5.5 and 7 keV. We fit first a power law, and then the sum of a power law and a Gaussian (we do not fit n_H as absorption is irrelevant at these energies).

We determine the best-fit central energy of the additional Gaussian and the resulting improvement in χ^2. These fits are summarised in Table 4.3 below and we also show explicitly the fit for the ACIS-S data in Figure 4.17.

Data set	Power law index	E_{line} (keV)	$\Delta\chi^2$ (Degrees of freedom)
ACIS-I Edge	$0.84^{+0.3}_{-0.1}$	$6.32^{+0.06}_{-0.05}$	6.8/2
ACIS-I Midway	$-0.14^{+0.03}_{-0.02}$	$6.33^{+0.05}_{-0.06}$	2.2/2
ACIS-S Central	$0.49^{+0.07}_{-0.11}$	$6.29^{+0.02}_{-0.03}$	12.2/2

Table 4.3: Fit parameters for a (power law + Gaussian) fit between 5.5 and 7 keV.

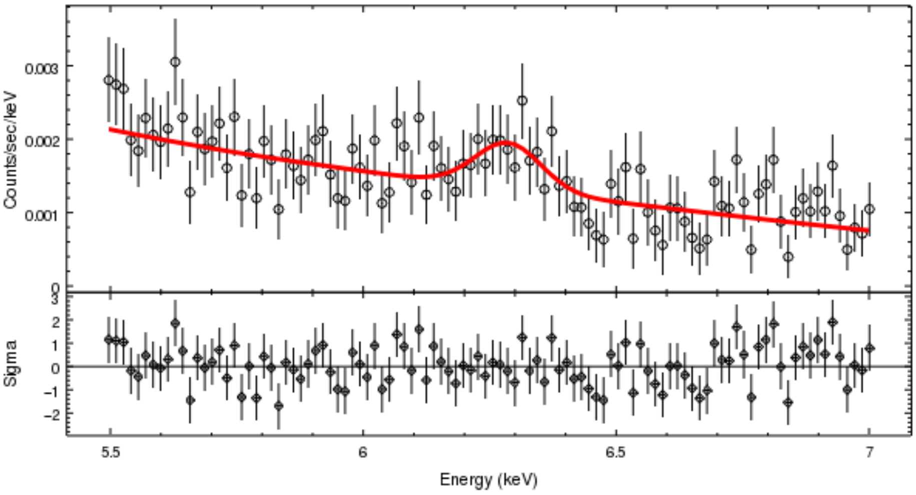

Figure 4.17: The Fe Kα line in the cleaned ACIS-S data.

These show that the iron line is clearly present in the ACIS-S and ACIS-I Edge data, and marginally present in the ACIS-I Midway data. As ACIS-S has better energy resolution than ACIS-I, it is unsurprising that the Fe line is found at highest significance there.

We note that all spectral indices are unphysically hard (and even negative for the ACIS-I Midway data). This is a clear consequence of pile-up. For energies above 5 keV, the *Chandra* effective area is falling off rapidly. There are also intrinsically far fewer source photons than are present at lower energies, given the power law spectrum. The effects of pile-up at these higher energies are therefore far more severe than at lower energies.

While these spectra have been cleaned, this does not mean that there is zero pile-up. For the bulk of the spectra, the effects of pile-up in the cleaned spectra are at the level of a few per cent. However, for $E > 5\,\mathrm{keV}$ the effects of pile-up are substantially more important for the above reasons: the combination of the intrinsic decline in the spectrum and the fall-off in the *Chandra* effective area implies that pile-up plays a proportionally far more severe role at high energies.

4.2.6 Bounds from *Chandra* data

From the perspective of constraints on ALPs, the overall summary of the data is that the spectrum from NGC 1275 is fit by an absorbed power law, with residuals from the power law not larger than 10%. We will further discuss residuals at the 10% level in Appendix 4.2.4 (in particular features around 2 - 2.2 keV and at 3.5 keV). However, the absence of any modulations in the spectrum at a level greater than 10% allows powerful constraints to be placed on ALP parameter space.

In particular, this allows us to say that on passing from NGC 1275 through the Perseus cluster and towards us, $\langle P(\gamma \to a) \rangle \lesssim 20\%$. It follows that ALP-photon couplings large enough to generate the saturated limit of $\langle P(\gamma \to a) \rangle = 1/3$ are robustly excluded, as these would produce larger residuals in the data than are actually observed.

Note that this approach to placing bounds is conservative. The simulated data (with ALPs) is idealised – it has no pileup present and assumes a perfect subtraction of the cluster background. Couplings are excluded when they lead to simulated data that would *still* be a worse fit than the actual data, even though the simulated data is observationally ideal while the real data contains effects such as pileup and cluster background that lead to increased residuals and make it far noisier than simulated data.

To obtain approximate bounds on $g_{a\gamma\gamma}$, we compare two models for the flux $F(E)$ observed from NGC 1275:

- Model 0: An absorbed power law $F_0(E) = AE^{-\gamma} \times e^{-n_H\sigma(E)}$, as described in Equation (4.1).
- Model 1: An absorbed power law multiplied by the photon survival probability assuming the existence of ALPs with coupling $g_{a\gamma\gamma}$. In this case the predicted flux also depends upon the magnetic field $\mathbf{B}$ along the line of sight. We have $F_1(E, \mathbf{B}) = AE^{-\gamma} \times e^{-n_H\sigma(E)} \times P_{\gamma\to\gamma}(E(1+z), \mathbf{B}, g_{a\gamma\gamma})$.

Although we have (limited) empirical estimates of the strength of the magnetic field in Perseus, the exact structure is unknown. In practice, we randomly generate many instances of the field from a given power spectrum. The parameter most relevant to ALP-photon conversion is the central magnetic field strength B_0, estimated as $25\,\mu$G in [123]. Based on results for the Coma cluster, we assume that B decreases with radius as $B \propto n_e^{0.7}$ [127]. The electron density n_e has the radial distribution found in [119],

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

We simulate each field realisation with 600 domains. The length l of each domain is between 3.5 and 10 kpc, randomly drawn from a power law distribution with minimum length 3.5 kpc and power 0.8. We therefore have:

$$P(l = x) = N \begin{cases} 0 & \text{for}\, x > 10\,\mathrm{kpc}\ , \\ x^{-1.2} & \text{for}\, 3.5\,\mathrm{kpc} < x < 10\,\mathrm{kpc}\ , \\ 0 & \text{for}\, x < 3.5\,\mathrm{kpc}\ , \end{cases} \tag{4.2}$$

with normalisation constant N.

The coherence length and power spectrum of the magnetic field in the centre of Perseus is not observationally determined. Instead, these parameters are motivated by those found for the cool core cluster A2199 [128], taking a conservative value for the magnetic field radial scaling. The magnetic field and electron density are constant in each domain, with $B(r)$ and $n_e(r)$ evaluated at the centre of the domain and the direction of $\mathbf{B}$ chosen at random.

We compute 95% confidence limits on $g_{a\gamma\gamma}$ by generating fake data from Model 1 and assessing how well it is fit by Model 0 i.e. how well the oscillations due to ALP-photon conversion can hide in the Poisson noise. As it has minimal pile-up, we use the clean ACIS-I edge observations for this analysis. We fit the spectrum between 1 and 4 keV (a region unaffected by pile-up) and bin such that there are 1000 counts in each energy bin. We use `Sherpa`'s Levenberg-Marquardt fitting method with Poisson errors derived from the value of the data in each bin. Our procedure to determine whether ALPs with coupling $g_{a\gamma\gamma}$ are excluded at the 95% confidence level is as follows:

1. Fit Model 0 to the real data and find the corresponding reduced χ^2, χ^2_{data}.
2. Randomly generate 50 different magnetic field realisations $\mathbf{B}_i$ for the line of sight to NGC 1275.
3. For each $\mathbf{B}_i$, compute $P_{\gamma\to\gamma}(E, \mathbf{B}_i, g_{a\gamma\gamma})$ by numerically propagating photons at different energies through $\mathbf{B}_i$, as described for example in [87]. We take 300 photon energies equally spaced between 1 and 4 keV.
4. For each $\mathbf{B}_i$, generate 10 fake data sets from Model 1, using `Sherpa`'s fake pha method.
5. Fit Model 0 to each of the 500 fake data sets and find the corresponding reduced χ^2, χ^2_i for each.
6. If fewer than 5% of the χ^2_i are lower than χ^2_{data}, $g_{a\gamma\gamma}$ is excluded at the 95% confidence level.

We scan over $g_{a\gamma\gamma}$ in steps of $10^{-13}\,\text{GeV}^{-1}$. For the $g_{a\gamma\gamma}$ value excluded, we also check that the three values above it in our grid are also excluded. For the magnetic field parameters described above, we find $g_{a\gamma\gamma} \lesssim 1.4 \times 10^{-12}\,\text{GeV}^{-1}$ which is shown in Figure 4.18.

If we consider a more pessimistic scenario with $B_0 = 15\,\mu\text{G}$ and a minimum coherence length of 0.7 kpc, we instead find $g_{a\gamma\gamma} \lesssim 2.7 \times 10^{-12}\,\text{GeV}^{-1}$. If we take an even more pessimistic scenario in which the central field is $B_0 = 10\,\mu\text{G}$ and the minimum coherence length is 0.7 kpc, the bound increases further to $g_{a\gamma\gamma} \lesssim 4.0 \times 10^{-12}\,\text{GeV}^{-1}$.

We finally note that if we excluded the feature at 2.2 keV from the fit, removing the 1.8 - 2.3 keV region based on this being the location of an effective area edge, the bounds would improve significantly, to $g_{a\gamma\gamma} \lesssim 1.1 \times 10^{-12}\,\text{GeV}^{-1}$ for the case of a central magnetic field of $B_0 = 25\,\mu\text{G}$ (see Appendix 4.2.4). This illustrates the conservative nature of the bounds compared to any mis-modelling of the actual data.

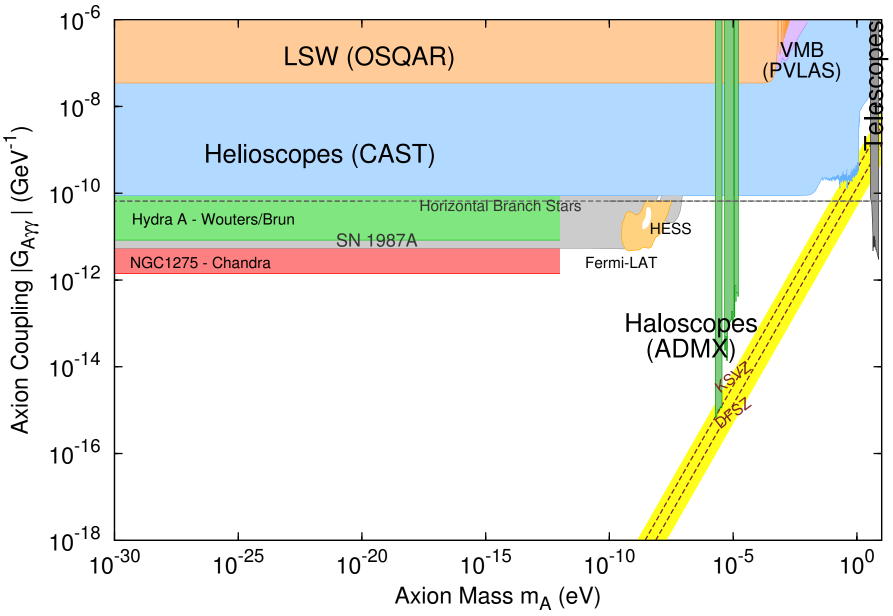

Figure 4.18: The 95% upper bounds on the ALP-photon coupling in comparison to previously obtained bounds on ALPs, for a central magnetic field strength $B_0 = 25\,\mu$G.

4.2.7 Validation with MARX simulation

As a validation of these bounds, we also perform MARX simulations ([143]) of the ACIS-I Edge observations to confirm that the results are robust against any pileup effects. Significantly piled-up data is known to produce complex effects, in particular with a bright background (see for instance [133]).

As the actual spectrum is a sum of the AGN power law and thermal emission from the cluster, we first need to determine the magnitude of thermal emission with the ACIS-I edge extraction region.

To extract the level of thermal emission within the ACIS-I Edge observations, we follow two strategies. First, we took the longest on-axis ACIS-S observation (*Chandra* obsid 4952) and used the same extraction region as for the ACIS-I edge data (also excluding the AGN with a central circle of radius 2 arcseconds). We fit an `apec` model to this, obtaining a temperature of $T \sim 4$keV, an abundance of $Z \sim 0.4$ and a normalisation of 0.0022. While such a single-temperature `apec` model is not an accurate description of the gas dynamics at the centre of the cluster ($\chi^2/N = 474/375$), it is adequate for our purposes as we are more concerned with the level of pileup, and thus the overall *magnitude* of thermal emission, rather than a perfect spectral shape. Using this model together with the exposure and responses for the ACIS-I edge observations, this suggests the presence of around $\sim$ 30k thermal counts present in the ACIS-I spectrum.

Secondly, as a check on this, we also used purely the ACIS-I edge data and directly fit it with a sum of an absorbed power law and absorbed thermal emission. In doing this, we froze the temperature and abundance of the thermal model to the values obtained from the *Hitomi* spectrum of the central region of the Perseus cluster, $T =$ 3.48 keV and $Z = 0.54$ ([144]). This produced a best-fit normalisation for the `apec` norm of 0.0027, comparable to that attained when using the ACIS-S data to determine the magnitude of thermal emission. This slightly higher value would correspond to a total of $\sim$ 36k counts from thermal emission within the ACIS-I Edge spectrum. As both results are consistent, they show that of the $\sim$ 270k counts in the ACIS-I edge spectrum only around 10-15% arise from thermal cluster emission.

We then implement a MARX simulation of the ACIS-I edge observations. For this, we used the observational parameters (date, on-axis location and roll angle) for the longest ACIS-I edge observation of the AGN (*Chandra* obsid 11713). As all the four actual ACIS-I Edge observations have similar properties, we set the exposure of this MARX simulation to 200ks, the total exposure of the actual data.

The data was first simulated without any axions present.

1. The thermal emission was simulated as a spatially extended Gaussian with a width of 30 arcminutes, centred on the AGN. It was modelled as an `apec` model with a temperature $T = 3.5$keV and an abundance $Z = 0.48$. The normalisation of the model was adjusted so that it had the same fitted strength as for real data (as determined above) when extracted within the region used for the ACIS-I Edge data.

2. The AGN emission was simulated as a point source power law. The normalisation of the power law was adjusted so that - after it had been combined with the thermal emission and the pileup processing applied - it had the same normalisation as the real data when extracted over the same extraction region.

3. The two individual simulations were then combined using `marxcat` and the MARX pileup processing applied to the combined simulation.

4. Counts were grouped to 500 and the combined spectrum was fitted with the sum of a power law plus thermal emission. The previous two steps were iterated until the fitted strength and index of the power law matched that of the real data.

We generated a total of 50 fake data samples in this way. To match the fitted strength of the AGN ($4.7 \times 10^{-3}\,\mathrm{ph\,cm^{-2}\,s^{-1}\,keV^{-1}}$ at 1 keV) required an intrinsic strength of $6.4 \times 10^{-3}\,\mathrm{ph\,cm^{-2}\,s^{-1}\,keV^{-1}}$ at 1 keV. Comparison of the number of counts within the extraction region in the `marx` event files pre- and post-pileup processing suggests an overall pileup fraction of 15% in the dataset. Nonetheless, the pileup is not sufficient to make the spectral fit unsatisfactory (the reduced χ^2 values cluster around unity for ~ 250 degrees of freedom, with the largest reduced χ^2 being 1.37).

As a test of an even more piled-up spectrum with a similar number of photon counts, we also simulated data using a 100ks MARX exposure but with the intrinsic AGN strength doubled. As this normalisation of $12.8 \times 10^{-3}\,\mathrm{ph\,cm^{-2}\,s^{-1}\,keV^{-1}}$ at 1 keV is much larger than that present in the 2016 *Hitomi* data, and as the AGN has been increasing in strength from 2001 onwards (see [122]), pileup is much greater than in the actual 2009 data. The fit now tends to be worse, but not terrible. Among the fifty fake data samples, the average reduced χ^2 is ~ 1.25 for ~ 250 degrees of freedom, but there are a few cases still with a reduced $\chi^2 < 1$.

We now repeat the MARX simulations including the effects of axions in the data. We multiply the AGN power law with a $P(\gamma \to \gamma)$ survival probability coming from photon-ALP conversion. As this leads to a net reduction in the number of photons present, we adjust the intrinsic normalisation upwards so that the fitted normalisation matches the actual data.

For a fixed ALP coupling, we generate 50 fake data samples. To compare these results to those described above using Sherpa's `fake_pha` command, we simulate three separate ALP couplings: $g_{a\gamma\gamma} = 1,\ 1.5,\ 2 \times 10^{-12}\,\mathrm{GeV}^{-1}$. In all cases we assume the 'optimistic' magnetic field model. We directly fitted the absorbed sum of a power law and `apec` thermal emission to the ACIS-I Edge extraction region for the fake data, and compared to the quality of the fit when doing so using the real data.[6] We first group counts to 500 and fit between 1 and 5 keV using the `chi2datavar` statistic. Applied to the real data, this gives a reduced χ^2 of 1.83 (coming from the excess around 2 keV).

We then apply this fit to data involving axions. Of the fifty fake data sets produced using $g_{a\gamma\gamma} = 2 \times 10^{-12}\,\mathrm{GeV}^{-1}$, the lowest reduced χ^2 was 4.07. When $g_{a\gamma\gamma} = 1.5 \times 10^{-12}\,\mathrm{GeV}^{-1}$, the average reduced χ^2 is still ~ 2.5, but now with three out of fifty spectra with better fits than the actual data ($\chi^2/N < 1.83$). When $g_{a\gamma\gamma} = 1.0 \times$

---

[6] This is marginally different from the analysis procedure described in earlier sections, as it aims at fitting the thermal emission rather than subtracting it. We also did this analysis subtracting background emission from the real data set, and fitting only with a power law. The results however are very similar, with no significant changes to the excluded coupling.

$10^{-12}\,\mathrm{GeV}^{-1}$, more than half the fake spectra have better fits than the actual data. Using 95% confidence level exclusions, this shows that $g_{a\gamma\gamma} = 2 \times 10^{-12}\mathrm{GeV}^{-1}$ is strongly excluded (the best scenario having a reduced χ^2 of 4.07), while $g_{a\gamma\gamma} = 1.5 \times 10^{-12}\,\mathrm{GeV}^{-1}$ is on the boundary of exclusion, and $g_{a\gamma\gamma} = 1 \times 10^{-12}\,\mathrm{GeV}^{-1}$ is clearly not excluded.

This extremely sharp behaviour of the reduced χ^2 with $g_{a\gamma\gamma}$ can be understood qualitatively. As $P(\gamma \to a) \propto g_{a\gamma\gamma}^2$, a linear increase in $g_{a\gamma\gamma}$ leads to a quadratic increase in the amplitudes of modulations in the data. For a Poissonian process, the amount of data required to detect a fixed fractional deviation grows quadratically with the size of the deviation. So, roughly, a decrease in $g_{a\gamma\gamma}$ by a factor of two requires a sixteen-fold increase in the quantity of data for the same statistical sensitivity.

We also repeated this using a 100 ks exposure with an AGN that is twice as bright (to give a data sample that has a similar number of photons, but is substantially more piled up than the actual data). It is again the case that $g_{a\gamma\gamma} = 2 \times 10^{-12}\,\mathrm{GeV}^{-1}$ is strongly excluded, $g_{a\gamma\gamma} = 1 \times 10^{-12}\,\mathrm{GeV}^{-1}$ is not excluded, and $g_{a\gamma\gamma} = 1.5 \times 10^{-12}\,\mathrm{GeV}^{-1}$ is marginal.

We note that there are of course several differences between the MARX simulation and actual data processing. The MARX treatment of pileup is simpler than the real physics of deposited electron clouds from the interaction of photons with CCD chips. The MARX simulation also does not take into account Charge Transfer Inefficiencies present on the actual chips, which obstruct the flow of charge to the readout. However, these results imply that, in terms of bounding $g_{a\gamma\gamma}$, uncertainties due to pileup appear to be far smaller than uncertainties due to the cluster magnetic field.

The above bounds used MARX's simulation of *Chandra*'s optics. The most advanced simulation of *Chandra*'s optics is through the *ChaRT* simulator ([134]), rather than MARX. As *ChaRT* requires the original source spectrum to be manually uploaded to a website, it is not possible to automate this process to produce axion bounds (as every different ALP conversion template represents a different source spectrum). However, for each of the couplings $g_{a\gamma\gamma} = 1, 1.5, 2 \times 10^{-12}\,\mathrm{GeV}^{-1}$, we simulated some individual *ChaRT* spectra. For the spectra simulated, the results are similar to those of the MARX simulations. We plot in figure 4.19 a sample spectrum for $g_{a\gamma\gamma} = 2 \times 10^{-12}\,\mathrm{GeV}^{-1}$. The ALP induced modulations are clearly visible leading to a large badness-of-fit.

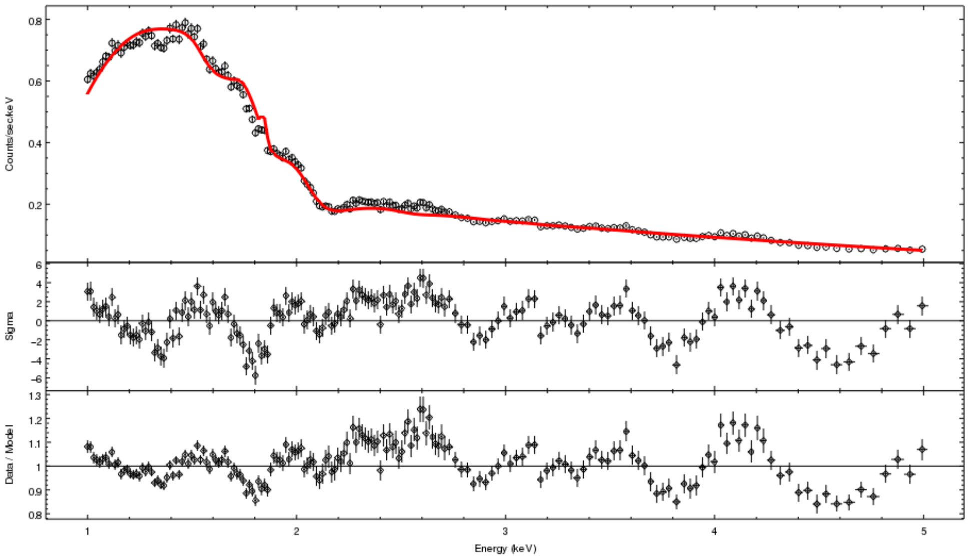

Figure 4.19: An example of simulated data using *ChaRT* and *MARX* reproducing the observational conditions for the ACIS-I edge observations, for an axion coupling of $g_{a\gamma\gamma} = 2 \times 10^{-12}\,\mathrm{GeV}^{-1}$. In this case the spectrum is directly fit with the sum of a power law and thermal `apec` emission (the ingredients to the simulation). Despite the pileup fraction of $\sim 15\%$, the axion-induced modulations are clearly visible and cannot be fit by any conventional astrophysical model.

4.2.8 *XMM-Newton* observations

There are two observations of NGC 1275 with significant exposure time, see Table 4.4. The first one, 0085110101, was taken in 2001 when the AGN had its lowest emissivity in observational history. The second observation, 0305780101, was taken when the emissivity was still relatively low but nevertheless almost twice as bright as in 2001. NGC 1275 is on-axis in both observations.

The 2001 observation was taken in full frame mode, while the 2006 observation was taken in extended full frame mode. This affects the frame time of the pn camera, which is 73.4 ms in full frame mode and 199.1 ms in extended full frame mode. For MOS, the frame time is 2.6 s. The pixel size of pn is 4.1 arcseconds and 1.1 arcseconds for MOS. This means that pn in extended full frame mode is more susceptible to pile-up than MOS. In general, pile-up is an issue for all *XMM-Newton* observations of NGC 1275 as we will discuss in the *XMM-Newton* analysis.

A significant difference between *XMM-Newton* and *Chandra* is the angular resolution. For both MOS and pn the radius of the disk containing 50% of the photons collected in the focal plane (Half Energy Width) is around 8.5 arcseconds at 1.5 keV,

Obs ID	Exposure [ks]	Year	Instrument	Location of NGC 1275
0085110101	53.08	2001	EPIC	Central
0305780101	123.3	2006	EPIC	Central

Table 4.4: The *XMM-Newton* observations used in this paper. The last column shows the location of NGC 1275 in the respective observation relative to the focal point.

while for *Chandra* it is much smaller at ~ 0.5 arcseconds. As the central region of the Perseus cluster is also intrinsically bright, this makes it harder to separate AGN and cluster emission for the *XMM-Newton* observations.

The effective area at 1 keV is $922\,\mathrm{cm}^2$ for MOS and $1227\,\mathrm{cm}^2$ for pn compared to $340\,\mathrm{cm}^2$ for *Chandra*. This allows MOS and pn combined to collect roughly 7 times more photons in a given observation time than *Chandra* (although this also increases the amount of pile-up). However, there is roughly 10 times more exposure time in the *Chandra* dataset and the significantly better angular resolution of *Chandra* allows a much better contrast to be attained between the AGN and the cluster emission. This leads to much better statistics in the *Chandra* dataset in terms of the total number of counts after background subtraction.

4.2.9 *XMM-Newton* Analysis

We reprocess the 2001 and 2006 datasets and removed flares. We find that there is a significant flare towards the end of the 2006 observational period and various shorter flares during the 2001 observation. Removing the polluted time intervals results in reducing the effective exposure time from 53 ks to 49 ks for MOS 2001, 25 ks to 7.4 ks for pn 2001, 123 ks to 117 ks for MOS 2006 and 76 ks to 49 ks for pn 2006.

For the extraction regions, we first choose circles of radii 13.8 arcseconds (MOS1 2001), 17.5" (MOS2 2001), 14.3" (MOS1 2006), 17.5" (MOS2 2006), 15.5" (pn 2001), and 19" (pn 2006). After checking with `epatplot` we find that pile-up is in general present in all the observations. Since the AGN was less bright by a factor ~ 2 in 2001, pile-up is slightly less of an issue for these observations. On the one hand, the smaller pixel size of 1.1 arcseconds of the MOS cameras makes it less susceptible for pile-up than the 4.1 arcseconds pixel size of the pn camera. On the other hand, pn's time resolution of 73.4 ms in full frame mode is advantageous with respect to pile-up compared to the 2.6 s time resolution of MOS. However, in extended full frame mode pn's 199.1 ms time resolution make it less advantageous with respect to pile-up compared to MOS. Note that after flare removal there are only 7.4 ks of pn data in full frame mode while there are 49 ks of more piled up data in extended full frame mode. As the MOS data is expected to be more sensitive, we do not present the pn data from either 2001 or 2006 and instead focus on the MOS data only.

In order to make a compromise between avoiding the most heavily piled-up regions and not too small signal over background ratio we choose the following inner annuli radii 2.5 arcseconds (MOS1 2001), 3" (MOS2 2001), 7.5" (MOS1 2006), 7.5" (MOS2 2006). For the background regions, we choose annuli of 13.8-22 arcseconds (MOS1 2001), 17.5-22.5" (MOS2 2001), 14.3-22.5" (MOS1 2006), 17.5-22.5" (MOS2 2006). The worse angular resolution of *XMM-Newton* compared to *Chandra* implies that the discrimination of the AGN emission against the cluster is relatively poor (for the 2006 data, the AGN:cluster data ratio is 1:3, whereas for *Chandra* spectra this ratio is not worse than 1.7:1).

We generate the spectra and response files, and use `epicspeccombine` to combine the MOS1 and MOS2 spectra of the 2001 and 2006 obervations, respectively. Comparison with using `mathpha`, `addrmf` and `addarf` shows that the results are within statistical errors. In order to avoid combining different systematic errors we do not combine the 2001 and 2006 data.

The spectra are fit using `Xspec` [135] and `Sherpa` [131], requiring a minimum count-rate per bin of about 1.5% of the overall count rate. The energy interval is 0.5–7.5 keV.

For NGC 1275, we use a spectral model of a power law as in our Chandra analysis

$$A(E) = K\,E^{-\gamma}\,, \tag{4.3}$$

where K is the normalisation in photons $\mathrm{keV}^{-1}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and γ is the photon index. To model a soft excess of the AGN, we add an `apec` model whose temperature we expect to be around 1 keV. The abundance parameter of this model is set to unity.[7] We also include the well known iron line at 6.4 keV in the cluster rest frame in our spectral fit.

To model photoelectric absorption, the source spectrum is multiplied by the `zwabs` model

$$M(E) = e^{-n_H \sigma(E)}\,, \tag{4.4}$$

where n_H is the hydrogen column density and σ is the photoelectric cross-section [145]. Photoelectric absorption is mostly relevant below approximately 1 keV if $n_H \gtrsim 10^{21}$ atoms cm^{-2} (the galactic column density is measured as $n_H = 1.5 \cdot 10^{21}$ atoms cm^{-2}).

---

[7] Note that around 1 keV this `apec` spectrum is dominated by atomic lines. Hence, there is a degeneracy between the abundance parameter and the normalisation of the `apec` model.

The fitted spectra are shown in Figures 4.20 and 4.21. While the fit is good (see Table 4.5), there are both fewer counts and a worse contrast against the cluster background than for the *Chandra* data. From the perspective of ALP constraints, we can say that modulations of the spectrum are not allowed beyond the 20% level – a weaker constraint than for the *Chandra* data.

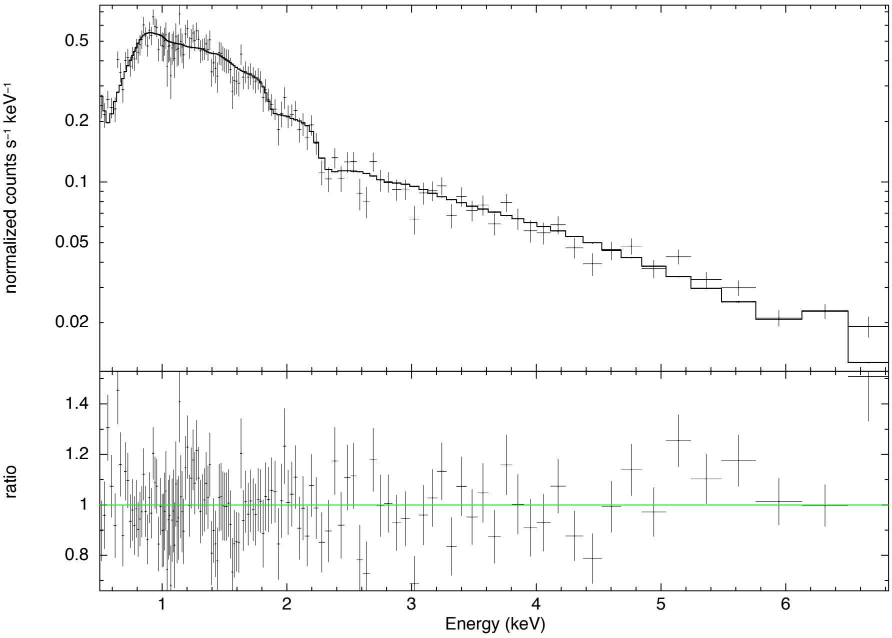

Figure 4.20: MOS 2001 spectral fit. The best-fit parameters are listed in Table 4.5.

4.3 Bounds from other point sources

The magnetic field model parameters we use for each cluster are shown in Table 4.6.

4.3.1 Quasar B1256+281 behind Coma

This quasar is located behind the Coma cluster at a redshift of $z = 0.38$. Its sightline passes through the entirety of the Coma intra-cluster medium (ICM). There are around 5000 counts from the source, of which around 10% can be attributed to the ICM (as the source is always off-axis, the *Chandra* Point Spread Function is degraded compared to an on-axis observation, increasing the level of contamination from thermal cluster emission).

Grouping counts so that there are at least 40 counts per bin, the quasar spectrum from 0.5 to 7 keV is well-fit by an unabsorbed power law with photon index 1.75 ± 0.04

	MOS 2001	MOS 2006
Exposure [ks]	49	118
Counts	$\sim$ 47000	$\sim$ 63000
γ	1.65 $\pm$ 0.03	1.65 $\pm$ 0.03
T_{soft}	0.84 $\pm$ 0.06	0.80 $\pm$ 0.06
n_H [10^{22} atoms cm^{-2}]	0.13 $\pm$ 0.03	0.13 $\pm$ 0.03
χ^2/ dof	154 / 131	205 / 178
Q-value	0.08	0.08

Table 4.5: Fit results for the *XMM-Newton* MOS observations from 2001 and 2006. These results are in agreement with the fit parameters found in [119] for the 2001 MOS data of NGC 1275.

Source	Cluster	$n_{e,0}$ (10^{-3}cm^{-3})	r_c (kpc)	β	B_0 μG	L_{total} (Mpc)
B1256+281	Coma	3.44	291	0.75	4.7	2
SDSS J130001.47+275120.6	Coma	3.44	291	0.75	4.7	2
NGC3862	A1367	1.15	308	0.52	3.25	1
IC4374	A3581	20	75	0.6	1.5	1
2E3140	A1795	50	146	0.631	20	1
CXOUJ134905.8+263752	A1795	50	146	0.631	20	2
UGC9799	A2052	35	32	0.42	11	1

Table 4.6: Parameters for the electron density and magnetic field models used for each of the clusters. All sources use $\eta = 0.7$, $L_{min} = 1\,\text{kpc}$ and $L_{max} = 17\,\text{kpc}$. For Coma the parameters are taken from [146]. For A1367 the β-model parameters come from [147] and the magnetic field from the article by M. Henriksen in [148]. For A1795 the central magnetic field is taken from [149], the β-model parameters from [150] and the central electron density from [151]. For A2052 the parameters are taken from [152] and [153] (correcting an error in the conversion of the core radius from arcseconds to kiloparsecs). For the poor cluster A3581 the central electron density is taken from [154, 155]. We could not find beta model parameters in the literature and have used illustrative values of $r_c = 75\,\text{kpc}$ and $\beta = 0.6$. For the central magnetic field we used the value for the poor cluster A194 [125] of $B_0 = 1.5\,\mu$G.

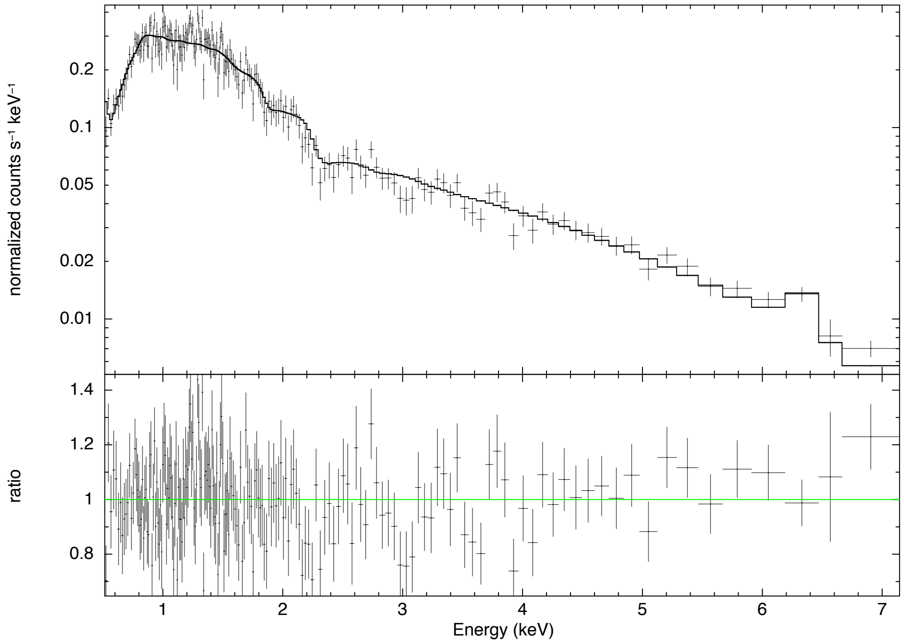

Figure 4.21: MOS 2006 spectral fit. The best-fit parameters are listed in Table 4.5.

(a reduced χ^2 of 0.88 for 96 degrees of freedom). There is no requirement for an Fe Kα line. The spectrum is plotted in Figure 4.22. No significant residuals are observed, and we can say that there are no ALP-induced modulations in the spectrum beyond the 20% level.

Simulating fake data with an ALP present in the spectrum as described in Chapter 3, the 95% confidence level bound is $g_{a\gamma\gamma} < 6 \times 10^{-12}\,\mathrm{GeV}^{-1}$.

4.3.2 Quasar SDSS J130001.47+275120.6 behind Coma

This is an even more distant quasar, at redshift 0.975. The sightline again passes through the entirety of the Coma ICM. There are around 3000 counts after background subtraction (around 3800 prior), and the resulting spectrum is well fit by the sum of an unabsorbed power law with photon index $\gamma = (1.80 \pm 0.05)$ and an Fe Kα line at 6.4 keV in the rest frame (inclusion of the Fe line gives an improvement of $\Delta\chi^2 = 5$ in the fit, and so is 2.2σ preferred). Fitting from 0.5 to 7 keV, and grouping counts so that there are at least 30 counts per bin, the reduced χ^2 is 0.99 for 95 degrees of freedom.

The spectrum is plotted in Figure 4.23. The relatively small number of counts means that we can only restrict ALP-induced spectral irregularities to the $\lesssim 30\%$ level, as there is not the statistical leverage to constrain beyond that.

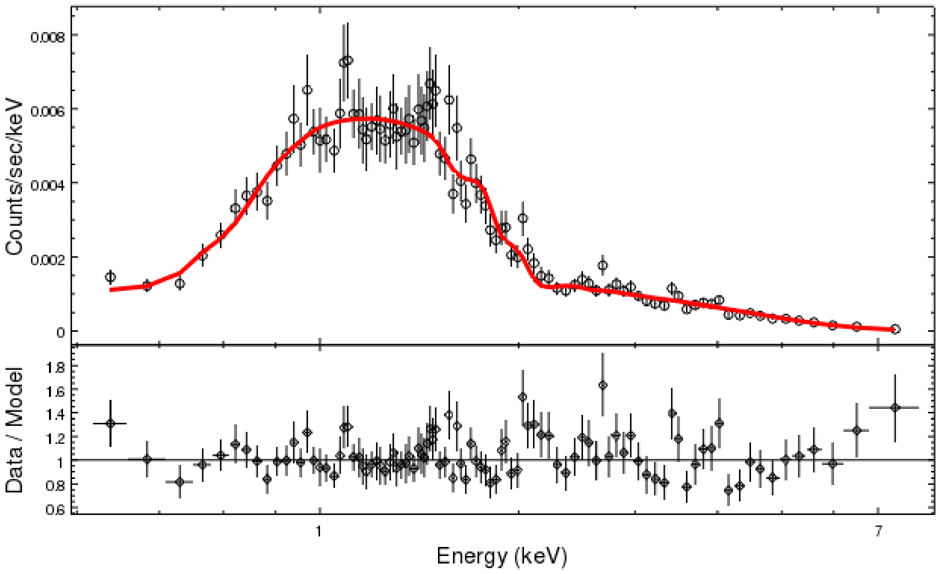

Figure 4.22: The stacked spectrum of the quasar B1256+281 behind Coma. The fit is an unabsorbed power law of photon index $\gamma = (1.75 \pm 0.03)$.

The small count allows no value of $g_{a\gamma\gamma}$ to be ruled out at the 95% confidence level. (A value of $g_{a\gamma\gamma} = 10^{-11}\,\mathrm{GeV}^{-1}$ is excluded at 87% confidence level.)

4.3.3 NGC3862 within A1367

The AGN NGC3862 within the cluster A1367 is characterised by a very soft power law (photon index 2.30 ± 0.03) absorbed by a column density of $N_H = 5 \times 10^{20}\,\mathrm{cm}^{-2}$, supplemented by a soft thermal component $T \sim 0.3\,\mathrm{keV}$). Grouping counts so that there are at least 50 counts per bin, the reduced χ^2 is 0.83 for 144 degrees of freedom, with a total of 21000 counts after background subtraction. The spectrum is plotted in Figure 4.24 and the resulting fit shows no sign of any significant spectral irregularities.

The low electron density within A1367 increases the efficiency of ALP-photon conversion (as it reduces the effective mass differential between the photon and the ALP). The large number of counts then allows good bounds to be obtained, $g_{a\gamma\gamma} < 2.4 \times 10^{-12}\,\mathrm{GeV}^{-1}$ at 95% confidence and $g_{a\gamma\gamma} < 2.9 \times 10^{-12}\,\mathrm{GeV}^{-1}$ at 99% confidence.

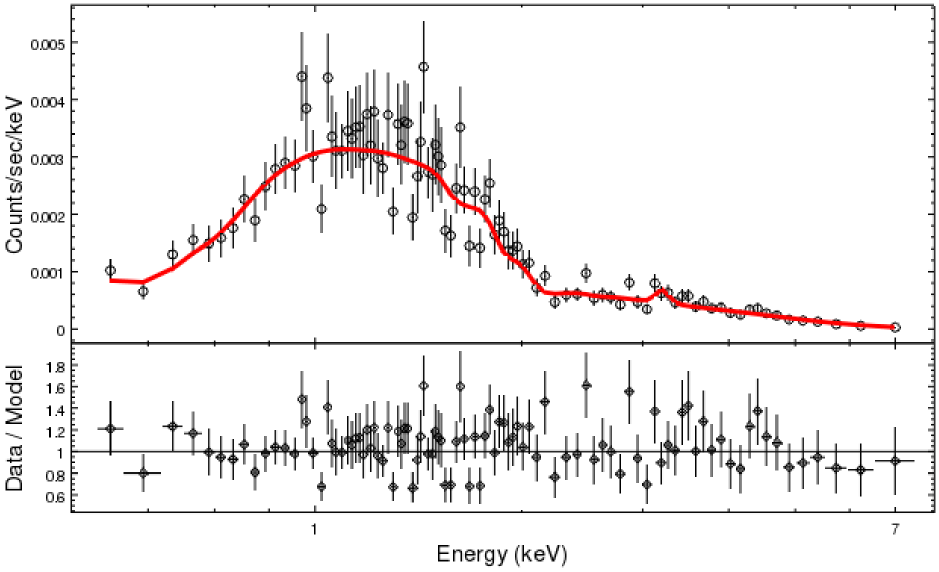

Figure 4.23: The stacked spectrum of the quasar SDSS J130001.48+275120.6 behind Coma at z=0.975. The fit is an unabsorbed power law of photon index $\gamma = (1.80 \pm 0.05)$ plus an Fe Kα line at 6.4 keV in the rest frame.

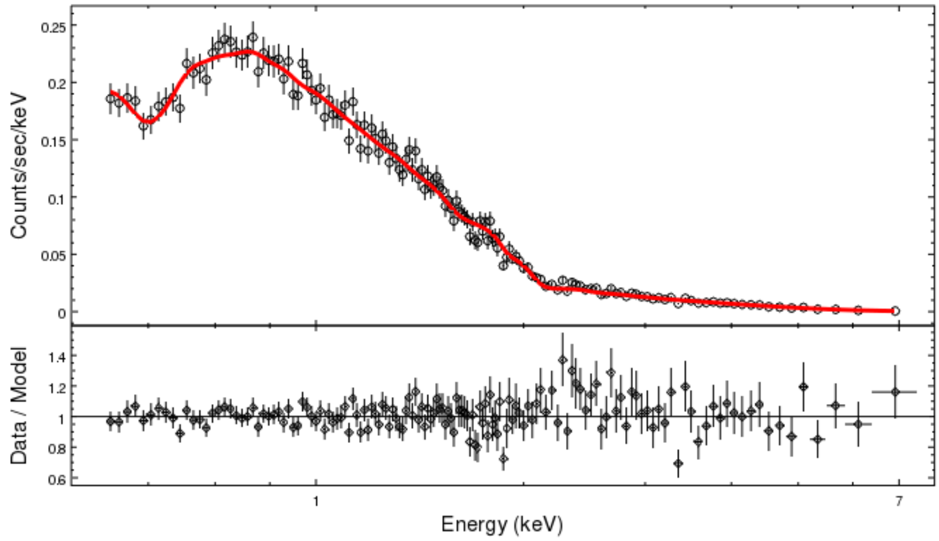

Figure 4.24: The stacked spectrum of the AGN NGC3862 in A1367. The fit is the sum of a power law of photon index $\gamma = (2.30 \pm 0.05)$ plus a soft thermal component with temperature 0.3 keV, absorbed by a column density of $N_H = 5 \times 10^{20}\,\mathrm{cm}^{-2}$.

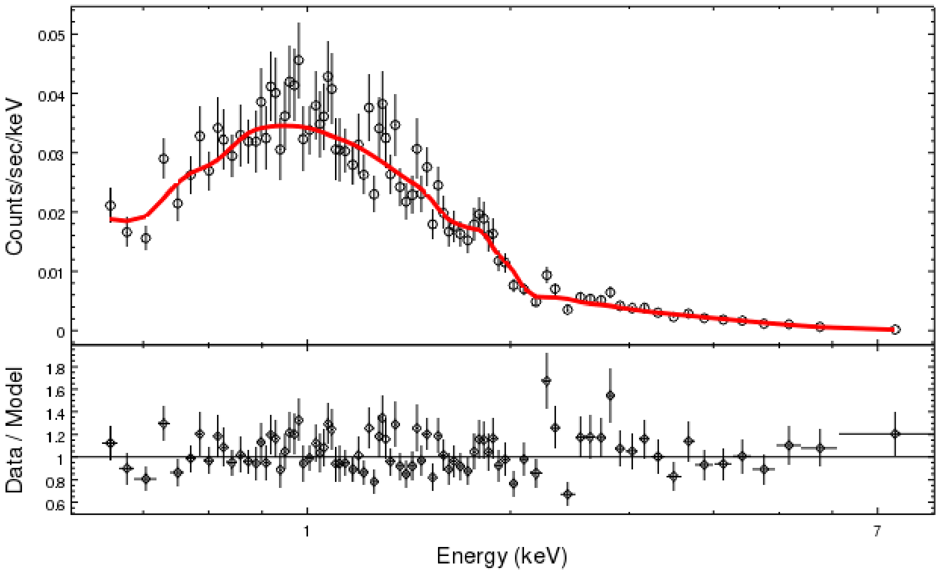

Figure 4.25: The stacked spectrum of the central cluster galaxy IC4374 located in A3581. A reasonably good fit is provided by an absorbed power law with $\gamma = 2.00 \pm 0.05$ and $N_H = (9 \pm 1.5) \times 10^{20}\,\mathrm{cm}^{-2}$. Grouping counts with at least 40 counts per bin, the reduced χ^2 is 1.13 for 82 degrees of freedom.

4.3.4 Central AGN IC4374 of A3581

This central AGN has around 4400 counts after background subtraction (4600 prior to background subtraction). A reasonably good fit is provided by an absorbed power law with $\gamma = 2.00 \pm 0.05$ and $N_H = (9 \pm 1.5) \times 10^{20}\,\mathrm{cm}^{-2}$. Grouping counts with at least 40 counts per bin, the reduced χ^2 is 1.13 for 82 degrees of freedom. The spectrum is plotted in Figure 4.25. As this is a very poor cluster, the central magnetic field is expected to be very weak. This reduces conversion efficiency, with the result that this source is unable to provide any constraints - a combination of the small number of counts and the weak magnetic field.

4.3.5 Central AGN of Hydra A

This system was studied in [112]. It is the central AGN (at redshift z=0.05) of the Hydra A galaxy cluster. There is a relatively poor contrast against the surrounding cluster, and we obtain around 2000 counts after background subtraction. There is a thick local absorbing column density, with a fit from 1.5 to 6 keV giving a best-fit value of $n_H = (2.64 \pm 0.47) \times 10^{22}\,\mathrm{cm}^{-2}$ and a power law index $\gamma = 1.63 \pm 0.20$

(consistent with the results of [112]). However, the low number of counts and weak contrast against the cluster means that ALP-photon interconversion is constrained only at the 30% level.

4.3.6 Seyfert galaxy 2E3140 in A1795

This is a bright unobscured AGN. Its redshift is 0.059, compared to a cluster redshift of 0.062. The radial velocity difference is $1000\,\mathrm{kms}^{-1}$, which is within the range of the A1795 velocity dispersion, and is consistent with 2E3140 being a bound member of the cluster A1795, with a sightline that passes through the intracluster medium.

However, we do not know the precise 3D location within the intracluster medium, and therefore the extent of the line of sight within the ICM. We assume a midway position.[8] The extracted spectrum contains around 78000 counts (of which around 1000 are ICM background). The spectrum from 0.5 to 6 keV is very well fit by the sum of a power law with photon index $\gamma = 2.11 \pm 0.01$, a soft thermal component with $T \sim 0.1$keV and a weak Fe Kα line at 6.4 keV in the rest frame (the Sherpa model `powlaw1d` + `xsapec` + `xszgauss`). No absorption is required in the fit (this is consistent with the small value of the galactic N_H value towards A1795). Grouping counts so that there are at least 500 in each bin, the overall fit is excellent with a reduced χ^2 of 0.98 for 103 degrees of freedom. The spectrum is plotted in Figure 4.26.

The large number of counts joined to the excellent fit results in strong bounds. Simulating fake data as in Chapter 3, we obtain $g_{a\gamma\gamma} < 1.5 \times 10^{-12}\,\mathrm{GeV}^{-1}$ at 95% confidence level and $g_{a\gamma\gamma} < 1.6 \times 10^{-12}\,\mathrm{GeV}^{-1}$ at 99% confidence level.

4.3.7 Quasar CXOUJ134905.8+263752 behind A1795

This $z = 1.3$ quasar is an obscured AGN with a total of around 5000 counts (5300 before background subtraction) arising from 985ks of *Chandra* observation time. It is fit by an absorbed power law, with a local absorption column density of $N_H = (1.0 \pm 0.1) \times 10^{22}\,\mathrm{cm}^{-2}$, and a power law photon index of $\gamma = 1.61 \pm 0.04$. No contribution from Milky Way absorption is required in the fit. Bins are grouped so that each bin has at least 40 counts. The overall fit between 0.5 and 7 keV is good, with a reduced χ^2 of 1.12 for 96 degrees of freedom. The spectrum is shown in Figure 4.27.

---

[8] This is perhaps supported by the large velocity relative to the cluster centre, as an object undergoing harmonic motion about a central source has maximal relative velocity at the midpoint of its oscillation.

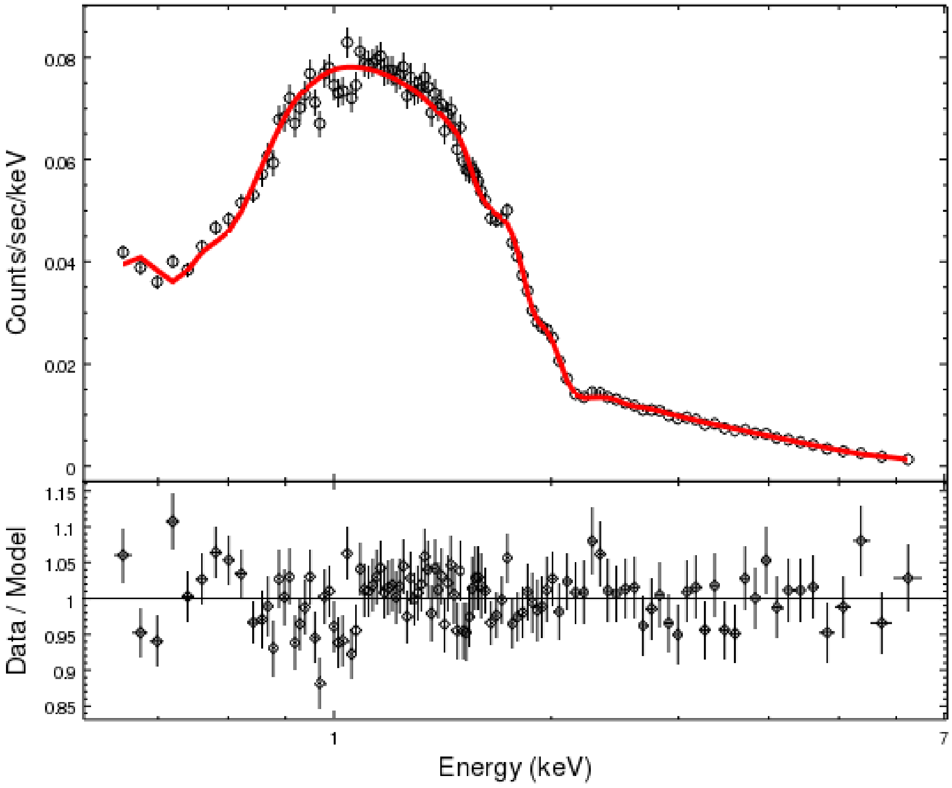

Figure 4.26: The stacked spectrum of the bright Seyfert 1 galaxy 2E3140 located in A1795. The fit is the sum of a power law with $\gamma = 2.11 \pm 0.01$, a soft thermal component with $T \sim 0.1\,\mathrm{keV}$ and a weak Fe Kα line at 6.4 keV in the rest frame (`powlaw1d` + `xsapec` + `xszgauss`). Grouping counts so that there are at leat 500 in each bin, the overall fit is excellent with a reduced χ^2 of 0.98 for 103 degrees of freedom.

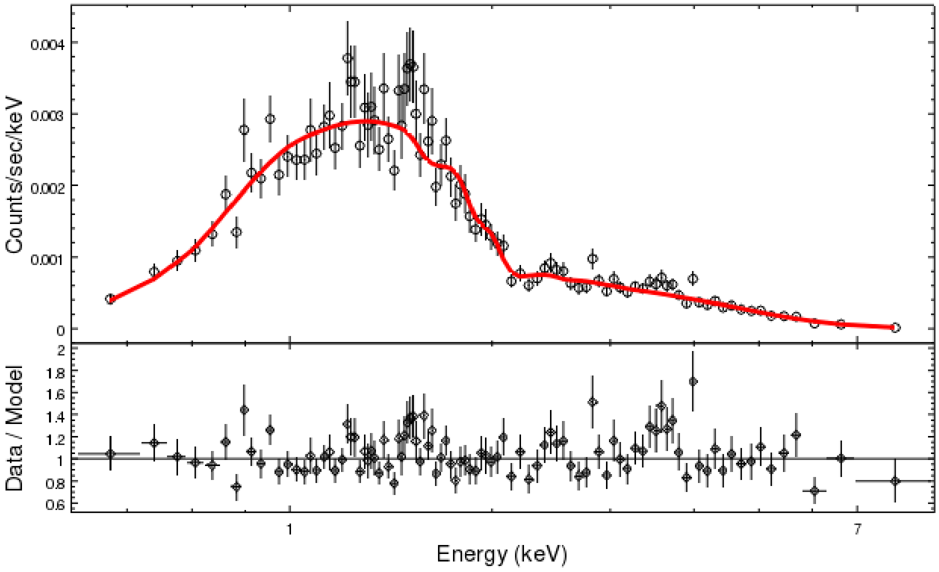

Figure 4.27: The stacked spectrum of the $z = 1.3$ quasar CXOU J134905.8+263752 behind A1795. The fit is to an absorbed power law, with a local absorption column density of $N_H = (1.0 \pm 0.1) \times 10^{22}\,\mathrm{cm}^{-2}$, and a power law photon index of $\gamma = 1.61 \pm 0.04$. Bins are grouped so that each bin has at least 40 counts. The overall fit between 0.5 and 7 keV is good, with a reduced χ^2 of 1.12 for 96 degrees of freedom.

The relatively small number of counts means that no strong bounds can be extracted for this source ($g_{a\gamma\gamma} < 10^{-11}\,\mathrm{GeV}^{-1}$ only at 75% confidence level).

4.3.8 UGC9799 in A2052

This central AGN of the cluster A2052 is well characterised by an unabsorbed power law with photon index $\gamma = 1.85 \pm 0.04$ supplemented by a soft thermal component with $T \sim 0.9 \pm 0.2\,\mathrm{keV}$. There are 4300 counts after background subtraction. Fitting between 0.5 and 7 keV, and grouping counts so that there are at least 30 counts per bin, the reduced χ^2 is 1.12 for 94 degrees of freedom. The soft spectrum, combined with the relatively low number of counts, implies that ALP-induced modulations are only excluded above the 20% level. The spectrum is plotted in Figure 4.28.

As with A3581, the relatively small number of counts means that this source is unable to provide useful constraints, and $g_{a\gamma\gamma} = 10^{-11}\,\mathrm{GeV}^{-1}$ is not excluded by this source.

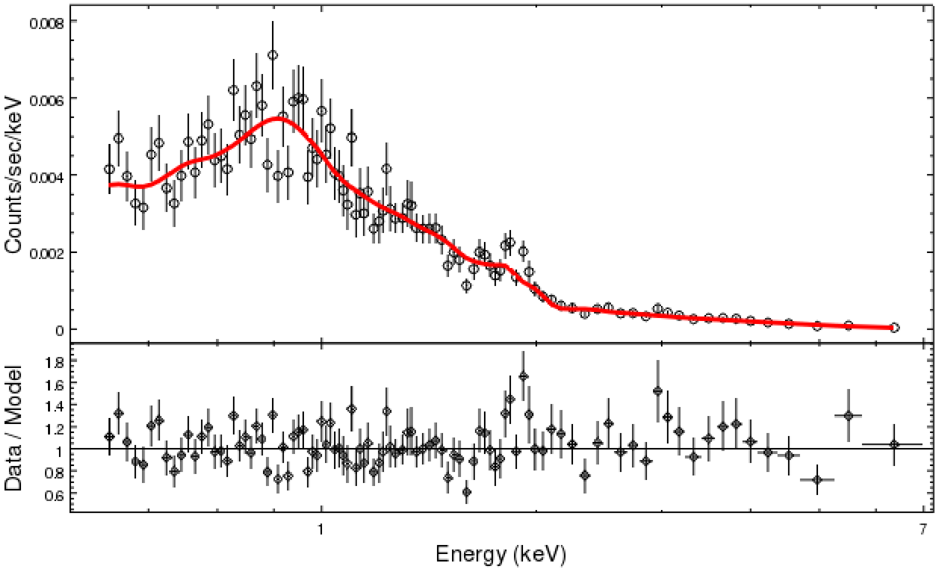

Figure 4.28: The stacked spectrum of the central cluster galaxy UGC9799 located in A2052. The fit is to an unabsorbed power law with photon index $\gamma = 1.85 \pm 0.04$ supplemented by a soft thermal component with $T \sim 0.9 \pm 0.2\,\mathrm{keV}$. There are 4300 counts after background subtraction. Fitting between 0.5 and 7 keV, and grouping counts so that there are at least 30 counts per bin, the reduced χ^2 is 1.12 for 94 degrees of freedom.

4.4 Simulated bounds from *Athena* Observations

In terms of estimating bounds on $g_{a\gamma\gamma}$ we use the same method as previously applied with *Chandra* data [1]. This allows for a direct comparison between the capabilities of *Chandra* and *Athena* in terms of placing bounds.

We simulate *Athena* observations of NGC 1275, using two models for the photon spectra of the AGN. The first is a standard spectrum without ALPs, and the second is a model with the same spectrum multiplied by the photon survival probability distribution as introduced in Chapter 3. Using simulations of the X-IFU detector response, we fit spectra generated with ALP-photon conversion to the model without ALPs (Model 0), and then we compare this fit to a fit of the spectrum generated without ALPs to the same model. To allow for the uncertainty in the magnetic field configuration along the line of sight, we repeat this analysis using many different randomly generated magnetic fields.

The two photon spectra that we model are:

1. Model 0: An absorbed power law plus thermal background:

$$F_0(E) = (AE^{-\gamma} + \texttt{BAPEC}) \times e^{-n_H \sigma(E,z)}, \tag{4.5}$$

   where A and γ are the amplitude and index of the power law, E is the energy, n_H is the equivalent hydrogen column, $\sigma(E,z)$ is the photo-electric cross-section at redshift z, and `BAPEC` is the standard plasma thermal emission model.

2. Model 1: An absorbed power law plus thermal background, multiplied by a table of survival probabilities for photons of different energies:

$$F_1(E, \mathbf{B}) = (AE^{-\gamma} + \texttt{BAPEC}) \times e^{-n_H \sigma(E,z)} \times P_{\gamma \to \gamma}(E(1+z), \mathbf{B}, g_{a\gamma\gamma}) \; . \tag{4.6}$$

The index of the power law was set based on the best-fit value from the cleanest *Chandra* observations of NGC 1275, and its normalisation was determined based on the *Hitomi* 230 ks observation of Perseus in 2016 [156]. As the AGN in 2016 was roughly twice as bright as in 2009 and it has previously exhibited large historical variation [122], it may be again much brighter (or dimmer) in 2028, which would affect both the contrast against the cluster background and also the observation time required to achieve a certain constraint on $g_{a\gamma\gamma}$.

The 2016 *Hitomi* observation also constrained the temperature, abundances and velocity dispersion of the cluster thermal emission to a high degree of accuracy [156]. For the spectral shape of the cluster background, we used the single-temperature

Model	parameter	symbol	value
`zwabs`	nH column density	n_H	$0.24 \times 10^{22}\text{cm}^{-2}$
	redshift	z	0.0176
`powerlaw`	index	γ	1.8
	normalisation	A	9×10^{-3}
`bapec`	temperature	kT	3.48 keV
	abundances		0.54 solar
	velocity dispersion	v	178 m s^{-1}
	normalisation	N	9×10^{-4}

Table 4.7: Parameters of the absorbed power law describing the spectrum of NGC 1275, and the thermal model of the cluster background.

`bapec` model that was a good fit to the *Hitomi* spectrum across its field of view. While this single-temperature model is unlikely to be a good fit for the background contiguous to the AGN, it represents a useful proxy for the actual background that can only be determined at the time. The normalisation of the background was set by extracting a circular region of the cluster emission close to the AGN from the *Chandra* observations, of radius equal to the angular resolution of *Athena*, and determining the best fit. All model parameters are shown in Table 4.7.

As for the study with *Chandra*, we take the central magnetic field value as $B_0 \sim 25\,\mu$G, following [123]. We also assume that B decreases with radius as $B(r) \propto n_e(r)^{0.7}$. As there is not a direct measurement of the power spectrum and coherence length for the Perseus magnetic field, we base the model on those inferred for the cool core cluster A2199 [128].

The electron density n_e has the radial distribution found in [119]:

$$n_e(r) = \frac{3.9 \times 10^{-2}}{[1 + (\frac{r}{80\,\text{kpc}})^2]^{1.8}} + \frac{4.05 \times 10^{-3}}{[1 + (\frac{r}{280\,\text{kpc}})^2]^{0.87}}\,\text{cm}^{-3}. \tag{4.7}$$

The magnetic field is generated over 300 domains, whose lengths are drawn from a Pareto distribution between 3.5 kpc and 10 kpc with power 2.8. In each domain the magnetic field and electron density are constant, with a random direction of $\mathbf{B}$. We then calculate the survival probability of a photon passing through this region, as described in [87].

The simulations were performed using the Simulation of X-ray Telescopes (`SIXTE`) code, a multi-instrument simulation package. It aims to offer an end-to-end simulation, i.e. the full detector chain from the source to the final data. It models the telescope's vignetting, ARF and PSF, and X-IFU's response, event reconstruction and pileup [157].

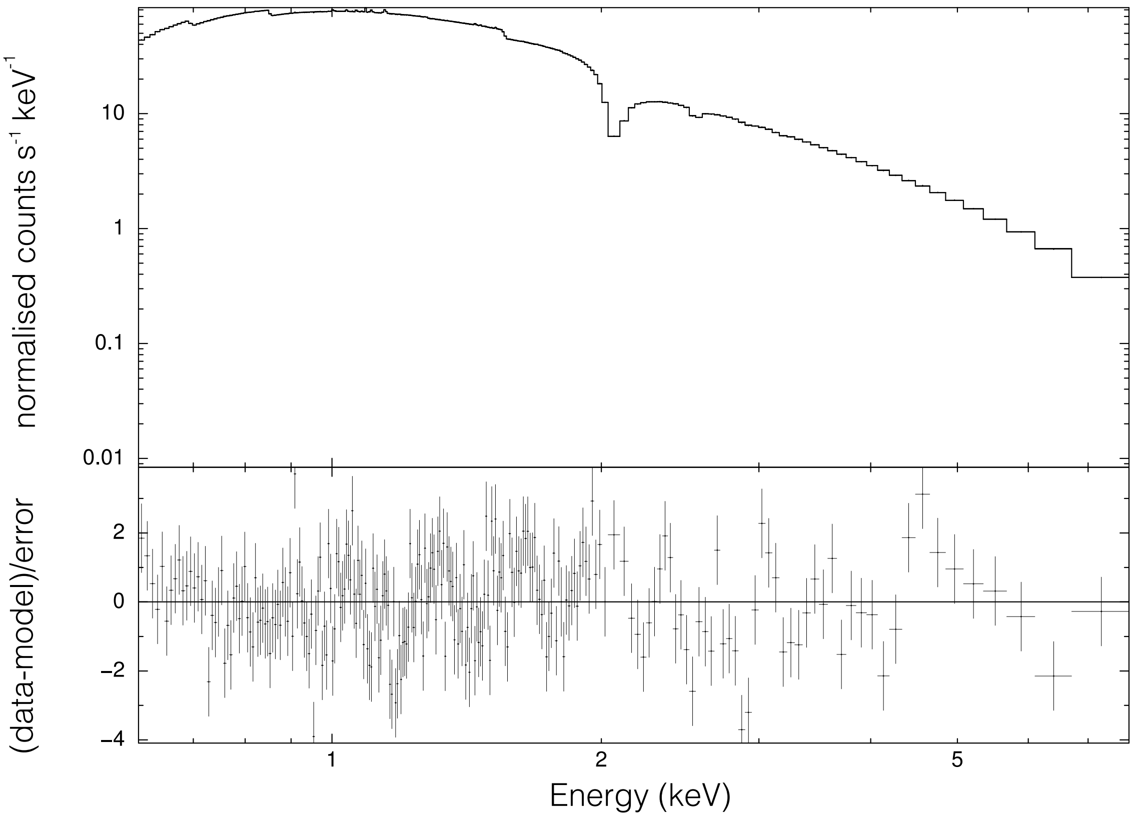

Figure 4.29: A simulated 200 ks dataset for NGC 1275 with $g_{a\gamma\gamma} = 3 \times 10^{-13}\,\mathrm{GeV}^{-1}$, and its fit to Model 0. The characteristic ALP-induced modulations are apparent.

The spectrum of NGC 1275, and the cluster background, were modelled in `XSPEC`[9] as an absorbed power law plus a thermal component, `zwabs` $*$ (`powerlaw` $+$ `bapec`). This spectrum, either multiplied by the photon survival probabilities or not, was converted to the SIMPUT[10] file format using the command `simputfile`. The mirror and detector response were modelled with `xifupipeline`, using the ARF file `athena_xifu_1469_onaxis_pitch249um_v20160401.arf`. We use the RMF file `athena_xifu_rmf_v20160401.rmf`. This generated an event FITS file, which was then converted into a PHA file using `makespec`. We produced a fit to this spectrum in `XSPEC`, using the Levenberg-Marquardt fitting method to calculate the reduced χ^2. Figure 4.29 shows one simulation for $g_{a\gamma\gamma} = 3 \times 10^{-13}\,\mathrm{GeV}^{-1}$ and its fit to an absorbed power law.

We use the following procedure to determine whether a particular value of $g_{a\gamma\gamma}$ is excluded: we varied the ALP-photon coupling $g_{a\gamma\gamma}$ from $g_{a\gamma\gamma} = 5 \times 10^{-13}\,\mathrm{GeV}^{-1}$ to

---

[9] https://heasarc.gsfc.nasa.gov/xanadu/xspec/manual/manual.html

[10] http://hea-www.harvard.edu/heasarc/formats/simput-1.0.0.pdf

$g_{a\gamma\gamma} = 1 \times 10^{-13}\,\mathrm{GeV}^{-1}$, with stepsize $0.5 \times 10^{-13}\,\mathrm{GeV}^{-1}$. As the bound is dependent on uncertainties in the magnetic field strength of a factor of 2, and we are only using simulated data, we do not consider step sizes smaller than this. For each $g_{a\gamma\gamma}$:

1. Generate 50 configurations of the magnetic field B_i.
2. Use the B_i to calculate the survival probability $P_{\gamma\to\gamma}$ along the line of sight for different photon energies (as done in [87]). We calculate for 8000 equally spaced photon energies in the range 0.01–10 keV.
3. Combine each $P_{\gamma\to\gamma}$ with the AGN spectrum.
4. Generate 10 fake PHAs for each spectrum, providing 500 fake data samples in total.
5. Fit the fake data to Model 0, and calculate the reduced chi-squareds χ_1^2.
6. Generate 100 fake PHAs based on Model 0, and compute the average of their reduced chi-squareds χ_0^2. Assuming the absence of ALPs, this represents the expected quality of the fit to the single real data set. If the actual data is a poor fit for some reason, then this will weaken the level of the resulting bounds that we can produce.
7. Determine the percentage of fake data sets that have a reduced chi-squared $\chi_1^2 < \max(\langle\chi_0^2\rangle, 1)$. If this is true for fewer than 5% of the data sets, the value of $g_{a\gamma\gamma}$ is excluded at 95% confidence.

For a simulation of 200 ks of data with the nominal mirror configuration, we derive a projected bound of $g_{a\gamma\gamma} \lesssim 1.5 \times 10^{-13}\,\mathrm{GeV}^{-1}$ at 95% confidence. At 99% confidence the bound is $g_{a\gamma\gamma} \lesssim 2.5 \times 10^{-13}\,\mathrm{GeV}^{-1}$. We plot these results in Fig. 4.30 alongside published data limits. This represents an order of magnitude improvement over the bound derived from the 200 ks of Chandra ACIS-I observations in [1]. We also find that even a short 10 ks observation will lead to an improved bound of $g_{a\gamma\gamma} \lesssim 4.5 \times 10^{-13}\,\mathrm{GeV}^{-1}$.

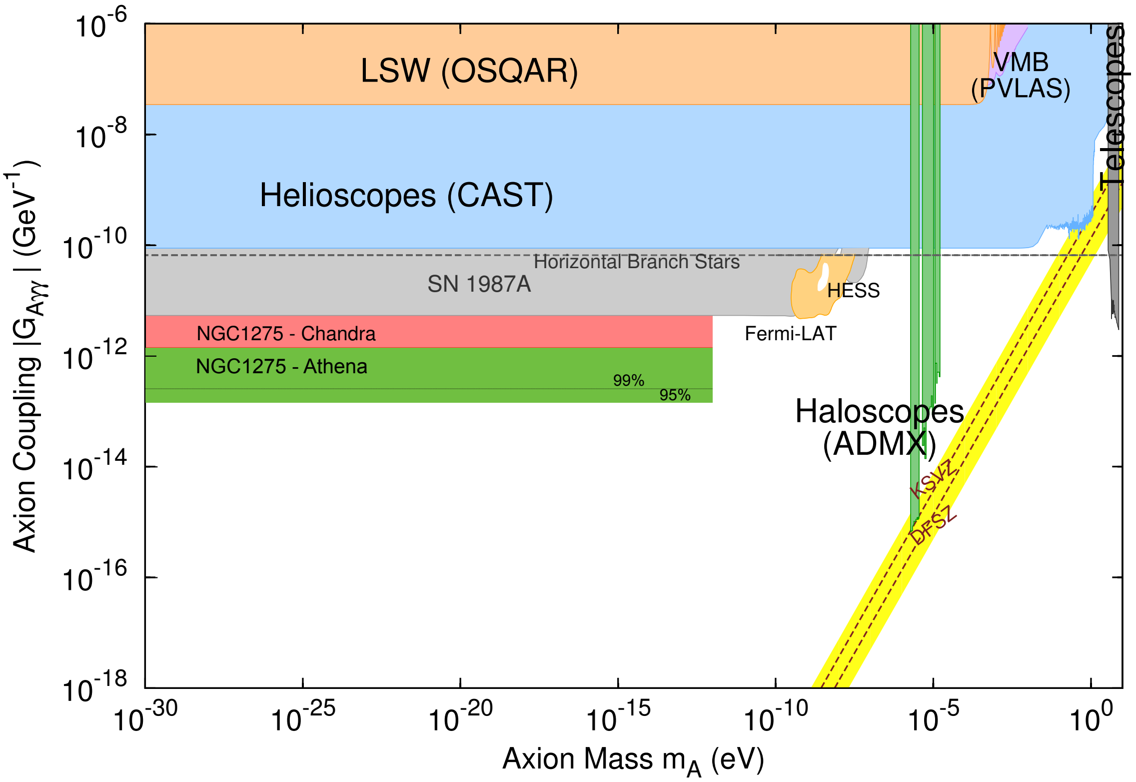

Figure 4.30: Overview of exclusion limits on axion couplings vs mass. For axion masses $m_a \sim 10^{-12}\,\mathrm{eV}$ then ALP-photon conversion can enter a resonant regime, with the potential of stronger bounds around this critical mass. We do not perform a detailed study of the resonant regime in this work and focus only on the low-mass region. Full references can be found in the Particle Data Group review on *Axions and other similar particles* [48].

Chapter 5

Discussion

We were able to derive bounds on $g_{a\gamma\gamma}$ in the mass range $m_a \lesssim 10^{-12}\,\mathrm{eV}$ from a number of X-ray observations of point sources in galaxy clusters. The *Chandra* observations of NGC 1275 in the Perseus Cluster produced a bound of $g_{a\gamma\gamma} \lesssim 1.4 \times 10^{-12}\,\mathrm{GeV}^{-1}$ with a central magnetic field strength of $25\,\mu\mathrm{G}$ in the model. This utilised the combined 200 ks of observations (ObsID 11713, 12025, 12033, 12036) where the image of the AGN was on the edge of the CCD detector chip. The *Chandra* observations of 2E3140 in the cluster A1795 produces a similar bound of $g_{a\gamma\gamma} \lesssim 1.5 \times 10^{-12}\,\mathrm{GeV}^{-1}$. The *Chandra* observations of the AGN NGC3862 behind the Coma Cluster produce a bound of $g_{a\gamma\gamma} \lesssim 2.4 \times 10^{-12}\,\mathrm{GeV}^{-1}$.

Simulations of the future *Athena* mission show that an observation of NGC 1275 of 200 ks would produce a bound of $g_{a\gamma\gamma} \lesssim 1.5 \times 10^{-13}\,\mathrm{GeV}^{-1}$, an order of magnitude better than the *Chandra* observations used.

5.1 Comments on NGC 1275

The ACIS-I and ACIS-S observations with Chandra generate over half a million X-ray counts from the central AGN – a dataset of extraordinary quality for searching for spectral irregularities. The 180 ks of observation time with *XMM-Newton* allows us to cross-check our results, but does not provide the same stringent bounds. This is despite a number of technical issues in the *Chandra* data of NGC 1275, including thermal components in the modelling, pileup, and residuals appearing in the data. We discuss these factors below.

5.1.1 Thermal component

Why is a soft thermal component necessary for a good fit in the ACIS-S observations but not in the ACIS-I observations? There are two main reasons. First, the effective area of ACIS-S has more support at low energies than for ACIS-I. The net result is that for ACIS-S the peak of the observed photon spectrum is around 1 keV, whereas for the ACIS-I observations the peak is around 1.6 keV. This implies that the ACIS-S observations have a proportionally higher fraction of low-energy photons, and so are more sensitive to the presence of a soft thermal component.

Secondly, the ACIS-S and ACIS-I observations were taken at different times - the former in 2002 and 2004 and the latter in 2009. Over this period the intrinsic brightness of the AGN, as measured by the power law component, has increased significantly. Assuming that the soft thermal component arises differently from the hard power law, and that the absolute strength of the soft thermal component has remained the same over this period, its relative significance would be much greater in the ACIS-S observations.

We have checked that one can add a soft thermal component to the ACIS-I spectra, of the same absolute magnitude as for the ACIS-S spectra, without affecting the quality of the fit. As for the ACIS-I observations, the thermal component neither significantly improves or worsens the fit, and does not affect the above results, we have not included it explicitly for the ACIS-I data.

5.1.2 Pileup and Bounds

Before we describe the data in detail, we comment on the relationship between the accuracy of the fit and the bounds on ALP parameters that we are able to produce. This is relevant as the data involves a bright point source at the middle of a bright and complex continuum background (the centre of the Perseus cluster).

Constraints on ALPs arise because they induce modulations that would make the conventional fit (an absorbed power law) a bad one. The constraints on ALPs come from requiring that the badness-of-fit that ALPs would induce is no worse than that present in the actual data – a coupling $g_{a\gamma\gamma}$ is excluded when it would lead to a fit that is worse than actually occurs. To place bounds on ALPs in this way, it is not necessary that the actual data be a perfect fit. It is true that the better the fit of the actual data to an absorbed power law, the stronger are the bounds that can be placed on ALPs. However, even if the actual data is only – for example – within 10%

of the standard fit, then this still implies the exclusion of ALP couplings large enough to give 30% modulations in the data.

This method of producing bounds is conservative. The simulated data we use when considering the effects of ALPs is more idealised than for real data. In particular, the simulated data does not account for any possible observational or instrumental inaccuracies arising from the complexity for the source, such as pileup from a high photon flux or possible mis-subtraction of the cluster background when obtaining the AGN spectrum.

On the other hand, in the actual data there are various observational, instrumental or astrophysical effects that could cause the standard fit to fail to be a perfect description of the measured data, even without any exotic physics. These include

1. As mentioned above, pile-up can lead to photons being registered at the wrong energy. This can lead to fake excesses at the location of effective area dips.
2. The X-ray emission at the centre of the Perseus cluster has an intensity that varies spatially away from the nucleus. As the background regions differ from the source region used, the process of background subtraction may lead to residual cluster emission being present in addition to the AGN spectrum.
3. The response files may not deal fully accurately with cases where the centre of the image has significant pile-up.

We will exclude ALP couplings for which idealised simulated data is still a worse fit compared to the noisier actual data. ‘Noise’ in the actual data reduces the bounds that can be placed on ALPs (as small genuine ALP-induced modulations could hide under the noise), but still robustly allows the exclusion of large ALP-induced modulations. The bounds that we place here are therefore conservative, but with the potential of further improvement from better modelling of the actual data or through observing NGC 1275 with cleaner settings (such as a shorter read-out time).

5.1.3 Analysis of residuals in *Chandra* NGC 1275 data

No residuals in the spectra exceed 10%. At the 10% level, the *Chandra* spectra presented in the main body of the paper show two main features departing from a power law: an excess at 2–2.2 keV, present in all *Chandra* observations with overwhelming statistical significance, and a deficit 3.4–3.5 keV that is not as strong but is still present at almost five (local) sigma in the ACIS-I edge data. These features ensure that an absorbed power law (plus thermal component) is not a good fit to the data.

Here we discuss potential instrumental, astrophysical or new physics explanations.

The 2.2 keV residual

Despite the overwhelming statistical significance of the feature at 2-2.2keV, it is at the same location of a sharp effective area edge from the Iridium coating of the mirrors. As pile-up arising from high flux levels can generate fake excesses at the location of such edges, we associate the feature with this edge and do not discuss it further.

The sharp drop in the effective area at this energy means that there is a danger of contamination from the many photons at neighbouring energies. We therefore carefully consider instrumental effects that could account for this excess. For completeness, we also briefly comment on possible astrophysical explanations.

Instrumental effects

The most obvious source of contamination in the spectra as a whole is pile-up. As noted in Section 3.5.1, pile-up contamination at energy E is dependent on the effective area at energies E_i, where E_i denotes the individual photon energies leading to a piled-up event at energy E. We therefore expect a higher pile-up fraction in an effective area dip.

We reiterate that accounting for pile-up contamination in data is highly non-trivial, as two completely different effects must be modelled (grade migration and energy migration). As we do not have knowledge of the "true" spectrum, we cannot do this analytically. We attempted to mitigate the effect of pile-up by cutting out the central region of the AGN image. This did not remove the excess, as it remained at the 10% level. However, as pile-up will alter the PSF of the detector, this does not disprove the pile-up explanation. Nor does the failure of the model *jdpileup* to account for the excess: fitting data to a power law plus a pile-up model is highly degenerate between a hard spectrum with little pile-up, or a soft spectrum with a lot of pile-up. We therefore retain pile-up as a possible explanation.

Another possible instrumental explanation is a miscalibration of the effective area at high fluxes. A miscalibration at the level of 10% would account for all of the residual. However, it is unclear why the miscalibration should be the same in ACIS-S and ACIS-I, and at different locations on the ACIS-I chip. Despite this, the possibility that effective area miscalibration is a contributory factor cannot be ruled out.

One final possibility that we consider is gain miscalibration. As there are many more photons at 1.8-2.0 keV than at 2-2.2 keV, it only requires a few photons at 1.8-2.0

keV to be misidentified as 2.0-2.2 keV to create a significant excess. However, thermal emission lines in the cluster background provide no evidence of gain miscalibration, and modelling gain directly in xspec produces no improvement in the fit. We therefore do not consider this to be a likely explanation.

Astrophysical explanations

One possibility is that the feature we observe is not a feature of the AGN, but instead arises from a mis-subtraction of the background. However, the brightness of the AGN, combined with the superb angular resolution of *Chandra*, means that we are able to obtain high ratios of AGN vs. cluster emission. In particular, in both ACIS-I edge and central observations we are able to achieve an AGN to cluster contrast reaching up to 15:1 (and even reaching 60:1 when we model pile-up directly) – making it impossible that a 10% localised excess could arise from the cluster emission.[1]

Another possibility is fluorescent emission from a neutral Kα line from Sulphur at 2.27 keV, arising from the accretion disc when illuminated by hard X-rays. However the energy of excess is too low, and its strength too large, to make this a reasonable explanation in comparison with the Fe Kα at ~ 6.3 keV.

The 3.5 keV deficit

We now consider possible explanations for the feature at 3.4–3.5 keV, present as a deficit of data compared to model. This is present at very high significance in the ACIS-I edge observations. Given the smaller data samples present there, the ACIS-I midway and ACIS-S observations are compatible but are not significant in themselves.

While this feature lacks the statistical significance of the one at 2–2.2 keV, it is harder to come up with instrumental or astrophysical explanations for such a deficit. Compared to the 2–2.2 keV region, at 3.4–3.5 keV the *Chandra* effective areas are smooth functions of energy, and in general this is a clean part of the spectrum, so the previous arguments for pile-up contamination or effective area miscalibration do not apply here.

The simplest explanation is as a statistical fluctuation – although as this deficit is present as a local $\sim 4.5\sigma$ deficit in the ACIS-I edge observation, such an explanation is problematic. It is also notable that this occurs at precisely the same energy as for the diffuse excess found in [141, 142].

---

[1] Although the spectra presented for the ACIS-I edge observations involve an AGN to cluster ratio of 6.4:1, by reducing the extraction region we are able to increase this ratio at the cost of only a small reduction in the overall photon count.

In terms of conventional astrophysical explanations, the most obvious candidate for this deficit would be an atomic absorption line. However it is difficult to see how this could work, as there are no strong lines around this energy. Even more seriously, absorption along the sightline to NGC 1275 dominantly arises from the Milky Way – this is evidenced by the fitted values of n_H, which are all consistent with the galactic value towards Perseus of $n_H \simeq 1.5 \times 10^{21}\text{cm}^{-2}$. This implies that, if the feature at 3.4–3.5 keV did arise from atomic absorption, a similar feature should also arise for the continuum cluster spectrum of Perseus, as the galactic n_H values have degree-scale gradients and so n_H is approximately the same for all sightlines to Perseus. However, no such absorption feature at 3.5 keV is detected in the cluster spectra of Perseus — indeed, precisely the opposite is found and instead a significant excess is found at 3.5 keV [141, 142] (for a review see [158]).

The most plausible conventional explanation of the 3.4–3.5 keV deficit is then as a localised detector effect or statistical fluctuation for the ACIS-I edge observations, combined with mild downwards statistical fluctuations for the ACIS-S and ACIS-I central observations. However, we also note that inspection of the background region for the ACIS-I edge observations does not show any deficit around 3.4–3.5 keV.

New physics explanations

Given the observations of [141, 142] of an unidentified line at 3.5 keV in diffuse emission from the Perseus cluster, it is an interesting fact that we observe a deficit at ~ 3.5 keV in our observations of the NGC 1275 AGN at the centre of the cluster. This of course may be just a coincidence, but in the context of new physics models we mention two ways that these facts could be more than simply a coincidence.

The first involves models of excited dark matter invoked for the 3.5 keV line, where dark matter has a resonance at an energy ~ 3.5 keV above its ground state (for example as in [159, 160, 161, 162, 163, 164]). In this case there is then an absorption cross-section of $E \sim 3.5$keV photons on dark matter. While for an isotropic initial distribution of photons such absorption and re-emission would not affect the photon spectrum, for a directional beam dark matter absorption will result in an absorption hole in the spectrum (the presence of an absorbing torus around an AGN ensures its outward radiation is indeed directional).

Would such an effect have been observed already elsewhere? We do not see why. In this scenario, the relevant quantity determining the fractional absorption rate is the dark matter column density along the line of sight. It is entirely plausible that the dark matter column density towards the NGC 1275 AGN is larger than for almost any other direction in the universe. This is because the emission all originates very close to the central AGN, and so the column density is sensitive to not just the Perseus cluster, but also the central cluster galaxy NGC 1275 right down to any sub-pc level dark matter spikes close to the central supermassive black hole. The effect we observe is not large – a 10% reduction over around 100 eV in width – and requires a spectrum with $\mathcal{O}(10^5)$ counts for a statistically significant detection. With a smaller dark matter column density, the effect would reduce to an unobservable $\mathcal{O}(1\%)$ effect.

There is a second possible connection to the 3.5 keV line. An attractive scenario for the 3.5 keV line involves decay of dark matter to 3.5 keV ALPs, which then convert to photons through axion-photon conversion in the cluster magnetic field [165, 166, 167, 168]. In this scenario the strength of the 3.5 keV line depends on the efficiency of ALP-photon interconversion – and so broadly is expected to be larger in regions with large magnetic fields extended over wide areas: for example, in galaxy clusters as opposed to galaxies. The 3.5 keV line is observed to be stronger towards the centre of Perseus than for other clusters; one way this could arise is if it fortuitously happens that ALP-photon interconversion is particularly efficient around 3.5 keV for sightlines towards the centre of Perseus. In this case, the presence of a deficit of $E \sim 3.5\,\text{keV}$ photons from NGC 1275 and an excess of $E \sim 3.5\,\text{keV}$ photons from the cluster as a whole could come from the same underlying physics – efficient photon-ALP interconversion at energies $E \sim 3.5\,\text{keV}$ along sightlines towards the centre of Perseus.

ALP Interpretation of Features

It is possible that these features are the result of the searched-for ALP oscillations in the data. While the presence of the largest feature at the Iridium edge forbids us from taking this interpretation too seriously, we estimated the strength of ALP-photon coupling that could induce oscillations like this. This gave a value $g_{a\gamma\gamma} \sim 1 \times 10^{-12}\,\text{GeV}^{-1}$. We can see therefore that *Athena* will comfortably rule out (or confirm) an ALP origin for these features.

5.2 Limitations of the study

A major source of uncertainty for this methodology is the magnetic field modelling. In particular, the experimental error on the central magnetic field strength produces a similar amount of error for $g_{a\gamma\gamma}$. Furthermore, a lack of 3D information about the magnetic field restricts our analysis to a statistical one, based on the likelihood that such a coupling would produce a given fit. Future measurements of the properties of the magnetic field from the Square Kilometer Array will substantially improve the error on the bounds. However, the statistical nature of the bound will remain without a specific 3D model of the magnetic field.

As we rely on archival data, the sources used here are not necessarily intrinsically optimal – for any source we rely on the existence of sufficient observational time to produce a usable number of photons. As one example, we use in this paper a $z = 1.3$ quasar that shines through the cluster A1795. This is usable solely because there is a total of 985ks of *Chandra* time including it in the field of view. A much shorter exposure would have contained insufficient photons to be useful.

The consequence of the increasing point spread function towards the edge of the chip for the *Chandra* satellite is that, despite having the shortest observation time, it is the third group (the 200 ks in which NGC 1275 is at the edge of the detector) that provides the cleanest data set. In the first case, the superb optics of *Chandra* work against it; almost all photons are concentrated onto a few pixels, and these central pixels are highly contaminated by pile-up. In this case a clean spectrum can only be obtained by extracting from the wings of the point spread function – which however reduces the photon count. In the last case however, being highly off-axis causes sufficient degradation of the optics that the arriving photons are scattered over many pixels, resulting in greatly reduced pile-up. The second grouping is intermediary in quality between these two.

The combination of pile-up and the intrinsic brightening of the AGN implies that, despite the shorter observation times, each of the stacked 200 ks and 300 ks ACIS-I edge and midway observations from 2009 has more counts than the 1Ms of ACIS-S observations taken in 2002 and 2004. Furthermore, when we consider cleaned spectra that exclude regions of high pile-up, the total counts in the ACIS-I edge spectra are more than in the ACIS-I midway and ACIS-S observations put together. These factors make the ACIS-I edge observations the optimal for analysis purposes, even though they involve the shortest observational time. Overall, the bounds derived from the *Chandra* observations of NGC 1275 are weaker than they would be if no artifacts in the data were restricting goodness of fit. This makes the bounds more conservative than they would be with a clean data set.

In terms of the satellites' capabilities and their effect on the constraints derived on $g_{a\gamma\gamma}$, the most important features are the angular resolution, effective area and energy resolution. The angular resolution must be good enough to resolve point sources sufficiently well against the cluster background, while a large effective area allows enough data to be collected to distinguish ALP oscillations from noise. Lastly, a good energy resolution is vital to ensure oscillations are not smeared out and lost, especially at lower energies.

5.3 Comparison with other studies

5.3.1 Chandra and XMM-Newton

The bounds from NGC 1275 are about a factor of 5 improvement over those derived from Hydra A in [112] of $g_{a\gamma\gamma} \lesssim 8.3 \times 10^{-12}\,\mathrm{GeV}^{-1}$. The improvement is consistent with the fact that the NGC 1275 data set contains far more counts than the Hydra A one. The ACIS-I and ACIS-S observations with Chandra generate over half a million X-ray counts from the central AGN. In contrast, the AGN in the centre of Hydra A is nowhere near as bright, and has relatively poor contrast with the intracluster emission. Despite the 238 ks of exposure time available, only about 2000 counts remain after background subtraction, compared to the 153 000 counts for the 'cleaned' ACIS-I edge observations. Therefore NGC 1275 represents a much improved source to search for ALP-photon conversion.

Our results from NGC 1275 are a factor of 3 improvement over the previous best limits, derived from gamma-ray observations from SN 1987A, in the mass range $m_a \lesssim 10^{-12}\,\mathrm{eV}$. For ALP masses greater than this, the SN 1987A bounds will remain superior, as our bound increases rapidly with increasing mass (I discuss the effect of increasing mass in more detail below). Once again the improvement can be explained by the large size of the data set at our disposal, in addition to the greater extent of the magnetic field in the Perseus Cluster than the Milky Way allowing stronger interconversion with ALPs. We have demonstrated that, even for our most conservative assumptions about the central magnetic field strength ($10\,\mu$G) and minimum coherence length (0.7 kpc) of the magnetic field in Perseus, our bound of $g_{a\gamma\gamma} \lesssim 4.0 \times 10^{-12}\,\mathrm{GeV}^{-1}$ is still an improvement over SN 1987A. No new data can be taken of SN 1987A, whereas (as we shall discuss below) new observations of NGC 1275 could improve our bound.

The fitted spectra from XMM-Newton show a good fit. However, there are both fewer counts and a worse contrast against the cluster background than for the *Chandra* data. From the perspective of ALP constraints, we can say that modulations of the spectrum are not allowed beyond the 20% level – a weaker constraint than for the *Chandra* data. This is unsurprising due to the poorer angular resolution of XMM-Newton leading to larger issues with pileup.

5.3.2 Athena

The estimated order of magnitude improvement in the bound from *Athena* is due to a combination of larger effective area, greatly improved energy resolution and reduced pileup contamination. These combined effects mean *Athena* has far more potential to detect ALP-induced oscillations than the best current satellites. It should be noted that this estimate is based on the calibration of the satellite being the same as is predicted from the `SIXTE` software, as well as there being no unexpected physics revealed in the data set.

These bounds are substantially better than any current experimental or astrophysical bound, and also go beyond the capabilities of the IAXO helioscope for ultralight ALP masses [169]. The proposed DM haloscope ABRACADABRA has the potential to explore $g_{a\gamma\gamma}$ down to $10^{-17}\,\text{GeV}^{-1}$ for $m_a \in [10^{-14}, 10^{-6}]\,\text{eV}$ [170], if ALPs constitute the Dark Matter. The existence of ALP-induced oscillations in galaxy clusters is independent of this. Proposed CMB experiments such as PIXIE [171] and PRISM [172] could produce a constraint $g_{a\gamma\gamma}B < 10^{-16}\,\text{GeV}^{-1}\,\text{nG}$ which might be competitive with bounds from galaxy clusters if the cosmic magnetic field is close enough to saturation $\sim$ nG [173]. Black hole superradiance also offers tentative constraints on ALPs in the mass range $m_a \in [10^{-14}, 10^{-10}]\,\text{eV}$, depending on measurements of black hole spin [111].

5.4 Constraints on ALPs as Dark Matter

An important question to ask is whether the parameter space excluded by our studies contains any viable ALP DM candidates. The short answer is no. Plugging the values for f_a that have been excluded ($\sim 10^{12}\,\text{GeV}$ for the NGC1275 *Chandra* study, and $\sim 10^{13}\,\text{GeV}$ for the NGC1275 *Athena* study) and the small mass limit $m_a \lesssim 10^{-12}\,\text{eV}$ into Equation 1.45, we find that the ALP density would be many orders of magnitude below the measured energy density for Dark Matter. We therefore conclude that such particles could at most contribute a tiny fraction to the DM density of the Universe.

5.5 Potential for improvement

There are a number of ways the analysis shown in this thesis can be extended and optimised. Firstly, calculating the ALP-photon conversion probability for non-zero masses allows us to extend the bounds above $m_a \lesssim 10^{-12}\,\text{eV}$. This work was conducted in [174], which extend the bounds in this thesis up to $m_a \lesssim 6 \times 10^{-12}\,\text{eV}$.

We have discussed how the data used in our analysis was not optimised for observations of AGN spectra. Satellite missions utilising shorter frame times, or gratings, could reduce pileup contamination and increase energy resolution, improving our determination of the AGN spectrum. This could potentially improve the constraints on $g_{a\gamma\gamma}$. In addition, given the statistical nature of these results, it is important to corroborate these findings with as many different galaxy clusters with magnetic fields as possible. Since the work shown in this thesis was published, a study of M87 was done in [136], finding a similar bound to our studies.

There could be potential to improve our constraints on $g_{a\gamma\gamma}$ using different statistical techniques. As the 3D structure of the magnetic field in a galaxy cluster is not known, neither are the precise energy and amplitude of ALP-induced oscillations in that cluster. Our bounds represent the point at which these oscillations on average become small compared to Poisson fluctuations. However, the discrete Fourier spectrum of ALP-oscillations is very different to noise, having greater amplitude and period at higher energies. Devising a statistical test that could discriminate these oscillations from noise could be exploited to push the bound down further.

Planned future polarimetry missions such as *XIPE*, *IXPE* and *PolStar* could also be used to constrain $g_{a\gamma\gamma}$. As the ALP-photon conversion only occurs for polarisations perpendicular to the magnetic field, previously unpolarised light would become polarised after passing through a large magnetic field. In addition, the degree of polarisation would be energy-dependent in the X-ray regime, similarly to the ALP-oscillations shown in this thesis, making it possible to differentiate signal from background. It is unclear whether this could produce a better constraint on $g_{a\gamma\gamma}$, but it merits detailed study.

Chapter 6

Conclusions and Outlook

In this thesis I have conducted a search for Axion-Like Particles, well-motivated extensions of the Standard Model, using X-ray astronomy. I have constrained their coupling to photons $g_{a\gamma\gamma}$ by looking for conversion in the large magnetic fields of galaxy clusters. By analysing bright point sources within or behind these clusters, we were able to constrain $g_{a\gamma\gamma}$ from the absence of quasi-sinusoidal modulations in their spectra induced by the energy-dependent ALP-photon conversion. The bounds derived from analysing the AGN in NGC1275 at the centre of the Perseus Cluster are about a factor of 5 improvement over those derived from a previous study of point sources in the Hydra A cluster. These bounds are a factor of 3 improvement over the previous best limits, derived from gamma-ray observations from SN 1987A, in the mass range $m_a \lesssim 10^{-12}\,\mathrm{eV}$. We have also demonstrated that the future satellite Athena could potentially improve this bound by a further order of magnitude, representing an improvement unmatched by any currently planned experiments for this mass range.

Bibliography

[1] Marcus Berg, Joseph P. Conlon, Francesca Day, Nicholas Jennings, Sven Krippendorf, Andrew J. Powell, and Markus Rummel. Constraints on Axion-Like Particles from X-ray Observations of NGC1275. *Astrophys. J.*, 847(2):101, 2017. doi: 10.3847/1538-4357/aa8b16.

[2] Joseph P. Conlon, Francesca Day, Nicholas Jennings, Sven Krippendorf, and Markus Rummel. Constraints on Axion-Like Particles from Non-Observation of Spectral Modulations for X-ray Point Sources. *JCAP*, 1707(07):005, 2017. doi: 10.1088/1475-7516/2017/07/005.

[3] Joseph P. Conlon, Francesca Day, Nicholas Jennings, Sven Krippendorf, and Francesco Muia. Projected bounds on ALPs from Athena. *Mon. Not. Roy. Astron. Soc.*, 473(4):4932–4936, 2018. doi: 10.1093/mnras/stx2652.

[4] Joseph P. Conlon, Francesca Day, Nicholas Jennings, Sven Krippendorf, and Markus Rummel. Consistency of Hitomi, XMM-Newton, and Chandra 3.5 keV data from Perseus. *Phys. Rev.*, D96(12):123009, 2017. doi: 10.1103/PhysRevD.96.123009.

[5] Steven Weinberg. The U(1) Problem. *Phys. Rev.*, D11:3583–3593, 1975. doi: 10.1103/PhysRevD.11.3583.

[6] Stephen L. Adler. Axial vector vertex in spinor electrodynamics. *Phys. Rev.*, 177:2426–2438, 1969. doi: 10.1103/PhysRev.177.2426. [,241(1969)].

[7] J. S. Bell and R. Jackiw. A PCAC puzzle: $\pi^0 \to \gamma\gamma$ in the σ model. *Nuovo Cim.*, A60:47–61, 1969. doi: 10.1007/BF02823296.

[8] Kazuo Fujikawa. Path Integral for Gauge Theories with Fermions. *Phys. Rev.*, D21:2848, 1980. doi: 10.1103/PhysRevD.21.2848,10.1103/PhysRevD.22.1499. [Erratum: Phys. Rev.D22,1499(1980)].

[9] Gerard 't Hooft. Symmetry Breaking Through Bell-Jackiw Anomalies. *Phys. Rev. Lett.*, 37:8–11, 1976. doi: 10.1103/PhysRevLett.37.8. [,226(1976)].

[10] Gerard 't Hooft. Computation of the Quantum Effects Due to a Four-Dimensional Pseudoparticle. *Phys. Rev.*, D14:3432–3450, 1976. doi: 10.1103/PhysRevD.18.2199.3,10.1103/PhysRevD.14.3432. [,70(1976)].

[11] R. Jackiw and C. Rebbi. Vacuum Periodicity in a Yang-Mills Quantum Theory. *Phys. Rev. Lett.*, 37:172–175, 1976. doi: 10.1103/PhysRevLett.37.172. [,353(1976)].

[12] C. A. Baker et al. An Improved experimental limit on the electric dipole moment of the neutron. *Phys. Rev. Lett.*, 97:131801, 2006. doi: 10.1103/PhysRevLett.97.131801.

[13] M. A. B. Beg and H. S. Tsao. Strong P, T Noninvariances in a Superweak Theory. *Phys. Rev. Lett.*, 41:278, 1978. doi: 10.1103/PhysRevLett.41.278.

[14] Howard Georgi. A Model of Soft CP Violation. *Hadronic J.*, 1:155, 1978.

[15] Rabindra N. Mohapatra and G. Senjanovic. Natural Suppression of Strong p and t Noninvariance. *Phys. Lett.*, 79B:283–286, 1978. doi: 10.1016/0370-2693(78)90243-5.

[16] Gino Segre and H. Arthur Weldon. Natural Suppression of Strong P and T Violations and Calculable Mixing Angles in SU(2) X U(1). *Phys. Rev. Lett.*, 42:1191, 1979. doi: 10.1103/PhysRevLett.42.1191.

[17] Stephen M. Barr and Paul Langacker. A Superweak Gauge Theory of CP Violation. *Phys. Rev. Lett.*, 42:1654, 1979. doi: 10.1103/PhysRevLett.42.1654.

[18] Ann E. Nelson. Naturally Weak CP Violation. *Phys. Lett.*, 136B:387–391, 1984. doi: 10.1016/0370-2693(84)92025-2.

[19] Stephen M. Barr. Solving the Strong CP Problem Without the Peccei-Quinn Symmetry. *Phys. Rev. Lett.*, 53:329, 1984. doi: 10.1103/PhysRevLett.53.329.

[20] Michael Dine and Patrick Draper. Challenges for the Nelson-Barr Mechanism. *JHEP*, 08:132, 2015. doi: 10.1007/JHEP08(2015)132.

[21] David B. Kaplan and Aneesh V. Manohar. Current Mass Ratios of the Light Quarks. *Phys. Rev. Lett.*, 56:2004, 1986. doi: 10.1103/PhysRevLett.56.2004.

[22] Sinya Aoki et al. Review of lattice results concerning low-energy particle physics. *Eur. Phys. J.*, C74:2890, 2014. doi: 10.1140/epjc/s10052-014-2890-7.

[23] R. D. Peccei and Helen R. Quinn. CP Conservation in the Presence of Instantons. *Phys. Rev. Lett.*, 38:1440–1443, 1977. doi: 10.1103/PhysRevLett.38.1440. [,328(1977)].

[24] R. D. Peccei and Helen R. Quinn. Constraints Imposed by CP Conservation in the Presence of Instantons. *Phys. Rev.*, D16:1791–1797, 1977. doi: 10.1103/PhysRevD.16.1791.

[25] Frank Wilczek. Problem of Strong p and t Invariance in the Presence of Instantons. *Phys. Rev. Lett.*, 40:279–282, 1978. doi: 10.1103/PhysRevLett.40.279.

[26] Steven Weinberg. A New Light Boson? *Phys. Rev. Lett.*, 40:223–226, 1978. doi: 10.1103/PhysRevLett.40.223.

[27] Cumrun Vafa and Edward Witten. Parity Conservation in QCD. *Phys. Rev. Lett.*, 53:535, 1984. doi: 10.1103/PhysRevLett.53.535.

[28] Y. Asano, E. Kikutani, S. Kurokawa, T. Miyachi, M. Miyajima, Y. Nagashima, T. Shinkawa, S. Sugimoto, and Y. Yoshimura. Search for a Rare Decay Mode K+ —> pi+ Neutrino anti-neutrino and Axion. *Phys. Lett.*, 107B:159, 1981. doi: 10.1016/0370-2693(81)91172-2. [,411(1981)].

[29] Ken Mimasu and Verónica Sanz. ALPs at Colliders. *JHEP*, 06:173, 2015. doi: 10.1007/JHEP06(2015)173.

[30] Mark B. Wise, Howard Georgi, and Sheldon L. Glashow. SU(5) and the Invisible Axion. *Phys. Rev. Lett.*, 47:402, 1981. doi: 10.1103/PhysRevLett.47.402.

[31] Jihn E. Kim. Weak Interaction Singlet and Strong CP Invariance. *Phys. Rev. Lett.*, 43:103, 1979. doi: 10.1103/PhysRevLett.43.103.

[32] Mikhail A. Shifman, A. I. Vainshtein, and Valentin I. Zakharov. Can Confinement Ensure Natural CP Invariance of Strong Interactions? *Nucl. Phys.*, B166: 493–506, 1980. doi: 10.1016/0550-3213(80)90209-6.

[33] Michael Dine, Willy Fischler, and Mark Srednicki. A Simple Solution to the Strong CP Problem with a Harmless Axion. *Phys. Lett.*, 104B:199–202, 1981. doi: 10.1016/0370-2693(81)90590-6.

[34] A. R. Zhitnitsky. On Possible Suppression of the Axion Hadron Interactions. (In Russian). *Sov. J. Nucl. Phys.*, 31:260, 1980. [Yad. Fiz.31,497(1980)].

[35] Peter Svrcek and Edward Witten. Axions In String Theory. *JHEP*, 06:051, 2006. doi: 10.1088/1126-6708/2006/06/051.

[36] Andreas Ringwald. Searching for axions and ALPs from string theory. *J. Phys. Conf. Ser.*, 485:012013, 2014. doi: 10.1088/1742-6596/485/1/012013.

[37] Edward Witten. Some Properties of O(32) Superstrings. *Phys. Lett.*, 149B: 351–356, 1984. doi: 10.1016/0370-2693(84)90422-2.

[38] Kiwoon Choi and Jihn E. Kim. Compactification and Axions in E(8) x E(8)-prime Superstring Models. *Phys. Lett.*, 165B:71–75, 1985. doi: 10.1016/0370-2693(85)90693-8.

[39] Michele Cicoli, Mark Goodsell, and Andreas Ringwald. The type IIB string axiverse and its low-energy phenomenology. *JHEP*, 10:146, 2012. doi: 10.1007/JHEP10(2012)146.

[40] Bobby Samir Acharya, Konstantin Bobkov, and Piyush Kumar. An M Theory Solution to the Strong CP Problem and Constraints on the Axiverse. *JHEP*, 11:105, 2010. doi: 10.1007/JHEP11(2010)105.

[41] Tom Banks, Michael Dine, Patrick J. Fox, and Elie Gorbatov. On the possibility of large axion decay constants. *JCAP*, 0306:001, 2003. doi: 10.1088/1475-7516/2003/06/001.

[42] Asimina Arvanitaki, Savas Dimopoulos, Sergei Dubovsky, Nemanja Kaloper, and John March-Russell. String Axiverse. *Phys. Rev.*, D81:123530, 2010. doi: 10.1103/PhysRevD.81.123530.

[43] Kiwoon Choi and Jihn E. Kim. Harmful Axions in Superstring Models. *Phys. Lett.*, 154B:393, 1985. doi: 10.1016/0370-2693(85)90416-2. [Erratum: Phys. Lett.156B,452(1985)].

[44] Stephen M. Barr. Harmless Axions in Superstring Theories. *Phys. Lett.*, 158B: 397–400, 1985. doi: 10.1016/0370-2693(85)90440-X.

[45] Tom Banks and Michael Dine. The Cosmology of string theoretic axions. *Nucl. Phys.*, B505:445–460, 1997. doi: 10.1016/S0550-3213(97)00413-6.

[46] Joseph P. Conlon. The QCD axion and moduli stabilisation. *JHEP*, 05:078, 2006. doi: 10.1088/1126-6708/2006/05/078.

[47] Michele Cicoli, Joseph P. Conlon, and Fernando Quevedo. Dark radiation in LARGE volume models. *Phys. Rev.*, D87(4):043520, 2013. doi: 10.1103/PhysRevD.87.043520.

[48] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016. doi: 10.1088/1674-1137/40/10/100001.

[49] David H. Weinberg, James S. Bullock, Fabio Governato, Rachel Kuzio de Naray, and Annika H. G. Peter. Cold dark matter: controversies on small scales. *Proc. Nat. Acad. Sci.*, 112:12249–12255, 2015. doi: 10.1073/pnas.1308716112.

[50] B. Moore, S. Ghigna, F. Governato, G. Lake, Thomas R. Quinn, J. Stadel, and P. Tozzi. Dark matter substructure within galactic halos. *Astrophys. J.*, 524: L19–L22, 1999. doi: 10.1086/312287.

[51] Anatoly A. Klypin, Andrey V. Kravtsov, Octavio Valenzuela, and Francisco Prada. Where are the missing Galactic satellites? *Astrophys. J.*, 522:82–92, 1999. doi: 10.1086/307643.

[52] Michael Boylan-Kolchin, James S. Bullock, and Manoj Kaplinghat. Too big to fail? The puzzling darkness of massive Milky Way subhaloes. *Mon. Not. Roy. Astron. Soc.*, 415:L40, 2011. doi: 10.1111/j.1745-3933.2011.01074.x.

[53] Andrew Pontzen and Fabio Governato. Cold dark matter heats up. *Nature*, 506:171–178, 2014. doi: 10.1038/nature12953.

[54] David N. Spergel and Paul J. Steinhardt. Observational evidence for self-interacting cold dark matter. *Phys. Rev. Lett.*, 84:3760–3763, 2000. doi: 10.1103/PhysRevLett.84.3760.

[55] Paul Bode, Jeremiah P. Ostriker, and Neil Turok. Halo formation in warm dark matter models. *Astrophys. J.*, 556:93–107, 2001. doi: 10.1086/321541.

[56] David J. E. Marsh and Pedro G. Ferreira. Ultra-Light Scalar Fields and the Growth of Structure in the Universe. *Phys. Rev.*, D82:103528, 2010. doi: 10.1103/PhysRevD.82.103528.

[57] Hsi-Yu Schive, Tzihong Chiueh, and Tom Broadhurst. Cosmic Structure as the Quantum Interference of a Coherent Dark Wave. *Nature Phys.*, 10:496–499, 2014. doi: 10.1038/nphys2996.

[58] Renée Hlozek, David J. E. Marsh, and Daniel Grin. Using the Full Power of the Cosmic Microwave Background to Probe Axion Dark Matter. *Mon. Not. Roy. Astron. Soc.*, 476(3):3063–3085, 2018. doi: 10.1093/mnras/sty271.

[59] Eric Armengaud, Nathalie Palanque-Delabrouille, Christophe Yèche, David J. E. Marsh, and Julien Baur. Constraining the mass of light bosonic dark matter using SDSS Lyman-α forest. *Mon. Not. Roy. Astron. Soc.*, 471(4):4606–4614, 2017. doi: 10.1093/mnras/stx1870.

[60] Edward W. Kolb and Igor I. Tkachev. Axion miniclusters and Bose stars. *Phys. Rev. Lett.*, 71:3051–3054, 1993. doi: 10.1103/PhysRevLett.71.3051.

[61] Kathryn M. Zurek, Craig J. Hogan, and Thomas R. Quinn. Astrophysical Effects of Scalar Dark Matter Miniclusters. *Phys. Rev.*, D75:043511, 2007. doi: 10.1103/PhysRevD.75.043511.

[62] Malcolm Fairbairn, David J. E. Marsh, and Jérémie Quevillon. Searching for the QCD Axion with Gravitational Microlensing. *Phys. Rev. Lett.*, 119(2):021101, 2017. doi: 10.1103/PhysRevLett.119.021101.

[63] Malcolm Fairbairn, David J. E. Marsh, Jérémie Quevillon, and Simon Rozier. Structure formation and microlensing with axion miniclusters. *Phys. Rev.*, D97 (8):083502, 2018. doi: 10.1103/PhysRevD.97.083502.

[64] Pierre Sikivie. The caustic ring singularity. *Phys. Rev.*, D60:063501, 1999. doi: 10.1103/PhysRevD.60.063501.

[65] Aravind Natarajan and Pierre Sikivie. The inner caustics of cold dark matter halos. *Phys. Rev.*, D73:023510, 2006. doi: 10.1103/PhysRevD.73.023510.

[66] M. Ted Ressell. Limits to the radiative decay of the axion. *Phys. Rev.*, D44: 3001–3020, 1991. doi: 10.1103/PhysRevD.44.3001.

[67] Eduard Masso and Ramon Toldra. New constraints on a light spinless particle coupled to photons. *Phys. Rev.*, D55:7967–7969, 1997. doi: 10.1103/PhysRevD.55.7967.

[68] P. Sikivie. Experimental Tests of the Invisible Axion. *Phys. Rev. Lett.*, 51:1415–1417, 1983. doi: 10.1103/PhysRevLett.51.1415,10.1103/PhysRevLett.52.695.2. [,321(1983)].

[69] W. Heisenberg and H. Euler. Folgerungen aus der Diracschen Theorie des Positrons. *Z. Phys.*, 98:714–732, 1936. doi: 10.1007/BF01343663.

[70] P. Sikivie. Detection rates for "invisible"-axion searches. *Phys. Rev. D*, 32: 2988–2991, Dec 1985. doi: 10.1103/PhysRevD.32.2988. URL `http://link.aps.org/doi/10.1103/PhysRevD.32.2988`.

[71] Georg Raffelt and Leo Stodolsky. Mixing of the Photon with Low Mass Particles. *Phys. Rev.*, D37:1237, 1988. doi: 10.1103/PhysRevD.37.1237.

[72] Jack W. Brockway, Eric D. Carlson, and Georg G. Raffelt. SN1987A gamma-ray limits on the conversion of pseudoscalars. *Phys. Lett.*, B383:439–443, 1996. doi: 10.1016/0370-2693(96)00778-2.

[73] J. A. Grifols, E. Masso, and R. Toldra. Gamma-rays from SN1987A due to pseudoscalar conversion. *Phys. Rev. Lett.*, 77:2372–2375, 1996. doi: 10.1103/PhysRevLett.77.2372.

[74] Mattias Christensson and Malcolm Fairbairn. Photon-axion mixing in an inhomogeneous universe. *Phys. Lett.*, B565:10–18, 2003. doi: 10.1016/S0370-2693(03)00641-5.

[75] Malcolm Fairbairn, Timur Rashba, and Sergey V. Troitsky. Transparency of the Sun to gamma rays due to axionlike particles. *Phys. Rev. Lett.*, 98:201801, 2007. doi: 10.1103/PhysRevLett.98.201801.

[76] Malcolm Fairbairn, Timur Rashba, and Sergey V. Troitsky. Photon-axion mixing and ultra-high-energy cosmic rays from BL Lac type objects - Shining light through the Universe. *Phys. Rev.*, D84:125019, 2011. doi: 10.1103/PhysRevD.84.125019.

[77] Clare Burrage, Anne-Christine Davis, and Douglas J. Shaw. Active Galactic Nuclei Shed Light on Axion-like-Particles. *Phys. Rev. Lett.*, 102:201101, 2009. doi: 10.1103/PhysRevLett.102.201101.

[78] Malcolm Fairbairn, Timur Rashba, and Sergey V. Troitsky. Photon-axion mixing and ultra-high-energy cosmic rays from BL Lac type objects - Shining light through the Universe. *Phys. Rev.*, D84:125019, 2011. doi: 10.1103/PhysRevD.84.125019.

[79] Alessandro Mirizzi and Daniele Montanino. Stochastic conversions of TeV photons into axion-like particles in extragalactic magnetic fields. *JCAP*, 0912:004, 2009. doi: 10.1088/1475-7516/2009/12/004.

[80] Fabrizio Tavecchio, Marco Roncadelli, Giorgio Galanti, and Giacomo Bonnoli. Evidence for an axion-like particle from PKS 1222+216? *Phys. Rev.*, D86:085036, 2012. doi: 10.1103/PhysRevD.86.085036.

[81] A. Payez, J. R. Cudell, and D. Hutsemekers. New polarimetric constraints on axion-like particles. *JCAP*, 1207:041, 2012. doi: 10.1088/1475-7516/2012/07/041.

[82] Dieter Horns, Luca Maccione, Manuel Meyer, Alessandro Mirizzi, Daniele Montanino, and Marco Roncadelli. Hardening of TeV gamma spectrum of AGNs in galaxy clusters by conversions of photons into axion-like particles. *Phys. Rev.*, D86:075024, 2012. doi: 10.1103/PhysRevD.86.075024.

[83] Manuel Meyer, Dieter Horns, and Martin Raue. First lower limits on the photon-axion-like particle coupling from very high energy gamma-ray observations. *Phys. Rev.*, D87(3):035027, 2013. doi: 10.1103/PhysRevD.87.035027.

[84] Denis Wouters and Pierre Brun. Constraints on Axion-like Particles from X-Ray Observations of the Hydra Galaxy Cluster. *Astrophys. J.*, 772:44, 2013. doi: 10.1088/0004-637X/772/1/44.

[85] Malcolm Fairbairn. Axionic Dark Radiation and the Milky Way's Magnetic Field. *Phys. Rev.*, D89(6):064020, 2014. doi: 10.1103/PhysRevD.89.064020.

[86] Joseph P. Conlon and M. C. David Marsh. Excess Astrophysical Photons from a 0.1-1 keV Cosmic Axion Background. *Phys. Rev. Lett.*, 111(15):151301, 2013. doi: 10.1103/PhysRevLett.111.151301.

[87] Stephen Angus, Joseph P. Conlon, M. C. David Marsh, Andrew J. Powell, and Lukas T. Witkowski. Soft X-ray Excess in the Coma Cluster from a Cosmic Axion Background. *JCAP*, 1409(09):026, 2014. doi: 10.1088/1475-7516/2014/09/026.

[88] Alexandre Payez, Carmelo Evoli, Tobias Fischer, Maurizio Giannotti, Alessandro Mirizzi, and Andreas Ringwald. Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles. *JCAP*, 1502(02):006, 2015. doi: 10.1088/1475-7516/2015/02/006.

[89] Andrew J. Powell. A Cosmic ALP Background and the Cluster Soft X-ray Excess in A665, A2199 and A2255. *JCAP*, 1509(09):017, 2015. doi: 10.1088/1475-7516/2015/09/017.

[90] Alexandra Dobrynina, Alexander Kartavtsev, and Georg Raffelt. Photon-photon dispersion of TeV gamma rays and its role for photon-ALP conversion. *Phys. Rev.*, D91:083003, 2015. doi: 10.1103/PhysRevD.91.083003.

[91] Martin Schlederer and Günter Sigl. Constraining ALP-photon coupling using galaxy clusters. *JCAP*, 1601(01):038, 2016. doi: 10.1088/1475-7516/2016/01/038.

[92] Joseph P. Conlon, M. C. David Marsh, and Andrew J. Powell. Galaxy cluster thermal x-ray spectra constrain axionlike particles. *Phys. Rev.*, D93(12):123526, 2016. doi: 10.1103/PhysRevD.93.123526.

[93] Yevgeny V. Stadnik and Victor V. Flambaum. New generation low-energy probes for ultralight axion and scalar dark matter. *Mod. Phys. Lett.*, A32: 1740004, 2017. doi: 10.1142/S0217732317400041.

[94] Pierre Pugnat et al. First results from the OSQAR photon regeneration experiment: No light shining through a wall. *Phys. Rev.*, D78:092003, 2008. doi: 10.1103/PhysRevD.78.092003.

[95] R. Ballou et al. New exclusion limits on scalar and pseudoscalar axionlike particles from light shining through a wall. *Phys. Rev.*, D92(9):092002, 2015. doi: 10.1103/PhysRevD.92.092002.

[96] Adrian Ayala, Inma Domínguez, Maurizio Giannotti, Alessandro Mirizzi, and Oscar Straniero. Revisiting the bound on axion-photon coupling from Globular Clusters. *Phys. Rev. Lett.*, 113(19):191302, 2014. doi: 10.1103/PhysRevLett.113.191302.

[97] Nicolás Viaux, Márcio Catelan, Peter B. Stetson, Georg Raffelt, Javier Redondo, Aldo A. R. Valcarce, and Achim Weiss. Neutrino and axion bounds

from the globular cluster M5 (NGC 5904). *Phys. Rev. Lett.*, 111:231301, 2013. doi: 10.1103/PhysRevLett.111.231301.

[98] Georg G. Raffelt. Axion Constraints From White Dwarf Cooling Times. *Phys. Lett.*, 166B:402–406, 1986. doi: 10.1016/0370-2693(86)91588-1.

[99] Alexandre Payez, Carmelo Evoli, Tobias Fischer, Maurizio Giannotti, Alessandro Mirizzi, and Andreas Ringwald. Revisiting the SN1987A gamma-ray limit on ultralight axion-like particles. *JCAP*, 1502(02):006, 2015. doi: 10.1088/1475-7516/2015/02/006.

[100] V. Anastassopoulos et al. New CAST Limit on the Axion-Photon Interaction. *Nature Phys.*, 13:584–590, 2017. doi: 10.1038/nphys4109.

[101] E. Arik et al. Probing eV-scale axions with CAST. *JCAP*, 0902:008, 2009. doi: 10.1088/1475-7516/2009/02/008.

[102] M. Arik et al. Search for Solar Axions by the CERN Axion Solar Telescope with ^{3}He Buffer Gas: Closing the Hot Dark Matter Gap. *Phys. Rev. Lett.*, 112 (9):091302, 2014. doi: 10.1103/PhysRevLett.112.091302.

[103] Stephen J. Asztalos et al. An Improved RF cavity search for halo axions. *Phys. Rev.*, D69:011101, 2004. doi: 10.1103/PhysRevD.69.011101.

[104] S. J. Asztalos et al. A SQUID-based microwave cavity search for dark-matter axions. *Phys. Rev. Lett.*, 104:041301, 2010. doi: 10.1103/PhysRevLett.104.041301.

[105] C. Hagmann, P. Sikivie, N. S. Sullivan, and D. B. Tanner. Results from a search for cosmic axions. *Phys. Rev.*, D42:1297–1300, 1990. doi: 10.1103/PhysRevD.42.1297.

[106] B. M. Brubaker et al. First results from a microwave cavity axion search at 24 μeV. *Phys. Rev. Lett.*, 118(6):061302, 2017. doi: 10.1103/PhysRevLett.118.061302.

[107] L. Zhong et al. Results from phase 1 of the HAYSTAC microwave cavity axion experiment. *Phys. Rev.*, D97(9):092001, 2018. doi: 10.1103/PhysRevD.97.092001.

[108] Ben T. McAllister, Graeme Flower, Eugene N. Ivanov, Maxim Goryachev, Jeremy Bourhill, and Michael E. Tobar. The ORGAN Experiment: An axion haloscope above 15 GHz. *Phys. Dark Univ.*, 18:67–72, 2017. doi: 10.1016/j.dark.2017.09.010.

[109] Asimina Arvanitaki and Sergei Dubovsky. Exploring the String Axiverse with Precision Black Hole Physics. *Phys. Rev.*, D83:044026, 2011. doi: 10.1103/PhysRevD.83.044026.

[110] Asimina Arvanitaki, Masha Baryakhtar, and Xinlu Huang. Discovering the QCD Axion with Black Holes and Gravitational Waves. *Phys. Rev.*, D91(8):084011, 2015. doi: 10.1103/PhysRevD.91.084011.

[111] Asimina Arvanitaki, Masha Baryakhtar, Savas Dimopoulos, Sergei Dubovsky, and Robert Lasenby. Black Hole Mergers and the QCD Axion at Advanced LIGO. *Phys. Rev.*, D95(4):043001, 2017. doi: 10.1103/PhysRevD.95.043001.

[112] Denis Wouters and Pierre Brun. Constraints on Axion-like Particles from X-Ray Observations of the Hydra Galaxy Cluster. *Astrophys. J.*, 772:44, 2013. doi: 10.1088/0004-637X/772/1/44.

[113] M. Ajello et al. Search for Spectral Irregularities due to Photon–Axionlike-Particle Oscillations with the Fermi Large Area Telescope. *Phys. Rev. Lett.*, 116(16):161101, 2016. doi: 10.1103/PhysRevLett.116.161101.

[114] A. Abramowski et al. Constraints on axionlike particles with H.E.S.S. from the irregularity of the PKS 2155-304 energy spectrum. *Phys. Rev.*, D88(10):102003, 2013. doi: 10.1103/PhysRevD.88.102003.

[115] Alfonso Cavaliere and Andrea Lapi. The Astrophysics of the Intracluster Plasma. *Phys. Rept.*, 533:69–94, 2013. doi: 10.1016/j.physrep.2013.08.001.

[116] Federica Govoni and Luigina Feretti. Magnetic field in clusters of galaxies. *Int. J. Mod. Phys.*, D13:1549–1594, 2004. doi: 10.1142/S0218271804005080.

[117] G. R. Blumenthal and R. J. Gould. Bremsstrahlung, synchrotron radiation, and compton scattering of high-energy electrons traversing dilute gases. *Rev. Mod. Phys.*, 42:237–270, 1970. doi: 10.1103/RevModPhys.42.237.

[118] A. Bonafede, L. Feretti, M. Murgia, F. Govoni, G. Giovannini, D. Dallacasa, K. Dolag, and G. B. Taylor. The Coma cluster magnetic field from Faraday rotation measures. *Astron. Astrophys.*, 513:A30, 2010. doi: 10.1051/0004-6361/200913696.

[119] E. Churazov, W. Forman, C. Jones, and H. Bohringer. XMM-newton observations of the perseus cluster I: the temperature and surface brightness structure. *Astrophys. J.*, 590:225–237, 2003. doi: 10.1086/374923.

[120] S. Yamazaki, Y. Fukazawa, M. Sasada, R. Itoh, S. Nishino, H. Takahashi, K. Takaki, K. S. Kawabata, M. Yoshida, and M. Uemura. X-Ray and Optical Monitoring of a Gamma-Ray-Emitting Radio Galaxy, NGC 1275. *Publications of the Astronomical Society of Japan*, 65, April 2013. doi: 10.1093/pasj/65.2.30.

[121] Barbara Balmaverde, Alessandro Capetti, and Paola Grandi. The Chandra view of the 3c/fri sample of low luminosity radio-galaxies. *Astron. Astrophys.*, 451:35, 2006. doi: 10.1051/0004-6361:20053799.

[122] A. C. Fabian, S. A. Walker, C. Pinto, H. R. Russell, and A. C. Edge. Effects of the variability of the nucleus of NGC 1275 on X-ray observations of the surrounding intracluster medium. *Mon. Not. Roy. Astron. Soc.*, 451(3):3061–3067, 2015. doi: 10.1093/mnras/stv1134.

[123] G. B. Taylor, N. E. Gugliucci, A. C. Fabian, J. S. Sanders, Gianfranco Gentile, and S. W. Allen. Magnetic fields in the center of the perseus cluster. *Mon. Not. Roy. Astron. Soc.*, 368:1500–1506, 2006. doi: 10.1111/j.1365-2966.2006.10244.x.

[124] Stephen Angus, Joseph P. Conlon, M. C. David Marsh, Andrew J. Powell, and Lukas T. Witkowski. Soft X-ray Excess in the Coma Cluster from a Cosmic Axion Background. *JCAP*, 1409(09):026, 2014. doi: 10.1088/1475-7516/2014/09/026.

[125] F. Govoni et al. Sardinia Radio Telescope observations of Abell 194 - the intracluster magnetic field power spectrum. *Astron. Astrophys.*, 603:A122, 2017. doi: 10.1051/0004-6361/201630349.

[126] P. Kuchar and T. A. Enßlin. Magnetic power spectra from Faraday rotation maps. REALMAF and its use on Hydra A. *Astron.Astrophys.*, 529:A13, May 2011. doi: 10.1051/0004-6361/200913918.

[127] A. Bonafede, L. Feretti, M. Murgia, F. Govoni, G. Giovannini, D. Dallacasa, K. Dolag, and G. B. Taylor. The Coma cluster magnetic field from Faraday rotation measures. *Astron. Astrophys.*, 513:A30, 2010. doi: 10.1051/0004-6361/200913696.

[128] V. Vacca, M. Murgia, F. Govoni, L. Feretti, G. Giovannini, R. A. Perley, and G. B. Taylor. The intracluster magnetic field power spectrum in A2199. *Astron. Astrophys.*, 540:A38, 2012. doi: 10.1051/0004-6361/201116622.

[129] Denis Wouters and Pierre Brun. Irregularity in gamma ray source spectra as a signature of axionlike particles. *Phys. Rev.*, D86:043005, 2012. doi: 10.1103/PhysRevD.86.043005.

[130] A. Fruscione, J. C. McDowell, G. E. Allen, N. S. Brickhouse, D. J. Burke, J. E. Davis, N. Durham, M. Elvis, E. C. Galle, D. E. Harris, D. P. Huenemoerder, J. C. Houck, B. Ishibashi, M. Karovska, F. Nicastro, M. S. Noble, M. A. Nowak, F. A. Primini, A. Siemiginowska, R. K. Smith, and M. Wise. CIAO: Chandra's data analysis system. In *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, volume 6270 of *Proceedings of the International Society for Optical Engineering*, page 62701V, June 2006. doi: 10.1117/12.671760.

[131] P. Freeman, S. Doe, and A. Siemiginowska. Sherpa: a mission-independent data analysis application. In J.-L. Starck and F. D. Murtagh, editors, *Astronomical Data Analysis*, volume 4477 of *Proceedings of the International Society for Optical Engineering*, pages 76–87, November 2001. doi: 10.1117/12.447161.

[132] J. E. Davis. Event Pileup in Charge-coupled Devices. *Astrophys. J*, 562:575–582, November 2001. doi: 10.1086/323488.

[133] D. E. Harris, C. C. Cheung, J. A. Biretta, W. B. Sparks, W. Junor, E. S. Perlman, and A. S. Wilson. The outburst of hst-1 in the m87 jet. *The Astrophysical Journal*, 640(1):211, 2006. URL `http://stacks.iop.org/0004-637X/640/i=1/a=211`.

[134] C. Carter, M. Karovska, D. Jerius, K. Glotfelty, and S. Beikman. ChaRT: The Chandra Ray Tracer. In H. E. Payne, R. I. Jedrzejewski, and R. N. Hook, editors, *Astronomical Data Analysis Software and Systems XII*, volume 295 of *Astronomical Society of the Pacific Conference Series*, page 477, 2003.

[135] K. A. Arnaud. XSPEC: The First Ten Years. In G. H. Jacoby and J. Barnes, editors, *Astronomical Data Analysis Software and Systems V*, volume 101 of *Astronomical Society of the Pacific Conference Series*, page 17, 1996.

[136] M. C. David Marsh, Helen R. Russell, Andrew C. Fabian, Brian P. McNamara, Paul Nulsen, and Christopher S. Reynolds. A New Bound on Axion-Like Particles. *JCAP*, 1712(12):036, 2017. doi: 10.1088/1475-7516/2017/12/036.

[137] Kirpal Nandra et al. The Hot and Energetic Universe: A White Paper presenting the science theme motivating the Athena+ mission. 2013.

[138] Luciano Gottardi et al. Development of TES-based detectors array for the X-ray Integral Field Unit (X-IFU) on the future x-ray observatory ATHENA. *Proc. SPIE Int. Soc. Opt. Eng.*, 9144:91442M, 2014. doi: 10.1117/12.2056920.

[139] Didier Barret et al. The Athena X-ray Integral Field Unit (X-IFU). *Proc. SPIE Int. Soc. Opt. Eng.*, 9905:99052F, 2016. doi: 10.1117/12.2232432.

[140] Marc Wenger et al. The simbad astronomical database. *Astron. Astrophys. Suppl. Ser.*, 143:9, 2000. doi: 10.1051/aas:2000332.

[141] Esra Bulbul, Maxim Markevitch, Adam Foster, Randall K. Smith, Michael Loewenstein, and Scott W. Randall. Detection of An Unidentified Emission Line in the Stacked X-ray spectrum of Galaxy Clusters. *Astrophys. J.*, 789:13, 2014. doi: 10.1088/0004-637X/789/1/13.

[142] Alexey Boyarsky, Oleg Ruchayskiy, Dmytro Iakubovskyi, and Jeroen Franse. Unidentified Line in X-Ray Spectra of the Andromeda Galaxy and Perseus Galaxy Cluster. *Phys. Rev. Lett.*, 113:251301, 2014. doi: 10.1103/PhysRevLett.113.251301.

[143] J. E. Davis, M. W. Bautz, D. Dewey, R. K. Heilmann, J. C. Houck, D. P. Huenemoerder, H. L. Marshall, M. A. Nowak, M. L. Schattenburg, N. S. Schulz, and R. K. Smith. Raytracing with MARX: x-ray observatory design, calibration, and support. In *Space Telescopes and Instrumentation 2012: Ultraviolet to Gamma Ray*, volume 8443 of *Proc. SPIE*, page 84431A, September 2012. doi: 10.1117/12.926937.

[144] F. A. Aharonian et al. *Hitomi* constraints on the 3.5 keV line in the Perseus galaxy cluster. *Astrophys. J.*, 837(1):L15, 2017. doi: 10.3847/2041-8213/aa61fa.

[145] R. Morrison and D. McCammon. Interstellar photoelectric absorption cross-sections, 0.03-10 keV. *Astrophys. J.*, 270:119, 1983. doi: 10.1086/161102.

[146] A. Bonafede, L. Feretti, M. Murgia, F. Govoni, G. Giovannini, D. Dallacasa, K. Dolag, and G. B. Taylor. The Coma cluster magnetic field from Faraday rotation measures. *Astron. Astrophys.*, 513:A30, 2010. doi: 10.1051/0004-6361/200913696.

[147] Torsten A. Ensslin, Peter L. Biermann, Ulrich Klein, and Sven Kohle. Cluster radio relics as a tracer of shock waves of the large - scale structure formation. *Astron. Astrophys.*, 332:395, 1998.

[148] H. V. Klapdor-Kleingrothaus and I. V. Krivosheina, editors. *Dark Matter in Astrophysics and Particle Physics - Proceedings of the 7th International Heidelberg Conference on Dark 2009*, 2010. doi: 10.1142/7610.

[149] J. P. Ge and F. N. Owen. Faraday rotation in cooling flow clusters of galaxies. I - Radio and X-ray observations of Abell 1795. *ApJ*, 105:778–787, March 1993. doi: 10.1086/116471.

[150] S. Ettori. β-model and cooling flows in X-ray clusters of galaxies. *MNRAS*, 318:1041–1046, November 2000. doi: 10.1046/j.1365-8711.2000.03664.x.

[151] T. J. Dennis and B. D. G. Chandran. Turbulent Heating of Galaxy-Cluster Plasmas. *ApJ*, 622:205–216, March 2005. doi: 10.1086/427424.

[152] E. L. Blanton, C. L. Sarazin, and B. R. McNamara. Chandra Observation of the Cooling Flow Cluster Abell 2052. *ApJ*, 585:227–243, March 2003. doi: 10.1086/345984.

[153] R. E. G. Machado and G. B. Lima Neto. Simulations of gas sloshing in galaxy cluster Abell 2052. *MNRAS*, 447:2915–2924, March 2015. doi: 10.1093/mnras/stu2669.

[154] R. M. Johnstone, A. C. Fabian, and G. B. Taylor. X-ray and radio observations of the poor cluster A3581 which hosts the radio galaxy PKS 1404-267. *MNRAS*, 298:854–860, August 1998. doi: 10.1046/j.1365-8711.1998.01675.x.

[155] R. M. Johnstone, A. C. Fabian, R. G. Morris, and G. B. Taylor. The galaxy cluster Abell 3581 as seen by Chandra. *MNRAS*, 356:237–246, January 2005. doi: 10.1111/j.1365-2966.2004.08445.x.

[156] F. A. Aharonian et al. Hitomi constraints on the 3.5 keV line in the Perseus galaxy cluster. *Astrophys. J.*, 837(1):L15, 2017. doi: 10.3847/2041-8213/aa61fa.

[157] J. Wilms, T. Brand, D. Barret, T. Beuchert, J.-W. den Herder, I. Kreykenbohm, S. Lotti, N. Meidinger, K. Nandra, P. Peille, L. Piro, A. Rau, C. Schmid, R. K. Smith, C. Tenzer, M. Wille, and R. Willingale. ATHENA end-to-end simulations. In *Space Telescopes and Instrumentation 2014: Ultraviolet to Gamma Ray*, volume 9144 of *Proc. SPIE*, page 91445X, July 2014. doi: 10.1117/12.2056347.

[158] Dmytro Iakubovskyi. Observation of the new line at 3.55 keV in X-ray spectra of galaxies and galaxy clusters. 2015.

[159] Douglas P. Finkbeiner and Neal Weiner. X-ray line from exciting dark matter. *Phys. Rev.*, D94(8):083002, 2016. doi: 10.1103/PhysRevD.94.083002.

[160] James M. Cline and Andrew R. Frey. Nonabelian dark matter models for 3.5 keV X-rays. *JCAP*, 1410:013, 2014. doi: 10.1088/1475-7516/2014/10/013.

[161] James M. Cline and Andrew R. Frey. Consistency of dark matter interpretations of the 3.5 keV x-ray line. *Phys. Rev.*, D90(12):123537, 2014. doi: 10.1103/PhysRevD.90.123537.

[162] Asher Berlin, Anthony DiFranzo, and Dan Hooper. 3.55 keV line from exciting dark matter without a hidden sector. *Phys. Rev.*, D91(7):075018, 2015. doi: 10.1103/PhysRevD.91.075018.

[163] Prateek Agrawal, Zackaria Chacko, Can Kilic, and Christopher B. Verhaaren. A Couplet from Flavored Dark Matter. *JHEP*, 08:072, 2015. doi: 10.1007/JHEP08(2015)072.

[164] Francesco D'Eramo, Kevin Hambleton, Stefano Profumo, and Tim Stefaniak. Dark matter inelastic up-scattering with the interstellar plasma: A new source of x-ray lines, including at 3.5 keV. *Phys. Rev.*, D93(10):103011, 2016. doi: 10.1103/PhysRevD.93.103011.

[165] Michele Cicoli, Joseph P. Conlon, M. C. David Marsh, and Markus Rummel. 3.55 keV photon line and its morphology from a 3.55 keV axionlike particle line. *Phys. Rev.*, D90:023540, 2014. doi: 10.1103/PhysRevD.90.023540.

[166] Joseph P. Conlon and Francesca V. Day. 3.55 keV photon lines from axion to photon conversion in the Milky Way and M31. *JCAP*, 1411:033, 2014. doi: 10.1088/1475-7516/2014/11/033.

[167] Joseph P. Conlon and Andrew J. Powell. A 3.55 keV line from DM $\to a \to \gamma$: predictions for cool-core and non-cool-core clusters. *JCAP*, 1501(01):019, 2015. doi: 10.1088/1475-7516/2015/01/019.

[168] Pedro D. Alvarez, Joseph P. Conlon, Francesca V. Day, M. C. David Marsh, and Markus Rummel. Observational consistency and future predictions for a 3.5 keV ALP to photon line. *JCAP*, 1504(04):013, 2015. doi: 10.1088/1475-7516/2015/04/013.

[169] Maurizio Giannotti, Jaime Ruz, and Julia K. Vogel. IAXO, next-generation of helioscopes. In *Proceedings, 38th International Conference on High Energy Physics (ICHEP 2016): Chicago, IL, USA, August 3-10, 2016*, 2016. URL `http://inspirehep.net/record/1498053/files/arXiv:1611.04652.pdf`.

[170] Yonatan Kahn, Benjamin R. Safdi, and Jesse Thaler. Broadband and Resonant Approaches to Axion Dark Matter Detection. *Phys. Rev. Lett.*, 117(14):141801, 2016. doi: 10.1103/PhysRevLett.117.141801.

[171] A. Kogut et al. The Primordial Inflation Explorer (PIXIE): A Nulling Polarimeter for Cosmic Microwave Background Observations. *JCAP*, 1107:025, 2011. doi: 10.1088/1475-7516/2011/07/025.

[172] Philippe André et al. PRISM (Polarized Radiation Imaging and Spectroscopy Mission): An Extended White Paper. *JCAP*, 1402:006, 2014. doi: 10.1088/1475-7516/2014/02/006.

[173] Hiroyuki Tashiro, Joseph Silk, and David J. E. Marsh. Constraints on primordial magnetic fields from CMB distortions in the axiverse. *Phys. Rev.*, D88(12):125024, 2013. doi: 10.1103/PhysRevD.88.125024.

[174] Linhan Chen and Joseph P. Conlon. Constraints on Massive Axion-Like Particles from X-ray Observations of NGC1275. 2017. doi: 10.1093/mnras/sty1591.