

RECEIVED: May 12, 2025

REVISED: August 20, 2025

ACCEPTED: October 20, 2025

PUBLISHED: December 1, 2025

Search for solar axions produced through the axion-electron coupling g_{ae} using a new GridPix detector at CAST

The CAST collaboration

K. Altenmüller,^a V. Anastassopoulos,^b S. Arguedas-Cuendis^{ID},^c S. Aune,^d J. Baier,^e K. Barth,^c H. Bräuninger,^{f,†} G. Cantatore^{ID},^g F. Caspers,^{c,h} J.F. Castel,^a S.A. Çetin^{ID},ⁱ F. Christensen,^j C. Cogollos,^{k,a} T. Dafni,^a M. Davenport^{ID},^c T.A. Decker,^l K. Desch^{ID},^m D. Díez-Ibáñez,^a B. Döbrich^{ID},^c E. Ferrer-Ribas^{ID},^d H. Fischer^{ID},^e W. Funk^{ID},^c J. Galán^{ID},^a J.A. García,^a A. Gardikiotis,ⁿ I. Giomataris,^d J. Golm^{ID},^{c,o} C.H. Hailey^{ID},^p M.D. Hasinoff^{ID},^q D.H.H. Hoffmann^{ID},^r I.G. Irastorza^{ID},^a J. Jacoby,^e A.C. Jakobsen,^j K. Jakovčić^{ID},^s J. Kaminski^{ID},^m M. Karuza^{ID},^{g,t} S. Kostoglou,^c C. Krieger^{ID},^{m,1} B. Lakić,^{s,†} J.M. Laurent,^c G. Luzón,^a C. Malbrunot,^c C. Margalejo^{ID},^a M. Maroudas^{ID},^u L. Miceli,^v H. Mirallas,^a P. Navarro,^w L. Obis,^a A. Özbel^{ID},^{i,x} K. Özbozduman^{ID},^{i,y} T. Papaevangelou^{ID},^d O. Pérez,^a M.J. Pivovaroff^{ID},^l M. Rosu^{ID},^z E. Ruiz-Chóliz,^a J. Ruz^{ID},^{l,a} T. Schiffer^{ID},^{m,*} S. Schmidt^{ID},^{m,*} M. Schumann^{ID},^e Y.K. Semertzidis^{ID},^{v,aa} S.K. Solanki^{ID},^{ab} L. Stewart,^c T. Vafeiadis^{ID},^c J.K. Vogel^{ID},^{l,a} and K. Zioutas^{ID},^{c,b}

^aCentro de Astropartículas y Física de Altas Energías (CAPA) & Departamento de Física Teórica, University de Zaragoza, 50009 — Zaragoza, Spain

^bPhysics Department, University of Patras, Patras, Greece

^cEuropean Organization for Nuclear Research (CERN), 1211 Geneva 23, Switzerland

^dIRFU, CEA, Université Paris-Saclay, 91191 Gif-sur-Yvette, France

^ePhysikalisches Institut, Albert-Ludwigs-Universität Freiburg, 79104 Freiburg, Germany

^fMax-Planck-Institut für Extraterrestrische Physik, Garching, Germany

^gUniversity of Trieste and Istituto Nazionale di Fisica Nucleare (INFN), Sezione di Trieste, Trieste, Italy

^hEuropean Scientific Institute, Archamps, France

ⁱIstanbul University, Dept. of Basic Sciences, 34396 Sariyer, Istanbul, Turkey

^jDTU Space, National Space Institute, Technical University of Denmark, 2800 Lyngby, Denmark

^kInstitut de Ciències del Cosmos, Universitat de Barcelona (UB-IEEC), Barcelona, Catalonia, Spain

^lLawrence Livermore National Laboratory, Livermore, California 94550, U.S.A.

*Corresponding author.

†Deceased.

¹Now at Universität Hamburg, Hamburg, Germany.

^m*Physikalisches Institut, University of Bonn, 53115 Bonn, Germany*

ⁿ*Institute of Quantum Computing and Quantum Technology NCSR “Demokritos”, Athens, Greece*

^o*Institute for Optics and Quantum Electronics, Friedrich Schiller University Jena, Jena, Germany*

^p*Physics Department and Columbia Astrophysics Laboratory, Columbia University, New York, New York 10027, U.S.A.*

^q*Department of Physics and Astronomy, University of British Columbia, Vancouver, Canada*

^r*Xi’An Jiaotong University, School of Science, Xi’An, 710049, China*

^s*Rudjer Bošković Institute, Zagreb, Croatia*

^t*Faculty of Physics and Center for Micro and Nano Sciences and Technologies, University of Rijeka, 51000 Rijeka, Croatia*

^u*Universität Hamburg, Hamburg, Germany*

^v*Center for Axion and Precision Physics Research, Institute for Basic Science (IBS), Daejeon 34141, Republic of Korea*

^w*Department of Information and Communications Technologies, Technical University of Cartagena, 30203 — Murcia, Spain*

^x*Istanbul University-Cerrahpasa, Vocational School of Technical Sciences, 34320 Avcılar, Istanbul, Turkey*

^y*Boğaziçi University, Physics Department, Bebek, Istanbul, Turkey*

^z*Extreme Light Infrastructure — Nuclear Physics (ELI-NP), 077125 Magurele, Romania*

^{aa}*Department of Physics, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 34141, Republic of Korea*

^{ab}*Max-Planck-Institut für Sonnensystemforschung, 37077 Göttingen, Germany*

E-mail: schiffer@physik.uni-bonn.de, phd@vindaar.de

ABSTRACT: We present a search for solar axions produced through the axion-electron coupling (g_{ae}) using data from a novel 7-GridPix detector installed at the CERN Axion Solar Telescope (CAST). The detector, featuring ultra-thin silicon nitride windows and multiple veto systems, collected approximately 160 hours of solar tracking data between 2017–2018. Using machine learning techniques and the veto systems, we achieved a background rate of $1.06 \times 10^{-5} \text{ keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1}$ at a signal efficiency of about 80% in the 0.2 to 8 keV range. Analysis of the data yielded no significant excess above background, allowing us to set a new upper limit on the product of the axion-electron and axion-photon couplings of $g_{ae} \cdot g_{a\gamma} < 7.35 \times 10^{-23} \text{ GeV}^{-1}$ at 95% confidence level for axion masses below 10 meV. This result improves upon the previous best helioscope limit and demonstrates the potential of GridPix technology for rare event searches. Additionally, we derived a limit on the axion-photon coupling of $g_{a\gamma} < 9.0 \times 10^{-11} \text{ GeV}^{-1}$ at 95% CL, which, while not surpassing CAST’s best limit, provides complementary constraints on axion models.

KEYWORDS: Beyond Standard Model, CP Violation, Dark Matter, Dark Matter and Double Beta Decay (experiments)

Contents

1	Introduction	1
2	CAST setup and detector	2
3	Data taking	6
4	Data analysis and background reduction	6
5	Limit calculation method	9
5.1	Signal and background in detail	11
5.2	Systematics	12
5.3	Evaluation of parameters	13
6	Data unblinding and g_{ae} limit	13
7	Summary and conclusion	15

1 Introduction

The strong CP-problem [1–3] has a solution in the form of the QCD axion [4–7], which arises as a pseudo Nambu-Goldstone boson if a new scalar field α is added to the Standard Model Lagrangian. This promotes the θ parameter of the θ -term

$$\mathcal{L}_\theta = \theta \frac{g_s^2}{32\pi^2} G_a^{\mu\nu} \tilde{G}_{a\mu\nu},$$

to a quantum field and leads to an effective coupling of the axion to photons via a quark-gluon loop. As the core of the Sun provides an environment of high rates of photon interactions, axions are expected to be produced there following a spectrum similar to that of blackbody radiation, with a rate proportional to $g_{a\gamma}^2$. These axions are searched for in axion helioscope experiments like the **CERN Axion Solar Telescope** (CAST) [8, 9] and in future the **International AXion Observatory** (IAXO) [10–12].

In the DFSZ [13, 14] invisible axion model, a tree level coupling of the axion to electrons g_{ae} is present $\mathcal{L} \supset g_{ae} \frac{\partial_\mu a}{2m_e} \bar{\psi}_e \gamma^5 \gamma^\mu \psi_e$. It can be shown [15] that this coupling can produce a dominant contribution — depending on the choice of coupling parameter values — to the solar axion flux by Compton, Bremsstrahlung and photon-nucleus interactions in the Sun. In contrast to the axion-photon flux with a peak energy at around 3 keV the axion-electron contributions peak at only around 1 keV. See figure 4(a) for comparison of the two spectra.

The possible parameter range for g_{ae} has been previously limited by experiments. Some experimental approaches are sensitive exclusively to the axion-electron coupling, while others rely on the axion-photon coupling either for production or reconversion. A recent summary of the theory of direct axion-electron sensitive experiments is found in [16], which also includes

an overview of the current best limits in different mass ranges. The strongest direct detection limit over all axion masses of interest comes from XenonNT [17]. From astrophysical processes the most stringent limits of g_{ae} over a wide range of axion masses can be obtained from the brightness of tip of the red-giant branch (TRGB) stars. Axion production would induce more cooling, leading to a larger core mass at helium burning ignition and thus brighter TRGB stars. Such limits are of the order of $g_{ae} < 1.3 \times 10^{-13}$ at 95 % CL [18, 19].

For narrow mass ranges much tighter limits can be set with different approaches. See also [20] for a very detailed overview of the current best limits.

Helioscope experiments are only sensitive to the product of $g_{ae} \cdot g_{a\gamma}$, due to the reliance on magnets to reconvert the solar axions back into photons in the X-ray regime. However, they provide limits for a wide range of axion masses up to 10 meV. The current best such limit is $g_{ae} \cdot g_{a\gamma} < 8.1 \times 10^{-23} \text{ GeV}^{-1}$ from CAST in 2013, [21].

In this paper we will present a new search for axions produced via the axion-electron coupling g_{ae} based on data taken at CAST in 2017 and 2018 using a 7-GridPix [22–27] detector for X-ray detection as the readout. First we will introduce the detector in section 2. Afterwards, we will give an overview of the data taking campaign, section 3, and then move to the data analysis and background reduction techniques employed in section 4. This leads to the limit calculation methodology used in section 5. The result of the limit calculation is then presented in section 6, before we conclude in section 7.

For more details about this work, see [28].

2 CAST setup and detector

The CAST experiment [8, 9] provides four experiment places to install detectors. In the last years of operation only the side of the magnet pointing towards sunrise were in operation, as there were X-ray optics behind each of the two magnet bores of the LHC prototype dipole magnet. One of these optics was a custom built NuSTAR-like X-ray optic, consisting of slumped glass layers, which are conical approximations [29] to the hyperbolic and parabolic mirrors of a Wolter type I optic. It was specifically designed for CAST, hence the effective area was optimized in the energy range of interest for axion searches [30]. The optic focuses the produced X-rays onto a focal plane 1500 mm behind the optic. The GridPix detector used for the described measurement was connected with a vacuum system to the optic such that the focal plane of the optic is located in the center of the gas volume. A render of the detector can be seen in figure 1. It uses a similar design in terms of dimensions as the MicroMegas (MM) detector used in [31]. A 78 mm diameter acrylic glass housing with an inner height of 30 mm to contain a gas mixture of Argon/Isobutane 97.7 %/2.3 % at 1050 mbar(a) chamber pressure. As a barrier between the gas and the vacuum, a 300 nm thick silicon-nitride SiN_x window is employed to provide X-ray transmission while withstanding the 1050 mbar pressure difference. Figure 4(b) shows the transmission efficiency of the window as well as the absorption efficiency of X-rays in the energy range of interest. The usage of ultra-thin SiN_x windows improves the X-ray transmission in the interesting range from 0–3 keV by 8.13 % compared to commonly used 2 μm Mylar windows in the energy range of interest for searches for g_{ae} . Figure 4(a) shows the differential solar axion flux expected when taking the axion-electron coupling into account. The used coupling constants are $g_{ae} = 10^{-13}$ and $g_{a\gamma} = 10^{-12} \text{ GeV}^{-1}$.

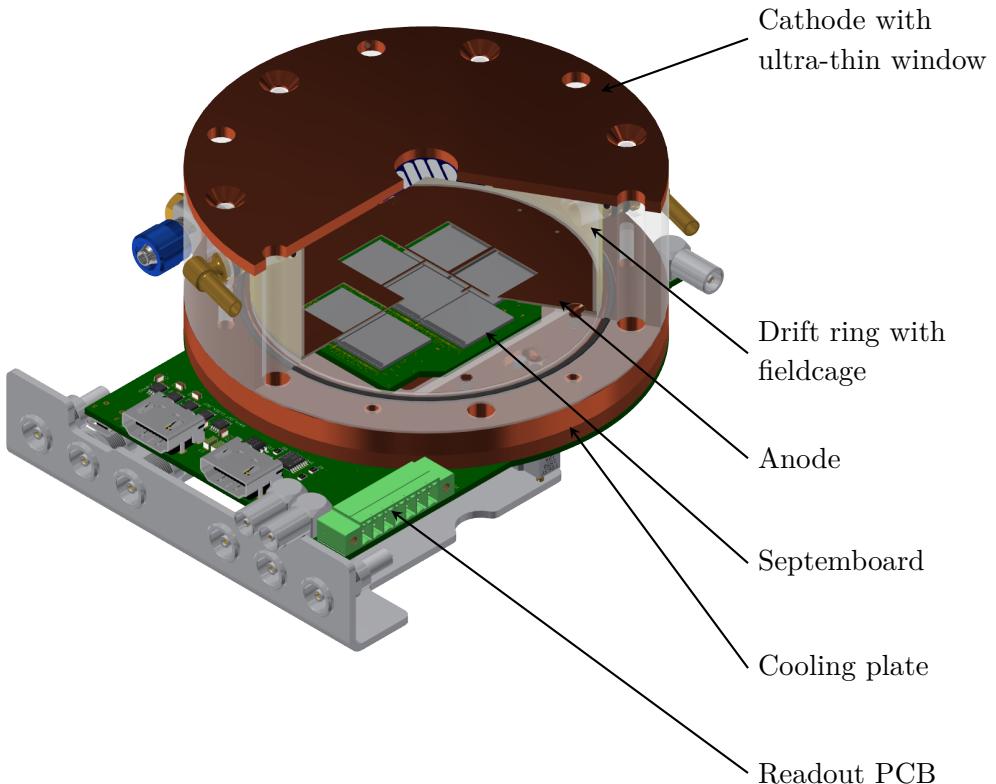


Figure 1. Render of the 7-GridPix detector, cut open. Visible is the layout of the ‘Septemboard’, with seven GridPixels as the readout. Not visible is a small SiPM-based scintillator installed below the PCB.

A render of the setup is shown in figure 2 and an annotated image of the real setup in figure 3. Since the expected image of the converted solar axions on the detector is of great importance for the data analysis, a dedicated raytracing simulation for X-rays was developed [32] to simulate the expected axion image, based on the expected axion emission in the Sun, conversion probability in the magnet, detailed description of the X-ray optics and detector. The raytracing code agrees with measurements of the telescope taken at the PANTER test facility within 2.6 % when comparing the half power diameter of three different X-ray lines, Al K α (1.49 keV), Ti K α (4.51 keV) and Fe K α (6.41 keV).

The detector itself is a gas ionisation chamber using the GridPix technology [22–25]. A GridPix consists of an integrated grid structure, as a gas amplification stage, on top of a Timepix ASIC. The grid is produced in a photolithographic process to perfectly align one hole with each pixel. This enables both sufficient amplification and spatial resolution to

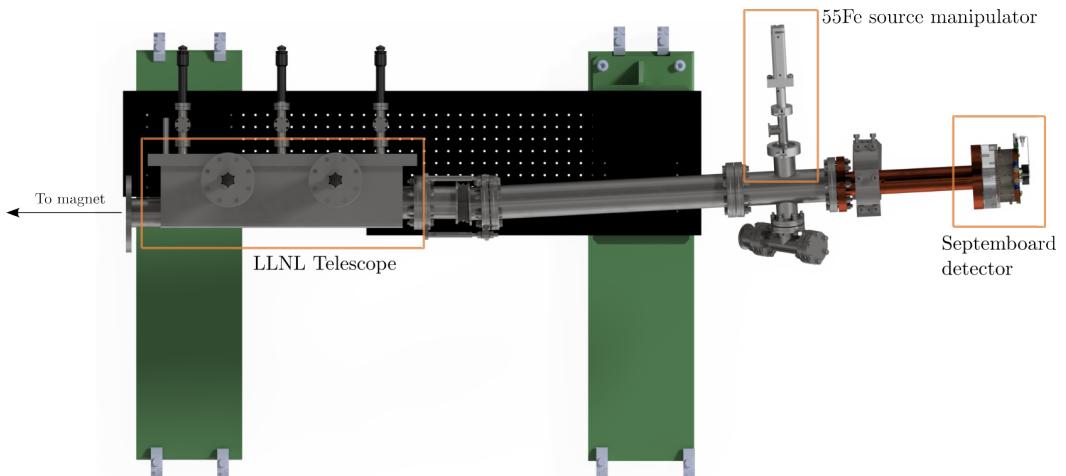


Figure 2. Render of the detector setup up to the magnet end cap as seen from above. The beamline kinks away from the other beamline ("below" in this image) to provide more space for two detectors at the same time.

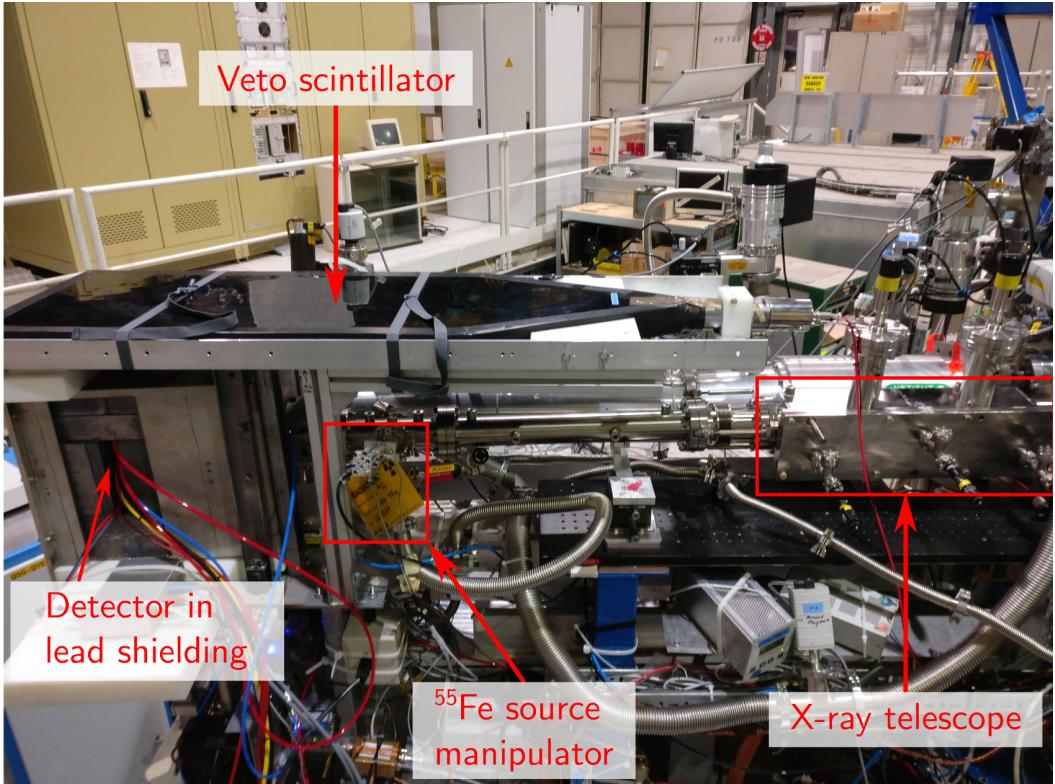


Figure 3. Annotated setup as installed in October 2017 for the first data taking campaign. The detector is housed within its lead shielding, with the veto scintillator covering a large angular portion above the detector. The ^{55}Fe source manipulator is seen head-on. On the right towards the magnet we see the housing of the X-ray telescope.

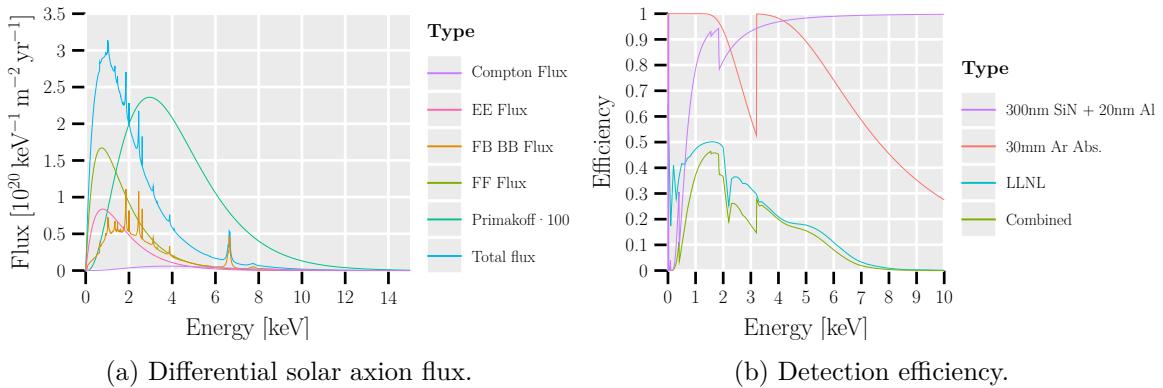


Figure 4. (a): differential solar axion flux assuming a distance to the Sun of 0.989 AU. ‘Primakoff·100’ corresponds to the pure axion-photon $g_{a\gamma}$ production in the Sun for $g_{a\gamma} = 10^{-12} \text{ GeV}^{-1}$, multiplied by a factor 100 to make this contribution visible for the choice of parameters. The other contributions are all via the axion-electron coupling at a value of $g_{ae} = 10^{-13}$. ‘EE’: electron-electron bremsstrahlung, ‘FB BB’: free-bound, bound-bound interactions (axio-recombination and axio-deexcitation), ‘FF’: free-free transition (electron-ion bremsstrahlung). (b): detection efficiency of the detector based on the window (SiN_4 and Al) transmission, gas absorption (‘30 mm Ar Abs.’) and X-ray telescope effective area (‘LLNL’).

detect individual primary electrons produced in the ionisation chamber. At the center of the detector is the ‘Septemboard’ — a PCB equipped with seven GridPixels. The detector is aligned such that the expected axion image is focused on the central GridPix. The outer six GridPixels extend the sensitive area to allow for additional background rejection. Because the seven GridPixels combined produce a significant amount of heat due to thermal losses a custom-built water cooling system is installed below the intermediate board. It is based on an oxygen-free copper body with 3 mm diameter channels for water transport. This setup provides sufficient cooling for the GridPixels operation.

The GridPixels used in this detector are based on the Timepix ASIC [33], which is read out in a frame based fashion. The Timepix consists of 256×256 pixels with $55 \mu\text{m}$ pitch and an active area of $1.4 \times 1.4 \text{ cm}^2$. Each pixel can either be used to record Time-over-Threshold (ToT) values for charge accumulation or Time-of-Arrival (ToA) information, but not both. For this reason an additional flash ADC (FADC) is connected to the grid of the central GridPix. It both acts as a readout trigger in case a fixed threshold is exceeded while also providing further charge as well as longitudinal time information. Reading out the seven GridPixels needs 175 ms due to the daisy-chained data output of the GridPixels. Thus, short shutter times would lead to a large dead time ratio. Finally, two scintillators behind and above the detector are installed to record potential muons traversing the detector. To reduce the amount of direct background from cosmics the detector sits inside of a 10 cm lead housing. A ^{55}Fe calibration source is available at CAST and can be inserted into the beamline in front of the detector using a pneumatic system. It provides a reference for detector stability via daily calibration runs, for energies of 5.9 keV (photopeak) and 3 keV (argon escape peak).

	Solar tracking [h]	Active s. [h]	Background [h]	Active b. [h]	Active [%]
Run-1	106.01	93.37	2391.16	2144.12	89.65
Run-2	74.30	67.01	1124.93	1012.68	90.02
Total	180.30	160.38	3516.09	3157.35	89.52

Table 1. Overview of the total data taken with the Septemboard detector at CAST in the time between October 2017 and December 2018. ‘Active s.’ and ‘Active b.’ refers to the total solar tracking and background time excluding the dead time due to readout of the detector.

3 Data taking

The detector described in section 2 was first installed at CAST in September 2017. The first data taking period started in Oct 2017 and went until April 2018 (Run-1). As CAST was typically shut down over the summer months, the data taking continued from Oct 2018 to Dec 2018 (Run-2). During this data taking about 3500 h of background data and 180 h of solar tracking data were recorded (see table 1). The detector ran with a dead time ratio of about 10% thanks to a long shutter time of about 2.2 s. Data was taken in ToT mode to later compute the recorded charge of each pixel. Due to a firmware bug, scintillator triggers were only saved in Run-2.

4 Data analysis and background reduction

The data is read out from the detector using the Timepix Operating Software (TOS) [34, 35]. Frames of hit pixels are saved as ASCII files containing signal or background clusters. Hit pixels originate typically from individual ionization electrons produced by the interacting particles, spread out by diffusion due to the drift in the gas volume. The files are parsed and stored in HDF5 [36] files using the TimepixAnalysis framework [37], which handles the entire analysis chain from raw data processing through final limit calculation. The data reconstruction is based on analyzing individual clusters of pixels. Expected signals from X-rays are roughly spherical in shape whereas the dominant background sources are less symmetrical (e.g. tracks from muons). The background rejection algorithms are therefore based on geometric properties of the clusters.

As part of the data reconstruction, first a distance based clustering algorithm is used to reconstruct individual clusters from raw data. From here, each cluster is handled separately. First the long and short axes are determined via optimization of the maximum pixel-to-pixel distance and its orthogonal extent. The long and short axes provide the basis for a variety of different geometric properties. The geometric properties derived from these axes, along with other cluster characteristics, are shown in table 2.

Next, the ToT values of each pixel are converted into a charge in electrons. Based on 90 min time intervals all charge values recorded are accumulated to a histogram, from which the gas gain can be extracted via a Pólya distribution fit (following [38]). This calibration works with any dataset — background or calibration — as it depends only on individual pixel charge distributions. Calibration runs are used to correlate potentially varying gas gains over time with a linear calibration function mapping recorded charges to target energies of the

Property	Meaning
Eccentricity	eccentricity of the cluster
RMS & Skewness & Kurtosis	RMS/skewness/kurtosis along long/short axis
Length & Width	size along long/short axis
LengthDivRmsTrans	length divided by transverse RMS
RotationAngle	rotation angle of long axis over chip coordinate system
FractionInTransverseRms	fraction of pixels within transverse RMS around center
TotalCharge	integrated charge of total cluster in electrons
D_T	transverse gas diffusion coefficient

Table 2. Table of all the properties of a single cluster used as input for the MLP. RMS, skewness and kurtosis are computed for each axis individually.

^{55}Fe calibration source. This allows to compute the energy of any cluster from its charge by using the associated gas gain from the 90 min interval it is part of.

With the data fully reconstructed and calibrated, each cluster needs to be classified as either background- or signal-like. The method employed to compute the final limit follows.

The initial classification is based on a machine learning approach using a Multi-Layer Perceptron (MLP) [39–41],¹ a fully connected feed-forward network. The specific network used has 14 input neurons, which correspond to twelve geometric properties of the input cluster, its total charge and the gas diffusion coefficient D_T during the time this cluster was taken. All properties are shown in table 2. D_T is determined by an iterative optimization approach comparing the transverse RMS distribution via the Cramér-von-Mises criterion [42–44] of simulated data with that of real data of a data taking run. The simulation is a Monte Carlo code, which takes a diffusion coefficient D_T , gas gain and target energy as an input. It samples a conversion point for an X-ray and models active pixels via expected electron drift to the readout using D_T . Each electron receives a charge based on sampling from a Pólya distribution of the given gas gain. The transverse RMS of a cluster is directly proportional to the transverse diffusion coefficient D_T . To determine D_T for a background run, we make use of the fact that along the short axis of a cluster both background tracks and X-ray clusters only undergo diffusion from a single line and point, respectively.

The MLP has two hidden layers with 30 neurons each and two output neurons, one for X-ray-like and one for background-like clusters. It uses `tanh` as an activation function and the mean squared error (MSE) as a loss function. Training was run over 83 000 epochs of 250 000 X-ray and background events each. The background events are taken from the extensive set of background data of the *outer chips*, while the X-ray events are fully simulated. It is the same event simulation as for the determination of the gas diffusion. X-rays of a wide range of gas diffusion coefficients, gas gains and uniform energies are generated. The distributions

¹Due to the rich and long history of artificial neural networks picking “a” or only a few citations is tricky. Rosenblatt [39] introduced the perceptron while Amari’s work [40] first combined a perceptron with a non-linear activation function and used gradient descent for training. See Schmidhuber’s recent overview for a detailed history leading up to modern deep learning [41].

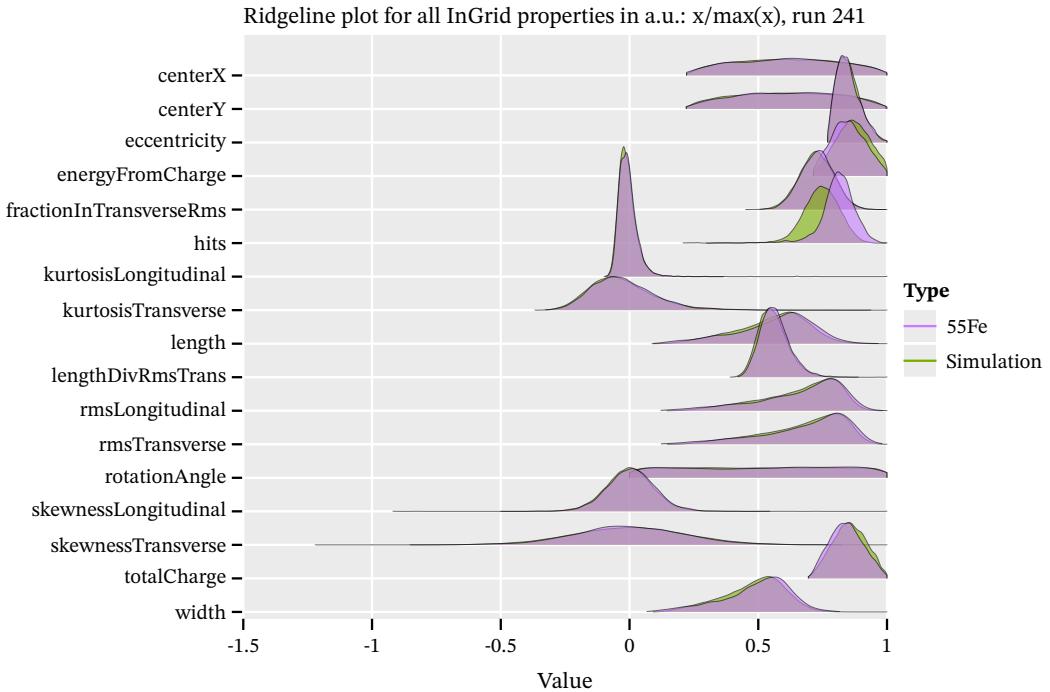


Figure 5. Comparison of all geometric properties of simulated and real data of the ^{55}Fe photopeak clusters of run 241 in a ridgeplot. Each ridge shows a KDE of the data. The gas gain and gas diffusion were first extracted from the real run data to simulate events of the photopeak. The two datasets agree very well with the exception of the number of hits (expected).

of geometric properties between simulated and real data with the same parameters agree very well, see figure 5. The total number of hits per cluster differs between simulation and real data, as the simulation uses a simplified model for neighboring pixel activation due to UV photons from the gas amplification. However, this does not affect the MLP training, since it only depends on the derived geometric properties which are well reproduced by the simulation. Based on the Kolmogorov-Smirnov test statistic, the largest deviation across the empirical distribution functions for simulated and real calibration data is below 5 % when averaged over all properties (except the ‘hits’ parameter) and all runs.

To differentiate signal from background events we define a cut on the output value of one of the neurons (as the two output neurons are complementary) by desiring a specific software efficiency for X-ray data:

$$\epsilon_S \equiv \epsilon_{\text{eff}} = \frac{\int_0^{x'} \text{MLP}(c) \, dx}{\int_0^{\infty} \text{MLP}(c) \, dx} = \text{EDF}(x')$$

where $\text{MLP}(c)$ corresponds to the MLP output for a given cluster c and the value x' then corresponds to the value to cut at. As indicated this is computed based on the empirical cumulative distribution function (EDF) in practice.

While the MLP cut value is determined using simulated events, the actual signal efficiency used in the limit calculation is calibrated using real X-ray data: both ^{55}Fe calibration data and measurements at eight different energies from an X-ray generator. Different energies are useful,

because any model working on similar gaseous detector data is inherently energy dependent. Lower energies mean fewer activated pixels, which reduces the possible differentiation between X-rays and background data.

The scintillators provide an additional veto: events are rejected if the center GridPix shows activity within $2.5\text{ }\mu\text{s}$ of a scintillator trigger, using the FADC as the GridPix readout trigger for X-ray energies above $1 - 2\text{ keV}$.

For clusters including FADC data, this data is also used to veto some more events. It effectively records the longitudinal shape information of the clusters. This is determined by the event shape and the longitudinal drift velocity and diffusion. These can be estimated from theory and the ^{55}Fe calibration datasets provide further validation. Based on this, all events outside the 1st and 99th percentile of the FADC signal rise time distribution are removed.

Finally, the outer ring of GridPixels is used as a veto. In case of an X-ray like cluster on the center chip, the data of the outer chips is considered. If a cluster on the outer chip is found whose long axis points towards the center cluster, the event is vetoed.

The last veto comes with an efficiency penalty due to potential for random coincidences. Shutter times are very long, yet only the center chip has a trigger. There is the chance of uncorrelated activity on the outer chip. To estimate this new events are bootstrapped from a combination of center GridPix clusters that pass the MLP cut and data from the outer chips of any other event. The rate of triggering this veto then is purely random coincidence. Using this approach the efficiency of this veto is estimated to 86 %.

Combining the MLP cut with all vetos, we achieve a background rate over the center $5 \times 5\text{ mm}^2$ area of the center chip (region of lowest background) of $1.056(69) \times 10^{-5}\text{ keV}^{-1}\text{ cm}^{-2}\text{ s}^{-1}$ within 0.2 to 8 keV, shown in figure 6. This is for the case of an effective MLP efficiency of 95 %.

5 Limit calculation method

To compute a limit on the axion-electron coupling constant we use a Bayesian approach based on finding the 95th percentile of the marginal posterior likelihood. Our initial likelihood function is derived from a ratio of two Poisson distributions, the signal plus background hypothesis over the pure background hypothesis:

$$\mathcal{L} = \prod_i \frac{P_{\text{pois}}(c_i; s_i + b_i)}{P_{\text{pois}}(c_i; b_i)}$$

which runs over all bins i and c_i are the number of candidates in each bin, s_i, b_i the signal and background, respectively. The signal s_i is the expected amount of signal in bin i based on the solar axion flux and all detection efficiencies included. The background is given by a background model constructed from the entire background dataset taken at CAST during non-tracking times. This likelihood is taken to the unbinned likelihood by choosing bins in time such that each bin only contains either 0 or 1 candidates.

The likelihood function simplifies to

$$\mathcal{L} = e^{-s_{\text{tot}}} \prod_i \left(1 + \frac{s_i}{b_i}\right)$$

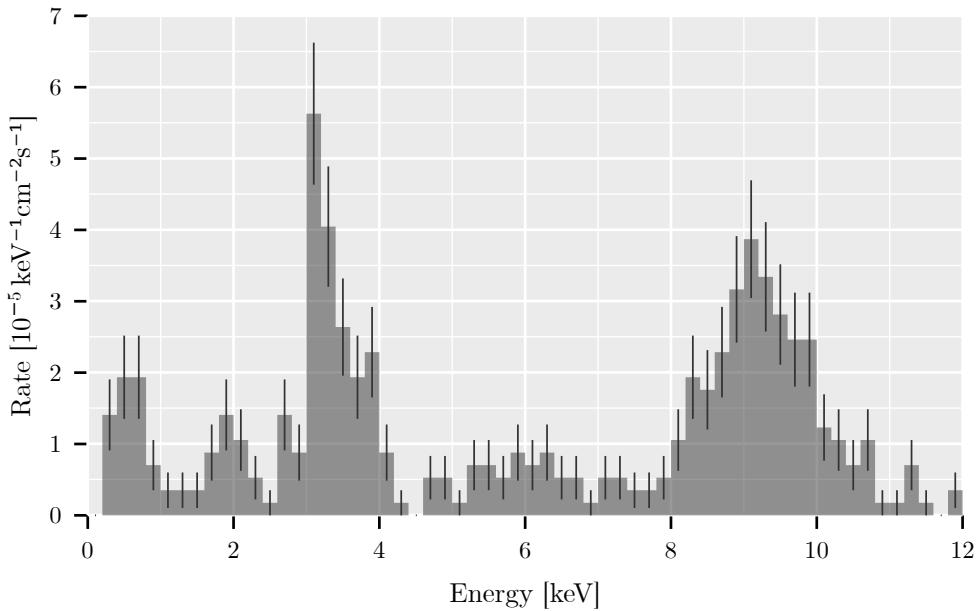


Figure 6. Background rate in the center $5 \times 5 \text{ mm}^2$ using the MLP at 95 % effective efficiency and all mentioned vetoes. Over the energy range of 0.2 to 8 keV an average background rate of $1.056(69) \times 10^{-5} \text{ keV}^{-1} \text{ cm}^{-2} \text{ s}^{-1}$ is achieved. The rate can be improved at the cost of signal efficiency.

in this case. Here s_{tot} represents the total expected signal over all signal sensitive data taking periods, i.e. a total number of expected axion-induced X-rays recorded by our detector.

Further, systematics are taken into account by adding four ‘nuisance parameters’, two for signal and background and two for the position. We multiply the likelihood function with one normal distribution for each nuisance parameter, which is normalized such that $\mathcal{N}(\theta = 0, \sigma) = 1$. The signal and background parameters are scaled by θ , such that a positive θ increases the parameter and a negative decreases it. At the same time the normal distribution acts as a penalty term. To compute the limit the explicit θ dependencies must be removed, which is done by marginalization, i.e. integrating them out

$$\mathcal{L}_M = \iiint_{-\infty}^{\infty} e^{-s'_{\text{tot}}} \cdot \prod_i \left(1 + \frac{s''_i}{b'_i}\right) \cdot \exp \left[-\frac{\theta_b^2}{2\sigma_b^2} - \frac{\theta_s^2}{2\sigma_s^2} - \frac{\theta_x^2}{2\sigma_{xy}^2} - \frac{\theta_y^2}{2\sigma_{xy}^2} \right] d\theta_b d\theta_s d\theta_x d\theta_y$$

with $a' = a(1 + \theta_a)$ and $a'' = a(1 + \theta_a)(1 + \theta_x)(1 + \theta_y)$ where $a \in \{s, b\}$. This keeps the possibility of variance due to systematics including the penalization embedded in the marginal likelihood, but restores a single variable likelihood function with a well defined single value for the 95th percentile.

The limit g'_{ae} is then defined by

$$0.95 = \frac{\int_{-\infty}^{g'_{ae}} \mathcal{L}(g_{ae}) \pi(g_{ae}) dg_{ae}}{\int_{-\infty}^{\infty} \mathcal{L}(g_{ae}) \pi(g_{ae}) dg_{ae}} \quad (5.1)$$

where $\pi(g_{ae})$ represents the prior probability distribution for the axion-electron coupling, taken as uniform in $g_{ae} > 0$ and zero elsewhere. It is computed from an empirical cumulative distribution function.

However, the evaluation of a four-fold integral where the integrand is expensive to evaluate numerically, is challenging due to the ‘curse of dimensionality’. As such the Metropolis-Hastings Markov Chain Monte Carlo algorithm [45, 46] is used to evaluate the integrand efficiently only in those regions of the parameter space where the integrand contributes to the integral.

5.1 Signal and background in detail

With an unbinned likelihood approach we use the position on the center GridPix and its energy as the parameters of interest. This means we need to be able to evaluate the expected signal as well as the background at the arbitrary positions and energies that each candidate may have.

The expected signal s_i can be expressed as

$$s_i(gae) = f(gae, E_i) \cdot A \cdot t \cdot P_{a \rightarrow \gamma, \text{vacuum}} \cdot \epsilon(E_i) \cdot r(x_i, y_i) \quad (5.2)$$

where $f(gae, E_i)$ is the solar axion flux at energy E_i as a function of g_{ae} , the area of the magnet bore A , the total tracking time t . $\epsilon(E_i)$ describes the energy dependence of the detection efficiency, which combines the telescope effective area ϵ_{tel} , detector window transparency ϵ_{window} , absorption probability of X-rays in the gas ϵ_{gas} , veto efficiency ϵ_{veto} and software efficiency ϵ_{S} ,

$$\epsilon(E_i) = \epsilon_{\text{tel}}(E_i) \cdot \epsilon_{\text{window}}(E_i) \cdot \epsilon_{\text{gas}}(E_i) \cdot \epsilon_{\text{veto}} \cdot \epsilon_{\text{S}}. \quad (5.3)$$

$P_{a \rightarrow \gamma, \text{vacuum}}$ is the conversion probability of the axion converting into a photon in the magnetic field B of length L , given by

$$P_{a \rightarrow \gamma, \text{vacuum}} = \epsilon_0 \hbar c^3 \left(\frac{g_{a\gamma} B L}{2} \right)^2.$$

Finally, $r(x_i, y_i)$ is the distribution of the axion flux over the detection area. This is computed using the raytracing simulation mentioned in section 2 based on the expected axion emission in the solar core and the X-ray telescope optics. This raytracing result is shown as part of figure 9 as a colored, partially transparent region.²

In order to achieve a smooth and continuous interpolation of the background over the entire chip, we construct a background interpolation based on all clusters found in the background dataset. This is done by defining b_i as a function of candidate position and energy using

$$b_i(x_i, y_i, E_i) = \frac{I(x_i, y_i, E_i)}{W(x_i, y_i, E_i)}$$

where I is an intensity defined over clusters within a range R and a normalization weight W . The intensity is given by (for clarity without arguments)

$$I(x, y, E) = \sum_{b \in \{\mathcal{D}(\vec{x}_b, \vec{x}) \leq R\}} \mathcal{M}(\vec{x}_b, E_b) = \sum_{b \in \{\mathcal{D}(\vec{x}_b, \vec{x}) \leq R\}} \exp \left[-\frac{1}{2} \mathcal{D}^2 / \sigma^2 \right],$$

²The exact numbers used are $B = 8.8 \text{ T}$, $L = 9.26 \text{ m}$, $d_{\text{bore}} = 43 \text{ mm}$, $t = 160.38 \text{ h}$.

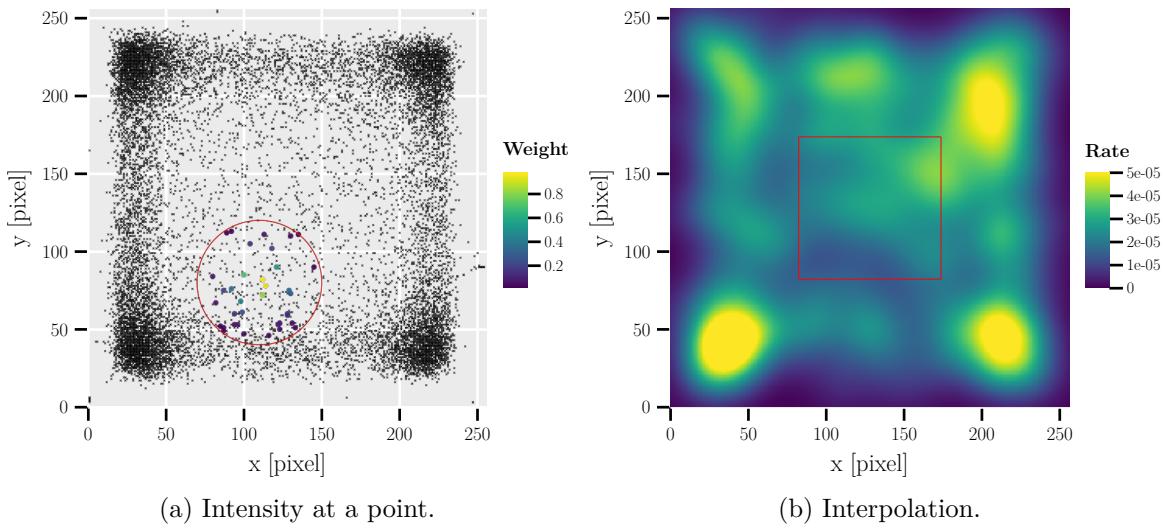


Figure 7. (a): calculation of intensity I at the center of the red circle. Black crosses indicate all background clusters. The red circle indicates cutoff R in the x-y plane. Only clusters with colored dots inside the circle are within 2.70 to 3.3 keV. Their color is the weight based on the gaussian measure \mathcal{M} . (b): example of the resulting background interpolation at 3 keV computed over the entire chip. A smooth, correctly normalized interpolation is obtained.

where we introduce \mathcal{M} to refer to the measure we use and \mathcal{D} to our metric:

$$\mathcal{D}((\vec{x}_1, E_1), (\vec{x}_2, E_2)) = \begin{cases} (\vec{x}_1 - \vec{x}_2)^2 & \text{if } |E_1 - E_2| \leq R \\ \infty & \text{if } (\vec{x}_1 - \vec{x}_2)^2 > R^2 \\ \infty & \text{if } |E_1 - E_2| > R \end{cases}$$

Finally, the normalization weight is the ‘volume’ of our measure within the boundaries set by our metric \mathcal{D} :

$$W(x', y', E') = \int_{\mathcal{D}(\vec{x}, \vec{x}) \leq R} \int_{E' - E_c}^{E' + E_c} \mathcal{M}(x', y') \, dx \, dy \, dE$$

The background clusters participating in the interpolation around a single point in range R and their associated weights for the interpolation are shown in figure 7(a). This leads to the anticipated interpolated background distribution, when applied to the whole chip at a specific energy, as shown in figure 7(b).³

5.2 Systematics

An overview on the systematics regarding the analysis can be found in [28]. The combined values for the systematic uncertainty (including simulation) for the background (σ_b) the signal (σ_s) hypothesis, and the mechanical positioning of the detector (σ_{xy}) used for the

³Note, corrections need to be added towards the chip edges where the radius of the metric is not fully contained on the chip. This is done by upscaling the value within the chip by the missing area.

limit calculation are:

$$\begin{aligned}\sigma_s &\leq 3.38\% \text{ (assuming } \sigma_{\text{software}} = 2\%) \\ \sigma_b &= 0.28\% \\ \sigma_{xy} &= 5\% \text{ (fixed, uncertainty numbers are bounds).}\end{aligned}$$

Note that the systematic of the signal rate σ_s depends on the exact choice of classifier and veto usage and hence the value there is an upper bound assuming a worst-case software systematic contribution of 2%.

5.3 Evaluation of parameters

As the limit calculation method is very sensitive to the combined software efficiency, position and energy distribution of possible candidates, the choice of parameters for the background suppression methods are not fixed. Instead a variety of methods are evaluated and expected limits are computed for each case. The set that yields the best expected limit is the setup used to compute the real limit. Data unblinding of the real tracking candidates was only performed after the expected limits were computed.

An expected limit $\langle L \rangle$ is defined by the median of sets of toy candidate limits L_{t_i} :

$$\langle L \rangle = \text{median}(\{L_{t_i}\}),$$

where each limit is computed based on equation (5.1) using a set of ‘toy candidates’ as inputs. Toy candidates are candidates drawn from the background distribution. As our background model is an interpolation in (x, y, E) , it is split into a grid of $(10, 10, 20)$ grid cells. Each cell then contains an expected (fractional) number of candidates if integrated over the tracking time. This is used as the mean of a Poisson distribution to be sampled from. The number of candidates sampled from said distribution is then distributed over the entire grid cell volume. This produces one set of toy candidates, for which a limit L_{t_i} is computed.

Figure 8 shows an example for the determination of an expected limit for one set of parameters. The blue line corresponds to the median of all computed toy limits. The color of each histogram indicates how many candidates were in the signal sensitive region (determined by $\ln(1 + \frac{s_i}{b_i}) > 0.5$).

Table 3 shows the best four parameter options, which only differ by their software efficiency targeted by the MLP ϵ_s , with the exception of the third row, which uses no vetoes. Other considerations (e.g. different sets of vetoes) produced worse limits.

This implies a best expected limit $g_{ae} \cdot g_{a\gamma} = 7.878(6) \times 10^{-23} \text{ GeV}^{-1}$ for an MLP software efficiency of 95% at a total combined efficiency of 79.7(7)%.

6 Data unblinding and g_{ae} limit

With the expected limits present for different setups, data unblinding was performed based on the setup with the best expected limit. The candidates obtained by this are shown in figure 9, with the axion image underlaid and the energy of each candidate given by the color as well as in text for those within a certain radius around the center. As evident by the figure there are no relevant candidates in terms of position (axion image location, hinted)

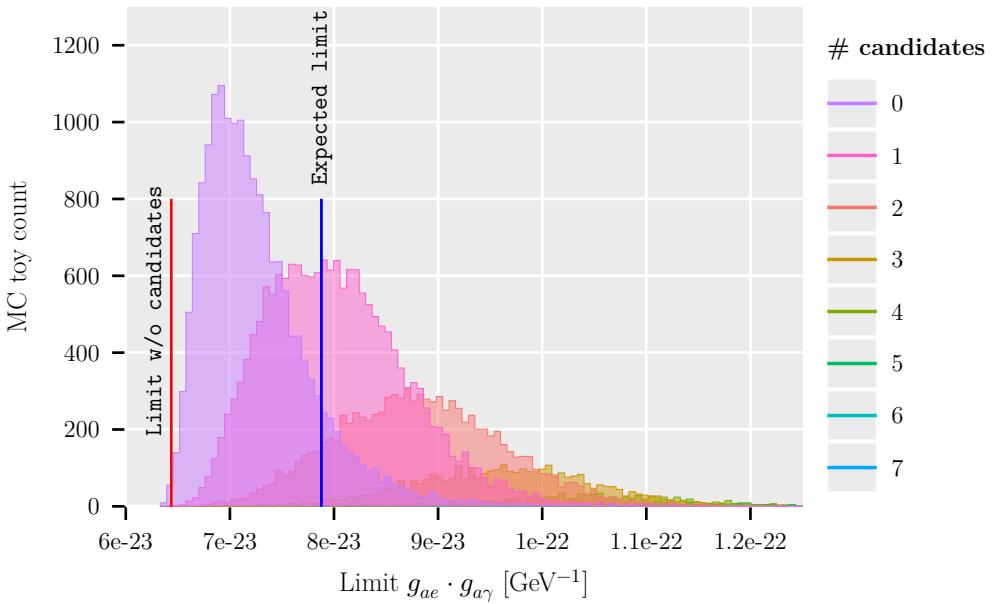


Figure 8. Distribution of limits from 50 000 toy candidate sets for the ‘background only’ hypothesis. The expected limit — the median — is shown as the blue line. The red line shows the limit for the case without any candidates. The different colored histograms correspond to toy sets with a different number of toy candidates in the signal sensitive region (not the *total* number of candidates), defined by $\ln(1 + s_i/b_i) > 0.5$. The most likely number of candidates in the sensitive region seems to be 0, 1 or 2.

Method	ϵ_S	Total eff.	Limit no signal [GeV $^{-1}$]	Expected limit [GeV $^{-1}$]
MLP	0.952(4)	0.797(7)	6.387×10^{-23}	$7.878(6) \times 10^{-23}$
MLP	0.980(3)	0.820(7)	6.20×10^{-23}	$7.88(3) \times 10^{-23}$
MLP †	0.859(6)	0.859(6)	6.11×10^{-23}	$7.95(5) \times 10^{-23}$
MLP	0.906(5)	0.758(8)	6.47×10^{-23}	$7.98(2) \times 10^{-23}$

Table 3. Overview of the efficiencies and expected limits for the best performing four setups. ϵ_S : software efficiency of the MLP, ‘Total eff.’: $\epsilon_S \cdot \epsilon_{\text{veto}}$. The first two entries and the last use the vetoes mentioned in the text, whereas the third entry MLP † is for the MLP only. Difference in uncertainty for first entry is due to larger number of samples.

and energy (axion spectrum, see figure 4(a)) for an axion signal (based on $\ln(1 + s_i/b_i) < 0.5$ for each candidate). This can also be seen in figure 10, which shows the signal rate against the background rate for the entire central chip. No visible excess is seen.

The best published helioscope limit on the axion-electron coupling is [21],

$$g_{ae} \cdot g_{a\gamma} < 8.1 \times 10^{-23} \text{ GeV}^{-1} \text{ at 95 \% CL.}$$

In contrast the observed limit computed with the presented setup in this work improves on this by 10 % to,

$$(g_{ae} \cdot g_{a\gamma})_{\text{observed}} < 7.35 \times 10^{-23} \text{ GeV}^{-1} \text{ at 95 \% CL.}$$

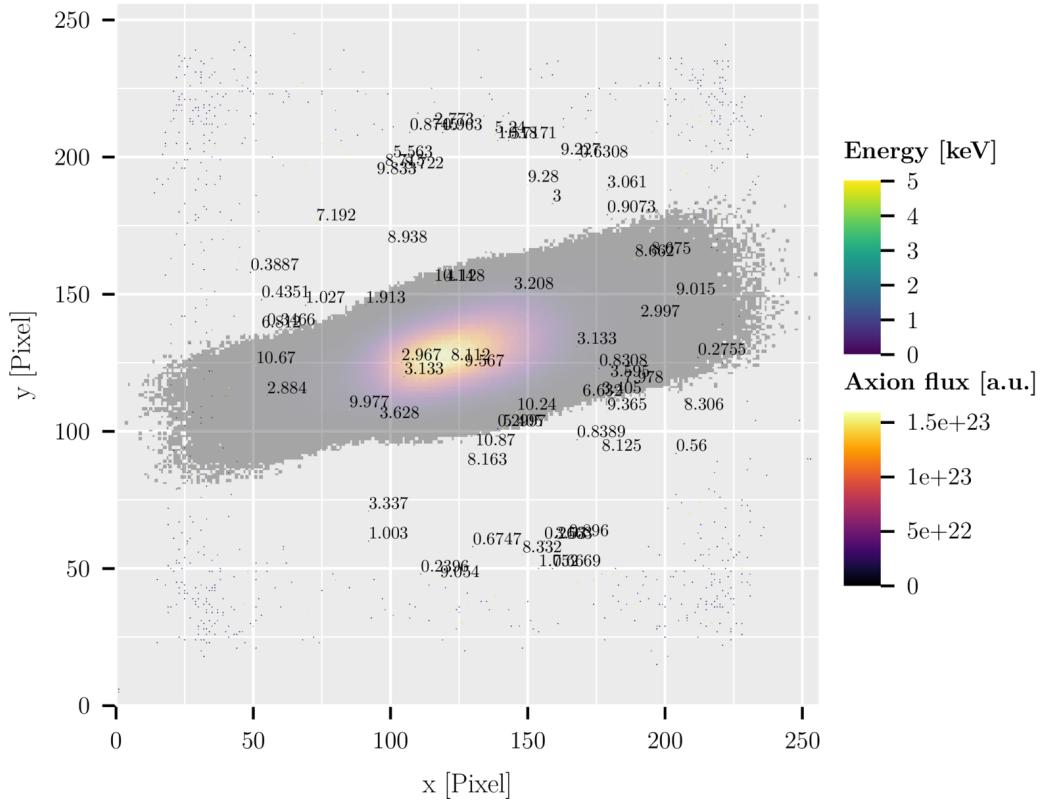


Figure 9. Position and energy of all candidates obtained by application of setup with best expected limit. No candidates in signal and energy sensitive region are found. The axion image from a raytracing calculation is shown in the background as a colored, partially transparent region. It corresponds to the image produced assuming a Sun-Earth distance of 0.989 AU at 0.3 cm behind the detector window. That means 1.2 cm in front of focal point of the X-ray telescope at the median conversion point for expected axion induced X-ray flux.

7 Summary and conclusion

Based on the 7-GridPix detector operated at CAST in 2017 and 2018, the limit on the axion-electron coupling could be improved to $(g_{ae} \cdot g_{a\gamma})_{\text{observed}} < 7.35 \times 10^{-23} \text{ GeV}^{-1}$ using approximately 160 h of solar tracking data. As the limit calculation method was designed to be generic under the signal source, a limit on the axion-photon coupling constant $g_{a\gamma}$ was also computed to

$$g_{a\gamma, \text{observed}} < 9.0 \times 10^{-11} \text{ GeV}^{-1} \text{ at 95\% CL.}$$

While this is significantly higher than the previous best limit from CAST [31] limit at $g_{a\gamma, \text{CAST}} < 6.6 \times 10^{-11} \text{ GeV}^{-1}$, this data has already been used for a new combined limit with Micromegas data from 2019–2021 [47] for the strongest helioscope limit on the axion-photon coupling at $g_{a\gamma, \text{CAST, 2024}} < 5.8 \times 10^{-11} \text{ GeV}^{-1}$.

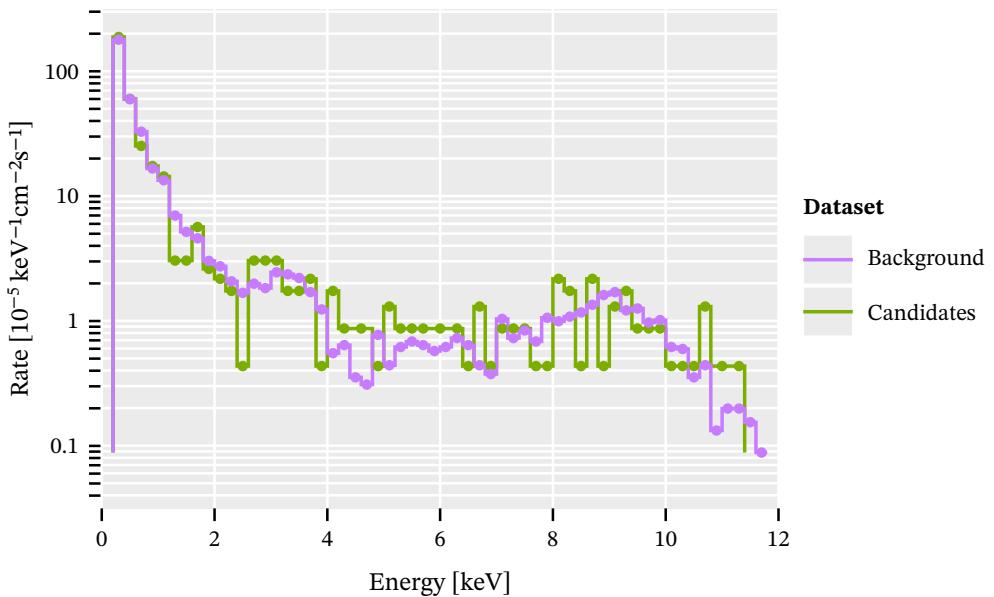


Figure 10. Comparison of the background and signal rate over the entire chip. No excess in signal visible compared to background model.

Data Availability Statement. This article has associated data in a data repository.

Code Availability Statement. This article has associated code in a code repository.

Open Access. This article is distributed under the terms of the Creative Commons Attribution License ([CC-BY4.0](#)), which permits any use, distribution and reproduction in any medium, provided the original author(s) and source are credited.

References

- [1] S. Weinberg, *The $U(1)$ problem*, *Phys. Rev. D* **11** (1975) 3583 [[INSPIRE](#)].
- [2] G. 't Hooft, *Symmetry breaking through Bell-Jackiw anomalies*, *Phys. Rev. Lett.* **37** (1976) 8 [[INSPIRE](#)].
- [3] G. 't Hooft, *How instantons solve the $U(1)$ problem*, *Phys. Rept.* **142** (1986) 357 [[INSPIRE](#)].
- [4] R.D. Peccei and H.R. Quinn, *CP conservation in the presence of instantons*, *Phys. Rev. Lett.* **38** (1977) 1440 [[INSPIRE](#)].
- [5] R.D. Peccei and H.R. Quinn, *Constraints imposed by CP conservation in the presence of instantons*, *Phys. Rev. D* **16** (1977) 1791 [[INSPIRE](#)].
- [6] S. Weinberg, *A new light boson?*, *Phys. Rev. Lett.* **40** (1978) 223 [[INSPIRE](#)].
- [7] F. Wilczek, *Problem of strong P and T invariance in the presence of instantons*, *Phys. Rev. Lett.* **40** (1978) 279 [[INSPIRE](#)].
- [8] K. Zioutas et al., *A decommissioned LHC model magnet as an axion telescope*, *Nucl. Instrum. Meth. A* **425** (1999) 480 [[astro-ph/9801176](#)] [[INSPIRE](#)].

[9] CAST collaboration, *First results from the CERN Axion Solar Telescope (CAST)*, *Phys. Rev. Lett.* **94** (2005) 121301 [[hep-ex/0411033](#)] [[INSPIRE](#)].

[10] I.G. Irastorza et al., *Towards a new generation axion helioscope*, *JCAP* **06** (2011) 013 [[arXiv:1103.5334](#)] [[INSPIRE](#)].

[11] J.K. Vogel et al., *IAXO — the International Axion Observatory*, in the proceedings of the *8th Patras workshop on axions, WIMPs and WISPs*, Chicago, U.S.A., 18–22 July 2012 [[arXiv:1302.3273](#)] [[INSPIRE](#)].

[12] E. Armengaud et al., *Conceptual design of the International Axion Observatory (IAXO)*, *2014 JINST* **9** T05002 [[arXiv:1401.3238](#)] [[INSPIRE](#)].

[13] A. Zhitnitsky, *Possible suppression of axion-hadron interactions*, *Sov. J. Nucl. Phys.* **31** (1980) 2.

[14] M. Dine, W. Fischler and M. Srednicki, *A simple solution to the strong CP problem with a harmless axion*, *Phys. Lett. B* **104** (1981) 199 [[INSPIRE](#)].

[15] J. Redondo, *Solar axion flux from the axion-electron coupling*, *JCAP* **12** (2013) 008 [[arXiv:1310.0823](#)] [[INSPIRE](#)].

[16] A. Berlin, A.J. Millar, T. Trickle and K. Zhou, *Physical signatures of fermion-coupled axion dark matter*, *JHEP* **05** (2024) 314 [[arXiv:2312.11601](#)] [[INSPIRE](#)].

[17] XENON collaboration, *Search for new physics in electronic recoil data from XENONnT*, *Phys. Rev. Lett.* **129** (2022) 161805 [[arXiv:2207.11330](#)] [[INSPIRE](#)].

[18] F. Capozzi and G. Raffelt, *Axion and neutrino bounds improved with new calibrations of the tip of the red-giant branch using geometric distance determinations*, *Phys. Rev. D* **102** (2020) 083007 [[arXiv:2007.03694](#)] [[INSPIRE](#)].

[19] O. Straniero et al., *The RGB tip of galactic globular clusters and the revision of the axion-electron coupling bound*, *Astron. Astrophys.* **644** (2020) A166 [[arXiv:2010.03833](#)] [[INSPIRE](#)].

[20] C. O’Hare, *cajohare/axionlimits: Axionlimits*, [Zenodo](#), July 2020.

[21] K. Barth et al., *CAST constraints on the axion-electron coupling*, *JCAP* **05** (2013) 010 [[arXiv:1302.6283](#)] [[INSPIRE](#)].

[22] X. Llopert et al., *Medipix2: a 64-k pixel readout chip with 55-/spl mu/m square elements working in single photon counting mode*, *IEEE Trans. Nucl. Sci.* **49** (2002) 2279.

[23] M. Campbell et al., *The detection of single electrons by means of a micromegas-covered MediPix2 pixel CMOS readout circuit*, *Nucl. Instrum. Meth. A* **540** (2005) 295 [[physics/0409048](#)] [[INSPIRE](#)].

[24] M. Chefdeville et al., *An electron-multiplying ‘micromegas’ grid made in silicon wafer post-processing technology*, *Nucl. Instrum. Meth. A* **556** (2006) 490 [[INSPIRE](#)].

[25] H. van der Graaf, *GridPix: an integrated readout system for gaseous detectors with a pixel chip as anode*, *Nucl. Instrum. Meth. A* **580** (2007) 1023 [[INSPIRE](#)].

[26] C. Krieger, J. Kaminski and K. Desch, *InGrid-based X-ray detector for low background searches*, *Nucl. Instrum. Meth. A* **729** (2013) 905 [[INSPIRE](#)].

[27] C. Krieger, J. Kaminski, M. Lupberger and K. Desch, *A GridPix-based X-ray detector for the CAST experiment*, *Nucl. Instrum. Meth. A* **867** (2017) 101 [[INSPIRE](#)].

[28] S.M. Schmidt, *Search for solar axions using a 7-GridPix IAXO prototype detector at CAST*, Ph.D. thesis, U. Bonn, Bonn, Germany (2024) [[INSPIRE](#)].

[29] R. Petre and P.J. Serlemitsos, *Conical imaging mirrors for high-speed X-ray telescopes*, *Appl. Opt.* **24** (1985) 1833.

[30] F. Aznar et al., *A Micromegas-based low-background X-ray detector coupled to a slumped-glass telescope for axion research*, *JCAP* **12** (2015) 008 [[arXiv:1509.06190](https://arxiv.org/abs/1509.06190)] [[INSPIRE](#)].

[31] CAST collaboration, *New CAST limit on the axion-photon interaction*, *Nature Phys.* **13** (2017) 584 [[arXiv:1705.02290](https://arxiv.org/abs/1705.02290)] [[INSPIRE](#)].

[32] S. Schmidt, *TrAXer — an interactive real-time X-ray raytracer*, <https://github.com/Vindaar/TrAXer> (2023).

[33] X. Llopert et al., *Timepix, a 65k programmable pixel readout chip for arrival time, energy and/or photon counting measurements*, *Nucl. Instrum. Meth. A* **581** (2007) 485 [[INSPIRE](#)].

[34] M. Lupberger, *The Pixel-TPC: a feasibility study*, Ph.D. thesis, Bonn U., Bonn, Germany (2015) [[INSPIRE](#)].

[35] *TOS — timepix operating software*, <https://github.com/Vindaar/TOS>.

[36] HDF group, *Hierarchical Data Format, version 5*, <https://www.hdfgroup.org/HDF5/>, (1997)–(2022).

[37] S.M. Schmidt, *TimepixAnalysis*, <https://github.com/Vindaar/TimepixAnalysis> (2022).

[38] G.D. Alkhazov, *Statistics of electron avalanches and ultimate resolution of proportional counters*, *Nucl. Instrum. Meth.* **89** (1970) 155 [[INSPIRE](#)].

[39] F. Rosenblatt, *The perceptron: a probabilistic model for information storage and organization in the brain*, *Psychol. Rev.* **65** (1958) 386.

[40] S. Amari, *A theory of adaptive pattern classifiers*, *IEEE Trans. Electron. Computers E* **C-16** (1967) 299.

[41] J. Schmidhuber, *Annotated history of modern AI and deep learning*, [arXiv:2212.11279](https://arxiv.org/abs/2212.11279).

[42] H. Cramér, *On the composition of elementary errors: first paper: mathematical deductions*, *Scand. Actuarial J.* **1928** (1928) 13.

[43] R. von Mises, *Wahrscheinlichkeit Statistik und Wahrheit: einführung in die Neue Wahrscheinlichkeitslehre und ihre Anwendung* (in German), Springer, Berlin, Heidelberg, Germany (1936) [[DOI:10.1007/978-3-662-41863-5](https://doi.org/10.1007/978-3-662-41863-5)].

[44] T.W. Anderson, *On the distribution of the two-sample Cramer-von Mises criterion*, *Annals Math. Statist.* **33** (1962) 1148.

[45] N. Metropolis et al., *Equation of state calculations by fast computing machines*, *J. Chem. Phys.* **21** (1953) 1087 [[INSPIRE](#)].

[46] W.K. Hastings, *Monte Carlo sampling methods using Markov chains and their applications*, *Biometrika* **57** (1970) 97 [[INSPIRE](#)].

[47] CAST collaboration, *New upper limit on the axion-photon coupling with an extended CAST run with a Xe-based micromegas detector*, *Phys. Rev. Lett.* **133** (2024) 221005 [[arXiv:2406.16840](https://arxiv.org/abs/2406.16840)] [[INSPIRE](#)].