Search for Axion-Like Particles produced in e^+e^- collisions and photon energy resolution studies at Belle II

Dissertation
zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Physik
der Universität Hamburg

vorgelegt von
Michael De Nuccio

Hamburg
2021

Gutachter/innen der Dissertation:	Dr. Torben Ferber Prof. Dr. Johannes Haller
Zusammensetzung der Prüfungskommission:	Dr. Torben Ferber Prof. Dr. Johannes Haller Prof. Dr. Jochen Liske Dr. Savino Longo Dr. Kai Schmidt-Hoberg
Vorsitzende/r der Prüfungskommission:	Prof. Dr. Jochen Liske
Datum der Disputation:	19.10.2021
Vorsitzender Fach-Promotionsausschusses PHYSIK:	Prof. Dr. Wolfgang Hansen
Leiter des Fachbereichs PHYSIK:	Prof. Dr. Günter H. W. Sigl
Dekan der Fakultät MIN:	Prof. Dr. Heinrich Graener

Abstract

Despite the great successes achieved by the Standard Model (SM) in explaining and predicting the behavior and existence of particles, multiple phenomena are yet to be given a satisfying explanation. Amongst these is Dark Matter (DM), a kind of matter that would permeate the whole Universe and that so far has been observed only via its gravitational interactions.

One possible extension of the SM, which may contribute to solve the mystery of DM and/or explain some astrophysical anomalies, are Axion-Like Particles (ALPs). The model taken into consideration in this thesis is of an ALP interacting with SM photons with a coupling strength $g_{a\gamma\gamma}$ and having mass m_a. This thesis describes a search for the direct production of such ALP via the process $e^+e^- \to \gamma a(a \to \gamma\gamma)$, in the mass range $0.2 < m_a < 9.7$ GeV/c^2. This search is performed using 0.445 fb^{-1} of data collected in 2018 by the Belle II detector.

No evidence for ALPs is found, and a 95%-confidence-level upper limit is set on the coupling constant $g_{a\gamma\gamma}$ at the level of 10^{-3} GeV^{-1}. These limits are the strongest to date for $0.2 < m_a < 1$ GeV/c^2.

Given that the final state of the $e^+e^- \to \gamma a(a \to \gamma\gamma)$ process is fully neutral, being made up by three photons, a proper kinematic fit with neutral particles may be a powerful tool to improve signal resolution. To achieve such a kinematic fit, a precise knowledge of the photon covariance matrix is needed. Such matrix is obtained from the results of photon resolution studies, whose status and results are presented in this thesis.

Trotz der großen Erfolge, welche das Standardmodell (SM) bei der Erklärung und Vorhersage des Verhaltens und der Existenz von Teilchen erzielt hat, fehlt für mehrere Phänomene noch eine zufriedenstellende Erklärung. Darunter ist Dunkle Materie (DM), eine Art Materie, die scheinbar das gesamte Universum durchdringt und die bislang nur durch ihre gravitative Wechselwirkung beobachtet wurde.

Eine mögliche Erweiterung des SM, die dazu beitragen könnte, das Mysterium der DM zu lösen und/oder einige astrophysikalische Anomalien zu erklären, sind Axion-ähnliche Teilchen (ALPs). Das in dieser Arbeit berücksichtigte Modell ist ein ALP, welches mit SM-Photonen mit einer Kopplungsstärke $g_{a\gamma\gamma}$ wechselwirkt, und eine Masse m_a besitzt. Diese Arbeit beschreibt eine Suche nach der direkten Produktion eines solchen ALPs im $e^+e^- \to \gamma a(a \to \gamma\gamma)$ Prozess im Massenbereich $0.2 < m_a < 9.7$ GeV/c^2. Diese Suche wird mit 0.445 fb^{-1} an Daten durchgeführt, welche 2018 vom Belle II Detektor gesammelt wurden.

Es wird kein Hinweis auf ALPs gefunden, und obere Ausschlussgrenzen von 10^{-3} GeV^{-1} werden mit einem Konfidenzniveau von 95% auf die Kopplungsstärke $g_{a\gamma\gamma}$ gesetzt. Diese Ausschlussgrenzen sind für $0.2 < m_a < 1$ GeV/c^2 die zurzeit stärksten.

Da der Endzustand des $e^+e^- \to \gamma a(a \to \gamma\gamma)$ Prozesses mit drei Photonen vollständig neutral ist, könnte ein geeigneter kinematischer Fit mit neutralen Teilchen ein wirksames Werkzeug zur Verbesserung der Signalauflösung sein. Um einen solchen kinematischen Fit zu ermöglichen, ist eine genaue Kenntnis der Kovarianzmatrix für Photonen erforderlich. Diese Matrix wird aus den Ergebnissen von Studien der Auflösung von Photonen gewonnen, deren Status und Ergebnisse in dieser Arbeit vorgestellt werden.

Disclaimer

The ALP search presented in this thesis was proposed to me by my supervisor, Dr. Torben Ferber. The analysis itself was designed by me and him. I performed the entirety of the analysis reported in this thesis, with the exception of the trigger efficiency (Sec. 5.4), timing (Sec. 6.1), and photon correction studies (Sec. 7.5.1), which were carried out by Dr. Ferber, who produced Figs. 5.13, 6.1, 6.2, 7.11, 7.12, 7.13, 7.19, and 7.20. All results are worded by me. The signal MC (Sec. 4.3.1) was produced by Dr. Sam Cunliffe, who produced Fig. 4.1. All results are worded by me. The background MC (Sec. 4.3.2) was produced by the Belle II collaboration, by Dr. Ferber, and by me. The Belle II collaboration, of which I am part since 2017, designed, built, and runs the Belle II experiment. The electron and positron beams are provided by the SuperKEKB collider. The data used for the ALP analysis were collected between April and July 2018. An article presenting the results of the ALP search, written by Dr. Ferber and me, has been published on Physical Review Letters on 14/10/2020 [1].

The photon resolution studies for building the photon covariance matrix, described in Chap. 8, were suggested to me by Dr. Ferber. I carried out the entirety of the studies reported in that Chapter, with the exception of the ISR modeling and the final convolutions fits (Sec. 8.3.3), which were performed by Dr. Savino Longo, who produced Figs. 8.7, 8.10, 8.11, 8.12, and 8.13 (the latter of which has been edited by me). All results are worded by me.

All pictures are produced by me, unless otherwise specified.

Contents

1 Introduction **13**

2 Physics motivation **15**
- 2.1 Standard Model overview . . . 15
- 2.2 Axions and axion-like particles . . . 17
- 2.3 The $e^+e^- \to \gamma a(a \to \gamma\gamma)$ process . . . 18
  - 2.3.1 Status of searches . . . 23
  - 2.3.2 The search at Belle II . . . 24

3 The Belle II experiment **27**
- 3.1 Brief history . . . 27
- 3.2 Purposes . . . 28
- 3.3 SuperKEKB . . . 29
- 3.4 The Belle II detector . . . 30
  - 3.4.1 CDC . . . 31
  - 3.4.2 ECL . . . 33
  - 3.4.3 Trigger . . . 37
  - 3.4.4 Software . . . 38
- 3.5 Beam background . . . 39

4 Analysis overview and data sample **41**
- 4.1 Analysis strategy outline . . . 41
- 4.2 Signal and background characteristics . . . 44
- 4.3 Data samples . . . 45
  - 4.3.1 Signal . . . 45
  - 4.3.2 Backgrounds . . . 46
  - 4.3.3 Cosmics . . . 47
  - 4.3.4 Data . . . 48

5 Event selection **49**
- 5.1 Optimization strategy . . . 49

5.2 Variables based on MC 51
5.2.1 Three most energetic photons 51
5.2.2 Angular acceptance 52
5.2.3 Photon energy 53
5.2.4 Invariant mass of the three photons 61
5.3 Variables based on data 64
5.3.1 Zernike of the most isolated photon 64
5.3.2 Time variable 64
5.3.3 Angular differences 65
5.3.4 Robustness variables 70
5.4 Trigger efficiency 70
5.5 Selection summary and performances 72

6 MC validation **77**
6.1 Timing from $e^+e^- \to \gamma(\gamma)$ and $e^+e^- \to e^+e^-\gamma$ events 77
6.2 Emulated signal selection with $e^+e^- \to e^+e^-\gamma$ 80
6.3 Signal sideband 81
6.4 Partial unblinding 85

7 Signal extraction and results **91**
7.1 Strategy 91
7.1.1 Negative significance check 93
7.1.2 Look-elsewhere effect 94
7.1.3 Diphoton to recoil mass transition 95
7.2 Signal PDF 98
7.2.1 Peaking component 98
7.2.2 Combinatorial component 100
7.3 Background PDF and fit range 104
7.4 Signal injection and bias studies 105
7.5 Systematic uncertainties 109
7.5.1 Mass resolution 109
7.5.2 Signal efficiency 112
7.5.3 Background shape and fit range 117
7.5.4 Other systematics and cross checks 117
7.6 Unblinding 118
7.6.1 Data/MC comparison 118
7.6.2 Upper limits extraction 126

8 Photon covariance matrix **131**
8.1 Motivations 131
8.1.1 Kinematic fit with photons 132

8.1.2 Mass resolution improvement 133
8.2 Single photon studies 138
8.3 Radiative dimuon studies 139
8.3.1 Radiative dimuon selection 141
8.3.2 Factor out the smearing effects 142
8.3.3 Results 149
8.4 Prospects 155

9 Summary and outlook 157

A Limits and systematics 161

Bibliography 167

Ringraziamenti / Acknowledgements 175

Chapter 1

Introduction

Particle physics is very well described by the Standard Model (SM). This theory has proved itself capable of outstandingly precise predictions, nonetheless, to date, it still presents some issues and lacks the capability of describing some aspects of reality. Amongst the most relevant open problems, there is the fact that gravity is not described by the theory; the neutrinos' masses are unexplained; the nature of Dark Matter (DM) has yet to be understood. DM is a hypothesized type of matter which seems to interact only gravitationally with standard matter, i.e. the one successfully described by the SM. Many cosmological evidences indeed suggest the existence of a kind of matter which does not interact electromagnetically, weakly, or strongly, but only gravitationally, and that would constitute the large majority of the matter present in the Universe.

Throughout the decades, many experiments have been performed to search for DM [2, 3, 4], following a variety of theoretical models, all resulting in no convincing evidence. The searches extend to *dark mediator*s too [5, 6], i.e. particles that would connect the SM with the DM particles. The hope is that actually DM or a dark mediator does indeed interact with SM matter, although very feebly, via any interaction other than the gravitational one.

DM candidates can be searched with a variety of experimental techniques. The analysis presented in this thesis, a search for Axion-Like Particles (ALPs), is a *collider search*, given that it is performed at Belle II, which is located on an electron-positron collider that runs at 10.58 GeV. The searched channel is $e^+e^- \to \gamma a(a \to \gamma\gamma)$, a being the ALP. In this case what would be produced is not directly a DM particle, but rather a DM-SM dark mediator. These are particles that interact feebly with SM particles and may have non-negligible couplings with DM particles. In the process analyzed in this thesis, the dark mediator ALP would couple with two SM photons: it would be produced, in association with a recoil photon, by the decay of the virtual photon deriving from the annihilation of the e^+e^- pair, and it would decay into a $\gamma\gamma$ pair. The final state are three resolved, high-energetic photons, and no missing energy. The searched ALP is not the so-called QCD axion, i.e. the particle hypothesized to solve the strong CP problem.

This thesis also includes a section about the status of the photon energy resolution studies. This is a necessary component to perform proper *photon kinematic fit*, which in turns could be a powerful tool to improve signal mass resolution and hence tighten the exclusion limits which can be set on the ALP process $e^+e^- \to \gamma a(a \to \gamma\gamma)$.

Chapter 2 presents a brief physics overview of the theoretical settings and the characteristics of the ALP channel $e^+e^- \to \gamma a(a \to \gamma\gamma)$.

Chapter 3 describes the Belle II experiment.

Chapter 4 provides a detailed overview of the analysis workflow and the details of the used data samples.

Chapter 5 describes the procedures used to determine the optimal event selection, and shows its performances.

Chapter 6 shows the results of various MC validation studies.

Chapter 7 describes how the signal is extracted, how the Upper Limits (ULs) are set, how the signal and backgrounds are modeled, how the systematic uncertainties are treated, and finally the results obtained.

Chapter 8 describes the potential usefulness of a kinematic fit. To perform kinematic fit with neutral particles, a properly built photon covariance matrix is needed, which is built starting from photon resolution studies results. The status and the prospects of these studies are reported in the Chapter.

Chapter 9 summarizes all the procedures and results of the analyses presented in this thesis, and provides outlooks for future improvements.

Chapter 2

Physics motivation

This Chapter provides an overview on the physics motivation leading to the ALP analysis presented in this thesis.

Section 2.1 summarizes the main features of the SM.

Section 2.2 explains the reasons why axions were hypothesized, their relationship with ALPs, and the characteristics of both.

Section 2.3 examines in details the process searched by the analysis presented in this thesis, focusing also on the current status of searches and on the features particularly exploitable in the Belle II experiment.

2.1 Standard Model overview

The SM is based on fundamental particles and their interactions (strong, weak, and electromagnetic). Particles' spin, in units of $\hbar$, can be integer (*bosons*), or half-integer (*fermions*). Each particle has its own anti-particle, that in some cases coincides with the particle itself. Fig. 2.1 shows the classification of the SM fundamental particles.

Fermions are classified in two categories: *quarks* take part in all the interactions, while *leptons* do not interact strongly. Both leptons and quarks are grouped into three doublets, called *generations*.

Each of the six quarks (u, d, s, c, b, t) comes in three different *colors*, that represent the strong charge. Particles that do not interact strongly, like the leptons, are *colorless*. The strong interaction is described by *quantum chromodynamics* (QCD). In each generation, one quark has an electric charge of $+\frac{2}{3}e$, while the other has $-\frac{1}{3}e$, e being the elementary charge.

The six leptons (e, ν_e, μ, ν_μ, τ, ν_τ), are colorless, so they can only interact

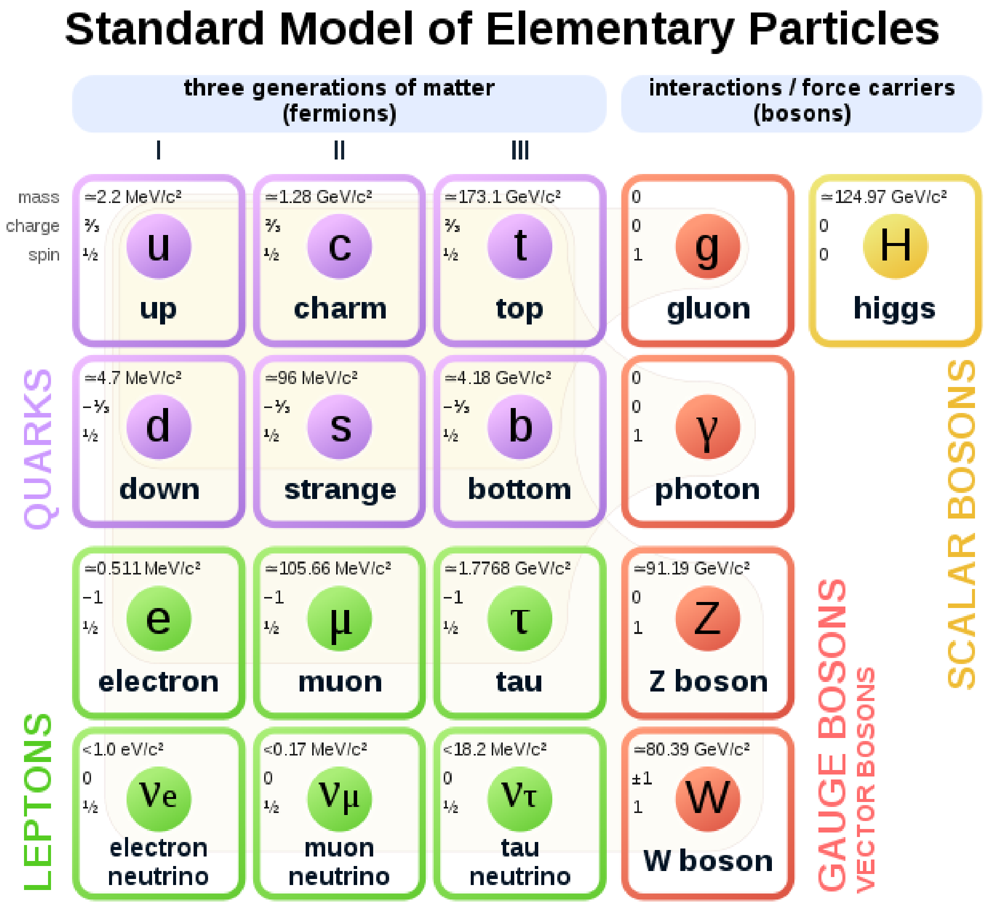

Standard Model of Elementary Particles

	three generations of matter (fermions)			interactions / force carriers (bosons)	
	I	II	III		
mass charge spin	≈2.2 MeV/c^2 ⅔ ½ u up	≈1.28 GeV/c^2 ⅔ ½ c charm	≈173.1 GeV/c^2 ⅔ ½ t top	0 0 1 g gluon	≈124.97 GeV/c^2 0 0 H higgs
QUARKS	≈4.7 MeV/c^2 −⅓ ½ d down	≈96 MeV/c^2 −⅓ ½ s strange	≈4.18 GeV/c^2 −⅓ ½ b bottom	0 0 1 γ photon	
	≈0.511 MeV/c^2 −1 ½ e electron	≈105.66 MeV/c^2 −1 ½ μ muon	≈1.7768 GeV/c^2 −1 ½ τ tau	≈91.19 GeV/c^2 0 1 Z Z boson	
LEPTONS	<1.0 eV/c^2 0 ½ ν_e electron neutrino	<0.17 MeV/c^2 0 ½ ν_μ muon neutrino	<18.2 MeV/c^2 0 ½ ν_τ tau neutrino	≈80.39 GeV/c^2 ±1 1 W W boson	

Figure 2.1: Fundamental particles included in SM. For each particle its mass, electric charge, and spin are reported. Fermions are in the first three columns: quarks on the top, leptons on the bottom. The force carriers are in the fourth column. The Higgs Boson, responsible for mass acquisition, is in the fifth column. Quarks are subject to all the three forces; charged leptons are subject only to electromagnetic and weak forces; neutrinos are subject uniquely to the weak force. Image from [7].

weakly and electromagnetically. In each of the three generations, one lepton has unitary electric charge, and the other is neutral. The neutral leptons, known as *neutrinos* (ν), can only interact weakly.

Each of the three fundamental forces described by the SM is associated with a *force carrier*. All of these carriers are vector (spin 1) bosons. The gluon g is the mediator of the strong interaction, and is the only colored boson; the photon γ is the mediator of the electromagnetic force; the Z

and the $W^{\pm}$ are the mediators of the weak interaction. The gluon and the photon are massless, while the Z and the W bosons are massive. The only charged boson is the W. The *Higgs boson* H is the scalar (spin 0), massive boson which is responsible for the acquisition of mass by particles through the mechanism of spontaneous symmetry breaking.

2.2 Axions and axion-like particles

While the SM has repeatedly been shown to be a very accurate and predictive theory, it still leaves a number of open questions. One of these is the so-called *strong CP problem* [8, 9, 10].

The operators C, P, and T represent, respectively, charge conjugation, inversion of spatial coordinates, and application of time-reversal. The CPT theorem states that CPT-symmetry is an exact symmetry, meaning that all processes are invariant under a CPT transformation. The C, P, and T symmetries can be individually violated, for example by the weak interaction. The CP symmetry is not preserved by the weak interaction. The first experiment showing this was carried out by Cronin and Fitch in 1964 and examined the decays of neutral kaons [11]. The predicted CP violation (CPV) in the B meson decays was first observed in 2001 by the first generation of B factories [12], see Sec. 3.1.

The CP symmetry can in principle be violated by the strong interaction, but such violation has never been observed. The QCD Langrangian [13, 14] contains terms which allow CPV. These can be parametrized with a parameter $\bar{\theta}$, that can be interpreted as an angle; as such, its value could be any number in the range $[0, 2\pi)$, a priori with equal probability.

The absence of experimental evidences of CPV by the strong interaction implies that the CP-violating parameter $\bar{\theta}$ of the QCD Lagrangian must be zero or very close to zero. The strongest limits, $|\bar{\theta}| < 10^{-10}$, come from the limits on the electric dipole moment of the neutron [15]. This is a so-called *fine-tuning* problem, i.e. a situation where a free parameter must be finely tuned to maintain consistency with the experiments. When such an event occurs, it often implies the presence of a still-unknown underlying physical phenomenon that constrains said parameter.

The most well-known solution to the strong CP problem was proposed by Peccei and Quinn [9, 10]. They promoted the parameter $\bar{\theta}$ to a dynamic field, by adding a new symmetry that is spontaneously broken, resulting in a new particle which they called *axion*. This particle, sitting at the bottom of its field potential, would naturally make the CP-violating parameter null, without requiring fine-tuning.

In this model, the coupling constant f_A for the axion-photon-photon interaction vertex and the mass m_A of the axion satisfy the following relation:

$$f_A m_A = f_\pi m_\pi, \tag{2.1}$$

with $f_\pi = 92$ GeV and $m_\pi = 135$ MeV (with c=1).

The physical mechanism that would allow the existence of the axion, i.e. the spontaneous breaking of an approximate symmetry, can also generate ALPs. Unlike axions, which are linked to the strong interactions and whose mass and coupling are determined by a single new parameter f_A, ALPs are less constrained, and their masses and couplings to photons are independent parameters. ALPs are pseudoscalar neutral massive particles that couple predominantly to gauge bosons, hence they can couple to pairs of bosons as gg, $\gamma\gamma$, ZZ, γZ, W^+W^-, etc., the specifics depending on the model.

Many models have been developed which contain ALPs. As any hypothetical, feebly interacting particle, they are candidates for being DM particles or DM mediators. Appropriately tuned ALPs could help solving astrophysical puzzles such as the anomalies in the energy loss of white dwarf stars [16] and the transparency of the Universe to TeV gamma rays [17].

ALPs in the sub-MeV mass range may affect cosmology and astrophysics, influencing for example the Big Bang Nucleosynthesis, the Cosmic Microwave Background, and star evolution [18, 19].

ALPs with masses in the MeV-GeV range are mostly irrelevant for astrophysic and cosmology, but they can become relevant for particle physics [20, 21]: for example, it has been suggested that ALPs may contribute to explain the anomalous magnetic moment of the muon [22, 23] or the exotic resonances in nuclear transitions [24].

There exist ALP models that are extensions of the original Peccei-Quinn model, like the *axiflavon* [25]. In this model the symmetry breaking also explains the *flavor puzzle*, i.e. why the fermion masses are hierarchical.

2.3 The $e^+e^- \to \gamma a(a \to \gamma\gamma)$ process

The search described in this thesis takes a benchmark model from [26]. The pseudoscalar ALP is labeled as a and interacts with SM gauge bosons. The Lagrangian to start from is the following:

$$\mathscr{L} = \frac{1}{2}\left[\partial^\mu a \partial_\mu a - m_a^2 a^2 - \frac{c_B}{2f_a} a B^{\mu\nu}\bar{B}_{\mu\nu} - \frac{c_W}{2f_a} a W^{i,\mu\nu}\bar{W}^i_{\mu\nu}\right] \tag{2.2}$$

where $B^{\mu\nu}$ and $W^{i,\mu\nu}$ represent the field strength of $U(1)_Y$ and $SU(2)_L$ respectively; the dual field strength tensors are defined as $\bar{B}_{\mu\nu} = \frac{1}{2}\epsilon_{\mu\nu\alpha\beta}B^{\alpha\beta}$; m_a and f_a represent the ALP mass and its decay constant, and given that the search is not about the QCD axion, for which Eq. 2.1 holds true, the two parameters are independent; the parameters c_B and c_W represent the coupling constant of the ALP to the fields $B^{\mu\nu}$ and $W^{i,\mu\nu}$ respectively.

The last two terms can be rewritten, after electroweak symmetry breaking, as it follows:

$$\mathscr{L} \supset -\frac{a}{4}\left[g_{a\gamma\gamma}F^{\mu\nu}\bar{F}_{\mu\nu} + g_{a\gamma Z}F^{\mu\nu}\bar{Z}_{\mu\nu} + g_{aZZ}Z^{\mu\nu}\bar{Z}_{\mu\nu} + g_{aWW}W^{\mu\nu}\bar{W}_{\mu\nu}\right], \tag{2.3}$$

while the kinematic and the mass terms, irrelevant for the following discussion, are omitted.

The Lagrangian above does not include all possible terms that could be present in a more general effective field theory description of ALPs. In particular, the interactions of a with fermions and/or gluons are neglected, because these generally lead to flavor-changing processes, which are strongly constrained by searches for rare decays [6]. Additional interactions generally increase the ALP production cross section, leading to stronger limits than the ones which would derive from considering only the Lagrangian above. Although, in some particular situations additional interactions may conversely weaken the limits: for example, additional decay modes could decrease the ALP lifetime, therefore weakening limits coming from analyses modeled to search for long-living ALPs. Similarly, additional interactions may reduce the flux of ALPs reaching the Earth from astrophysical objects, via interaction with the interstellar medium.

The last two components of Eq. 2.3 can be ignored, as SuperKEKB (see Sec. 3) operates at energies of the order of 10 GeV, much lower than the ones necessary to produce the $W^\pm$ and Z bosons. The couplings of interest are therefore:

$$g_{a\gamma\gamma} = \frac{c_B\cos^2\theta_W + c_W\sin^2\theta_W}{f_a} \tag{2.4}$$

and

$$g_{a\gamma Z} = \frac{\sin 2\theta_W(c_W - c_B)}{f_a}, \tag{2.5}$$

where θ_W is the Weinberg angle. Two cases of particular interest can be considered. If $c_W \sim c_B$, then $g_{a\gamma Z} \ll g_{a\gamma\gamma}$: this case is called *photon coupling*. If $c_W \ll c_B$, then $g_{a\gamma\gamma} \approx \frac{1}{2}\cot\theta_W g_{a\gamma Z} \approx -0.94 g_{a\gamma Z}$: this case is called *hypercharge coupling*.

The same low-energy argument as above can be applied again, hence the *hypercharge coupling* can be ignored, and only the *photon coupling* case will be considered.

ALPs for which the *photon coupling* is the relevant coupling to SM particles can be produced with two different processes in an e^+e^- experiment like Belle II (Chap. 3): *ALP-strahlung* ($e^+e^- \to \gamma^* \to \gamma a$) and *photon fusion* ($e^+e^- \to e^+e^-a$), both depicted in Fig. 2.2.

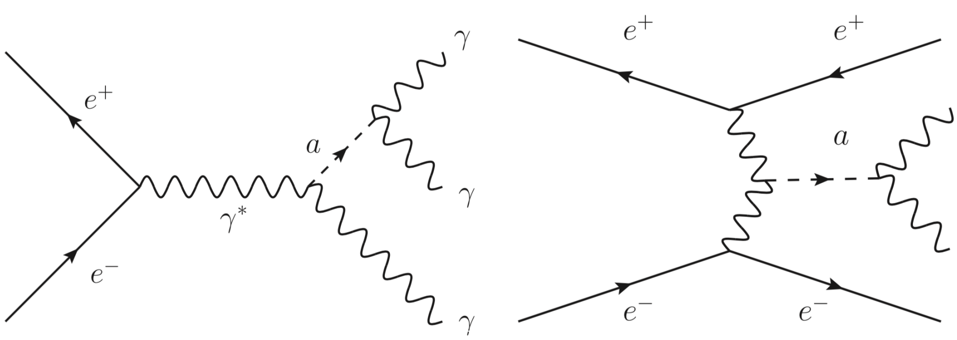

Figure 2.2: The two main processes for production of ALP with photon coupling, at e^+e^- colliders. Left: ALP-strahlung, right: photon fusion. In both cases the ALP decays into two photons. Image from [26].

ALPs produced via photon fusion generally carry little energy, making them hard to detect. The two final-state electrons also tend to be produced close to the beams direction (polar angle θ close to 0 or π), and thus out of the detector acceptance. These two factors are the reasons why the first analysis at Belle II considers only the ALP-strahlung process.

The differential production cross section for this process is given by [23]:

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\cos\theta} = \frac{g_{a\gamma\gamma}^2\alpha}{128}(3+\cos 2\theta)\left(1-\frac{m_a^2}{s}\right)^3, \tag{2.6}$$

where θ is the angle between the beams direction and the recoil photon in the center-of-mass system (cms), α is the fine-structure constant ($\alpha \approx 1/137$), and $\sqrt{s}$ is the collision energy in the cms ($\sqrt{s} \approx 10.58$ GeV for Belle II).

Integrated over all the polar angle, the expression becomes the following:

$$\sigma_{e^+e^- \to a\gamma} = \frac{g_{a\gamma\gamma}^2 \alpha}{24} \left(1 - \frac{m_a^2}{s}\right)^3 . \tag{2.7}$$

In the ALP-strahlung process, the ALP is produced in association with the aforementioned recoil photon. This is monoenergetic in the cms, with energy given by:

$$E^{CM}(\gamma_{recoil}) = \frac{s - m_a^2}{2\sqrt{s}}. \tag{2.8}$$

Given that only the photon-coupling-only case is considered, in this model the ALP can decay only in two photons, therefore its lifetime τ_a is given only by the $g_{a\gamma\gamma}$ coupling. In particular, the following relation holds true [27, 28]:

$$\Gamma_a = \tau_a^{-1} = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi}, \tag{2.9}$$

where Γ_a is the ALP decay width, which gets smaller for light and/or feebly interacting ALPs.

For a detector of length L, the fraction F_a of ALPs that decay within the detector itself can be computed as:

$$F_a = 1 - \exp\left(-\frac{L}{\gamma_a \beta_a c \tau_a}\right) \equiv 1 - \exp\left(-\frac{L}{L_a}\right) \tag{2.10}$$

where γ_a and $\beta_a c$ are the Lorentz boost and the speed of the ALP, and $L_a \equiv \gamma_a \beta_a c \tau_a$ is the characteristic ALP decay length in the laboratory reference frame.

Depending on the mass and lifetime of the ALP (so, ultimately on mass and coupling, given that the lifetime is given by mass and coupling), different final state configurations are possible. These are explained in the following, and schematically summarized in Fig. 2.3:

- if the ALP is heavy, then it is not particularly boosted and the two photons coming from its decay are more energetic, thus they end up far away from each other when hitting the detector: this the *resolved* configuration, where the three photons (the recoil one and the two from the ALP decay) are separate from each other, high energetic, and produced within the detector;

- if the ALP is light, then its boost is relevant and the two photons coming from its decay are less energetic, thus they tend to be collimated, i.e. close to each other, hence they tend to merge and be seen by the detector as one single cluster: this is the *merged* configuration, where *de facto* two photons are seen (the recoil one, and the fusion of the two coming from the ALP decay);

- if the ALP is long lived, it decays outside of the detector volume, so the two decay photons are not detectable, thus only the recoil photon can be detected: this is the *invisible* configuration;
- if the ALP's lifetime is long enough to not decay promptly, but not long enough to decay outside of the detector volume, then the two photons coming from its decay are not originating from the Interaction Point (IP): this is the *displaced* configuration, where the recoil photon originates from the IP and the two photons coming from the ALP originate from a displaced vertex.

Each of these configurations requires a specific analysis, with different selection techniques. The search presented in this thesis is about the *resolved* configuration, where three high-energy resolved photons are present in the event.

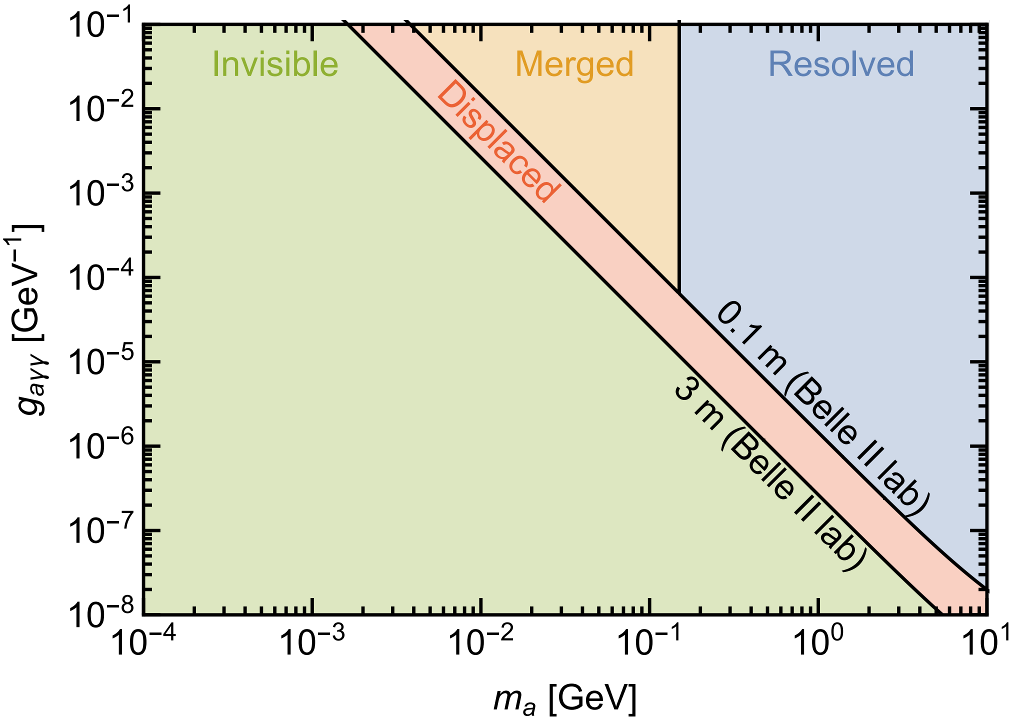

Figure 2.3: Different final-state configurations in the coupling-mass ALP parameters space. Image from [26].

2.3.1 Status of searches

In the following, limits on ALPs with photon coupling, pre-existing to the Belle II search described later in this thesis, are briefly summarized. A more detailed descriptions can be found in [26] and in the papers relative to each experiment or class of experiments.

- From e^+e^- colliders:
  - mono-photon searches at LEP [20] and recast of [29] by following [30]: single photon and missing energy, equivalent to the ALP invisible configuration;
  - radiative $\varUpsilon$ decay [31]: reinterpretations of radiative decays of $\varUpsilon(nS)$, $n = 1, 2, 3$ as $\varUpsilon(nS) \to \gamma^* \to \gamma a$. Equivalent to an ALP decaying invisibly. Recast from a BaBar search for $\varUpsilon(3S) \to \gamma A^0$;
  - Dark photon searches at BaBar [32]: similar to the above search, but with e^+e^- colliding at the $\varUpsilon(4S)$ energy, final state being a photon and missing energy. Recast performed by scaling the geometrical acceptance for the dark photon final state to the one for the ALP final state, and by converting the kinetic mixing to $g_{a\gamma\gamma}$;
  - recast of multi-photon searches at LEP-II, OPAL experiment [33]: searches for $\gamma\gamma$ resonances in $e^+e^- \to \gamma\gamma(\gamma)$ processes.
- From beam dump experiments: they all exploit Primakoff production, i.e. the conversion of a photon into an ALP when the photon interacts with a nucleus' electric field:
  - SLAC E141 [27]: electron beam dump, rescaling for different ALP masses from the original results;
  - SLAC E137 [34]: electron beam dump, a dedicated search for ALPs coupling only to photons. Rescaling from the original paper to take into account the exponential suppression of the number of ALPs reaching the detector;
  - CHARM and NuCal [35, 36, 28]: proton beam dump experiments;
  - NA64 [37]: electron-beam-on-target experiment, direct search for ALPs.
- From ultra-peripheral heavy-ion collisions from CMS [38]: ALPs would be produced via the fusion of two coherently emitted photons (similarly to the photon-fusion process shown in the right side of Fig. 2.2). The

photons emitted from the decay of the ALPs are very low energetic and back-to-back. The limits are for $m_a \gtrsim 5$ GeV/c^2;

- PrimEx [39]: photon-beam-on-target experiment to measure π^0 decay width reinterpreted in terms of exclusions limits for the presence of ALPs.

All of these limits are summarized in Fig. 2.4.

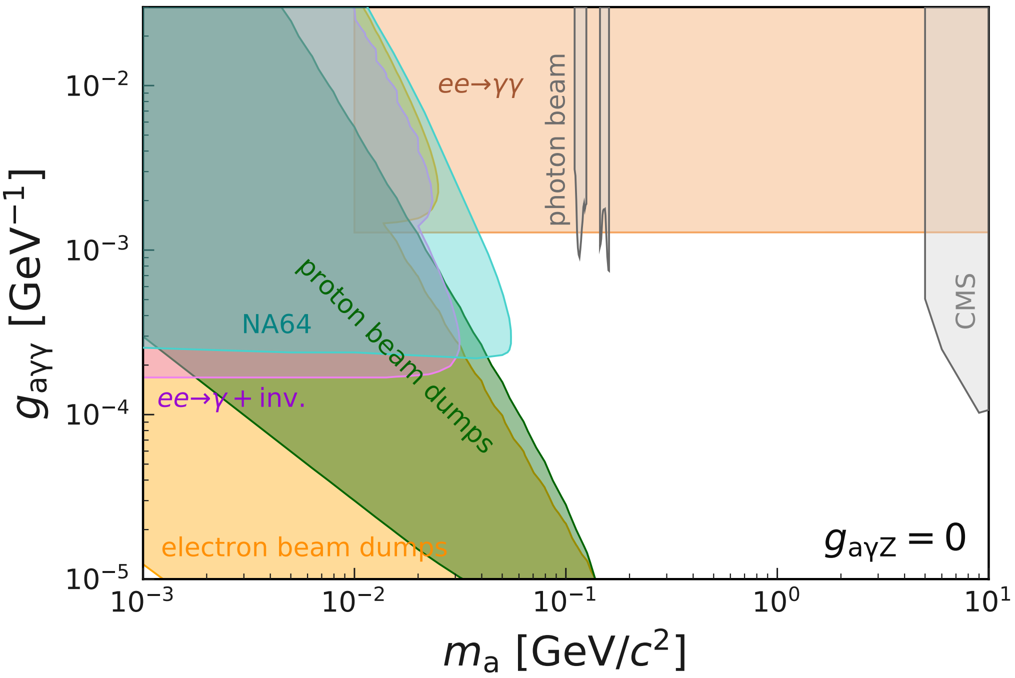

Figure 2.4: Summary of existing limits on ALPs with photon coupling, before the ones obtained from this analysis. The sources of limits are described in the text. Adapted from [1].

2.3.2 The search at Belle II

This Section provides a brief recap of the main experimental characteristics of the ALP analysis presented in this thesis. Attention is drawn also to its features and constraints deriving from working with Belle II, on an e^+e^- collider. The whole detailed analysis strategy is described in Sec. 4.1.

The search presented in this thesis is for ALPs, produced in association with a high-energy recoil photon, and decaying into two high-energetic, resolved photons. One of them is the recoil photon, see Eq. 2.8. The other two come from the decay of the ALP. This is called *resolved* configuration (see Fig. 2.3).

Belle II runs on an asymmetric (boost factor $\beta\gamma \approx 0.28$) e^+e^- collider (more in Sec. 3). This implies that the initial-state conditions are well known, differently from hadronic collider. This, together with the presence of exactly three photons in the final state, allows for the computation of the ALP mass in two different ways. It can be computed directly from the two photons constituting the ALP candidate, i.e. the so-called *diphoton mass*:

$$m_{\gamma\gamma} \equiv m\,(\gamma\gamma) \equiv m_{dip} \equiv m_a^{diphoton} = ||\boldsymbol{P}_1 + \boldsymbol{P}_2||, \tag{2.11}$$

where $\boldsymbol{P}_{1,2}$ are the four-momenta of the two photons. It can also be computed with the energy of the recoil photon in the cms and $\sqrt{s}$, i.e. the so-called *recoil mass*:

$$m^2_{recoil} \equiv m^2(recoil) = s - 2\sqrt{s}E^{CM}(\gamma_{recoil}). \tag{2.12}$$

The idea is that the two of them can provide different ALP mass resolution in different mass regimes: the diphoton mass has a better resolution for low ALP masses, and vice versa for the recoil mass.

The dataset used to performed this search is the one collected in the so-called 2018 *Phase 2*, a preliminary data-taking period mainly appointed for calibration and tuning purposes, and corresponds to an total integrated luminosity of 497 ± 3 pb^{-1}.

Chapter 3

The Belle II experiment

This Chapter provides a description of the Belle II experiment and its accelerator SuperKEKB, together with a brief general introduction on B factories.

Section 3.1 shortly summarizes the historical reasons behind the construction of the first B factories.

Section 3.2 explains the past and present main Physics goals of the B factories, and what is needed to achieve them.

Section 3.3 provides an overview on the characteristics of the SuperKEKB accelerator, which supplies the beams to the Belle II detector, together with a description of the beams' characteristics and an overview on the needs of the detectors.

Section 3.4 describes the Belle II detector.

Section 3.5 lists the sources of beam background that affect the Belle II experiment.

3.1 Brief history

Before the experiment by Cronin and Fitch in 1964 [11], it was known that the Universe is not symmetrical under C or P, but there were no evidences for violations of the composed operator CP. This experiment provided the first evidence for CPV.

In 1973, Kobayashi and Maskawa demonstrated that CPV could be included in the SM assuming the existence of six different quarks, because of an irreducible complex phase in the weak interaction quark-flavor-mixing matrix, called Cabibbo-Kobayashi-Maskawa (CKM) matrix [40]. At that time, though, only the three lightest quarks u, d, and s were known. By 1980, after the discoveries of the c and the b quarks, the CKM matrix was included in the SM, and predictions about the magnitude of the CPV in charm and beauty sectors were performed.

The goal of measuring the CPV in the $B\bar{B}$ system could only be achieved with a very large amount of B mesons, and was because of this that the first B factories were conceived and the first related experiments, *BABAR* (1999-2008) and Belle (1999-2010) [12], were built. Their successor is Belle II, which started taking data in 2018.

This kind of experiments require a high luminosity, $\sim 10^{33}$ cm^{-2} s^{-1}, to produce the required huge amount of B mesons, which is the reason why these machines are called *B factories*.

3.2 Purposes

One of the main goals of the experiments located at B factories was and is to study B physics, in particular CPV in the B mesons.

B mesons can be produced both in e^+e^- and in hadronic interactions. The production via hadronic interactions at high energy, for instance in $p\bar{p}$ collisions at 13 TeV at LHC, offers a large cross section and a large spectrum for beauty mesons and baryons. On the other hand, electron-positron colliders provide a clean environment; a well-known initial state, thus the possibility to set up kinematical constraints and exploit the quantum mechanical correlation between the two B mesons; a high signal-over-background ratio; the possibility to adopt relatively simple trigger strategies; an almost complete angular acceptance, $\sim 260°$ in θ and $360°$ in ϕ; a relatively small number of tracks per event, as shown in Fig. 3.1, which reduces the combinatorial background and often allows a full event reconstruction.

To optimize the $B\bar{B}$ production, the B factories run at a center-of-mass (CM) energy of 10.58 GeV, corresponding to the invariant mass of the $\Upsilon(4S)$ resonance: this is a bound state of quark b and antiquark $\bar{b}$ (*bottomonium*). This resonance is of particular interest because it is the first bottomonium state whose mass is high enough to allow it to decay into a B mesons pair: indeed $m_{B^\pm} = 5.279$ GeV/c^2, $m_{B^0} = 5.280$ GeV/c^2, and $m\left(\Upsilon(4S)\right) = \sqrt{s}/2 = 5.29$ GeV/c^2. Thus, each of the two B mesons carries about 10 MeV of kinetic energy in the cms, making it almost at rest in that reference frame.

B mesons have a short lifetime ($\tau_{B^\pm} = 1.638 \times 10^{-12}$ s, $\tau_{B^0} = 1.519 \times 10^{-12}$ s), hence if the CM and the laboratory frame are the same, the relative displacement is of the order of 60 μm. This small vertex distance is barely distinguishable, therefore *asymmetric* accelerators (where electron and positron beams have different energies) are used, so that the resulting products are boosted in the laboratory frame and the relative displacement is enhanced.

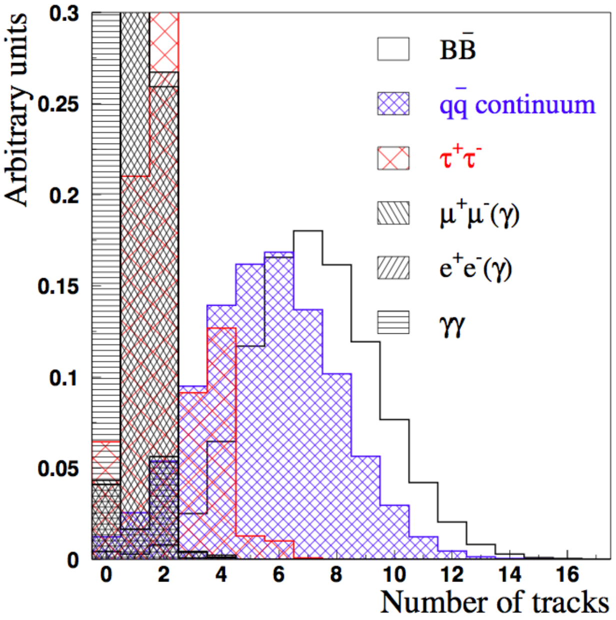

Figure 3.1: The number of charged tracks per event, for different processes. Image from [12].

This way, each of the two B mesons can be reconstructed independently.

Beside the study of B physics, these experiments also allow the study of the charm and tau physics (the cross sections for these processes are comparable with the one for $b\bar{b}$), as well as two-photon physics (e.g. [41]).

Searches for dark matter (also known as *dark sector*) candidates are one of the main physics goals of Belle II. The first two Belle II physics publications concern dark matter searches: [42] and [1], the latter being the analysis described in this thesis.

3.3 SuperKEKB

To provide beam collisions to the detector, an accelerator and a storage ring are required: *BABAR* was located on PEP-II and Belle on KEKB; Belle II is located on SuperKEKB [43], an upgrade of the pre-existing KEKB. A schematic is in Fig. 3.2. The accelerator is designed to provide world-highest instantaneous luminosity, $\sim 8 \cdot 10^{35}$ cm^{-2} s^{-1} at the maximum planned performance. This luminosity is $\sim$ 40 times higher than that of KEKB, and this is achieved by increasing the beam currents (factor 2) and squeezing the beams (factor 20), via an improved final-focus magnet system [44].

The electron beam runs at 7 GeV, the positron one at 4 GeV. This translates into a boost of the CM frame with respect to the laboratory equal to $\beta\gamma = 0.28$. Asymmetric beams need a dedicated storage ring for each of them. One single interaction region is present.

SuperKEKB is planned to mainly run at the $\Upsilon(4S)$ energy (as PEP-II and KEKB did), collecting so-called *on-resonance* events. A small but non-negligible fraction of the data will be taken at different energies, centered at other Υ resonances, and at other energies, also referred to as *off-resonance*, to collect *continuum* events.

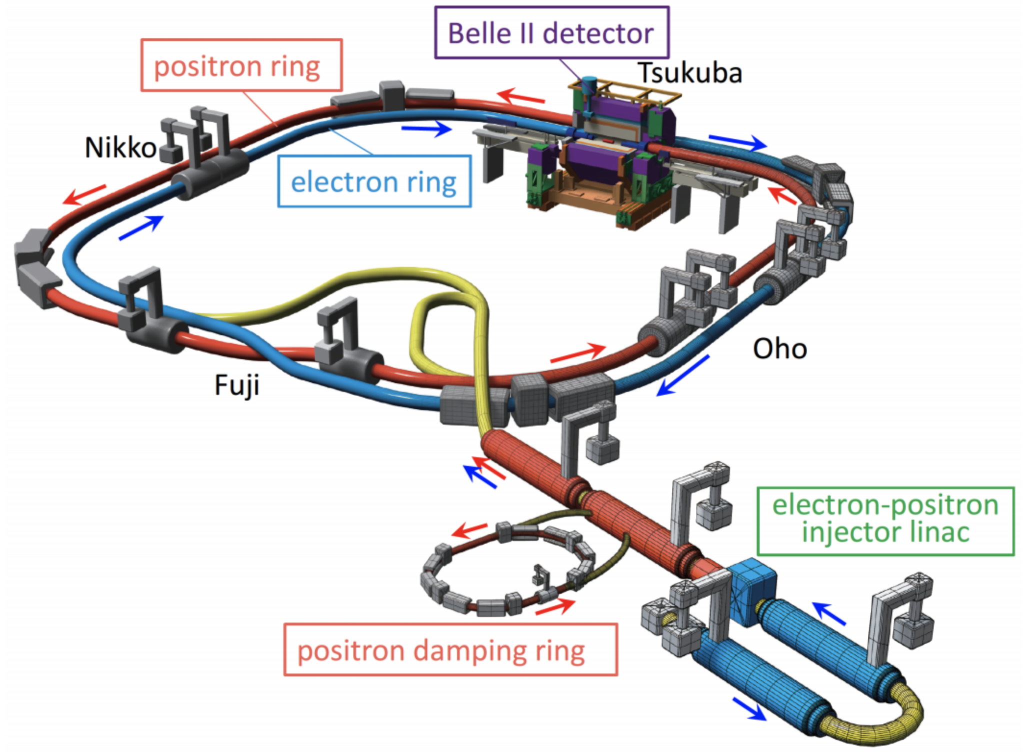

Figure 3.2: Schematic of the SuperKEKB accelerator. Image from [43].

3.4 The Belle II detector

The Belle II detector is an upgrade of Belle. It is comprised of several subdetectors. The tracking system is formed by the Pixel Detector (PXD), Silicon Vertex Detector (SVD), and the Central Drift Chamber (CDC). The PXD and the SVD together are referred to as Vertex Detector (VXD). The PXD was not fully installed in 2018, and neither the PXD nor the SVD are particularly relevant for this analysis. The dedicated particle identification system is comprised of the Aerogel Ring-Imaging Cherenkov detector (ARICH) and the Time-Of-Propagation counter (TOP). The return-yoke of the magnet, which provides a 1.5 T magnetic field used for tracking purposes, is instrumented with scintillators and resistive-plate chambers which are used to detect muons and K_L^0s (KLM). These subdetectors are also not used for the analysis described in this thesis and are therefore not discussed further. Detailed descriptions of all the subdetectors above can be found in the Technical Design Report [44] and in the Belle II Physics Book [45]. Of relevance to this analysis are the Electromagnetic Calorimeter (ECL) and the CDC. They are described in all the relevant details in the following.

A scheme with the top view of Belle II can be seen in Fig. 3.3.

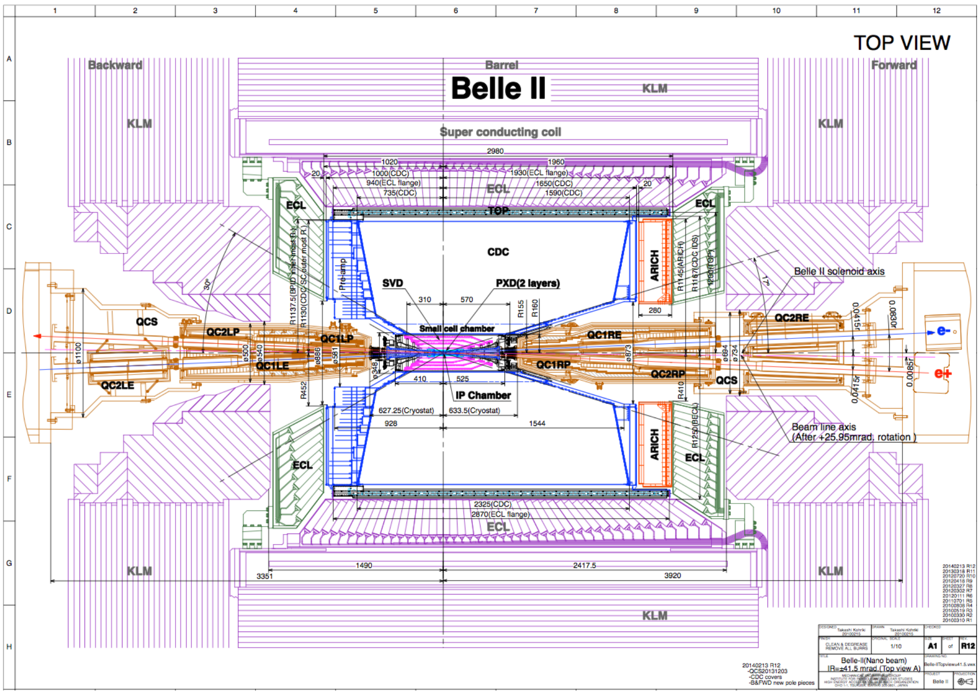

Figure 3.3: Scheme of Belle II detector. Image from [45].

3.4.1 CDC

The Belle II Central Drift Chamber (CDC) (see Fig. 3.4) has to satisfy three important requirements: reconstruct tracks of charged particles and precisely measure their momenta; provide PID information using measurements of energy loss dE/dx within its volume; provide efficient and reliable trigger signals for charged particles.

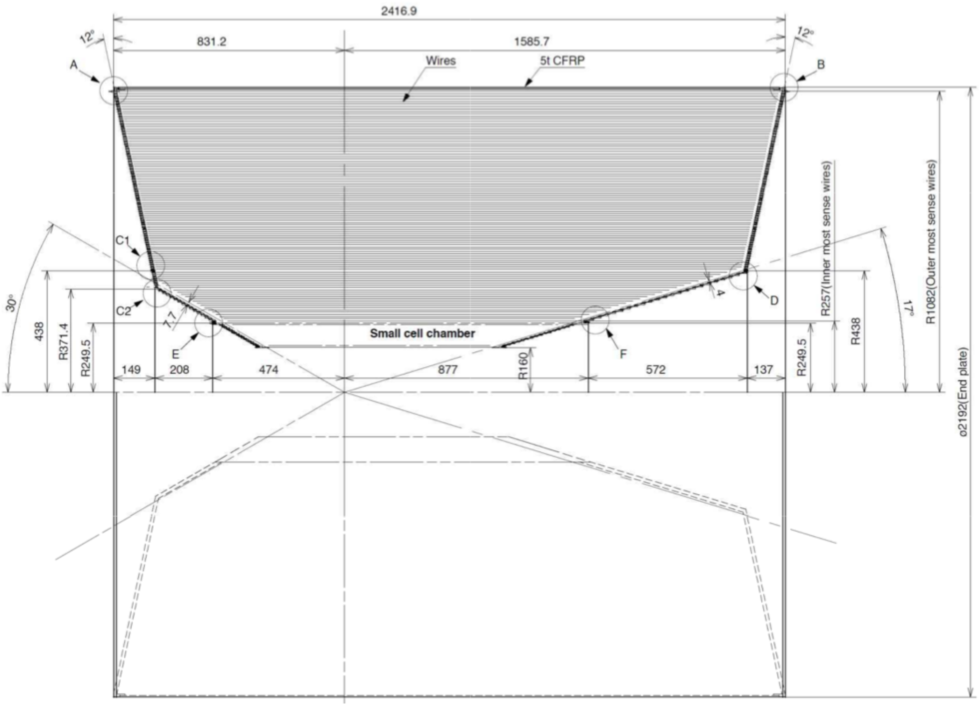

Figure 3.4: Scheme of the Belle II CDC. Lengths are given in mm. Image from [44].

The Belle II CDC is based on cells of wires immersed in a helium-ethane gas mixture and an alternate configuration of differently oriented superlayers, each made up by a certain number of layers, as shown in Fig. 3.5. The polar angle coverage, or *acceptance*, of this detector is $[17°, 150°]$. Spatial resolution is $\lesssim 130\ \mu$m.

Further numerical details can be found in the aforementioned references [44, 45].

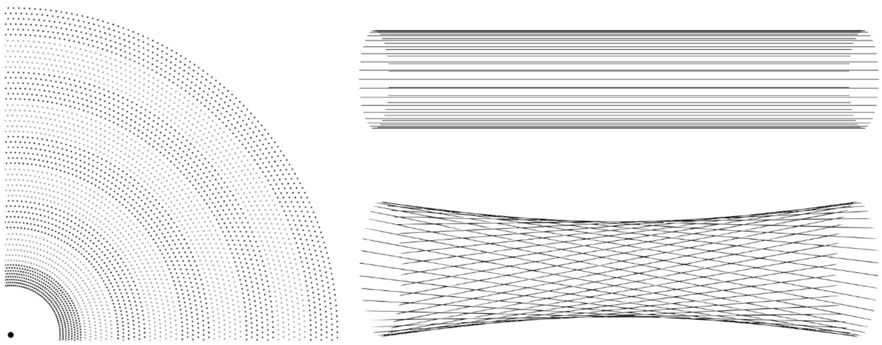

Figure 3.5: Left: a quadrant of the transversal view of the Belle II CDC. The innermost superlayer has eight layers, the others have six. Right: a visualization of stereo wires (bottom) relative to axial wires (top). The skew is exaggerated. Image from [46].

3.4.2 ECL

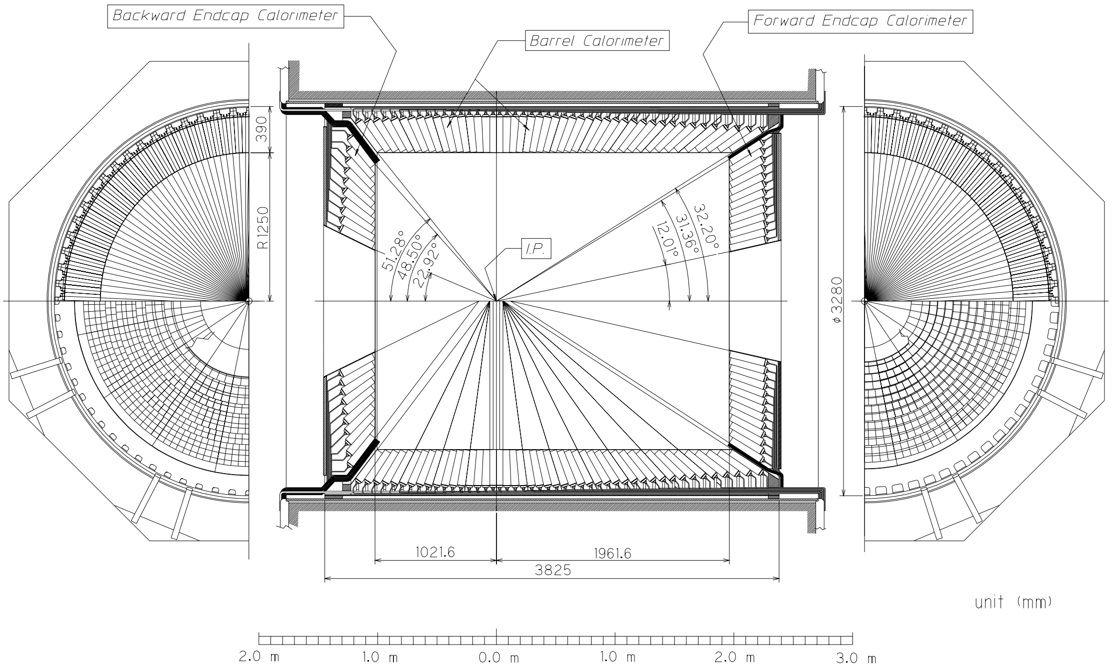

Figure 3.6: Scheme of the Belle II ECL. Image from [47].

The Belle II Electromagnetic Calorimeter (ECL), shown in Fig. 3.6, uses the same crystals and photodetectors as Belle's, but with upgraded readout electronics, to cope with the higher background rate. Given that the final state of the analysis presented in this thesis is fully neutral, this is the main and most relevant subdetector used for this search.

The ECL is able to efficiently measure photon energies between 20 MeV and 7 GeV; it provides angular coordinates of the detected photons; it performs electron, muon, and hadron identification; it generates a proper trigger signal; it performs measurement of the luminosity, both online and offline; together with the KLM, it performs K_L^0 detection.

The ECL is made up of three main components: the barrel section (3 m long, inner radius of 1.25 m) and the two endcaps, in the forward ($z = 1.96$ m) and in the backward ($z = -1.02$ m) regions. It covers the polar angle between 12.4° and 155.1°, with the exceptions of two gaps approximately 1° wide between the barrel and the endcaps. More precisely, the forward endcap covers $[12.0°, 31.4°]$, the barrel covers $[32.3°, 128.7°]$, and the backward endcap covers $[131.5°, 155.0°]$.

The ECL barrel is built with 6624 CsI(Tl) truncated pyramid crystals, arranged in 46 θ rings, each of them covering the entire ϕ angle. The forward [backward] endcap consists of 1152 [960] CsI(Tl) crystals, arranged in 13 [10] θ rings. The average surface and length of the crystals are $\sim 6 \times 6$ cm^2 and 30 cm (16.1 radiation lengths) respectively. Each crystal is read by two photodiodes with a surface of 10×20 mm^2, placed at the back of the crystal.

To avoid missing particles because they fly between the crystals, the ECL is built with a quasi-projective structure, which means that the crystals do not point directly to the IP, but are slightly tilted with respect to the axial direction.

The VXD, the CDC, the TOP, the ARICH, and some instrumentation material are located between the IP and the ECL. Particles can interact with this material, which is hence called *material budget*. Material budget may induce *pair conversions*, i.e. the process, possible only in presence of material, where a photon converts into an electron-positron pair: $\gamma \to e^+e^-$. If this happens beyond the main tracking volume but before the ECL, it may happen that the two tracks associated with the e^+ and the e^- are not reconstructed by the tracking (i.e. the tracking does not reconstruct the displaced tracks), while the magnetic field bends them so that they hit the ECL far enough from each other that they are reconstructed as two different particles. Although, from the moment that there are no tracks associated with them, they are considered neutral particles rather than an e^+e^- pair. The consequence of this is that a single photon is reconstructed as two separate neutral particles, most often as two photons. This has consequences for some analyses, including the one presented in this thesis, as it is further explained in Sec. 5.3.3.

In the following Sections, the ECL reconstruction and clustering process and the ECL-related variables used in the analysis presented in this thesis are explained in all the relevant details.

3.4.2.1 Reconstruction and clustering

Particles interacting electromagnetically with the CsI(Tl) crystals produce electromagnetic showers. These are collected by the two photodiodes connected to each crystals. Each of their signals is amplified by a preamplifier, then they are summed together. Upon further processing, the resulting signal is represented by 31 samples of amplitude as a function of time, what is called a *waveform*. These 31 samples are fit with a template fit, performed by a Field Programmable Gate Array (FPGA): this returns, after calibration, the energy deposited in the crystal by the particle, and the time when this happened. These quantities are stored in an object called `ECLCalDigit`, that is saved for each crystal.

A particle generally releases its energy in more than one crystal. A *clustering algorithm* is deployed to group together all the crystals belonging to the same physical electromagnetic shower, and associate each crystal to its corresponding shower.

To begin with, crystals whose `ECLCalDigit` energy is below a given readout threshold (approximately 0.1 MeV) are ignored. Any `ECLCalDigit` whose energy is above 10 MeV is considered a possible *seed*. Starting from each of these seeds, its neighbors, i.e. its adjacent `ECLCalDigits`, are attached to it if their energy is above 0.5 MeV. If the same crystal was to be attached simultaneously by more than one seed, then a merge occurs, and the two forming clusters are grouped together. If the newly attached crystal has an energy above 1.5 MeV, then the procedure is reiterated and its neighbors are attached, until no more suitable neighbors are available. When this process ends, each of the isolated structures is called *connected region (CR)*.

A local maximum (LM) is defined as a crystal having energy greater than 10 MeV and having no neighbors with higher energy than itself. Within a CR, all LMs are identified. Then, for each CR, all crystals outside of a 5x5 matrix centered around the LM are ignored for the following. If the CR has only one LM, then the process is concluded and the CR is treated as an `ECLCluster`[1]. If more than one LM is present, a *splitting* happens: crystals

---

[1] Technically, at this stage, another internal dataobject is created, which is later con-

in a CR with more than one LM are given weights ω_i, which depend on their energy and geometrical distance from the LMs. At the end of this process, each crystal in the CR is split between the multiple LMs, and a number of `ECLClusters` equal to the number of LMs is created. Such a splitting is the way an `ECLCluster` can have a non-integer number of crystals associated to it.

After all `ECLClusters` have been reconstructed, their energies and position are computed from their own crystals' energies and positions. For both quantities, up to 21 crystals contained within a 5x5 matrix centered on the LM, with the four corners removed, are used for these computations.

The position reconstruction returns the θ and ϕ angles for the reconstructed cluster. These are mildly biased toward the central crystal due to the granularity of the calorimeter.

The energy reconstruction uses the n most energetic crystals, where n is a number which depends on a first-order cluster energy estimation (n decreases at low energies) and on the background conditions (n decreases with high background). This procedure is tuned so to optimize the photon energy resolution.

3.4.2.2 ECL Variables

`clusterTime`: time of the ECL cluster, calculated as the `ECLCluster` timing minus the time of event, t_0. The time t_0 is the time of the event relative to the trigger time. It can be measured by different subdetectors, mostly CDC and/or ECL. The value derives from the combined information of all of them. The `ECLCluster` timing is given by the fitted time of the waveform of the most energetic crystal in the cluster. After all calibrations and corrections, including the time-of-flight one, photons from the IP should have an `ECLCluster` timing that corresponds to t_0. This implies that the `clusterTime` of a prompt photon should be compatible with zero.

`dt99`: the ECL cluster timing uncertainty, such that it contains 99% of true photons. The photon timing uncertainty is determined using MC. This is done by studying the timing distributions in multiple bins of true deposited energy in the highest-energetic crystal of the cluster and local beam background level. For each of these bins, the value `dt99` is computed, being defined such that 99% of the true photons coming from the IP have

---

verted into an `ECLCluster`, but for all and any purposes of this thesis the two terms are interchangeable. The only difference worth to be mentioned is that only showers with energies above 20 MeV are saved and converted into actual `ECLClusters`.

abs(timing)<`dt99`. Very large values of `dt99` are an indication of failed waveform fits in the ECL.

`ZernikeMVA`: the output of a neural network (`MVA` stands for Multi-Variate Analysis) which uses 11 Zernike moments of the cluster as inputs. Zernike moments are calculated per shower in a plane perpendicular to the shower direction as:

$$|Z_{nm}| = \frac{n+1}{\pi} \frac{1}{\sum_i \omega_i E_i} \left| \sum_i R_{nm}(\rho_i) e^{-im\alpha_i} \omega_i E_i \right|, \tag{3.1}$$

where n, m are integers of same parity satisfying $5 \leq n \leq m \leq 0$; i runs over all the crystals in the cluster; E_i, ρ_i, ω_i, α_i are respectively the energy, the scaled radial distance from the cluster center, the energy weight, and the polar angle in the perpendicular plane of the i-th crystal; R_{nm} is the radial component of the corresponding Zernike polyomial. The neural network is trained to discriminate between photon-like clusters and non-photon-like ones. The training is performed on single-particle MC samples of photons and K_L^0. Further details can be found in [48]. `ZernikeMVA` is a shower-shape variable: the closer its value is to 1, the more photon-like the cluster is; vice versa, the closer is to 0, the less photon-like the cluster is.

`clusterNHits`: the number of crystals in an `ECLCluster`. As mentioned in Sec. 3.4.2.1, this number can be not integer in case more than one LMs is present in a CR. This results in each crystal in the CR to be shared amongst the various resulting `ECLClusters`. Thus, `clusterNHits` is equal to the sum of weights ω_i of all crystals in an `ECLCluster`.

`E1oE9`: the ratio between the energy contained in the central crystal (`E1`) and the total energy contained in the 9 crystals in the 3x3 matrix centered around the central crystal (`E9`).

3.4.3 Trigger

In order to reduce the data rate to a level that is manageable by the data acquisition, a two-tier triggering system is employed: a hardware Level 1 (L1) trigger and a software high-level trigger (HLT). In particular, Bhabha and diphoton events need to be vetoed as they highly contribute to the background. The requirements of the trigger are the following:

- Maximum average trigger frequency of 30 kHz;

- Trigger latency $\lesssim 5\ \mu$s;
- Timing precision $\lesssim 10$ ns.

Also a $\sim 100\%$ efficiency for physics processes ($\Upsilon(4S) \to B\bar{B}$ and continuum) is required, as well as for hypothetical Beyond Standard Model (BSM) events.

Amongst the L1 triggers active during the period when the data for the analysis presented in this thesis were acquired, at least one was 100% efficient (see Sec. 5.4). The HLT was turned off at the time, hence it will not be discussed further. In the following, some extra details about the L1 trigger are provided.

The L1 trigger is comprised of various sub-trigger systems, and each summarizes the trigger information coming from its corresponding subdetector. These information are sent to the final-decision logic, made by the Global Reconstruction Logic (GRL) and Global Decision Logic (GDL), the first being an input for the second. The main technological improvements with respect to Belle are the introduction of high-speed serial link cables instead of parallel ones and the usage of FPGA for each component, so that the trigger logic is configurable rather than hard-wired.

The sub-trigger systems are the following: CDC, ECL, BPID (Barrel PID, so from the TOP subdetector), KLM. The CDC provides charged track information: momentum, position, charge, multiplicity. The ECL provides energy deposit and cluster information, and Bhabha and cosmic ray identification. The BPID provides precise timing and hit topology information. The KLM provides muon track information.

The two main sub-trigger systems are the CDC and the ECL ones. The CDC sub-trigger finds and characterizes the charged tracks detected in the CDC, both with 2D and 3D tracking, the last one being particular useful to reject beam-gas and beam-wall background thanks to the z information.

The ECL sub-trigger uses two schemes: total energy and isolated-cluster counting triggers. The first is sensitive to physics events with high electromagnetic energy deposit, while the second is sensitive to multi-hadron events with low-energy clusters and/or minimum ionizing particles (MIPs). The ECL sub-trigger also identifies Bhabha and $\gamma\gamma$ events, needed to measure luminosity.

3.4.4 Software

The Belle II analysis software framework (basf2) [49] is used in this analysis. Specifically, it is used to generate the background and signal events, for ntuple processing and production, and for data processing. The selection optimization has been carried out mostly with python [50] and ROOT [51].

3.5 Beam background

Broadly speaking, *beam background* is an umbrella term that includes all kinds of processes producing undesired accelerator-generated particles that act like background for any search or measurement. Such particles can be produced in single-beam as well as in two-beam processes.

This Section lists the sources of beam background relevant for Belle II. Most of the listed processes are characterized by fewer than two charged tracks in the CDC and no more than two clusters in the ECL. These signatures are similar to the *low-multiplicity* physical events, amongst the topics of interest for Belle II, and thus constitute a challenge for such studies. More details about beam background can be found in [45].

Touschek scattering: it consists of the intra-bunch Coulomb scattering process, between two electrons (or two positrons) belonging to the same bunch. After the scattering, the two particles have their energies changed with respect to the nominal values (7 or 4 GeV), and they fly away from the bunch. The Touschek-scattered particles end up hitting the beam pipe, and if this happens close to the detector, the generated showers can be detected, and act as background. This background is mitigated with the usage collimators (horizontal in particular) and metal shielding.

Beam-gas scattering: particles from the beams can interact with the residual gas present in the pipe. This may happen via Coulomb scattering (which changes the directions of the particles) or via bremsstrahlung scattering (which decreases the energy of the particles).

The rate of beam-gas bremsstrahlung losses within the detector is suppressed with the usage of horizontal collimators. It is small with respect to Touschek background.

Beam-gas Coulomb losses are instead not negligible. The scattered particles fly until they hit the pipe, like for the Touschek background, and the showers they induced can be detected by Belle II. The same counter-measures as for the Touschek component are deployed, in particular the usage of vertical collimators.

Synchrotron radiation: the beams, by circulating in the storage rings (and thus accelerating), emit radiation. Photons from synchrotron radiation have typical energies between the keV and the tens of keV, i.e. in the X-ray range. The VXD is protected from them by a layer of gold present in the inside of the beam pipe close to the IP, and by the shape of the pipe itself, which helps deflecting the produced particles away from the acceptance of Belle II.

Radiative Bhabha: the photons produced in radiative Bhabha events, $e^+e^- \to e^+e^-\gamma$, can interact with the iron of magnet. In photon-nucleus interactions, neutrons can be produced, and such neutrons are the main background source for the KLM. Additionally, the e^+e^- that interacted have less energy than their nominal expected values, and thus may be over-bent by the magnets used to quench the beams when they exit from the IP. If this happens, also the e^+e^- can hit the magnet, so generating electromagnetic showers, that contribute to the background for the KLM.

Two-photon process: it consists of the process $e^+e^- \to e^+e^-e^+e^-$. One of the two e^+e^- pair has very low momentum, thus can spiral extensively within the detector and induce a huge number of hits in the subdetectors. Additionally, the more energetic e^+e^- pair can be over-bent once exiting the IP, in a similar fashion as for the radiative Bhabha background.

Chapter 4

Analysis overview and data sample

This Chapter provides a detailed overview and introduction to the analysis strategy, as it was briefly described in Sec. 2.3.2.

Section 4.1 explains the various steps necessary for performing the analysis, which are later described in thorough details in the relative Chapters.

Section 4.2 describes the characteristics of the signal signature and of the processes which constitute a background for this analysis.

Section 4.3 contains the details of the Monte Carlo and data samples used.

4.1 Analysis strategy outline

The analysis described in this thesis is a so-called *blind analysis* [12]. This means that the analysts avoid to access the relevant distributions of the full experimental data sets until the whole analysis strategy, mainly the selection optimization, has been finalized. This is done to avoid any experimenter's bias, which could lead to adapt the selection to enhance a statistical fluctuation in such a way to make it look like a statistically significant signal-like excess.

Many techniques can be used to blind an analysis. In this case, all the optimization procedures have been performed on MC samples and on data samples which do not include the signal region, i.e. *orthogonal* data sets or *sidebands*, the latter being data sets which are identical to the signal data set except for one or few different variable cuts, which make them completely non-overlapping.

Given that the mass of the ALP is a free parameter of the model, the search is carried out by performing a *mass scan* in the mass range between 0.2 and 9.7 GeV/c^2, in steps of half of the σ parameter of the Crystal Ball (CB) modeling the signal peak (see Sec. 7.2.1) at the mass under exam.

The mass range roughly corresponds to the kinematically allowed range, i.e. up to about 10.6 GeV/c^2. The lower limit is because very light ALPs are highly boosted, so the two photon daughters are close to each other and tend to form overlapping clusters in the calorimeter. This means that often they are reconstructed as one single photon, and therefore do not pass the selection, which requires the presence of three distinct photons in the event. The upper limit is because of the photon energy cut included in the selection (see Sec. 5.2.3), which rejects the soft recoil photons associated with heavy ALPs, whose energy is given by Eq. 2.8. The recoil photon is monoenergetic in the center-of-mass frame, which is boosted with respect to the laboratory frame (see Sec. 3.3), and can be produced in any direction, thus it is not monoenergetic in the laboratory frame. This implies that the photon energy cut induces a roughly linear decrease in signal detection efficiency above a certain ALP mass, rather than a sharp cut-off.

The signal selection has been performed with a multi-dimensional rectangular cut approach. The cut applied to each *selection variable* was optimized via maximization of a figure of merit, the Punzi figure of merit (PFM) [52]. For each variable, the cut that maximizes the PFM is found, checking also if it is mass dependent. This was the case for only one variable. The procedure has been iterated, and repeated with different cut orders, to validate the stability of the cuts found.

Most of the selection variable cuts derive from studies performed on MC samples, while some derive from studies performed on orthogonal non-unblinding data sets or from studies performed on the sidebands.

The optimal selection criteria result in some remaining background, which is then modeled. To search for signal peaks over the locally smooth background, a 1D simultaneous fit of signal+background is performed to the ALP mass distribution.

As discussed further in Sec. 7.2, the signal is made up of two components: a peaking component, modeled with a CB, and a combinatorial component. Both of them are fixed for the simultaneous fit, with parameters obtained from MC. The background component is modeled with a polynomial. Its order and the fit range, which are ALP-mass dependent, are chosen such to make the reduced χ^2 of the background-only fit as close as possible to one, and also to make the reduced χ^2 as a function of the ALP mass as smooth as possible.

The mass scan is divided into two regions: at low ALP masses the fit is performed using the diphoton mass, see Eq. 2.11, while at high ALP masses the recoil mass is used, see Eq. 2.12. Two mass definitions are deployed because each of them has better (i.e. smaller) resolution for different ranges of ALP masses. In particular, the diphoton mass provides a better resolution for light ALP, and vice versa for the recoil mass.

This is because the photon energy resolution is roughly directly proportional to the photon energy itself. Thus, diphoton mass is used at low ALP masses because it is built upon two low-energy photons, and the higher the ALP mass, the higher the energies, the higher the uncertainties. Conversely, the recoil mass is used at high ALP masses because it is computed using the recoil photon, whose energies diminishes at higher ALP masses, so the higher the mass, the lower the energy, the lower the uncertainty. The method to determine the exact transition point ($m_a = 6.85$ GeV/c^2), based the exclusion upper limit expected from a background-only sample, is detailed in Sec. 7.1.3.

The squared recoil mass can assume negative values, and in that case the recoil mass itself would be non-real. For this reason, the fit is not performed directly to linear diphoton and recoil mass distributions, but on the squared distributions.

A kinematic fit of the three photons in the event with the beam energy constraint is not used because the covariance matrix of the photon uncertainty, needed for such a fit, was not properly tuned with at the time of the data taking. The status of the studies on the photon covariance matrix is described in Chap. 8.

When performing the mass scan, both a signal+background and a background-only fit are performed for each mass hypothesis. The significance of the signal hypothesis is quantitatively tested by computing the local significance (see Chap. 7).

A signal-like event excess is considered statistically significant if the *global* significance is higher than 3. The global significance takes into account the fact that testing a sufficiently high number of times a false hypothesis, there will be some tests which will result in a positive result, just because of statistical fluctuations in the testing procedure. In order to pass from local to global significance, the Look-Elsewhere Effect (LEE) has to be taken into account; although, as stated in the abstract of this thesis, no *local* significance higher than 3 has been found, therefore this procedure was not necessary.

If no statistically significant signal-like event excess is found, as it is the case for the analysis presented in this thesis, an UL is computed and set.

The convention is to set an UL with a Confidence Level (CL) of 95%. This is done with a Negative Logarithmic Likelihood (NLL) approach [53], using the CLs method [54].

Finally, systematic uncertainties (or simply *systematics*) have to be taken into account, as they affect the UL, weakening it, i.e. making it higher. Systematic uncertainties [55] arise from a variety of sources, like assumptions made by the experimenters, models used to infer from data, or uncertainties associated with the measurement devices, e.g. calibration constants. Statistical uncertainties associated with different measurements are generally uncorrelated and scale with the inverse of the square root of data taken. Systematic uncertainties often do not have such a scaling and are generally correlated between different measurements.

For the analysis presented in this thesis, the main sources of systematics are associated with the parameters of the functions used to model the background, the mass range used for the signal+background fit, the signal efficiency, and the signal width. Other sources are taken in consideration and examined, but are negligible. More details are in Sec. 7.5.

4.2 Signal and background characteristics

The search presented in this thesis is for ALPs, produced in association with a recoil photon, and decaying into two hard, resolved photons, via the process $e^+e^- \to \gamma a(a \to \gamma\gamma)$. This is a three-photon final state with no missing energy and no tracks. These features alone already allow to set up some selection cuts useful to reject many background sources.

Background is anything which can mimic this signature; in particular, the following five sources are considered the main ones:

- radiative diphoton events, $e^+e^- \to \gamma\gamma(\gamma)$, i.e. diphoton events with an extra beam background photon. The three photons need to add up to the total energy, which is equal to 10.58 GeV in the cms, in order to mimic the signal. This is the main source of background;

- radiative Bhabha events, $e^+e^- \to e^+e^-(\gamma)$, where both tracks are not reconstructed and therefore they appear as photons to the reconstruction software. This process has the highest cross section, but the probability of the tracking algorithm missing both tracks is very low, order of 10^{-4}, so this is only the second-largest background;

- peaking background, $e^+e^- \to P\gamma(\gamma)$, $P = (\pi^0, \eta, \eta\prime)$, $P \to \gamma\gamma$. These processes are identical to the signal one, but they have a very low cross section, and furthermore the Branching Ratios (BRs) for the η resonances to decay into a photon pair are small.

Table 4.1 lists the cross sections σ for the background processes considered. The technical details and the settings of the generators used for the MC production of the background samples are in Sec. 4.3.2.

Background process	σ $[nb]$
$e^+e^- \to \gamma\gamma(\gamma)$	4.01
$e^+e^- \to e^+e^-(\gamma)$	124.5
$e^+e^- \to \pi^0\gamma(\gamma)$, $\pi^0 \to \gamma\gamma$	0.243×10^{-2}
$e^+e^- \to \eta\gamma(\gamma)$, $\eta \to \gamma\gamma$	0.96×10^{-3}
$e^+e^- \to \eta\prime\gamma(\gamma)$, $\eta\prime \to \gamma\gamma$	0.3×10^{-4}

Table 4.1: Cross section σ for the background processes of interest for the analysis presented in this thesis. As it will be shown later in the text, the rejection power of the selection against the different channels varies considerably. The processes are listed in order from the most to the least relevant, after having applied the selection.

4.3 Data samples

This Section lists the technical details and the settings of the generators used for the MC production of the signal and background samples, as well as the information about the data taken by Belle II and used for this analysis.

4.3.1 Signal

The signal samples have been produced with `MadGraph5_aMC@NLO` [56] version v2.6.3.2. The version of the MadGraph model, `ALP_UFO` [57, 33], is included in the generator package. An initial-state radiation (ISR) correction [58] is included. All the samples have been produced with simulated nominal beam background conditions at the time of the Phase 2 data taking in 2018, to better reproduce the situation in the real detector.

The only non-zero ALP coupling is $g_{a\gamma\gamma}$ (see Eq. 2.4). Across the whole mass spectrum, and for reasonable choices for the coupling $g_{a\gamma\gamma}$, the decay width of the ALP is negligible compared to experimental resolution. This is shown in Fig. 4.1, where the trend of the signal (Eq. 2.9) as a function of the ALP mass is plotted for $g_{a\gamma\gamma} = 10^{-4}$ GeV^{-1}.

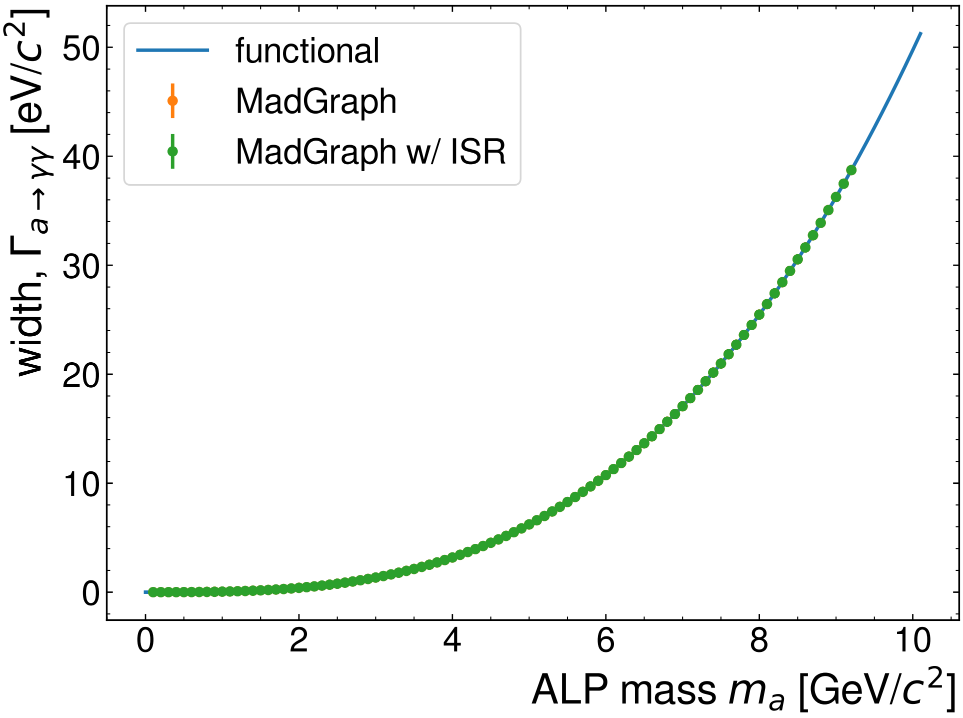

Figure 4.1: Signal width as a function of the ALP mass, for $g_{a\gamma\gamma} = 10^{-4}$ GeV^{-1}. The values obtained from the theoretical formula (see Eq. 2.9) and from the MC generator, with and without ISR, are in agreement.

About 500 samples have been produced with these settings, each for a different simulated ALP mass between 0.1 and 10.5 GeV/c^2, and each having 20k (20 thousand) generated signal events. The mass distance between each of them roughly corresponds to half of the signal width. The majority of the samples have been used for the UL extraction studies, while a smaller subset of benchmark mass values have been used for the procedure of selection optimization.

4.3.2 Backgrounds

The $e^+e^- \to \gamma\gamma(\gamma)$ and $e^+e^- \to e^+e^-(\gamma)$ processes are simulated using `BABAYAGA.NLO` [59, 60]. `BABAYAGA.NLO` generates multiphoton initial ISR, final state radiation (FSR), and the interference of initial and final state radiation. Z exchange and γ-Z interference are included at the leading order. Narrow resonances and vacuum polarization corrections are included. All events are generated assuming the IP position being at its nominal value, $(0, 0, 0)$, and beam energy spread is included.

$e^+e^- \to \gamma\gamma(\gamma)$ events are generated in the polar angle range $0^\circ < \theta < 180^\circ$. Then, a preselection is applied on MC truth level, requiring that they have at least one true γ with $E > 1.5$ GeV within $17^\circ < \theta_{\text{lab}} < 150^\circ$. The effective cross section of this process, after preselection, is $\sigma_{eff} = 4.01$ nb. The sample is comprised of 60M (60 millions) events. The theoretical uncertainty for both the total and the differential cross section is about 1%.

$e^+e^- \to e^+e^-(\gamma)$ events, i.e. radiative Bhabha, are generated in the polar angle range $15^\circ < \theta_{\text{cms}} < 165^\circ$. With this requirement, the effective cross section is 124.5 nb. The sample is comprised of 100M (100 millions) events. The theoretical uncertainty for both the total and the differential cross section is about 1%. Tracking inefficiency, i.e. the probability that a track is not reconstructed, was simulated to be zero in MC, meaning that all the charged tracks are correctly reconstructed by the tracking algorithm in the simulation. From early evaluations on data, although, it was found that the tracking inefficiency was $2.5 \pm 2.5\%$ per track. Therefore, a 2.5% inefficiency per track was manually added for radiative Bhabha events. Technically, the following was done: all simulated $e^+e^- \to e^+e^-(\gamma)$ events are processed with no tracking inefficiency; the tracking-related quantities are ignored and only the ECL-related variables are considered for the purposes of the selection (Chap. 5); each event is weighted with a weight $w_T = 0.025^{N_T}$, where N_T is the number of tracks in the event (so, in the majority of cases, $w_T = 0.025^2 = 6.25 \cdot 10^{-4} = 0.0625\%$).

The $e^+e^- \to P\gamma(\gamma)$, $P = (\pi^0, \eta, \eta\prime)$ processes are simulated with `PHOKHARA 9.3` [61]. Each sample is comprised of 45k events. The cross sections and the BRs for the meson to decay into two photons, $P \to \gamma\gamma$, for these processes are:

- $e^+e^- \to \pi^0\gamma(\gamma)$, $\pi^0 \to \gamma\gamma$: $\sigma = 0.243 \times 10^{-2}$ nb, BR=98.8%;
- $e^+e^- \to \eta\gamma(\gamma)$, $\eta \to \gamma\gamma$: $\sigma = 0.96 \times 10^{-3}$ nb, BR=39.4%;
- $e^+e^- \to \eta\prime\gamma(\gamma)$, $\eta\prime \to \gamma\gamma$: $\sigma = 0.3 \times 10^{-4}$ nb, BR=2.2%.

4.3.3 Cosmics

A *cosmic* sample, that is a data set of events from cosmic rays, has been used to assess the relevance of this kind of event for this analysis, i.e. to estimate how many signal-like events could be mimicked by cosmics passing through the Belle II detector during the data acquisition time. The data set used to estimate the amount of cosmics in the actual signal data set corresponds to a data taking time of 86 400 s = 23.5167 h, during which all detectors and the magnetic field were on while the beams were off.

The amount of signal-like events found in this data set is then scaled according to the duration of the data taking of the data set used for this analysis.

The final selection (see Sec. 5.5) is applied to the cosmic sample. Using the $E(\gamma)$ cut of 1000 [650] MeV, only 3 [6] candidates are selected, and amongst these only 1 [4] is in the proper mass range $M^2 < [>]$ 4 GeV^2/c^4. Scaling the data-taking time of the cosmic sample to the data-taking time of Phase 2, which is 277.6075 h, only 12 [47] candidates are expected. This is negligible with respect to the main backgrounds (hundreds to thousands of background events are expected in most fits, see again Sec. 5.5), therefore the contribution from cosmics is henceforth ignored.

4.3.4 Data

The data used for the analysis presented in this thesis were taken during *Phase 2*, the preliminary data taking that took place from April to July 2018. All events were taken at the $\Upsilon(4S)$ resonance. Only runs where the CDC, ECL, and TRG subdetectors (see Sec. 3.4) were active are included.

The total available data set corresponds to an integrated luminosity of 497 pb^{-1}.

Chapter 5

Event selection

This Chapter describes in detail the selection applied to discriminate signal events from background events.

Section 5.1 explains the selection optimization procedure.

Section 5.2 lists the variables whose choice has been based upon MC-only studies, and explains each of them.

Section 5.3 lists the variables included in the selection due to studies performed on other data sets.

Section 5.5, finally, summarizes all the cuts and shows the performances of the selection.

5.1 Optimization strategy

The discrimination between signal and background is performed with a multi-dimensional rectangular cut. The value of each selection cut, i.e. the cut applied to each *selection variable*, has been optimized via the maximization of a figure of merit, the PFM [52]:

$$\mathrm{PFM} = \frac{\epsilon_S}{\frac{a}{2} + \sqrt{B}}, \tag{5.1}$$

where $\epsilon_S \equiv \frac{S}{S_{gen}}$ is the signal efficiency (S is the number of signal events and S_{gen} = 20k, see Sec. 4.3.1), a is the target significance in number of standard deviations, set to $a = 5$, corresponding to the accepted "discovery" threshold, and B is the number of background events. Both S and B are counted within an asymmetric window W_{PFM} around the ALP mass value under exam, with $W_{PFM} = [m_a - 5\sigma_{CB}, m_a + 3\sigma_{CB}]$, where m_a is the ALP mass and σ_{CB} is the σ parameter of the CB used to model the signal peak at that mass value (more details on the signal fit are in Sec. 7.2.1).

The procedure is the following: fix the cuts to all variables but the one being optimized to a very loose value, or not apply those cuts at all; for multiple benchmark ALP mass hypotheses, compute the PFM as a function of the variable cut; find the cut value(s) that maximize the PFM; repeat this process for all variables.

This procedure has then been repeated by fixing the variable cuts which were not under study to the values just obtained, to cross-check that the resulting cut value for the variable under study remained the same.

For additional cross-check, the procedure has been performed also by applying only some of the cuts, to check that the optimal selection point for the variable under study remained the same.

All the steps above were performed simultaneously for multiple ALP-mass benchmarks to verify if it was needed to adopt mass-dependent cuts. All cuts but one are indeed independent from the ALP mass, the exception being the cut on the photon energy, which is changed from $E_\gamma > 1000$ MeV to $E_\gamma > 650$ MeV at $m_a = 4.0$ GeV/c^2 to avoid *sculpting effects* (see Sec. 5.2.3).

Given how small the contributions from $e^+e^- \to P\gamma(\gamma)$, $P = (\pi^0, \eta, \eta\prime)$ are (see Sec. 4.3.2), the selection optimization has been carried out using only the two main background sources, $e^+e^- \to e^+e^-(\gamma)$ and $e^+e^- \to \gamma\gamma(\gamma)$. It is verified a posteriori that the contributions from $P = (\pi^0, \eta, \eta\prime)$ are indeed negligible.

All the selection studies are performed on data sets that have undergone a very loose pre-selection (simply referred to as *loose selection* in the text) to allow a quicker data processing. The loose selection is the following:

- Only the three most energetic photons are considered to reconstruct the event;
- $E_\gamma \geq 100$ MeV for each photons;
- $0.80\sqrt{s} \leq m_{\gamma\gamma\gamma} \leq 1.05\sqrt{s}$;
- $\theta_\gamma \in [17°, 150°]$, that corresponds to the CDC acceptance.

Only truth-matched (TM) signal events are used, i.e. only events where the algorithm that associates a reconstructed particle to the original MC particle generated during the simulation successfully matched all the three reconstructed photons to MC-generated photons (that is, those photons are not background photons), and where the ALP too has been successfully truth-matched. The matching algorithm distinguishes between *physical particles*, i.e. particles produced by the physical process requested to the generator (in this case the recoil photon and the ALP, and subsequently the two decay photons), and *non-physical particles*, which derive from the simulated beam background. Only physical particles can successfully be truth-matched.

5.2 Variables based on MC

In this Section, all cuts resulting from MC studies are listed and explained. For each variable, the distributions for signal and background are shown, both for the loose and for the final selection. Only events belonging to the mass windows W_{PFM} previously defined are shown. The distributions are shown for three benchmark mass values (low, intermediate, and high mass, i.e. 0.5, 5.0, and 9.2 GeV/c^2 respectively).

Plots summarizing the PFM studies are shown as well, with a finer mass scan, in the form of the trend of the PFM as a function of the cut value on the variable under study. They are shown both when the other cuts are fixed at the values of the loose selection, and when all other cuts are fixed at the values of the final selection.

5.2.1 Three most energetic photons

Any event with three or more photons passing the applied selection cuts constitutes a possible signal event. In an event with N reconstructed photons, the number of possible candidates is given by $C(N)$, defined as:

$$C(N) \equiv 3\binom{N}{3} = \frac{N(N-1)(N-2)}{2}, \tag{5.2}$$

where the binomial factor $\binom{N}{3}$ gives the number of triplets that can be made out of N elements, and the additional factor 3 accounts for the distinct possibilities within each triplet (each photon can be the recoil photon). Note that $C(N)$ increases quickly: for example, $C\,(5) = 30$ and $C\,(10) = 360$. For $N = 3$, the minimal case, we have $C\,(3) = 3$ candidates per event.

Due to the rapid growth of $C(N)$, only the three most energetic photons passing all other photon-related cuts are used. In this way, the combinatorial background is minimized and the number of candidates per event is fixed and known. No best-candidate selection is applied, i.e. *all the three candidates for each event are kept.* This implies that in case of a signal event, only one out of three candidates is the actual ALP, the other two being wrong combinations. This has some implications for modeling, as discussed in Sec. 7.2.

Table 5.1 shows that by selecting the three most energetic photons, only a negligible amount of signal events is lost, less than 0.4 percent points.

m_a [GeV/c^2]	Selection efficiency with all photons [%]	Selection efficiency with the three most energetic photons [%]
0.5	69.3	68.9
3.0	66.1	65.8
9.3	68.2	67.9

Table 5.1: Percentage efficiency of the loose selection for different ALP masses, not requiring truth-matching. Selecting only the three most energetic photons affects the signal efficiency by less than 0.4%.

5.2.2 Angular acceptance

The optimization of the selection was initially performed in two different polar angle acceptances: using the full ECL acceptance (see Sec. 3.4.2), and using only a subset of the ECL barrel acceptance, i.e. the barrel deprived of the first and last 5°, approximately corresponding to three rings of crystals, to avoid potential problems related with improper cluster reconstruction due to shower leakage outside of the barrel itself. The *reduced barrel* acceptance is thus defined as $[37.3°, 123.7°]$. Fig 5.1 shows an example of this study.

For each ALP mass benchmark, the PFM is higher when using the reduced barrel acceptance rather than the full CDC; this is because the main backgrounds are peaked toward the beams axis, i.e. in the endcaps areas. In all of the Figures of the following Sections, only the PFM trend for the reduced barrel acceptance is shown.

The polar angle cut is therefore:

$$37.3° < \theta(\gamma) < 123.7°. \tag{5.3}$$

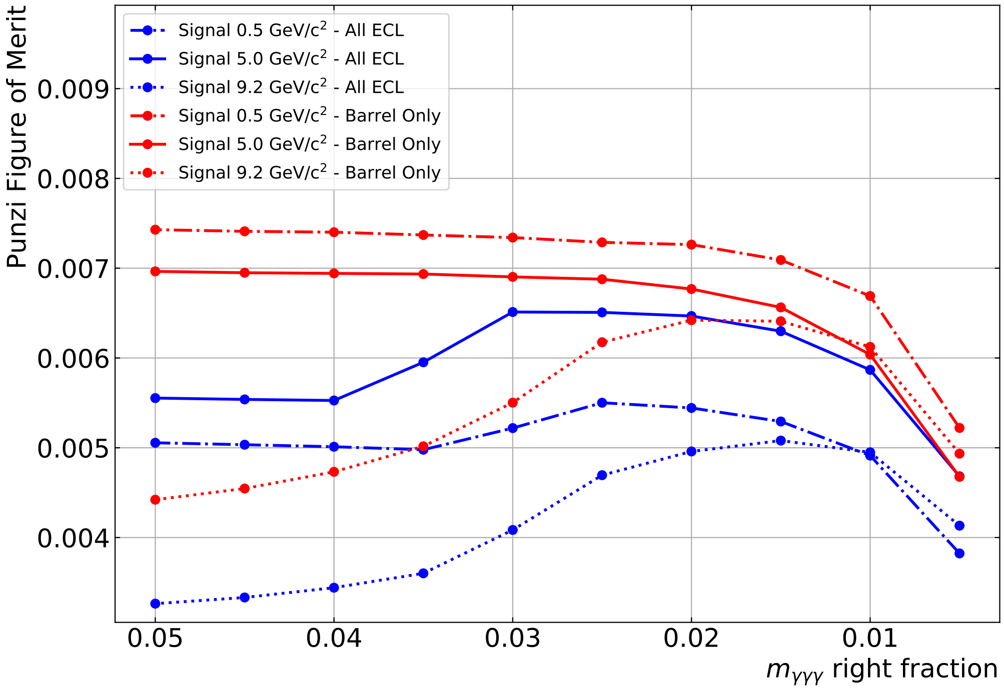

Figure 5.1: PFM as a function of the selection on the invariant mass of the three photons constituting the event candidate. Before selections on other variables are applied. Comparison between the two different acceptances. See Sec. 5.2.4 for the exact variable definition.

5.2.3 Photon energy

The cut on photon energy has the strongest discriminating power, together with the angular acceptance cut. It reduces the contribution deriving from beam background photons, also known as *non-physical photons*, and it helps rejecting many $e^+e^- \to \gamma\gamma(\gamma)$ and $e^+e^- \to e^+e^-(\gamma)$ events, where often the radiative photon or the particle it was radiate from has low energy.

This cut effectively sets up the upper limit of the mass scan: the energy of the recoil photon, see Eq. 2.8, decreases while the ALP mass increases. So, by accepting only photons whose energy is above a given threshold, the low-energy recoil photons associated with high-mass ALPs are rejected. For low-mass ALPs all photons are relatively high energetic.

Due to the sorting algorithm of *basf2*, the three photons making up each event are sorted as shown in Table 5.2. This is because in each of the three candidates built for each event, each photon assumes a different role, as it may either be considered one of the two photons deriving from the decay of the ALP or the recoil photon. This naming scheme is used throughout the rest of this thesis.

Candidate	Low E γ	Medium E γ	High E γ
0	daughter 0	daughter 1	recoil
1	daughter 0	recoil	daughter 1
2	recoil	daughter 0	daughter 1

Table 5.2: Scheme showing how the analysis package reconstructs the three signal candidates.

Figs. 5.2, 5.3, 5.4 show the distributions of the photons' energies, and Fig. 5.5 shows the PFM studies, with the final value of the cut (for high ALP masses, see the following) highlighted as well.

The photon energy cut is the only one that is mass dependent. This is done to avoid *sculpting effects*, i.e. the effects that a given selection may induce on the background shape, in particular on the distribution of the squared diphoton mass $m^2_{\gamma\gamma}$ of the ALP candidate. This effect is shown in Fig. 5.6, where the $m^2_{\gamma\gamma}$ distribution for the main background, $e^+e^- \to \gamma\gamma(\gamma)$, is shown for different values of the cut on $E(\gamma)$, and in Fig. 5.7, which shows systematic excesses for the significance, on multiple MC-only simulated extractions, at different values of $m_{\gamma\gamma}$, depending on the cut on the photon energy. More detailed information on the procedures used to extract these values are in Chap. 7. In Fig. 5.6 it can be seen that for the energy cut $E(\gamma) > 0.65$ GeV the mass distribution shows a step at around $m^2_{\gamma\gamma} = 13$ $\mathrm{GeV}^2/\mathrm{c}^4 \approx \left(3.5\ \mathrm{GeV/c}^2\right)^2$, while this step moves to higher mass values when increasing the photon energy cut, as it can be seen in the same plot and in Fig. 5.7. This is the reason why the cut on $E(\gamma)$ is made energy dependent.

The photon energy cut is therefore:

$$\begin{aligned} E(\gamma) &\geq 1.00\ \mathrm{GeV\ for}\ m_a < 4.0\ \mathrm{GeV/c}^2, \\ E(\gamma) &\geq 0.65\ \mathrm{GeV\ for}\ m_a \geq 4.0\ \mathrm{GeV/c}^2. \end{aligned} \tag{5.4}$$

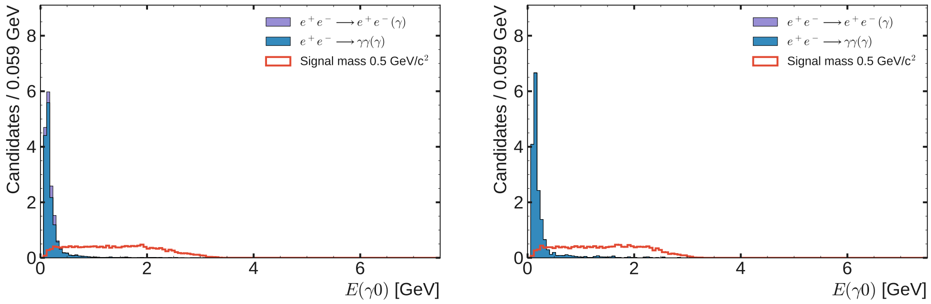

(a) Energy of the daughter 0 of the ALP candidate. Loose selection.

(b) Energy of the daughter 0 of the ALP candidate. Final selection.

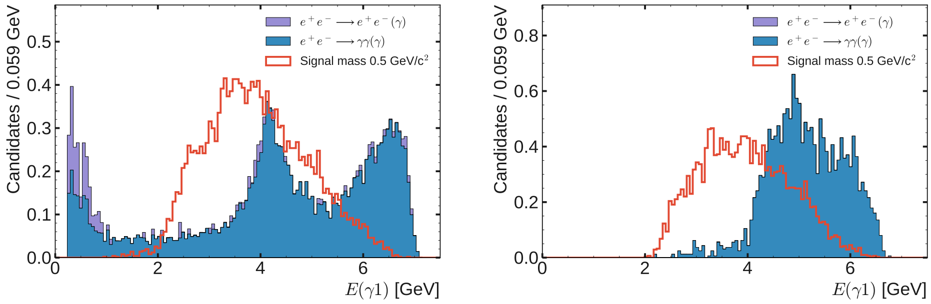

(c) Energy of the daughter 1 of the ALP candidate. Loose selection.

(d) Energy of the daughter 1 of the ALP candidate. Final selection.

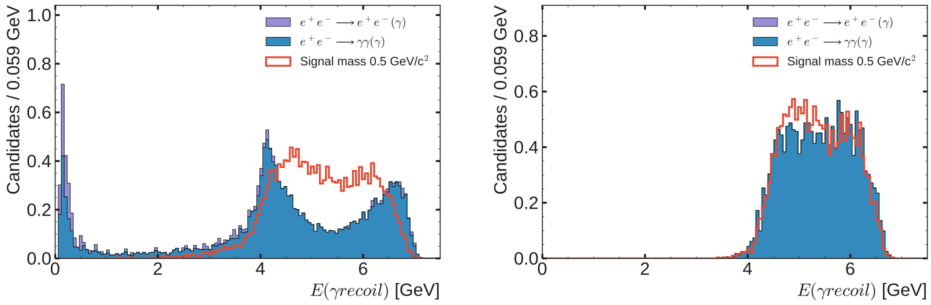

(e) Energy of the recoil photon. Loose selection.

(f) Energy of the recoil photon. Final selection.

Figure 5.2: Energy distributions for the three photons for signal ($m_a = 0.5$ GeV/c^2), $e^+e^- \to e^+e^-\,(\gamma)$, and $e^+e^- \to \gamma\gamma\,(\gamma)$. Left: after applying the loose selection; right: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Normalized to unity.

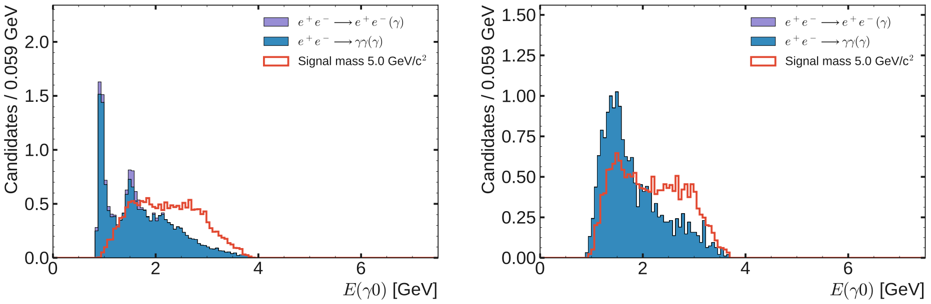

(a) Energy of the daughter 0 of the ALP candidate. Loose selection.

(b) Energy of the daughter 0 of the ALP candidate. Final selection.

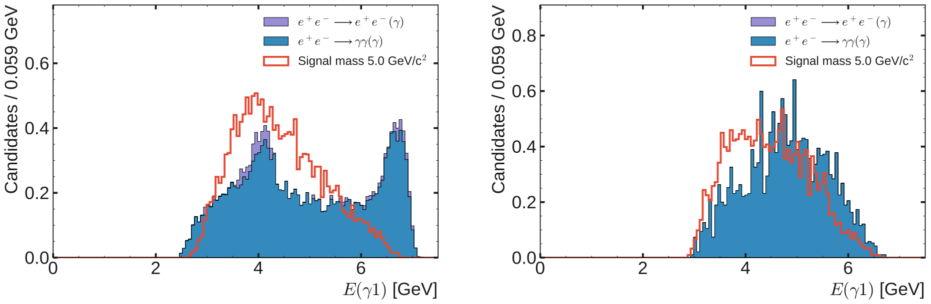

(c) Energy of the daughter 1 of the ALP candidate. Loose selection.

(d) Energy of the daughter 1 of the ALP candidate. Final selection.

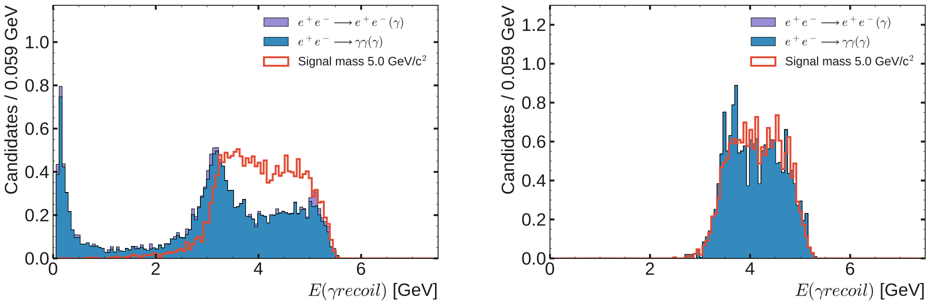

(e) Energy of the recoil photon. Loose selection.

(f) Energy of the recoil photon. Final selection.

Figure 5.3: Energy distributions for the three photons for signal ($m_a = 5.0$ GeV/c^2), $e^+e^- \to e^+e^-(\gamma)$, and $e^+e^- \to \gamma\gamma(\gamma)$. Left: after applying the loose selection; right: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Normalized to unity.

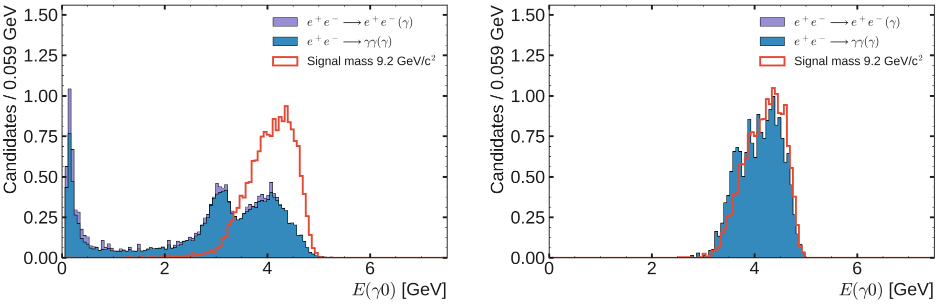

(a) Energy of the daughter 0 of the ALP candidate. Loose selection.

(b) Energy of the daughter 0 of the ALP candidate. Final selection.

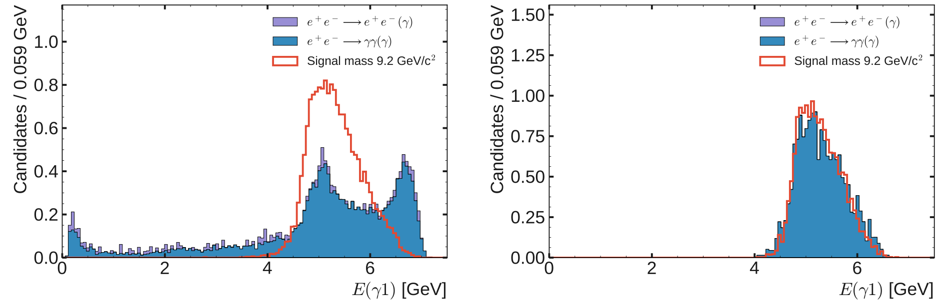

(c) Energy of the daughter 1 of the ALP candidate. Loose selection.

(d) Energy of the daughter 1 of the ALP candidate. Final selection.

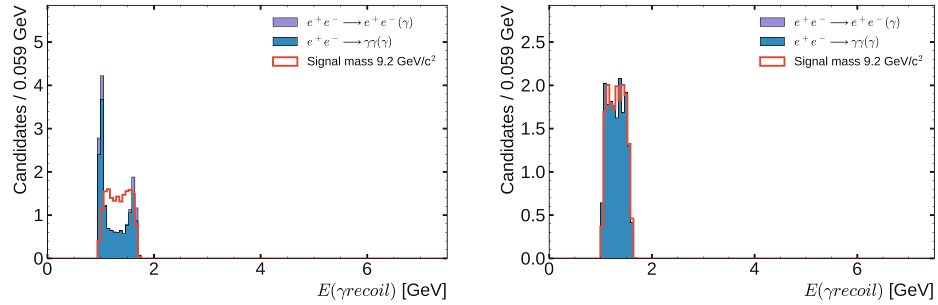

(e) Energy of the recoil photon. Loose selection.

(f) Energy of the recoil photon. Final selection.

Figure 5.4: Energy distributions for the three photons for signal ($m_a = 9.2$ GeV/c^2), $e^+e^- \to e^+e^-\,(\gamma)$, and $e^+e^- \to \gamma\gamma\,(\gamma)$. Left: after applying the loose selection; right: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Normalized to unity.

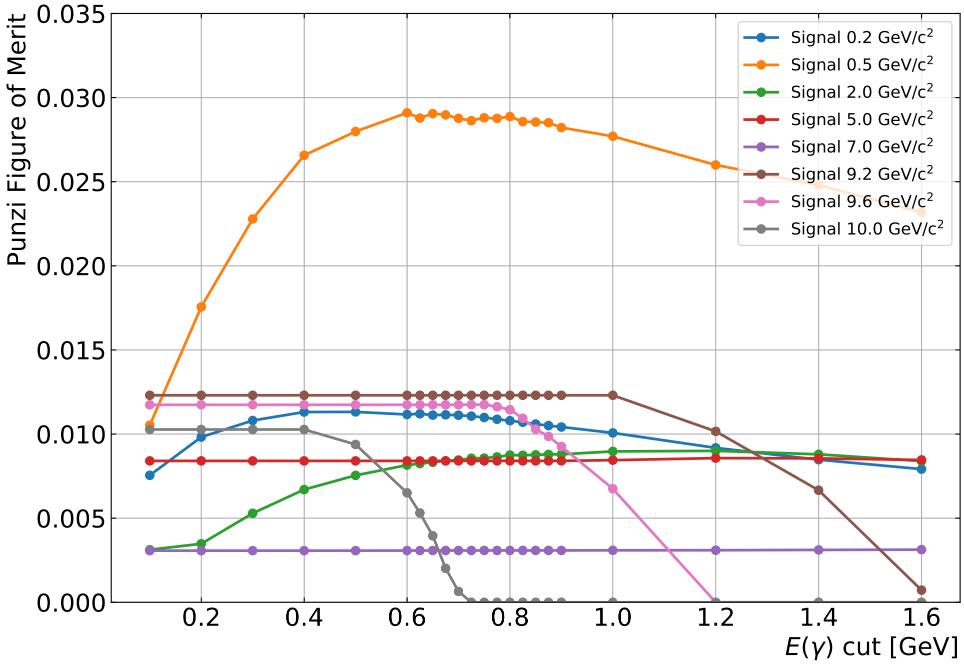

(a) Before selections on other variables are applied.

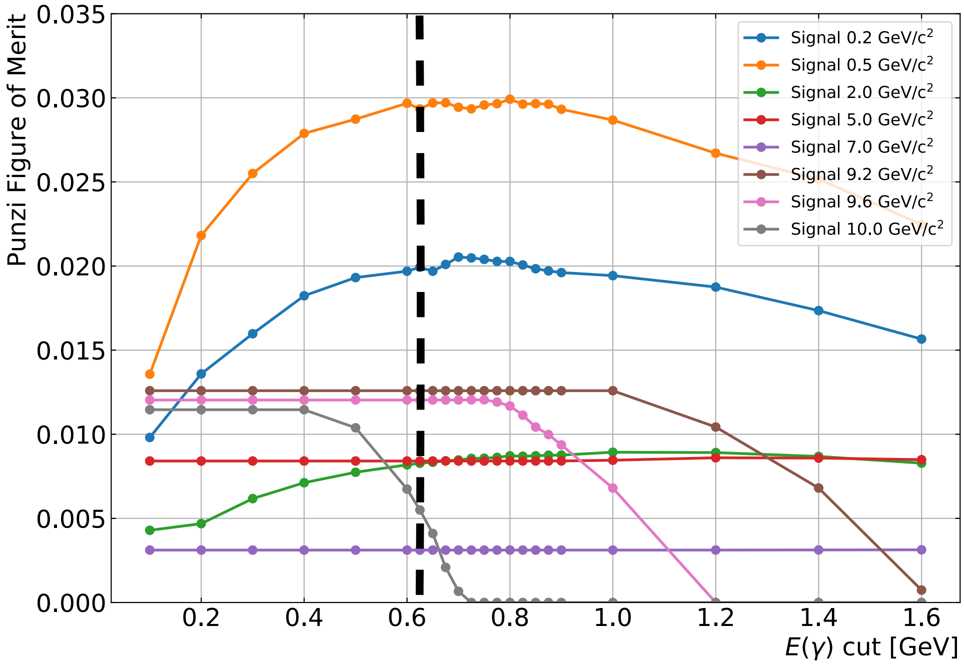

(b) After selections on other variables are applied.

Figure 5.5: Trend of the PFM as a function of the cut on the energy of the three photons constituting the event candidate.

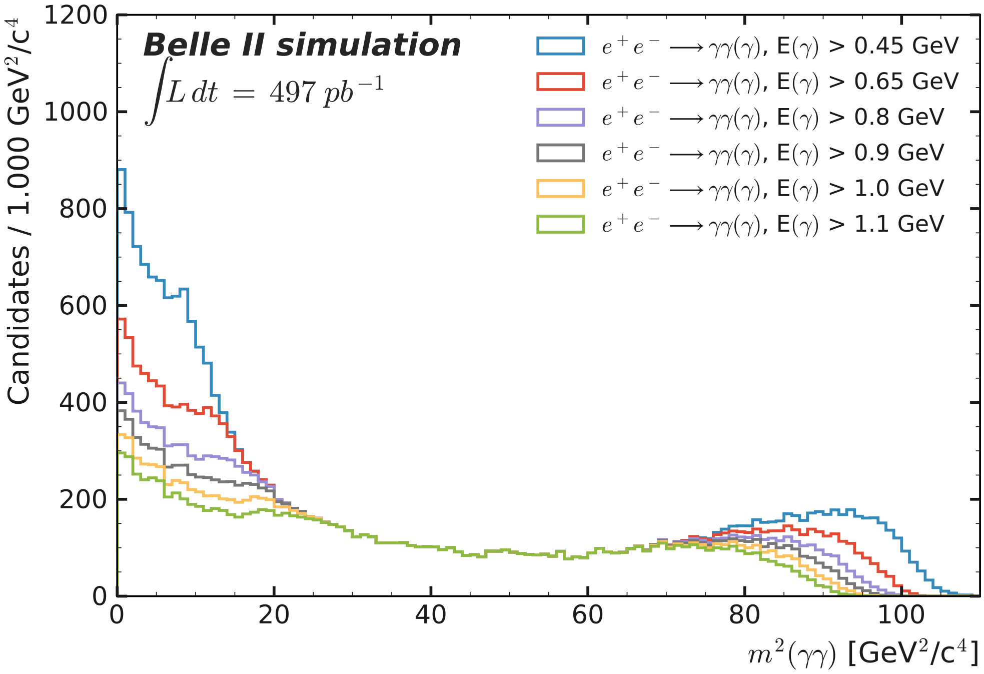

Figure 5.6: Squared diphoton mass distribution for the $e^+e^- \to \gamma\gamma\,(\gamma)$ background, for different values of the photon energy cut. The red line (650 MeV) corresponds to the selection used for $m_a > 4.0$ GeV/c^2, while the yellow line (1000 MeV) corresponds to the selection used for $m_a < 4.0$ GeV/c^2.

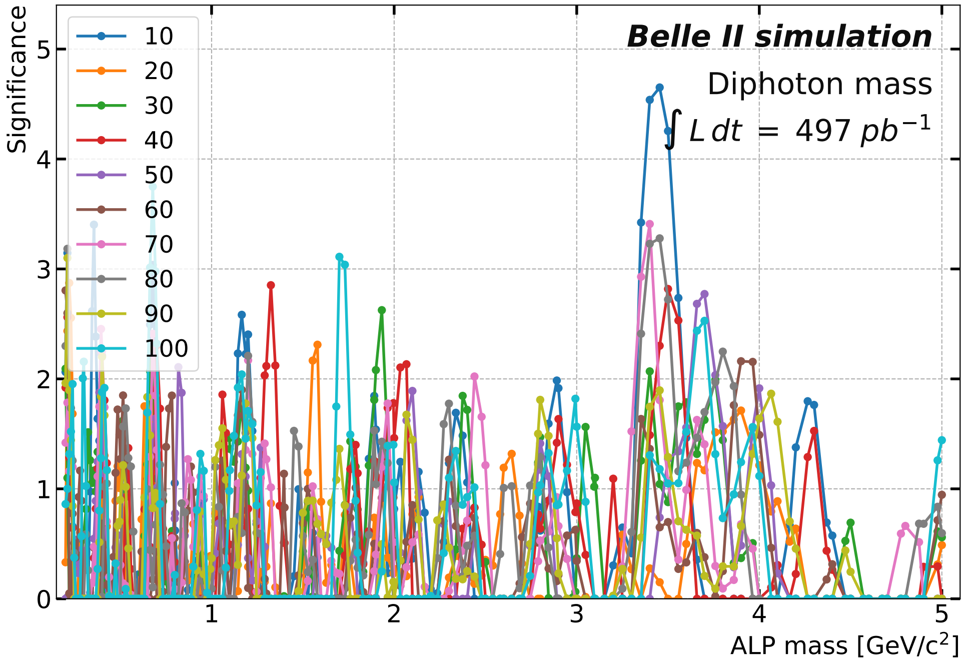

(a) $E(\gamma) > 650$ MeV.

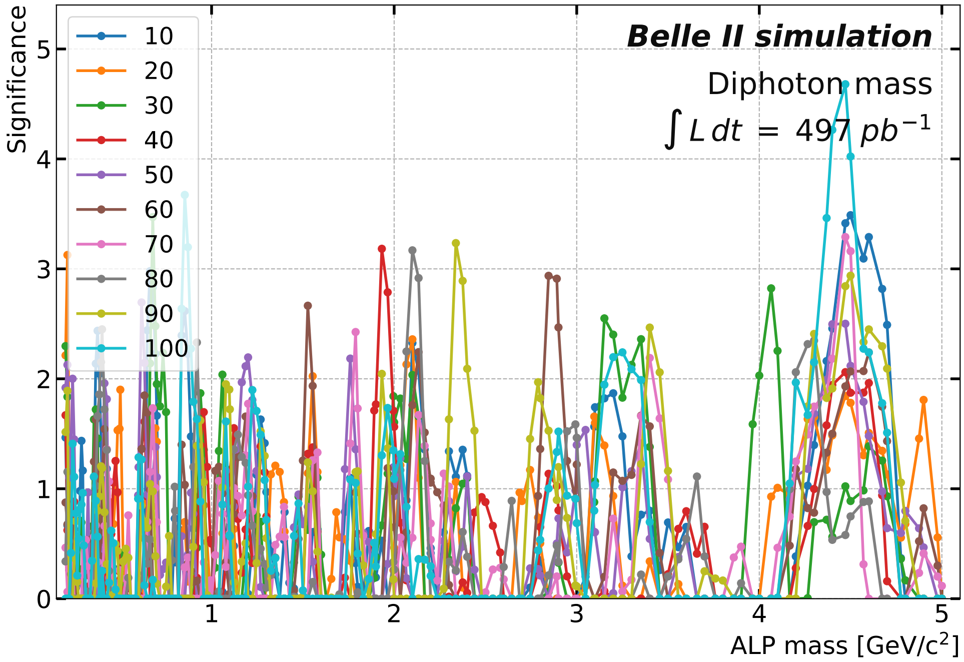

(b) $E(\gamma) > 1000$ MeV.

Figure 5.7: Significance trend, based on 10 random extractions on background-only MC (the numbers in the legend correspond to the different seeds used, see Sec. 7.1). Top: $E(\gamma) > 650$ MeV; bottom: $E(\gamma) > 1000$ MeV. Systematic excesses are present at around 3.5 and 4.5 GeV/c^2 respectively.

5.2.4 Invariant mass of the three photons

The invariant mass of the three photons is close to $\sqrt{s}$ for signal events. Applying a selection on the invariant mass is equivalent to requiring no missing energy in the event.

The three signal candidates in each event are comprised of the same three photons, thus each candidate has an identical three-photon invariant mass. Therefore, in the following plots each value enters three times, one for each candidate for each event.

Fig. 5.8 shows the distributions of the invariant mass of the three photons, and Fig. 5.9 shows the PFM studies, with the final value of the cut highlighted as well. In the Figures below, the PFM trend is shown as a function of the so-called E_{cms} *right fraction*, RF, which is used to define the cut in the following way:

$$(1 - 4 \cdot RF)\sqrt{s} \leq m_{\gamma\gamma\gamma} \leq (1 + RF)\sqrt{s}. \tag{5.5}$$

The invariant mass cut is therefore:

$$0.88\sqrt{s} \leq m_{\gamma\gamma\gamma} \leq 1.03\sqrt{s}. \tag{5.6}$$

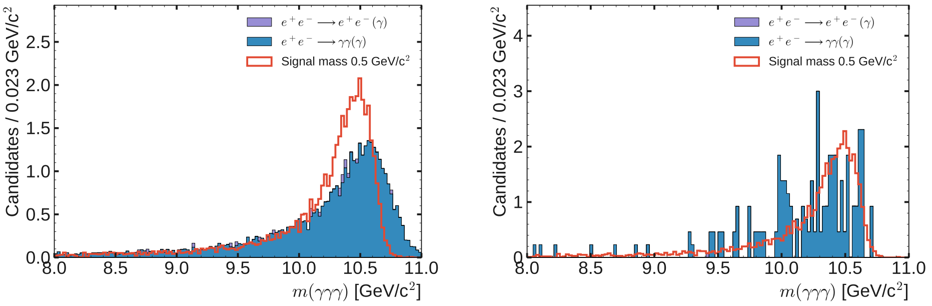

(a) $m_a = 0.5$ GeV/c^2. Loose selection. (b) $m_a = 0.5$ GeV/c^2. Final selection.

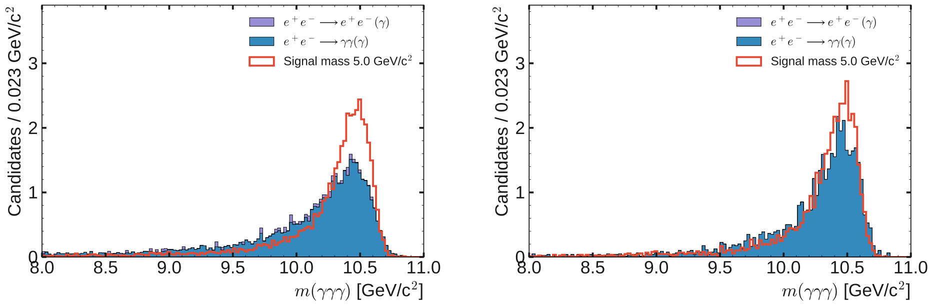

(c) $m_a = 5.0$ GeV/c^2. Loose selection. (d) $m_a = 5.0$ GeV/c^2. Final selection.

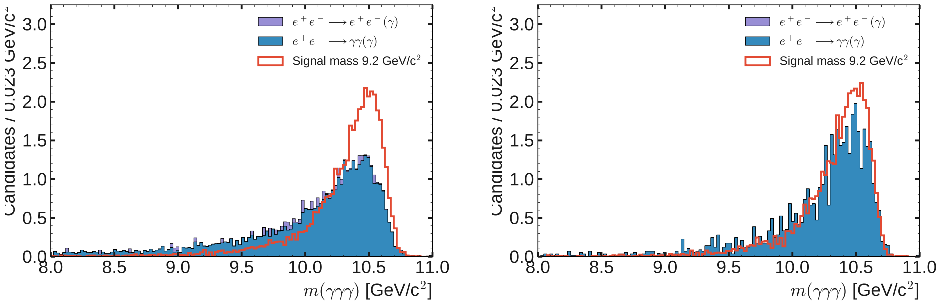

(e) $m_a = 9.2$ GeV/c^2. Loose selection. (f) $m_a = 9.2$ GeV/c^2. Final selection.

Figure 5.8: Invariant mass of the three photons system for the signals ($m_a = 0.5, 5.0, 9.2$ GeV/c^2), $e^+e^- \to e^+e^-(\gamma)$, and $e^+e^- \to \gamma\gamma(\gamma)$. Left: after applying the loose selection; right: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Normalized to unity.

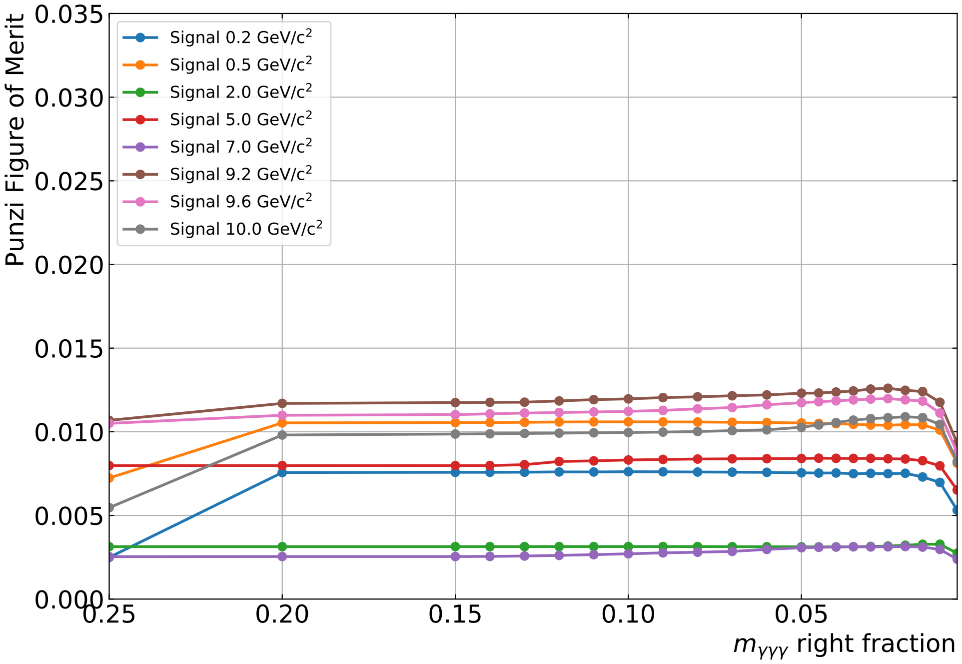

(a) Before selections on other variables are applied.

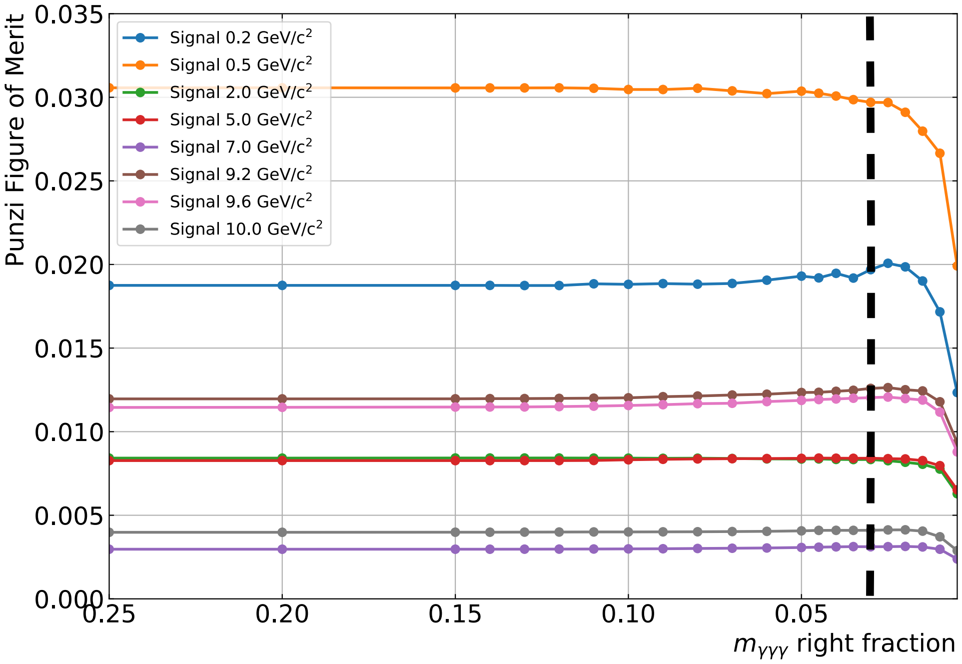

(b) After selections on other variables are applied.

Figure 5.9: Trend of the PFM as a function of the cut on the invariant of the three photons constituting the event candidate.

5.3 Variables based on data

Some additional variable cuts have been developed based on studies of control-channel data, or for robustness reasons.

5.3.1 Zernike of the most isolated photon

The MC simulation has been validated on a signal *sideband* dataset (see Sec. 6.3 for more details), which has been obtained applying a selection identical to the final one except for one cut, which has been reversed in order to keep the analysis blinded. So, the selection is the same as the one summarized in Sec. 5.5, except:

- $m_{\gamma\gamma\gamma} < 0.88\sqrt{s}$ (instead of $0.88\sqrt{s} \leq m_{\gamma\gamma\gamma} \leq 1.03\sqrt{s}$ as stated in Sec. 5.2.4);
- absence of the cut on the `ZernikeMVA` variable.

The `ZernikeMVA` variable (Sec. 3.4.2.2) has indeed been added to the selection after the studies performed on the sideband sample. The closer this variable is to 1, the more photon-like the cluster is.

The `ZernikeMVA` cut is applied only on the most isolated photon. This is because by applying a cut on the `ZernikeMVA` of *all* the three photons, the signal efficiency and the sensitivity for light ALPs would be reduced by more than an order of magnitude: light ALPs are very boosted, thus the two photons deriving from their decay are very close to each other, hence the two clusters share some ECL crystals, and so the `ZernikeMVA` of both of these photons have lower values. The most isolated photon is defined as the one for which the sum in quadrature of the distances in polar and azimuthal angle with respect to the other two is maximum. Applying the cut only on the most isolated photon still eliminates many unmodeled background events in the sideband region, which are mostly comprised of three non-photon-like clusters, without reducing the efficiency for signal events.

The `ZernikeMVA` cut is therefore:

$$\texttt{ZernikeMVA}\text{ of the most isolated photon} > 0.6. \tag{5.7}$$

5.3.2 Time variable

A cluster-timing variable has been defined to guard against potential sources of background where the three photon candidates are out-of-time with respect to each other: in signal events, the three photons should be simultaneous. The variable is the *maximum weighted distance from the average ECL time of the three clusters*, defined as:

$$t_c = max\left\{\left|\frac{t_j - <t>}{dt99_j}\right|\right\}_{j=0,1,2}, \tag{5.8}$$

where t_j is the `clusterTime` of the j-th photon candidate, $dt99_j$ its associated `deltaTime99` (Sec. 3.4.2.2), and $<t>$ is defined as it follows:

$$<t> = \frac{\sum\limits_{j=0}^{2} \frac{t_j}{dt99_j^2}}{\sum\limits_{j=0}^{2} \frac{1}{dt99_j^2}}. \tag{5.9}$$

Note that t_c is dimensionless. This variable has large values if there is a photon candidate which is highly out-of-time with respect to the average time of the three photons. After all other cuts have been applied, a cut on t_c is not found to be particularly powerful in MC. However, it is included in the selection to protect from backgrounds not modeled in the simulation. This is motivated by a data-driven study described in Sec. 6.1. Given that this variable is not included due to studies performed on MC samples, its MC distributions for backgrounds and signal are not shown, and neither are the PFM studies.

The t_c cut is therefore:

$$t_c < 10. \tag{5.10}$$

5.3.3 Angular differences

The angular separation between the two photons that constitute the ALP candidate is part of the selection. Pair conversion, $\gamma \to e^+e^-$, can be induced by the electric fields of the materials through which the photons travel. If this happens outside the tracking volume, the clusters associated with the e^+e^- pair cannot be matched to charged tracks, hence they are misidentified as photons. The two charged particles are still affected by the magnetic field, but while the angular separation $\Delta\theta$ in the polar angle will stay small, the separation $\Delta\phi$ in the azimuthal angle can be quite large. These variables are not found to be particularly powerful on simulation, although, they are kept in the selection in order to make the analysis robust against pair conversion events from outside tracking, as pair conversions may not be correctly modeled in the MC.

This is the only used variable used in the selection that is candidate-dependent and not event-dependent. In case one of the three candidates is rejected, then the entire event, i.e. all the three candidates, is rejected.

Figs. 5.10, 5.11, and 5.12 show the distributions for the three sample signal masses and the stacked distributions for the main backgrounds.

The $\Delta\theta - \Delta\phi$ cut is therefore:

$$\Delta\theta > 0.014 \text{ rad or } \Delta\phi > 0.40 \text{ rad.} \tag{5.11}$$

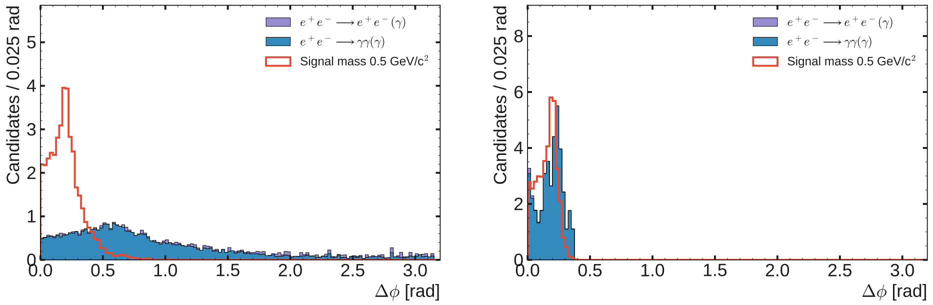

(a) $\Delta\phi$. Loose selection.

(b) $\Delta\phi$. Final selection.

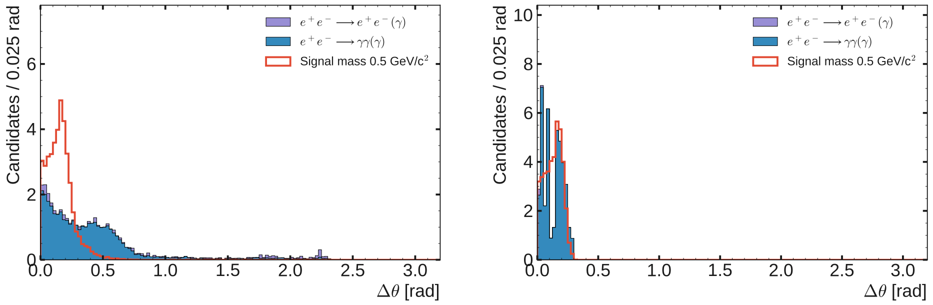

(c) $\Delta\theta$. Loose selection.

(d) $\Delta\theta$. Final selection.

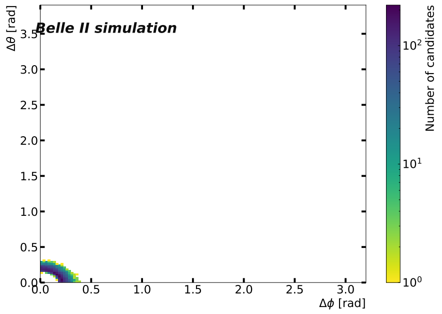

(e) $\Delta\theta$ as a function of $\Delta\phi$. Final selection.

Figure 5.10: $\Delta\theta$ and $\Delta\phi$ distributions for signal ($m_a = 0.5$ GeV/c^2), $e^+e^- \to e^+e^-(\gamma)$, and $e^+e^- \to \gamma\gamma(\gamma)$. Left: after applying the loose selection; right and bottom: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Top four plots normalized to unity. The 2D plot is for the TM signal only.

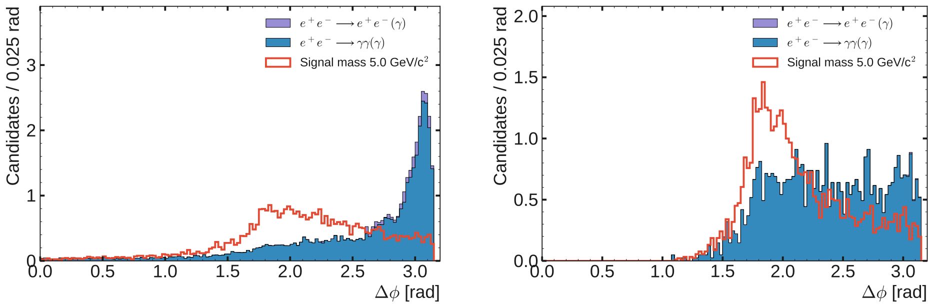

(a) $\Delta\phi$. Loose selection.

(b) $\Delta\phi$. Final selection.

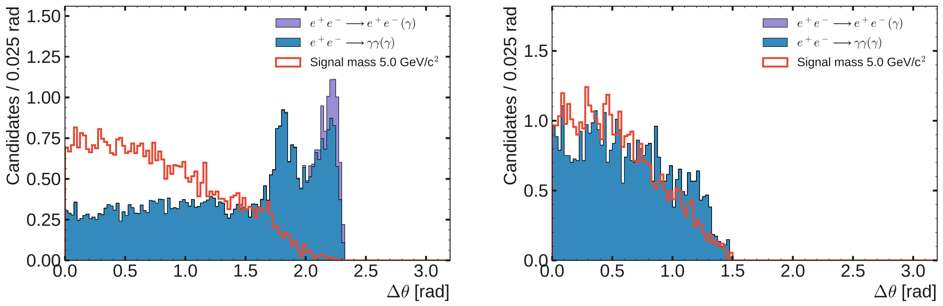

(c) $\Delta\theta$. Loose selection.

(d) $\Delta\theta$. Final selection.

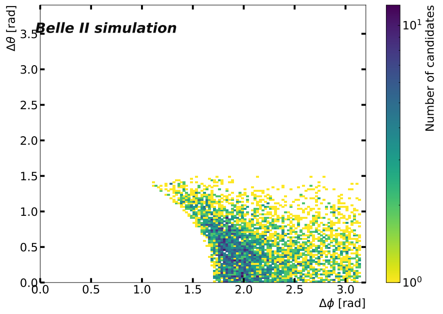

(e) $\Delta\theta$ as a function of $\Delta\phi$. Final selection.

Figure 5.11: $\Delta\theta$ and $\Delta\phi$ distributions for signal ($m_a = 5.0$ GeV/c^2), $e^+e^- \to e^+e^-(\gamma)$, and $e^+e^- \to \gamma\gamma(\gamma)$. Left: after applying the loose selection; right and bottom: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Top four plots normalized to unity. The 2D plot is for the TM signal only.

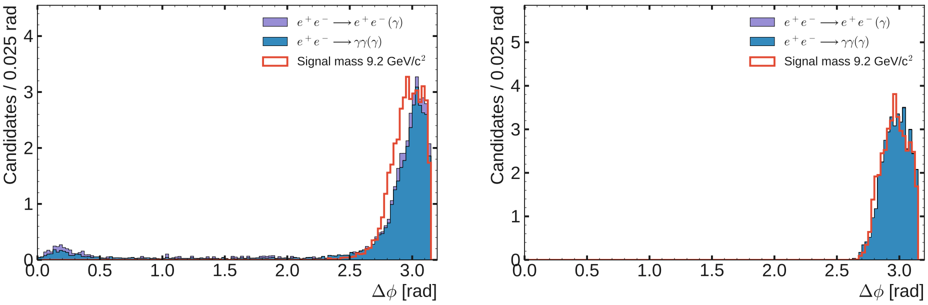

(a) $\Delta\phi$. Loose selection.

(b) $\Delta\phi$. Final selection.

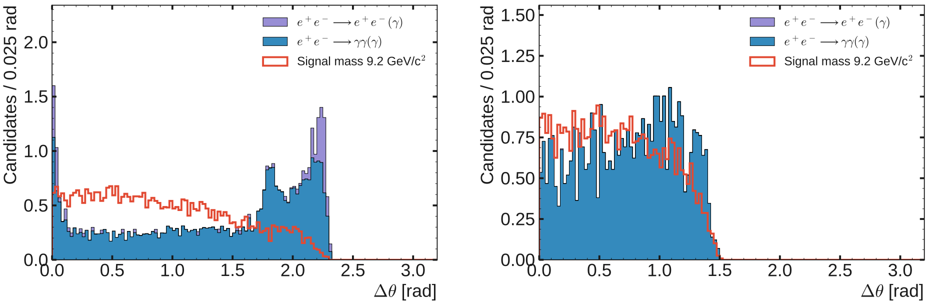

(c) $\Delta\theta$. Loose selection.

(d) $\Delta\theta$. Final selection.

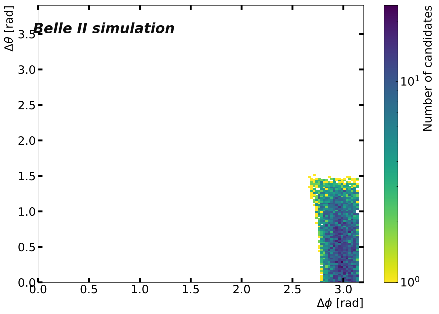

(e) $\Delta\theta$ as a function of $\Delta\phi$. Final selection.

Figure 5.12: $\Delta\theta$ and $\Delta\phi$ distributions for signal (m_a = 9.2 GeV/c^2), $e^+e^- \rightarrow e^+e^-(\gamma)$, and $e^+e^- \rightarrow \gamma\gamma(\gamma)$. Left: after applying the loose selection; right and bottom: after the final selection, except for the cut on the plotted quantity. Signal ALP candidates are TM. Top four plots normalized to unity. The 2D plot is for the TM signal only.

5.3.4 Robustness variables

Based on recommendations from neutral and tracking internal working groups, two extra cuts are added to the selection to make it robust against background.

An ECL cluster can be formed by any number of crystals. Clusters made up only by one single crystal are generally caused by electric noise in the readout electronic connected to the crystal itself. To avoid this background, it is required that `clusterNHits` > 1.5 (Sec. 3.4.2.2). Non-integer amount of crystals may happen when a crystal is shared between overlapping clusters.

Events with *good track*s are vetoed. A good track is defined as having a non-zero number of hits in the CDC and having its origin close enough to the IP, meaning no more than 1 cm in the radial direction ($|d0| < 1.0$ cm) and 5 cm in the beam axis direction ($|z0| < 5$ cm). Events with tracks that do not fulfill these requirements (*bad tracks*) are allowed, as most bad tracks originate from beam background, which would also be present in true ALP signal events.

5.4 Trigger efficiency

Given that the final state of this analysis is fully neutral, it relies solely on ECL-based L1 triggers (Sec. 3.4.3). No vetoes were applied during Phase 2, so it is not necessary to select a specific trigger channel, but it is possible to just include all events in this analysis. Consideration of the HLT is not necessary, as it was not used to reject events in Phase 2[1].

For reasons that are not fully understood (possibly issues in the data acquisition), some data-taking runs have a large fraction of their events without any available L1 trigger information. There is no indication that events without any L1 triggers are correlated with a certain trigger line. Only runs that have L1 trigger information for at least 95% of their events are used for this study.

The ALP signal event signature consists of three high-energy neutral ECL clusters, whose energies sum up to the beam energy. Trigger efficiency is measured by using a non-unblinding control sample of radiative Bhabha events $e^+e^- \to e^+e^-\gamma$. A method of orthogonal triggers is deployed. In a sample where all the events should be selected by both triggers, the number of events selected by a CDC-only trigger and those that are selected by both ECL and

---

[1] The HLT was in so-called *monitoring* or *passthrough* mode.

CDC triggers are counted, and treated as denominator and numerator respectively in the following efficiency definition:

$$\varepsilon = \frac{\text{triggered by ECL and CDC}}{\text{triggered by CDC}}. \tag{5.12}$$

The data sample is made up by events with two tracks, each associated with an ECL cluster, and one additional neutral cluster, with the following selections (in the following, `clusterE`[`Theta`] is the energy [polar angle] of the cluster, as reconstructed by the ECL only). Tracks are selected if $p > 1.0$ GeV, $|dz| < 4.0$ cm and $|dr| < 1.0$ cm. The ECL clusters associated with the tracks must satisfy `clusterE` > 0.65 GeV and $0.651 <$ `clusterTheta` $<$ 2.159 rad. Events must contain an additional neutral ECL cluster with `clusterE` > 0.65 GeV and $0.651 <$ `clusterTheta` < 2.159 rad. The invariant mass of the three particles must be such that $0.8\sqrt{s} < m_{ee\gamma} < 1.1\sqrt{s}$. If more than one $ee\gamma$ candidate passes this selection, the whole event is discarded.

The trigger efficiency ε has been calculated, with two specific trigger lines[2], as function of the invariant mass $m(e\gamma\gamma)$ (see Fig. 5.13a), and as function of the sum of the three used ECL cluster energies (see Fig. 5.13b). Out of 18 244 events that activate the CDC trigger, all events are also selected by the ECL L1 trigger. This results in a trigger efficiency of

$$\varepsilon = 1^{+0}_{-6\times10^{-5}}. \tag{5.13}$$

To conclude, one of the active trigger lines is found to be 100% efficient with a negligible statistical uncertainty.

---

[2] In particular, trigger *ff* as a fully efficient CDC trigger, and trigger *lume* as a fully efficient ECL trigger.

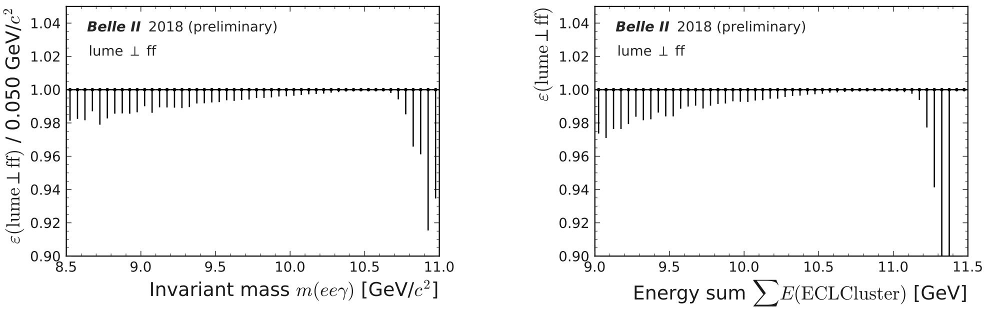

(a) As a function of the invariant mass $m(e\gamma\gamma)$.

(b) As a function of the sum of the three cluster energies.

Figure 5.13: Trigger efficiency as function of invariant mass $m(e\gamma\gamma)$ and sum of the three used ECL cluster energies of the selected radiative electron-pair events.

5.5 Selection summary and performances

The following summarizes the whole selection:

- $E_\gamma \geq 1000$ MeV if $m_a < 4$ GeV/c^2, $E_\gamma \geq 650$ MeV if $m_a \geq 4$ GeV/c^2;
- $0.88\sqrt{s} \leq m_{\gamma\gamma\gamma} \leq 1.03\sqrt{s}$, $\sqrt{s} = 10.58$ GeV/c^2;
- $\Delta\theta \geq 0.014$ rad or $\Delta\phi \geq 0.40$ rad, with $\Delta\theta$ and $\Delta\phi$ computed between the two photons constituting the ALP candidate;
- $\theta_\gamma \in [37.3°, 123.7°]$, that correspond to the ECL barrel area without the first and last 5°;
- The absence of any *good track*, defined as tracks with more than zero hits in the CDC, $|d0| < 1.0$ cm, and $|z0| < 5$ cm is required;
- ECL clusters must have at least 1.5 crystals;
- `ZernikeMVA` > 0.6 for the most isolated photon;
- $t_c < 10$, i.e. the three photon must be simultaneous;
- Only the three most energetic photons are considered (additional photons in the event are ignored).

In Fig. 5.14, the signal efficiency for TM events and the number of background events are reported as a function of the true ALP mass.

To compute these figures, the number of signal and background events at a given ALP mass point must be *counted*. In order to do so, it is necessary to define a *mass window*, the width of which depends on the mass hypothesis. Two different mass windows are defined: one in terms of the diphoton mass, the other in terms of the recoil mass. The full definitions are discussed in Sec. 7.3.

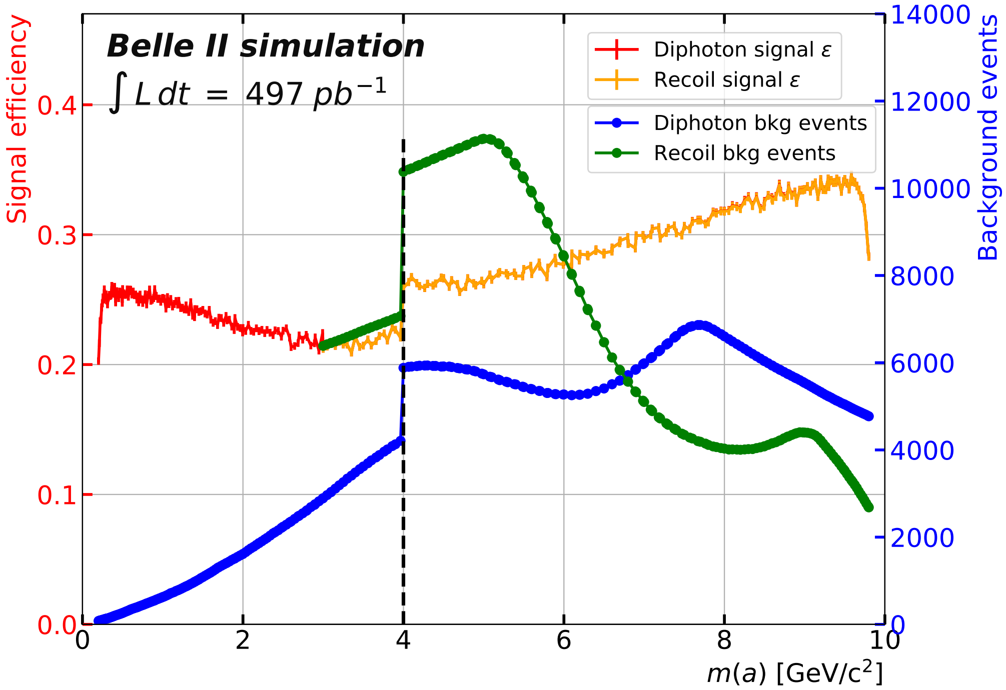

Figure 5.14: Signal efficiency, computed with TM events only, and number of background events, for diphoton and recoil mass definition, as a function of the ALP mass. The number of background events is computed using the fit windows used for the upper limit computation, see Chap. 7. The vertical dashed line corresponds to the mass at which the energy cut passes from 1000 MeV to 650 MeV.

In Figs. 5.15 and 5.16 there are the diphoton and recoil squared mass distributions for all of the background sources considered.

The background from radiative Bhabha $e^+e^- \to e^+e^-(\gamma)$ is included with the caveat discussed in Sec. 4.3.2: a 2.5% tracking inefficiency per track is assumed, and only events where both tracks are not reconstructed are included.

The background from $e^+e^- \to S\gamma$, $S = \pi^0, \eta, \eta\prime$ is shown in Fig. 5.17. The second peak in Fig. 5.17a is at the ω resonance mass, and comes mainly from the π^0 sample. This is a feature of the PHOKARA generator [61]: it provides samples with a π^0 and at least one photon in the final state, regardless of the intermediate states. One channel to produce such a final state is $e^+e^- \to \omega\gamma$, $\omega \to \pi^0\gamma$. The π^0 from the decay of the ω is highly boosted, and the two photons coming from its decay are so close to each other that are often reconstructed as one single photon. These two photons, reconstructed as a single cluster interpreted as a single photon, combined with the other photon from the ω decay, are reconstructed as an ALP candidate. This ALP candidate has the invariant mass of the ω, and this explains the second peak in Fig. 5.17a.

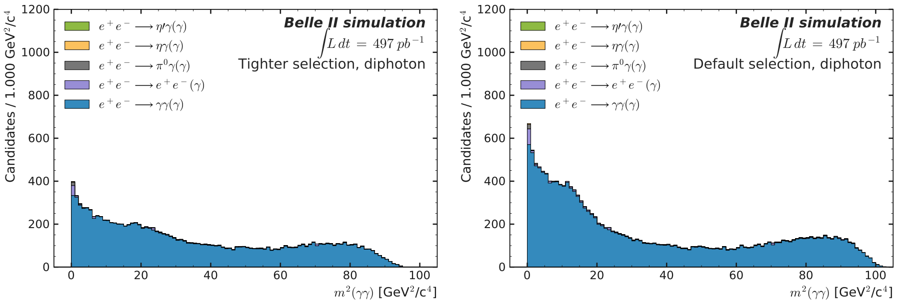

(a) Selection used below $m(a) = 4.0$ GeV/c^2, i.e. with $E(\gamma) > 1000$ MeV.

(b) Selection used above $m(a) = 4.0$ GeV/c^2, i.e. with $E(\gamma) > 650$ MeV.

Figure 5.15: Diphoton squared mass distribution, stacked background, both selections.

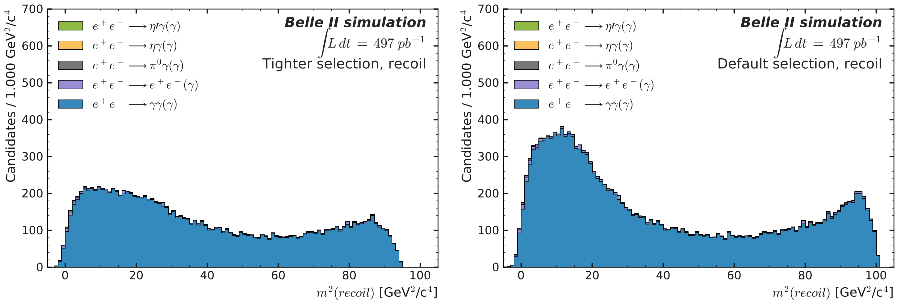

(a) Selection used below $m(a) = 4.0$ GeV/c^2, i.e. with $E(\gamma) > 100$ MeV.

(b) Selection used above $m(a) = 4.0$ GeV/c^2, i.e. with $E(\gamma) > 650$ MeV.

Figure 5.16: Recoil squared mass distribution, stacked background, both selections.

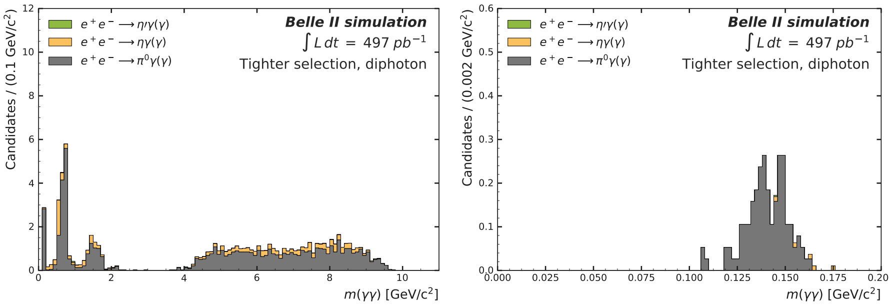

(a) All mass spectrum.

(b) Zoom around π^0 mass.

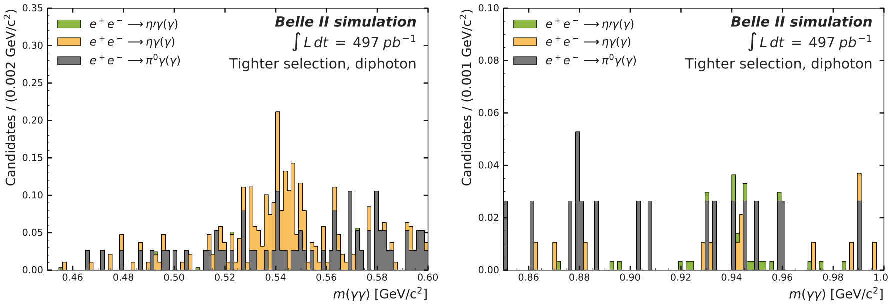

(c) Zoom around η mass.

(d) Zoom around $\eta\prime$ mass.

Figure 5.17: Background from $e^+e^- \to S\gamma$, $S = \pi^0, \eta, \eta\prime$, after final selection.

Chapter 6

MC validation

In this Chapter, the studies performed to validate the goodness of the MC samples and MC-based procedures are shown. This has been done while avoiding full unblinding. This can be achieved by using orthogonal data sets (which must contain good tracks and/or fewer than three ECL clusters), sidebands, or, at the end, by performing only a *partial* unblinding.

Section 6.1 details a study performed on a sample made up of events with fewer than three ECL clusters which would pass the signal selection, and subsequently on a sample of radiative Bhabha events, which is used to demonstrate the usefulness of a cut on a timing variable, as already described in Sec. 5.3.2.

Section 6.2 describes a simulated signal extraction using radiative Bhabha events.

Section 6.3 shows the results of the studies performed on the sidebands, which resulted in the addition of a cut to the selection, as anticipated in Sec. 5.3.1.

Section 6.4, finally, illustrates the results of a 10% preliminary partial unblinding, performed to check the goodness of the MC simulation.

6.1 Timing from $e^+e^- \to \gamma(\gamma)$ and $e^+e^- \to e^+e^-\gamma$ events

A dataset orthogonal to the signal sample, made up of events with up to two neutral clusters, has been studied to check in data the presence of out-of-time and high-energy clusters not present in MC. These clusters may enhance the SM background from $e^+e^- \to \gamma\gamma$ where a third photon in the event comes from an unknown background source.

All events with one or two photon candidates with `clusterE` > 250 MeV within the CDC acceptance are saved for this study. The presence or absence of tracks is not required. The `E1oE9` variable (Sec. 3.4.2.2), energy, polar angle θ, and azimuthal angle ϕ as functions of cluster time are shown in Fig. 6.1. There is a clear indication for the presence of photon candidates at around $t \approx -500$ ns.

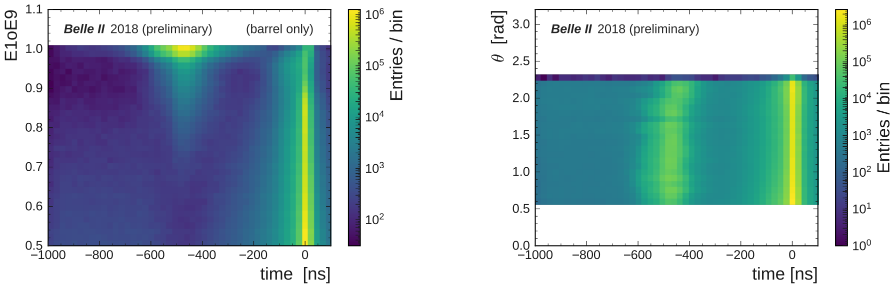

(a) E1oE9 variable.

(b) Polar angle θ.

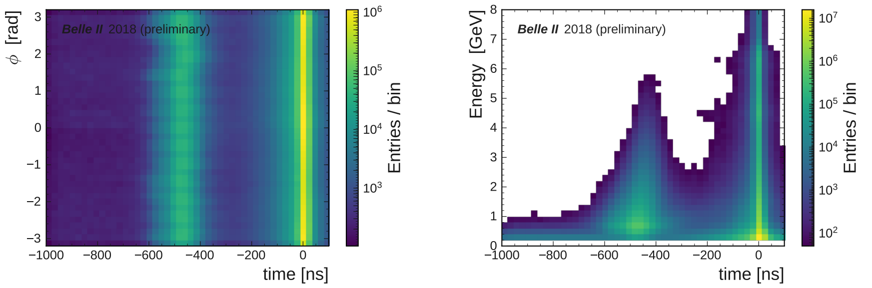

(c) Azimuthal angle ϕ.

(d) Energy.

Figure 6.1: `E1oE9`, energy, polar angle, and azimuthal angle as functions of cluster time. Only clusters in the ECL barrel.

These *anomalous events* with highly negative timing activated single-photon L1 triggers with a high rate, supporting the hypothesis that they are not correlated with physics events, because such a high number of out-of-time single-photon events are not expected from any known physical source. The most supported hypothesis is that these are events where a neutron directly interacts with one of the photomultipliers connected to an ECL crystal, which can result in an output interpreted as an high-energy, out-of-time ECL cluster. These anomalous events occur with an high rate. Therefore, even if most of them are triggered by single-photon triggers, it is possible for a non-negligible fraction of them to overlap with actual physics events.

This possible contamination has been studied using $e^+e^- \to e^+e^-\gamma$ events, selected by requiring exactly two oppositely charged tracks, both matched to an ECL cluster with `clusterE` > 3 GeV. Within this sample, events are labeled as *anomalous* if they have at least one additional neutral cluster with

$$\texttt{E1oE9} > 0.8,\ \texttt{clusterE} > 0.1\ \text{GeV},\ -800 < t < -300\ \text{ns}, \qquad (6.1)$$

and are labeled *normal* if they have at least one additional neutral cluster with

$$\texttt{E1oE9} < 0.8,\ \texttt{clusterE} > 0.1\ \text{GeV},\ t > -300\ \text{ns}. \qquad (6.2)$$

Di-photon events with one or more anomalous neutral clusters could constitute a background for the ALP search. They can be rejected using the t_c variable (see Sec. 5.3.2). The distribution of the variable t_c for events with and without anomalous clusters, within the radiative Bhabha sample defined above, is shown in Fig. 6.2.

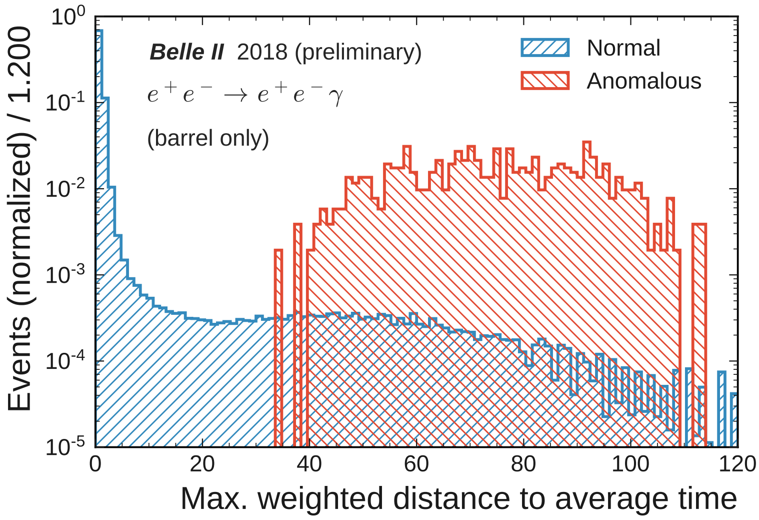

Figure 6.2: Maximum weighted distance from the average cluster time t_c for control-channel Bhabha candidates in data with an additional neutral cluster.

6.2 Emulated signal selection with $e^+e^- \to e^+e^-\gamma$

Radiative Bhabha $e^+e^- \to e^+e^-\gamma$ events are used in order to validate the whole selection procedure and reconstruction chain. These events and the ALP ones share a similar topology, as both of them present three high-energy clusters whose energies sum up to the beam energy $\sqrt{s}$. Although, radiative Bhabha and ALP events are two completely orthogonal data sets, because for the former it is required the presence of two good tracks, which is instead vetoed for the latter. Hence, the radiative Bhabha sample can be fully reconstructed without unblinding. The selection is the following:

- Only the most energetic photon and the two good tracks with the highest associated cluster energy are considered to reconstruct the event;
- `clusterE` ≥ 650 MeV for each particle;
- $0.88\sqrt{s} \leq m_{e^+e^-\gamma} \leq 1.03\sqrt{s}$;
- $\Delta\theta \geq 0.003$ rad or $\Delta\phi \geq 0.30$ rad, with $\Delta\theta$ and $\Delta\phi$ computed between the two particles (this is a proxy to the ALP candidate in the signal events);
- `clusterTheta` $\in [37.3°, 123.7°]$.

The selection is slightly different than the one reported in Sec. 5.5 because is based on early optimization studies. The idea is to treat the electron tracks as photons, so that they look like a three-photon final state with no missing energy, exactly as for the studied signal process. To better emulate the signal reconstruction procedure, ECL variables (energy and angles) are used for the $e^\pm$, instead of the ones coming from the tracking algorithm. Similarly, $m_{\gamma\gamma\gamma}$ is computed using ECL variables rather than the tracking ones.

This selection has been applied to early available data and MC. The MC has been scaled to the luminosity available in data[1], 472 pb^{-1}, and the comparisons are shown in Fig. 6.3.

The agreement in terms of both absolute normalization and variable shape is excellent.

---

[1] Only runs where the proper trigger information for this kind of process was available were used.

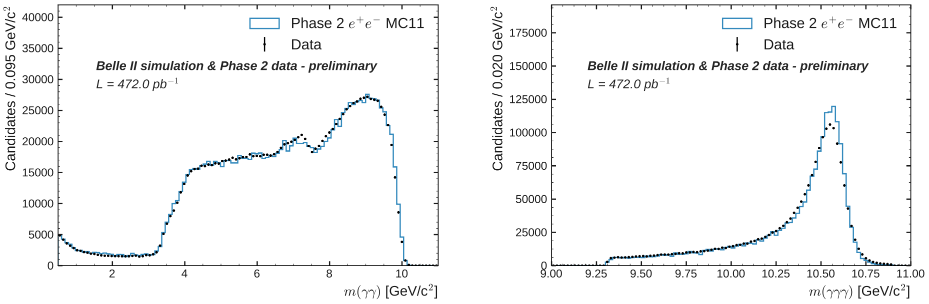

(a) Mass of the ALP candidates. Distributions start from 0.5 GeV/c^2.

(b) Invariant mass of the three-particle system.

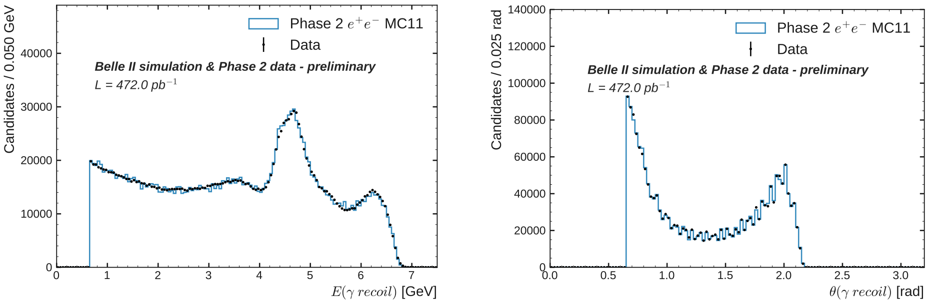

(c) `clusterE` of the recoil photon.

(d) `clusterTheta` of the recoil photon.

Figure 6.3: Variables distribution, data and MC comparison, for the radiative Bhabha selection.

6.3 Signal sideband

As anticipated in Sec. 5.3.1, a sideband study has been performed. From this study, the importance of an additional cut on `ZernikeMVA` has been observed.

The sideband selection differs from the final signal selection (Sec. 5.5) only for the reversed $m_{\gamma\gamma\gamma}$ cut and for the absence of the cut on `ZernikeMVA`:

- $E_\gamma \geq 1000$ MeV if $m_a < 4$ GeV/c^2, $E_\gamma \geq 650$ MeV if $m_a \geq 4$ GeV/c^2;
- $m_{\gamma\gamma\gamma} < 0.88\sqrt{s}$, $\sqrt{s} = 10.58$ GeV/c^2;
- $\Delta\theta \geq 0.014$ rad or $\Delta\phi \geq 0.40$ rad, with $\Delta\theta$ and $\Delta\phi$ computed between the two photons constituting the ALP candidate;
- $\theta_\gamma \in [37.3°, 123.7°]$, that correspond to the ECL barrel area without the first and last 5°;

- The absence of any *good track*, defined as tracks with more than zero hits in the CDC, $|d0| < 1.0$ cm, and $|z0| < 5$ cm is required;
- ECL clusters must have at least 1.5 crystals;
- $t_c < 10$, i.e. the three photon must be simultaneous;
- Only the three most energetic photons are considered (additional photons in the event are ignored).

The `ZernikeMVA` cut has little to no effect on MC-only studies, but it helps in reducing unmodeled background in the sideband region, thus makes the selection more robust against possible contamination. The distribution of `ZernikeMVA` for the sideband sample is shown in Fig. 6.4. The distribution of $m_{\gamma\gamma\gamma}$ for the sideband sample before and after a cut on `ZernikeMVA` is shown in Fig. 6.5. The goal was to make the data-MC agreement for the highest $m_{\gamma\gamma\gamma}$, i.e. the closest to the signal region, as good as possible.

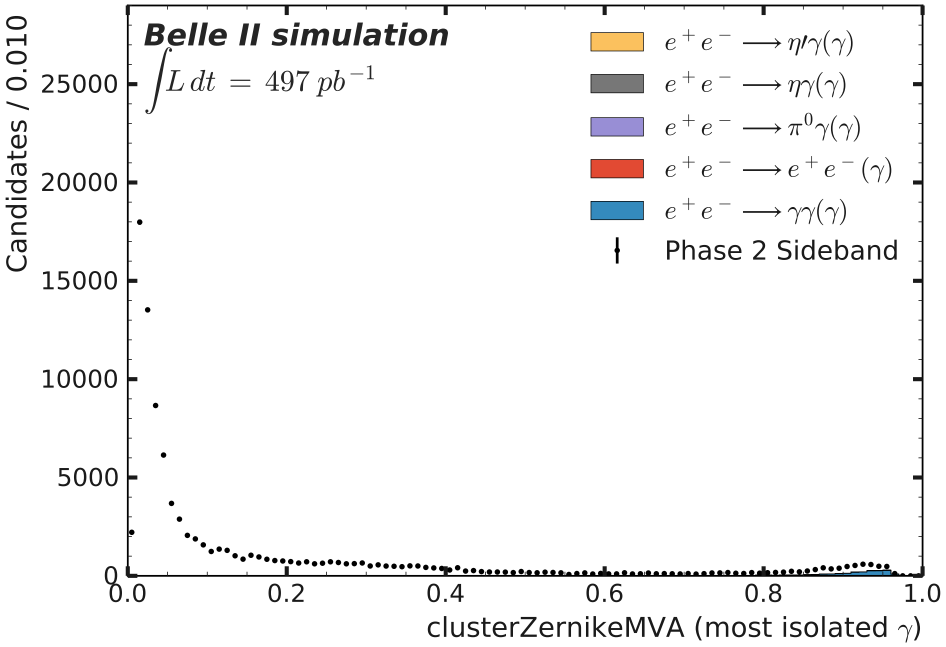

(a) Normalized to luminosity.

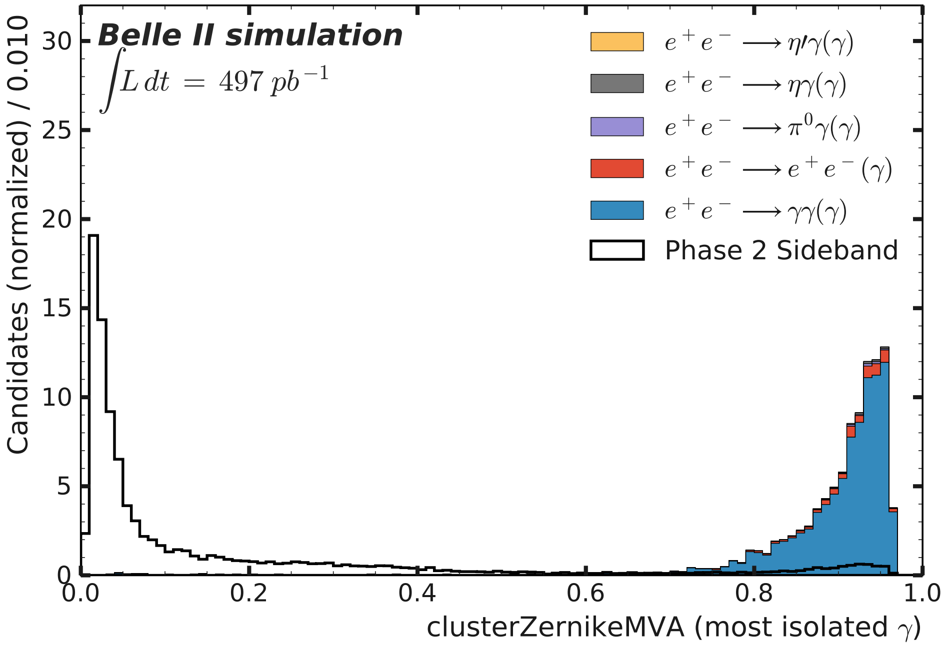

(b) Normalized to unity.

Figure 6.4: Distribution of `ZernikeMVA` in the sideband region. The second plot is normalized to unity.

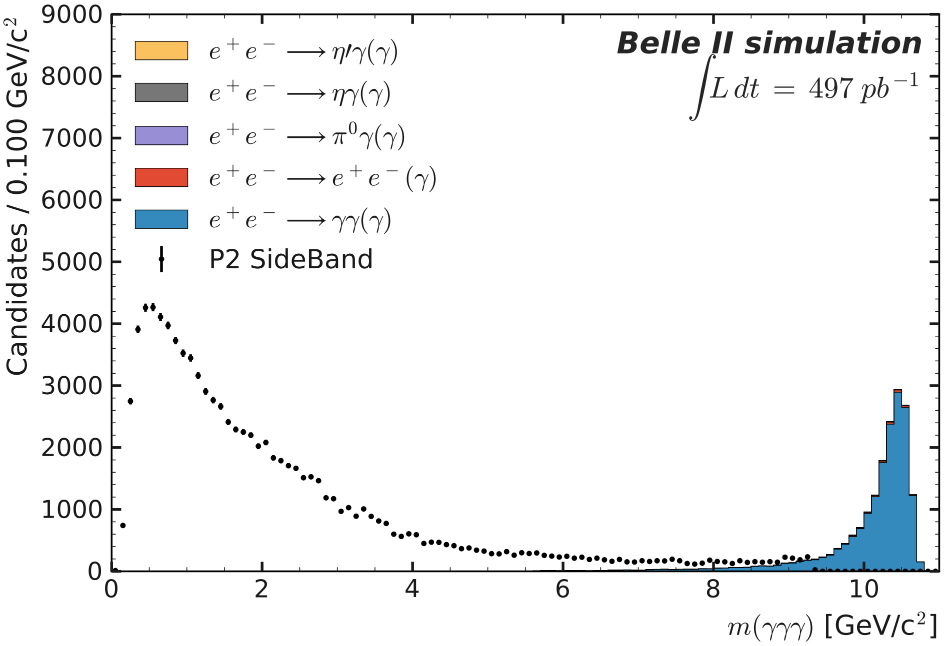

(a) Without cut on `ZernikeMVA`. MC signal region is shown too.

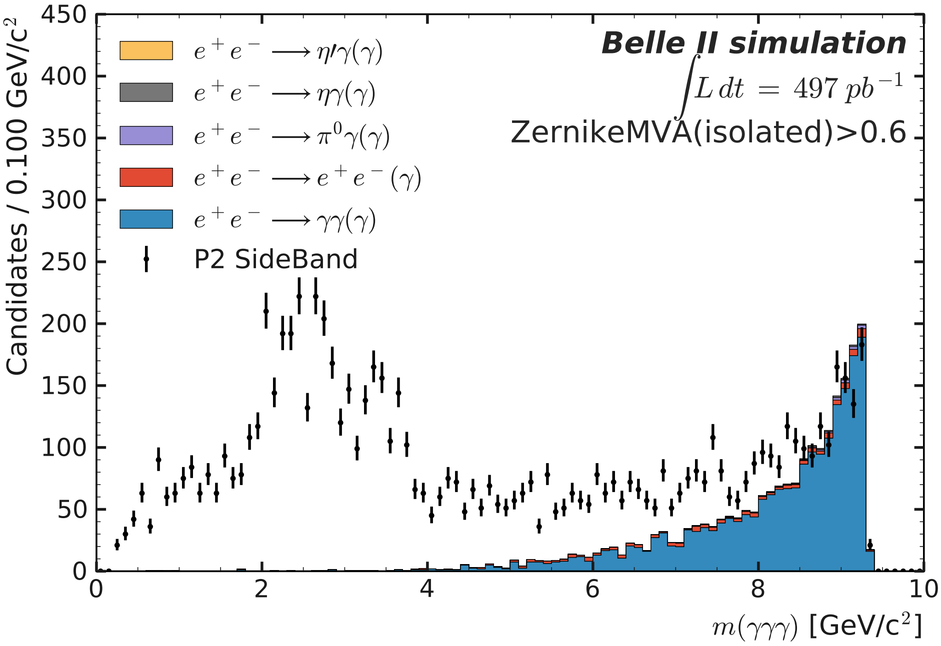

(b) With cut on `ZernikeMVA`.

Figure 6.5: Distribution of $m_{\gamma\gamma\gamma}$ in the sideband region. The first plot, with no cut on `ZernikeMVA`, shows also the MC distribution in the signal region. The second plot shows the distribution after applying the cut on `ZernikeMVA`. Note the different ranges of the x axes.

6.4 Partial unblinding

Prior to the full data set unblinding, one tenth of the final dataset was used to perform a partial unblinding, in order to check if any major data-MC discrepancy was present. The events were chosen at random to avoid possible bias related to certain data-taking periods. Due to random number generation, the partially unblinded data sample was, in fact, 10.4% of the data. The data unblinded for this study are not used for the actual ALP search.

1836 [1326] data events pass the selection with the 650 [1000] MeV photon energy cut. 1772 [1298] MC events, after scaling to 10.4% of the total luminosity, i.e. 52 pb^{-1}, pass the selection with the 650 [1000] MeV photon energy cut. In both cases, the agreement is within 1.5σ.

All photon energies are corrected with the bias correction factor described in Sec. 7.5.1. This affects the data distributions of photon energies, ALP candidates' mass, and $m_{\gamma\gamma\gamma}$, but not the MC distributions. All MC photon energies are smeared with the resolution smearing factor described as well in Sec. 7.5.1. Figs. 6.6 to 6.15 show the data-MC comparison for some of the variables of interest. With the exception of t_c and `clusterNHits`, all variables show excellent agreement between data and MC. The exceptions are somewhat expected as timing is hard to model in simulation and `clusterNHits` was known to be mismodeled. The discrepancies in these variables are far from the cut values applied, therefore they are of no concern for this analysis.

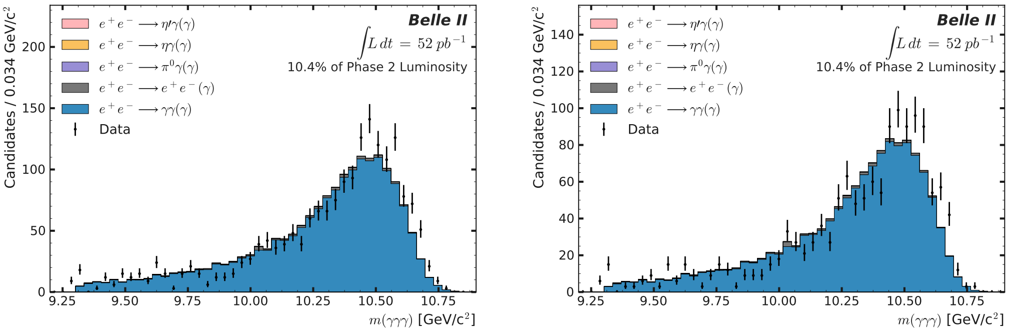

(a) E(γ)>650 MeV.

(b) E(γ)>1000 MeV.

Figure 6.6: Invariant mass $m_{\gamma\gamma\gamma}$ for data and MC in the unblinded sub-dataset.

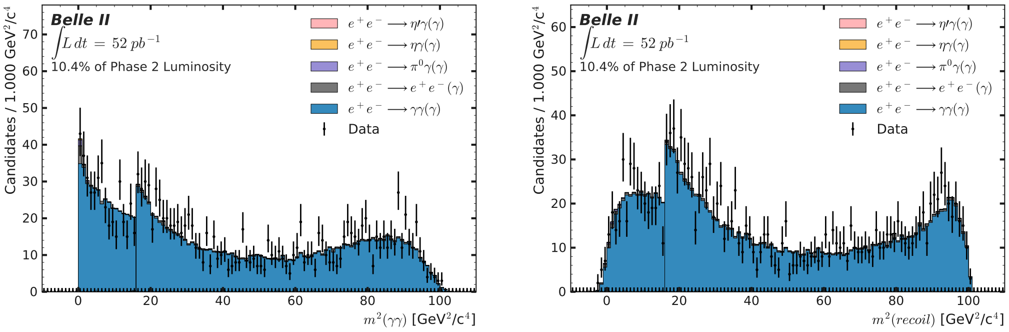

(a) Diphoton mass $m^2_{\gamma\gamma}$.

(b) Recoil mass m^2_{recoil}.

Figure 6.7: Squared invariant masses (left: diphoton, right: recoil) for data and MC in the unblinded sub-dataset. Both energy selection shown, with selection change point at 16 GeV^2/c^4.

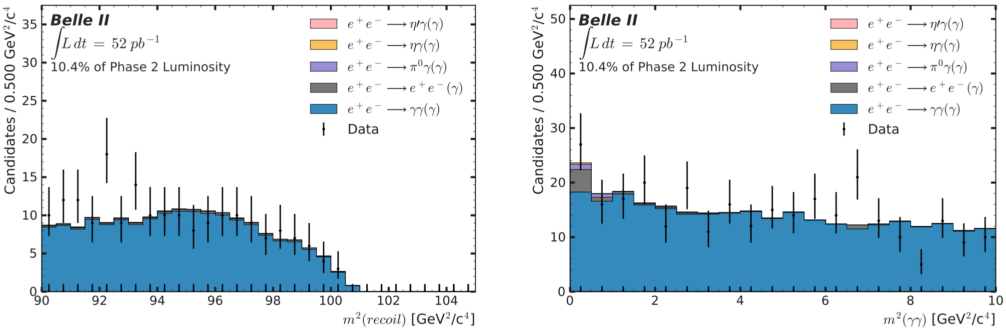

(a) Squared recoil mass endpoint. E(γ)>650 MeV.

(b) Squared diphoton mass startpoint. E(γ)>1000 MeV.

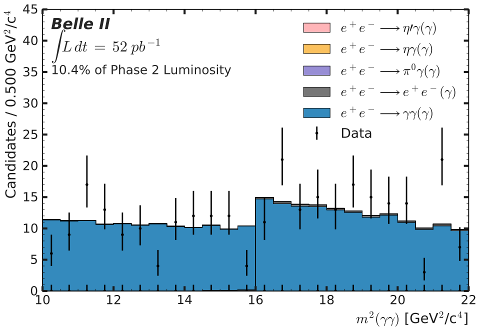

(c) Squared diphoton mass transition region.

Figure 6.8: Zooms of the squared mass startpoint and endpoint, and the selection change point at 16 GeV^2/c^4, for data and MC in the unblinded sub-dataset.

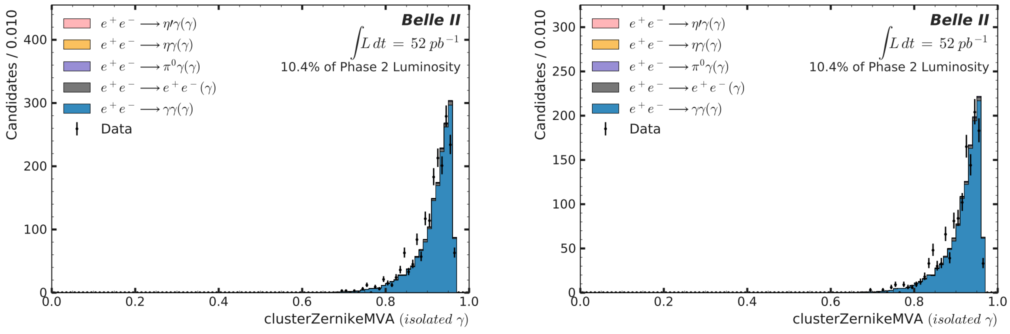

(a) E(γ)>650 MeV.

(b) E(γ)>1000 MeV.

Figure 6.9: `ZernikeMVA` of the most isolated photon distributions for data and MC in the unblinded sub-dataset.

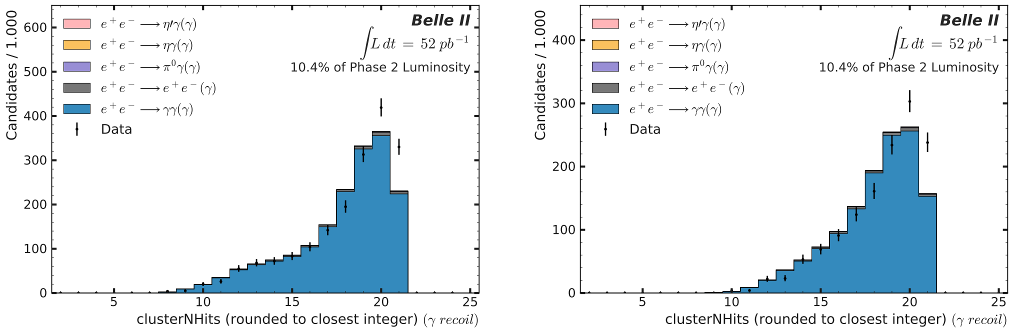

(a) E(γ)>650 MeV.

(b) E(γ)>1000 MeV.

Figure 6.10: `clusterNHits` distributions for data and MC in the unblinded sub-dataset.

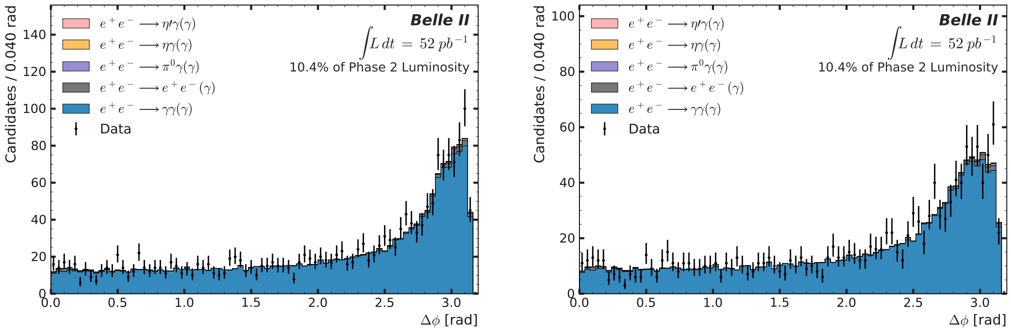

(a) E(γ)>650 MeV.

(b) E(γ)>1000 MeV.

Figure 6.11: $\Delta\phi$ distributions for data and MC in the unblinded sub-dataset.

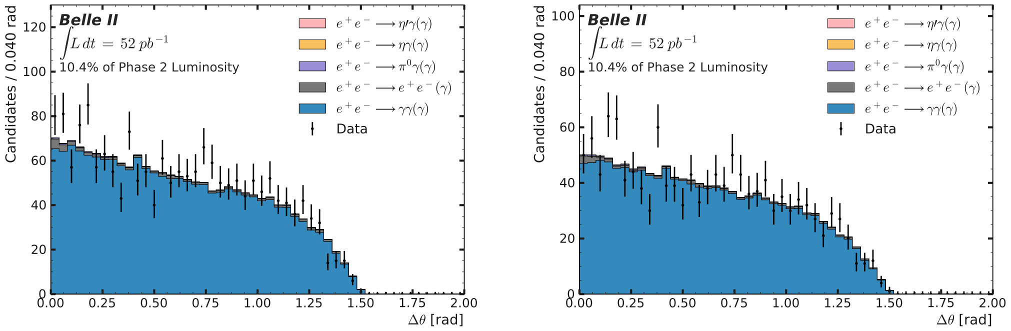

(a) E(γ)>650 MeV.

(b) E(γ)>1000 MeV.

Figure 6.12: $\Delta\theta$ distributions for data and MC in the unblinded sub-dataset.

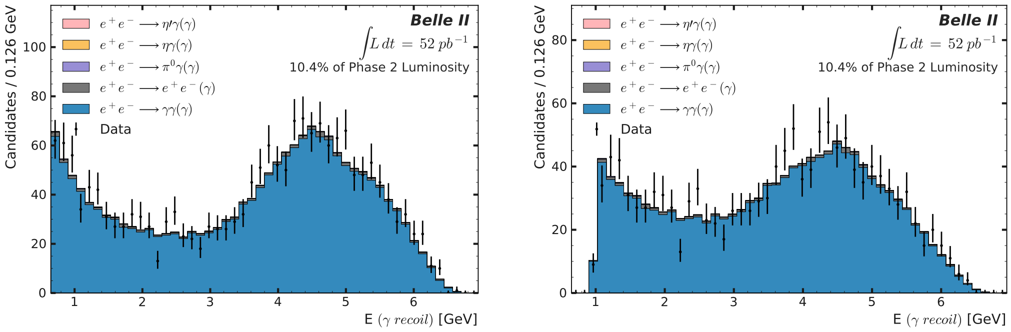

(a) E(γ)>650 MeV.

(b) E(γ)>1000 MeV.

Figure 6.13: $E(\gamma\ recoil)$ distributions for data and MC in the unblinded sub-dataset.

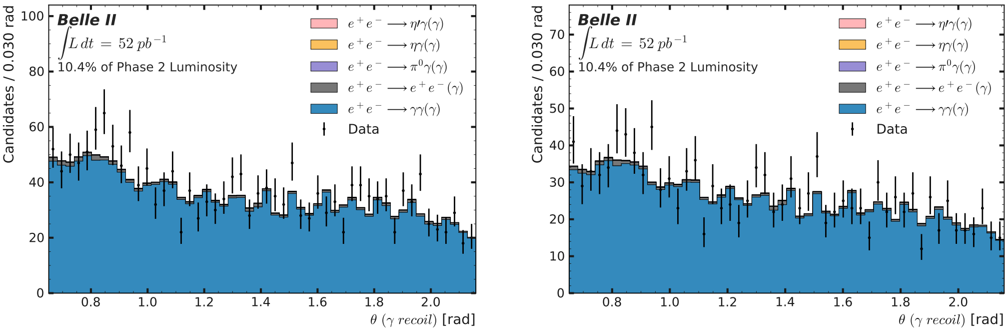

(a) E(γ)>650 MeV. (b) E(γ)>1000 MeV.

Figure 6.14: $\theta(\gamma\ recoil)$ distributions for data and MC in the unblinded sub-dataset.

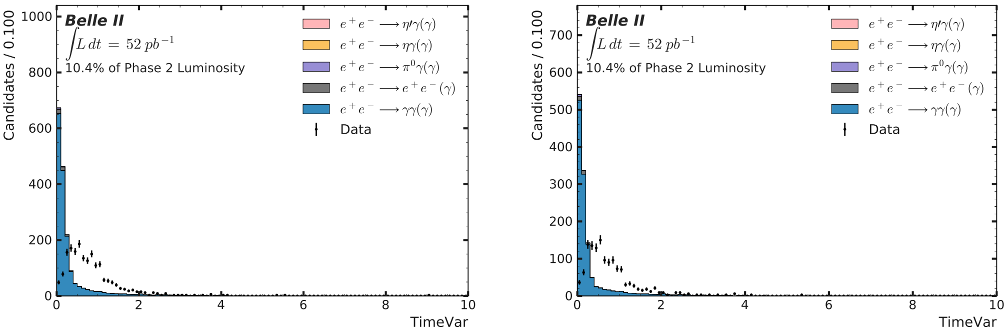

(a) E(γ)>650 MeV. (b) E(γ)>1000 MeV.

Figure 6.15: Time variable t_c distributions for data and MC in the unblinded sub-dataset.

Chapter 7

Signal extraction and results

This Chapter describes the procedures undertaken to perform the signal extraction, and the results of the mass scan and UL extraction on unblinded data.

Section 7.1 overviews the signal extraction strategy.

Sections 7.2 and 7.3 describe how the probability distribution functions (PDFs) used to model signal and background have been determined.

Section 7.4 describes the studies performed to evaluate the signal extraction biases, via signal injection.

Section 7.5 explains the sources of systematic uncertainties (or *systematics* for short) and how they have been propagated.

Section 7.6 shows the results of the full unblinding and the results, including the obtained ULs.

7.1 Strategy

The search for an ALP candidate has been carried out performing a mass scan on the squared masses. Both the diphoton (Eq. 2.11) and the recoil (Eq. 2.12) mass definitions are used, the choice of the transition point between the two is explained in Sec. 7.1.3.

The squared masses are used instead of the linear masses because it is possible, in principle, for the squared recoil mass to take negative values, so the squared root would result in an imaginary number, which are non-physical values deriving from reconstruction smearing. In hindsight, this precaution was unnecessary, given that the recoil mass definition is used for high mass values, where this issue is irrelevant. The mass scan must be carried out in small steps to avoid missing a possible peak. At the same time, the steps should not be too small to avoid performing too many unnecessary fits too close to each other. The chosen compromise is to perform fits in steps of 0.5 σ of the CB modeling the signal at that mass value (see Sec. 7.2.1).

The significance is computed as the square root of two times the difference of the negative logarithm of the likelihoods (NLL) of the fits with and without the signal component [12]:

$$S = \sqrt{2\left(NLL_{signal+bkg} - NLL_{bkg}\right)}. \tag{7.1}$$

For each fit, the local significance is computed. The systematic uncertainties (Sec. 7.5) are taken into account at this step. The *signed significance* is studied, i.e. the significance multiplied by the sign of the signal yield. Negative peaks, or dips, are not allowed in the final fit, so a dedicated scan was performed purposely for this study. If a fit with *local* significance higher than 3 had been found, the *global* significance would have been estimated taking into account the LEE (Sec. 7.1.2). If a fit with *global* significance higher than 3 had been found, an evidence would have been claimed, otherwise ULs to the signal strength would have been computed and set. In this analysis, no *local* significance higher than 3 was found, so the global significance was not computed and ULs are set.

In the following, all signal-extraction-related studies preceding the unblinding are performed on MC-only samples. Different *random extractions* are obtained by randomly sampling from the MC background samples a number of events equal to the one expected by the cross section of the relative process (Sec. 4.3.2) and the collected integrated luminosity. The process of randomly extracting a MC background sample is made repeatable by saving the *randomizing seed* used for that particular sampling. Given the involved cross sections, the data integrated luminosity $\mathcal{L}$, and the sizes of the MC samples, the different random extractions have very little overlap:

$$\begin{aligned} N^{\text{data}}_{e^+e^-\to\gamma\gamma(\gamma)} &= \mathcal{L}\cdot\sigma_{e^+e^-\to\gamma\gamma(\gamma)} \approx \\ &\approx 1.8\cdot 10^6 \ll 60\cdot 10^6 = N^{\text{MC}}_{e^+e^-\to\gamma\gamma(\gamma)} \\ N^{\text{data}}_{e^+e^-\to e^+e^-(\gamma)} &= \mathcal{L}\cdot\sigma_{e^+e^-\to e^+e^-(\gamma)}\cdot < w_T > \approx \\ &\approx 3500 \ll 100\cdot 10^6 = N^{\text{MC}}_{e^+e^-\to e^+e^-(\gamma)}, \end{aligned} \tag{7.2}$$

where $< w_T > \approx \epsilon^2_{\text{tracking}} = 0.025^2$ is the average of the weights w_T, defined in Sec. 4.3.2, over the sample. For the purpose of plotting (e.g. Sec. 7.6.1), in order to have a smoother distribution, every radiative Bhabha event in the sample is considered and weighted with its corresponding w_T.

The UL is computed with a NLL approach [53], using the CLs method [54]. The `RooFit` package [62, 63] from `ROOT` [51] was used, in particular the *asymptotic calculator* [64, 65]. The signal yield was parametrized as a function of the signal cross section[1], so the mass scan returned a cross section value for each fit point. The signal cross section and the signal yield are related by the following simple formula:

$$Y_S = \sigma_{e^+e^- \to a\gamma} \mathcal{L} \epsilon_S, \tag{7.3}$$

where Y_S is the signal yield, $\sigma_S \equiv \sigma_{e^+e^- \to a\gamma}$ is the signal production cross section, and ϵ_S is the signal efficiency for a given mass value. Only positive signal yields are allowed, i.e. only positive cross sections: $\sigma_S \geq 0$. Part of the systematics (Sec. 7.5), in particular the ones relative to signal efficiency and signal resolution[2], are applied by propagating them as nuisance parameters in the fit. Such systematics are modeled with a Gaussian constraint centered on the nominal value and with standard deviation equal to the uncertainty itself.

The procedure is repeated for each scanned mass point, and a set of values are returned for each of them:

- Expected UL, i.e. the median UL expected in the absence of a signal;
- Expected UL $\pm 1\ [2]\,\sigma$, i.e. the above expected UL increased and decreased by one [two] σ of uncertainty;
- Observed UL, i.e. the UL effectively obtained on the actual dataset.

With these values, the so-called *Brazilian bands*[3], i.e. the bands defined by the expected UL ± 1 and 2 σ, can be built by plotting the results for each scanned mass values. On top of this, the observed UL is plotted as well.

7.1.1 Negative significance check

In the final extraction, only positive signal yields are allowed. Nonetheless, a test allowing negative signal yields is performed to check that the fitter

---

[1] Which can be later converted into coupling strength using Eq. 2.7.

[2] Which translates into an uncertainty on the σ parameter of the CB used to model the signal PDF.

[3] Called this way because conventionally the 1-σ band is green and the 2-σ one is yellow, resulting in yellow-green-yellow horizontally oriented strips, which resemble the flag of Brazil.

behaves correctly. In particular, it was checked that negative signal-like peaks coming from statistical fluctuations are properly fitted, and that the significance trend as a function of the ALP candidate mass is smooth. This is done by showing the significance multiplied by the sign of the fitted cross section, shown in Fig. 7.1, where the significance scan for a normal extraction (only positive signal yields allowed) and one with negative yields allowed are shown. The latter is continuous, and the negative peaks are qualitatively like the positive ones, confirming the robustness of the extraction procedure.

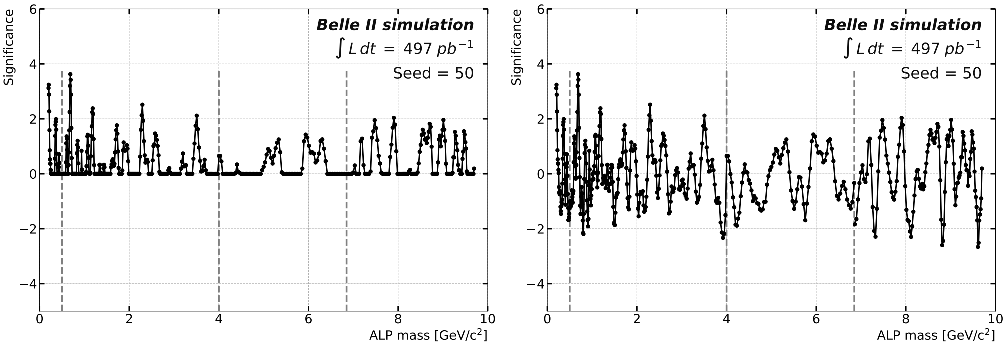

(a) Only positive signal yields. (b) Allowing negative signal yields.

Figure 7.1: Significance scan. Left: default extraction where the signal yield is forced to be positive; right: extraction allowing negative signal yields. MC only, for the same given random extraction.

7.1.2 Look-elsewhere effect

The mass of the hypothetical ALP is unknown. This is the reason why a mass scan is performed. Consequently, multiple fits (around 500 in this analysis) are performed. Therefore, the significance of observing a local excess from background fluctuations anywhere in the search range must be taken into account to get the global significance.

The most conservative way [66] to estimate the LEE is the following:

$$p_{global} = p_{local} \cdot N, \tag{7.4}$$

where p_{global} is the global p-value, p_{local} is the local one, and N is the number of mass fits performed.

It is possible to better approximate p_{global} in the asymptotic limit of high local signficance by introducing trial factors [67], such that $p_{global} = p_{local} \cdot N_{trial}$:

$$N_{trial} = 1 + \frac{1}{p_{local}} \langle n_{up}(Z_{test}) \rangle e^{\frac{Z^2_{test} - Z^2_{local}}{2}} \tag{7.5}$$

where $\langle n_{up}(Z_{test}) \rangle$ is the average number of times where the significance Z is larger than a reference significance Z_{test}, determined using multiple background-only MC samples. The global significance is thus given by:

$$p_{global} = p_{local} + \langle n_{up}(Z_{test}) \rangle e^{\frac{Z^2_{test} - Z^2_{local}}{2}}. \tag{7.6}$$

By using a relatively small test significance Z_{test}, for example between 0.5 and 2, the number of up-crossings $n_{up}(Z_{test})$ can be estimated with good accuracy. An estimation can be performed directly on data, without the need to resort to multiple background-only MC samples, by counting $n_{up}(Z_{test})$ with a small Z_{test}.

Alternatively, the more classic, precise, and resource-consuming option is to use the toy-MC approach [68]. In this case, a large number of independent MC toy background-only distributions are sampled from the PDF that is used to model the background shape. For each of them, a full mass scan is run, and the highest significance S^{max}_{MC} is saved. The distribution of the highest simulated local significances S^{max}_{MC} is built.

When the actual extraction on data is performed, the highest observed local significance S^{max}_{obs} is compared with the cumulative probability distribution of S^{max}_{MC}: the probability L that $S^{max}_{MC} > S^{max}_{obs}$ is computed, and L is the *global* p-value used to compute the global significance.

Although, given that the maximum local significance found in this analysis is lower than 3, none of the above approaches has been deployed.

7.1.3 Diphoton to recoil mass transition

In this analysis two ALP mass definitions, diphoton and recoil, are used. They perform differently at different ALP mass regimes: diphoton is narrower at low ALP masses, and vice versa. This double definition could be avoided by performing a *kinematic fit*, i.e. applying the constraint that the invariant mass of the three photons is equal to the collision energy $\sqrt{s}$. Although, at this early stage of the experiment, and with early data, the kinematic fitting strategy cannot be employed. This is due to the fact that the photon covariance matrix is not precisely known, and therefore the uncertainties on photon position and energy can not be correctly propagated to the constraint.

At the time of writing, the photon covariance matrix is in development: further details, including motivations, current status, and prospects, are described in Chap. 8.

To determine the transition point, i.e. the ALP candidate mass, where to stop using the diphoton mass definition and start using the recoil one, a method based on the upper limits computed on background-only MC samples is deployed. This consists of plotting the limits obtained with the two mass definitions, and determine the crossing point. In particular, the median of 10 random extractions is plotted and studied. This is shown in Fig. 7.2b. The significance is also taken in consideration, to avoid regions of the masses distributions where the modeling of the background is such that there is a systematic excess. This would happen for the diphoton mass definition in the region between ~ 6.9 GeV/c^2 to ~ 7.7 GeV/c^2 (see Fig. 7.2a). To avoid this area, the transition point is set to be lower than what would come from a study of the expected upper limit only. This decision is corroborated by the study of the observed upper limit, also extracted from background-only MC samples. As can be seen in Fig. 7.2c, the chosen transition point corresponds to the mass value where the limits obtained with the recoil mass are stronger than the ones obtained with the diphoton mass. The resulting transition point is

$$m_a(\text{threshold diphoton to recoil}) = 6.85\ \text{GeV/c}^2. \tag{7.7}$$

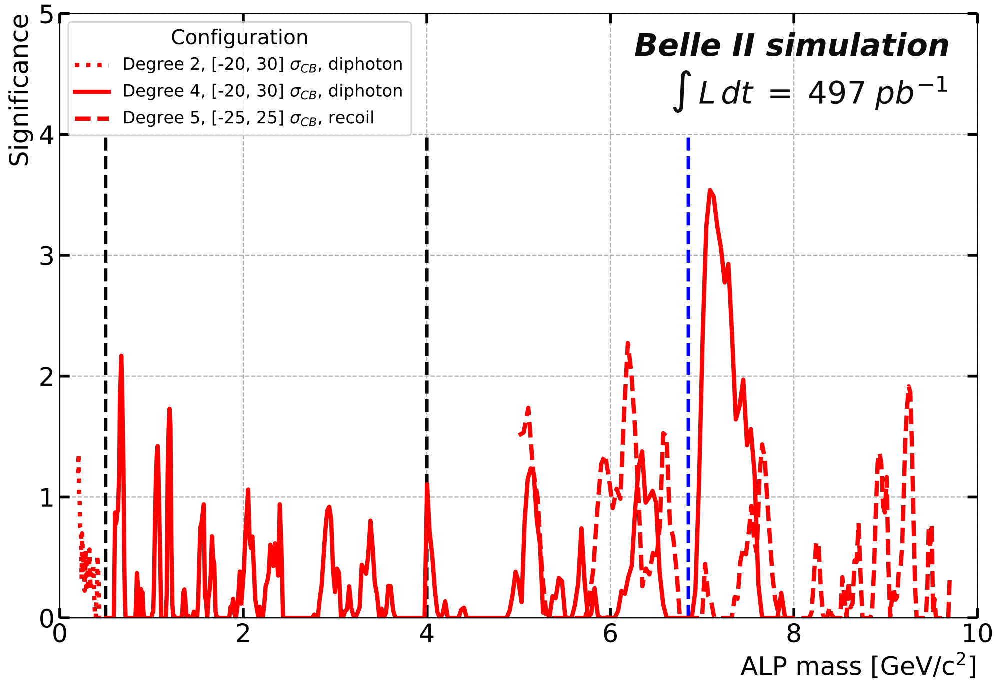

(a) Significance.

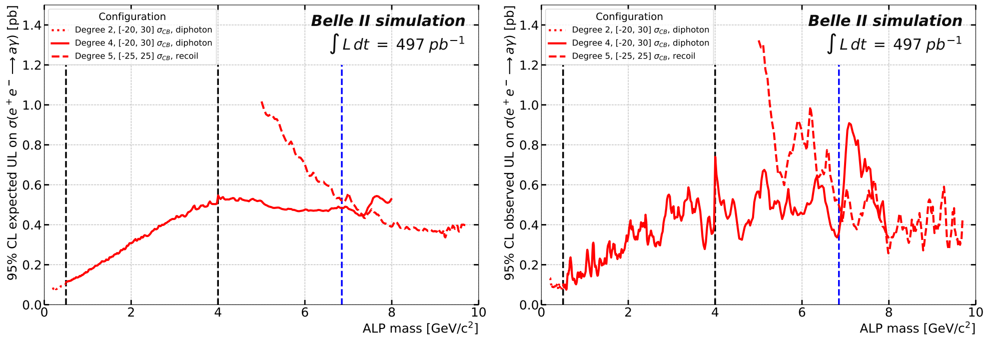

(b) Expected upper limit on cross section. (c) Observed upper limit on cross section.

Figure 7.2: Significance and expected and observed upper limit on cross section, with diphoton and recoil mass definitions, from background-only MC extractions. The red lines are the median of 10 extractions, each corresponding to a different random sampling. The left dashed black line corresponds to the point where the order of the polynomial is increased, 0.5 GeV/c^2; the right black one to the point where the photon energy cut changes from 1000 MeV to 650 MeV, 4.0 GeV/c^2; the blue one to the transition value of 6.85 GeV/c^2.

7.2 Signal PDF

Each event passing the selection contains three photons. As explained in Sec. 5.2.1, this means that there are three ALP candidates per event. In case of a true signal event, only one of the three candidates is the real ALP, while the other two form the *combinatorial component* of the signal.

Real ALP candidates have a peaking distribution in the mass variable, while the candidates from the combinatorial component have a wider and non-peaking mass distribution; see Fig. 7.6. Both of these components have to be properly modeled to obtain the overall signal PDF. The signal PDF is fixed for the mass scan, based upon results obtained from the MC studies; only the signal yield is allowed to float. The signal PDF modeling is described in the following Sections.

7.2.1 Peaking component

The signal peak model is built to describe both the diphoton and the recoil mass definitions. High mass ALPs are produced in association with a soft recoil photon and decay into hard photons. The ECL relative resolution on photon energy is approximately constant. Therefore the ECL absolute resolution is proportional to the photon energy, thus the uncertainty on the recoil mass decreases when the ALP mass increases. This implies that there is a threshold point where the width of the recoil mass distribution becomes smaller than the width of the diphoton mass distribution.

The squared diphoton mass distributions have been fitted from 0.2 to 9.7 GeV/c^2. The squared recoil mass distributions have been fitted from 3.0 to 9.7 GeV/c^2. The fit function used is a CB for both distributions. All of the CB parameters α, n, σ, and μ are free to float in the fit. Only the TM ALP candidates are used for these fits, i.e. the combinatorial component has not been taken into account. Some fit examples, out of the hundreds performed at different ALP masses hypotheses, are shown in Figs. 7.3 and 7.4.

The threshold mass m_a(threshold) above which the recoil mass distribution gets narrower than the diphoton mass distribution is at approximately

$$m_a \simeq 6.5 \text{ GeV/c}^2, \tag{7.8}$$

which is in agreement with the transition point found in Eq. 7.7, which does take into account also the background shape and not only the signal width.

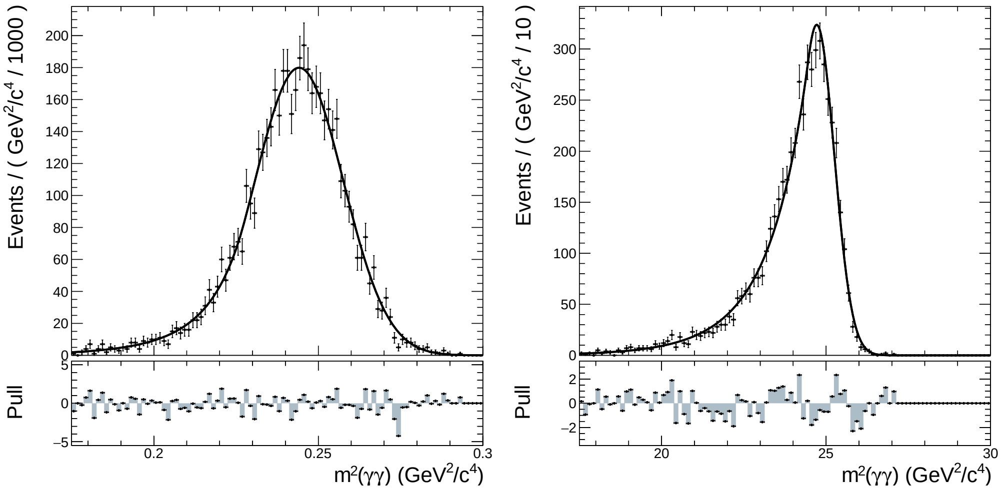

(a) $m_a = 0.5$ GeV/c^2. (b) $m_a = 5.0$ GeV/c^2.

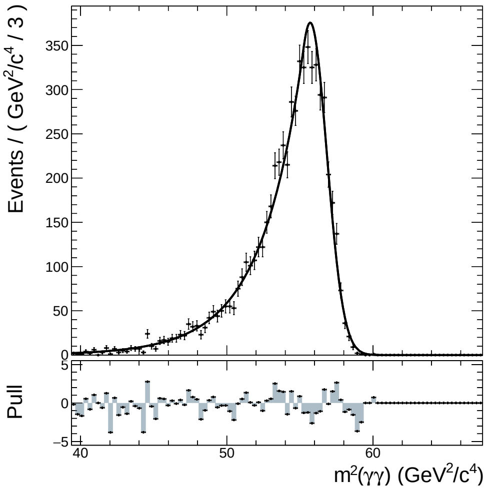

(c) $m_a = 7.5$ GeV/c^2.

Figure 7.3: Fitted distributions of squared diphoton mass, for some benchmark points. Only TM ALPs.

The CB parameter α and the signal efficiency ϵ_S, obtained from all different MC signal samples, are smoothed, and these smoothed values are used to interpolate between points for the mass scans. The CB resolution σ is fit as described in Sec. 7.5.1, and the fit results are used for the mass scans.

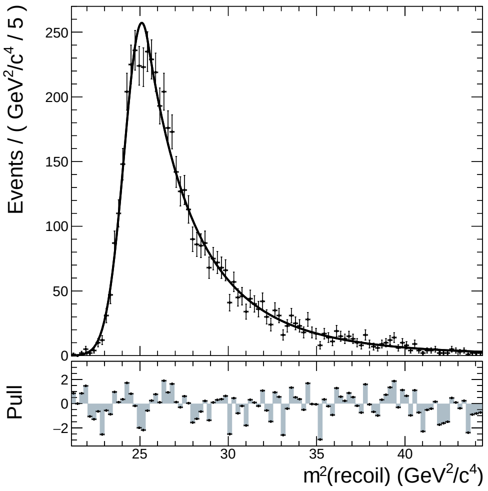

(a) $m_a = 5.0$ GeV/c^2.

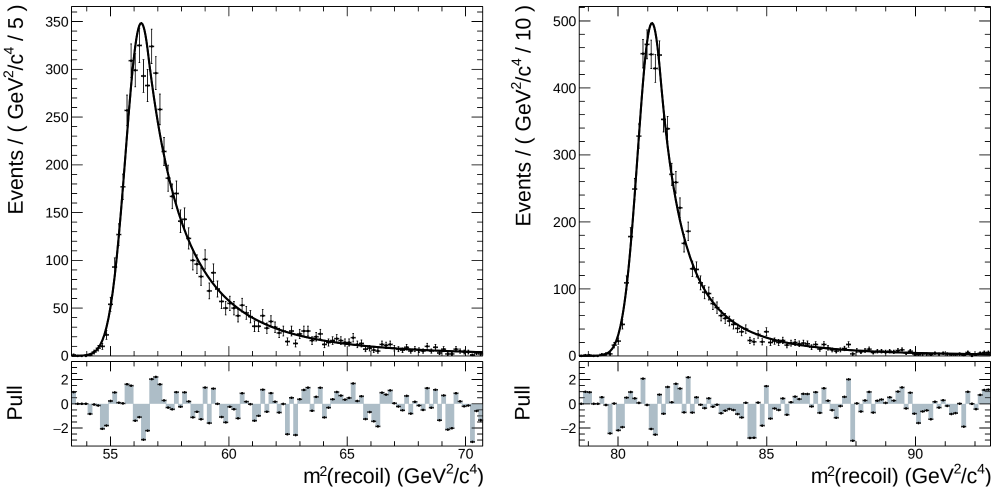

(b) $m_a = 7.5$ GeV/c^2.

(c) $m_a = 9.0$ GeV/c^2.

Figure 7.4: Fitted distributions of squared recoil mass, for some benchmark points. Only TM ALPs.

7.2.2 Combinatorial component

The ratio between the combinatorial and the true component of the signal, normalized over the latter, is reported in Fig. 7.5. These numbers are computed within the fit windows defined in Sec. 7.3. The combinatorial component is included in the total signal PDF only if this ratio is above 1%.

The combinatorial component is modeled using a Kernel Density Estimator (KDE)[4] [70], with a smoothing parameter of 1.5, applying a mirroring to both sides of the window. Fig. 7.6 shows some stacked histograms for the true and combinatorial components of the signal. Fig. 7.7 shows some fits to the combinatorial component, for ALP masses where distributions have their endpoints, using multiple smoothing parameters for the fitting KDE.

The mass scan is performed in steps of half signal width. The KDE for the analyzed mass, as well as the selection efficiency on the combinatorial component, are computed using the closest available sample, which is no more far away than half a signal width from the analyzed mass point.

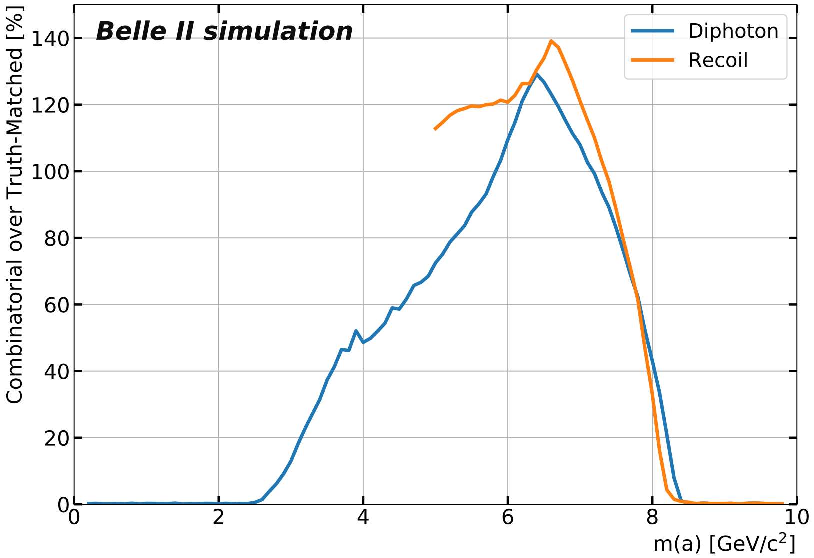

Figure 7.5: Ratio of the number of combinatorial candidates over number of true candidates, within the fit window, normalized to the number of true candidates.

---

[4] Technically, the `RooKeysPdf` [69] class of `ROOT` has been used.

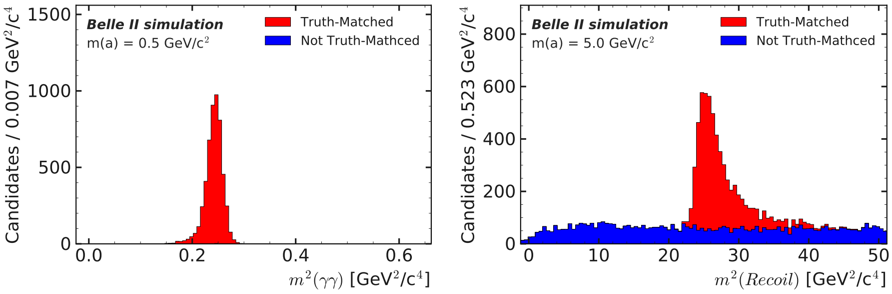

(a) Diphoton invariant mass distribution, $m_a = 0.5$ GeV/c^2.

(b) Recoil invariant mass distribution, $m_a = 5.0$ GeV/c^2.

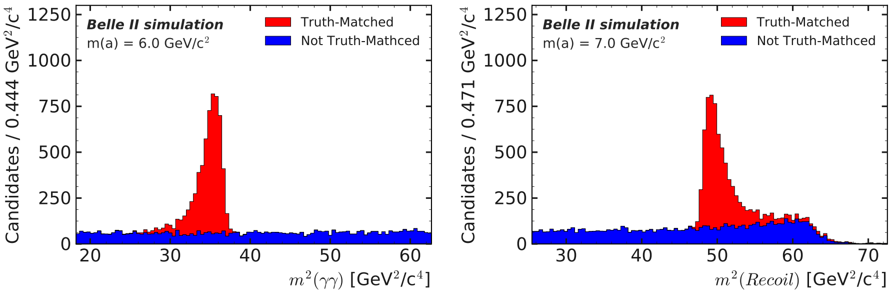

(c) Diphoton invariant mass distribution, $m_a = 6.0$ GeV/c^2.

(d) Recoil invariant mass distribution, $m_a = 7.0$ GeV/c^2.

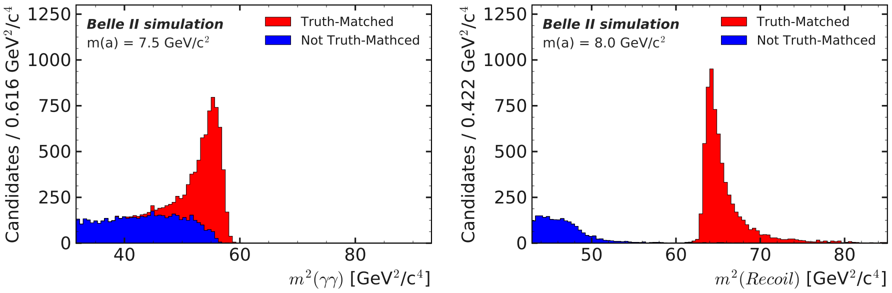

(e) Diphoton invariant mass distribution, $m_a = 7.5$ GeV/c^2.

(f) Recoil invariant mass distribution, $m_a = 8.0$ GeV/c^2.

Figure 7.6: TM and non-TM ALP candidates masses distributions, stacked.

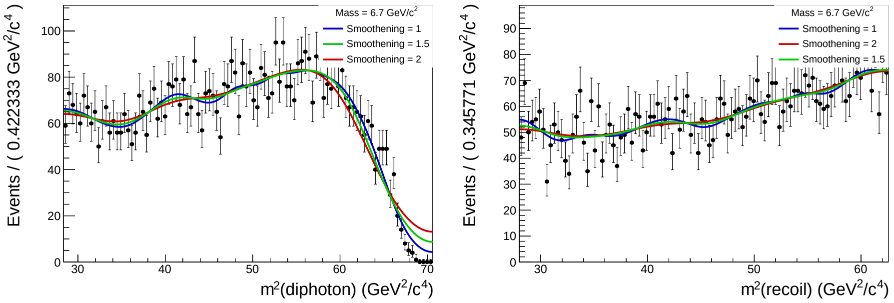

(a) Diphoton invariant mass distribution, $m_a = 6.7$ GeV/c^2.

(b) Recoil invariant mass distribution, $m_a = 6.7$ GeV/c^2.

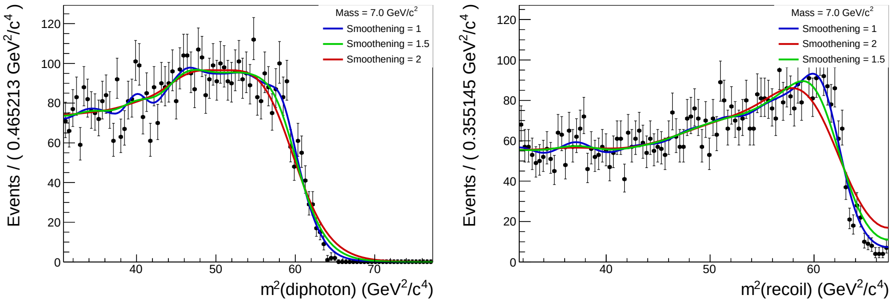

(c) Diphoton invariant mass distribution, $m_a = 7.0$ GeV/c^2.

(d) Recoil invariant mass distribution, $m_a = 7.0$ GeV/c^2.

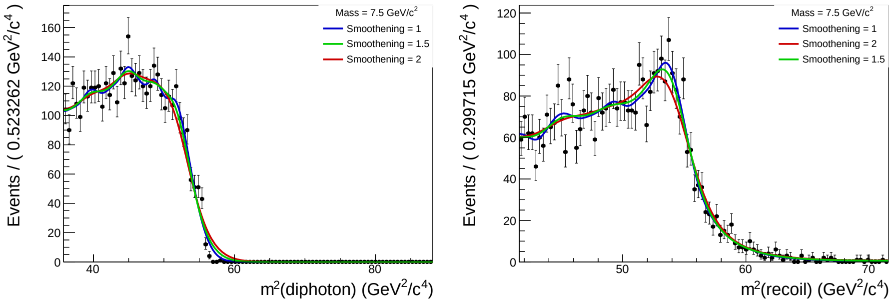

(e) Diphoton invariant mass distribution, $m_a = 7.5$ GeV/c^2.

(f) Recoil invariant mass distribution, $m_a = 7.5$ GeV/c^2.

Figure 7.7: Combinatorial component of the signal, fitted with KDE, using different smoothing parameters 1, 1.5, and 2. The chosen value is 1.5.

7.3 Background PDF and fit range

The background PDF shape parameters are left floating during the mass scan, to adapt to the actual background. Before unblinding, the fit range (also called *fit window*) and the PDF family used were decided and fixed. Given the relative smoothness of the background shape, orthogonal polynomial PDFs are used. In particular, `RooChebychev` PDFs [71] are used. Although, the particular kind of polynomials has little impact on the extraction, as long as it is capable of adequately describe the background shape. The degree of the polynomial and the fit range are fixed, but the actual parameters of the polynomials are left floating during the unblinding mass scan.

The fit range is defined as a function of numbers of σ_{CB} around the mean value μ_{CB} of the CB modeling the signal at the analyzed value of m_a. The range is hence defined by two numbers, n_{Left} and n_{Right}, such that the fit window is defined as

$$[-n_{Left} \cdot \sigma_{CB}, +n_{Right} \cdot \sigma_{CB}]. \tag{7.9}$$

Following preliminary tests, studies are performed by varying the fit ranges within reasonable extremes, namely using $n_{Left} = 20 \pm 5$ and $n_{Right} = 30 \pm 5$ for diphoton, and $n_{Left} = 25 \pm 5$ and $n_{Right} = 25 \pm 5$ for recoil. Different background polynomial orders are tested, depending on the mass range: $n = [1, 2, 3]$ for diphoton with $m_a < 0.5$ GeV/c^2, $n = [3, 4, 5]$ for diphoton with $0.5 \leq m_a < 8.0$ GeV/c^2, $n = [4, 5, 6]$ for recoil with $m_a \geq 5.0$ GeV/c^2. The tests are performed for different values of m_a, in steps of approximately half of the signal width.

The resulting reduced χ^2 distributions as a function of the ALP mass for both the diphoton and the recoil squared masses are shown in Fig. 7.8. The choice is for the combination of polynomial order and fit range that makes the reduced χ^2 the closest possible to 1. Also, it is required that the reduced χ^2 is smooth across m_a, for the interval of interest. The chosen background orders and fit ranges are shown in Tab. 7.1 and Tab. 7.2.

m_a [GeV/c^2]	**Background order**	n_{Left} [σ_{CB}]	n_{Right} [σ_{CB}]
0.2-0.5	2	20	30
0.5-8.0	4	20	30

Table 7.1: Background orders and fit ranges for the di-photon mass.

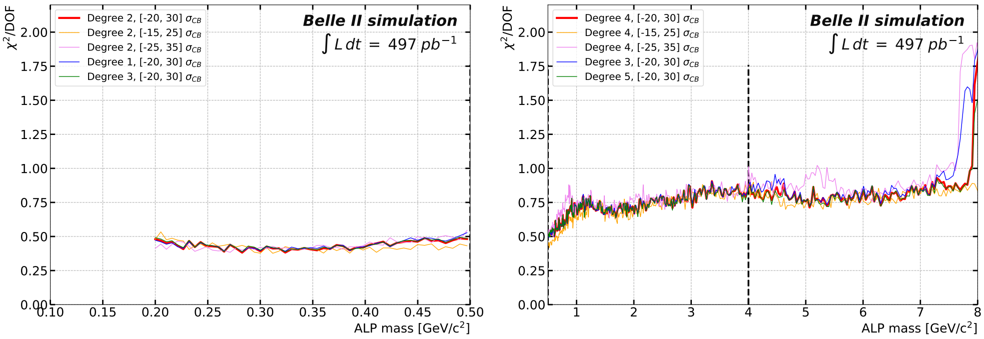

(a) Diphoton, low masses.

(b) Diphoton, high masses.

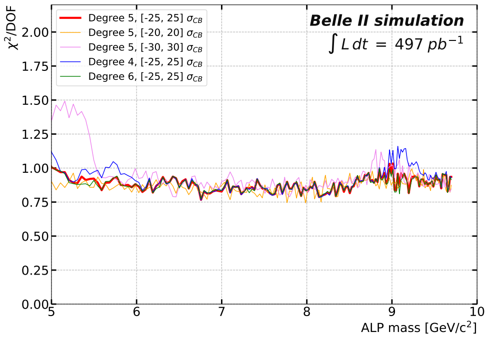

(c) Recoil.

Figure 7.8: Trend of the reduced χ^2 for the different combinations of fit widths and polynomial orders.

m_a **[GeV/c^2]**	**Background order**	n_{Left} $[\sigma_{CB}]$	n_{Right} $[\sigma_{CB}]$
5.0-9.7	5	25	25

Table 7.2: Background orders and fit range for the recoil mass.

7.4 Signal injection and bias studies

Before unblinding, the validity of the fit procedure in the presence of a non-zero signal was tested. MC signals of different strengths were injected, and the signal+background fits were performed. The goal was to check if the cross section measured by the fit was in agreement with the injected one, and to quantify the presence of possible biases.

The bias studies were conducted in the following way.

The MC background samples were scaled to four times the data integrated luminosity. These events were fitted with the background-only PDFs described in Sec. 7.3, and the resulting PDFs were saved. These parent PDFs are used to randomly sample background events in the next steps.

For 6 different injected cross sections and 15 different ALP masses, 250 toys were run. For each toy, background events were sampled from the parent PDFs according to a Poisson distribution with mean equal to the number of events expected in data, and signal events were sampled from the signal PDF at the injected ALP mass (Sec. 7.2) according to a Poisson distribution with mean equal to the number of signal events expected from the injected cross section under study. The signal+background fit was run (an example is in Fig. 7.9), and it returned a fitted cross section value. The fitted cross section value is compared with the injected one by computing their pull: $P(\sigma) = \left(\sigma^{\text{fit}} - \sigma^{\text{injected}}\right) / \sigma^{\text{fit}}\text{err}$. For each injected cross section at each tested ALP mass, the 250 pulls obtained from the 250 toys are plotted and fit with a Gaussian. The mean and the width of these pulls are expected to be equal to zero and one respectively, if no bias is present and the uncertainties are correctly estimated.

The results of this study are summarized in Fig. 7.10. All values agree with their expectations within uncertainties, so the fit procedure is unbiased within the required precision.

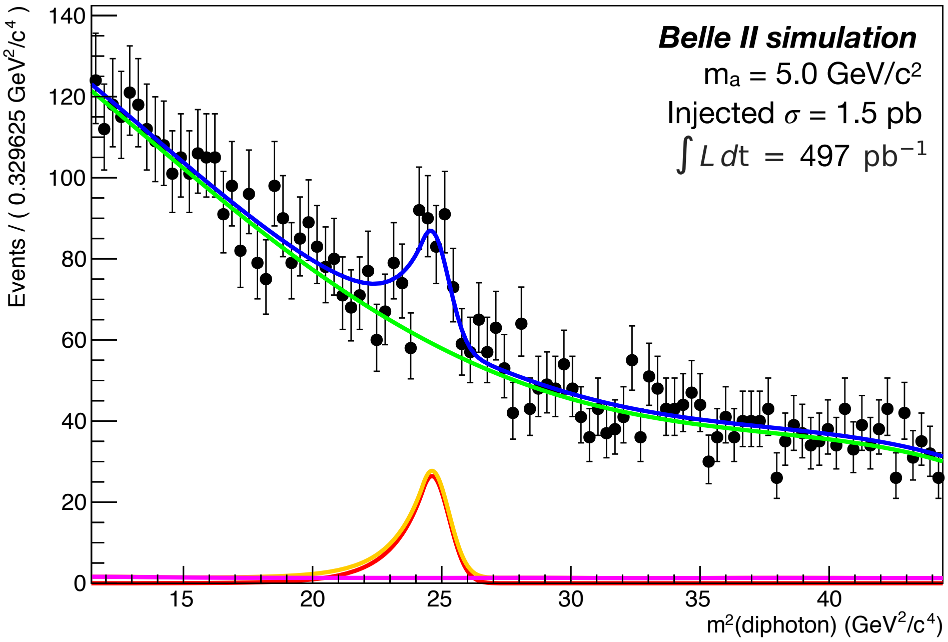

Figure 7.9: Example of MC-only signal+background fit. The signal corresponds to an injected cross section of 1.5 pb at ALP mass $m_a = 5.0$ GeV/c^2.

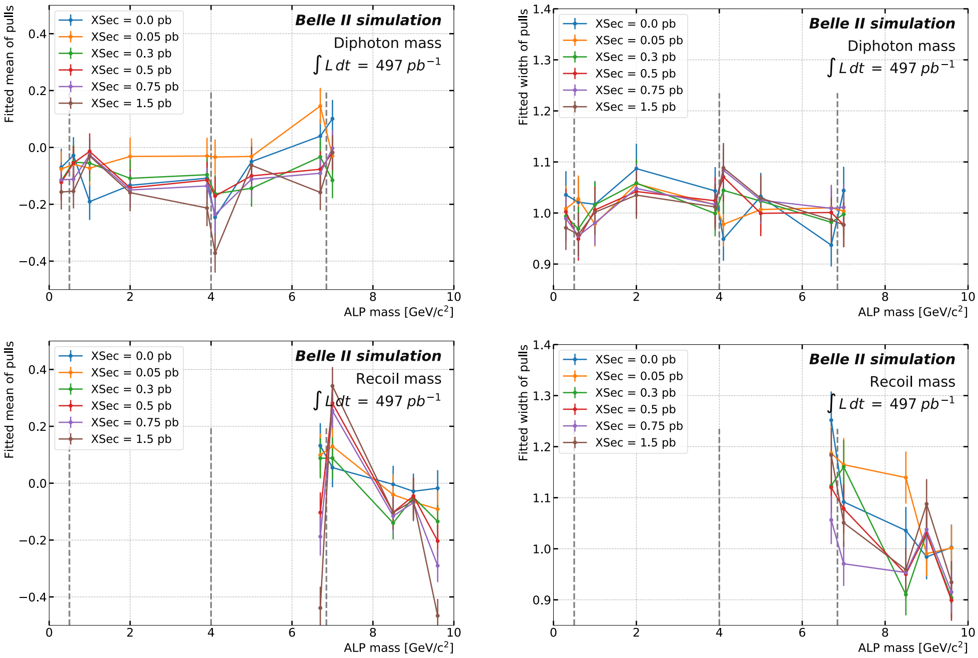

Figure 7.10: Trend of the mean (left) and width (right) of the Gaussian fits to the distributions of the pulls of the fitted cross section, for diphoton (top) and recoil (bottom) mass definitions. Differently colored lines correspond to different cross sections (*XSec* in the legends). The mean [width] values are consistent with the expected value of zero [one].

7.5 Systematic uncertainties

This analysis is affected by three main sources of systematics:

- Uncertainty regarding the signal width (Sec. 7.5.1);
- Uncertainty about signal efficiency (Sec. 7.5.2);
- Uncertainty deriving from the choice of background PDF shape and fit range (Sec. 7.5.3).

The first two contributions are included by smearing the signal width and the signal efficiency with two Gaussians centered on zero and having width equal to the uncertainty itself. The last source of uncertainty is propagated and estimated by re-performing the extraction with different background PDF shapes and fit ranges.

7.5.1 Mass resolution

Studies on radiative dimuon events $e^+e^- \to \mu\mu\gamma$ have been performed to study the photon reconstruction performance. Events are selected if they contain exactly two good tracks and one high-energy photon. Good tracks here are defined by $p > 1.0$ GeV/c, $|z0| < 4.0$ cm, $|d0| < 1.0$ cm, $0.651 < \theta < 2.159$ rad, and by having at least 4 hits in the CDC, also they must be matched to an `ECLCluster` with $0.15 <$ `clusterE` < 0.4 GeV. Both muon tracks are used to perform a vertex fit, after which the kinematic of both tracks is updated. The photon must satisfy `clusterE` > 0.1 GeV, $0.651 <$ `clusterTheta` < 2.159 rad, and it cannot be farther away than 0.25 rad from the direction of the recoil momentum of the two muons. The event is rejected if it contains other good tracks or other photons with $E > 0.2$ GeV. The kinematic of the matched photon is then compared to the kinematic of the predicted $p_{recoil}(\mu\mu)$. The p_{recoil} quantities are fully independent from the ECL, as they are wholly determined from tracking.

Once these events are selected, the quantity $\frac{E(\gamma)-p_{recoil}(\mu\mu)}{p_{recoil}(\mu\mu)}$, i.e. the difference between the photon energy and the recoil momentum of the dimuon pair, normalized to the recoil momentum itself, is computed in bins of p_{recoil}, both for data and for MC samples. This quantity is expected to be distributed around zero for $e^+e^- \to \mu\mu\gamma$ events, and its width is correlated with the photon resolution. All these distributions are fitted with a CB, and the mean value μ and the Full Width Half Maximum (FWHM) are extracted for each of them. The obtained results are summarized in Fig. 7.11.

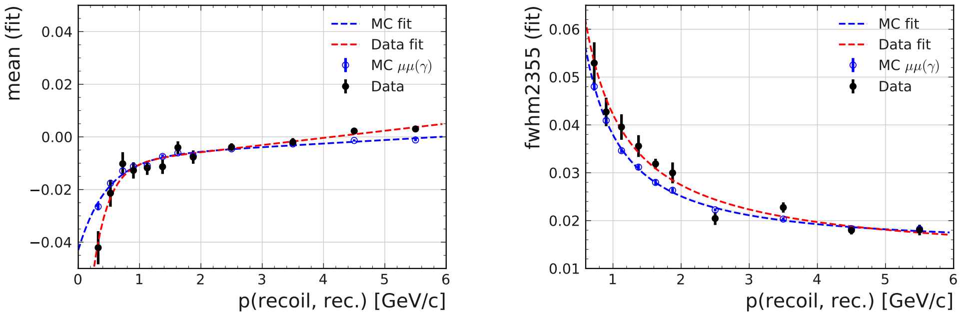

(a) Fitted mean values as function of p_{recoil}.

(b) Fitted FWHM/2.355 values as function of p_{recoil}.

Figure 7.11: Fit parameters of the CB fits on data and MC. The FWHM is divided by 2.355, the same value needed to convert a FWHM into a standard deviation for a Gaussian.

For each of these plots, the difference between data and MC values is computed to correct for discrepancies. In particular, a *bias correction factor* (the one also used for the partial unblinding, see Sec. 6.4) is extracted from Fig. 7.11a, and a *resolution smearing factor* is extracted from Fig. 7.11b.

The *bias correction factor* between data and MC is computed as the difference of the fitted μ values. The correction is small for low-energy photons, and goes up to about 0.5% for high-energy ones. This correction is applied to data. The remaining bias towards non-zero means is typically due to additional radiative photons which were not reconstructed, thus making the recoil momentum of the muon pair not coincide with the reconstructed photon's energy. This effect is described by the simulation and the distributions are consistent between data and MC.

The *resolution smearing factor*, i.e. the resolution difference between data and MC, is computed as the difference between the FWHM/2.355 values obtained from the two samples. The resolution smearing factor is about 2% for low-energy photons and negligible for high-energy ones. This correction is applied to MC by smearing the MC photon energies and re-computing the quantities that depend on them, see below.

The original and corrected results are summarized in Fig. 7.12.

The polar angle difference $\theta(\gamma) - \theta_{recoil}(\mu\mu)$ and the azimuthal angle difference $\phi(\gamma) - \phi_{recoil}(\mu\mu)$ have been studied too. Great agreement has been found for these variables. Additionally, the diphoton invariant mass is more dependent on the energies of the two γ than on their position, and possible

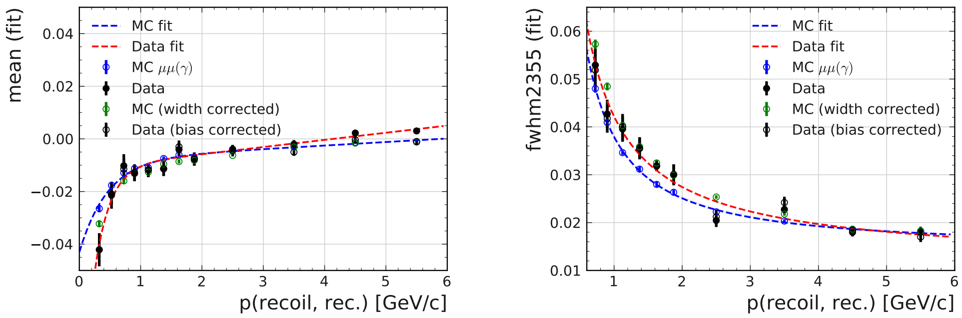

(a) Fitted mean values as function of p_{recoil}. (b) Fitted FWHM/2.355 values as function of p_{recoil}.

Figure 7.12: Fit parameters of the CB fits on data (after bias correction) and MC (after resolution smearing).

mis-modelings have indirectly taken into account, from the point of view of uncertainty propagation, with the studies performed by shifting the IP (see Sec. 7.5.2).

Using the results summarized in Fig. 7.12, it is possible to propagate the uncertainty on photon energy resolution to the mass resolution. The signal resolution is first computed using the squared mass distributions obtained using the nominal values of the photon energies. Then, a correction to the photon energies is applied by smearing the photon energies in MC with a Gaussian with mean zero and standard deviation equal to the difference between data and MC of the FWHM/2.355 (from Fig. 7.11b). The new signal mass distributions are fit, and half of the difference between the newly obtained width values and the default ones is taken as an estimate of the size of the systematic uncertainty. These results are finally fitted to obtain a smooth distribution, and they are shown in Fig. 7.13. As explained at the beginning of this Section, this uncertainty is propagated via profiling as a nuisance parameter during the final fits for the UL extractions.

The same Gaussian smearing is applied to all variables that depend on photon energy, as mentioned above. This means that E_γ itself, $M^2_{\gamma\gamma}$, M^2_{recoil}, and $M_{\gamma\gamma\gamma}$ are all affected by this smearing.

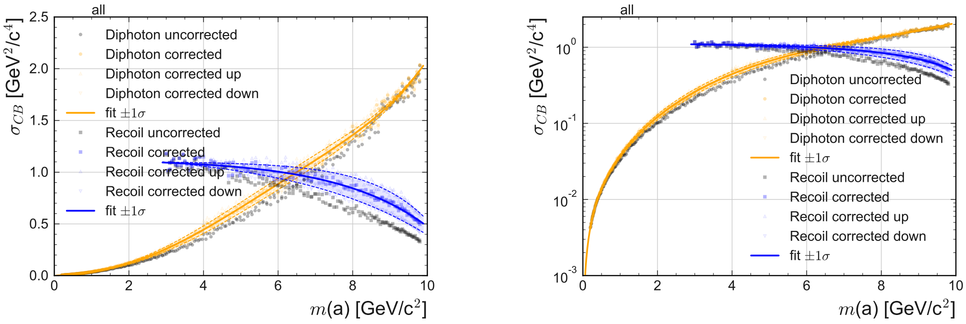

(a) Linear y-axis.

(b) Logarithmic y-axis.

Figure 7.13: Fitted squared mass resolutions and resulting uncertainties from photon energy uncertainty. Linear and logarithmic y scales.

7.5.2 Signal efficiency

The uncertainty on the signal efficiency is also propagated via profiling as a nuisance parameter during the final fits for the UL extractions. There are multiple sources of uncertainty on this quantity.

A component of the signal reconstruction efficiency uncertainty is obtained directly from data/MC yield comparison after unblinding (see Sec. 7.6.1). In particular, the data and MC yields have been compared, separately for $m_a < 4$ GeV/c^2 (corresponding to $E_\gamma > 1000$ MeV) and for $m_a > 4$ GeV/c^2 (corresponding to $E_\gamma > 650$ MeV). The actual yields are reported in Table 7.3; to compute the uncertainty, the discrepancy[5] between data and MC is linearly summed (to be the most conservative) to its own uncertainty. This results in a 5.3% uncertainty for $m_a < 4$ GeV/c^2, and 4.5% for $m_a > 4$ GeV/c^2. This is the largest systematic uncertainty for the signal reconstruction efficiency.

The other sources of systematics on the signal reconstruction efficiency are derived from varying the selection cuts' values (Sec. 5.5). Each variable cut is increased and reduced by a value which is of the same order of magnitude of the data/MC discrepancy for that variable or of the uncertainty on that variable; the actual values are reported in the list below. In the following, the variations in signal reconstruction efficiency are shown as a function of the ALP mass, for both directions of the cut variation. Plots for the following

---

[5] In particular, the discrepancy has been evaluated on the entire squared mass spectrum rather than per single mass window, to be the most conservative.

relevant variables are shown:

- Angular acceptance: the polar angle selection is varied by propagating the effect of 2 mm shift of the ECL; effects on the signal efficiency are shown in Fig. 7.14;
- Photon energy: the energy selection is varied by propagating the effect of a 1% shift in the energy cut; effects on the signal efficiency are shown in Fig. 7.15;
- Zernike of most isolated photon: the selection cut of the `ZernikeMVA` variable of the most isolated photon is varied by ±0.05 (over a range between 0 and 1), consistent with shape differences between data and MC; effects on the signal efficiency are shown in Fig. 7.16;
- Angular separation: the selection cut of the angular difference variables is varied[6] by 5 mrad, consistent with position resolution of the ECL; effects on the signal efficiency are shown in Fig. 7.17;
- Three-body invariant mass: from internal studies on the data set used for this analysis, it resulted that the beam energy $\sqrt{s}$ was known with an uncertainty of 0.2%. Therefore, the $m_{\gamma\gamma\gamma}$ cut is varied by $\pm 0.002\sqrt{s}$; effects on the signal efficiency are shown in Fig. 7.18.

Regarding the last point, the uncertainty on $\sqrt{s}$ may, in principle, affect other aspects of the analysis: it enters in the equation defining the recoil mass (Eq. 2.12) and is an input for the computation of $g_{a\gamma\gamma}$ (Eq. 2.7).

For what concerns the recoil mass calculation, if the beam energy is shifted, then the mass values would be shifted as well. This effect is smaller than the energy resolution and is ignored during limit extraction.

For what concerns the dependence of the $g_{a\gamma\gamma}$ coupling to $\sqrt{s}$, the coupling has been recalculated with changed beam energy (see Fig. 7.19). The resulting shift is less than 3% at the highest masses, so it is negligible with respect to statistical uncertainty.

Modifying selection requirements on t_c and the number of ECL crystal hits `clusterNHits` is found to have a negligible effect.

All the different sources of uncertainty on the signal efficiency described above are combined via summing them in quadrature, and then propagated to the final fits. The combined uncertainty on the signal efficiency is shown in Fig. 7.20.

---

[6] Simultaneously, using the combination with maximal differences.

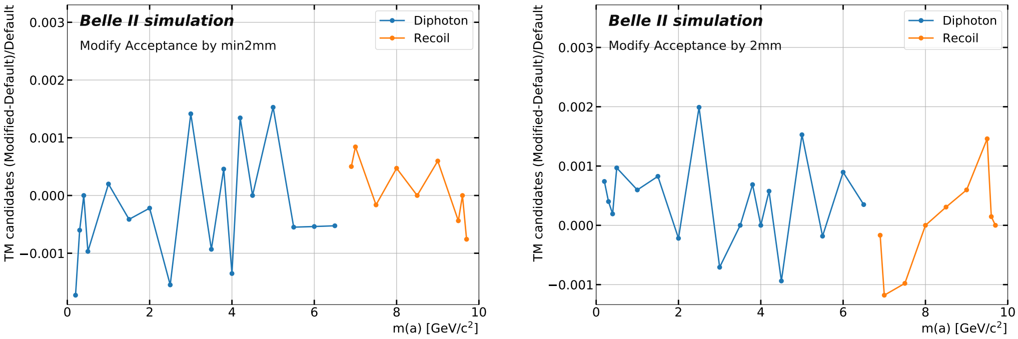

(a) Selection -2 mm.

(b) Selection +2 mm.

Figure 7.14: Change of signal efficiency by indirect variation of the angular selection, evaluated on MC signal samples.

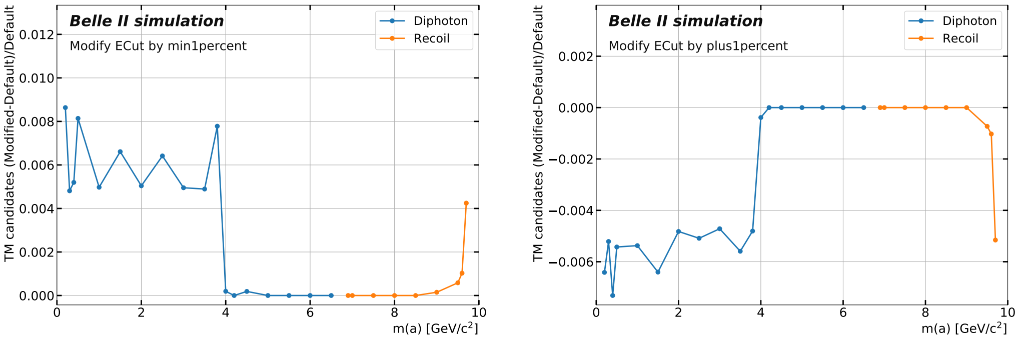

(a) Selection -1%.

(b) Selection +1%.

Figure 7.15: Change of signal efficiency by variation of the energy selection, evaluated on MC signal samples.

All the cut variations discussed in the previous Sections that contribute to the total uncertainty on the signal efficiency are applied only to the TM component of the signal. The effects on the combinatorial component of the signal and on the backgrounds are neglected, as they are second-order effects on one of the three components of the systematic uncertainty.

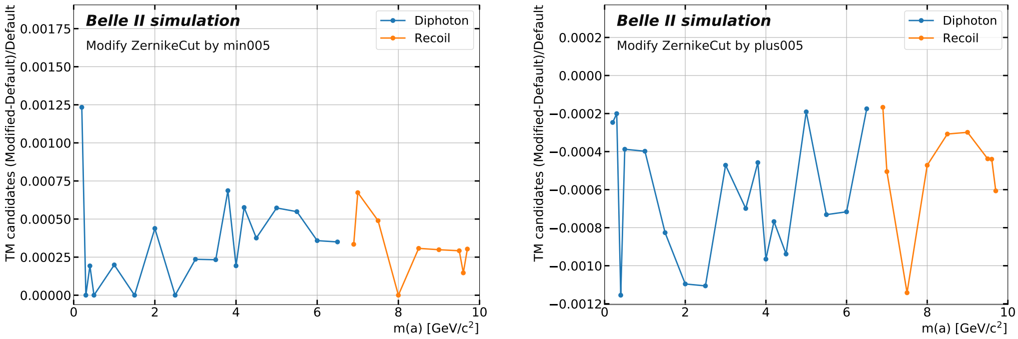

(a) Selection -0.05. (b) Selection +0.05.

Figure 7.16: Change of signal efficiency by variation of the `ZernikeMVA` selection, evaluated on MC signal samples.

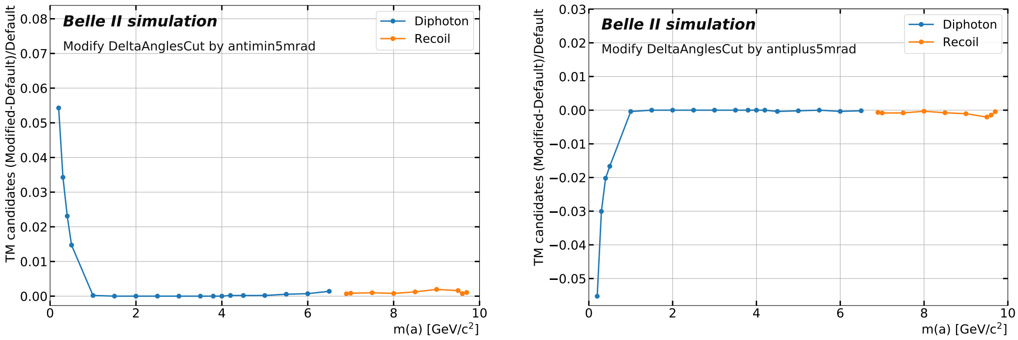

(a) $\Delta\theta$ -5 mrad, $\Delta\phi$ +5 mrad. (b) $\Delta\theta$ +5 mrad, $\Delta\phi$ -5 mrad.

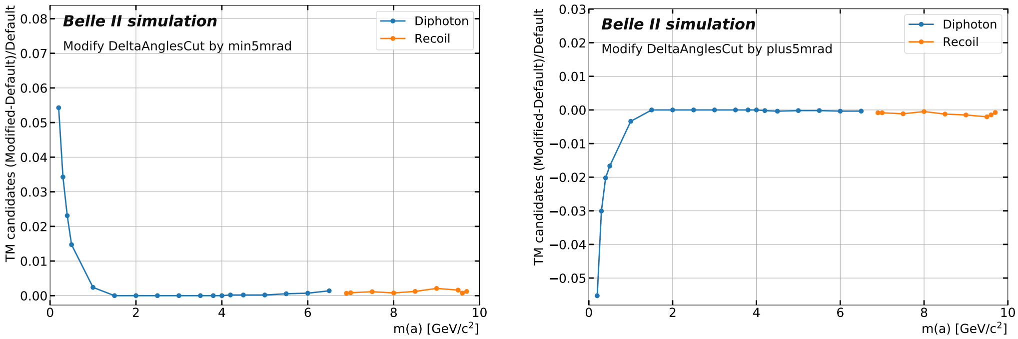

(c) $\Delta\theta$ -5 mrad, $\Delta\phi$ -5 mrad. (d) $\Delta\theta$ +5 mrad, $\Delta\phi$ +5 mrad.

Figure 7.17: Change of signal efficiency by variation of the angular difference ($\Delta\theta$ and $\Delta\phi$) selection, evaluated on MC signal samples.

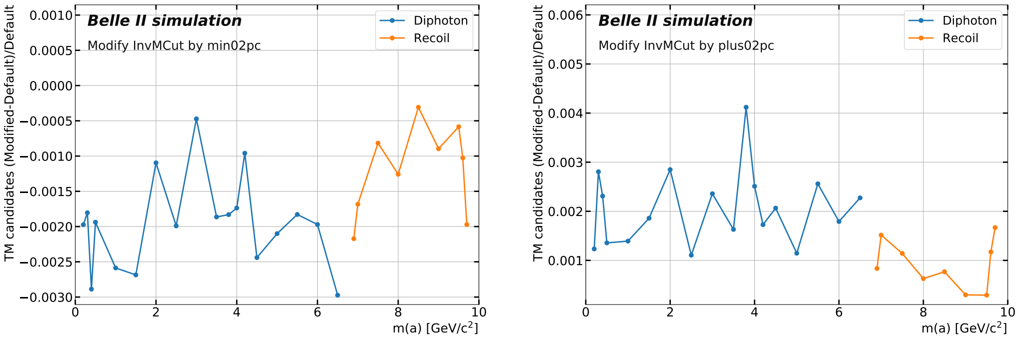

(a) Selection -0.2%.

(b) Selection +0.2%.

Figure 7.18: Change of signal efficiency by variation of the $m_{\gamma\gamma\gamma}$ selection, evaluated on MC signal samples.

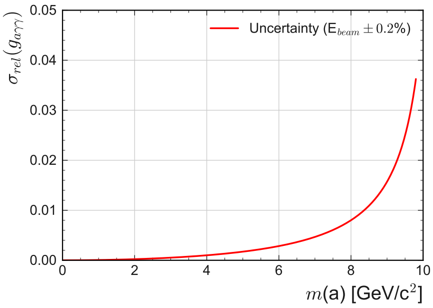

Figure 7.19: Relative change of the coupling $g_{a\gamma\gamma}$ from a variation of the beam energy by 0.2%.

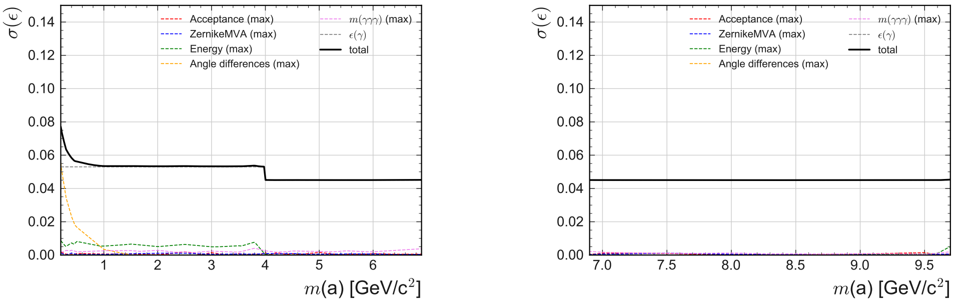

(a) Diphoton (max).

(b) Recoil (max).

Figure 7.20: All systematic uncertainties of the signal efficiency using the average absolute plus and minus uncertainties.

7.5.3 Background shape and fit range

To estimate the systematics associated with background shape (order of the polynomial used to fit the background) and fit range (width of the mass window used to perform the fit), other significance scans and UL extractions are performed with modified vales of these parameters. In particular, following from the discussion in Sec. 7.3, the polynomial order is modified by ±1 and the fit range by $\pm 5\sigma_{CB}$ (per side). This means that the significance scan and UL extraction are performed nine times:

- Default polynomial order, default fit range;
- Polynomial order reduced by one, default fit range;
- Polynomial increased by one, default fit range;
- Default polynomial order, fit range enlarged by $5\sigma_{CB}$ per side;
- Polynomial order reduced by one, fit range enlarged by $5\sigma_{CB}$ per side;
- Polynomial increased by one, fit range enlarged by $5\sigma_{CB}$ per side;
- Default polynomial order, fit range shrunk by $5\sigma_{CB}$ per side;
- Polynomial order reduced by one, fit range shrunk by $5\sigma_{CB}$ per side;
- Polynomial increased by one, fit range shrunk by $5\sigma_{CB}$ per side.

During each of these 9 extractions, the nuisance parameters to model the uncertainties on signal efficiency and resolution, discussed in the previous Sections, are always applied. Ultimately, for each scanned mass point, the worst extracted value (i.e., lowest significance and highest UL) amongst the nine is taken. The significance scan and the ULs shown in Sec. 7.6.2 are obtained with this procedure.

7.5.4 Other systematics and cross checks

The following sources of systematics have been investigated, but they resulted in negligible contributions to the total uncertainty.

- KDE smoothing parameter (Sec. 7.2.2): extractions with a smoothing parameter modified by ±0.25 have been performed;
- Luminosity: the uncertainty associated to the recorded luminosity in the unblinded data set is 0.7%. This uncertainty is negligible;

- Beam energy $\sqrt{s}$: as discussed in Sec. 7.5.2, the effects of the uncertainty on the beam energy on the signal extraction (in form of recoil mass and coupling computation) are negligible;

- Beam background levels: the MC simulation for early *Phase 2* was known to not perfectly reproduce the measured background distributions. Tests were performed with higher-nominal-background MC samples, and slightly worse energy resolution were found, as expected. This uncertainty is covered by the floating background model and by the signal reconstruction efficiency uncertainty;

- Trigger efficiency: the trigger efficiency had been found to be 100% with an uncertainty of sub-permill level (Eq. 5.13) so this contribution is negligible;

- Cluster efficiency: an additional cross-check on the efficiency of clustering has been performed. This study uses radiative electron pairs and deploys a tag-and-probe method[7]. The results are compatible with the uncertainty on signal efficiency, and the uncertainty on the data/MC ratio is about 1%, hence a negligible correction;

- Event generator: the dominant background contribution is due to $e^+e^- \to \gamma\gamma(\gamma)$ events generated with the `BABAYAGA.NLO` generator (Sec. 4.3.2). An uncertainty of 1% for the cross section value (internal studies) is used, and this uncertainty is negligible.

7.6 Unblinding

As explained in Sec. 6.4, 89.6% of the total available data have been used for the signal extraction. This corresponds to an integrated luminosity of 445 pb^{-1}. In the following, the data/MC comparison and the results of the signal extraction, in the form of significance scan and extracted UL, are shown.

7.6.1 Data/MC comparison

The data/MC comparison for all selection variables is shown in Figs. 7.21 to 7.30.

---

[7] One of the two electron tracks must be associated to a cluster in the ECL: this is the tag. The other track may, but doesn't have to, be associated to another ECL cluster: this is the probe. The clustering efficiency is estimated as the ratio between the number of events with a track reaching the ECL that also have an associated ECL cluster, and the total number of events with a track hitting the ECL.

The corrections described in Sec. 7.5.1, that is the bias correction factor and the resolution smearing factor, are applied to all variables which are computed starting from photon energies (E_γ, $M^2_{\gamma\gamma}$, M^2_{recoil}, $M_{\gamma\gamma\gamma}$). This means that, as done for the partial unblinding (Sec. 6.4), the data photon energies are corrected using the bias correction factor, and the MC photon energies are smeared with a Gaussian smearing having mean zero and width equal to the resolution smearing factor.

The agreement is excellent for all variables except the timing variable (Fig. 7.29) and the `clusterNHits` (Fig. 7.23), as expected, given that the simulation of these variables was not good in early data taking. Although, the discrepancies have no effects on the selection, because the areas excluded by the cuts have zero or a negligible amount of events.

Additionally, the first bin in the squared diphoton ALP mass distribution (Fig. 7.21), corresponding to masses below 1 GeV^2/c^4, is underestimated by the MC. Zooms of this area, with different binning, are shown in Fig. 7.31. Although, the signal extraction does not depend on the background predictions, given that the background parameters are left floating and fitted during the extraction instead of being fixed by the MC (Sec. 7.3), and any data/MC discrepancy on the background yield has little impact on the final result. Also, and more importantly, the shape of the data/MC discrepancy is *not* ALP-like, i.e. is far wider than an ALP of those masses would be. This can be seen in Fig. 7.32 (with the same binnings as Fig. 7.31), where an arbitrarily huge MC signal is injected on top of MC backgrounds, overlayed with data, to show that indeed the ALP peak is narrower than the broad structure. This confirms that this discrepancy cannot be due to an ALP.

Data and MC agree well in terms of total normalization, with a 4% discrepancy at most. The data/MC ratios are summarized in Table 7.3, for the two different cuts on photon energy, and by looking both at the whole mass range and only at $M^2(\gamma\gamma) > 1$ GeV$^2/c^4$, i.e. beyond the discrepancy at low mass values.

To account for the yield discrepancies, these are propagated as uncertainty for the photon efficiency, as described in Sec. 7.5.2. The propagation is performed via the linear sum of the central value of the discrepancy and its uncertainty, to be the most conservative. They are used the values obtained on the full $M^2(\gamma\gamma)$ range, which are bigger, to be the most conservative. This results in a step at $m_a = 4$ GeV/c^2, where the selection changes at that mass value: for $m_a < 4$ GeV/c^2 there is a 5.3% total uncertainty on the signal efficiency, while for $m_a > 4$ GeV/c^2 it is 4.5%.

	$E(\gamma) > 650$ MeV	$E(\gamma) > 1000$ MeV
Whole mass range	Data: 15777 ± 126 MC: 15225 ± 54 Data/MC: 1.036 ± 0.009	Data: 11625 ± 108 MC: 11156 ± 46 Data/MC: 1.042 ± 0.011
$M^2(\gamma\gamma) > 1$ GeV$^2/c^4$	Data: 15062 ± 123 MC: 14626 ± 53 Data/MC: 1.030 ± 0.009	Data: 11175 ± 106 MC: 10798 ± 45 Data/MC: 1.035 ± 0.011

Table 7.3: Yields and ratios for data and MC, for the two energy cuts and two considered mass regions.

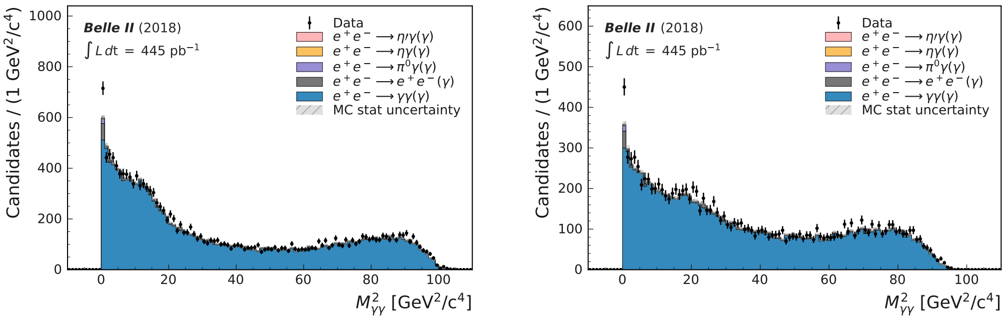

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.21: Data/MC comparison for squared diphoton ALP mass.

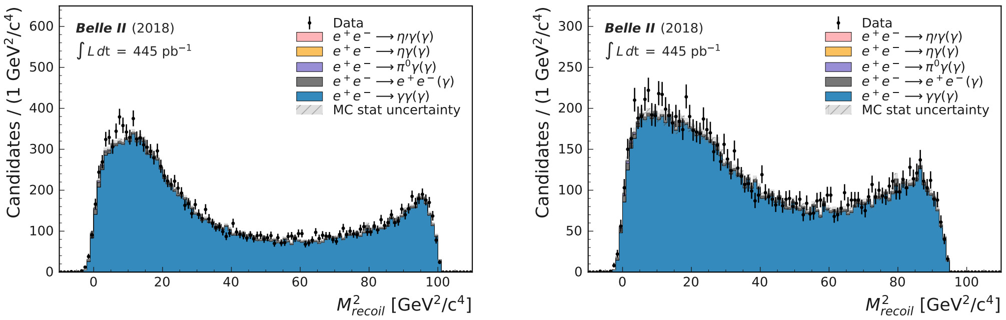

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.22: Data/MC comparison for squared recoil ALP mass.

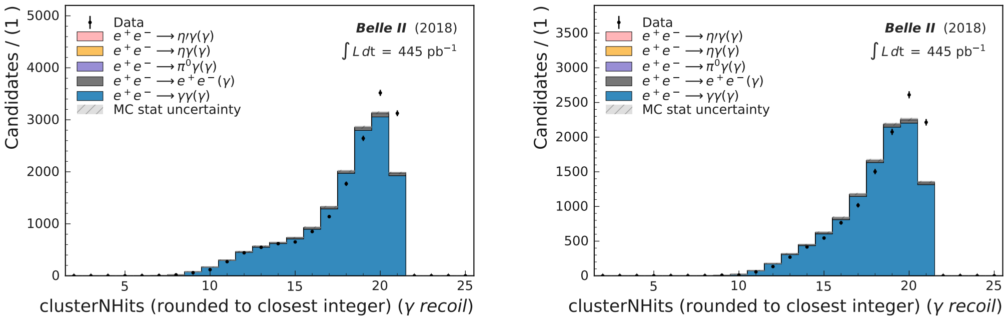

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.23: Data/MC comparison for `clusterNHits`, recoil photon.

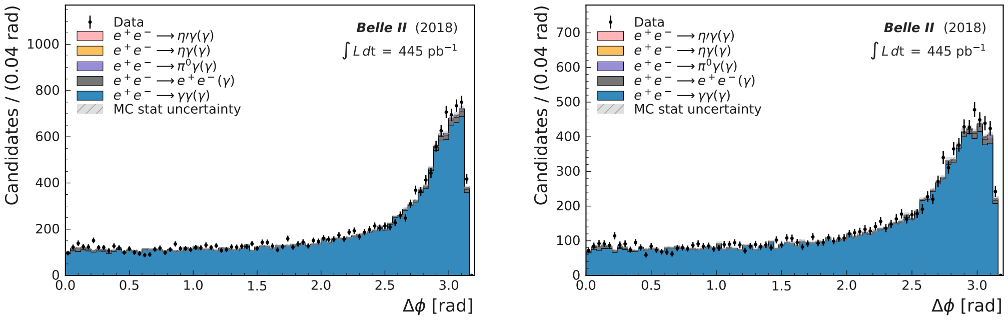

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.24: Data/MC comparison for $\Delta\phi$.

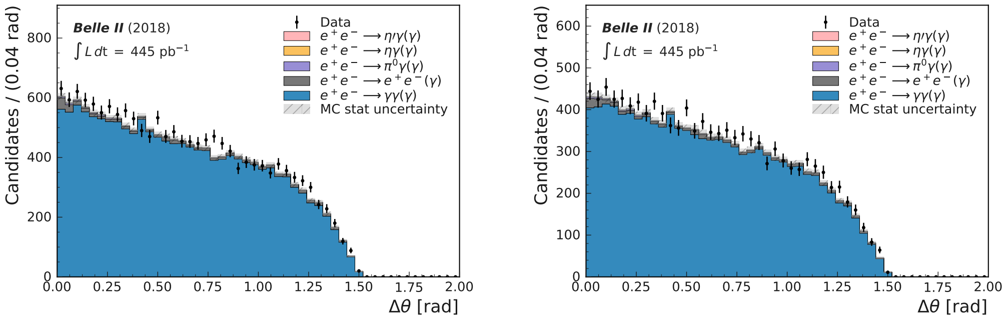

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.25: Data/MC comparison for $\Delta\theta$.

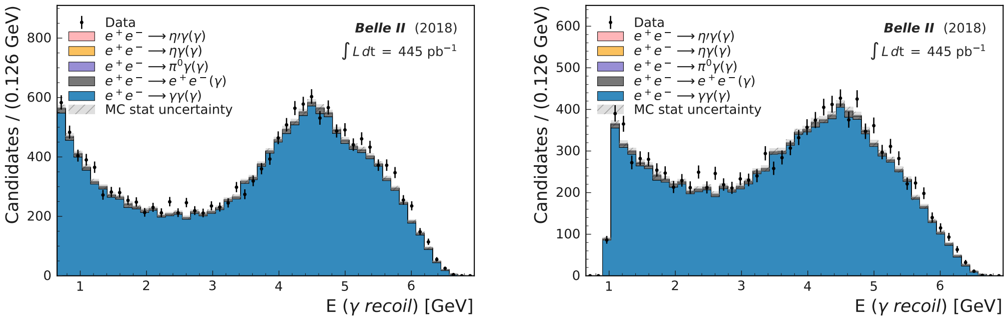

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.26: Data/MC comparison for photon energy, recoil photon.

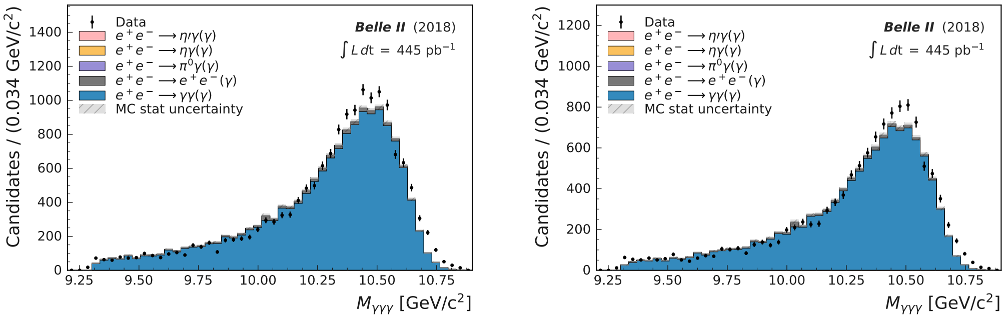

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.27: Data/MC comparison for $m_{\gamma\gamma\gamma}$.

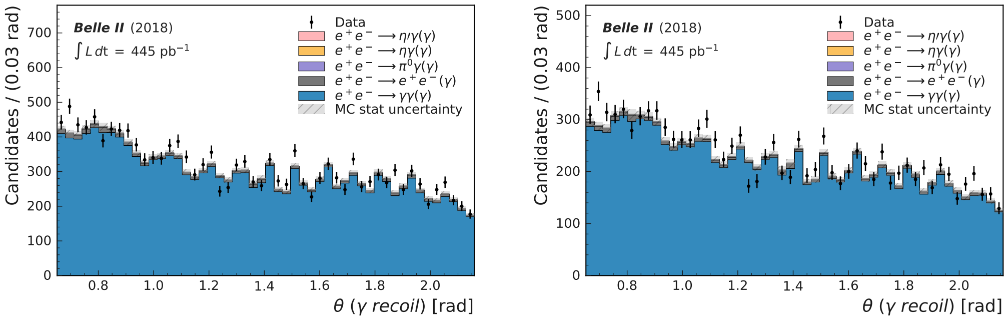

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.28: Data/MC comparison for polar angle, recoil photon.

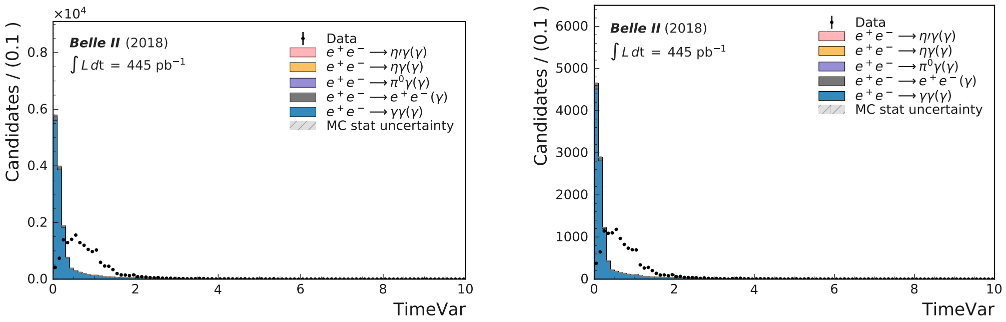

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.29: Data/MC comparison for time variable.

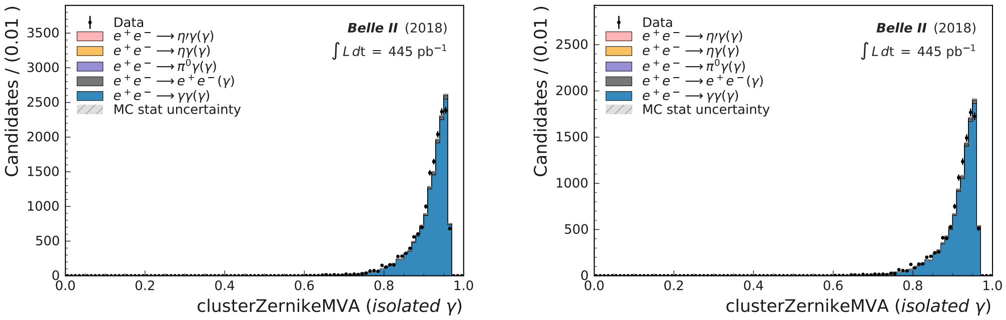

(a) $E_\gamma \geq 0.65$ GeV selection.

(b) $E_\gamma \geq 1$ GeV selection.

Figure 7.30: Data/MC comparison for `ZernikeMVA` of the most isolated photon.

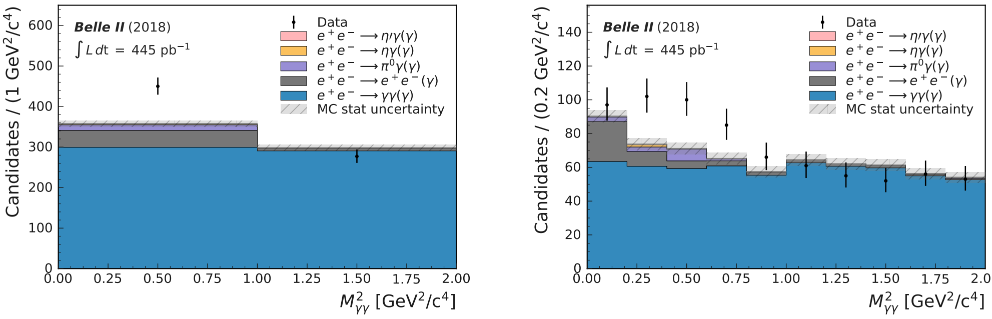

(a) 2 bins. Same binwidth as Fig. 7.21.

(b) 10 bins.

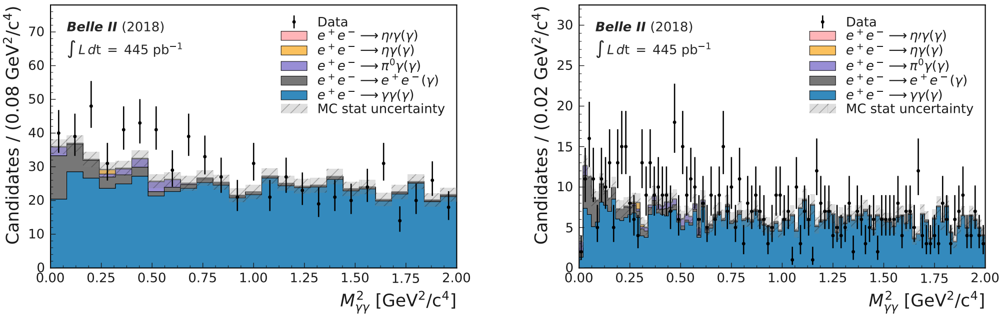

(c) 25 bins.

(d) 100 bins.

Figure 7.31: Data/MC comparison for squared diphoton ALP mass. $M_{\gamma\gamma}^2 < 2$ GeV^2/c^4.

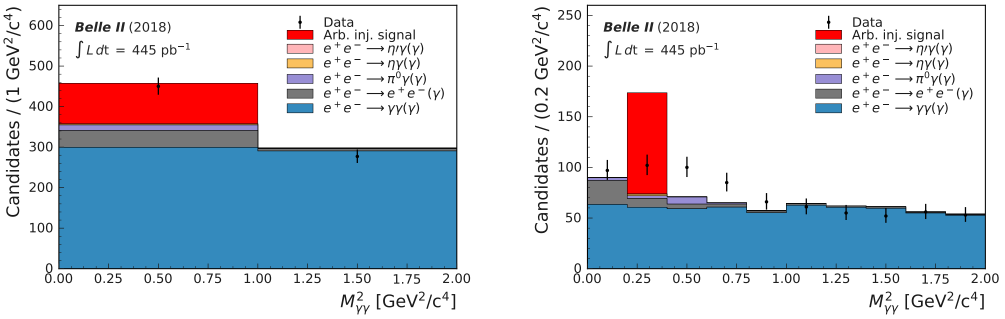

(a) 2 bins. Same binwidth as Fig. 7.21..

(b) 10 bins.

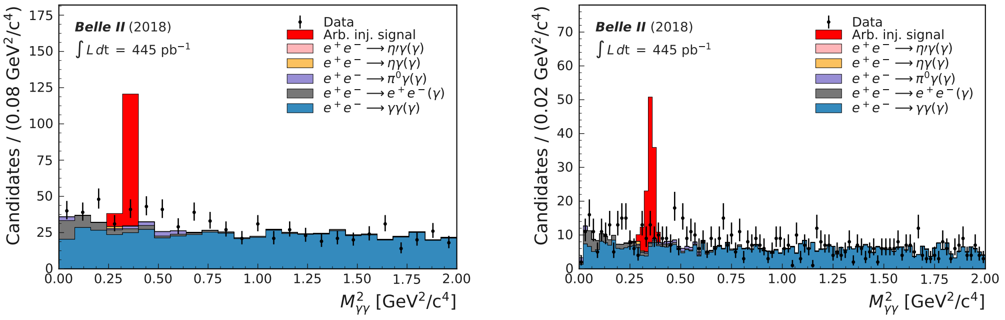

(c) 25 bins.

(d) 100 bins. Useful to better see the narrow signal peak structure.

Figure 7.32: Data/MC comparison for squared diphoton ALP mass. $M^2_{\gamma\gamma} < 2$ GeV^2/c^4. Arbitrarily large MC signal injected.

7.6.2 Upper limits extraction

In the following, the results of the mass scan are reported, in the form of significance scan and UL extraction. Systematic uncertainties are propagated as described in 7.5. More details on their effect on the significance and ULs are shown in Appendix A, in Figs. A.1 to A.4.

The maximum local significance, after taking systematics into consideration, is

$$\text{max significance} = 2.8\sigma \tag{7.10}$$

for

$$m_a\,(\text{max significance}) = 0.477\ \text{GeV/c}^2. \tag{7.11}$$

Fig. 7.33 shows the signed significance as a function of the ALP mass m_a. This plot has been obtained performing an extraction while allowing negative signal yields, after taking systematic uncertainties into account. The values plotted correspond to the significance multiplied by the sign of the cross section obtained for that mass point. As previously mentioned, this plot is produced to verify the goodness of the fit, by checking that there are no discontinuities in the signal extraction. To set limits, the default extraction where only positive signal yields are allowed is used. Given that there are no points with a local (let alone a global) significance of 3 or more sigmas, a 95%-confidence-level upper limit is set for the channel investigated in this analysis. Fig. 7.34 shows the 95%-confidence-level UL on the cross section. Fig. 7.35 shows the same 95%-confidence-level UL converted into coupling constant, and is compared with previous results obtained by other experiments.

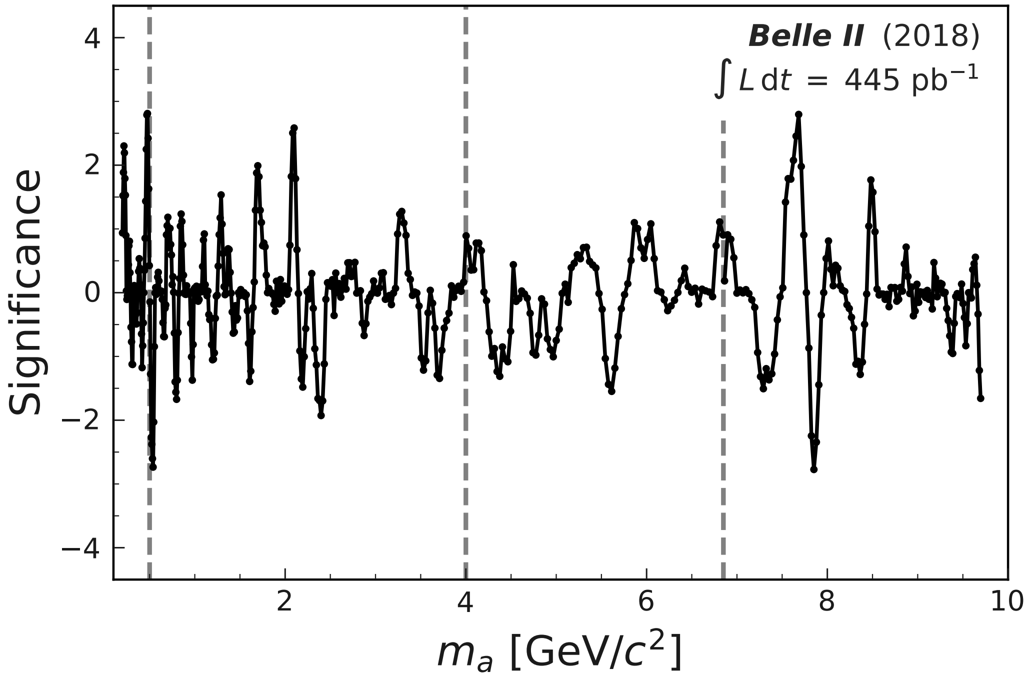

Figure 7.33: Signed significance, after considering systematics, as a function of the fitted ALP mass.

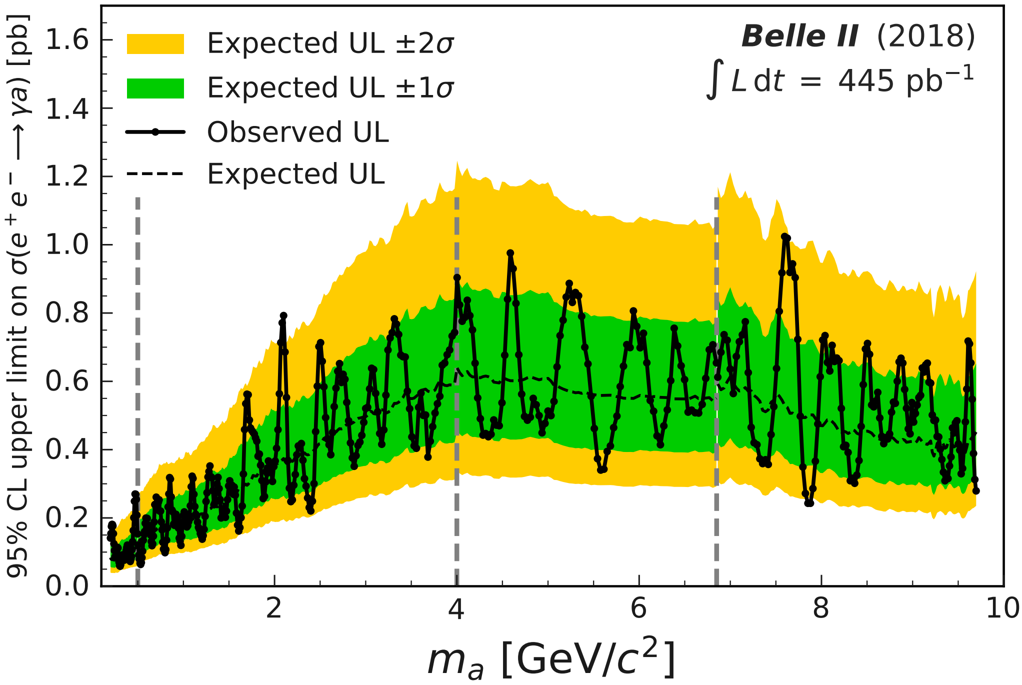

Figure 7.34: 95%-confidence-level upper limit on the cross section for $e^+e^- \to \gamma a$.

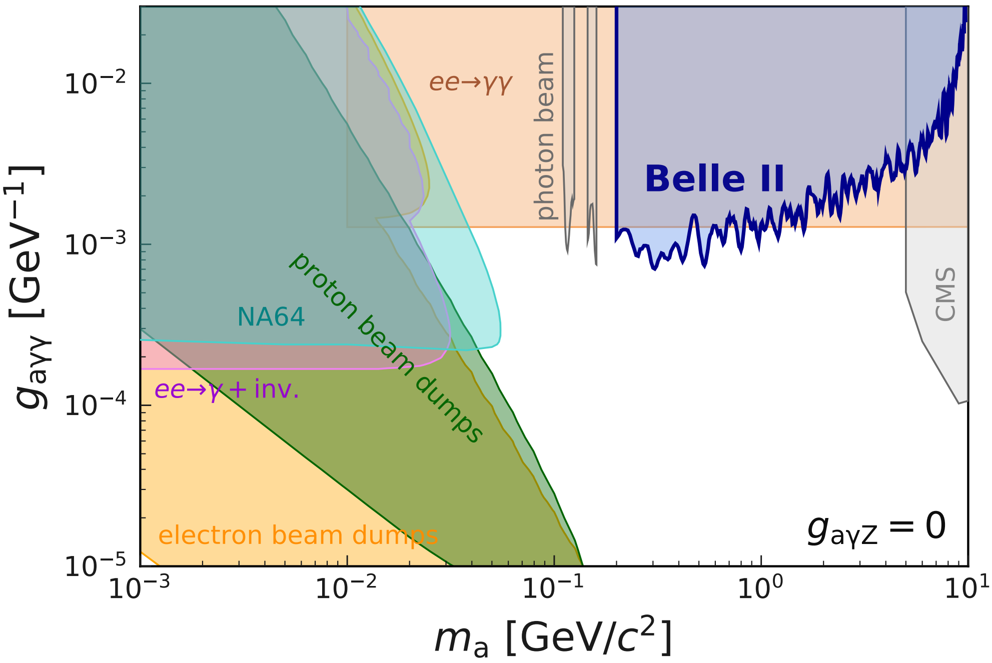

Figure 7.35: 95%-confidence-level upper limit on the ALP-photon coupling $g_{a\gamma\gamma}$ from Belle II and previous constraints from electron-beam-dump experiments [27, 34], $e^+e^- \to \gamma + \text{invisible}$ [20, 30, 29, 31, 32], NA64 [37], proton-beam-dump experiments [35, 36, 28], $e^+e^- \to \gamma\gamma$ [33], photon-beam experiment [39], and heavy-ion collisions [38]. Compare with Fig. 2.4. Adapted from [1] by updating the limits from proton-beam-dump experiments due to an erratum on [26], that appeared after [1].

Chapter 8

Photon covariance matrix

As anticipated in Sec. 7.1.3, by performing a *kinematic fit* of the three photons of the event it would be possible to avoid the double mass definition. The kinematic fit would constrain the invariant mass of the three photons to be equal to $\sqrt{s}$. Although, in order to achieve so, a precise knowledge of the so-called *photon covariance matrix (PCM)* is needed. This was not available for the data used for the ALP analysis, due to the little data collected. The study is ongoing at the time of writing. This Chapter presents its details, status, and prospects. In this Chapter, *smearing* and *resolution* are used interchangeably, unless otherwise specified.

Section 8.1 explains the structure of the PCM, the study workflow, and the results needed to be achieved.

Section 8.2 illustrates the single-photon study approach, which is MC only.

Section 8.3 shows the results of the studies performed on radiative dimuon samples, $e^+e^- \to \mu^+\mu^-\gamma$, which can be used to extract photon resolution from data rather than from MC.

Section 8.4 summarizes the status of the studies and the prospects.

8.1 Motivations

Belle II is placed on an e^+e^- collider, so the initial state is well known. In particular, the four-momentum of the collision is known: in the cms, this is a motionless virtual photon with energy equal to $\sqrt{s} = 10.58$ GeV/c^2. In any fully contained event, i.e. without missing energy, the decay products all together must have an invariant mass equal to $\sqrt{s}$. In the ALP analysis, this constraint was indirectly applied by requiring that $m_{\gamma\gamma\gamma}$ is close enough to $\sqrt{s}$ (Secs. 5.2.4 and 5.5). This knowledge is also used in the definition of the

ALP recoil mass, Eq. 2.12, which includes $\sqrt{s}$.

In the ALP analysis, the diphoton mass definition was used for $m_a < 6.85$ GeV/c^2; for higher masses, the recoil mass definition is used (see Sec. 7.1.3). The transition point is chosen such that the sensitivity is maximized throughout all the mass range. This is given both by the narrowness of the signal, parametrized with the value σ_{CB} (Fig. 7.13), and by the background shape and yield. About the former, the diphoton mass signal peak gets wider with increasing ALP mass because it is constituted by two photons, whose energies increase with the ALP mass: the photon energy uncertainty increases roughly linearly with the photon energy, so their propagated uncertainty increases together with the ALP mass. The recoil mass signal peak, instead, gets narrower with increasing ALP mass because it is computed via the recoil photon, whose energy, and thus uncertainty, gets smaller at higher ALP masses.

The narrowness of the signal peak can be improved by using a kinematic fit, particularly in the intermediate region (around $m_a \sim 5$ GeV/c^2), where neither the diphoton nor the recoil definitions excel (see Figs. 8.1c and 8.1d). Incidentally, for the same reason, the improvement in ALP mass resolution is not as relevant for very light and very heavy ALPs. It must be noted that the improvement to the overall signal sensitivity is not straightforwardly related to the improvement in signal mass resolution, because it also depends on the background shape and yield. Hence, for a full study, the kinematic fit has to be applied to the background too, and the background modeling described in Sec. 7.3 has to be performed again, and with it the whole signal extraction. Although, as can be seen in Fig. 8.2, the background shape after applying the kinematic fit, for as much as a preliminary version, is locally similar to the shape obtained without kinematic fit, so little differences are expected for the background.

8.1.1 Kinematic fit with photons

The principle behind the kinematic fit is to apply the constraint that, in an event with no missing energy, the decay particles must have a total invariant mass equal to $\sqrt{s}$. This is technically done by varying the particles' four-momenta so that the total invariant mass exactly equals $\sqrt{s}$. The allowed variation is bigger for four-momenta with higher uncertainty. Hence, precise knowledge of the uncertainties is needed to perform the kinematic fit.

The uncertainty on the components of the four-momentum, i.e. the variances, as well as the covariances, are grouped in the PCM. This PCM is filled with the variances and covariances of the *residuals*, i.e. the difference between the *reconstructed* and the *generated* (or *true*) components of the four-momentum. The covariance between two variables A and B of a sample is defined as:

$$cov(A, B) = \frac{1}{N-1} \sum_{i=1}^{N} \left(A_i - \bar{A}\right) \left(B_i - \bar{B}\right) , \tag{8.1}$$

where N is the number of events in the sample, the index i runs over them, and $\bar{X}$ is the mean value of the variable X. The variance of a variable is the covariance of that variable with itself:

$$var(A) = cov(A, A) = \frac{1}{N-1} \sum_{i=1}^{N} \left(A_i - \bar{A}\right)^2 . \tag{8.2}$$

Given that photons have mass zero, their four-momenta have effectively only three free parameters. The used parametrization is thus the photon energy E, polar angle θ, and azimuthal angle ϕ, which fully characterize a photon. The PCM is symmetric because the covariance is an even operator: $cov(A, B) = cov(B, A)$. This 3x3 matrix is the object technically needed to perform the kinematic fit.

The residual of the quantity A is labeled with ΔA :

$$\Delta A = A^{\text{rec}} - A^{\text{true}} \tag{8.3}$$

where $A^{\text{rec[true]}}$ is the reconstructed [true] value of A, i.e. the value reconstructed by the detector, be it the real or the simulated one [the value possessed by the real particle or generated at simulation level]. The diagonal elements of the PCM are the variances of ΔE, $\Delta\theta$, and $\Delta\phi$, or, in other terms, the *resolutions* of E, θ, and ϕ. Given that the PCM is symmetric, only half of out-of-diagonal elements are relevant, the other half being redundant: $cov(\Delta E, \Delta\theta)$, $cov(\Delta E, \Delta\phi)$, and $cov(\Delta\theta, \Delta\phi)$.

The most relevant elements for the proper error handling in the kinematic fit are the diagonal ones. Amongst these, the energy resolution is the most impactful. Therefore, the *photon energy resolution* is the first quantity studied.

8.1.2 Mass resolution improvement

At the time of writing, the PCM values are not calculated exactly in the `basf2` code, but a basic version with approximate values is available. This approximate preliminary PCM can be used to qualitatively show the effects of a kinematic fit. Proof-of-concept fits are shown in Fig. 8.1, where the masses distributions for different m_a values are shown with and without application of the kinematic fit. It can be seen that the distributions obtained using the kinematic fit are indeed narrower, particularly for high values of m_a when compared to the diphoton mass, and vice versa for the recoil one.

Fig. 8.2 shows the (squared, to better compare with Figs. 7.21, 7.22, and similar) masses distribution for the main background, $e^+e^- \to \gamma\gamma(\gamma)$, using the diphoton and recoil definition and using the available kinematic fit. As expected, for low ALP candidate masses the distribution obtained applying the kinematic fit is close to the diphoton distribution, while for high masses is close to the recoil distribution.

The metric to quantify the optimization of the PCM is the following. Let's first define the pull of a quantity A as:

$$P(A) = \frac{A^{\text{rec}} - A^{\text{true}}}{\sigma_A} = \frac{\Delta A}{\sigma_A} \tag{8.4}$$

where σ_A is the uncertainty on that quantity, and the second equality uses the resolution definition from Eq. 8.3. If all uncertainties were Gaussian and properly estimated, the pulls of a variable should be distributed *normally*, meaning following a Gaussian with mean zero ($\mu = 0$) and standard deviation one ($\sigma = 1$).

The ultimate goal of this study is to build a set of PCMs that make the pulls $P(m_a)$ as normal as possible for all values of the ALP mass m_a. The ALP model $e^+e^- \to \gamma a(a \to \gamma\gamma)$ is an excellent testing ground, because it satisfies the condition of having no missing energy and it has a fully neutral final state, meaning that the effects of the PCM are enhanced; also, the various testable m_a values offer a variety of E_γ values to verify that the PCM is properly built for any value of photon energy.

To verify the normality of the obtained $P(m_a)$ distributions, a simple monitoring plot is deployed. The so-called *normality monitor plot* for the approximate preliminary PCM is shown in Fig. 8.3. This plot is to be redrawn every time the PCMs are updated.

The goal is to parametrize the PCM as a function of E_γ, and if necessary of θ_γ, ϕ_γ, and/or beam background levels, so that the pulls $P(m_a)$ are as normal as possible. The starting values of the PCM will be the actual photon resolution and covariances, which will be then modified in an iterative way to achieve the desired goal. This is why Secs. 8.2 and 8.3 are dedicated to study the photon resolution.

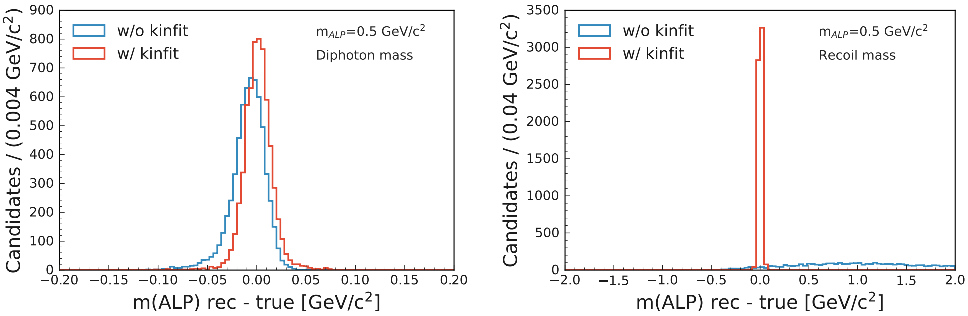

(a) Diphoton, $m_a = 0.5$ GeV/c^2. (b) Recoil, $m_a = 0.5$ GeV/c^2.

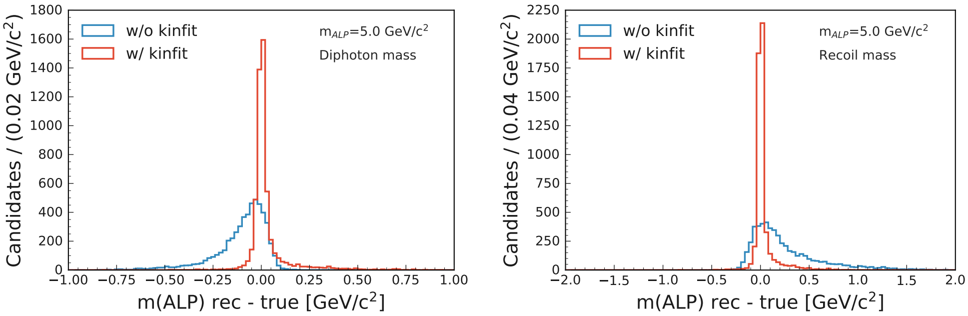

(c) Diphoton, $m_a = 5.0$ GeV/c^2. (d) Recoil, $m_a = 5.0$ GeV/c^2.

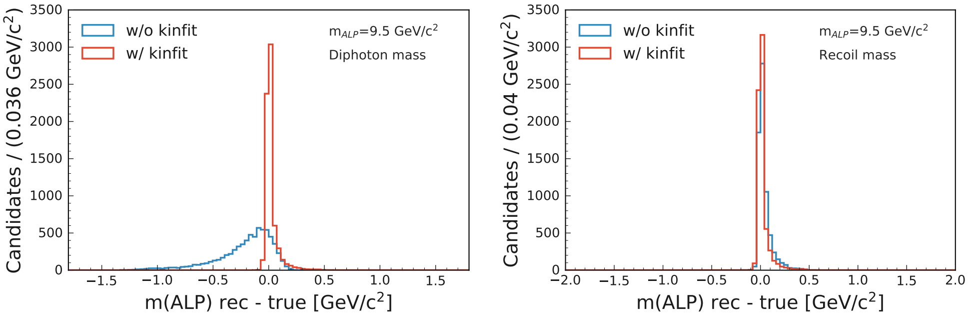

(e) Diphoton, $m_a = 9.5$ GeV/c^2. (f) Recoil, $m_a = 9.5$ GeV/c^2.

Figure 8.1: ALP mass resolution Δm_a, for three mass benchmark values. Comparison between diphoton (left) or recoil (right) mass definitions and ALP mass obtained with the kinematic fit, performed with the approximate preliminary PCM.

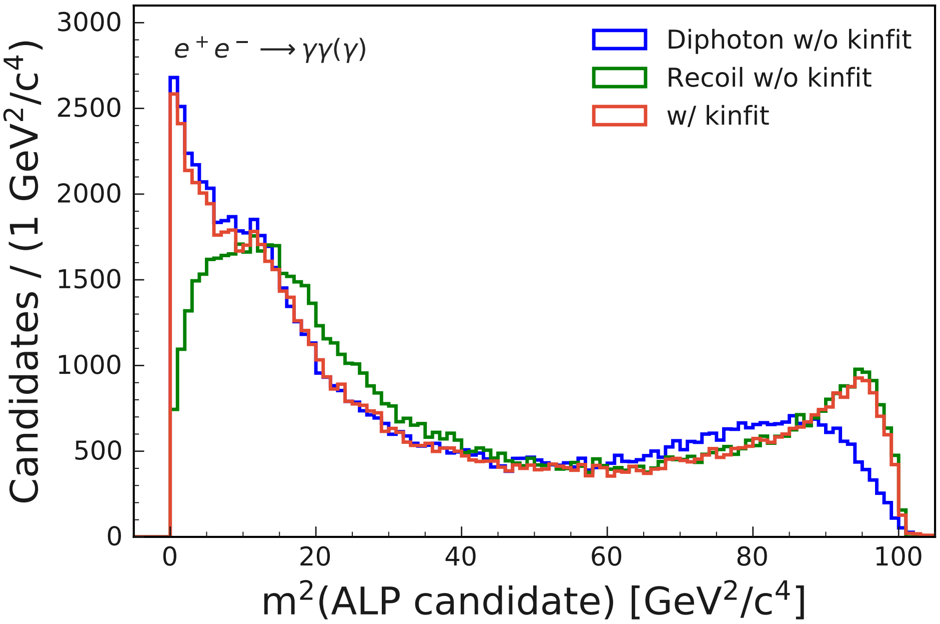

Figure 8.2: ALP candidate squared mass distribution for $e^+e^- \rightarrow \gamma\gamma(\gamma)$ background. Diphoton and recoil masses distributions are obtained without the kinematic fit and compared with the mass distribution obtained with the kinematic fit, performed with the approximate preliminary PCM.

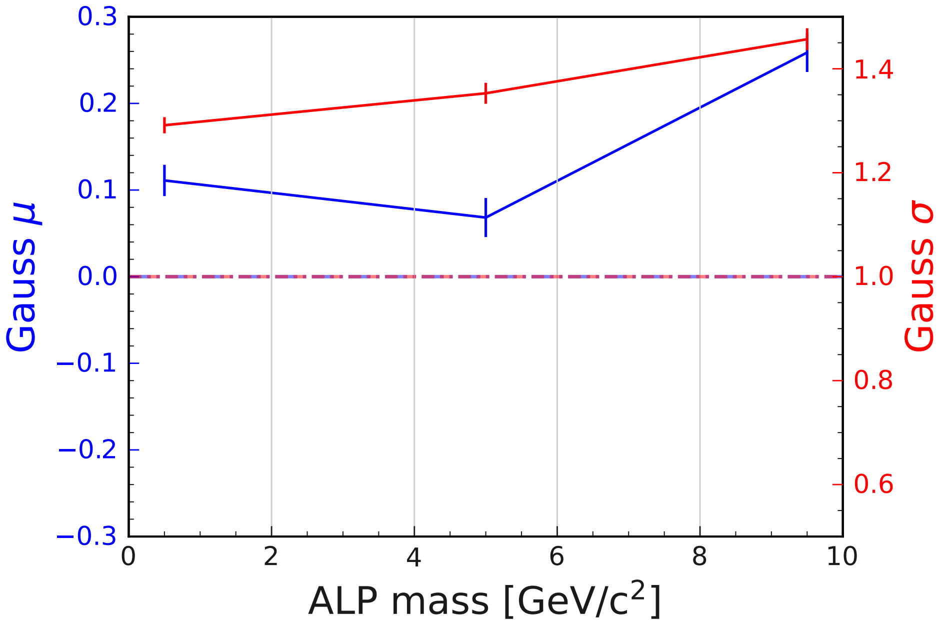

Figure 8.3: Normality monitor plot obtained with the approximate preliminary PCM.

8.2 Single photon studies

Preliminary, MC-only studies can be carried out with single-photon particle-gun MC samples. Each sample is generated at a given fixed E_γ^{true} and in a small polar angle range (within a θ ring, also called `thetaID`, see Sec. 3.4.2). These samples have been used to study directly the true resolutions ΔE, $\Delta\theta$, and $\Delta\phi$. These are the exact quantities, because the true values are accessible, but they cannot take into account any possible data/MC discrepancy. Thus, these studies may work as a first step to obtain some reference values, but other samples need to be used to extract the photon resolution first, and the PCM then, from data (see next Section).

The results from this simple, MC-only study can be used to have a reference value of what to expect for the actual photon resolution. Additionally, they provide a simplified but easily accessible way to define qualitatively the dependence of the photon resolution on the θ and ϕ angles. The results of this preliminary study are summarized in Fig. 8.4.

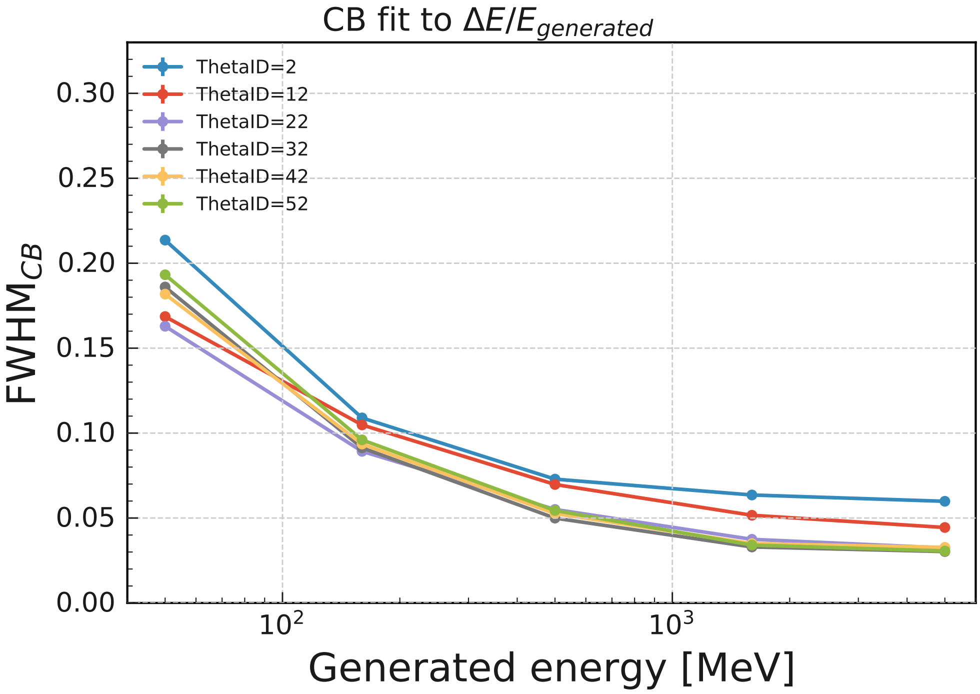

(a) Normalized E resolution.

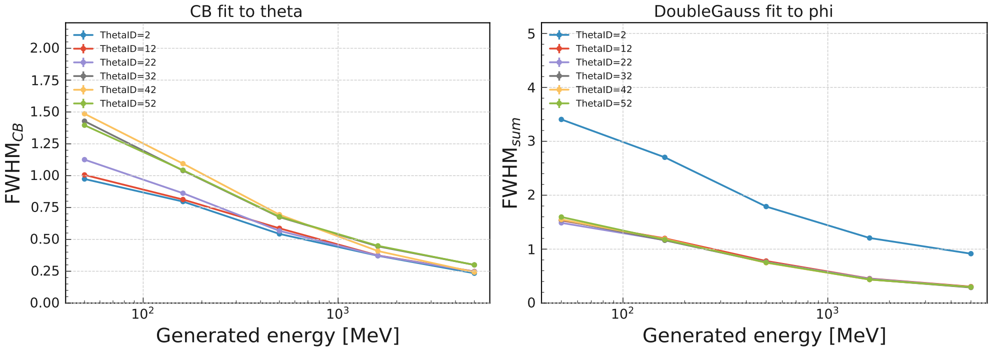

(b) θ resolution.

(c) ϕ resolution.

Figure 8.4: Photon E, θ, and ϕ resolution from single photon MC-only studies.

8.3 Radiative dimuon studies

As mentioned in the previous Section, it is necessary to find a data sample and an analysis procedure that can be used in data too. In data it is not possible to access the true values of photon E, θ, and ϕ. It is therefore necessary to find a process that provides a quantity that acts as a proxy of these values and does not rely on ECL. The channel used for the study shown in this Section is $e^+e^- \to \mu^+\mu^-\gamma$, i.e. radiative dimuon events (see Fig. 8.5 for a sketch of the process).

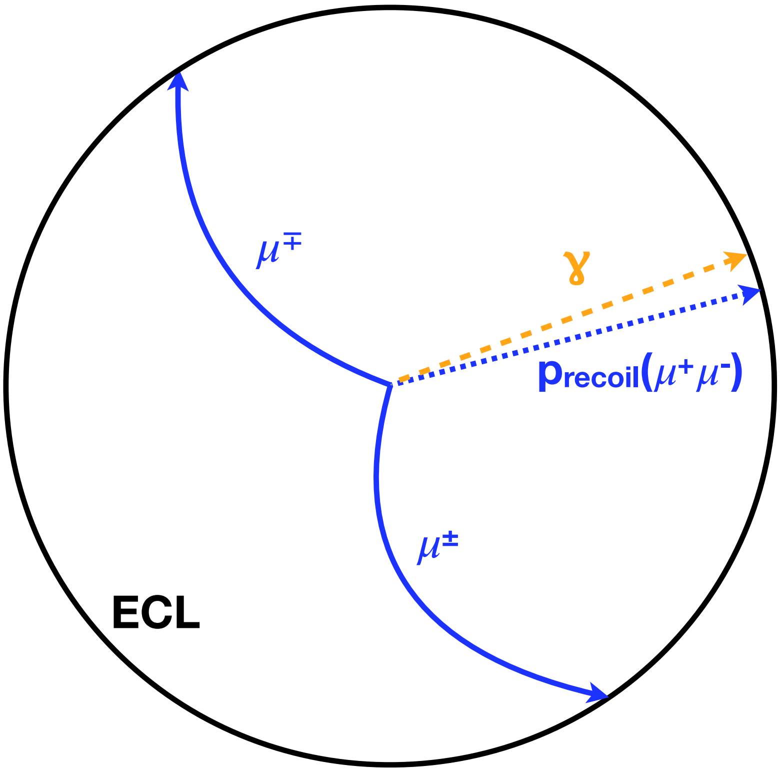

Figure 8.5: Sketch of the radiative dimuon process $e^+e^- \to \mu^+\mu^-\gamma$.

The event is considered to be a *pure radiative* dimuon event if there are no extra ISR [FSR] photons, i.e. photons emitted by the initial [final] state particles, in this case e^+e^- $[\mu^+\mu^-]$, by bremsstrahlung processes. Under the assumption of having a pure radiative dimuon event, the *recoil momentum* $\overrightarrow{p}_{\text{recoil}}(\mu\mu)$ of the dimuon pair equals the photon momentum. $\overrightarrow{p}_{\text{recoil}}(\mu\mu)$ is computed as:

$$\overrightarrow{p}_{\text{recoil}}(\mu\mu) = \overrightarrow{p}(\text{beam}) - \left[\overrightarrow{p}\left(\mu^+\right) + \overrightarrow{p}\left(\mu^-\right)\right], \tag{8.5}$$

where $\overrightarrow{p}$ (beam) is the momentum of the beams. $\overrightarrow{p}$ (beam) is known and equal to $(0,0,0)$ GeV/c in the cms, and $\overrightarrow{p}(\mu^\pm)$ are measured by tracking, and thus is $\overrightarrow{p}_{\text{recoil}}(\mu\mu)$. In the assumption of no extra photons in the event, and given that photons are massless, the following relation holds true:

$$E(\gamma) = p(\gamma) = p_{\text{recoil}}(\mu\mu) = \left|\overrightarrow{p}_{\text{recoil}}(\mu\mu)\right|. \tag{8.6}$$

Tracking resolution is better for low-momentum tracks, because they are bent more by the magnetic field and thus their momenta can be better measured. In $e^+e^- \to \mu^+\mu^-\gamma$ events, the higher is the energy of the photon, the lower is the momenta of the μ tracks. Therefore, the $p_{\text{recoil}}\,(\mu\mu)$ precision improves with higher photon energies. This allows, for $E(\gamma) \gtrsim 2$ GeV, to use it as a proxy for $E^{\text{true}}(\gamma)$.

More specifically, $p_{\text{recoil}}\,(\mu\mu)$ is different from $E^{\text{true}}(\gamma)$ due to tracking resolution effects and the presence of ISR/FSR photons. If these smearing effects are known and can be factored out, the true photon resolution can be extracted. The method to factor out these smearing is detailed in the following.

8.3.1 Radiative dimuon selection

The $e^+e^- \to \mu^+\mu^-\gamma$ selection is the following:

- For the photon(s) in the event:
  - $E(\gamma) > 1.0$ GeV for the most energetic photon in the event;
  - $E(\gamma) < 0.5$ GeV for other photons, if any;
  - `clusterNHits`$>$ 1.5, for the same robustness reasons as for the ALP selection (Sec. 5.3.4);
  - abs(cluster time) $<$ 200 ns;
  - The minimum distance between a photon and the point where a track hits the ECL must be greater than 50 cm;
  - Photon must be in the ECL barrel deprived of the first and last 5°, as for the ALP selection (Sec. 5.2.2);
- For the muons in the event:
  - $0.05 \leq E(\texttt{ECLCluster}(\mu)) \leq 0.25$ GeV;
  - $p(\mu) > 2$ GeV/c;
- $9.8 \leq m_{\mu\mu\gamma} \leq 10.8$ GeV/c^2;
- $m_{\mu\mu} \geq 1.0$ GeV/c^2;
- When comparing $\vec{p}(\gamma)$ and $\vec{p}_{\text{recoil}}\,(\mu\mu)$, the following are required:
  - $-0.10 < \Delta\theta\left(\vec{p}(\gamma), \vec{p}_{\text{recoil}}\,(\mu\mu)\right) < 0.10$ rad;
  - $-0.10 < \Delta\phi\left(\vec{p}(\gamma), \vec{p}_{\text{recoil}}\,(\mu\mu)\right) < 0.10$ rad;
  - $0.5 < \frac{p(\gamma)}{p_{\text{recoil}}(\mu\mu)} < 1.5$;

8.3.2 Factor out the smearing effects

The quantity that can be measured in data is $E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu)$. The quantity needed to evaluate the true ECL photon resolution is $E^{\text{rec}}(\gamma) - E^{\text{true}}(\gamma)$. It is thus necessary to factor out both the tracking resolution smearing and the ISR/FSR smearing (henceforth simply called *ISR smearing*).

The true ECL photon resolution $E^{\text{rec}}(\gamma) - E^{\text{true}}(\gamma)$ is modeled with a floating CB. The ISR smearing is modeled with a template, directly from MC, sampling the distribution $p^{\text{true}}_{\text{recoil}}(\mu\mu) - E^{\text{true}}(\gamma)$, see Sec. 8.3.2.1. The tracking smearing is modeled with a procedure applicable to data, using the uncertainties on $p^{\text{rec}}_{\text{recoil}}(\mu\mu)$, see Secs. 8.3.2.2 and 8.3.2.2.1. The two smearing contributions are then combined together by convolving them. The resulting *measured total smearing*[1] function, which is fixed for the final fit, is convoluted with the aforementioned floating CB and used to fit the $E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu)$ distribution. The width of the fitted CB is the measurement of the true ECL photon energy resolution. For comparison purposes, the $E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu)$ distribution too is fitted directly with a CB, to evaluate the magnitude of the effects of the total smearing.

The following equation summarizes how the overall fit function is split in terms of its components:

$$\begin{aligned} E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu) &= \left(E^{\text{rec}}(\gamma) - E^{\text{true}}(\gamma)\right) * \left(p^{\text{rec}}_{\text{recoil}}(\mu\mu) - E^{\text{true}}(\gamma)\right) = \\ &= \left(E^{\text{rec}}(\gamma) - E^{\text{true}}(\gamma)\right) * \\ &* \left[\left(p^{\text{true}}_{\text{recoil}}(\mu\mu) - E^{\text{true}}(\gamma)\right) * \left(p^{\text{rec}}_{\text{recoil}}(\mu\mu) - p^{\text{true}}_{\text{recoil}}(\mu\mu)\right)\right] = \\ &= \text{true ECL resolution} * [\text{ISR smearing} * \text{tracking smearing}]\,, \end{aligned} \tag{8.7}$$

where $*$ represents convolution.

Each of the steps above is repeated for multiple bins of $p_{\text{recoil}}(\mu\mu)$, equivalent to $E(\gamma)$, in steps of 0.50 GeV/c between 1.25 GeV/c and 6.25 GeV/c. The results are discussed in Sec. 8.3.3.

8.3.2.1 ISR smearing

The ISR smearing is due to the presence of additional photons in the event. In this case, the relation $p_{\text{recoil}}(\mu\mu) = E(\gamma)$ is no longer exact, and the

---

[1] The word *measured* is used to distinguish from the *true* smearing, be that the tracking component or the total one. The true smearing function can be accessed only on MC and it is used to validate the goodness of the measured one (Figs. 8.7, 8.10, and 8.11), which is built to be used on data.

discrepancy increases with the number and energy of the additional photons. Using MC information it is possible to plot $p^{\text{true}}_{\text{recoil}}\,(\mu\mu) - E^{\text{true}}(\gamma)$, i.e. the difference between the true recoil momentum and the true photon energy. It is necessary to rely on MC information to obtain this quantity, which cannot be otherwise reliably accessed with data.

The distribution is plotted for each bin in $p_{\text{recoil}}\,(\mu\mu)$, and modeled by sampling its amplitude for different values of $p^{\text{true}}_{\text{recoil}}\,(\mu\mu) - E^{\text{true}}(\gamma)$. Some examples can be seen in Fig. 8.7. As it can be seen in the Figure, the ISR contribution is much smaller than the tracking one, although it is not negligible.

The ISR smearing is later convoluted with the tracking smearing (see next Section) to obtain the measured total smearing.

8.3.2.2 Tracking smearing

The true tracking smearing is given by $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu) - p^{\text{true}}_{\text{recoil}}\,(\mu\mu)$, i.e. the difference between the reconstructed and the true value of the recoil momentum. This quantity is independent from the presence of ISR photons, and it is due to tracking only. This distribution cannot be directly accessed in data, because $p^{\text{true}}_{\text{recoil}}\,(\mu\mu)$ is not available. It is therefore necessary to find a proxy that models it accurately.

For each event, the quantity $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)$ is computed, together with its uncertainty $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}$. The values $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}$ are computed using the so-called *tracking covariance matrix*, the matrix of variances and covariances of the tracking parameters, used to propagate the uncertainty on the track parameters. Some examples of the distribution of $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}$ can be seen in Fig. 8.6. For each event a normalized Gaussian G is defined:

$$G = \text{Gauss}\left(\mu = 0, \sigma = p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}\right). \tag{8.8}$$

G has mean 0 and width equal to $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}$. All the Gaussians G, one per event of the sample examined, are summed together:

$$T = \sum_j G_j = \sum_j G\left(0, p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}_j\right), \tag{8.9}$$

with j running over all the events in the sample.

The function T is a good proxy of $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu) - p^{\text{true}}_{\text{recoil}}\,(\mu\mu)$ if the errors $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}$ are accurately calculated. At the time of writing the tracking covariance matrix is known to have proper proportions between its elements, but the absolute scaling must be validated. It is necessary to independently

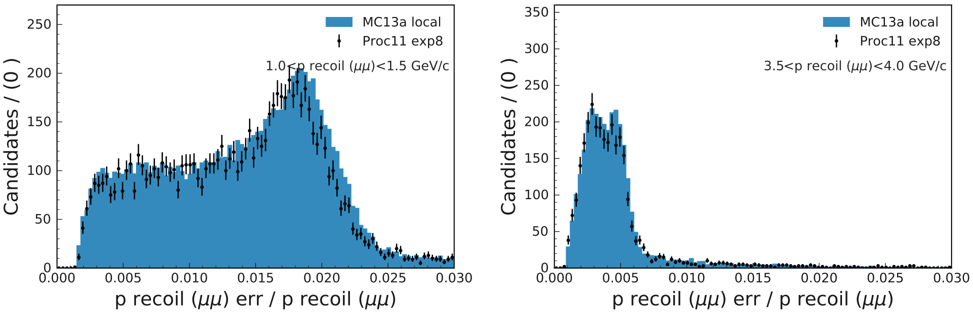

(a) $p_{\text{recoil}}\,(\mu\mu) = 1.25$ GeV/c. (b) $p_{\text{recoil}}\,(\mu\mu) = 3.75$ GeV/c.

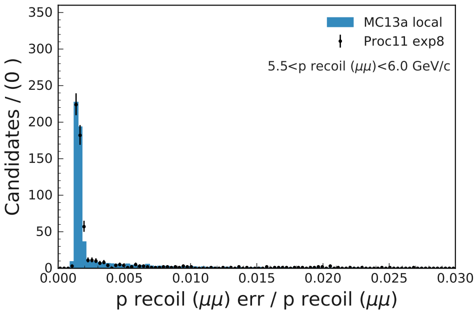

(c) $p_{\text{recoil}}\,(\mu\mu) = 5.75$ GeV/c.

Figure 8.6: Distributions of normalized $p_{\text{recoil}}\,(\mu\mu)$ err, for data and MC. The three plots have the same x range.

evaluate the goodness of the $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)$ err scaling, and apply a correction, if needed. The method to extract these correction factors, henceforth called *scaling factors* S_i, is explained in the following Section, 8.3.2.2.1.

Once such scaling factors are obtained, the proxy function for the tracking resolution can be computed using the scaled $p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}^{\text{ scaled}}$:

$$p^{\text{rec}}_{\text{recoil}}\,(\mu\mu) - p^{\text{true}}_{\text{recoil}}\,(\mu\mu) \approx T^{\text{scaled}} = \sum_j G\left(0, p^{\text{rec}}_{\text{recoil}}\,(\mu\mu)\,\text{err}^{\text{scaled}}_j\right). \tag{8.10}$$

The left term is the true tracking smearing, the right term is its proxy usable in data, the measured tracking resolution (also labeled as p_{recoil} *resolution*).

In Fig. 8.7 the comparison between the true and measured tracking smearing is shown for some values of $p_{\text{recoil}}\,(\mu\mu)$. The ISR smearing (Sec. 8.3.2.1) is shown as well, to better compare the different magnitudes of the two components of the total smearing. The total smearing is shown in Sec. 8.3.3, together with the results of the final fits.

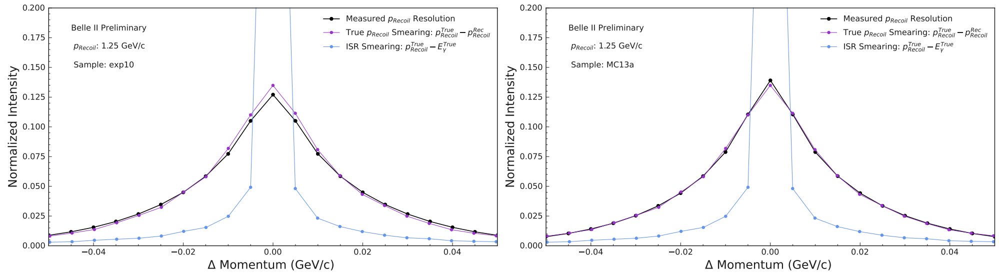

(a) Data, $p_{\text{recoil}}\,(\mu\mu) = 1.25$ GeV/c. (b) MC, $p_{\text{recoil}}\,(\mu\mu) = 1.25$ GeV/c.

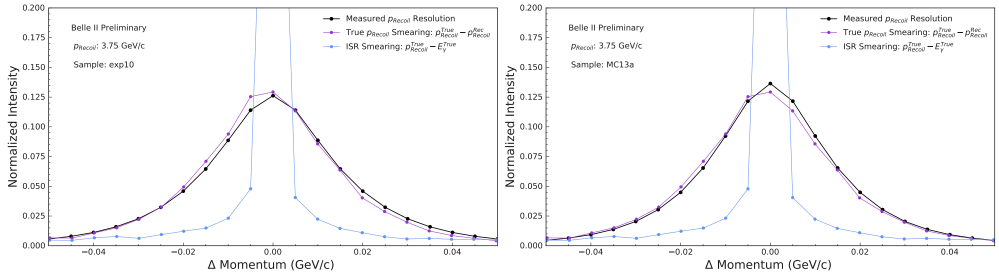

(c) Data, $p_{\text{recoil}}\,(\mu\mu) = 3.75$ GeV/c. (d) MC, $p_{\text{recoil}}\,(\mu\mu) = 3.75$ GeV/c.

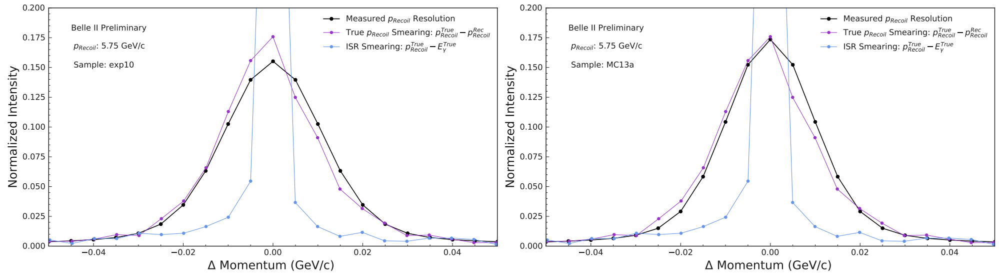

(e) Data, $p_{\text{recoil}}\,(\mu\mu) = 5.75$ GeV/c. (f) MC, $p_{\text{recoil}}\,(\mu\mu) = 5.75$ GeV/c.

Figure 8.7: Comparison of ISR smearing (azure) and true and measured tracking smearing (purple and black) for different bins of $p_{\text{recoil}}\,(\mu\mu)$ (different rows), for data (left) and MC (right).

8.3.2.2.1 Scaling from pure dimuon

To correct the scaling factors S_i, *pure dimuon* $e^+e^- \to \mu^+\mu^-$ events are used. The selection is the following:

- $0.05 \leq$ `clusterE`$(\mu) \leq 0.25$ GeV;

- $p(\mu) > 2$ GeV/c;
- $9.8 \leq m_{\mu\mu} \leq 10.8$ GeV/c^2.

Pure (i.e. non-radiative) dimuon events are used because it is known that $p_{\text{recoil}}(\mu\mu) = 0$.

If the uncertainties are properly calculated, the pulls of $p_{\text{recoil}}(\mu\mu)$ should be normally distributed. The pulls of the x, y, and z components are studied. Given the assumption that $p^{\text{true}}_{\text{recoil}}(\mu\mu)_i = 0$, $i = x, y, z$, the pulls can be rewritten as:

$$P(p_{\text{recoil}}(\mu\mu)_i) = \frac{p_{\text{recoil}}(\mu\mu)_i}{p_{\text{recoil}}(\mu\mu)_i\text{err}}. \tag{8.11}$$

Each of the three pull distributions is fitted with a Gaussian. Correcting for a possible bias in the mean values is out of the scope of this study, hence if $\mu_i \neq 0$ this is ignored for the time being. Instead, if $\sigma_i \neq 1$, the $p_{\text{recoil}}(\mu\mu)_i\text{err}$ are multiplied by a factor S_i such that all σ_i are made equal to 1.

The Gaussian fit is performed in the range $[-3, 3]$. Modifying the values of $p_{\text{recoil}}(\mu\mu)_i\text{err}$ modifies the distribution of the pulls: if the correction is greater than one, the distribution becomes narrower, and events which were previously outside of the fit range may now enter. Therefore, this process is reiterated until the σ_i parameters of the Gaussian fits converge to one. Technically, this has been reiterated until $|\sigma_i - 1| < 10^{-3}$.

The pulls before and after applying the correcting scalings S_i, with the corresponding fits, are shown in Fig. 8.8.

These would be the scaling factors, if no photons at all were present in the events considered. Although, some radiative photons will always contaminate the sample. This effect is particularly important for the z component, where the ISR radiation is more relevant. It is thus necessary to eliminate the ISR contamination.

On MC sample, a cut on the $p^{\text{true}}_{\text{recoil}}(\mu\mu)$ value is applied, requiring it to be smaller than a given threshold, since in pure dimuon events it should be zero. The scaling factors S_i are then recomputed. The S_i are plotted as a function of the $p^{\text{true}}_{\text{recoil}}(\mu\mu)$ cut in Fig. 8.9. For low values of the cut, the S_i plateau at the values that they would have if no ISR contamination was present. Tightening the $p^{\text{true}}_{\text{recoil}}(\mu\mu)$ cut too much reduces the statistics excessively, thus making the fits meaningless. At the end of this procedure, the S_i corresponding to the five smallest $p^{\text{true}}_{\text{recoil}}(\mu\mu)$ cuts applied, corresponding

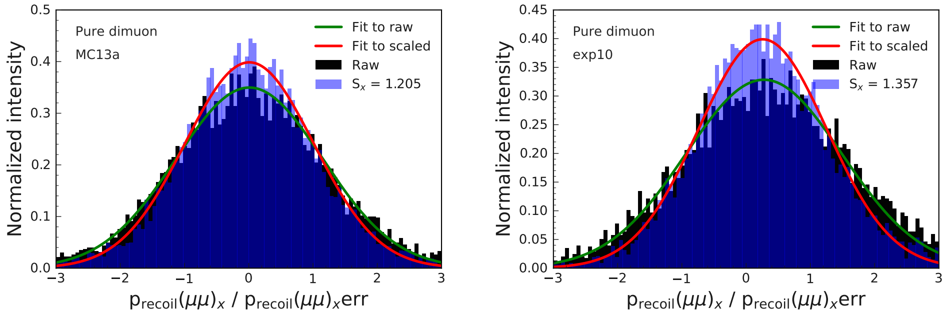

(a) Data, $p_{\text{recoil}}\left(\mu\mu\right) = 1.25$ GeV/c.

(b) MC, $p_{\text{recoil}}\left(\mu\mu\right) = 1.25$ GeV/c.

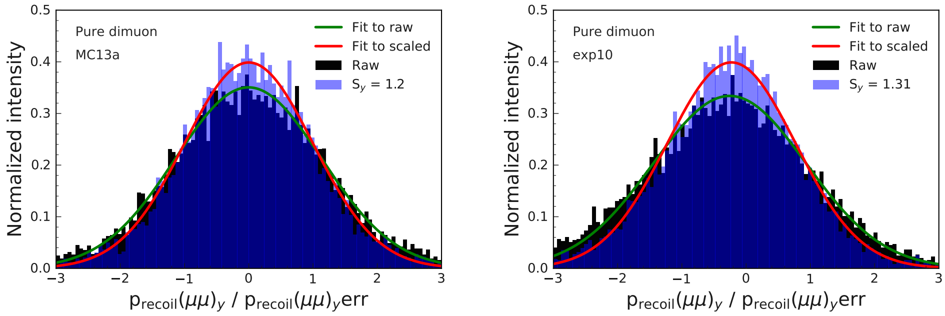

(c) Data, $p_{\text{recoil}}\left(\mu\mu\right) = 3.75$ GeV/c.

(d) MC, $p_{\text{recoil}}\left(\mu\mu\right) = 3.75$ GeV/c.

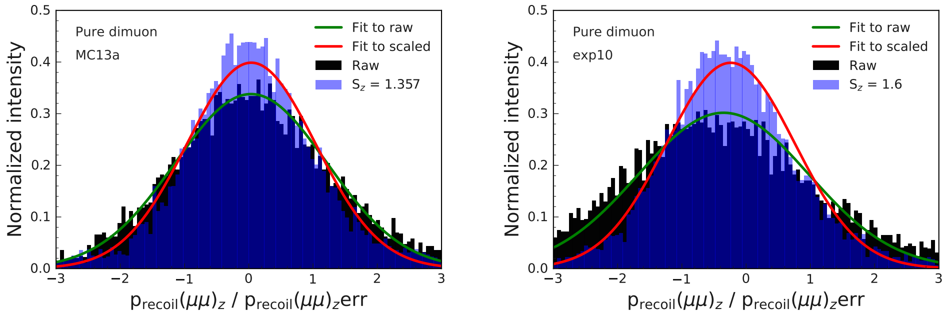

(e) Data, $p_{\text{recoil}}\left(\mu\mu\right) = 5.75$ GeV/c.

(f) MC, $p_{\text{recoil}}\left(\mu\mu\right) = 5.75$ GeV/c.

Figure 8.8: $p_{\text{recoil}}\left(\mu\mu\right)$ pulls. Each plot shows the histrogrammed distribution, and the corresponding Gaussian fit, before ("raw") and after applying the (uncorrected, see the text) scaling factor S_i. Different rows correspond to different x, y, z components. MC on left, data on right.

to $p^{\text{true}}_{\text{recoil}}\left(\mu\mu\right) < 10^{-5}$ GeV/c, are used to compute a mean value:

$$S_i^{\text{corr}} = \text{mean}\left[S_i\left(p^{\text{true}}_{\text{recoil}}\left(\mu\mu\right) < 10^{-5}\ \text{GeV/c}\right)\right]. \tag{8.12}$$

The values S_i^{corr} are the corrected scaling factors used to compute the measured tracking smearing (right side of Eq. 8.10) in MC. For data, this procedure is not possible, so the previously obtained S_i^{data} are corrected with the same ratio by which the MC values are corrected:

$$S_i^{\text{data, corr}} = S_i^{\text{data, uncorr}} \frac{S_i^{\text{MC, corr}}}{S_i^{\text{MC, uncorr}}}. \tag{8.13}$$

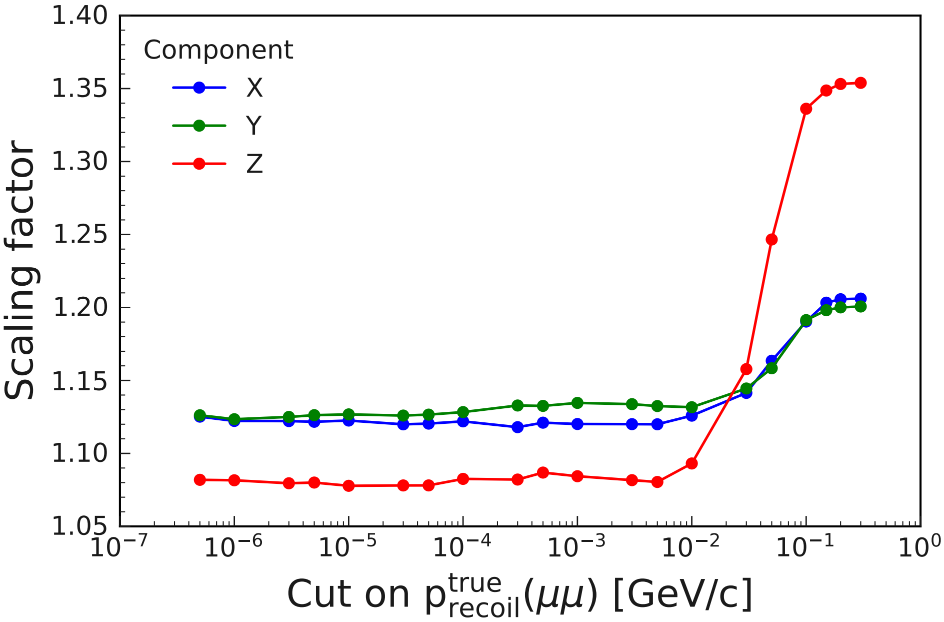

Figure 8.9: Trend of the scaling factors S_i, computed with reconstructed MC events for which $p_{\text{recoil}}^{\text{true}}(\mu\mu)$ is smaller than the cut shown on the x axis, indicating the absence of ISR photons. All of the S_i plateau for low values of the cut.

These are the scaling factors used to compute the $p_{\text{recoil}}^{\text{rec}}(\mu\mu)\,\text{err}^{\text{scaled}}$ values that appear in Eq. 8.10:

$$p_{\text{recoil}}^{\text{rec}}(\mu\mu)\,\text{err}^{\text{scaled}} = p_{\text{recoil}}^{\text{rec}}(\mu\mu)\,\text{err}\sqrt{\sum_i^{x,y,z}\left(\frac{p_{\text{recoil}}^{\text{rec}}(\mu\mu)_i}{p_{\text{recoil}}^{\text{rec}}(\mu\mu)}S_i^{\text{corr}}\right)^2}. \tag{8.14}$$

8.3.3 Results

Before showing the results of the convolution procedure, the main steps are briefly summarized.

The overall fit function, used for the final fit to $E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu)$, is given by the convolution of a floating CB, which models the true ECL resolution, with the measured total smearing function. The measured total smearing is given by the convolution of the ISR smearing (from MC, see Sec. 8.3.2.1) with the measured tracking smearing (from data, see Sec. 8.3.2.2). The measured tracking smearing is given by the right side of Eq. 8.10. This function uses the scaled uncertainty on the recoil momentum, $p^{\text{rec}}_{\text{recoil}}(\mu\mu)\,\text{err}^{\text{scaled}}$, computed for each event. These are obtained from the raw values by applying a scaling factor, as described in Eq. 8.14. The scaling factors S_i, defined for data and MC in Eqs. 8.13 and 8.12, are obtained via the pure dimuon studies described in Sec. 8.3.2.2.1.

The measured total smearing function can be compared with the true total smearing function, of which it is an approximation. Some comparisons can be seen in Fig. 8.10. The agreement is very good, except for some points at the extreme values of recoil momentum, partially due to lower statistics.

In Fig. 8.11 the goodness of the whole convolution procedure is evaluated. The raw results are shown, i.e. with no convolution procedure applied, obtained by just fitting the $E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu)$ distributions with a CB (green points). The target results, which use MC information, are included, obtained by fitting the $E^{\text{rec}}(\gamma) - E^{\text{true}}(\gamma)$ with a CB (red points). The unfolded results are obtained from the convolution procedure, by plotting the width of the CB used to model the true ECL resolution. Results are shown both when the true total smearing is used, i.e. using directly templates from $p^{\text{rec}}_{\text{recoil}}(\mu\mu) - E^{\text{true}}(\gamma)$ (azure points), and when the convolution of true tracking smearing and ISR smearing is used (purple points).

The excellent agreement between the latter two (azure and purple) confirms that the overall smearing is indeed due to only the ISR and the tracking components.

The agreement between these two and the target results (red) confirms that the convolution procedure is a good method to extract the true ECL resolution, if the tue total smearing function is properly approximated by the measured total smearing function.

The distance between the raw (green) and the target values (red) shows the size of the smearing contribution which needs to be factored out, almost a factor two for low values of photon energy. As just discussed, this huge smearing contribution is indeed efficiently factored out if the smearing function is properly measured.

Fig. 8.12 shows the unfolded results of the convolution procedure, applied both to data and MC. The red, azure, and green points have the same meaning as in Fig. 8.11. The unfolded results obtained with the measured total smearing are shown for both data (black points) and MC (blue points); examples of the fits are shown in Fig. 8.13. The raw results, obtained by fitting the $E^{\text{rec}}(\gamma) - p^{\text{rec}}_{\text{recoil}}(\mu\mu)$ distributions with a CB, are shown for data too (grey points).

The great agreement between the results obtained with the measured (blue) and true (azure) total smearing confirms the goodness of the approximation performed in Eq. 8.10, i.e. the proxy function used to model the total smearing.

The distance between raw and unfolded results (green and blue for MC, grey and black for data) shows the magnitude of the smearing which has been factored out.

A non-negligible data/MC discrepancy is present for the unfolded results (black vs red). Given that a similar discrepancy is present also for the raw results (grey vs green), it cannot be ascribed to the convolution procedure, but rather to some underlying imperfection in the MC modeling of the reconstruction.

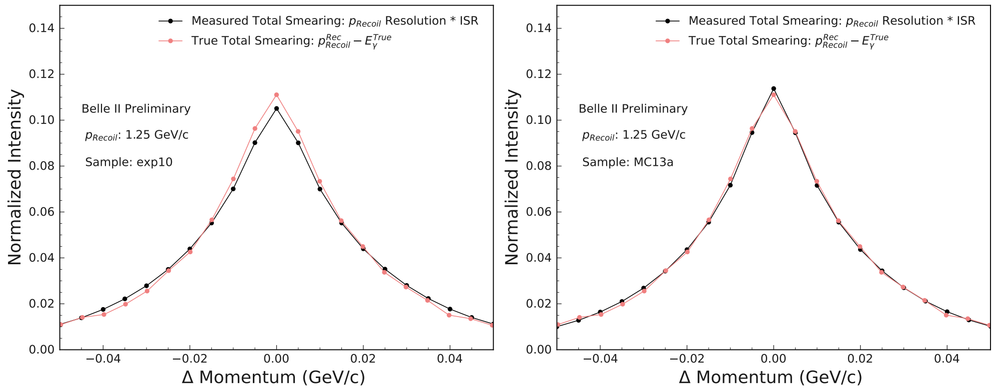

(a) Data, $p_{\text{recoil}}(\mu\mu) = 1.25$ GeV/c. (b) MC, $p_{\text{recoil}}(\mu\mu) = 1.25$ GeV/c.

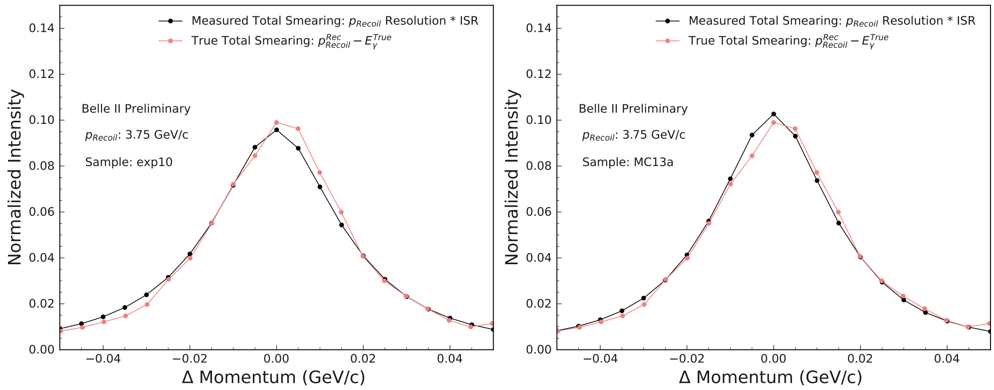

(c) Data, $p_{\text{recoil}}(\mu\mu) = 3.75$ GeV/c. (d) MC, $p_{\text{recoil}}(\mu\mu) = 3.75$ GeV/c.

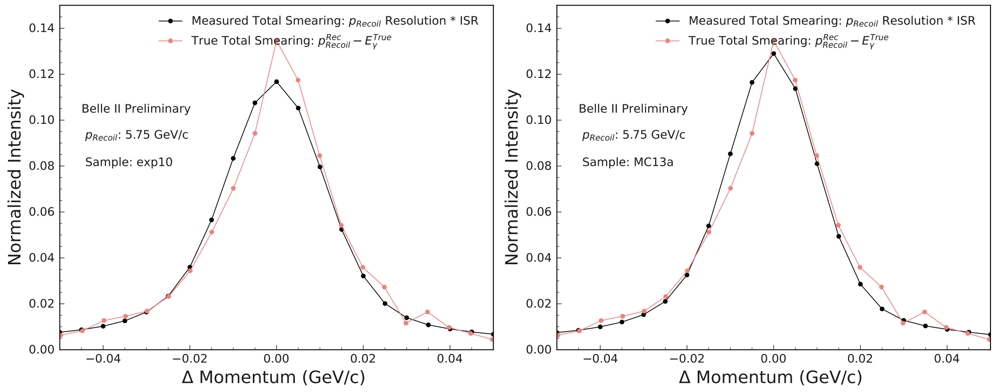

(e) Data, $p_{\text{recoil}}(\mu\mu) = 5.75$ GeV/c. (f) MC, $p_{\text{recoil}}(\mu\mu) = 5.75$ GeV/c.

Figure 8.10: Comparison of true (pink) and measured (black) total smearing for different bins of $p_{\text{recoil}}(\mu\mu)$ (different rows), for data (left) and MC (right).

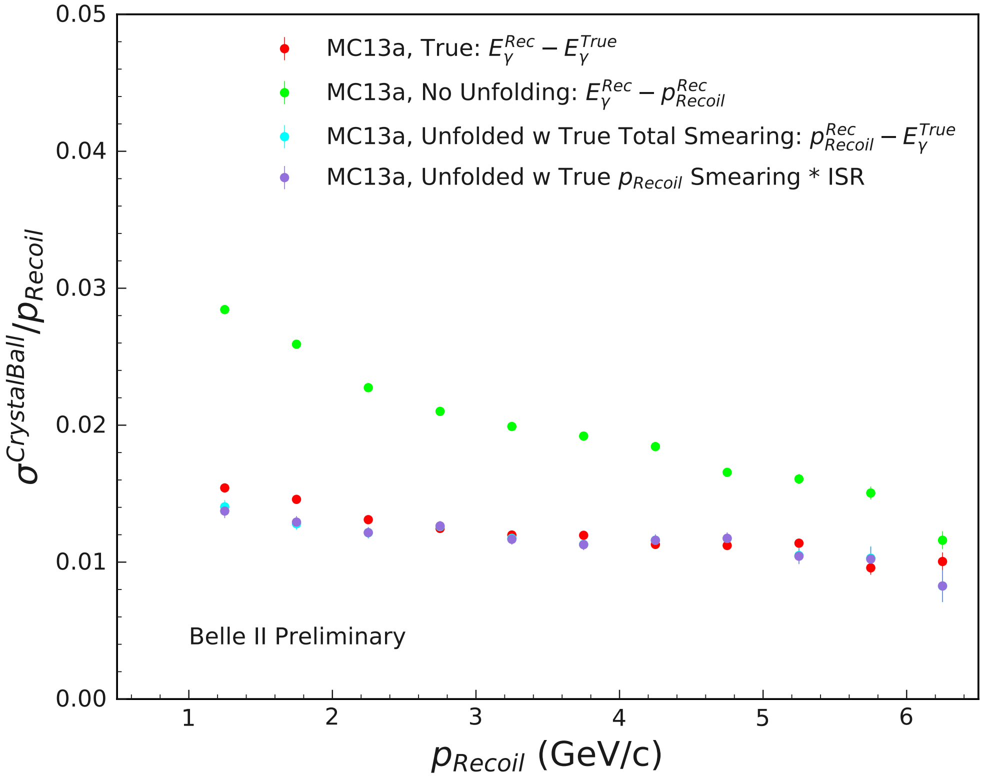

Figure 8.11: Photon energy resolution as a function of $p_{\mathrm{recoil}}\,(\mu\mu) \approx E(\gamma)$. Parametrized as the σ parameter of the CB used for the modeling of the ECL resolution, normalized with respect to $p_{\mathrm{recoil}}\,(\mu\mu)$. Red: true MC resolution; green: MC without convolution procedure applied; azure: MC with convolution procedure applied using true total resolution instead of the measured one; purple: MC with convolution procedure applied using true tracking resolution instead of the measured one.

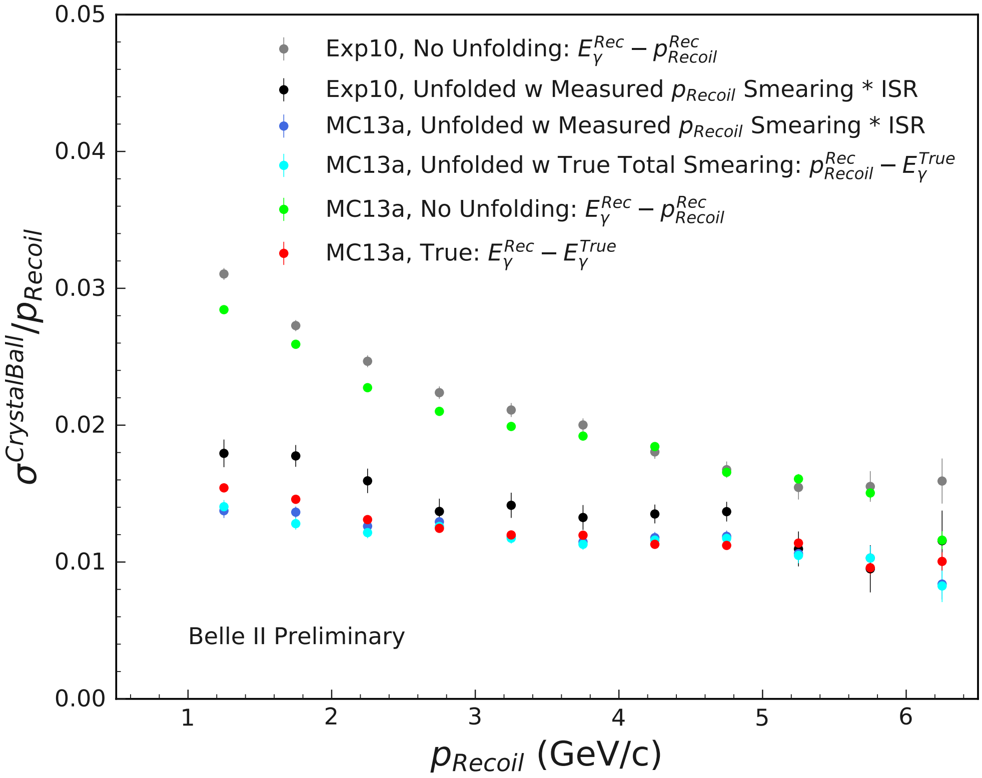

Figure 8.12: Photon energy resolution as a function of $p_{\text{recoil}}(\mu\mu) \approx E(\gamma)$. Parametrized as the σ parameter of the CB used for the modeling of the ECL resolution, normalized with respect to $p_{\text{recoil}}(\mu\mu)$. Red: true MC resolution; green [grey]: MC [data] without convolution procedure applied; blue [black]: MC [data] with convolution procedure applied; azure: MC with convolution procedure applied using true total resolution instead of the measured one.

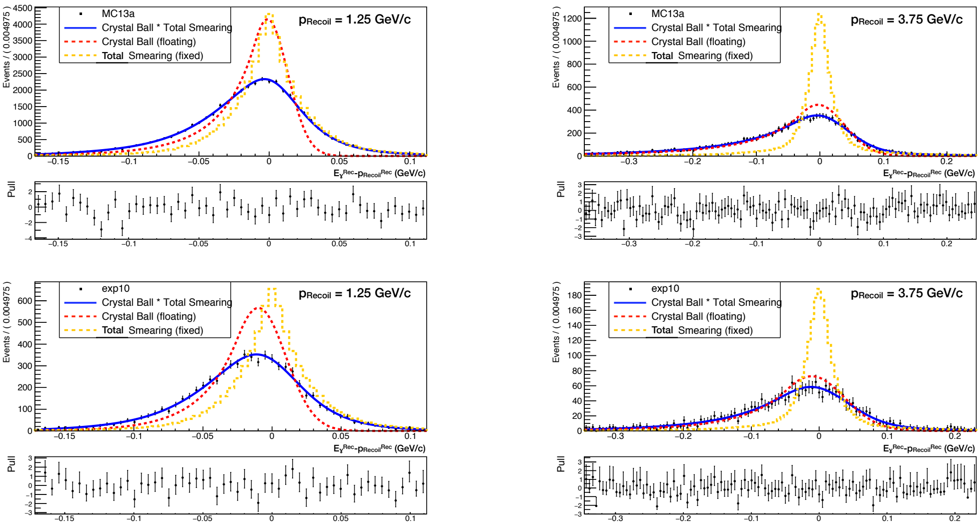

Figure 8.13: Examples of the convolution fits. Top row: MC (blue points in Fig. 8.12), bottom row: data (black points in Fig. 8.12); left: $p_{\text{recoil}}\,(\mu\mu) = 1.25$ GeV/c, right: $p_{\text{recoil}}\,(\mu\mu) = 3.75$ GeV/c. The red dashed lines is the floating CB, the yellow dashed line is the fixed measured total smearing, the blue solid line is the convolution of the two, used to fit the black point.

8.4 Prospects

To finalize the photon covariance matrix and photon resolution studies presented in this Chapter, the following step may be taken:

- Also the θ and ϕ angular resolutions are to be extracted, following the same procedure described above. Summary plots similar to Fig. 8.12 are to be obtained;

- It may be necessary to split the photon energy resolution studies into bins of θ, as hinted by Fig. 8.4. In general, the resolution of each of the three photon variables E, θ, and ϕ may be dependent on all of them;

- All these resolution values are to be inserted into the PCM. This would produce a PCM for each value of E, θ, and ϕ. The energy resolution may be interpolated using the calorimeter resolution function, $\frac{\sigma(E)}{E} = \frac{A}{\sqrt{E}} \oplus B \oplus \frac{C}{E}$ [72], where A, B, and C are parameters to be determined. As for the angular variables, a bin approach may be more suited;

- These photon resolutions would act as starting values for the diagonal elements of the PCM;

- The covariances between these would be used as starting values of the out-of-diagonal elements;

- Deploying an iterative approach, possibly using a multi-dimensional grid, the corrections that make the ALP mass pulls as normal as possible (using the normality monitor plot in Fig. 8.3) are to be found. They should be either set in a lookup table, or defined via parameterized empirical functions;

- These PCMs should be computed from data and updated any time there is a relevant change in the beam settings, particularly when the beam background increases as a consequence of the higher luminosity;

- One or more samples to cross-check the goodness of the updated PCM in data, rather than in MC via the ALP samples, should be studied. A possibility are the $e^+e^- \to P\gamma(\gamma)$, $P = (\pi^0, \eta, \eta\prime)$ processes that constituted a peaking background for the ALP search, described in Sec. 4.3.2. Some non-fully neutral decays may be used too, for example decays including π^0 and/or $K^{0[*]}$.

Chapter 9

Summary and outlook

This thesis presented a search for Axion-Like Particles, known as ALPs, in the mass range $0.2 < m_a < 9.7$ GeV/c^2 (Secs. 4 to 7) and a study of the photon covariance matrix and photon energy resolution (Sec. 8) at Belle II (Sec. 3).

Axions and ALPs are considered amongst the most promising candidates for Beyond Standard Model (BSM) Physics (Sec. 2). ALPs may be Dark Matter (DM) particles or mediators between Standard Model (SM) and DM.

The model taken into consideration in this analysis, $e^+e^- \to \gamma a (a \to \gamma\gamma)$, is of the latter kind. Its free parameters are the ALP mass m_a and the coupling constant to photons $g_{a\gamma\gamma}$. This photophilic ALP could be produced in e^+e^- collisions in association with a recoil photon, and would decay into a photon pair. The search presented in this thesis has been performed with 0.445 fb^{-1} of early calibration data, collected at Belle II during the first data-taking period, in 2018. The signature of this process is no missing energy, three high-energetic resolved photons, and absence of any charged tracks. The search has been carried out in the mass range $0.2 < m_a < 9.7$ GeV/c^2.

This was a blind analysis, meaning that the experimenters didn't access the variables' distributions of the whole data set before the analysis strategy was finalized. As m_a is a free parameter of the model, a mass scan was performed.

The selection has been optimized via maximization of the Punzi figure of merit (PFM). This has been done in an iterative way for all variables.

In order to perform a mass scan, both signal and background need to be modeled. The signal is modeled using a Crystal Ball (CB) for the peaking component, and a Kernel Density Estimator (KDE) for the combinatorial one. The background is modeled with polynomials of degrees two to four.

For each mass hypothesis, a signal+background and a background-only fit are performed. From these, the local significance is computed. No mass point, after including systematics, returned a local significance higher than three. Thus, the global significance has not been calculated, no evidence has been claimed, and 95%-confidence-level upper limits on the coupling constant $g_{a\gamma\gamma}$ have been set using the CLs method.

The obtained limits, at the level of 10^{-3} GeV^{-1}, are, to date, the most stringent in the mass range $0.2 < m_a < 1.0$ GeV/c^2.

The mass scan was performed with two different definitions of ALP mass: diphoton up to $m_a = 6.85$ GeV/^{2}, and recoil onward. This was done because a kinematic fit was not available, due to the absence of a properly calibrated photon covariance matrix (PCM).

Photon resolution values are needed as starting points of the PCM, which is comprised of the variances and covariances of $E(\gamma)$, $\theta(\gamma)$, and $\phi(\gamma)$ residuals. The variances are more important than the covariances for the error propagation, and the most relevant one is the energy $E(\gamma)$; thus, the photon resolution studies started with this variable.

Radiative dimuon events $e^+e^- \to \mu^+\mu^-\gamma$ are used to access photon resolution in data. The recoil momentum of the $\mu\mu$ pair is compared with $E(\gamma)$. To access the true calorimeter resolution, it is necessary to factor out the tracking and the ISR smearings.

The results of these studies are the starting point of the PCM. The ultimate goal, yet to be achieved, is to build a set of PCMs calibrated such that the ALP mass pulls are distributed normally, i.e. according to a Gaussian with mean zero and standard deviation one.

Starting from the results obtained in this thesis, further evolutions can be implemented in future iterations of the ALP search. These will be possible thanks to better knowledge of the detector and with the upgraded data sample of 50 ab^{-1} that Belle II will collect in the upcoming years. This is about 10^5 times the data used for the analysis presented, and can lead to upper limits on $g_{a\gamma\gamma}$ about 20 times stronger than the current ones.

The ALP analysis can be extended to lower masses. This will present a challenge: light ALPs are highly boosted, so the two decay photons are merged together and reconstructed as a single cluster by the calorimeter. Such events thus look like $e^+e^- \to \gamma\gamma$, which will be largely rejected in the near future due to high background, deriving from the frontier-high luminosity of Belle II, already the world highest and still increasing. Dedicated trigger lines and improved clustering, capable of better separating close clusters, will therefore be necessary to improve the sensitivity for this search. Additionally, the higher background will require the selection to be updated.

In particular, a neural-network-based approach can be investigated. Another way to search for low-mass ALPs is to require that one of decay photons undergoes pair conversion, $\gamma \to e^+e^-$, before reaching the calorimeter. The resulting e^+e^- pair is bent by the magnetic field, thus resulting in four resolved clusters, two neutrals and two associated with oppositely charged tracks.

With higher integrated luminosity, there will be the possibility to improve the background modeling: with better knowledge of the detector, the MC simulation of the background processes will be further improved, possibly to the point where the background PDF could be fixed before unblinding rather than left floating.

The kinematic fit can be used to avoid the double ALP mass definition once the PCM is properly built. This will be achieved upon completion of the photon resolution studies, and subsequent computation of the set of PCMs such that the ALP mass pulls are made normal.

ALPs at e^+e^- colliders can also be produced via photon fusion: $e^+e^- \to e^+e^-a(a \to \gamma\gamma)$. In such processes, ALPs carry little energy, so do the two decays photons, and the e^+e^- pair is often lost due to acceptance. A two-fold dedicated search can be designed, both requiring the presence of the e^+e^- pair within acceptance, hence sacrificing the cross-section, and by studying the invariant mass distributions of soft $\gamma\gamma$ pairs, with missing energy, over the beam background.

Appendix A

Limits and systematics

This Appendix contains plots regarding the Upper Limits (ULs) and the magnitude of the systematic uncertainties. See Chap. 7 for more details.

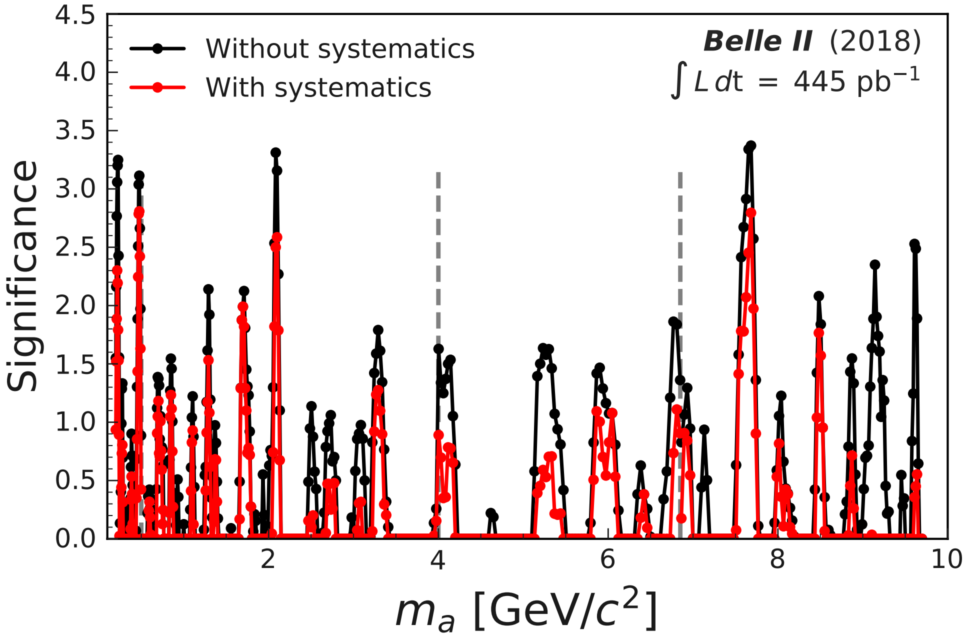

Figure A.1: Significance, with and without applying systematics, as a function of the fitted ALP mass.

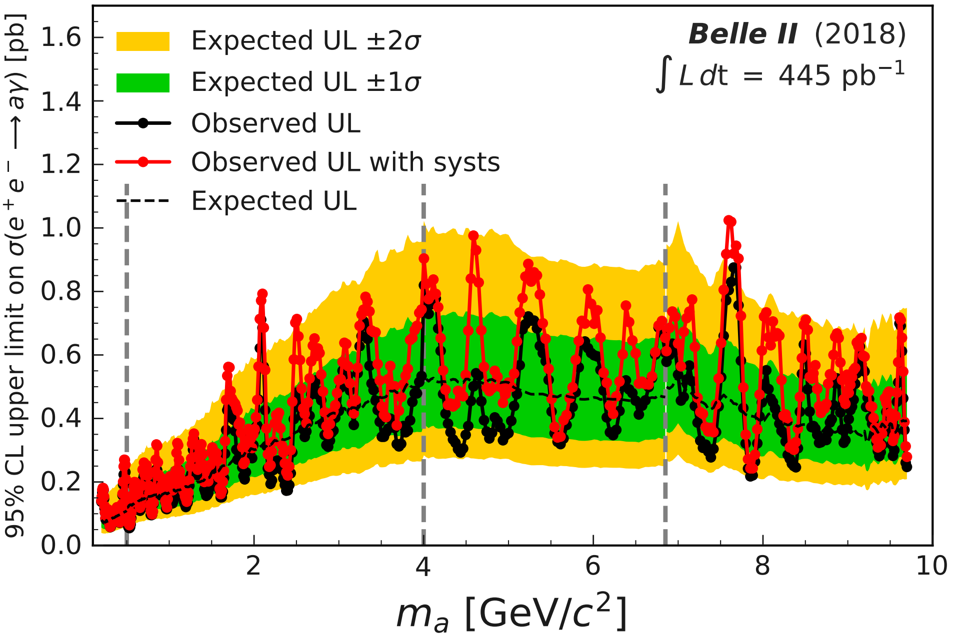

Figure A.2: 95%-confidence-level UL on the cross section for $e^+e^- \to \gamma a$, with and without systematics included. The red line corresponds to the black solid line of Fig. 7.34, i.e. the UL with systematics included. The black lines and the Brazilian bands correspond to the limits obtained without applying the systematics.

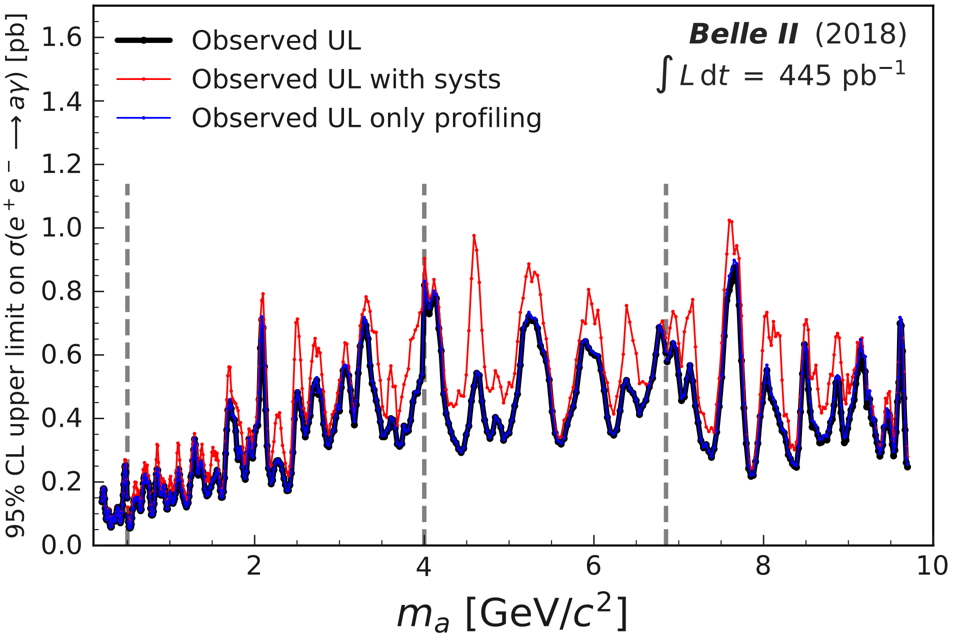

Figure A.3: 95%-confidence-level UL on the cross section for $e^+e^- \to \gamma a$. Only observed limits. Without systematics (black line), with all systematics (red line), only with systematics propagated via profiling (blue line). It can be seen that the effect of the latter is almost negligible with respect to the systematics deriving from the choice of background shape and fit range.

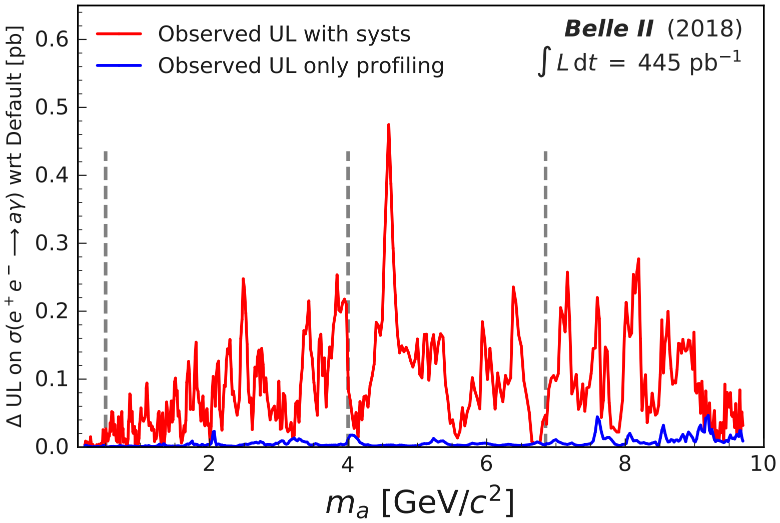

(a) Linear y scale.

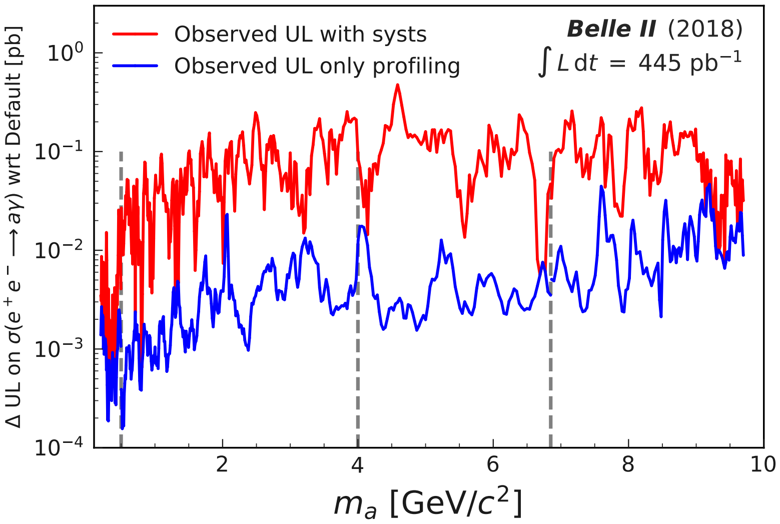

(b) Logarithmic y scale.

Figure A.4: Differences of the 95%-confidence-level UL on the cross section for $e^+e^- \to \gamma a$, obtained applying systematics, with respect to the UL obtained without applying systematics. Blue line: only applying systematics propagated via profiling; red line: applying all systematics. Linear x scale on the top, logarithmic on the bottom.

Acronyms

SM Standard Model

BSM Beyond Standard Model

ALP axion-like particle

CPV CP violation

DM Dark Matter

FCNC flavor-changing neutral current

cms center-of-mass system

CM center-of-mass

`basf2` Belle II analysis software framework

HLT high-level (software) trigger

ISR initial-state radiation

L1 Level 1

PFM Punzi figure of merit

TM truth-matched

FWHM Full Width Half Maximum

UL Upper Limit

KDE Kernel Density Estimator

IP Interaction Point

CB Crystal Ball

LEE Look-Elsewhere Effect

CL Confidence Level

NLL Negative Logarithmic Likelihood

BR Branching Ratio

FPGA Field Programmable Gate Array

LM local maximum

CR connected region

PCM photon covariance matrix

PDF probability distribution function

Bibliography

[1] Belle II collaboration, F. Abudinén et al., *Search for Axion-Like Particles produced in e^+e^- collisions at Belle II*, *Phys. Rev. Lett.* **125** (2020) 161806 [2007.13071]. (cited on pages 7, 24, 29, and 129)

[2] V. Trimble, *Existence and Nature of Dark Matter in the Universe*, *Ann. Rev. Astron. Astrophys.* **25** (1987) 425. (cited on page 13)

[3] G. Bertone and D. Hooper, *History of dark matter*, *Rev. Mod. Phys.* **90** (2018) 045002 [1605.04909]. (cited on page 13)

[4] D. Clowe, M. Bradac, A. H. Gonzalez, M. Markevitch, S. W. Randall, C. Jones, et al., *A direct empirical proof of the existence of dark matter*, *Astrophys. J. Lett.* **648** (2006) L109 [astro-ph/0608407]. (cited on page 13)

[5] J. Alexander et al., *Dark Sectors 2016 Workshop: Community Report*, 8, 2016, 1608.08632. (cited on page 13)

[6] M. J. Dolan, F. Kahlhoefer, C. McCabe, and K. Schmidt-Hoberg, *A taste of dark matter: Flavour constraints on pseudoscalar mediators*, *JHEP* **03** (2015) 171 [1412.5174], [Erratum: JHEP 07, 103 (2015)]. (cited on pages 13 and 19)

[7] Wikimedia Commons, *File:Standard Model of Elementary Particles.svg — Wikimedia Commons, the free media repository*, https://commons.wikimedia.org/w/index.php?title=File:Standard_Model_of_Elementary_Particles.svg&oldid=570200339, last accessed 2021/07/09. (cited on page 16)

[8] R. D. Peccei, *The Strong CP problem and axions*, *Lect. Notes Phys.* **741** (2008) 3 [hep-ph/0607268]. (cited on page 17)

[9] R. D. Peccei and H. R. Quinn, *CP Conservation in the Presence of Pseudoparticles*, *Phys. Rev. Lett.* **38** (1977) 1440. (cited on page 17)

[10] R. D. Peccei and H. R. Quinn, *Constraints Imposed by CP Conservation in the Presence of Pseudoparticles*, *Phys. Rev. D* **16** (1977) 1791. (cited on page 17)

[11] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Evidence for the* 2π *Decay of the* K_2^0 *Meson*, *Phys. Rev. Lett.* **13** (1964) 138. (cited on pages 17 and 27)

[12] BaBar, Belle collaboration, A. J. Bevan et al., *The Physics of the B Factories*, *Eur. Phys. J. C* **74** (2014) 3026 [`1406.6311`]. (cited on pages 17, 28, 29, 41, and 92)

[13] C.-T. Chan, *On the symmetry constraints of CP violations in QCD*, `hep-ph/9704427`. (cited on page 17)

[14] Particle Data Group collaboration, C. Patrignani et al., *Review of Particle Physics*, *Chin. Phys. C* **40** (2016) 100001. (cited on page 17)

[15] B. Yoon, T. Bhattacharya, and R. Gupta, *Neutron Electric Dipole Moment on the Lattice*, *EPJ Web Conf.* **175** (2018) 01014 [`1712.08557`]. (cited on page 17)

[16] J. Isern, S. Catalan, E. Garcia-Berro, M. Salaris, and S. Torres, *Axion cooling of white dwarfs*, 4, 2013, `1304.7652`. (cited on page 18)

[17] A. De Angelis, M. Roncadelli, and O. Mansutti, *Evidence for a new light spin-zero boson from cosmological gamma-ray propagation?*, *Phys. Rev. D* **76** (2007) 121301 [`0707.4312`]. (cited on page 18)

[18] D. Cadamuro and J. Redondo, *Cosmological bounds on pseudo Nambu-Goldstone bosons*, *JCAP* **02** (2012) 032 [`1110.2895`]. (cited on page 18)

[19] M. Millea, L. Knox, and B. Fields, *New Bounds for Axions and Axion-Like Particles with keV-GeV Masses*, *Phys. Rev. D* **92** (2015) 023010 [`1501.04097`]. (cited on page 18)

[20] K. Mimasu and V. Sanz, *ALPs at Colliders*, *JHEP* **06** (2015) 173 [`1409.4792`]. (cited on pages 18, 23, and 129)

[21] E. Izaguirre, T. Lin, and B. Shuve, *Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes*, *Phys. Rev. Lett.* **118** (2017) 111802 [`1611.09355`]. (cited on page 18)

[22] M. Bauer, M. Neubert, and A. Thamm, *LHC as an Axion Factory: Probing an Axion Explanation for $(g-2)_\mu$ with Exotic Higgs Decays*, *Phys. Rev. Lett.* **119** (2017) 031802 [`1704.08207`]. (cited on page 18)

[23] W. J. Marciano, A. Masiero, P. Paradisi, and M. Passera, *Contributions of axionlike particles to lepton dipole moments*, *Phys. Rev. D* **94** (2016) 115033 [`1607.01022`]. (cited on pages 18 and 20)

[24] U. Ellwanger and S. Moretti, *Possible Explanation of the Electron Positron Anomaly at 17 MeV in ^{8}Be Transitions Through a Light Pseudoscalar*, *JHEP* **11** (2016) 039 [`1609.01669`]. (cited on page 18)

[25] L. Calibbi, F. Goertz, D. Redigolo, R. Ziegler, and J. Zupan, *Minimal axion model from flavor*, *Phys. Rev. D* **95** (2017) 095009 [`1612.08040`]. (cited on page 18)

[26] M. J. Dolan, T. Ferber, C. Hearty, F. Kahlhoefer, and K. Schmidt-Hoberg, *Revised constraints and Belle II sensitivity for visible and invisible axion-like particles*, *JHEP* **12** (2017) 094 [`1709.00009`], [Erratum: JHEP 03, 190 (2021)]. (cited on pages 18, 20, 22, 23, and 129)

[27] B. Döbrich, *Axion-like Particles from Primakov production in beam-dumps*, *CERN Proc.* **1** (2018) 253 [`1708.05776`]. (cited on pages 21, 23, and 129)

[28] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald, and K. Schmidt-Hoberg, *ALPtraum: ALP production in proton beam dump experiments*, *JHEP* **02** (2016) 018 [`1512.03069`]. (cited on pages 21, 23, and 129)

[29] DELPHI collaboration, J. Abdallah et al., *Search for one large extra dimension with the DELPHI detector at LEP*, *Eur. Phys. J. C* **60** (2009) 17 [`0901.4486`]. (cited on pages 23 and 129)

[30] P. J. Fox, R. Harnik, J. Kopp, and Y. Tsai, *LEP Shines Light on Dark Matter*, *Phys. Rev. D* **84** (2011) 014028 [`1103.0240`]. (cited on pages 23 and 129)

[31] BaBar collaboration, B. Aubert et al., *Search for Invisible Decays of a Light Scalar in Radiative Transitions $\Upsilon(3S) \to \gamma A^0$*, in *Proceedings of the 34th International Conference on High Energy Physics, 2008*, 2008, `0808.0017`. (cited on pages 23 and 129)

[32] BaBar collaboration, J. P. Lees et al., *Search for Invisible Decays of a Dark Photon Produced in e^+e^- Collisions at BaBar*, Phys. Rev. Lett. **119** (2017) 131804 [1702.03327]. (cited on pages 23 and 129)

[33] S. Knapen, T. Lin, H. K. Lou, and T. Melia, *Searching for Axionlike Particles with Ultraperipheral Heavy-Ion Collisions*, Phys. Rev. Lett. **118** (2017) 171801 [1607.06083]. (cited on pages 23, 45, and 129)

[34] J. D. Bjorken, S. Ecklund, W. R. Nelson, A. Abashian, C. Church, B. Lu, et al., *Search for Neutral Metastable Penetrating Particles Produced in the SLAC Beam Dump*, Phys. Rev. D **38** (1988) 3375. (cited on pages 23 and 129)

[35] CHARM collaboration, F. Bergsma et al., *Search for Axion Like Particle Production in 400-GeV Proton - Copper Interactions*, Phys. Lett. B **157** (1985) 458. (cited on pages 23 and 129)

[36] J. Blumlein et al., *Limits on neutral light scalar and pseudoscalar particles in a proton beam dump experiment*, Z. Phys. C **51** (1991) 341. (cited on pages 23 and 129)

[37] NA64 collaboration, D. Banerjee et al., *Search for Axionlike and Scalar Particles with the NA64 Experiment*, Phys. Rev. Lett. **125** (2020) 081801 [2005.02710]. (cited on pages 23 and 129)

[38] CMS collaboration, A. M. Sirunyan et al., *Evidence for light-by-light scattering and searches for axion-like particles in ultraperipheral PbPb collisions at $\sqrt{s_{\rm NN}} = 5.02$ TeV*, Phys. Lett. B **797** (2019) 134826 [1810.04602]. (cited on pages 23 and 129)

[39] D. Aloni, C. Fanelli, Y. Soreq, and M. Williams, *Photoproduction of Axionlike Particles*, Phys. Rev. Lett. **123** (2019) 071801 [1903.03586]. (cited on pages 24 and 129)

[40] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, Prog. Theor. Phys. **49** (1973) 652. (cited on page 27)

[41] V. P. Druzhinin, *Recent results on two-photon physics at BABAR*, PoS **EPS-HEP2009** (2009) 051 [0909.3148]. (cited on page 29)

[42] Belle-II collaboration, I. Adachi et al., *Search for an Invisibly Decaying Z' Boson at Belle II in $e^+e^- \to \mu^+\mu^-(e^\pm\mu^\mp)$ Plus Missing Energy Final States*, Phys. Rev. Lett. **124** (2020) 141801 [1912.11276]. (cited on page 29)

[43] SuperKEKB collaboration, K. Akai, K. Furukawa, and H. Koiso, *SuperKEKB Collider*, *Nucl. Instrum. Meth. A* **907** (2018) 188 [1809.01958]. (cited on pages 29 and 30)

[44] Belle-II collaboration, T. Abe et al., *Belle II Technical Design Report*, 1011.0352. (cited on pages 29, 31, and 32)

[45] Belle-II collaboration, W. Altmannshofer et al., *The Belle II Physics Book*, *PTEP* **2019** (2019) 123C01 [1808.10567], [Erratum: PTEP 2020, 029201 (2020)]. (cited on pages 31, 32, and 39)

[46] Belle II Tracking Group collaboration, V. Bertacchi et al., *Track finding at Belle II*, *Comput. Phys. Commun.* **259** (2021) 107610 [2003.12466]. (cited on page 33)

[47] H. Ikeda, *Development of the CsI(Tl) Calorimeter for the Measurement of CP Violation at KEK B-Factory*, Ph.D. thesis, Nara Women's University, 1999. `https://webhepl.cc.nara-wu.ac.jp/old_HP/thesis/doctor/hitomi/main/`. (cited on page 33)

[48] A. Hershenhorn, *Search for Axion Like Particles with the BABAR detector and photon hadron separation using Zernike moments at Belle II*, Ph.D. thesis, The University of British Columbia, May, 2021. `https://open.library.ubc.ca/cIRcle/collections/ubctheses/24/items/1.0398055`. (cited on page 37)

[49] Belle-II Framework Software Group collaboration, T. Kuhr, C. Pulvermacher, M. Ritter, T. Hauth, and N. Braun, *The Belle II Core Software*, *Comput. Softw. Big Sci.* **3** (2019) 1 [1809.04299]. (cited on page 38)

[50] G. Van Rossum and F. L. Drake Jr, *Python tutorial*. Centrum voor Wiskunde en Informatica Amsterdam, The Netherlands, 1995. (cited on page 39)

[51] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, *Nucl. Instrum. Meth. A* **389** (1997) 81 See also "ROOT" [software], release 6.14/06, 2018/11/05, `https://root.cern/releases/release-61406/`. (cited on pages 39 and 93)

[52] G. Punzi, *Sensitivity of searches for new signals and its optimization*, *eConf* **C030908** (2003) MODT002 [physics/0308063]. (cited on pages 42 and 49)

[53] I. J. Myung, *Tutorial on maximum likelihood estimation*, *Journal of Mathematical Psychology* **47** (2003) 90. (cited on pages 44 and 93)

[54] A. L. Read, *Presentation of search results: The CL(s) technique*, *J. Phys. G* **28** (2002) 2693. (cited on pages 44 and 93)

[55] P. Sinervo, *Definition and Treatment of Systematic Uncertainties in High Energy Physics and Astrophysics*, *eConf* **C030908** (2003) `https://www.slac.stanford.edu/econf/C030908/papers/TUAT004.pdf`. (cited on page 44)

[56] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [`1405.0301`]. (cited on page 45)

[57] Knapen, Simon and Lin, Tongyan and Lou, Hou Keong and Melia, Tom, *ALP madgraph model*, `https://github.com/belle2/basf2/blob/main/generators/madgraph/models/ALP_UFO/CREDITS.md`, last accessed 2021/07/09. (cited on page 45)

[58] Q. Li and Q.-S. Yan, *Initial State Radiation Simulation with MadGraph*, `1804.00125`. (cited on page 45)

[59] C. M. Carloni Calame, G. Montagna, O. Nicrosini, and F. Piccinini, *Status of the BabaYaga event generator*, *EPJ Web Conf.* **218** (2019) 07004. (cited on page 46)

[60] T. Ferber and P. Urquijo, *Overview of the Belle II Physics Generators*, 2015/06/22, internal note, `https://docs.belle2.org/record/282?ln=en`, last accessed 2021/07/12. (cited on page 46)

[61] H. Czyż, P. Kisza, and S. Tracz, *Modeling interactions of photons with pseudoscalar and vector mesons*, *Phys. Rev. D* **97** (2018) 016006 [`1711.00820`]. (cited on pages 47 and 74)

[62] W. Verkerke and D. P. Kirkby, *The RooFit toolkit for data modeling*, *eConf* **C0303241** (2003) MOLT007 [`physics/0306116`]. (cited on page 93)

[63] W. Verkerke and D. P. Kirkby, *RooFit Users Manual v2.91*, `https://root.cern.ch/download/doc/RooFit_Users_Manual_2.91-33.pdf`, last accessed 2021/07/12. (cited on page 93)

[64] *RooStats::AsymptoticCalculator Class Reference*, `https://root.cern/doc/master/classRooStats_1_1AsymptoticCalculator.html`, last accessed 2021/07/12. (cited on page 93)

[65] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554 [`1007.1727`], [Erratum: Eur.Phys.J.C 73, 2501 (2013)]. (cited on page 93)

[66] S. Algeri, J. Conrad, D. A. van Dyk, and B. Anderson, *On methods for correcting for the look-elsewhere effect in searches for new physics*, *JINST* **11** (2016) P12010 [`1602.03765`]. (cited on page 94)

[67] E. Gross and O. Vitells, *Trial factors for the look elsewhere effect in high energy physics*, *Eur. Phys. J. C* **70** (2010) 525 [`1005.1891`]. (cited on page 95)

[68] CMS collaboration, G. L. Bayatian et al., *CMS technical design report, volume II: Physics performance*, *J. Phys. G* **34** (2007) 995. (cited on page 95)

[69] *RooKeysPdf Class Reference*, `https://root.cern.ch/doc/master/classRooKeysPdf.html`, last accessed 2021/07/12. (cited on page 101)

[70] K. S. Cranmer, *Kernel estimation in high-energy physics*, *Comput. Phys. Commun.* **136** (2001) 198 [`hep-ex/0011057`]. (cited on page 101)

[71] *RooChebychev Class Reference*, `https://root.cern/doc/master/classRooChebychev.html`, last accessed 2021/07/21. (cited on page 104)

[72] C. W. Fabjan and F. Gianotti, *Calorimetry for particle physics*, *Rev. Mod. Phys.* **75** (2003) 1243. (cited on page 155)

Ringraziamenti / Acknowledgements

Dopo tutti questi mesi passati in Germania a parlare in inglese (provando a imparare giapponese e spagnolo) non so se sono ancora in grado di scrivere in italiano, ma proviamo!

Prima di tutto, vorrei ringraziare i miei **genitori**, senza i quali non sarei qui. Beh, certo, ovviamente anche per una mera questione biologica, ma senza la loro spinta alla curiosità, all'analisi critica, al mettere in dubbio e in discussione ogni preconcetto e a non dare niente per scontato, senza il loro supporto e il loro amore, non sarei la persona che sono oggi, incluso un* scienziat*, per quanto in erba. Grazie. Vi voglio bene.

Voglio ringraziare anche il resto del parentando: le mie **nonne** e l'amore e la pazienza che hanno avuto e hanno per il loro ”nipote sciagurato” che chiama e scrive *quasi spesso*, mia **cugina**, le mie **zie** e **zii**. Grazie a tutte, anche se sono così discontinu* nelle comunicazioni!

Another enormous ”thank yous” goes to my colleagues, with whom I worked throughout these last three and a half years, and even before.

Firstly, **Torben**, that supervised me and taught me so so much about how to survive a PhD, how to be a better scientist, and even a better person. His knowledge, support, expertise, and great balance between a demanding and an understanding side all served as a great motivation and push to move forward. He has a rare and precious combination of human and working side that I deeply appreciate.

Thanks to **Sam** now I am able to type faster, google better, and I have a wider knowledge of git and metal -although I think I miserably failed in my attempts to make him like folk and power-.

Thanks to **Savino** for his friendliness, help, and support, and for not being mad at me for having dragged him into my PhD committee.

A collective, huge thanking goes to my ”peers”: **Sascha**, **Munira**, **Abtin**,

Cyrille, **Filippo**, **Stella**, **Henrikas**; you all made these years fun, entertaining, and interesting, and even if in hindsight I may regret not having spent more time with you all during the less-constrained periods of the pandemic, out of isolation-induced inertia, I am deeply happy to have bounded with yous, and I look so much forward to maintain contacts and see each other at the next in-person B2GMs!

Also, a shout-out to all new colleagues I've met shortly, to my former colleagues in Pisa, and an apology to anyone that I forgot and I'll remember of when reading the printed version -lately my memory is underperforming, I meant no offense-.

Voglio salutare e ringraziare i miei "amici storici": **Sandro**, **Gianfranco**, **Niccolò**, che nonostante la distanza geografica, la pandemia, le difficoltà nelle vite di ognuno, e le ovvie e naturali divergenze che avvengono quando l'esistenza ci fa mutare in cambiamenti non sempre necessariamente sincroni, mi sono rimasti vicini e cari. Vi voglio un mondo di bene, e nei miei anni futuri ci immagino sempre vicini e legati, nonostante le distanze geografiche o di altra natura. Siamo amici da due decadi, e spero sia solo l'inizio.

Un altro grande grazie va a **Daniele**, con la sua pazienza, i suoi infiniti messaggi, le gigantesche discussioni sulle lingue, il farsi beffe dell'inglese per la lingua barzelletta che è, le riflessioni personali, filosofiche, politiche, esistenziali, tutti i birbs che mi manda. Non vedo l'ora di essere "vicini di casa" a Vancouver, burocrazia permettendo (perché il fato può essere oltremodo bizzarro, e ancora non mi capacito del fatto che senza metterci d'accordo finiremo con l'abitare nella stessa città, dall'altra parte del pianeta rispetto a dove ci siamo conosciuti).

Un grazie pieno di abbracci e cuoricini a **Claudio**. Le nostre chiamate Skype/Discord saranno state poco frequenti, ma indubbiamente corpose e immensamente gradite. Grazie per entusiasmarci insieme riguardo a Pokémon, all'ultimo videogioco che non abbiamo avuto tempo di giocare, all'ultimo fumetto che siamo riusciti a stento a leggere; per arrabbiarci insieme riguardo i difetti del sistema accademico e per tutti gli abbracci e le nerdate.

Grazie anche a **Carlo** e **Davide**, per quanto soprattutto in periodo pandemico i nostri contatti siano stati più diradati (come avrete notato, è una mia costante non essere particolarmente brav* in queste cose...). Spero di riuscire a ritrovare un buon equilibrio per sentirci meno raramente nei prossimi mesi, soprattutto quando saremo negli stessi continenti.

Grazie a **Bianca** per l'aiuto nel mio primo periodo teutonico quando tutto era alieno e strano, per il supporto nei periodi più difficili, per tutte le slide revisionate e le email corrette, per avermi ricordato di respirare e che molte cose sono sopravvalutate, e per le tante lettere e pensieri adorabili.

Grazie ad **Angela** per aver condiviso con me questa avventura amburghese, per le chiacchierate e il supporto.

Un grazie anche ai miei "amici di Yugi" **Matteo**, **Andrea**, **Francesco**. Tutte le volte che apro la chat e trovo 700 nuovi messaggi inviati fra l'una e le quattro di notte di amodio tantissimo. È sempre interessante confrontarmi con voi. E comunque, Plunder beste decke.

Talking about Plunder, yeah, I do wish to drop a non-negligible thank you to the **Plunder server** and its amazing people, that I grew to know and appreciate a lot during the last months, and that supported me during the tough times of the last months. That place is a little golden web oasis, brimming with memes, shitposting, but most importantly nice, welcoming, friendly people, willing to learn and teach and not to judge. Oh yeah, sporadically we also happen to talk about children card games, sometimes on motorcycles.

Also, a lil silly thanks to the **birbs** that visited my birbalcony over the last half year, and that brought me many serendipitous moments during the lone times of the thesis writing. I'll miss feeding yous.

A huge xiè xie to **Lux**, for the love and the splendid time we shared. You made me a better person on so many level, helping me discovering and understanding so much of my identity, my personality, and of the world around me, particularly its more queer and political side. You're a lovely banana, so kind and gentle, full of passion and will to help. Thank you so much for your immense support. I wish you all the best in your adventure in life, and I hope, in some way, to be a part of it.

A gigantic vielen Dankeschön goes to **Amelia**. Thanks so infinitely much for your fundamental support in this very intense year. You made me discover sides of me I didn't know of yet, made me understand more how I wish to live a relationship, and you make me a better person everyday. I am so happy to be with you. Your kindness, determination to grow better and stronger, your force of will, the strength with which you fight your battles for your rights, the passion, your approach to relationships, your self-caring: I love all of this, and much more. I love sharing our passions, arguing about our tastes for videogames, teaching you Yugi and learning about your job and meditation and your games and your channels. I love feeding the birbs together, teasing each other, cooing at you and cocoonify with you. I love our scooter-bike rides, buying grocery, hunting for GeoCaches, and BBQing with you. I want to keep sharing my life with you. I look so much forward for our big fridge. Ich liebe dich, meine Cocoon.

E infine, ultimo ma decisamente non ultimo, voglio ringraziare mio **nonno**. Non mi aspettavo di riuscire a scrivere queste righe senza piangere, e difatti non ci sto riuscendo. Grazie per aver instillato in me l'amore per la scienza. Grazie per le lunghe chiacchierate a gambe accavallate in cucina su questo o quell'argomento scientifico, o sulle storie della tua vita. Grazie per avermi aspettato ogni giorno per pranzo quando tornavo tardi da scuola. Grazie per avermi spiegato nei dettagli il funzionamento della bomba atomica per la mia tesina di terza media. Grazie per aver ascoltato pazientemente le mie bizzarre idee sul mio libro ancora da pubblicare ("Ci scusiamo con il lettore, ma il libro finisce per mancanza di personaggi"). Grazie di aver sopportato e avermi supportato nella mia ansia per la patente, e di aver festeggiato con me. Grazie di avermi amato. Sono ormai quasi dieci anni, e mi manchi ogni giorno, anche quando non me ne rendo conto. Avrei voluto mi vedessi raggiungere questo traguardo. Avrei voluto celebrare ed essere felice insieme a te. Ti voglio bene. Questa tesi è dedicata a te.

Eidesstattliche Versicherung / Declaration on oath

Hiermit versichere ich an Eides statt, die vorliegende Dissertationsschrift selbst verfasst und keine anderen als die angegebenen Hilfsmittel und Quellen benutzt zu haben.

I hereby affirm, in lieu of an oath, that I have written this dissertation myself and that I have not used any aids or sources other than those stated.

Hamburg, den 10.11.2021

Michael De Nuccio