



GHENT UNIVERSITY

DOCTORAL THESIS

Measurement of top quark pair-production in
association with a photon using the CMS detector
and Effective Field Theory interpretation

Author:
Gianny Mestdach

Supervisors:
prof. dr. Didar Dobur
dr. Kirill Skovpen

*A thesis submitted in fulfillment of the requirements for the degree of
doctor in the sciences: physics*

Faculty of sciences
Department of Physics and Astronomy

November 6, 2022

Supervisors:	prof. dr. Didar Dobur (UGent) dr. Kirill Skovpen (UGent)
Chairman of the Examination Board:	prof. dr. Christophe Detavernier (UGent)
Members of the Examination Board:	prof. dr. Robert Schoefbeck (HEPHY) prof. dr. Dirk Ryckbosch (UGent) prof. dr. Archisman Ghosh (UGent) prof. dr. Steven Lowette (VUB) prof. dr. Ilse De Looze (UGent)

“You know, in my opinion it’s sometimes even good to be ridiculous, if not better: it’s easier to forgive each other, easier to be humble; we can’t understand everything at once, we can’t start right out with perfection. To achieve perfection, one must first begin by not understanding many things. And if we understand too quickly, we may not understand well. ”

The idiot - Fjodor Dostojewski

GHENT UNIVERSITY

Abstract

Faculty of sciences
Department of Physics and Astronomy

Doctor of Philosophy

Measurement of top quark pair-production in association with a photon using the CMS detector and Effective Field Theory interpretation

by Gianni Mestdach

The standard model (SM) of particle physics is one of the most stringently tested scientific theories, and provides precise predictions for a wide range of processes involving fundamental particles. Strong theoretical and experimental arguments pointing towards its incompleteness do however exist, leading experimentalists to search for new particles predicted by proposed extensions to it. As these direct searches have until now been fruitless, and large increases in the energies probed in collider experiments are not expected any time soon, attention has shifted to precision measurements of the couplings of SM particles. Deviations of these couplings from SM predictions could indirectly indicate the existence of new particles, even when these can not yet be directly produced themselves. Interpreting these deviations in a useful way can be performed using the model independent framework offered by standard model effective field theory.

The cross section for the production of a top quark pair in association with a photon, with the W bosons from both quarks decaying leptonically, is measured using 138 fb^{-1} of proton-proton collision data. This data was collected in the 2016-2018 run period of the CMS experiment, with the LHC producing collisions at $\sqrt{s} = 13\text{ TeV}$. All leptonic final states are included ($e^{\pm}\mu^{\mp}$, e^+e^- , and $\mu^+\mu^-$), and the signal is defined so that it includes photons radiated by initial-state particles, top quarks, or any of their decay products. A fiducial phase space is defined with exactly one photon with a transverse momentum of above 20 GeV , and at least one jet identified as coming from the hadronization of a bottom quark. A fiducial cross section of $175.2 \pm 2.5\text{ (stat)} \pm 6.3\text{ (syst)}\text{ fb}$ is obtained using a profile likelihood fit to the measured transverse photon momentum distribution. Within this same phase space differential cross sections are measured as a function of kinematic observables of the final state particles, and compared to theoretical predictions. Both this measurement and a combination with the CMS measurement of the same process in a final state with just one lepton are interpreted in the framework of standard model effective field theory.

Acknowledgements

I have measured $t\bar{t}\gamma$, and it has measured me. There are many people who have supported me throughout these past few years, both on a personal and professional level, and for that I am grateful. Starting off with the person who started it all off, Didar. Thank you for opening the door to particle physics for me twice, and for the support and guidance over the years. Throughout my PhD I was supported by three of the best postdocs one can have to fortune to of working with. Tom, if there's any good code in my work, it's because of you. Thank you for your patience in transferring some of your skills to a mere mortal. Kirill, thank you for your tireless efforts, your highly Slavic words of wisdom, and the coordinated escapes from useless meetings. Joscha, thank you for your decisiveness, your surgical use of words, and the genuine excitement for physics you brought with you. I'll be surprised if I ever meet a more reasonable man.

I am also grateful to have been surrounded by some great colleagues, people both clever and kind. A special thanks goes out to Liam, who consistently managed to turn frustrations into jokes. In no particular order I would like to thank Tu Thong, Luka, Illia, Martina, Basile, Willem, Daniele, Marek, Niels, and anyone else roaming the hallways at INW. Stef and Tom, may your woks be full and your drawers filled with secrets. I also want to thank Christos and Ivan, for somehow brightening people's day by yelling at them from across the hallway. I have had the opportunity to work with some charming people from several other institutions as well. For the sake of brevity I will not name all of them, but I do wish to specifically thank Lukas for his consistent willingness to discuss problems and help out.

Of course there are plenty of people outside university who have supported me. These past few years I have often lived like a bear, leaving my cave mostly to gather food while growling half of the time. I have been lucky to have a family who made such a lifestyle possible, and tolerated it. For this help and patience I thank you. Moe, bedankt voor alles. Emma, the positivity I lost at some point I found with you (neatly folded in your pocket, as you do). Thank you for your endless support, your love, and your tolerance for my highly incoherent rants on physics. Last but not least I want to thank my local gang, most notably Yenn, Stijn, and Alexander, for keeping me around. Did I meet up enough with you guys? No. But did I at least become a wiser man in the meantime? Also no.

Contents

Abstract	vii
Acknowledgements	ix
Acronyms and names	xiii
1 Outline of the thesis	1
2 The standard model	3
2.1 Particles	3
2.1.1 Fermions	3
2.1.2 Bosons	4
2.2 Couplings and Feynman diagrams	5
2.3 On abstraction	6
2.4 Fields and matrix elements	7
2.5 Strengths and shortcomings of the standard model	9
2.6 The top quark	13
2.6.1 Properties of the top quark	13
2.6.2 Top quark production	14
2.7 The Higgs boson discovery	17
3 The LHC and the CMS experiment	21
3.1 Basic principles of particle detection	21
3.2 CERN	21
3.3 The large Hadron Collider	22
3.4 Beams, collisions, and luminosity	24
3.5 The Compact Muon Solenoid	25
3.5.1 Wish list for a general purpose detector	25
3.5.2 The CMS coordinate system	26
3.5.3 Detector components	27
3.5.4 Trackers	27
3.5.5 ECAL	28
3.5.6 HCAL	29
3.5.7 Solenoid magnet	29
3.5.8 Muon system	29
3.6 Physics object and event reconstruction	30
3.6.1 Triggering and data acquisition	30
3.6.2 Event reconstruction	30
3.6.3 Track reconstruction	31
3.6.4 Muon reconstruction	32
3.6.5 Calorimeter cluster identification	32
3.6.6 Electron and photon reconstruction	33
3.6.7 Hadronic jet reconstruction and heavy flavour identification	34
3.6.8 Global variables	35
3.7 Monte Carlo event simulation	35
3.8 Maximum likelihood fitting	38
4 Effective field theory	41
4.1 Introduction	41
4.2 Operator selection	44
4.3 Anomalous couplings	45
4.4 SMEFT event simulation and reweighting	46
4.4.1 The dim6top model	46
4.4.2 Event reweighting	46

4.4.3	Multi-dimensional parametrizations	47
4.5	Interpretations of standard model measurements	48
	Measurements of the $t\bar{t}Z$ cross section	48
	Search for new physics in top quark production with additional leptons	50
4.6	Increasing sensitivity through multivariate approaches	50
	Probing EFT operators in the associated production of top quarks with a Z boson	50
4.7	Interpreting variables sensitive to new physics	51
	Measurement of the energy asymmetry in $t\bar{t}j$ events by ATLAS	51
4.8	Global fits	53
5	Measurement of $t\bar{t}\gamma$ production	55
5.1	$t\bar{t}\gamma$ production and decay	56
5.2	Data samples and simulations	57
5.2.1	Data	57
5.2.2	Signal simulation and normalization	57
5.2.3	Background process simulation	58
5.2.4	Removing overlap between simulated samples	58
5.3	Particles and objects	59
5.3.1	Leptons	59
	Lepton reconstruction and identification efficiency calibrations	59
	ECAL scale and resolution corrections	60
	Muon Rochester corrections	60
5.3.2	Jets	61
5.3.3	Photons	61
5.4	Event selection	63
5.5	Trigger selection and efficiency measurement	65
5.6	Background prediction	68
5.6.1	$Z\gamma$ background	69
5.6.2	Nonprompt photon categories	70
5.6.3	Estimation of nonprompt photons	73
5.6.4	Other backgrounds	79
5.7	Systematic uncertainties	79
5.8	Fiducial region definition	86
5.9	Inclusive cross section measurement	87
5.10	Differential cross section measurement	92
5.10.1	Unfolding methods and tools	92
5.10.2	Differential cross sections	95
5.10.3	Normalized differential cross sections	96
5.10.4	Comparison with theoretical predictions	99
5.11	Effective field theory interpretation	102
6	Conclusion and outlook	107
	Summary	112
	Samenvatting	113
A	Charge misidentification prediction for a $t\bar{t}t\bar{t}$ search	115
A.1	Production and measurement of 4 top quarks	115
A.2	Charge misidentification of electrons and muons	116
A.3	Charge misidentification background estimate	116
	Bibliography	129

Acronyms and names

ALICE	A Large Ion Collider Experiment
ATLAS	A Toroidal LHC ApparatuS
BSM	beyond the standard model
CDF	Collider Detector at Fermilab
CERN	conseil européen pour la recherche nucléaire
CKM	Cabibbo–Kobayashi–Maskawa
CL	confidence level
CLIC	Compact Linear Collider
CMS	Compact Muon Solenoid
CP	charge and parity symmetry
CTF	Combinatorial Track Finder
D0	D0 is the label of the location along the Tevatron accelerator
DAQ	data acquisition
DY	Drell-Yan
ECAL	Electromagnetic Calorimeter
EDM	electric dipole moment
EFT	effective field theory
EWK	electroweak
FCC	Future Circular Collider
FSR	final state radiation
GSF	Gaussian sum filter
HB	HCAL barrel section
HCAL	Hadronic Calorimeter
HE	HCAL endcap section
HEL	Higgs Effective Lagrangian
HF	HCAL forward section
HLT	High Level Trigger
HO	HCAL outer section
ID	a set of identification criteria
ILC	International Linear Collider
IP	interaction point
ISR	initial state radiation
KF	Kalman filter

LEP	Large Electron Positron (collider)
LHC	Large Hadron Collider
LHCb	Large Hadron Collider beauty
LO	leading order
MC	Monte Carlo
MDM	magnetic dipole moment
ME	matrix element
MVA	multivariate discriminator
NLO	next-to-leading order
PD	primary data set
PDF	parton distribution functions
PF	Particle Flow
PL	particle level
PS	Proton Synchrotron
PSB	Proton Synchrotron Booster
PV	primary vertex
QCD	quantum chromodynamics
RF	radio frequency
RMS	root mean square
SF	same-flavour
SLD	SLAC Large Detector
SM	standard model
SMEFT	standard model effective field theory
SPS	Super Proton Synchrotron
SUSY	supersymmetry

Chapter 1

Outline of the thesis

The main research topic described in this thesis consists of the measurement of the cross section of top quark pair production in association with a photon ($t\bar{t}\gamma$). The first experimental evidence for $t\bar{t}\gamma$ production was provided by the CDF experiment, measured in $p\bar{p}$ collisions at a centre of mass energy of $\sqrt{s} = 1.96$ TeV [1]. Subsequently the process was observed by the ATLAS collaboration using pp collisions at $\sqrt{s} = 7$ TeV [2], followed by a measurement by CMS using 8 TeV pp collision data [3]. The first measurements at $\sqrt{s} = 13$ TeV were also already performed by the ATLAS collaboration, one using the lepton+jets final state [4], and one using the $e\mu$ final state in a combined measurement of $t\bar{t}\gamma$ and $tW\gamma$ [5]. Within the CMS collaboration two measurement of $t\bar{t}\gamma$ production using 3 years of proton-proton collision data collected at $\sqrt{s} = 13$ TeV were started in parallel around 2018. The first focusing on the case where one of the W bosons coming from the top quarks decays into leptons and one decays into quarks [6], and the second on the case where both W bosons decay into leptons [7]. The latter publication corresponds to the work documented in this thesis, where all three leptonic final states, namely $e^\pm\mu^\mp$, e^+e^- , and $\mu^+\mu^-$, are used. This includes the most precise inclusive cross section measurement of $t\bar{t}\gamma$ production to date.

Inclusive and differential measurements are performed within a fiducial phase space defined at the particle level. Both the particle level and reconstruction level selections involve requiring the presence of two oppositely charged leptons, exactly one photon, and at least one jet identified as originating from a b quark decay. The $t\bar{t}\gamma$ process is sensitive to the coupling between the top quark and the photon, which makes its measurement a test of standard model predictions and provides sensitivity to new physics phenomena beyond the standard model. The photon selection is however not aimed specifically at photons originating from top quarks, instead allowing them to be radiated from a top quark, an incoming quark, or any of the charged decay products of the top quarks. The remaining dominant source of background after the event selection stems from events with nonprompt photons, which can either correspond to misidentified jets, or real photons originating from particles inside hadronic jets or from secondary interaction vertices. This background is estimated from a control region in data, while other background sources are estimated using Monte Carlo (MC) simulations.

A profile likelihood fit is used to obtain the measured inclusive cross section from the distribution of the transverse momentum p_T of the reconstructed photon. The various systematic uncertainties, for example those associated with corrections applied to the simulated samples, are treated as nuisance parameters in the fit. For the differential cross section measurements no fit is performed, instead relying on the subtraction of the various backgrounds before removing detector distortions using an unfolding procedure.

By modelling the effect of relevant standard model effective field theory (SMEFT) operators on the signal prediction a similar fit can be used to obtain constraints on the Wilson coefficients corresponding to these operators [8, 9]. The use of the differential photon p_T information as input to the fit offers greater sensitivity than inclusive information would offer, thanks to the energy dependence of the new physics couplings. Even more stringent constraints are obtained from the combination of these results with those from the measurement in the lepton+jets channel.

Before delving into the specifics of the analysis a brief introduction to the standard model and the CMS experiment are provided. In the first part the emphasis will be put on providing an insight into how couplings appear in the calculation of production cross sections. Then, being one of the stars of the show, some context will be given to the top quark and its production in proton-proton collisions. At this point I could describe the $t\bar{t}\gamma$ measurement, its interpretation in the SMEFT framework, and subsequently sketch the context in which it should be seen in an outlook section. This measurement however forms a piece of a broader SMEFT puzzle, so that it seems more appropriate to describe the pieces that are already laid first. I will therefore first describe the framework of effective field theory (EFT), both theoretically and practically, while choosing examples most relevant to the $t\bar{t}\gamma$ process.

More information will be given than is strictly required to understand the EFT interpretation of the $t\bar{t}\gamma$ measurement. I wrote my master thesis aiming to provide transparent information on the basics of EFT for other students looking for an introduction, and I wish to extend upon this here. In order to not make assumptions of the reader's knowledge of maximum likelihood fitting the basic principles of this technique will also already be described early on. Once this is done an overview of published EFT results will be provided, showing the evolution of methods throughout the years, and giving examples of exciting more recently emerging ideas.

Publications contributed to by the author:

Peer-reviewed:

[7] A. Tumasyan et al. "Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the dilepton channel and effective field theory interpretation in proton-proton collisions at $\sqrt{s}=13$ TeV". in: (Jan. 2022). arXiv: [2201.07301](https://arxiv.org/abs/2201.07301) [[hep-ex](#)]

With the corresponding hepData record:

[10] CMS Collaboration. *Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the dilepton channel and effective field theory interpretation in proton-proton collisions at $\sqrt{s}=13$ TeV*. HEP-Data (collection). <https://doi.org/10.17182/hepdata.113657>. 2022

[11] D. Dobur et al. "Photon radiation effects in kinematic reconstruction of top quarks". In: *JINST* 16.12 (2021), P12001. DOI: [10.1088/1748-0221/16/12/P12001](https://doi.org/10.1088/1748-0221/16/12/P12001). arXiv: [2107.07586](https://arxiv.org/abs/2107.07586) [[hep-ex](#)]

Not peer-reviewed:

[12] G. Mestdach. "Inclusive and differential $t\bar{t}\gamma$ measurement in the dilepton channel and effective field theory interpretation". In: *14th International Workshop on Top Quark Physics*. Dec. 2021. arXiv: [2112.12541](https://arxiv.org/abs/2112.12541) [[hep-ex](#)]

Presentations by the author:

" $t\bar{t}\gamma$ production: dilepton channel" at the CMS TOP Group Workshop 2019, 20–21 Nov 2019

"Measurement of inclusive and differential $t\bar{t}\gamma$ cross section" Invited "CMS Joker talk" at TOP2021: 14th International Workshop on Top Quark Physics, 13–17 Sep 2021

Chapter 2

The standard model

The efforts of particle physicists in the second half of the 20th century have culminated in what is now commonly referred to as the standard model (SM) of particle physics. This theory boils down to a list of particles with certain features, and the interactions between them. These particles are quanta of what we call “fields”, as will be explained later. Even though it is the fields which are considered to be the fundamental substrate of the universe (within the picture of the standard model), it is usually both more intuitive and more useful to think in terms of particles. I will therefore start by giving a simple overview of the set of known fundamental particles and their interactions, and sketch the underlying mathematical framework afterwards.

Despite the fact that the SM theory has proven to be extraordinarily successful, the shared goal of experimental particle physicists is to find ways in which it falls short. Any significant deviation of a measurement from the corresponding theoretical prediction would be a hint in the direction of a more complete theory than the standard model. The model is known to be incomplete, as it fails to predict certain astronomical observations, and is fundamentally incompatible with general relativity. A more elaborate discussion of the strengths and shortcomings of the standard model is given after a brief explanation of the model itself. It will become clear later on that while the work documented here centres around the measurement of $t\bar{t}\gamma$ production, the main reason for interest in this process is the probing of the coupling of the photon to the top quark. Within the language of Feynman diagrams, which will also be introduced later on, these couplings take on the form of vertices. I will therefore centre the text around these vertices, both when it comes to the standard model and when discussing extensions to it. When it comes to the field theory underpinning particle physics I will limit the mathematical component to the bare necessities needed to link visual descriptions of diagrams to the calculation of probabilities.

2.1 Particles

In the context of particle physics, the term “particle” is implicitly used to refer to fundamental particles. That is to say, particles which cannot be further subdivided into even smaller components. This means that in terms of spatial dimensions they correspond to the mathematical concept of a point, and are otherwise characterized by certain fixed features. In the standard model these features consists of the particle’s mass, charges, spin, energy, and momentum. The spin is associated with the main division of the types of particles: particles with half integer spin values are referred to as fermions, and they are colloquially called the building blocks of the universe. Particles with integer spin values are called bosons, and they are for the most part interpreted as carriers of the various forces.

2.1.1 Fermions

Within the category of the fermions particles can be further distinguished by whether or not they carry a colour charge. The colour charge can take on 3 values, which are labelled as red, green, and blue, and the fermions which carry them are called quarks. This colour charge makes it so that quarks can never be found in isolation. Quarks either appear in groups of three, with each one carrying a different colour charge, or paired with another quark, with one carrying a colour charge, and the other carrying the corresponding anti-colour charge. In this way one can say that composite particles made up of quarks are always colour neutral, or “white”, as the naming borrowed for the mixing of actual colours implies. Protons and neutrons are examples of such groups of three quarks, containing two “up” quarks and one “down” quark, and one “up” quark and two “down” quarks, respectively. The full set of bound states of 3 quarks are called the baryons. The bound states consisting of 2 quarks are called mesons, which typically appear in the decays of heavier particles before decaying into stable particles themselves.




FIGURE 2.1: Overview of the particles in the standard model [13].

Figure 2.1 lists the 6 known types of quarks, along with their charges and masses. This figure contains all known particles of the standard model, including the so-called antiparticles of the quarks just described. All particles have a corresponding antiparticle, with the exact same properties except for an opposite electrical charge. Some particles, namely those with zero electrical charge, are simply the same as their antiparticle, and are therefore only listed once.

The subcategory of fermions without colour charges is called the leptons, the most commonly known of which is the electron. The electron and its heavier variants, the muon and the tau, each have a corresponding neutrino. Instead of the word “type” the term “flavour” is used to refer to specific fermions. The flavours of both the quarks and the fermions can be sorted into three generations, with each generation having a higher mass than the one before it. Ignoring several technicalities, as one necessarily does when describing particle physics outside its mathematical formulation, one can say that all common matter is made up of just the up quarks, down quarks (forming the protons and neutrons inside nuclei), and electrons. To understand the forces that bind them together and all necessary all non-gravitational phenomena to occur, one has to look at the bosons.

2.1.2 Bosons

The bosons are particles with integer isospin values, and can be thought of as being carriers or units of the forces of nature. Photons are massless particles which carry the electromagnetic interaction, and can be thought of as units of light. It is the field of which photons are the excitations that allows for the electromagnetic attraction between opposite charges, and causes electrons to “orbit” the positively charged nuclei. With only the electromagnetic interaction present nuclei would however spontaneously disintegrate, as the positively charged protons would repel each other. A different force, called the strong nuclear force, is mediated by particles called gluons and couples to particles with a colour charge. These are also massless particles, of which several subtypes exist with different colour charge combinations. The strong nuclear force has the peculiar property of getting stronger as the distance between particles increases, and the corresponding field causes the colour neutrality described in the previous subsection. The theory describing the strong force is known as quantum chromodynamics or QCD.

These two boson types provide a picture, albeit a very simplistic one, of how atoms and molecules and by extension most of the physical world is constructed: Nuclei made up of quarks, bound together

by the exchange of gluons, surrounded by clouds of electrons, which are also involved the interactions which bind atoms together into molecules, and set the rules though which chemical interactions occur. For a more realistic picture one has to step away from seeing the gluons as merely mediators of the force that ties the quarks together, and consider them as entirely relevant degrees of freedom. Around 99% of the rest mass of the proton is for example present in the form of kinetic and potential energy of the gluons and sea quarks, virtual quark pairs continuously being formed and annihilated, with only a small fraction corresponding to the rest masses of the valence quarks themselves. A less prominent force in everyday life is called the weak nuclear force, and its force carrying bosons are not massless but in fact much heavier than the quarks and electrons. There is the electrically neutral Z boson, which is its own antiparticle, and the W boson, of which a charge $+1$ and -1 version exists. These particles appear frequently in particle physics experiments, but in nature they are essentially the particles which allow certain types of nuclear decay to occur. The electromagnetic and weak nuclear forces are unified in the electroweak (EW) force. One final type of boson is the Higgs or “Brout-Englert-Higgs” boson, which is not directly associated with a force of nature, but can be thought of as the by-product of the functioning of its field. It is the Higgs field which is responsible for giving mass to nearly all fundamental particles.¹ [14, 15] The coupling between the Higgs boson and fermions is described by the Yukawa interaction, which is proportional to the particles’ mass.

2.2 Couplings and Feynman diagrams

To complete this simplified picture of physics’ current understanding of nature a sketch of how the particles listed above interact with one another is needed. A Feynman diagram, as shown in figure 2.2, has a clear mathematical meaning within quantum field theory calculations. It can however also be used as an intuitive way of thinking about the processes occurring in particle physics experiments. The vertical axis can be thought of as a spatial dimension or what is informally called a direction, while the horizontal axis represents time. The first diagram shows two gluons merging into one, which splits or “decays” into a top quark and an antitop quark, after which the antitop quark radiates a photon. Top quarks almost always decay into a W boson and a b quark, as is also shown in the diagram. The resulting particles can decay further until only stable particles are left. While such decays quickly happen spontaneously, the inverse process in which high mass particles are created is very rare as it requires gluons, or other particles, to collide at energies high enough to produce the heavy top quarks. Points where the lines representing particles meet are called vertices, and the rules regarding which vertices can exist and how commonly they occur is exactly what the SM theory encodes. An overview of all SM vertices is shown in figure 2.3.

The job of a theoretical particle physicist is now to calculate the probability of a certain combination of particles to be produced based on the rules established so far. Meanwhile it is the task of the experimentalist to measure these probabilities (as there is no way of forcing nature to produce a certain outcome), and through comparison with the predictions to test the rules that produced these predictions.

Notice that a listing of the SM particles and their properties and this brief description of their couplings are fundamentally the same thing. The various features of the particles, from their charge and spin to their mass and momentum, only have a meaning in light of their interactions with other particles, and hence their couplings. In other words any feature ascribed to a particle corresponds to one of its interactions. This can be thought of as an inherent aspect of a fundamental theory, as if an object were to have a feature that is not relevant to its behaviour, this feature could also not be measured in any way. In classical physics abstractions are typically made, as for example the colour of an object is not relevant to the calculation of how quickly it rolls down a ramp. Any fundamental theory can inherently not have such distinctions between relevant and irrelevant, as the abstraction and the theory are one and the same.

Additionally one could now imagine that additional particles exist, which means that their corresponding vertices could be used to construct additional Feynman diagrams. On a physical level this means that the probabilities for any process in which they appear would be modified with respect to the SM picture. This idea will be explored further in the section on effective field theory, where

¹It is not yet clear whether the neutrinos obtain their mass through coupling to the Higgs field.




FIGURE 2.2: A Feynman diagram in which a top and antitop quark are produced, one of which radiates a photon, after which both decay decay into a charged lepton, a neutrino, and a bottom quark (left), and a simple lepton pair production diagram (right) [16].

such new particles and interactions will be introduced in a way that explains why their hypothetical existence is not necessarily immediately apparent, even when performing measurements using highly precise detectors. Before moving on however I will give a few thoughts on abstractions themselves, before expanding on the ones I have been describing here.

2.3 On abstraction

This thesis gives a brief description of the SM theory, as well as the tools used to study it. The Large Hadron Collider (LHC), the Compact Muon Solenoid (CMS) detector, and the computers used to perform analyses are the obvious tools needed to perform this work, and will be described later on. Nevertheless, the mental tool called “abstraction” is what allows one to build all others.

The use of abstraction is exceptionally clear in the context of programming. It would be impossible to produce anything but the simplest of programs if one had to be mindful of everything that needs to happen down to the level of the flow of electrons on a chip. Instead, transistors were developed up to a level where their behaviour is so reliable their operation might as well be described as a mathematical operation on binary information. Upon this level logic gates can be built, and the underlying level is once again abstracted away. After climbing several such layers thanks to advances in electronics and computer science this generation of programmers find themselves harnessing a set of functions like “loadFile(fileName)”, whose intuitive behaviour is reflected in how close its naming choices come to everyday language. It is only when an error shows up that the complexity of the underlying layers shows up again. Aside from true mistakes this is usually the case when the code underlying the function is not designed to handle a certain kind of input, or is being used outside the scope in which it has been tested. One then has to descend one or more of the layers.

Physics itself relies on similar abstractions, but aimed at processing abstract and complex ideas using a less perfect computer contained inside an ape. An additional use of abstraction in this context comes in the form of approximations. For some matters like the description of a wheel rolling down a slope approximating the situation to that of a perfect circle on a flat plane is mostly a practical consideration. In other cases, common in all but the earliest branches of physics, approximation is not only convenient, but also necessary. In quantum field theory it is impossible to calculate the production probability of any process to absolute precision, even if the constants needed for the calculation were known with absolute certainty. Instead, a series expansion is used, allowing for the calculation of a cross section up to some limited level of certainty. To each of the terms in these series expansions corresponds a Feynman diagram, which is the example of abstraction most relevant to particle physicists. One can think of physical processes in terms of these diagrams, communicate using them, and some of them are even named. This is all true despite the fact that they have no one-to-one link to what is happening in nature. Nevertheless the intuitive conclusions that can be drawn from thinking of processes in terms of these approximations, even if only those of the first order are considered, are generally correct. So much so that it often has to be clarified to students that they have no direct

correspondence to reality, and are in fact nothing more than mathematical tools.

2.4 Fields and matrix elements

In classical mechanics, the evolution of a system can be described by constructing its lagrangian, given by:

$$\mathcal{L} = T - V \quad (2.1)$$

Here, T and V are respectively the kinetic and potential energy of the system, both of which can have any number of terms with their own dependences. The integral of this lagrangian between two points in time is called the action. The principle of stationary action now states that the trajectory an object will take (or how the system will evolve) is the one for which the action is stationary, which means that its functional differential is zero.² Or mathematically:

$$\delta \int_{t_1}^{t_2} \mathcal{L} dt = 0 \quad (2.2)$$

The general solution to this comes in the form of the Euler-Lagrange equation

$$\frac{d}{dt} \frac{\partial \mathcal{L}}{\partial \dot{x}} = \frac{\partial \mathcal{L}}{\partial x} \quad (2.3)$$

which is in essence a generalized version of Newton's second law.

The same principle extends into quantum field theory, where the fundamental components are not objects but fields, and the lagrangian is therefore replaced by the lagrangian density. Values like momentum and energy are now represented by operators acting on these fields. A quantum field theory is therefore a theory consisting of a set of these fields, which are invariant under certain transformations, and whose evolution is described by the minimal action principle applied to their total lagrangian density. These invariances under certain transformations are called symmetries, and each one corresponds to the conservation of some value.³

The simplest symmetries are those for which the lagrangian is invariant under the same transformation being applied everywhere in space. In gauge theories, of which the field theory underlying the standard model is an example, there is an additional form of symmetries which have a dependence on the location in space. These so-called local symmetries correspond to transformations which leave the lagrangian density unaffected exactly the way global symmetries do, but the exact transformation they entail is spacetime dependent.

A more rigorous mathematical picture of the quantum field theory underlying the SM is of lesser relevance to sketching the context needed for this work. I will therefore move on to describing what the calculation of a production cross section for a process looks like, or in other words what constitutes a prediction in the context of particle physics.

As a simple example consider the right-hand diagram of figure 2.2 taken from "Feynman diagrams for beginners" [16], which shows the annihilation of an electron and a positron, forming a photon decaying into a muon-antimuon pair.

Following the so-called Feynman rules this diagram corresponds to the matrix element

$$\mathcal{M} = \frac{e^2}{(p_1 + p_2)^2} [\bar{u}_3 \gamma_\mu v_4] [\bar{v}_2 \gamma^\mu u_1] \quad (2.4)$$

Where the p parameters represent the momentum of the incoming particles, and u and v correspond to the incoming and outgoing particles, which have a momentum and spin dependence. The differential cross section this corresponds to can then be written as

$$d\sigma = \frac{1}{u_\alpha} \frac{1}{2E_1} \frac{1}{2E_2} |\overline{\mathcal{M}^2}| dLips_2, \quad (2.5)$$

²This principle is a universal truth observed in any physical system, and is in itself not derived from anything else.

³Which is another general truth concerning symmetries known as Noether's theorem.

where the line over M refers to the fact a summation over all physically allowed spin permutations has been performed. $dLips_2$ represents the two-particle phase space to be integrated over. After integration over the full phase space a value for the total cross section associated with this process is obtained. A full calculation of the cross section for the $e^+ + e^- \rightarrow \mu^+ + \mu^-$ process however requires that matrix elements for all possible Feynman diagrams are taken into account.

Luckily in electroweak processes more complicated diagrams involve more electroweak vertices, each of which makes the contribution of the diagram smaller given the small value of the electroweak coupling constant. Because of this higher order diagrams always have smaller contributions to the cross section than lower order ones. Outside the electroweak processes this is not necessarily the case. The specifics of this example are not relevant to this work, but principles that are important from the experimentalist's perspective can be highlighted here:

- In an experimental setup only the type of particle to be accelerated, and the energy with which they are made to collide can be controlled. This is a very basic idea, but explains why it is necessary to produce billions of collisions to measure just a few hundred events of a process of interest.
- The quadratic dependence of the cross section on the matrix element (or rather the sum of individual matrix elements) is readily apparent.
- Predictions in particle physics provide only probabilities for the production of a certain combination of particles, within a certain kinematic phase space. Similarly, if a leading order Feynman diagram shows a virtual Z boson forming two leptons measurements of their kinematic features will for the most part correspond to what one might expect from a Z boson decaying into two particles, but it is unphysical to say that one knows that they were produced in a Z boson decay. Nevertheless people discuss processes using this kind of language, since it makes for an intuitive way to communicate ideas.

*“There’s no earthly way of knowing
Which direction we are going
There’s no knowing where we’re rowing
Or which way the river’s flowing”*

- Willy Wonka

2.5 Strengths and shortcomings of the standard model

The standard model is without a doubt one of the best tested theories of any scientific field. At the same time it allows to make predictions for a wide range of phenomena, over a wide energy range, with no significantly large discrepancies being found so far. An obvious test of the predictive power of the SM is the comparison between precise cross section measurements and the corresponding predictions. As different processes have different dependences on electroweak and QCD couplings cross section measurements form indirect checks of the SM description of these couplings. Figure 2.4 summarizes a wide range of such comparisons, showing good agreement with precise theory predictions. What makes this even more impressive is the fact that the cross section values being shown span multiple orders of magnitude.




FIGURE 2.4: Summary of several cross section measurements by CMS compared to theory predictions. For most processes results are shown for several centre-of-mass energy values, with the 7 and 8 TeV cases corresponding to run periods of the LHC before 2016 [18].

More direct measurements of SM parameters are also continuously performed as new data is collected. Central to the SM theory is the strong coupling strength, for which nearly all of the major collider experiments have published measurements. Figure 2.5 shows a summary of these measurements, and the corresponding global average obtained from them. Masses of SM particles also form a target for precision measurements.




FIGURE 2.5: Measurements of the strong coupling constant at the Z boson mass and their combined average, compared to the SM theory prediction [18].

A precision measurement of the W boson mass by the CDF collaboration recently garnered attention because of a highly significant deviation with respect to the value predicted based on global electroweak fits [19]. Previous measurements by other experiments were however found to be in agreement with the SM prediction, as shown in figure 2.6. Given the complexity of modern particle physics analyses it is not unthinkable for a mistake to sneak past the various review steps. It will therefore remain to be seen whether W boson mass measurements should be seen as a proof of the strength of the standard model, or rather a first significant indication towards its shortcomings from the collider physics side, based on this new result.




FIGURE 2.6: Measurements of the W boson mass by various experiments, most notably the result by the CDF collaboration which shows a strong tension with the SM prediction [19].

Despite the wide range of evidence for the validity of the SM, there are several strong yet indirect signs the framework of the standard model must be incomplete. First of all it stands entirely separate from the theory of general relativity, and attempts to combine the two have so far been unsuccessful. Any so-called “theory of everything” should be able to merge the two.

Within the standard model the mass of the Higgs boson, measured to be 125.25 ± 0.17 GeV, also raises questions [20–22]. While its measured mass, which the standard model does not directly predict, is not in disagreement with the theory, it is nevertheless peculiar that its value is so low. The observed mass of the Higgs boson is affected by loop corrections involving any particles that couple to it, most notably the top quark. Because of this the calculation of the Higgs mass contains several quadratically divergent terms. This essentially comes down to adding and subtracting numbers of multiple orders of magnitude larger than the Higgs boson mass, and their values cancelling out almost perfectly to leave just the value of around 125. This calculation is mathematically acceptable, but seems unnatural, and is therefore called a “naturalness problem” or more specifically the “hierarchy problem”. A theory, or rather a group of theories, that offers a solution to this is supersymmetry (SUSY). Within SUSY a supersymmetric partner exist for every SM particle, and their presence in loops cancels out the SM terms in a natural way. Unfortunately despite many efforts to do so no SUSY particles have been discovered, and the subset of models which have not yet been excluded by measurements is becoming very small.

Also outside the realm of particle physics observations have been made which the standard model cannot explain. Cosmological models show that only around 4% of the energy present in the universe corresponds to the particle content of the standard model [23]. Another component that is most commonly thought to be made up of particles is that of dark matter. When comparing the rotational curves of spiral galaxies with mass measurements based on the light emitted from them a discrepancy is consistently found [24]. This could be explained by the presence of a cloud of particles that interact gravitationally, but do not otherwise couple significantly to any of the SM particles. The same conclusion is drawn when looking at the spacetime curvature around distant galaxies. As massive objects curve spacetime around them they bend the trajectory of light moving past them. The result of this is an optical phenomenon called gravitational lensing, which can be exploited to measure the total mass present in these galaxies. Once again the mass measured in this way exceeds the mass that is obtained from measurements of the emitted radiation [25]. In some cases the centre of mass of these objects is even found to be shifted with respect to the visual centre [26]. An image of the typical example of a cluster exhibiting this feature, called the “Bullet cluster”, is shown in figure 2.7.




FIGURE 2.7: Image of a merging galaxy cluster (“Bullet cluster”) taken using the Chandra X-ray observatory, overlaid with green lines showing where gravitational lensing observations indicate the mass of the cluster is located. The difference between where X-rays are emitted and hence most regular matter is located and the centre of mass areas seem to indicate that while the regular matter is “colliding” some other massive component which does not emit radiation has continued on its path uninterrupted [26].

While many additions to the SM that could explain these observations have been proposed, direct production of these particles at colliders have not been observed (for example in the form of additional missing transverse energy) [27, 28], and results from experiments focused on dark matter measurements have been either unfruitful or not sensitive enough [29].

Another question from the field of cosmology left unanswered is the asymmetry between matter and antimatter seen in the universe. If antimatter were common in the universe, there would be a constant emission of high energy radiation coming from regions where matter and antimatter annihilate upon coming into contact. No such sources are observed, so that one can be quite certain that antimatter is very rare. If no processes were to exist in which CP-symmetry (invariance under the simultaneous inversion of the charge and parity of all particles in a process), such an asymmetry would be completely unexplained. In reality some CP violation has been observed in electroweak processes, but the rate at which this occurs makes it insufficient to fully explain the observed asymmetry [30]. There is no fundamental reason for QCD not to violate CP symmetry, but nevertheless it has been measured to be either very small or zero. This fact in itself once again raises questions regarding why the CP violation in the QCD sector is so small, similar to the situation regarding the Higgs boson mass. In general unexplained “coincidences” like these are referred to as finetuning problems.

In collider experiments the most straightforward approach to looking for new physics is simply to generate higher and higher energy collisions, so that heavier particles can potentially be produced. So far however no particles with masses above that of the top quark have been observed. By performing precision measurements of the couplings of particles that are known to exist, most notably those of the top quark, one can however look for signs of physics beyond the standard model (BSM) without having to reach the energies required to produce new particles directly. The idea behind this will be explored in detail later on.

2.6 The top quark

2.6.1 Properties of the top quark

While there is no hard line that separates it from the other quark flavours, the top quark is of great interest from a theoretical perspective due to its large mass of 172.76 ± 0.30 GeV [31]. The existence of the top quark was first proposed in 1973 by Kobayashi and Maskawa, a year before even the charm quark was discovered. Adding a third generation of quarks, thereby extending the principle of the 2x2 Cabibbo matrix to the 3x3 CKM matrix, was a way to describe the measured CP violations that could not be explained in the Cabibbo model [32]. The bottom quark was subsequently discovered in 1977 at Fermilab [33]. It took until 1995 however for the top quark to be discovered by both the CDF [34] and D0 [35] experiments at Fermilab (using the Tevatron collider). Measurements of its mass have since been performed by CDF, D0, CMS, and ATLAS, and the value quoted above is a combination of several of these measurements.

Since a particle's Yukawa coupling is directly proportional to its mass, the top quark also has the largest Yukawa coupling with a value of close to one. This large coupling is what makes the top quark theoretically interesting, as many proposed extensions to the standard model feature similar mass dependent couplings. The top is therefore a probe for new physics, and deviations in its couplings from SM predictions would be indicative of physics beyond the standard model.

In addition to this its large mass means that it has a significantly shorter lifetime than other quarks, of the order of 10^{-25} seconds. Because of this it has no time to hadronize before decaying, and even short-lived bound states involving top quarks do not exist [36]. These decays can only occur through the weak interaction, where the result is nearly always a W boson and a bottom quark. Top quarks therefore do not form jets themselves, and are instead characterized by the b quark jet and the decay products of the W boson coming from their decay. Decay into a W boson and either a strange or down quark is not forbidden, but is very rare due to the fact that the CKM matrix is nearly diagonal [31].

Precise measurements of the top quark mass are non-trivial, and can be performed in two main ways: The most straightforward approach consists of reconstructing the top quarks from their decay products in $t\bar{t}$ events. This comes down to performing a maximum likelihood fit over the possible ways in which the reconstructed final state particles can be combined, while taking into account probability distributions for the slight mismeasurement of their kinematic features. Less direct methods involve inclusive or differential cross section measurements of processes known to be directly dependent on the top quark mass. Direct approaches typically lead to the highest experimental precision, with the main drawback being the strong reliance on the MC simulation.

The highest precision obtained so far (not including combinations of measurements) was achieved in a CMS analysis using one year of data (2016) collected at $\sqrt{s} = 13$ TeV looking at lepton+jets events in a $t\bar{t}$ enriched selection [37]. This measurement uses a direct kinematic reconstruction approach as described above, but supplements this information with variables like the invariant mass of the lepton and b quark jet system. In cases where the fit used in the kinematic reconstruction results in a low goodness-of-fit value the reconstructed top quark mass cannot be relied on, so that these simpler variables can provide valuable additional information in these cases. The top quark mass and the corresponding uncertainty extracted from this measurement ends up being 171.77 ± 0.38 GeV.

The earlier statement about the Yukawa coupling being directly proportional to the top quark mass is of course also one of the tenets proposed by the standard model, so that measuring the Yukawa coupling directly forms another test of the SM and the Higgs mechanism. Measurements of top quark pair production in association with a Higgs boson ($t\bar{t}H$) are essentially also measurements of this coupling, with the observations by both CMS and ATLAS finding agreement with the SM predicted cross sections [38, 39]. In $t\bar{t}$ events the invariant mass of the top quark pair and the rapidity difference between the top quarks are sensitive to the Yukawa coupling, so that this process can also be used to measure the Yukawa coupling of the top quark directly. The CMS collaboration has published such an analysis [40], with a measured coupling of $Y_t = 1.16^{+0.24\%}_{-0.35\%}$ where the SM prediction is quoted to be $Y_t = 1^{+0.30\%}_{-0.57\%}$.

2.6.2 Top quark production

Despite the fact that their large mass has historically been the limiting factor in the production of top quarks, once the energy threshold for their production is achieved they are predominantly produced in pairs. The reason for this is that top-antitop pairs can be produced through the strong interaction, while single top production has to occur through the weak interaction. This is illustrated in figure 2.8 where the production cross section for both options are shown as a function of the centre of mass energy $\sqrt{s}$. Considering all production channels together top quark pair production has a cross section 3 times as large as single top production at LHC energies.




FIGURE 2.8: The production cross section for $t\bar{t}$ production [41] (left) and the different single-top production channels [42] (right) as a function of the centre of mass energy in proton-proton collisions.

With the LHC operating at $\sqrt{s} = 13$ TeV, $t\bar{t}$ events are produced at 100 times the rate of Tevatron, which has earned it the title of “top quark factory”. This has allowed for not only inclusive cross section measurements, but also differential and multi-differential cross section measurements of $t\bar{t}$ production to be performed. The process through which top quark pairs are produced is also somewhat different between Tevatron and the LHC. At the operating centre of mass energy of Tevatron about 85% of the production occurs through quark-antiquark annihilation, while around 15% occurs through gluon-gluon fusion. At LHC energies these ratios are flipped. More relevant to this work is the fact that the high energy and integrated luminosity have made it possible to perform precision measurements of $t\bar{t}$ production in association with a boson, collectively referred to as $t\bar{t}X$ measurements. Representative leading order diagrams for $t\bar{t}$, $t\bar{t}Z$, $t\bar{t}W$, $t\bar{t}H$, and $t\bar{t}\gamma$ are given in figure 2.9 along with their approximate predicted cross sections. The fact that the photon can be reconstructed directly sets $t\bar{t}\gamma$ production apart from the other $t\bar{t}X$ processes. $t\bar{t}Z$, $t\bar{t}W$, and $t\bar{t}H$ form backgrounds to each other, so that uncertainties in measurements of any of these processes are inherently linked. Meanwhile $t\bar{t}\gamma$ forms a background to all of these processes though photons converting into electron pairs, but the requirement of a high p_T reconstructed photon in a $t\bar{t}\gamma$ measurement efficiently removes backgrounds from the other $t\bar{t}X$ processes. The photon being massless also gives $t\bar{t}\gamma$ production a significantly larger production cross section.




FIGURE 2.9: Example leading order diagrams for $t\bar{t}$ and the various $t\bar{t}X$ processes along with their predicted cross sections. In cases where two relative uncertainties are listed these respectively correspond to the uncertainty in the factorization and normalization scales, and the uncertainty in the parton distribution function.

The couplings between the top quark and the vector bosons is much less accurately known than for the lighter quarks, so that there is great interest in finding out whether deviations from theoretical predictions are hiding in this relatively unexplored corner of the standard model. Undiscovered high mass particles could manifest as modifications to these couplings at energies far below their invariant mass, which can be parametrized in a model independent way using effective field theory, as will be shown later.

Since the top quark nearly always decays into a bottom quark and a W boson the final states of the $t\bar{t}$ and $t\bar{t}X$ processes are categorized based on what the W bosons decay into. The W boson has a $67.41 \pm 0.27\%$ measured probability of decaying into a pair of hadrons, and a $3 \times (10.86 \pm 0.09\%)$ probability to decay into a lepton and a neutrino [31]. All other decay channels have a negligible branching fraction. Because of this $t\bar{t}$ decays result in either 6 jets (hadronic channel), 4 jets plus a lepton and a neutrino (semi-leptonic or lepton+jets channel), or 2 jets, 2 leptons, and 2 neutrinos (dilepton channel). Independent of the decay channel two of the jets come from b quarks. For $t\bar{t}X$ analyses the final states are then combinations of these channels and the decay products of the additional boson. Often the measurement of a process is limited to a specific decay channel, or the work is divided between students or groups based on the channel and combined later on. The reason for this is that different channels of the same process do not necessarily come with the same challenges. The higher branching fractions for hadronic decays typically lead to better statistics in channels with fewer leptons, but are associated with final states that are more difficult with respect to background rejection. Meanwhile fully leptonic final states can more easily be separated from backgrounds, but

suffer from reduced statistics.

While it fits less obviously in the category of $t\bar{t}X$ processes, the production of four top quarks is interesting for similar reasons. An example leading order diagram for its production is also shown in figure 2.9. The large Yukawa coupling of the top quark makes it so that diagrams involving intermediate Higgs bosons form a large contribution to its cross section at leading order. Because of this the process is also a sensitive probe of the top quark Yukawa coupling.

As mentioned the cross section for single top production is smaller than that of top quark pair production at LHC energies. Its three production modes, namely s-channel, t-channel, and tW -channel, are typically considered separately. Representative leading order diagrams for each of them are shown in figure 2.10 along with the corresponding cross sections. The smaller cross sections and larger backgrounds for single top production processes in association with a boson make them somewhat less interesting to study in the context of top-boson couplings than associated top quark pair production. All three production channels are however directly proportional to the value of $|V_{tb}|$ from the CKM matrix [42]. In this regard measurements of single top production are complementary to $t\bar{t}$ measurements, where $|V_{tb}|$ is obtained indirectly by measuring $R_b = \frac{|V_{tb}|^2}{|V_{td}|^2 + |V_{ts}|^2 + |V_{tb}|^2}$ and making the assumption that $|V_{td}|^2 + |V_{ts}|^2 + |V_{tb}|^2 = 1$. In single top production this assumption is not necessary, and therefore it can also be used to test this assumption.




FIGURE 2.10: Leading order diagrams for single top quark production in the t-channel (left), s-channel (middle), and tW -channel (right), and their associated cross sections [42, 46, 47]. The relative uncertainties respectively correspond to the uncertainty in the factorization and normalization scales, and the uncertainty in the parton distribution function.

Despite the challengingly small predicted fiducial cross section of 406^{+25}_{-32} fb, ATLAS recently presented the observation of single top quark production in association with a photon with a significance of 9.1σ [48]. The analysis selection targets the t-channel and uses neural networks to separate the signal from the various backgrounds. Up to this point observations of the production of a single top quark in association with a boson had been limited to the Z and W bosons [49–52]. The dominant background to this process comes from $t\bar{t}\gamma$ production.

Aside from the $t\bar{t}\gamma$ measurement documented in this thesis the CMS group at the university of Ghent is also involved in many more analyses involving top quarks. This includes a simultaneous measurement of the $t\bar{t}Z$ and $t\bar{t}W$ processes [53], and subsequent separate measurements of $t\bar{t}Z$ and $t\bar{t}W$ using larger data sets [54, 55]. The first experimental observation of single top quark production in association with a Z boson also resulted from work mainly performed by members of the UGent group [49], and was followed up by a publication of inclusive and differential tZq cross section measurements using a larger data set [56]. A search for four top production using multiple final states is ongoing, and could potentially lead to the first experimental observation of this process.

2.7 The Higgs boson discovery

A wide range of measurements have been performed using the LHC and its detectors, but the crowning work is still the first observation of the Higgs boson by the CMS and ATLAS collaborations in 2012 [20, 21]. Before moving on it seems only appropriate to describe how the last main piece of the SM puzzle was laid. To a large extent the detector designs of CMS and ATLAS were aimed at making the discovery of the Higgs boson possible, and some of the tools developed in the context of this analysis are used to this day. The fitting procedure described in the previous section is for example nearly always performed using the “Higgs Combine tool”.

The analysis by CMS used two data sets consisting of 5.1 fb^{-1} and 5.3 fb^{-1} worth of proton-proton collisions, respectively measured at centre-of-mass energies of $\sqrt{s} = 7 \text{ TeV}$ and $\sqrt{s} = 8 \text{ TeV}$. Higgs bosons were expected to decay into pairs of particles with branching ratios dependent on the masses of the decay products, and of course the mass of the Higgs boson itself. The mass of the Higgs boson is a largely free parameter of the standard model, so that its mass needs to be obtained from measurements. The ZZ , $\gamma\gamma$, WW , $\tau\tau$, and $b\bar{b}$ decay channels were used, with the first being considered as the “golden channel” thanks to the high mass resolution its 4 lepton final state can provide. Despite its small branching ratio the $\gamma\gamma$ decay channel was also considered to be highly important, especially for the lower mass values of the Higgs boson. The dependence of the branching ratios on the Higgs mass is shown in the left-hand graph of figure 2.11. The right-hand graph shows the invariant mass distribution of the 4-lepton system within an event selection focused on the $H \rightarrow ZZ \rightarrow 4l$ decay channel. This channel alone delivered an observed significance of 3.2σ , while the combination of all five final states resulted in an observed significance of 5σ , making it significant enough to speak of a discovery. The observed bump was found to be most compatible with a particle with a mass of 125.3 GeV . Further measurements were needed to confirm that this particle is indeed the scalar boson hypothesised by Brout, Englert, and Higgs [14, 15]. The fact that it decays into a pair of photons did however already establish that it could not be a vector boson.




FIGURE 2.11: Left: Branching ratios of the Higgs boson as a function of its mass [57]. Right: The distribution of the four-lepton invariant mass from CMS Higgs discovery paper, within a selection enriched in $H \rightarrow ZZ \rightarrow 4l$ events [20]. The points represent data, the filled histograms corresponds to the background, and the open histogram corresponds to the background expectation with the predicted peak from a 125 GeV Higgs boson added on top.




FIGURE 2.3: All vertices encoded by the standard model [17]. The symbol m_B refers to any massive boson, and m can be any massive particle excluding neutrinos.

A word of thanks

I have limited the acknowledgements section to the people with whom I have lived or worked directly. The work that went into an experiment of this size, operating in a field of science that has been growing for many centuries, can however not be reduced down to a handful of people. It has taken many generations of mathematicians, theoreticians, experimentalist, engineers, and supporting personnel to make this thesis possible. I find it most appropriate for me to sing their praises between the lines attempting to describe the tools they have given us. While I have had the privilege to work in comfort, backed by many people and institutions, this is not the case for everyone. I therefore want to thank everyone who has contributed to this field in any way, to the capacity their environment has allowed them. A special thanks should go out to the engineers and physicists involved in the development of the detector and accelerators systems. Having had the chance to get to know some of them, I have seen the work and brilliance that goes into every bit of detector performance.

Finally, I wish to point out a source of support perhaps too often taken for granted. The work documented in this thesis coincided with a pandemic, the threat of a large scale military conflict on European soil, and an economic crisis. Nevertheless I saw fundamental research proposals receive funding, overheard people who dread studying physics have nuanced discussions about the benefits of nuclear and renewable energy, and got to read about the James Webb Space Telescope finally aiming for the stars. As the field of physics becomes ever more abstract, the accessible communication of these ideas will become more challenging. Yet it is this we will need in order to continue on this journey, as we tell society ‘bear with us, we’re onto something’.

Chapter 3

The LHC and the CMS experiment

The measurements described in this work are performed using data collected using the CMS detector. The various sub-detectors of this device, their features, and how the information they collect is unified and stored is described in the following few pages. Unlike experiments in which particles originating from cosmic rays or nuclear decays of unstable elements are of interest, in CMS so-called “events” are produced by colliding protons at high energies, and therefore the methods by which these protons are accelerated and made to collide are first described. The accelerator used to produce the collisions which CMS attempts to reconstruct is the LHC, located at the European Laboratory for Particle Physics (CERN) in Geneva, Switzerland. First, I will describe a few simple examples of how particles can be detected, a concept which I have found is often puzzling to non-physicist. I will then give a brief history of CERN and the collaborations that made the construction of these devices possible, followed by an overview of the LHC machine itself. Finally, the CMS detector, its components, the computing infrastructure supporting it, and the performance resulting from this construction will be described.

3.1 Basic principles of particle detection

Many different kinds of particle detectors exist, and in fact several different types will be mentioned in the description of the CMS sub-detector systems. At their core however all of them rely on the fact that photons and electrically charged particles can ionize atoms, which both disturbs their environment slightly and locally reduces the electrical resistance of the material. A simple particle tracking detector can be constructed by producing a chamber in which for example a superheated liquid in a metastable phase is present.¹ When an ionizing particle passes through this medium the slight disturbance causes bubbles to form along its path, which can simply be photographed. This type of technology was used in early particle physics experiments, and an event captured using such a setup is shown in figure 3.1. The particles in the image have curved paths due to the fact that a magnetic field is applied, so that their charge and momentum can be deduced from the curvature their tracks. In modern detectors it is typically an electrical current induced by the ionization that is measured. This allows for layered detectors to be built (so that three-dimensional tracks can be reconstructed), and makes it so that the information can more easily be processed digitally. Even biological eyes and camera sensors rely on the fact that photons can knock electrons off their atoms, so that by applying a voltage a current can be measured.²

3.2 CERN

CERN, originally an acronym for “conseil européen pour la recherche nucléaire” or “European organisation for nuclear research”, was officially founded in 1954 with the signing of a convention by 12 european countries with the intention of stimulating international collaboration in physics research. This original name reflects the fact that research was at first focused on the study of atomic nuclei, with the foundations of particle physics and the standard model being established slightly later. Establishing this collaboration is reflective of the fact that fundamental research required increasing amounts of resources and personpower, as the relevant energy scales increased and technology would have to follow suit. In line with the founding principles of CERN, publications of its experiments are made publicly available, and data is made accessible to people outside the experimental collaborations once they have had time to analyse it. The greatest achievement of CERN as a collaboration is

¹A low-tech example of such a state sometimes occurs when distilled water is heated in a microwave. The liquid can be heated to slightly above its boiling point without actually boiling, after which it suddenly and violently boils once the container is disturbed or another substance is added.

²Charged particles in principle continuously radiate photons which ionize atoms in their path, so that the detection of photons and ionizing particles rely on the same principles.




FIGURE 3.1: An image of an event recorded inside a bubble chamber at the Lawrence Berkley National Laboratory (left) [58], and a drawing showing the particles the tracks correspond to (right) [59]. The present magnetic field is oriented such that it “points into the page”. The right-hand diagram of figure 3.2 shows a leading order Feynman diagram corresponding to the top half of the event shown here. An incoming high energy photon (which leaves no track) interacts with an atomic electron, which is ejected with a high momentum with a slight right-hand curve to its track. The produced electron and positron have lower momenta, as evidenced by their more sharply curving tracks, and spiral inward as they further lose energy due to bremsstrahlung. The forked shape on the bottom with no additional track could correspond to a similar process where the interaction takes place with a nucleus instead of with an atomic electron, so that no additional charged particle is emitted besides the produced pair.

probably the construction of the LHC, which has allowed for large amounts of data to be collected at an unprecedented centre of mass energy of $\sqrt{s} = 13 \text{ TeV}$.

3.3 The large Hadron Collider

The Large Hadron Collider is the largest hadron collider ever constructed, and is able to reach the highest collision energies ever achieved. It consists of a 27 km long circular tunnel, through which beams of protons or other particles can be made to travel inside high vacuum tubes, and sits around a hundred metres below ground level. Superconducting magnets are used to steer and collimate particle beams moving in opposite directions in separate vacuum tubes. At four points along the circle these beams are allowed to intersect, so that collisions can take place. It is in these locations, called interaction points (IP), that the detectors are installed. Around two of these interaction points so-called general purpose detectors are located, respectively named “CMS” and “ATLAS”. The CMS detector is the device used to collect the data used later on in this work, and was designed with a wide range of physics topics in mind. The main goal was the discovery of the Higgs boson, but several searches for new physics and tests of the standard model were also envisioned during its design phase. The ATLAS (A Toroidal LHC ApparatuS) detector was designed with the same goals in mind, but was intentionally based around a different layout. The groups associated with each of the detectors were also intentionally organized to have minimal communication between them. Given the fact that there are no other experiments around that could cross-check the results obtained by the LHC experiments, these two were intentionally made to serve this function to each other.³ The other two interaction

³And perhaps to stir up a sense of competition in the race towards the new discoveries the high LHC energies would allow.




FIGURE 3.2: Leading order Feynman diagrams corresponding to the processes of bremsstrahlung (left) and pair production by a photon (right).

points are occupied by the LHCb and ALICE detectors, which have distinctly different goals from the general purpose detectors. LHCb is an asymmetrical detector designed to study CP violation in the EW theory that could explain the matter and anti-matter asymmetry in the universe. ALICE is a detector optimized to study collisions between heavy ions, which the LHC is capable of accelerating as well.

While this will become clear later on, it should already be noted that the beams of particles are not significantly “consumed” in these collisions, so that they can be made to intersect repeatedly. The acceleration of the particles itself is done by radio-frequency (RF) cavities placed around the beam paths, each increasing the velocity of the particles moving through them slightly. This then explains the circular shape of the collider: Beams can be made to pass by the same RF cavities multiple times, to incrementally increase their kinetic energy. Particles moving with a higher kinetic energy take a stronger magnetic field to produce the same curvature of their trajectory however, so that the maximum energy is limited by the strength of the bending magnets and the radius of the tunnel. Charged particles also lose energy as their paths are made to bend, with the rate of loss depending on the sharpness of the turn. Beams of light particles like electrons are therefore additionally limited by how much kinetic energy can be added compared to how much is lost due to this effect. The energy lost through this phenomenon, called synchrotron radiation, steeply declines as the particle mass increases, so that it does not form a limiting factor for proton beams.

The LHC is not the first accelerator to occupy the tunnel it is located in. The tunnel itself was constructed for the installation of the Large Electron Positron collider (LEP), which operated between 1989 and 2000. Plans to construct a new hadron collider to take its place were approved in 1994, which became operational by 2008. Proton-proton colliders bring several advantages, and a few disadvantages with respect to electron-positron colliders. The large mass of protons means they lose much less energy to synchrotron radiation, so that much higher beam energies can be achieved. This makes proton-proton colliders excellent “discovery machines” as previously undetected particles are expected to have high masses, and therefore lurk in the high energy ranges. Protons are however composite particles, so that the choice of which of the fundamental particles takes part in the interaction, and what fraction of the kinetic energy of the proton it carries, is governed by probability density functions. Proton-proton collisions are also associated with the production of extra particles from lower energy collisions, leading to less “clean” events than those produced in lepton colliders. Such challenges can however be overcome using high resolution detectors and advanced reconstruction algorithms, so that CMS and ATLAS can be used for high precision measurements despite this additional complexity.

CERN has constructed linear and circular colliders with an increasing size and maximum beam energy over the years, where many of the old accelerators are now used to pre-accelerate protons before they are injected into the LHC. An overview of the accelerator chain, as well as the position of the four experiments installed along the LHC ring is shown in figure 3.3. After the pre-acceleration is complete the LHC accelerates the protons up to a kinetic energy of 6.5 TeV, which results in a centre of mass energy of 13 TeV in the collisions between two protons moving in opposite directions. At this point the protons are moving at what is practically the speed of light.




FIGURE 3.3: An overview of the LHC pre-accelerator chain [60]. Protons are produced by ionizing hydrogen gas after which they are accelerated to a kinetic energy of 50 MeV by the linear accelerator Linac 2, using RF cavities similar to the ones used in the LHC itself. The figure shows Linac 5, which is being constructed in a tunnel parallel with the one housing Linac 2, and will replace Linac 2 in future runs of the LHC. The beam is then sequentially accelerated further in the PSB, PS, and SPS accelerators, after which they are injected into the LHC with an energy of 450 GeV.

3.4 Beams, collisions, and luminosity

The proton beams in the LHC do not form a continuous stream, nor are protons being injected individually. Protons are injected and accelerated in bunches of around 10^{11} particles, and each of the two beams consist of maximally 2808 of these bunches. In this configuration bunches are separated by 25 ns or about 7.5 m. Inside the beam pipes protons travel though a high vacuum is produced, so that collisions only take place intentionally and beams are not slowed down by stray particles.

At the interaction points the beams moving in opposite directions are allowed to intersect so that head-on collisions can take place. Given the high energy and the large number of protons in each bunch multiple inelastic collisions take place in each bunch crossing. It is rare for more than one of these collisions to correspond to an interesting physics process, so that only the most promising one, called the primary vertex, is selected for physics analysis using a method explained later. Other vertices associated with inelastic collisions in the same bunch crossing are sometimes referred to as pileup vertices. The number of vertices per bunch crossing is called “pileup”, with average values of around 30 during the 2016-2018 data-taking period. Most of these interactions are uninteresting QCD processes, so that the most efficient use of the detector is to generate them in excess and select the interesting ones using the rapid trigger system described later.

The number of interactions (N), in a given amount of time (t), is dependent on the total cross section (σ_{total}) of whatever is defined as an interaction, and the so-called instantaneous luminosity ($\mathcal{L}_{\text{instant.}}$), which is dependent on various features of the two beams [31]. Two similar equations are of relevance here:

$$\frac{dN}{dt} = \mathcal{L}_{\text{instant.}} \cdot \sigma_{\text{total}} \qquad N_{t\bar{t}\gamma} = \mathcal{L}_{\text{int.}} \cdot \sigma_{t\bar{t}\gamma} \qquad (3.1)$$

The first links the total frequency of interactions to the luminosity at a given instant, and is proportional to the average pileup at that moment (the actual value being of course randomly distributed). The second involves the time integral of the instantaneous luminosity over the total data taking period, and is used to estimate the number of e.g. $t\bar{t}\gamma$ events to be expected in the total data set. Inversely, this means that by counting the number of $t\bar{t}\gamma$ events in the data set and dividing it by the integrated luminosity a measurement of the $t\bar{t}\gamma$ cross section ($\sigma_{t\bar{t}\gamma}$) is obtained. The instantaneous luminosity itself is given by

$$\mathcal{L}_{\text{instant.}} = f_{\text{coll}} \frac{n_1 n_2}{4\pi\sigma_x\sigma_y} \mathcal{F}, \quad (3.2)$$

where f_{coll} is the number of bunches per beam times the frequency with which they pass by the interaction point, and n_1 and n_2 are the number of protons in the two beams, respectively. σ_x and σ_y are the quadratic means of the beam sizes in the horizontal and vertical directions, where the bunches are assumed to have the exact same shape and to have a constant transverse width along the longitudinal direction. Deviations from this idealized description are corrected using the factor $\mathcal{F}$, which has a value of close to unity. This value takes into account things like the crossing angle of the beams, the finite bunch length, and the mutual focusing of the two beams as they cross.

3.5 The Compact Muon Solenoid

3.5.1 Wish list for a general purpose detector

The CMS detector was constructed around one of the interaction points of the LHC, on the opposite side of the circular tunnel from the ATLAS detector. CMS and ATLAS were designed to be general purpose machines, but with the primary goal of either discovering the Higgs boson, or excluding its existence over a wide range of masses. The centre of mass energy, statistics, and object resolution these goals require imply a certain set of features any suitable detector needs:

- The large amount of useful data required implies a large integrated luminosity is needed, which can only be achieved in a reasonable amount of time through a high event rate made possible by a trigger system which rapidly identifies events worth storing and processing further.
- A high granularity is needed to deal with the high amounts of pileup associated with these high luminosities, since tracks from different vertices need to be distinguished. Additionally, detector systems with good time resolution and low dead time are needed in order to distinguish events from subsequent bunch crossing. This granularity is of course also important with regards to the angular resolution of the individual measured objects themselves.
- The high luminosities mean that sub-detectors need to be able to resist large amounts of radiation in order to perform consistently over long run times. This is especially true for sub-detectors which are placed as close as possible to the beamline.
- In order to fully reconstruct events and be able to extrapolate the missing transverse momentum vector from measurable objects a large geometric coverage is needed. This means a detector which extends as far as possible into the forward regions (high pseudorapidity value), with as few gaps as possible in the sub-detectors is required.
- Muons need to be reconstructed with efficiencies of close to unity, with high momentum resolution and the associated correct assignment of the muon charge. The same goes for electrons, and is achieved in part by producing a sufficiently powerful magnetic field to curve the particle tracks, but muons pose the additional challenge of producing less ionization as they pass through matter.
- The tracking system needs to perform well enough to be able to distinguish particles which decay after travelling a small but non-negligible distance inside the detector. This is important for the identification of b quarks and tau leptons, as well as for potential undiscovered particles which might show similar behaviour. The performance of the tracking system is also directly reflected in the reconstructed particle momenta, which are calculated based on the curvature of the tracks.




FIGURE 3.4: Diagram illustrating the CMS coordinate system centred around the interaction point [61]. The diagram also gives a rough idea of the position and orientation of the CMS detector.

3.5.2 The CMS coordinate system

In principle two coordinate systems are used in the context of the CMS detector, but both are centred around the interaction point. The first is a standard cartesian coordinate system, where the z-axis is defined perpendicular to the beamline, pointing towards the Jura mountains, or in other words pointing in an anticlockwise direction when looking at the LHC from above. The x-axis lies in the same plane as the LHC, pointing towards the centre of the ring. The y-axis lies perpendicular to the first two, such that a right-handed system is obtained, or in other word the y-axis “points up”. A diagrammatic representation of both coordinate systems is shown in figure 3.4.

The second coordinate system is reminiscent of cylindrical coordinates, and is the more useful of the two from a physics analysis perspective. In this second system $r = \sqrt{x^2 + y^2}$ is the distance to the z-axis and therefore the beamline near the interaction point. Within this same xy-plane ϕ is the azimuthal angle with respect to x-axis, with positive values in the positive direction of the y-axis.⁴ Finally, the information regarding the z-direction is parametrized using the pseudorapidity variable defined as $\eta = -\ln(\tan(\theta/2))$, where $\theta = \arctan(y/z)$ is the polar angle with respect to the beamline. The main reason for using the pseudorapidity is the fact that it is both Lorentz invariant and additive, making it mathematically convenient when switching between relativistic coordinate systems. It is zero for points in the xy-plane, and approaches infinity as a point gets closer to the beamline. Measurements and positions of sub-detectors and other detector features are also often expressed in terms of this second coordinate system, as these are much more relevant than a length in metres when it comes to e.g. the coverage of the detector.

Since it is the partons carrying a random fraction of the proton energy that undergo interactions the total momentum in the z direction is not zero. Therefore the conservation of momentum, and the assumption that its total is zero, can only be relied on in the xy-plane. Because of this often only the transverse component of e.g. the momentum of a particle (p_T) is used. These coordinates are also used when describing the angular distances between objects, for which variables like the angular distance $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ are defined.

⁴The often quoted definition of $\phi = \arctan(x/y)$ is insufficient, as it gives the same values for points in opposite quadrants. Instead ϕ corresponds to the “2-argument arctangent” used in programming, which unambiguously returns values in the range $[-\pi, \pi]$.




FIGURE 3.5: A drawing showing the main components of the CMS detector, along with their main features [62].

3.5.3 Detector components

A drawing showing the various main components and sub-detectors of the CMS detector is shown in figure 3.5. What follows now is a brief description of the subsystems starting from the centre of the detector and working outwards, in accordance with the order in which a particle that makes it through all of the layers would encounter them. The name CMS refers to the 3 most prominent design features of the detector: Its tracking and calorimetry sub-detectors sit almost entirely inside the solenoidal magnet, making it significantly smaller than the ATLAS detector, explaining the “Compact” and “Solenoid” parts of the name. As was already mentioned in the wish list for a detector of this kind, special focus is put on the detection and reconstruction of muons, hence the word “Muon” in the name.

3.5.4 Trackers

Situated mere centimetres from the beamline the tracking system is one of the most important components when it comes to disentangling the tracks from the large number of pileup events. The list of features it requires in order to perform its task well is long. Given that particles haven’t “spread out” yet at this distance means a very high resolution is needed in order for their tracks to be distinguishable. Because of this the inner part of the tracker consists of a silicon pixel layout, similar to a digital camera. For the outer part a silicon strip structure is used, which reduces both costs and the number of readout channels required. A setup like this also limits the amount of material in this region, so that effects like bremsstrahlung and photon conversion are kept to a minimum. Its position close to the interaction point also means that the tracker is subjected to the highest density of radiation damage. The pixel and strip trackers were designed to withstand multiple years of the radiation levels they are subject to, and are kept at -10°C during operation to minimize the effects of the damage they do sustain on the readout signal. Together the silicon pixel and strip detectors fill the space between $r = 4.4\text{ cm}$ and $r < 1.2\text{ m}$. Like most sub-detectors the trackers consist of a cylindrical barrel region, and 2 circular endcaps, so that an area up to $|\eta| = 2.5$ is covered. In the barrel region the fact that the electron drift occurs perpendicular to the magnetic field spreads the energy of a hit out over two




FIGURE 3.6: Image showing the layout of the CMS ECAL and pre-shower, their η ranges, and the gap between the barrel and endcap sections [63].

to three neighbouring pixels, which can be exploited to improve the spatial resolution. In general a resolution between 10 and 20 μm is achieved.

3.5.5 ECAL

The task of the Electromagnetic Calorimeter (ECAL) is to measure the energy of photons and electrons/positrons with a high precision, and with a granularity that is high enough for the signatures of individual particles to be distinguishable. Additionally, as for all CMS sub-detectors, a fast response time and good radiation resistance is required. One low-tech feature a calorimeter needs is the presence of high-density materials, so that the electrons and photons passing through it are induced to form showers thereby depositing their energy to be measured. The energy released of course needs to be detected, but due to the relatively large area to be covered this cannot be done using microelectronics as was used in the trackers. Instead, lead tungstate ($PbWO_4$) crystals are used, which have a high density, but also scintillate when charged particles move through them. The light emitted is then amplified and converted into an electrical signal using photo-multipliers. This material has a short scintillation decay time so that signal overlap in time is not an issue, and is inherently very robust against radiation damage.

The ECAL is made up of 61200 lead tungstate crystals in the barrel regions, and 7324 in the endcaps, so that a pseudorapidity range of up to $|\eta| = 3$ is achieved. An overview of this configuration is shown in figure 3.6. As shown on this drawing there is a small gap between the barrel and endcaps of the ECAL, which corresponds to a region where support structures sit and readout cables pass through. In the r direction the ECAL fills the space between 1.24 m and 1.75 m. Also drawn is the so-called pre-shower detector, which sits in front of the endcap sections of the ECAL. In the endcap regions the angle between two photons coming from a π^0 decay is likely to be small, so that the ECAL would pick up just a single signal. In order to improve this resolution the pre-shower consists of two lead plates followed by a set of silicon strips similar to those used in the tracker. The 2 mm silicon strips offer a much better resolution than the 3 cm wide lead tungstate crystals, so that the pre-shower silicon strips can help distinguish the two photons, which the lead layers induce into showering. This addition was in part made with the Higgs discovery in mind, as its decay into photon pairs was expected to offer one of the cleanest signatures for Higgs boson decays.

The signal is read out in groups of 5×5 crystals called supermodules. Already at this level signal noise is reduced using zero suppression, which essentially means that only energy spikes above a certain (low) threshold are considered. The energy resolution for photons is highly dependent on the pseudorapidity and the time of conversion. For unconverted photons or photons that convert late values of around 1% are achieved, while in the endcaps the resolution becomes around 2.5%. Photons that convert early on have a worse energy resolution, with values between 1.3% and 2.5% in the barrel region, and up to 4% in the endcaps.

3.5.6 HCAL

The hadronic Calorimeter or HCAL sits around the ECAL, and is used to measure the position and energy of hadronic showers. This is the last main layer of the detector that is placed inside the solenoid magnet volume, which will be described in the next section. It consists of a barrel section (HB), endcap section (HE), and a forward (HF) section which makes it so that the pseudorapidity range extends up to $|\eta| = 5$. This additional coverage in the eta direction is required because the information from the HCAL is also used to calculate the total missing transverse energy, and any undetected jets would affect this value. Because of this the gap between the barrel and endcap regions is also oriented in such a way that, while the performance in this area will be less than ideal, no fully dead areas are present. Because of the limited space inside the solenoid the HB is insufficiently wide in the r direction to fully stop all hadrons. To compensate for this an additional HCAL layer (outer HCAL or HO) sits just outside the magnet. The HB, HE, and HF are made up of alternating layers of brass plates and plastic scintillators. Brass is used here because of its high density, and therefore high stopping power, while being non-magnetic. The forward part of the HCAL sits very close to the beamline, so that it is subject to higher levels of radiation and needs to rely on a different design. There steel plates are used instead, while the actual detection relies on quartz fibres.

The obtained energy resolution, measured using test beams, is highly dependent on the energy of the jet and the sub-detector it is measured by. For the barrel and endcap sections it can be written as:

$$\frac{\sigma_E}{E} = 90\% \frac{\text{GeV}^{1/2}}{\sqrt{E}} \oplus 4.5\% \quad (3.3)$$

While in the HF section this becomes:

$$\frac{\sigma_E}{E} = 172\% \frac{\text{GeV}^{1/2}}{\sqrt{E}} \oplus 9\% \quad (3.4)$$

Where the symbol $\oplus$ is used to imply the quadratic sum is taken, as the two components of the uncertainty are uncorrelated.

3.5.7 Solenoid magnet

The superconducting solenoid magnet inside of which the previously described sub-detectors are installed is used to generate the 3.8 T magnetic field needed to curve the tracks of highly energetic charged particles moving through the detector volume. Generating a magnetic field this strong is achieved by running a current of around 20 kA through four layers of wound niobium-titanium (NbTi) wires, cooled down to just 4.6 K. At this temperature the material is superconductive, so that it has essentially no resistance. The solenoidal shape of the magnet means that the magnetic field inside its volume is nearly homogeneous. Because the magnet is a relatively dense and wide object it additionally serves the function of partially stopping hadrons from reaching the muon systems. Outside the solenoid, in between the layers of the muon system systems described in the next section, sit large steel return yokes, which help control the magnetic field outside the magnet volume. Within this region the field lines run in the opposite direction with respect to inside the magnet.

3.5.8 Muon system

Because muons lose far less energy to bremsstrahlung than electrons, they are for the most part the only charged particles that make it outside the solenoid volume. While the inner tracker system is efficient at reconstructing their tracks, this feature allows them to be distinguished from other particles by placing additional dedicated muon tracking systems outside the magnet volume. These sub-detectors also allow for more accurate reconstruction of very high energy muon tracks. At this distance from the interaction point the surface area that needs to be covered is very large, so that inexpensive detectors need to be used. Luckily this distance also means particles are already much more spread-out, so that overlapping tracks are much less of an issue. The muon system is once again made up of a barrel region and two endcaps, each of which is made up of layers interleaved with the steel plates of the return yoke. In the barrel region these layers consist of drift tubes, which are essentially small volumes of gas with a positively charged wire running through them. Electrons freed through ionization by charged particles passing through collect on these wires and can be read

out as an electrical signal. This system results in a spatial resolution of around $170\text{ }\mu\text{m}$ in both directions, as measured using test beams. In the two endcaps cathode strip chambers are used, which are similar in function to the drift tubes but use both negatively and positively charged wires running perpendicular to each other so that two coordinates are obtained in each detection. The performance of these chambers far exceeds the design goal, with a spatial resolution of around $80\text{ }\mu\text{m}$, while the goal was set at $150\text{ }\mu\text{m}$. The η range of the muon system extends up to 2.4, matching that of the ECAL.

3.6 Physics object and event reconstruction

3.6.1 Triggering and data acquisition

In order to deal with the 40 MHz collision rate the high instantaneous luminosity of the LHC provides a two-tiered trigger system is used. This system rapidly identifies events with potentially interesting tracks, so that only these need to be passed on to more detailed reconstruction steps. Events that make it through the trigger step are passed on to the so-called data acquisition system (DAQ). The first trigger step “Level-1” (L1) trigger consists of a set of field-programmable gate arrays which can process information faster than regular microcontrollers can. These L1 triggers look at calorimeter and muon system information separately, and reduce the event rate to a more manageable 100 kHz. Full resolution event information is in the meantime held in a buffer for $3.2\text{ }\mu\text{s}$, before either being thrown away or passed on to the subsequent High Level Trigger (HLT) system. The HLT in turn uses the full detector information to perform a simplified version of the offline reconstruction that will be performed if the event makes it through this step. The more complete information available to the HLT system allows it to impose requirements on global variables like the total missing transverse energy and the distance of closest approach between leptons. Requirements on the presence of electrons and muons with transverse momentum values above some threshold can once again be imposed, this time benefiting from the much higher quality objects obtained from the online reconstruction procedure. After the HLT the event rate has dropped down to around 1 kHz, which corresponds to a manageable data stream of 1 GB/s which is permanently stored for later offline reconstruction. The HLT trigger definitions need to be tuned so that a maximal amount of interesting physics data of various kinds is stored, while not overwhelming the write-out system.

3.6.2 Event reconstruction

Before further analysis can take place the information recorded by the sub-detectors needs to be combined into physics objects with nearly direct correspondence to the particles produced in the collisions. In CMS this is performed using the Particle Flow algorithm [64], which reconstructs particles and jets, but also calculates associated values like the total missing transverse momentum indicative of neutrinos. It does this by combining information from the various sub-detectors (including the reconstructed tracks) and outputs collections of particles, which are referred to as PF candidates. This not only improves the identification efficiency for these objects with respect to using only sub-detector information individually, but also improves the spatial and momentum resolution of the reconstructed particles. The algorithm is highly complex, and its detailed functioning is of lesser relevance to the analyser. Instead, I will describe the specifics of which information is used to reconstruct the most important particle categories, and provide context to the analysis-specific object identification criteria that will be explained later on. The object categories this algorithm produces are not exclusive, meaning that e.g. a certain set of hits can be reconstructed as both an electron and a muon. The main goal for the algorithm is to reconstruct objects with a high efficiency, while removing mis-reconstructed objects or overlap needs to be performed later by applying selection criteria. To this end most low-level information of the sub-detectors is kept and assigned to the reconstructed objects, so that the analyser is not limited to predetermined identification criteria.

A first step consists of the reconstruction of charged particle tracks, the curvature of which is used to estimate the particle momentum and charge. Electrons and charged hadrons additionally deposit their energy into the ECAL and HCAL systems, so that their energy can be measured. Photons and neutral hadrons also deposit their energy into the calorimeters, but have no signal in the trackers. As far as the ECAL is concerned electrons and photons are the same kind of object, so that the tracking information is highly important to distinguish them. Muons, being charged particles, typically have

a signal in the trackers, but are not stopped by the calorimetry systems. They make it to the outer parts of the detector where they produce additional hits in muon systems. Neutrinos have no signal in any of the sub-detectors, but their presence in an event can nevertheless be detected. The total transverse momentum vector of reconstructed objects in an event without any undetected particles can be assumed to be close to zero. In any event the inverse of this vector is therefore an estimate of the vector sum of the neutrinos. The particle candidates are reconstructed on a per-type basis, starting with the muons, followed by the electrons and photons, and finally the hadrons. The order is of importance here since after each step the sub-detector information assigned to the reconstructed particles is removed before moving forward.

3.6.3 Track reconstruction

The first step in reconstructing an event consists of identifying the tracks of individual particles, which is non-trivial given the large track density close to the IP. A track finding algorithm called the Combinatorial Track Finder (CTF) is therefore applied to the hit patterns measured in the pixel and strip trackers, which is optimized to achieve a high efficiency while limiting the number of fake tracks. Figure 3.7 shows the reconstructed tracks in a typical 2016 bunch crossing. By extrapolating to where tracks intersect the z-axis interaction vertices can then be identified, where the primary vertex (PV) is defined as the vertex with the largest value of the summed physics-object p_T^2 . Other vertices are not considered for physics analysis.

The hits that form the input to the tracking algorithm are reconstructed with an efficiency of around 99.5%, where each hit has a resolution between 10 μm and 50 μm depending on the tracker (pixel vs strip tracker) and the angle of incidence. Higher instantaneous luminosities do of course lead to slightly lower efficiencies due to the higher detector occupancy, but good performance is achieved for even the highest pileup events. Aside from the resolution the spatial accuracy of the hits is however also limited by the positioning of the sensors. A small discrepancy exists between where the sensors are assumed to be positioned, and where they really are positioned, which is only relevant because the individual pixels/strips are so small. Both tracks from collisions and from cosmic rays can be used to calibrate this alignment, after which the alignment uncertainty is reduced to 1 μm and 10 μm for the pixel and strip trackers, respectively.

The CTF uses an iterative procedure where each iteration starts with the identification of so-called track seeds. These track seeds define an initial track, which is then traced outward through the layers of the tracker, each time adding compatible hits to the track before adjusting the trajectory based on this new information. This iterative system of repeatedly predicting and updating the track is known in a more general context as a Kalman filter (KF) [65], which is an estimation method designed to deal with “noisy” data with a time-like dependence. Once the full track is reconstructed a check of how well it matches its corresponding hits is performed using a goodness-of-fit calculation, so that insufficiently compatible tracks can be rejected. If a track is accepted, its corresponding hits are removed before the next iteration is started. The iteration starts with the “most promising” tracks, namely those with high p_T values originating close to the IP. In each subsequent iteration the goodness-of-fit requirement is made less stringent so that a high reconstruction efficiency is also obtained for lower-quality tracks. In the final two iterations the muon systems are taken into account as well. For electrons the KF tracks are used as input to a so-called Gaussian sum filter (GSF) algorithm, which is too computationally expensive to run on the full set of tracker hits, from which more reliable track parameters are obtained. While the KF approach models the bremsstrahlung energy loss using a single Gaussian distribution, the GSF algorithm uses a mixture of several Gaussians, thereby better approximating the electron energy loss theoretically described by the Bethe-Bloch formula [66].

The reconstruction efficiency obtained this way for electrons and charged pions is highly dependent on their pseudorapidity, with values ranging between 80 and 99%, and fake rates between 5 and 15% [67]. For muons the efficiency is better, with values of above 99% over the entire pseudorapidity range. It should be noted that the values mentioned here assume a detector that is in pristine condition, with no dead zones, which in reality appear due to radiation damage. The reconstruction efficiency for vertices is also very high, with an efficiency of 98% for vertices with just 2 tracks, and practically 100% when more tracks are present.




FIGURE 3.7: Image showing the reconstructed tracks in a single bunch crossing during the 2016 data taking period of CMS [68]. Also shown are the interaction vertices associated with these tracks.

3.6.4 Muon reconstruction

Muons are characterized by a series of hits in the trackers, compatible with a track in the muon systems. The PF algorithm manages to reconstruct muons with a high efficiency by using both this combined information, and the information from the muon systems separately. Three types of muon objects can be defined:

- Tracker muons are found by extrapolating an inner track to the muon systems, where a matching segment means that the track corresponds to a muon candidate.
- Global muons are constructed by starting at a track in the muon chambers, and trying to find a compatible inner track. A fit to the combined information then defines the global muon and its measured properties.
- Standalone muons are found using only information from the muon systems, where a procedure similar to the normal tracking algorithm can be used. At this distance from the IP the density of tracks is much lower, so that the relatively low resolution is not an issue.

These types have a lot of overlap in terms of the muons they manage to reconstruct, with the global muon approach resulting in the highest efficiencies over most of the p_T range. Since for low energy muons the muon systems are of lesser importance the tracker muon approach does outperform the global one for $p_T(\mu) < 10 \text{ GeV}$.

The collection of isolated muons (i.e. muons not produced inside jets) is finally obtained by applying further requirements on the reconstructed global muons. This can simply be done by calculating the transverse energy and momentum of all objects inside a cone of $\Delta R < 0.3$ around the muon. If these parameters exceed 10% of those of the muon itself it can be assumed to be part of a jet, and would therefore be rejected. As mentioned earlier once isolated muons have been reconstructed their associated sub-detector information is removed before moving on to the electron and photon reconstruction.

3.6.5 Calorimeter cluster identification

Aside from muons all final state particles deposit most of their energy into the ECAL and HCAL sub-detectors.⁵ This sub-detector information is especially important for neutral particles, which produce no tracks, and whose energy and position can therefore only be reconstructed based on calorimeter

⁵Not mentioning neutrinos here to avoid confusion, which are not detected by any of the sub-detectors.

signals. Much like the tracker hits need to be combined into actual tracks before being used to reconstruct particles, these energy deposits need to be merged into distinct clusters. Several technicalities need to be taken into account when doing this: Energy deposits can fully or slightly overlap, so that care needs to be taken when figuring out which cluster energy deposits are assigned to. In the case of nearly fully overlapping showers distinguishing them is of course impossible, and they will unavoidably be reconstructed as a single object. Furthermore the calorimetry systems are designed to collect as much of the deposited energy as possible, while also suppressing signals originating from for example thermal noise. Since such a system can never be 100% efficient, calibrations need to be performed so that the reconstructed energies match the total energy carried by the detected particle.

Clustering is performed using slightly different algorithms for the barrel and endcap regions of each sub-detector, but are collectively referred to as the PF clustering algorithm. First, a cluster seed is identified, corresponding to a group of detector cells with an energy above a certain threshold. Neighbouring cells are then added onto the seed to form topological clusters, to which a maximum likelihood fit is performed to extract the energy of the cluster.

3.6.6 Electron and photon reconstruction

The reconstruction of both electrons and photons starts off with the merging of ECAL clusters into superclusters. These superclusters are groups of ECAL clusters merged together with similar η values, but slightly more spread out ϕ values. As electrons pass through the detector their tracks are curved in the xy -plane due to the presence of the strong magnetic field, all the while emitting bremsstrahlung photons which in turn can convert into electron-positron pairs. This conversion also explains why photons and electrons have practically identical ECAL signatures: Both deposit their energy in a cascade of photon conversion and bremsstrahlung photon radiation, so that it matters little which of the two types of particles first entered the calorimeter volume. The distinction then needs to be made based on the presence of inner tracker hits compatible with the ECAL supercluster.

The reconstruction uses an ECAL cluster as a seed and traces a track back inwards under the hypothesis of having been produced by an electron. If a compatible track is found the object is labelled as an electron, while if no such track is found it is assumed to be a photon. If however more than two matching tracks are found the electron hypothesis is also rejected, as a large number of matching tracks is indicative of a non-isolated electron or a jet. The energy deposited in the HCAL around the ECAL cluster is also checked for both the electron and photon hypothesis, which is required to be below 10% of the energy of the ECAL cluster. For photons the energy and direction can only be extracted from the ECAL information, to which corrections for measurement inefficiencies are applied. In the case of an electron information from the ECAL is combined with track information to obtain the energy, while the directional information is only based on the track.

Due to the presence of bremsstrahlung the reconstruction of the electron charge is less trivial than for muons. Photons from bremsstrahlung can convert into electron pairs, which lead to hits in the tracker that can be hard to disentangle from the main electron track, especially if the conversion takes place early on [69, 70]. Three different methods are used to reconstruct the charge, the information from which can then be combined to perform the final assignment. The most obvious approach consists of looking at the curvature of the GSF track of the electron. This method alone is quite sensitive to the effects of conversions, with charge misidentification rates of up to 10% for high pseudorapidity values. This kinematic dependence is a result of the larger amount of material the electron has to traverse as its pseudorapidity increases, in addition to the track being less bent due to the electron moving more parallel to the magnetic field. As a secondary check the curvature of the KF track associated with the GSF track (sharing a hit in the inner layer of the pixel detector) can also be checked. Finally, the charge can be obtained by comparing the position of the supercluster associated with the electron and its innermost hit in the pixel detector. The sign of the ϕ angle between these two positions once again implies a track curving in a direction perpendicular to the magnetic field, and therefore the probable charge of the electron. The electron is then assigned the charge at least two of these methods agree on, an approach referred to as the “majority method”. In cases where the electron charge is highly important a charge consistency requirement can be applied, in which case electrons are only used if all three methods obtain the same sign for the electron charge. While this requirement significantly improves the charge misidentification rate, as illustrated in figure 3.8, it does reduce the reconstruction




FIGURE 3.8: Probability of correctly assigning the charge for electrons using the majority method (left), and when applying the charge consistency requirement (selective method, right) as measured in both real and simulated $Z \rightarrow ee$ events with some non-stringent object identification requirements applied [69].

efficiency for the electrons themselves from 100% to between 93% and 97%, so that it is not applied by default. For example in the $t\bar{t}\gamma$ analysis described in chapter 5 it is therefore not applied, as there is no reason to incur this loss in efficiency there. Meanwhile for the measurement of the production of four top quarks this requirement is highly important, as one of the signal regions requires the presence of two leptons with the same charge. There the misidentification of electron charges causes background events with two oppositely charged leptons, mainly from $t\bar{t}$ production, to enter this signal region. A method used to estimate this background is described in appendix A.

3.6.7 Hadronic jet reconstruction and heavy flavour identification

Electromagnetic and hadronic showers are similar from a structural perspective, but the term jet is specifically used to refer to the signatures produced by the hadronization of quarks and gluons. In the hadronization of these particles and their interaction with the detector material all kinds of particles can be produced including photons and leptons, but for the most part jets are made up of hadrons, both neutral and charged. The photons and leptons produced within these jets are typically not considered as individual objects at the analysis level, and they are not what is being referred to when a final state with a certain lepton or photon multiplicity is defined. Instead, jets are seen as objects themselves, which are characterized by global variables like their width and charged particle content.

As collective objects jets are reconstructed using the anti- k_T clustering algorithm with distance parameter $R = 0.4$ [71]. Both HCAL and ECAL information is used here, since jets typically already deposit a small percentage of their energy in the ECAL volume. A basic filter on the jet quality (or in other words the fraction of fake jets) is already applied by requiring jets to consist of at least two objects, with at least one being a charged PF candidate. A baseline requirement is also set on the fraction of the jet energy coming from charged particles.

Jets produced by b quarks can to a good degree be distinguished from jets produced by other hadrons and gluons due to their relatively long lifetime and high mass. Because of this lifetime they can travel up to 10 mm before decaying and forming a jet. The tracks corresponding to these jets therefore intersect in a vertex which is displaced with respect to the primary vertex, called the “secondary vertex”. A deep neural network can be used to turn information about the tracks, the secondary vertex, and object kinematics into a single output variable which parametrizes how b jet-like a jet is. One such neural network used in CMS is called DeepCSV, and will later on be used to classify events based on their reconstructed b quark content. This is in practice done by choosing an output value of the neural network, called a “working point”, above which a jet is assumed to have originated from

a b quark decay. In actual measurements there is of course no certainty whether this label is correct, and each working point is associated with both an efficiency and a misidentification rate.

3.6.8 Global variables

Once the physics objects in an event are reconstructed variables pertaining to the entire event can be calculated, the two most commonly used of which are the missing transverse energy E_T^{miss} and the scalar sum of the jet transverse momenta H_T :

- The missing transverse energy E_T^{miss} is a variable which carries information about the neutrinos in the event. In the case of an event in which only one neutrino was present it is in fact an approximation of its transverse momentum vector. The E_T^{miss} is calculated as the negative vector sum of the transverse momentum of all PF candidates, including those not classified as physics objects.⁶ Overlap between leptons and jets is however first removed, and jet energy corrections are already applied at this point [72].
- H_T is defined as the scalar sum of the p_T of all jets with $p_T > 10 \text{ GeV}$ and $|\eta| < 3$ associated with the primary vertex [73]. This variable is used in the trigger system, where the fact that it is uncorrelated with the lepton activity in the event proves to be useful for performance measurements. In section 5.5 of the analysis presented later a data set produced using H_T -based triggers is used to calculate the efficiency of the sets of lepton triggers used in the measurement.

3.7 Monte Carlo event simulation

Measurements rely on the comparison between observations and predictions, both for inclusive cross sections and for differential results. These predictions come in the form of simulated events, typically produced for each physics process separately, and sometimes even divided up into for example separate decay channels. Contrary to the cross section calculation example given earlier most events produced in proton-proton collisions are dominated by QCD interactions, so that factorization of perturbative QCD calculations is required [74]. Simulations consist of three main steps at the software level:

- A matrix element (ME) level calculation is performed to leading (LO), next-to-leading (NLO), or even higher order with respect to the Feynman diagrams. The process to be simulated needs to be defined in an unambiguous way, and various settings with direct or indirect equivalents in the standard model need to be specified. Such examples include the mass and decay width of the top quark, the beam energies, and the factorization and normalization scales. Since it is the partons which form the true initial state particles in these processes parton distribution functions (PDF) also need to be defined when simulating proton-proton collisions. One is also not limited to SM predictions, as events can also be generated under for example alternative hypotheses for the top quark mass. The entire model including its particles and their couplings can in fact be replaced, which will be the case when exploring effective field theory models.
- At the ME level it is not possible to include all possible decays and radiation effects, and this information is therefore added in a second, less precise step. Decays and hadronizations, including the radiative cascades initiated by quarks and gluons known as parton showers, are simulated in a probabilistic way so that final state objects are obtained. Initial and final state radiation is added to events, and in cases where some radiation was already added in the previous step overlap is removed through a matching procedure. Besides the hard interaction other soft QCD parton-parton interactions tend to occur in proton-proton collisions. These interactions are simulated using phenomenological models tuned to match data. This information is also added in this step, and is referred to as the underlying event. At the end of the showering step the simulation also needs to ensure that only colour neutral objects are left, and hence quark with colours that can form colour neutral objects need to be connected. Various models exist for this so-called colour reconnection procedure, but the phenomenon is in general not well understood.
- Finally, the interaction between the generated particles and the detector is simulated. This simulation uses a description that is as accurate as possible with respect to responses, material

⁶Where a physics object is any of the PF particle types described earlier.

position and density, and potential imperfections in the sub-detectors. Perfectly matching this simulation to reality is however impossible, which is why the performance is typically adjusted later by applying correction factors to simulated events. A full simulation of the detector also allows for the study of the detector response itself. This means that this response can be parametrized, which can then be inverted to remove detector effects from measured distributions, as will be explained further in the section on differential cross section measurements.

Figure 3.9 shows a visual illustration of the event simulation steps, not including the interactions with the detector material. For each step multiple software packages are available, but at their core all of them make use of Monte Carlo methods [75]. The specific tools used for the production of samples used in the $t\bar{t}\gamma$ analysis will be described later in section 5.2.2, where the samples themselves are also listed.

The ME level calculation relies on the fact that cross section calculations for collisions between protons can essentially be factorized between terms which describe the probability of a certain parton taking part in the interaction (the PDFs), and the ME terms corresponding to each of the possible parton interactions. This does however lead to a dependence on the chosen factorization scale of the obtained result. Similarly, cross section calculations and event simulations have a dependence on the chosen renormalization scale. This scale is introduced in the renormalization procedure used to remove divergent terms in cross section calculations, at the cost of making the coupling strengths dependent on this scale. This dependence is typically referred to as the “running” of a coupling constant, and is the reason measurements of these couplings are always to be quoted along with the scale at which they were determined.

The full simulation chain including the detector simulation step leads to collections of reconstructed particles very closely approximating the reconstructed particles obtained in the real experiment. The information produced in the intermediate steps of the simulation is however also of major importance to most analyses, as it can for example be used to check whether particles have been correctly identified, or to trace back which particles a final state particle was produced by (sometimes called the parentage or parentage chain of the particle). The collection of particles corresponding to the ME level calculation can be used to investigate which part of the phase space is covered by a specific sample. The particle collection corresponding to the situation just before the parton shower simulation is performed is typically referred to as the parton level information. Another level corresponds to the particle collection that is present just before the detector simulation step is performed, consisting of stable final state particles, and is referred to as the particle level (PL). At this point the partons have formed showers and have hadronized. For physically well-defined comparisons, the particle-level objects need to be clustered into jets, for which the same clustering algorithm as is used for the reconstructed objects is applied. This particle level also corresponds to what is referred to as the generator level, with the term “particle level” typically implying a set of requirements has been imposed on the full generator level particle collection to obtain a nearly one-to-one correspondence with the collection of reconstructed particles.

When using data each event contributes the same amount, or is in other words assigned a weight of 1. In simulated events this is not the case: The aim is usually not to generate the same number of events as expected in data. Whenever possible much more events are generated for all processes, which are then assigned weights in order for the total yield to correspond to the predicted yield. If the number of events is indeed larger than the expected yield each event is assigned a weight w_i smaller than 1, making the total uncertainty in N events smaller than the $\sqrt{N}$ value obtained if weights were 1. In general the total statistical uncertainty, or the statistical uncertainty in a certain bin, can be written as $\sigma = \sqrt{\sum w^2}/N$, so that for example 100 events with a weight of 1 have a total statistical uncertainty of 10, while 400 events with a weight of 0.25 (and hence the same total yield) have an uncertainty of 5. While this does not hold when weights vary event by event, as a rule of thumb an increase in the number events by a factor of x reduces the statistical uncertainty by a factor of $1/\sqrt{x}$. In simulations of higher than leading order events can take on negative weights, which necessarily arise when taking into account interference effects between different contributions to the total cross section. Weights will be used in clever ways later on to adjust samples to alternative physics hypotheses without generating new events. Many corrections will also come in the form of weights to be multiplied with the simulation weight, once again to morph distributions without having to reproduce entire samples.




FIGURE 3.9: Illustration of the various steps involved in the simulation of proton-proton collisions, in this case for the production of a top quark pair together with a Higgs boson and an additional gluon [76]. The image shows two incoming protons, drawn in green, from which two partons interact and initial state radiation is produced as drawn in blue. Red lines show the produced particles, the non-electroweak decays, and particles produced as a result of final state radiation from these decay products. The green shapes towards the outside of the figure correspond to the hadronization and hadron production steps, along with the decays of the unstable hadrons. The bottom half of the image shows the underlying event resulting from the proton remnants not involved in the primary interaction, which similarly leads to hadron production and decay. Finally, as shown in yellow, the electroweak decay products are produced and photons from final state radiation are added.

3.8 Maximum likelihood fitting

In a cross section measurement the final main step of the analysis typically consists of performing a fit in which the signal strength is extracted. Simply put this signal strength is the factor by which the predicted signal needs to be scaled in order for it to optimally agree with the observed data, taking into account the various backgrounds and uncertainties. I will already introduce the general principles here, so that the terminology is already defined when describing fitting procedures in the wider context of effective field theory analyses, and when discussing the uncertainties relevant to the $t\bar{t}\gamma$ measurement described later.

We can start with a situation in which two histograms with N bins are filled with respectively observed and predicted events, which have passed the full event selection. Each bin i will correspond to a certain number of observed events o_i and a total number of predicted events given by $rs_i^{\text{pred.}} + b_i$.⁷ Here $s_i^{\text{pred.}}$ and b_i are respectively the predicted signal and background yields in bin i , and r is the signal strength we wish to extract. For this example just a single background is used, but extending to multiple backgrounds is straightforward and simply adds notational complexity. Such event counts are distributed following a Poisson distribution, so that a likelihood function for this observation and prediction takes on the following form:

$$\mathcal{L}(r) = \prod_i^N \frac{(rs_i^{\text{pred.}} + b_i)^{o_i}}{o_i!} e^{-(rs_i^{\text{pred.}} + b_i)} \quad (3.5)$$

All bins are considered to be uncorrelated, but the signal strength r is common between all of them. For any value of r this equation gives us a likelihood value, with larger values corresponding to a better agreement between observation and prediction.

Other uncertainties can now be added in, which are implemented as so-called nuisance parameters, collectively written as $\theta = (\theta_s, \theta_b)$. Both the signal and the background depend on these nuisances, so that we could in principle write them as $s_i^{\text{pred.}}(\theta_s)$ and $s_i^{\text{pred.}}(\theta_b)$. The normalization of the background(s) and the statistical uncertainties in the signal and background predictions themselves will also be implemented through these nuisance parameters.⁸ These nuisance parameters are values to be varied in the fit in order to maximize the likelihood, but are only of secondary interest in the measurement, and unlike the signal strength r they do not float freely in the fit. Their observed or measured values $\hat{\theta}_i$ are distributed around their assumed θ_i values following a collective probability distribution which can be written as $\rho(\tilde{\theta}|\theta)$. Here $\rho(\tilde{\theta}|\theta)$ essentially corresponds to the multiplication of the probability distribution functions of all added nuisances. Adding these into the likelihood gives us:

$$\mathcal{L}(r, \theta) = \prod_i^N \left[\frac{(rs_i^{\text{pred.}} + b_i)^{o_i}}{o_i!} e^{-(rs_i^{\text{pred.}} + b_i)} \right] \rho(\tilde{\theta}|\theta) \quad (3.6)$$

The functions individual nuisances are represented by are typically log-normal distributions, which are similar to regular normal distribution but describe values which can only be positive.⁹ In practice a value related to the likelihood $\mathcal{L}$ is typically used instead, defined as:

$$-2\Delta \ln \mathcal{L}(r) = -2 \ln \left(\frac{\mathcal{L}(r, \hat{\theta}_r)}{\mathcal{L}(\hat{r}, \hat{\theta})} \right) = 2 \left(\ln \mathcal{L}(\hat{r}, \hat{\theta}) - \ln \mathcal{L}(r, \hat{\theta}_r) \right) \quad (3.7)$$

⁷The word “number” is somewhat of a misnomer here as predicted events have weights differing from unity, but in the end the total yield does represent the prediction of a discrete number of events.

⁸The signal strength does not have an uncertainty directly associated with it, so while one generally expects it to be close to 1, this is in no way enforced in the fit. If the signal prediction were to be scaled by a factor of X the same signal strength would be obtained multiplied by $1/X$. Nuisance parameters on the other hand are tied to their nominal value through their probability distribution appearing in the likelihood, so that deviating from this value is penalized in the fit.

⁹There is of course a good justification for why this is the case. The distribution of a value obtained through the multiplication of a number of random positive variables distributed with a finite variance approaches a log-normal distribution as the number of variables increases. This principle is known as the “multiplicative central limit theorem” or “Gibrat’s law”. Many variables in nature implicitly arise as a multiplication of the many factors which determine them, so that their distribution is well described as a log-normal.

Where $\hat{r}$ and $\hat{\theta}$ are the values of the signal strength and the nuisances which globally maximize the likelihood. This test statistic is larger for worse fitting predictions and reaches its minimum of 0 in the best fit point $\hat{r}$. The fraction $\mathcal{L}(r, \hat{\theta}_r) / \mathcal{L}(\hat{r}, \hat{\theta})$ appearing in this equation is referred to as the profile likelihood ratio. In addition to a best-fit point this test statistic also provides a measure of how well the predicted signal in the bins matches the observation for any value of the signal strength. When these values are calculated the values of the nuisance parameters are optimized to minimize the test statistic for the given value of the signal strength, which is referred to as “profiling” over the nuisance parameters. In the notation above these profiled values of the nuisances for a given signal strength are written as $\hat{\theta}_r$.

Within the example described here the signal strength is the parameter of interest. The fitting procedure can however be used more generally. A simple variation would be the case in which the signal or the signal strength is dependent on some other variable, where this secondary variable is the parameter of interest. This will be the case in the context of effective field theory described later, where the signal will depend on a certain coefficient, with a different dependence in each of the bins.

In addition to the best-fit value for the parameter of interest and the nuisances confidence level (CL) intervals are typically quoted. These are ranges of the parameter of interest which enclose a certain area of the total probability distribution. In other words the $X\%$ CL interval is defined such that such ranges obtained in repeated experiments have a probability of $X\%$ of enclosing the true value of the parameter of interest. Typically these values are quoted for the probabilities of 68% and 95%, which respectively correspond to one and two standard deviations of a normal distribution. In case the reader is unfamiliar with likelihood functions it can be useful to look ahead to figure 5.37 on page 105 for a few example graphs, the full meaning of which is not important at this point. If multiple parameters of interest affect the signal two kinds of CL intervals can be defined: First of all a single parameter of interest can be considered, while all others are fixed to some assumed value, typically corresponding to the SM situation. As a second option the CL interval for a parameter can be evaluated while all other parameters are allowed to float to the value for which the likelihood is maximized, similar to the profiling over nuisance parameters. The CL intervals obtained in this way are wider and therefore more conservative if they correspond to some measured constraint on the parameter of interest, and are referred to as “marginalized” CL intervals. The values of $-2\Delta \ln \mathcal{L}(r)$ which correspond to the CL interval edges for a certain confidence level are given by the inverse cumulative distribution function of the chi-squared distribution. For estimations involving a single degree of freedom the 68% and 95% confidence levels respectively correspond to $-2\Delta \ln \mathcal{L}(r) = 1$ and $-2\Delta \ln \mathcal{L}(r) = 3.84$. For 2 degrees of freedom these values become 2.30 and 5.99, respectively.¹⁰

¹⁰The scipy library for python provides the ppf method under `scipy.stats.chi2` using which these values can be calculated for any percentage and number of degrees of freedom.

Chapter 4

Effective field theory

4.1 Introduction

The conventional direct search approach to looking for new physics involves choosing some theory that extends upon the standard model, identifying a final state that one of the added particles could decay into, then specifically searching for this signature in the data. This kind of search can be partially parametric, when for example the mass of the particle is not fixed by the theory, so that a search over some mass range is performed. A popular example is supersymmetry, which proposes the existence of multiple undiscovered particles, of which the masses are unknown. Searches of this kind have thus far been unsuccessful in finding signs of new physics, with constraints on the existence of these particles becoming ever more stringent [77]. A different approach comes in the form of standard model effective field theory (SMEFT), often also simply called effective field theory (EFT).

The central idea of EFT is that the existence of new particles, at energy scales well below their mass, should affect the yield and kinematics of known processes. This effect corresponds to the addition of new vertices or the modification of existing couplings in the SM theory. A straightforward example would be the existence of a Z' boson, with a mass of above 5 TeV but otherwise similar to the Z boson of the standard model [78]. At energies below this invariant mass its existence would manifest itself as a 4-fermion interaction vertex, similar to the 4-fermion coupling of the Fermi theory below the W boson mass.

Figure 4.1 shows leading order diagrams for the decay of a muon into an electron and two neutrinos as described by the SM, and as it is described by the Fermi theory [79]. In the framework of the SM this process occurs through a virtual W boson, and each of the weak interaction vertices is associated with the appearance of a factor of the weak coupling constant g in the matrix element. The fermi model, an earlier attempt to describe beta decay and related processes, features a direct four-fermion interaction associated with a coupling given by the Fermi constant G_F . At energies below the W boson mass these descriptions lead to very similar predictions, and a relationship between the two couplings strengths exists:

$$G_F = \frac{\sqrt{2}}{8} \frac{g^2}{M_W^2 c^4}, \quad (4.1)$$

where M_W is the mass of the W boson and c is the speed of light in vacuum. When calculations involving Fermi coupling terms are however performed cross section values are found that are quadratically dependent on the energy in the interaction, and predictions lose validity at energies exceeding 100 GeV. Because of this the Fermi theory is said to be a low-energy effective field theory for the weak interaction model of the SM.




FIGURE 4.1: Leading order Feynman diagrams for muon decay as described by the SM (left) and the Fermi theory (right).




FIGURE 4.2: An exaggerated picture illustrating how directly searching for new physics, sometimes also called “bump hunting”, differs from the EFT approach in terms of the assumed energy scale of new physics processes.

The effect of any hypothetical heavy particle can be modelled by the addition of extra operator terms to the SM lagrangian, and changes induced by the interactions this introduces are expected to appear in the high energy tails of distributions. However, in EFT the attention is shifted from specific new physics particles to these additional operators themselves. By adding the full set of possible operators any deviation from SM predictions can be interpreted in a model independent way. A simple illustration of the two different approaches to searching for new physics is shown in figure 4.2. A lagrangian for an EFT theory can be written as

$$\mathcal{L} = \mathcal{L}_{SM} + \sum_i \frac{c_i}{\Lambda^2} \mathcal{O}_i, \quad (4.2)$$

where a summation over all operators $\mathcal{O}_i$ of dimension six is shown. In this context the term “dimension” is used to refer to the exponent of the unit of energy in the operator. The standard model operators themselves are of dimension four or lower, and the terms added here match that as the operators are suppressed by the mass scale Λ squared. This mass scale loosely corresponds to the scale of the new physics particles, and is for practical purposes usually set to 1 TeV.¹ Operators with a dimension of higher than 6 should in principle be added for a fully complete picture, but they are suppressed by higher orders of this mass scale Λ . Their effects are therefore of secondary importance, and they are omitted at this point. There is only one operator of dimension 5, which generates the Majorana neutrino masses, which is also not included here. Dimension 5 operators violate lepton number conservation, and Majorana masses are already strongly constrained, with further improvements likely to come from direct searches and neutrino oscillation experiments. With these constraints in mind only the dimension six operators will be of interest here.

The c_i parameters shown in the lagrangian are called Wilson coefficients, which can be thought of as a measure of the strength of the couplings, and are in general complex numbers.² When performing interpretations of measurements within the EFT framework the goal will be to find constraints on these Wilson coefficients. In total, there are 59 independent operators of dimension six per quark and lepton generation, where only those which do not lead to baryon and lepton number violation are considered. While the underlying physics is independent of the choice of operator basis, the most commonly used set of definitions corresponds to the Warsaw basis, the operators of which are exclusively constructed out of SM fields. Fifteen of them are relevant to top physics in general, but for a specific physics process the number of relevant operators is typically much lower. Examples of these operators will be given later on, and a full listing of the Warsaw basis can be found in Ref. [9]. A theory constructed in this way has several favourable theoretical features:

- The gauge symmetries, Lorentz invariance, and unitarity of the standard model are kept intact.

¹In practice this mass scale acts as a lower limit to the mass scales relevant to new physics, below which the framework of EFT loses validity.

²Named after Kenneth Wilson, who was one of the first people to conceive of the idea of an effective field theory and worked to lay the foundations of this framework.

- In the high mass or low energy limit the standard model is recovered.
- The theory is general enough to model any new physics effects within the reasonable constraints listed earlier.
- Despite the model independence, the operators can be matched to new physics models which could produce their respective couplings. (Or to put it crudely, the EFT framework is not a purely mathematical trick used to justify fitting a large amount of parameters.)

In addition to this the EFT approach is inherently sensible with regards to unitarity violation. As the added EFT terms grow with increasing values of the energy, they eventually lead to unitarity violation. At these energies however the initial assumption of considering only energies for which $E < \Lambda$ is also violated, so that the theory does not claim to be reliable beyond this range. If a high mass particle were to be discovered a new set of operators would need to be constructed, incorporating the fields associated with this new particle, and the energy scale Λ should be set to a value beyond the mass of this new heavy particle.

The addition of EFT operators to the SM lagrangian affects both the overall yield and the kinematics of processes in which the corresponding couplings are relevant. Based on equation 4.2 it is clear that the general cross section dependence for any set of operators will take on the form

$$\sigma = \sigma_{SM} + \sum_i \frac{c_i}{\Lambda^2} \sigma_i + \sum_{i,j} \frac{c_i c_j}{\Lambda^4} \sigma_{ij} \quad (4.3)$$

or in the case of a single operator:

$$\sigma = \sigma_{SM} + \frac{c_i}{\Lambda^2} \sigma_i + \frac{c_i^2}{\Lambda^4} \sigma_{ii}. \quad (4.4)$$

Here, the first term is the cross section obtained using purely the SM lagrangian, the second corresponds to the interference terms between SM diagrams and EFT diagrams, and the last term is a pure EFT term. In the case of multiple EFT operators being added the last summation includes terms corresponding to the interference between different operators, so that Wilson coefficients can a priori not be studied and constrained independently. These equations also only hold in the case where at most 1 EFT vertex is allowed to be inserted into the diagrams for the process being considered. This is once again a reasonable simplification to make, given that any diagrams containing two such vertices are suppressed by a factor of Λ^{-4} . The principle is exactly the same for differential distributions, where the cross section for a specific bin will show the same dependence. For each of the bins the parametrization will of course be different, so that the shape of distributions can change as a function of the Wilson coefficients. In fact, as will be discussed later, the simulation weights associated with individual events share this same dependence.³

In most analyses the number of coefficients considered is relatively small, given that a few practical assumptions can be made. Any operator that would modify a well-constrained interaction for any reasonably large value of its Wilson coefficient can be considered to be irrelevant ($c = 0$). This means that for a specific process not all operators are necessarily relevant at first or even higher orders, so that these too can be ignored within the scope of process-specific analyses.

At this point a picture of the past, present, and future of EFT from an experimental perspective can be sketched, so that once the EFT interpretation of the measurement of $t\bar{t}\gamma$ is arrived at it can immediately be appreciated as a piece of a larger puzzle. The operators relevant to the signal will be described first, so that in subsequent sections chosen examples can be related to this work. Once this is done an alternative to EFT in the form of anomalous couplings can be explored, which will prove to be founded on very different principles, while in some regards being equivalent to the EFT approach. The practical implementation of these ideas is then discussed, along with the considerations which need to be made when it comes to using larger numbers of operators. Finally, an overview is provided of several analyses involving similar approaches to the one documented here, an analysis making clever use of multivariate techniques to increase its sensitivity to EFT, and an approach in which results from many existing analyses can be collectively interpreted in the EFT framework.

³This is a rather obvious remark, but its relevance to the feasibility of using differential information in an EFT interpretation will become clear later on.

4.2 Operator selection

Before moving forward we need to get an idea of how these operators look. All of them are constructed out of SM fields, so that looking at a few relevant examples should suffice for our purposes. Also in an analysis it is pragmatic to focus on operators to which the process under study is especially sensitive, while leaving others to be constrained using other measurements. This is also the case for $t\bar{t}\gamma$ production, where such considerations narrow the scope down to just two operators:

$$\mathcal{O}_{uW}^{(33)} = (\bar{q}_3 \sigma^{\mu\nu} t) \tau^i \tilde{\phi} W_{\mu\nu}^i \quad (4.5)$$

$$\mathcal{O}_{uB}^{(33)} = (\bar{q}_3 \sigma^{\mu\nu} t) \tilde{\phi} B_{\mu\nu} \quad (4.6)$$

Here (33) indicates that the operators are specific to the third generation of quarks, since lighter generations are essentially implicitly constrained in any SM measurement involving quark-boson couplings. This is reflected in the definition only involving the left-handed third generation quark doublet written as q_3 and the right-handed top quark singlet t . $W^{1,2,3}$ and B are the electroweak fields of the unbroken standard model, and ϕ the Higgs doublet. Finally, τ^i are the Pauli matrices, and $\sigma^{\mu\nu} = \frac{i}{2}[\gamma^\mu, \gamma^\nu]$ with γ referring to the Dirac gamma matrices. A full mathematical description of the EFT operators is of limited use here, but can be found in Ref. [80]. A linear combination of these operators can be made to isolate the modifications of the t - γ and t - Z couplings, and their real and imaginary components are considered separately, resulting in the set of Wilson coefficients that will later be used in the analysis:

$$c_{tZ} = \text{Re} \left(-\sin \theta_W c_{uB}^{(33)} + \cos \theta_W c_{uW}^{(33)} \right) \quad (4.7)$$

$$c_{tZ}^I = \text{Im} \left(-\sin \theta_W c_{uB}^{(33)} + \cos \theta_W c_{uW}^{(33)} \right) \quad (4.8)$$

$$c_{t\gamma} = \text{Re} \left(\cos \theta_W c_{uB}^{(33)} - \sin \theta_W c_{uW}^{(33)} \right) \quad (4.9)$$

$$c_{t\gamma}^I = \text{Im} \left(\cos \theta_W c_{uB}^{(33)} - \sin \theta_W c_{uW}^{(33)} \right) \quad (4.10)$$

Here θ_W is the Weinberg angle, and “Re” and “Im” represent taking the real and imaginary components of the terms. Finally, it can be assumed that the Wtb vertex is well constrained from measurements of W helicity fractions [81], which corresponds to setting $c_{uW}^{(33)} = 0$ in the description above. This makes the two pairs of operators above differ only by a constant factor, so that one can choose either one to proceed with. In order to make comparison to other analyses later on straightforward the c_{tZ} and c_{tZ}^I coefficients are chosen. Finally, it is interesting to note that these equations illustrate the difference in sensitivity to the $\mathcal{O}_{uB}^{(33)}$ operator between the t - γ and t - Z vertices, with the former being $\cos \theta_W / \sin \theta_W \approx 1.83$ times more sensitive than the latter.

Narrowing the scope down to just a few operators is typically a necessary evil in analyses not aimed directly at performing a comprehensive study of all dimension six operators. In fact even in such analyses a subset of the available operators is typically chosen, albeit a larger one than for example considered here. The process of selecting operators to which the process at hand is most sensitive, and dropping those on which strong constraints have already been obtained elsewhere, is in itself perfectly justifiable. It does however make the implicit assumption that operators which minimally affect a certain process don’t exhibit a larger effect through interference terms involving couplings from other EFT operators. The fact that stringent constraints on a given Wilson coefficient exist should also not be simplified down to stating that these are demonstrably zero. For example when deriving CL intervals for Wilson coefficients while profiling over additional coefficients much wider ranges are obtained compared to the case of setting all other Wilson coefficients to zero, even if existing tight constraints on these other operators are taken into account. When reporting measured CL intervals for individual coefficients values are usually quoted for the case of fixing all other coefficients to zero, and profiling over the other coefficients being considered, but leaving all other coefficients still set to zero. The takeaway from this is that all constraints on EFT operators should to some degree be interpreted as being sensitivity study results, with constraints from analyses with a wider scope being closer to what one might consider to be actual constraints, but perhaps never fully so because of the unbounded range of operators that could potentially be added.

4.3 Anomalous couplings

The framework of EFT is not the only way to study new physics through potential modifications of SM couplings. A similar approach under the name of anomalous couplings was introduced around the same time as EFT in the late 1970's [82]. The idea of this is simply to allow for deviations from the gauge theory predictions for vector boson couplings. At the time the spontaneous breaking of gauge symmetry was not yet confirmed, so that such hypotheses were not unreasonable. Electroweak precision measurements have since then made it overly clear that the gauge symmetry of the standard model should also be present in any extension to it. As one can guess based on the fact that this work includes an interpretation in the framework of EFT, and not in terms of anomalous couplings, this alternative approach has several drawbacks. Besides the fact that the SM gauge symmetry is not incorporated in an anomalous coupling model, its main problems are:

- No consistent way of treating loop corrections when anomalous couplings are allowed. This means that predictions at the level of precision needed at this point in time are not possible, which was less of a pressing issue before the turn of the century.
- While the SM is of course recovered when the relevant couplings are set to their SM values, the way in which a SMEFT theory approaches the SM situation for $\Lambda \rightarrow \infty$ is much more natural.
- The anomalous couplings correspond to a limited range of new physics scenarios, and the method therefore cannot be considered to be truly model independent.

It is nevertheless interesting to look at an example of such an anomalous coupling, as it is in the end still the couplings we are interested in. We can look at the lagrangian describing the interaction between the top quark and the photon, in which we can illustrate the anomalous magnetic dipole moment (MDM) and electric dipole moments (EDM) [83]:

$$\mathcal{L}_{t\bar{t}\gamma} = e\bar{t} \left(Q_t \gamma_\mu A^\mu + \frac{1}{4m_t} \sigma_{\mu\nu} F^{\mu\nu} (\kappa + i\tilde{\kappa}\gamma_5) \right) t \quad (4.11)$$

Here A^μ is the photon field, Q_t is the charge of the top quark and hence $= 2/3$, m_t is the top quark mass, γ refers to the Dirac gamma matrices, and $\sigma_{\mu\nu} = \frac{i}{2} [\gamma_\mu, \gamma_\nu]$.⁴ $F^{\mu\nu}$ is the electromagnetic field tensor, or what can in this context be thought of as representing the kinetic energy of the photon. The factors κ and $\tilde{\kappa}$ are now respectively related to the anomalous MDM and anomalous EDM of the top quark, where the first is CP even and the latter is CP odd. The actual anomalous dipole moments can be defined as a function of these as:

$$a_t = \kappa/Q_t \quad (4.12)$$

$$d_t = e\tilde{\kappa}/2m_t \quad (4.13)$$

Given that the SMEFT framework is supposed to be able to model any new physics effects it should also provide operators which give rise to these anomalous dipole moments. This is indeed the case, and it turns out that it is precisely the $\mathcal{O}_{uB}^{(33)}$ and $\mathcal{O}_{uW}^{(33)}$ operators mentioned earlier which do this.⁵ After going through the calculations, which the authors of Ref. [87] have kindly done for us, we find:

$$\kappa = -\frac{2\sqrt{2}}{e} \frac{vm_t}{\Lambda^2} \text{Re} \left(\sin \theta_W c_{uW}^{(33)} + \cos \theta_W c_{uB}^{(33)} \right) \quad (4.14)$$

$$\tilde{\kappa} = -\frac{2\sqrt{2}}{e} \frac{vm_t}{\Lambda^2} \text{Im} \left(\sin \theta_W c_{uW}^{(33)} + \cos \theta_W c_{uB}^{(33)} \right) \quad (4.15)$$

Which means that the anomalous dipole moments can be written as a linear combination of EFT operators. In other words, if we wish, we could identify this combination as a new operator and use it instead.⁶

⁴Some new physics theories propose other values for the charge of the top quark, but both searches for such exotic top quarks [84] and direct measurements of the top charge [85, 86] have practically excluded this possibility.

⁵I have of course structured the text for this revelation to take appear here, but a revelation every now and then doesn't hurt.

⁶This is to some degree a matter of choice, and since $c_{uB}^{(33)}$ and $c_{uW}^{(33)}$ are typically the coefficients for which constraints are reported there is no reason to deviate from this.

4.4 SMEFT event simulation and reweighting

The practical implementation of the methods described here was inspired by the work of Robert Schöfbeck and Lukas Lechner in the context of their measurement of the $t\bar{t}\gamma$ cross section in the lepton+jets channel [6]. This setup was in turn based on methods first described in the context of Ref. [88]. I therefore include the general description of the framework here, and only provide analysis-specific details later on in the chapter on the analysis work I performed.

4.4.1 The dim6top model

As hinted at in the introduction to Monte Carlo event simulation, the ME level calculations are not restricted to purely SM settings. The so-called **dim6top** model [80, 89] implements baryon and lepton number conserving EFT operators derived from the Warsaw basis mentioned earlier, and allows for the simulation of events with any reasonable values for the Wilson coefficients. The maximum number of EFT vertices to be inserted into process diagrams can be chosen, but this and most analyses allow for at most 1 new physics vertex. While the value of Λ can be arbitrarily large, within this framework it is by default set to $\Lambda = 1$ TeV. This means that the only difference between quoted values of the Wilson coefficients c_i and the factors c_i/Λ^2 is the multiplication by the units TeV^{-2} . A few reasonable assumptions and limitations should be mentioned for this model:

- Simulations are limited to leading order in QCD. A more recently released alternative model under the name of **smeft@nlo** should make it possible to include one-loop corrections in the predictions used in upcoming analyses [90].
- The CKM matrix is approximated as the unit matrix, and the couplings of the third generation quarks are treated as completely independent from the first two generations. The latter assumption is typically referred to as minimal flavour violation [80].
- The masses and Yukawa couplings of the e and μ fermions, as well as those of the light quarks, are neglected. Whether the bottom quark is considered among these light quarks can be chosen, and other settings of the simulation are automatically adjusted accordingly.
- As would be the case with any ME calculation model, the new physics effects are not considered in the showering step. This is not a problem as long as the process definition ensures none of the vertices of interest come into play past the first simulation step, but does imply the assumption that the decays past this point occur in line with the SM description.

4.4.2 Event reweighting

The main idea of an EFT interpretation is to compare measured yields and/or distributions to new physics scenarios, attempting to find the hypothesis that best aligns with the observations. These hypotheses correspond to a continuous range of Wilson coefficient values, each parametrizing the strength of the couplings their corresponding operators introduce. A priori this would mean that samples need to be produced over a range of Wilson coefficient values, and preferably with a fine granularity. Doing so is impractical when even just a single operator is considered, and becomes impossible as more operators are added. For the inclusive cross section it was already apparent that the dependence on the Wilson coefficients can be parametrized as a second order polynomial. The fact that this is also the case for the individual event weights now allows us to avoid performing dedicated simulations for every hypothetical scenario by reweighting samples from one point in the parameter space to another.

The parametrization of the event weights takes on the following familiar form:

$$w = w_{SM} + \sum_i \frac{c_i}{\Lambda^2} w_i + \sum_{i,j} \frac{c_i c_j}{\Lambda^4} w_{ij} = s_0 + \sum_i c_i s_i + \sum_{i,j} c_i c_j s_{ij}, \quad (4.16)$$

where I have introduced the structure constants s which need to be obtained to describe the dependence. This can be done by performing dedicated calculations for several points of the parameter space, and using these to extract the values of the structure constants. To illustrate this method it is

convenient to look at the simplest case of a single Wilson coefficient, where the parametrization can be written as the product of two vectors:

$$w(c_x) = s_0 + s_1 c_x + s_2 c_x^2 = \begin{bmatrix} 1 & c_x & c_x^2 \end{bmatrix} \cdot \begin{bmatrix} s_0 \\ s_1 \\ s_2 \end{bmatrix}, \quad (4.17)$$

where in this case c_x refers to some specific value of the Wilson coefficient, and hence a point in the phase space. With the weight calculated at 3 points in the phase space the structure constants can be determined by solving a quadratic equation that can once again be written using matrices:

$$\begin{bmatrix} w(c_x) \\ w(c_y) \\ w(c_z) \end{bmatrix} = \begin{bmatrix} 1 & c_x & c_x^2 \\ 1 & c_y & c_y^2 \\ 1 & c_z & c_z^2 \end{bmatrix} \cdot \begin{bmatrix} s_0 \\ s_1 \\ s_2 \end{bmatrix} \quad (4.18)$$

Here, each row in the square matrix corresponds to a probed point of the phase space. Inverting this to obtain the expression for the structure constants then comes down to inverting this matrix:

$$\begin{bmatrix} s_0 \\ s_1 \\ s_2 \end{bmatrix} = \begin{bmatrix} 1 & c_x & c_x^2 \\ 1 & c_y & c_y^2 \\ 1 & c_z & c_z^2 \end{bmatrix}^{-1} \cdot \begin{bmatrix} w(c_x) \\ w(c_y) \\ w(c_z) \end{bmatrix} \quad (4.19)$$

Notice that while the structure constants s are different for every event, or in other words for every weight to be parametrized, the square matrix depends only on the values of the Wilson coefficients. This matrix is therefore the same for all events, assuming they are generated at the same points in the parameter space, so that the inversion needs to be performed only once.⁷ With this the parametrization for any event can be easily obtained, and the weight $w(c_m)$ can be calculated for any value of c_w .

It is important for the sample being reweighted to cover the phase space that is being reweighted to, which is not necessarily the case when extrapolating from a SM sample to an EFT hypothesis given that operators can introduce new couplings. The sample to be used is therefore typically generated with all relevant Wilson coefficients set to some relatively large value. The ability of the reweighting procedure to reproduce the yield and kinematics corresponding to a specific point in the phase space can simply be checked by directly generating samples at a few phase space points and comparing their distributions to those obtained from the reweighted sample. This method was first developed by the authors of Ref. [88], and the excellent explanations in the analysis notes corresponding to Ref. [6] and [91], which use similar approaches, were of great use to me in understanding this implementation.

4.4.3 Multi-dimensional parametrizations

The description of the reweighting procedure given in the previous section was limited to a single Wilson coefficient for simplicity. In reality the implementation was created to work with an arbitrary number of coefficients, as the mathematical procedure can be generalized. The generalized dependence of a cross sections on N Wilson coefficients can be written as:

$$\sigma(c_1, c_2, \dots) \propto |\mathcal{M}_{SM} + c_1 \mathcal{M}_1 + c_2 \mathcal{M}_1 + \dots|^2 \quad (4.20)$$

Which results in 1 SM term, N linear terms, N quadratic terms, and $\frac{1}{2}N(N-1)$ interference terms.⁸ In other words the total number of structure constants N_s is given by $N_s = 1 + 2N + \frac{1}{2}N(N-1)$.

⁷While reasonable choices need to be made regarding which points of the parameter space are probed, there is no reason to not keep them fixed for all events, so that no trade-off is made here.

⁸This last number is nothing more than the number of possible unique pairs given by the binomial coefficient $\binom{N}{k}$ with $k = 2$, so that one finds $\binom{N}{2} = \frac{N!}{2!(N-2)!} = \frac{1}{2}N(N-1)$.

Looking at the example of $N = 3$ should also make this clear:

$$\sigma(c_1, c_2, c_3) = \underbrace{s_0}_{\text{SM}} + \underbrace{s_1 c_1 + s_2 c_2 + s_3 c_3}_{\text{linear}} + \underbrace{s_4 c_1^2 + s_5 c_2^2 + s_6 c_3^2}_{\text{quadratic}} + \underbrace{s_7 c_1 c_2 + s_8 c_1 c_3 + s_9 c_2 c_3}_{\text{interference}} \quad (4.21)$$

In which $s_0 = \sigma_{SM}$ and the other structure constants s are terms of the form σ_i/Λ^2 or σ_{ij}/Λ^4 as they appeared in eq. 4.3.

As mentioned earlier these structure constants cannot be determined in a direct way, and therefore need to be extracted from a few directly calculated points in the phase space using the procedure described earlier. The minimum number of points needed to extract this parametrization is then also equal to N_s , assuming the points are chosen appropriately. While the number of calculations this requires increases about quadratically as a function of the number of operators, this is still much more feasible than performing fine grained scans over the full parameter space.

4.5 Interpretations of standard model measurements

What will follow in this section is an overview of several analyses by CMS and ATLAS which either include an EFT interpretation or are directly relevant to EFT. Some of these analyses are not specifically constructed around the EFT interpretation, which for authors means that the publication does not stand or fall based on whether or not the extra interpretation works out. This has been a common approach over the past few years, as a collective understanding of EFT was still developing within the collaborations, and corresponding tools were being created. The $t\bar{t}\gamma$ measurement, which will be the main topic of this work, falls within this category. We will also see an analysis where the main result is not a cross section, but an observable sensitive to new physics (energy asymmetry), which is then also interpreted in the context of EFT. Finally, we will have a look at an analysis fully aimed at constraining Wilson coefficients. I believe that moving forward analyses purely focused on EFT will become more common, and that those which provide both “standard” cross section results and EFT interpretations will put a larger emphasis on the latter.

The end goal of any EFT interpretation is to obtain constraints on the values of Wilson coefficients relevant to the signal being considered. In a typical cross section measurement a maximum likelihood fit is used to find the cross section for the signal which best agrees with the observations in data, taking into account the various uncertainties in the signal and backgrounds. The value that is extracted in such a fit is the signal strength, which represents the factor by which the predicted number of events needs to be scaled according to the fit. Along with this best-fit signal strength a total uncertainty is also obtained. In an EFT interpretation context this fit is instead performed as a function of one or more Wilson coefficient values. The obtained results are the best-fit values for these Wilson coefficients and the 1 and 2 σ confidence intervals.

Measurements of the $t\bar{t}Z$ cross section

Thanks to the large amounts of data provided by the LHC, several precision measurements of $t\bar{t}Z$ production have been performed by both CMS and ATLAS over the past few years. Due to the fact that $t\bar{t}Z$ and $t\bar{t}W$ form backgrounds to each other their cross sections are sometimes measured simultaneously. A chronological overview of these measurements and their approach to EFT interpretations can give us an overview of how methods and interests have shifted over the years.

In a paper published by the CMS collaboration in 2016 we find a simultaneous measurement of the $t\bar{t}Z$ and $t\bar{t}W$ cross sections using $\sqrt{s} = 8$ TeV of proton-proton collision data [92]. After the main results section the authors describe how the $t\bar{t}Z$ process could be affected by anomalous vector and axial couplings between the top quark and the Z boson. These anomalous couplings are written as a function of EFT operators, and a fit as a function of these two couplings is performed. The attention is then shifted to the EFT operators themselves, where 5 operators of interest are identified. These operators are not implemented through the `dim6top` model, but come from the so-called `HEL` (Higgs

Effective Lagrangian) model which makes use of a different basis [93, 94]. This HEL model is very similar to the `dim6top` one, but was developed in the context of Higgs physics, which is reflected in the operators it includes. One of the operators considered is $\mathcal{O}_{uB}$, where at this point the quark generations are implicitly assumed to be affected to the same degree.

In a later paper (2018) by CMS in which another simultaneous measurement of $t\bar{t}Z$ and $t\bar{t}W$ is described a first improvement is made, this time using $\sqrt{s} = 13$ TeV collision data collected in 2016 [53].⁹ While the HEL model is still being used, the authors have manually modified the implementation so that the EFT couplings of the 3 quark generations are independent. This means that the operators described earlier, like $\mathcal{O}_{uB}^{(33)}$, are used instead of the generation-agnostic ones like $\mathcal{O}_{uB}$ used in the first publication. Anomalous couplings are also no longer explicitly mentioned, reflecting how the attention has shifted to the EFT paradigm.

In 2019 ATLAS also published a simultaneous $t\bar{t}Z$ and $t\bar{t}W$ measurement with an added EFT interpretation, in which one year of data recorded at $\sqrt{s} = 13$ TeV is used [95]. The operators considered here are also specific to the third generation of top quarks.¹⁰ In this analysis and those before it only one operator is considered at a time. It is also consistently only the real components of Wilson coefficients which are considered, even though this is typically not mentioned explicitly.

Another step forward is found in the $t\bar{t}Z$ measurement by CMS using 2016 and 2017 data, published in 2020 [54]. The analysis provides differential cross section measurements, and also uses differential information for its new physics interpretation. Constraints on both anomalous couplings and Wilson coefficients are provided, and both the real and imaginary components of Wilson coefficients are considered. The dependence of the yield and shapes on the Wilson coefficients is obtained by performing separate simulations for each point of a 2D grid of the Wilson coefficients. A maximum likelihood fit can then be performed for every point on this grid in order to find the global minimum. A fine grained grid is used so that the results do not suffer from the fact that the best fit point as a function of the coefficients is not obtained in a continuous way. This kind of method would however not extend well to higher numbers of simultaneous operators, which is the reason the reweighting method described earlier has since been developed.

Most recently a measurement of $t\bar{t}Z$ production by ATLAS using 3 years of data at $\sqrt{s} = 13$ TeV was published [96]. This paper does not include an EFT interpretation in any way. This should however not be interpreted as a lack of interest in application of a SMEFT interpretation. Detailed differential results are included for several variables that have been established to be highly sensitive to EFT effects. This means that an EFT interpretation can later be performed by a third party, most notably in the context of a combination with other processes.¹¹ An example of such a combined interpretation performed by theorists will be given later on. Additionally the $t\bar{t}Z$ process might be more interesting to study in a wider context of t - Z coupling sensitive processes when it comes to EFT effects, so that this study might be separated from the cross section measurements entirely. The CMS collaboration has such an analysis published, which is given as an example in the next subsection.

With this we have seen an example of the evolution of EFT interpretations in top physics. Early on more attention was given to anomalous couplings, and EFT interpretations were “added on top”. Attention shifted to EFT operators, and increasing thought and effort was put into the implementations. At this point in time dedicated EFT analyses are gaining traction, and increases in sensitivity and scope are continually being worked on. It should also be noted that while the CMS and ATLAS collaborations have their own internal analysis tools, most event simulation tools and models like `dim6top` are developed independently, and can therefore be used by both parties. This means that combinations are possible, as long as the same models and conventions are used so that all parties are “speaking the same language”.

⁹This work was performed within the UGent group, with Illia Khvastunov and Deniz Poyraz as the main analysers.

¹⁰It is typically possible to compare constraints obtained on Wilson coefficients from different models, but this is non-trivial and great care needs to be taken in making sure that all factors are matched between the two. For converting between some common models converters are available.

¹¹The authors themselves can of course also publish an EFT interpretation of these results later on, for which they would have even more detailed information in hand than is provided to the public.

Search for new physics in top quark production with additional leptons

An EFT interpretation built on top of a cross section measurement of a specific process inherently makes assumptions about the backgrounds to this processes. The effect of the operators on these backgrounds can be included, and control regions can be defined to help constrain them, but that is about the extent to which backgrounds can be focused on. The specificity of the selection also means that the measurement is only relevant to a small number of operators. Other Wilson coefficients are therefore assumed to be zero, as they could not be constrained anyway. While this is a reasonable assumption, it does additionally introduce the assumption that these operators do not become relevant through interference terms. On the other hand if multiple measurements are combined based on published information the correlations between processes and uncertainties are necessarily either neglected or approximated.

In a recent publication by CMS a broad final state definition of a top quark with additional leptons is used instead of a process-specific selection [88]. Within this selection $t\bar{t}H$, $t\bar{t}l\bar{l}$, $t\bar{t}l\nu$, $t\bar{t}lq$, and tHq are all relevant EFT-sensitive signal processes. The broad signal region is then subdivided based on the number leptons and their charges, the number of jets and whether or not they are identified as coming from a b quark, and whether or not the leptons have an invariant mass close to the Z boson mass. Each of these regions is enriched in different signal processes, so that they are to some degree disentangled. Given the wide range of signal processes a large number of Wilson coefficients are relevant, with 15 being considered in the analysis. The non-specific signal definition also results in sensitivity to operators which introduce four-fermion operators that lead to the production of leptons without the involvement of H, W, or Z bosons.¹²

With such a large number of Wilson coefficients, and the explicit goal of constraining all of them simultaneously taking into account their interference terms, a parametrization method is needed that extends well to larger numbers of coefficients. It should therefore not come as a surprise that the reweighting method described earlier was developed in the context of this analysis. The final results are extracted using a maximum likelihood fit to the various sub-signal regions, in which several Wilson coefficients are varied simultaneously. CL intervals are provided for all coefficients, both while setting all others to zero, and profiling over all other Wilson coefficients.

4.6 Increasing sensitivity through multivariate approaches

Probing EFT operators in the associated production of top quarks with a Z boson

The suppression of EFT operators by the mass scale Λ makes it so that their contributions are most apparent towards the high energy tails of certain distributions. Because of this a higher sensitivity to the Wilson coefficients is obtained when such a distribution is used to constrain them, compared to the case where only inclusive information or a variable uncorrelated with the energy in the event is used. In the EFT interpretation of the measurement of $t\bar{t}\gamma$ production the obvious variable to look at will be the transverse momentum of the photon, where simulations will show that the high $p_T(\gamma)$ tail is much more affected by the values of c_{tZ} and c_{tZ}^I than the bulk of the yield. This does however not mean that other kinematic variables are not affected, or that all kinematic variables correlated with the energy in the event carry the same information. It is therefore possible to use multivariate techniques to optimize for the discrimination between SM and EFT contributions. In doing so the output value of such a discriminator essentially becomes a variable which is maximally sensitive to EFT effects.

This idea was implemented for the first time in a CMS analysis in a recent publication, in which a signal involving the production of one or two top quarks in association with a Z boson is used to obtain constraints on various Wilson coefficients [91]. Additionally, this publication serves as an example of an analysis aimed specifically at an EFT interpretation, in contrast to analyses which mainly aim at inclusive or differential cross section measurements, and add an EFT interpretation on top of those results. As is commonly done only the operators most relevant to the measured processes are studied, and operators that can be better constrained using other measurements are dropped. The authors

¹²For example in a $t\bar{t}Z$ measurement any excess in the top quark + leptons final state would implicitly be interpreted as a larger than predicted $t\bar{t}Z$ yield. Similar final states could arise from four-fermion operators, which would essentially form a background not present in the standard model.

also decided to focus on CP-conserving new physics hypotheses, so that only the real parts of the Wilson coefficients are used.

Multivariate techniques are used at two levels in the analysis, with the latter being the innovative one: First of all a multi-class classifier is trained to distinguish the two signal processes (tZq and $t\bar{t}Z$) from each other and from the various backgrounds. The outputs of this classifier are used to define three regions, each enriched in one of the categories. Second of all binary classifiers are trained to distinguish SM events from events where at least one Wilson coefficient has a non-zero value. For each of the relevant operators a SM vs EFT discriminator is trained, the output of which is then used as the input distribution for the fit as a function of the corresponding Wilson coefficient. When this is done all other EFT operators are disabled by fixing their coefficients to zero. Additionally, a simultaneous fit for all 5 relevant operators is performed. In this fit the input distribution is the output of a classifier trained to distinguish SM events from events related to any of the 5 operators.

In the training the EFT data sets consist of events generated for a wide range of values of the Wilson coefficient, where the range is chosen wider than existing constraints, so that bias is avoided and the neural network receives clear exaggerated examples of EFT scenarios to train on.¹³ A single training is performed on the combined data set of all samples with a non-zero Wilson coefficient. The sums of the weights are also normalized to unity, and the event weights are not a parameter used in the training, so that the network cannot rely on the normalization changes EFT operators produce. The truth information is supplied in a binary way: Events generated under a SM hypothesis are given class label 0, and events generated using a model with at least one non-zero Wilson coefficient are given label 1. In other words no continuous spectrum between “small EFT contributions” and “large EFT contributions” is specified.

This classification task is non-trivial for several reasons. First of all the most noticeable effect of most EFT operators is the change in the normalization they induce. The shape changes they produce for reasonably small values of the Wilson coefficients are small, and for some operators even completely negligible. Second of all the shape changes induced by the operators at the generator level don’t necessarily translate into changes of similar magnitude at the reconstruction level. The authors mention two operators showing a small shape effect at the simulation level, which are found to have negligible shape effects at the reconstruction level due to imperfect reconstruction. Nevertheless of the order of 50% better constraints are obtained when using classifier output compared to only inclusive information. An improvement with respect to using only inclusive information would of course also be obtained when simply using a well-chosen kinematic variable instead. The authors themselves point out that most of the discriminating power of the classifiers comes from the variable defined as the reconstructed transverse momentum of the Z boson. As they illustrate however, the contributions from other variables are not to be ignored, as even the least important variable is found to have a non-negligible separation power. Two highlight results qualitatively showing the separation power of the classifiers trained for 5 operators simultaneously are shown in figure 4.3.

4.7 Interpreting variables sensitive to new physics

Measurement of the energy asymmetry in $t\bar{t}j$ events by ATLAS

Another avenue in the application of EFT is the interpretation of measurements of new physics sensitive variables by the analysers themselves. This application is not entirely distinct from what we have seen earlier, where indeed the use of especially EFT-sensitive variables had already come into play. The example I will describe here is a measurement of the charge asymmetry in $t\bar{t}j$ events by ATLAS and its interpretation in the SMEFT framework [97]. It is different from the analyses mentioned earlier in that the main result, if there were no EFT interpretation added, would not be an inclusive or differential cross section.

The charge asymmetry of the top quarks is a value which essentially parametrizes to what extent the top quark energy tends to differ from the energy of the antitop quark in a $t\bar{t}$ event. The presence of an extra jet in the $t\bar{t}$ production process makes it so that charge asymmetry is already relevant at

¹³Referring to bias produced by assuming the best fit values for the Wilson coefficients fall within the ranges given by existing constraints.




FIGURE 4.3: Data to simulation comparisons for the classifiers trained to distinguish SM events from events involving any of 5 different EFT operators [91]. The left-hand and right-hand graphs respectively show this for signal regions enriched in $t\bar{t}Z$ and tZq events, where each classifier is also optimized for that specific process. Distributions are shown as they are after a maximum likelihood fit is performed where all 5 Wilson coefficients are simultaneously allowed to float. The middle panels display the ratios of the observed event yields to their post-fit expected values, while the lower panels show the change of the yield with respect to the SM post-fit expectation for two benchmark EFT scenarios. This is shown for both the main signal process in the region (thick lines) and for the total prediction (dashed lines).

tree-level, which would otherwise only be relevant starting at NLO. This value is highly dependent on the chirality of the top and antitop quarks, which can be modified by four-quark EFT operators with chiral dependencies.

In the paper, a measured value for this charge asymmetry is obtained and compared to the SM prediction. A priori this would be the endpoint for an experimentalist, where the parameter of interest has been measured and any further scrutiny would be left to interested theorists. Despite the fact that much effort is put into making results accessible, the provided uncertainties are however typically limited to total statistical and systematic ones. It is also not possible to make further adjustments if the theory interpretation were to call for it.

Within the framework of EFT the authors can however perform an interpretation taking into account all systematics and correlations in full detail. In this publication the authors describe selecting six relevant four-quark operators, for which best-fit values and constraints are obtained. This not only makes this measurement useful in the broader context of the dimension six SMEFT model, but also established the charge asymmetry as a variable that might be interesting for others to use in the future in an EFT context.

4.8 Global fits

An ideal collective end goal for EFT interpretations would be a simultaneous fit of a wide range of Wilson coefficients to an equally wide range of measurements, in which all processes and uncertainties are taken into account in full detail. It is however unmanageable to coordinate the huge number of people across different experiments that would require. Furthermore, for results published even just a few years ago the relevant code and expertise might not even be present anymore. What is however possible is the reinterpretation of results using just the information that was published.¹⁴ A recent publication by John Ellis et al. does just that, using information from not only the CMS and ATLAS experiments, but also from smaller and older experiments like CDF, D0, LEP, and SLD [99]. The main types of measurements used in the interpretation are as follows:

- Fourteen electroweak precision observables are used. This includes features of the Z bosons like its width and branching ratio measured by the LEP and SLD experiments, and measurements of the W boson mass measurements by CDF, D0, and ATLAS.
- A total of 118 diboson measurements, including several measured differential distributions, obtained by LHC experiments.
- A wide range of Higgs-related measurements performed by the CMS and ATLAS collaborations, mostly cross section measurements in the various decay channels, for a total of 72 results.
- Top-related measurements including processes with one or two top quarks and additional bosons, 137 in total, also obtained by CMS and ATLAS.

These different categories of observables each have Wilson coefficients they are most sensitive to, with plenty of overlap between them. Without giving a full description of all relevant processes and operators figure 4.4 shows a clean overview of the constraining potential each category has to offer. In total 34 operators from the Warsaw basis are considered, of which a subset of 20 are kept when focusing on top-specific processes. Care is taken to ensure only statistically independent measurements of the same observable are used. For LHC experiments this comes down to selecting just one measurement per final state, per run, from both CMS and ATLAS. Any correlations are taken into account through the covariance or correlation matrices, assuming they are provided. All forms of data are converted to a format corresponding to a signal strength as it is used in a cross section measurement. These various signal strengths can then be parametrized as a function of the Wilson coefficients as described earlier.

While the authors mention that the inclusion of quadratic terms is possible within their framework, the documented interpretation is performed using a parametrization in which only linear terms are considered. This means that the quadratic and EFT interference terms (cfr. equation 4.21) are dropped, so that only a single structure constant needs to be found for every added operator. These linear structure constants are also completely independent from each other. Because of this a single EFT sample generated for each operator suffices to obtain the full parametrization, each time setting one Wilson coefficient to a non-zero value. It should be noted, as the authors also explicitly mention, that neglecting quadratic terms is not always justified. Most notably for the top-related operators, including the c_{tZ} operator considered in this work, the quadratic terms are actually dominant with respect to the linear ones. Constraints obtained for these operators are therefore automatically on the conservative side, and should be considered as provisional. Nevertheless including them is important when performing a simultaneous fit, as allowing their coefficients to vary affects the CL intervals for other operators.¹⁵ Including quadratic (and interference) terms would significantly increase the number of simulations to be performed from 35 (SM case + 34 EFT scenarios) to 630 (adding N quadratic terms and $\frac{1}{2}N(N-1)$ interference terms as explained earlier, with $N = 34$). It would also imply terms suppressed by a factor of Λ^4 are considered to be relevant, at which point the linear terms generated by dimension-8 operators would also have to be considered.

¹⁴This is not limited to the numbers and figures actually given in the paper document. For many analyses additional material including tabulated data and correlation matrices is also made publicly accessible. For LHC experiments it has become the norm for this information to be published to an online platform called hepData [98].

¹⁵To be precise, it makes these intervals more conservative. In the context of a global analysis which aims to make minimal assumptions about the values of coefficients it seems fair to call these more conservative results more correct as well.




FIGURE 4.4: The 34 Wilson coefficients considered in the global fit paper, and the types of processes that can be used to constrain them [99].


The results of this analysis, as one would expect, come in the form of best fit values for the Wilson coefficients and corresponding CL intervals. These are provided for both the general case of using all inputs and all 34 operators, and a top-specific case in which only measurements involving top quarks are used, and the set of operators is restricted to the ones most relevant to top quark physics. In addition to this, however, the authors note that large levels of correlation are present among the Wilson coefficients, which is shown in the form of a correlation matrix. Using principal component analysis linear combinations of these operators can be constructed which are uncorrelated (eigenvectors). The constraints on these eigenvectors then indicate for which directions in the parameter space no good handles are available yet.

In the spirit of the level of indirect collaboration this global interpretation entails the authors have made the software framework developed for it publicly available.

Chapter 5

Measurement of $t\bar{t}\gamma$ production

The analysis described in this chapter leads to three main results: An inclusive measurement of the $t\bar{t}\gamma$ cross section, several differential cross sections, and an interpretation in the EFT framework. There is of course a large overlap in the main steps required to obtain these three results, so that the first and longest part of the analysis description will be relevant for all of them. In order to have some foresight regarding what is being worked towards the flowcharts below summarize the steps to be taken for each result, with colours indicating where intermediate results from one procedure also serve as input elsewhere.



5.1 $t\bar{t}\gamma$ production and decay

Despite the fact that the top quark is the heaviest particle in the standard model, the high beam energies at the LHC have made the production and detection of large numbers of events involving top quarks possible, earning it the colloquial title of "top factory". This has in the past few years lead to the publication of precision measurements of $t\bar{t}$ production, from which precise constraints on the top quark mass itself can be extracted [100].¹ Additionally, many processes involving top quark pairs produced in association with other particles have been experimentally established and measured, sometimes collectively referred to as $t\bar{t}X$ measurements. It is also in this category that $t\bar{t}\gamma$ production is found: the process is in principle nothing more than the production of a pair of top quarks, in which one of the initial, intermediate, or final state particles radiates a photon. Given the fact that the photon is massless, and can be radiated by any electrically charged particle, the production rate for $t\bar{t}\gamma$ is high compared to other $t\bar{t}X$ processes. The way photons are abundantly radiated does however produce its own set of challenges, which this analysis aimed to mitigate.

What sets the $t\bar{t}\gamma$ process apart from $t\bar{t}$ production is the requirement of a prompt photon being radiated somewhere in the process.² Example production diagrams are shown in figure 5.1. In this analysis photons from various origins are included in the signal definition, including those radiated by the incoming quarks (initial state radiation), and those radiated by one of the top quarks or its decay products. An overview of the different photon origins and the slight differences in their kinematic features is shown in figure 5.2. A type of real photons not included in the signal definition are those radiated inside jets, or in other words those who have a meson in their parentage list. Signal events can also be mimicked when a jet is misidentified as a photon, or when a photon from a pileup vertex is wrongly assigned to the primary vertex. All these cases are classified as nonprompt photon background, and while the main process contributing to this is $t\bar{t}$ any other process can enter the selection in similar ways.

In simulations the ME level $t\bar{t}\gamma$ process can either be defined up to the production of the top quark pair and the photon, or up to the point where the top quarks have already decayed. The latter is more correct, as photons from final state radiation are properly treated in the ME definition, but is far more challenging to implement. Studies of the first approach have been available for a while [101–107], while publications on the more complete option have been appearing more recently [108–110]. The approach used for this work, as a result of current technical limitations, is discussed in the next section. Modifications of the top quark couplings have been studied in the context of anomalous couplings [83, 111, 112], and in the framework of EFT [113]. While no such studies are performed here, the possibility of studying the charge asymmetry in $t\bar{t}\gamma$ events has also been investigated [114, 115]. The interest in the top-photon coupling from the theory side, and therefore the $t\bar{t}\gamma$ process, has already been discussed in detail in the introductions to top quarks physics and effective field theory, so that we can move on to the analysis description without further ado.




FIGURE 5.1: Example LO Feynman diagrams for $t\bar{t}\gamma$ production with two leptons in the final state, where the photon is radiated by a top quark (left), by an incoming quark (middle), and by one of the charged decay products of a top quark (right).

¹This is not the only method for measuring the top quark mass, and direct measurements looking at top quark decays are also commonly performed.

²The exact definition of a prompt photon will be given later, but for the most part it aligns with the concept of an isolated photon. The technical definition will rely on a matching procedure to nearby generator particles, where a match to a generator photon originating from leptons, quarks, or bosons will mean that the photon is prompt.




FIGURE 5.2: Photon origins in the $t\bar{t}\gamma$ signal sample within the event selection defined later on. The top left figure shows the fractions of photons produced as initial state radiation (ISR), radiated by the top quarks, by the b quarks, and by the W bosons or their decay products. Since the contribution from b quarks is negligible it is merged with the top quark category in the other three graphs, which show the normalized distributions as functions of the photon transverse momentum (top right), the angle between the photon and the nearest lepton (bottom left) and the angle between the photon and the nearest b quark jet (bottom right).

5.2 Data samples and simulations

5.2.1 Data

The data used for this measurement was collected between 2016 and 2018 and adds up to 138 fb^{-1} of proton-proton collision events at a centre-of-mass energy of $\sqrt{s} = 13\text{ TeV}$. This total period consists of 3 intervals, referred to by their respective years, with each having slightly different running conditions. The integrated luminosity values corresponding to these three periods are 36.33 fb^{-1} , 41.53 fb^{-1} , 59.74 fb^{-1} for 2016, 2017, and 2018, respectively. Simulated samples are therefore also generated for each year individually, and any calibrations and corrections are also provided on a per-year basis. While often only combined results will be shown, the three years were processed and analysed separately.

5.2.2 Signal simulation and normalization

The $t\bar{t}\gamma$ signal process is simulated for each lepton multiplicity channel separately at LO accuracy in QCD using the MADGRAPH5_amc@NLO 2.6.0 [116]. The parton distribution functions used in the simulation correspond to the NNPDF3.1 [117] PDF set. For the dilepton channel the ME level process is defined up to the decay products of the W bosons, i.e. $pp \rightarrow b\ell^+\nu b\ell^-\nu\gamma$. Using this definition photons originating from initial state radiation, the top quarks, the W bosons, or their decay products are included in the ME calculations. Given their central importance in this analysis it is preferable for the photons to be included in the ME calculation, instead of being added in later steps which rely on phenomenological models instead of full calculations. Since high energy photons are important in this

analysis but appear in only a small fraction of $t\bar{t}\gamma$ events separate samples are generated with $p_T(\gamma)$ between 100 and 200 GeV, and with $p_T(\gamma) > 200$ GeV. This ensures that the statistical uncertainties associated with the simulated signal samples are also sufficiently small in the last few bins of the photon p_T distribution, which will turn out to be the most sensitive to EFT effects. The overlap with the normal signal samples is easily removed by requiring that the ME level photon has a p_T falling below this range.

Calculating the $t\bar{t}\gamma$ cross section at NLO is currently not possible using the more correct process definition described earlier ($pp \rightarrow bl^+\nu\bar{b}l^-\nu\gamma$), in which the process is defined up to the decay products of the top quarks in the matrix element. In order to correct its normalization to NLO correction factors, called K-factors, are calculated using MADGRAPH5_aMC@NLO using the simpler $pp \rightarrow t\bar{t}\gamma$ definition. This K-factor is the ratio of the cross sections calculated at NLO and LO, by which the cross section obtained using the more complete process definition can then be scaled.

5.2.3 Background process simulation

Background events from $t\bar{t}$, $t\bar{t}H$, and single top quark t -channel and tW production are produced at NLO accuracy in QCD using POWHEG2 [118–125]. LO samples for $gg \rightarrow ZZ$ production are generated using the MCFM 7.0.1 event generator [126, 127]. Finally, all other background processes are simulated using MADGRAPH5_aMC@NLO, either at LO or NLO accuracy in QCD. For 2016 samples the NNPDF3.0 [128] PDF set is used, while for 2017 and 2018 NNPDF3.1 is used.

For all samples the showering and hadronization step is performed using PYTHIA8 [129], with version 8.226 being used for 2016, and 8.230 for 2017 and 2018. For all 2017 and 2018 samples, and most 2016 samples, the underlying event is modelled using the CP5 tune [130]. In some cases for 2016 no samples with this tune are available, in which case samples generated using the tunes CUETP8M1, CUETP8M2, or CUETP8M2T4 [131–133] are used. In all of the samples the interactions between particles and the detector material are simulated using the GEANT4 software [134], which relies on a detailed description of the detector geometry.

Samples for nearly all processes relevant to CMS analyses are centrally produced, as is the case for the background samples mentioned above. As an analyst I could therefore simply pick out samples as needed, and only needed to verify whether the chosen samples were appropriate for this specific analysis. The signal samples and the corresponding cross sections were generated for the $t\bar{t}\gamma$ analysis specifically by the authors of the lepton+jets analysis, which not only benefited me, but also ensured the combination described later would be possible.

5.2.4 Removing overlap between simulated samples

In the showering step of sample simulation radiated photons are added to events, even if they do not include photons in their ME level definition. This means that $t\bar{t}$ and Drell-Yann (DY) samples respectively contain events that fall under the process definitions of $t\bar{t}\gamma$ and DY. Events for which the photon is part of the process definition are of course better simulated than those in which the photons are merely included as extra radiation, so that for an analysis of this kind using both is required. Simulated events from the $t\bar{t}$ sample are therefore removed if they fall into the phase space covered by the $t\bar{t}\gamma$ sample, which means it contains a photon fulfilling the following requirements:

- falls inside the kinematic range of $p_T > 10$ GeV and $|\eta| < 5$
- does not originate from the decay of a hadronic particle
- is isolated, which in this context means it is separated from any other stable particle by $\Delta R > 0.1$, not including photons and neutrinos³

This corresponds to the generator level definitions and thresholds used in the production of the $t\bar{t}\gamma$ samples. While all $t\bar{t}\gamma$ samples should in principle fall into this phase space, events containing no such photon are removed from the $t\bar{t}\gamma$ samples.

³The word “stable” has an unambiguous definition in generator terms, meaning the particle has a lifetime of longer than 3 ns [135].

5.3 Particles and objects

5.3.1 Leptons

Sets of identification criteria are typically labelled using names like “loose”, “medium”, or “tight”, with tighter criteria providing lower efficiencies but also reduced misidentification rates. These criteria can be more complex than a simple set of requirements, and can for example correspond to a trained multivariate discriminator (MVA). The trade-off between efficiency and purity however always remains and therefore different choices, called “working points”, are almost always available. The set of criteria, or the multivariate discriminator itself in combination with its working point, is called an “ID”, and this convenient wording will be used going forward.

The main purpose of these identification criteria is to distinguish leptons coming from the primary interaction process, referred to as prompt leptons, from those produced later on inside of jets or resulting from an erroneous reconstruction, called nonprompt leptons. Light leptons produced through the decay of prompt τ leptons are also included in this prompt lepton definition, and no attempt is made to distinguish these leptons from those produced directly. In this analysis leptons are selected using an MVA trained to discriminate between prompt and nonprompt leptons, as described in the next subsection. In the event selection described later the presence of at least 2 leptons passing a very loose working point of the MVA will be required. The leading lepton is required to have a transverse momentum of at least 25 GeV, while that of the subleading lepton must exceed 15 GeV. In addition to the ID criteria described above, electrons within $\Delta R < 0.02$ of a muon passing the baseline selection are removed. Both leptons must satisfy $-2.4 < \eta < 2.4$, and electrons with a supercluster $|\eta|$ between 1.4442 and 1.566 are also not used because of the gap in the ECAL in this range.

The MVA used in this analysis is based on the one trained for the tZq analysis [136, 137], with some improvements made and trained on data from all 3 years by UGent colleagues [138]. The idea behind the new training was to produce an MVA that can be used in a wide range of analyses involving top quarks. The discriminator was trained using prompt leptons from simulated tZq , ttZ , and ttW samples, and nonprompt leptons from simulated semileptonic $t\bar{t}$ events. This training is performed independently for muons and electrons, and per data-taking year.⁴ The set of electrons and muons on which this discriminator is trained already passed some loose set of quality criteria, called a baseline selection. Receiver operating characteristic (ROC) curves for both electrons and muons are shown in figure 5.3, and compared to the performance of the discriminator used in the tZq analysis. These figures also illustrate the trade-off between efficiency and purity mentioned earlier. Because of the relatively low contribution from nonprompt leptons in this analysis a very loose working point can be used, with an overall selection efficiency of around 95% for prompt leptons.

Lepton reconstruction and identification efficiency calibrations

Scale factors used to correct for reconstruction efficiency differences between data and simulation for electrons are applied. The corrections are obtained by measuring the performance of electron reconstruction in $Z \rightarrow ee$ events using a tag-and-probe method, which shows efficiencies in data of above 95% in the E_T range from 10 to 500 GeV [69]. The ratio of the efficiencies in data and simulation stay within 2% of unity over this entire range as well. For photons there is no obvious process that can be used to perform such tag-and-probe measurements, but as far as ECAL information is concerned electrons and photons are the same object and these corrections therefore apply to both. For muons a reconstruction efficiency of above 99% is found, so that no correction factors need to be applied and any small deviations from 100% efficiency should be covered by other uncertainties applied for the identification of muons.

The performance of the MVA is slightly different between data and simulation, so that correction factors need to be applied. These scale factors are obtained by measuring the selection efficiencies using a tag-and-probe method in DY events from both data and simulation. Ignoring the many details which go into this calibration the scale factors are then obtained by taking the ratio between the two. Two types of uncertainties are associated with these corrections: Statistical uncertainties

⁴In principle one could say there are 6 separate MVAs being trained and used, but for convenience they are typically referred to in the singular form.




FIGURE 5.3: ROC curves showing the performance of the MVA discriminator used for the lepton selection in this analysis. The performance is shown for electrons (left) and muons (right), for values of the lepton $p_T > 25$ GeV (top), and $p_T < 25$ GeV (bottom), and compared to the older iteration of the MVA. The points along the curves corresponding to the very loose working points defined for both discriminators are indicated using coloured dots. This example was produced using a $t\bar{t}\gamma$ sample for 2016.

due to the finite sample sizes, and systematic uncertainties related to e.g. model choices for the signal and backgrounds. These corrections are binned as a function of the p_T and $|\eta|$ of the leptons.

ECAL scale and resolution corrections

For electrons and photons corrections for the energy scale and resolution are applied using recipes provided for use in all CMS analyses [139]. The scale corrections compensate for the inefficiency inherent to any energy measurement, and are applied in data and simulation to make them agree in this regard. In simulations the resolution is intentionally made slightly better than in data, so that the signals can later be smeared to match the lower resolution observed in data. Since scale correction variations are also provided for the simulated samples, where the up variation in simulation corresponds to the down variation in data and vice-versa, the systematic uncertainties associated with both the scale and resolution uncertainties can be assessed by varying the simulated predictions.

Muon Rochester corrections

Rochester corrections are applied to the muons used in the analysis, as described in [140]. These corrections are meant to compensate for misalignments of the CMS detector, which are not present in the idealized simulated detector. The impact of this correction is small for muons with a p_T below 200 GeV, which holds for nearly all muons in this analysis. The uncertainty associated with this correction was found to be negligible, and was therefore not taken into account.

5.3.2 Jets

Particle candidates found by the PF algorithm are clustered into jets using the anti- k_T algorithm [71] with distance parameter $R = 0.4$. The influence of pileup is mitigated by the Charged Hadron Subtraction technique which removes charged particle candidates which are matched to any vertex other than the main primary vertex. Jets are calibrated in both data and simulation to account for energy deposits from pileup and the imperfect detector response. Only jets with $p_T > 30$ GeV and $|\eta| < 2.4$ are used and are additionally required to pass a set of standard jet identification criteria [141], sometimes also referred to as the Tight jet ID. This Tight ID has an efficiency of above 99% everywhere so no looser set of identification criteria needs to be considered. A selected jet may still overlap with the selected leptons, as these can be clustered into jets as well. To prevent such cases, jets that are found within a cone of $R = 0.4$ around any of the selected signal leptons are removed from the set of selected jets. Similarly, jets are removed from the event when they are found in a cone of $R = 0.1$ around the selected photon. Jets originating from the hadronization of b quarks are identified using the DeepCSV discriminator [142]. The DeepCSV algorithm combines information from track impact parameters and secondary vertices identified within a given jet. In this analysis, a jet is identified as originating from a b hadron when its DeepCSV value passes a medium working point, chosen to result in a balance between efficiency (about 70%) and low misidentification rates (12% for jets from c quarks, and 1% for jets from light quarks or gluons). Scale factors are applied to all simulated samples to account for discrepancies in efficiency and purity with respect to data, and this for each of the jets in the event. These scale factors were determined for the specific working point as a function of the jet p_T and η . The tagging efficiency for b quark jets, c quark jets, and light jets in simulation needed to apply these scale factors is measured in $t\bar{t}$ events with the same lepton and photon selection applied as used in the rest of the analysis.

The energy of the jets in both data and simulations are calibrated using a series of corrections collectively referred to as jet energy scale corrections, the aim of which is to align the reconstructed jet energy with the true jet energy. In a first step the additional energy clustered into the jets because of soft jet activity coming from the underlying event, detector noise, and pileup is removed. To do this the area corresponding to the jet cone and a pseudorapidity dependent estimate map of this soft jet energy is used, which is calculated on an event-by-event basis due to the pileup dependence of this map. In simulated events the jet clustering algorithm is then applied to generator level particles excluding neutrinos, and the resulting objects are matched to their corresponding reconstructed jets. Based on these matched objects the detector response can then be calculated as a function of the pseudorapidity and transverse momentum of the jets, the inverse of which, when applied to reconstructed jets, corrects their energy to match the true energy of the jet. As this correction is however accounting for particles which were not really measured it is more correct to say that the reconstructed energy is corrected to the true jet energy on average. These corrections derived in simulated events are applied both in simulation and in data, as there is no way to measure this true jet energy in data. Final residual corrections make use of data-driven techniques in well understood events with clean signatures. One option is the use of Z +jet events, in which the leptonically decaying Z boson can be reconstructed with high precision, and no E_T^{miss} is intrinsically present. Any E_T^{miss} can therefore be assumed to be caused by a mismeasurement of the recoil of the jet, and can therefore be used to correct the measured energy. Additional calibrations make use of events in which only two jets are produced, with the requirement that they are produced back-to-back. Once again any energy imbalance in the transverse plane can only be caused by a mismeasurement of the jet energies, so that tuning the jet energies to make their vector sum vanish leads to a correction.

5.3.3 Photons

Photons with $p_T > 20$ GeV are selected using a set of photon identification requirements with an average efficiency of about 80% [143], corresponding to a medium ID. Such a “cut-based” ID consists of a relatively simple set of requirements on the object features, in contrast to more complex multivariate discriminators. Such multivariate identification algorithms are available, and offer slightly better performance, but do not allow for the flexibility of removing or inverting one or more of the requirements as will be done later on. Originally photons with $p_T > 15$ GeV were considered, but this threshold was later increased as the 15 to 20 GeV range is much more background dominated. The set of photon identification requirements for the barrel region of the detector is described in table 5.1 and involves the following variables:

TABLE 5.1: The set of photon identification requirements, corresponding to the medium photon ID and the additional pixel seed veto.

	barrel
$H/E <$	0.02197
$\sigma_{i\eta i\eta} <$	0.01015
charged isolation (ρ -corrected) $<$	1.141
neutral isolation (ρ -corrected) $<$	$1.189 + 0.01512 * p_T + 2.259 * 10^{-5} * p_T^2$
photon isolation (ρ -corrected) $<$	$2.08 + 0.004017 * p_T$
pixel seed veto	yes

- The **ratio of the hadronic energy to the ECAL supercluster energy** within a cone radius of $\Delta R < 0.15$, denoted as H/E .
- A variable parametrizing the **shower shape of the ECAL cluster** is computed with logarithmic weights as:

$$\sigma_{i\eta i\eta}^2 = \frac{\sum_i^{5 \times 5} w_i (\eta_i - \bar{\eta}_{5 \times 5})^2}{\sum_i^{5 \times 5} w_i}, w_i = \max \left(0, 4.7 + \ln \frac{E_i}{E_{5 \times 5}} \right) \quad (5.1)$$

where $E_{5 \times 5}$ and $\bar{\eta}_{5 \times 5}$ are the energy and pseudorapidity of the entire 5×5 matrix around the seed of the cluster, and E_i and η_i are the energy and pseudorapidity of the i th crystal within the 5×5 matrix. This variable is mostly useful for rejecting neutral pions decaying into a photon pair. The signature of such a decay is a single wide jet since the two photon showers tend to overlap, so that imposing a well-chosen upper limit on the shower width efficiently removes them.

- The PF **photon isolation** defined as the transverse energy sum of all particles identified as photons by the PF algorithm within $\Delta R < 0.3$. An effective-area based ρ correction is applied to take into account the contribution of pileup. This method is used in all isolation variables, and comes down to estimating the energy inside of the shower cone coming from pileup and the underlying event before subtracting it. This estimate uses a pileup dependent map of the average energy from these sources, parametrized as a function of the pseudorapidity factorized by the area corresponding to the cone of the shower.
- The PF **charged isolation** is the transverse energy sum of all charged hadrons identified by the PF algorithm within $\Delta R < 0.3$ and compatible with the primary vertex. An effective-area based ρ correction is applied to take into account the contribution of pileup.
- The PF **neutral isolation** is the transverse energy sum of all neutral hadrons identified by the PF algorithm within $\Delta R < 0.3$. The amount of energy in the cone due to pileup is estimated and subtracted, once again using an effective-area based ρ correction.

In addition to the medium ID, the **pixel seed veto** requirement is added which ensures there is no track seed identified in the pixel detector. A pixel veto is typically used in analyses which are sensitive to an electron to photon fake rate. Even though the measurement is not sensitive to electrons misidentified as photons, adding the pixel seed veto seems to be a good handle to reduce the hadronic and fake photons (defined later on in section 5.6.2), rejecting up to 35% of $t\bar{t}$ events in the signal regions while $t\bar{t}\gamma$ is only affected by about 5%.

Photons are also required to be separated by $\Delta R > 0.4$ from the selected leptons. Only photons in the ECAL barrel ($|\eta| < 1.4442$) are being considered for this measurement. Both signal efficiency and statistics are lower for events with photons in the endcap region. Endcap photons also feature differently distributed shower shape variables, so that an entirely separate nonprompt photon background estimate would be required.⁵ Corrections for small discrepancies between data and simulation of the efficiency of the photon identification criteria and the pixel seed veto are applied in the form of scale factors, parametrized as a function of the photon p_T and $|\eta|$ in simulated events. Similar to electrons, ECAL scale and smearing corrections are also applied to the photons.

⁵The very limited statistics in the endcap regions also make it doubtful that the same estimation method could even be used.

TABLE 5.2: The summary of the baseline selection in each channel. Note that the last two selections are applied only to the ee and $\mu\mu$ channels.

Variable	Requirement
$p_T^{\text{leading}}(l)$	$> 25 \text{ GeV}$
$p_T^{\text{trailing}}(l)$	$> 15 \text{ GeV}$
N_{leptons}	$= 2$
Charge sum of the leptons	$= 0$
$p_T(\gamma)$	$> 20 \text{ GeV}$
N_γ	$= 1$
N_γ looser ID	$= 1$
N_b	≥ 1
$m_{\ell\ell}$	$> 20 \text{ GeV}$
$ m_{\ell\ell} - m_Z $	$> 15 \text{ GeV}$ (SF channels)
$ m_{\ell\ell\gamma} - m_Z $	$> 15 \text{ GeV}$ (SF channels)

As will be discussed in section 5.6.3, the charged isolation and $\sigma_{\text{in}}\eta$ requirements will be temporarily relaxed for the purpose of estimating the nonprompt photon background. Other requirements will be kept as in the medium ID. In other parts of the analysis, unless stated otherwise, the full medium ID is used.

5.4 Event selection

Once electrons, muons, photons and jets are identified in a collision, a selection based on the presence of certain numbers of these object and their kinematics can be defined, aimed at defining a signal region highly enriched in $t\bar{t}\gamma$ signal events. An event level selection can however still include further requirements on features of individual objects. In such an event level selection a balance needs to be found between reducing the the contributions from background processes, while removing as few signal events as possible.

Exactly one photon passing the ID detailed in the photon selection description is required. Events containing additional photons passing a looser ID are vetoed. This looser ID is obtained by relaxing the $\sigma_{\text{in}}\eta$ and charged hadron isolation requirements from the medium ID. Only a negligible fraction of events is removed in this way, but this ensures that an event cannot contribute to multiple regions in the nonprompt photon background estimation.

Two leptons (ee , $\mu\mu$ or $e\mu$) with opposite charges are required, with a minimum dilepton invariant mass $m_{\ell\ell}$ of 20 GeV. This invariant mass requirement successfully rejects all events from QCD multi-jet processes. The leading lepton is required to pass a transverse momentum threshold of $p_T > 25 \text{ GeV}$ to avoid a bias from the online selection, while the subleading lepton only has to pass $p_T > 15 \text{ GeV}$ as defined in the lepton selection. Once the leptons have been selected channel-dependent triggers are also applied, which will be explained in the next section. For same-flavour (SF) dilepton events (ee or $\mu\mu$), events with an $m_{\ell\ell}$ value falling within a 15 GeV range from the Z boson mass are rejected in order to avoid Drell-Yan background contamination. Furthermore, also for SF dilepton events, a similar Z boson mass window is applied on the dilepton plus photon invariant mass, i.e. $|m_{\ell\ell\gamma} - m_Z| > 15 \text{ GeV}$, which rejects a large part of the $Z\gamma$ background. Finally, at least one b-tagged jet is also required, identified using the medium working point of the DeepCSV tagger mentioned earlier.

The baseline selection described in the paragraph above is summarized in table 5.2. With the event selection applied the $t\bar{t}\gamma$ signal process becomes the dominant contribution to the total yield. This is illustrated in figure 5.4 where the various contributions are shown as a function of the lepton channel, and the jet and b jet multiplicities. A priori twice as much $t\bar{t}\gamma$ events are expected in the $e\mu$ channel compared to the ee and $\mu\mu$ channels.⁶ This difference is somewhat enlarged because of the additional invariant mass requirements in the SF channels. The naturally high purity of the $e\mu$ channel is immediately obvious. The only significant background there comes from nonprompt photons, which

⁶This is simply because of the combinatorics of the decay options, or in other words because the $e\mu$ final state contains both $e^+\mu^-$ and $e^-\mu^+$ events.

predominantly stems from $t\bar{t}$ events and is therefore hard to disentangle from the signal. This high purity explains why other measurements of $t\bar{t}\gamma$ in the dilepton channel choose to neglect the other final states. As the final results will show however, in this analysis the SF channels do contribute to the reduction of the total uncertainty, thanks to careful control of the backgrounds in these regions.




FIGURE 5.4: The observed (points) and predicted (shaded histograms) signal and background yields as functions of the lepton flavours (top), the number of jets (bottom left) and number of b-tagged jets (bottom right). Distributions are shown after the event selection for the three lepton flavour channels combined, with all relevant corrections applied. The vertical bars on the points show the statistical uncertainty in the data, and the hatched bands the systematic uncertainty in the predictions. The lower panels show the ratio of the event yields in data to the sum of the predictions.

Distributions for various kinematic variables are shown in figures 5.5 and 5.6. These include the p_T and absolute value of the pseudorapidity of the selected photon, the p_T of the leading jet $p_T(j_1)$, the angular separation between the selected photon and the closest selected lepton, jet, or b jet (respectively $\Delta R(\gamma, l)$, $\Delta R(\gamma, j)$, and $\Delta R(\gamma, b)$), the absolute value of the pseudorapidity difference between the two selected leptons $|\Delta\eta(l\bar{l})|$, and the azimuthal angle difference between the two selected leptons $\Delta\varphi(l\bar{l})$. Also shown are the transverse momentum component of the four-vector sum of the two selected leptons $p_T(l\bar{l})$, and the linear sum of the p_T of the leading and subleading leptons $p_T(l_1) + p_T(l_2)$. It is these variables which will later be considered for the differential cross section measurements. Differential measurements of these variables can be interesting from a theory perspective, as for example the $\Delta\varphi(l\bar{l})$ variable carries information regarding the spin correlation between the top quarks [144, 145]. A potential future addition would be the inclusion of variables related to the reconstructed top quarks, which would also allow for measurements of the charge asymmetry. Since the momentum carried by photons in $t\bar{t}\gamma$ events is substantial this would benefit from a top quark reconstruction procedure which explicitly takes this radiation into account, like the one documented in Ref. [11], which was developed alongside this $t\bar{t}\gamma$ measurement but in the end not applied in it.

Some basic insights can be extracted from the first two graphs: Simulations predict harder photon p_T spectra for the prompt photons from $t\bar{t}\gamma$ and $Z\gamma$ than for the nonprompt photons mostly originating from the $t\bar{t}$ and Drell-Yan backgrounds (included in the nonprompt γ category in the figures). As a result, the $t\bar{t}\gamma$ signal has a high purity for high values of the photon p_T . The $|\eta|$ distribution shows that most $t\bar{t}\gamma$ events will have their photon produced in the barrel region of the detector. Nonprompt photons, mostly originating from $t\bar{t}$ events, have a flatter $|\eta|$ distribution leading to a lower signal purity for endcap photons. This is another reason for only including barrel photons in the object selection. Two graphs showing the signal and background yields for events with photons in the endcaps are shown in figure 5.7. The exact distinction between prompt and nonprompt photons, and therefore the difference between contributions from $t\bar{t}$ and the signal is specified in section 5.6.2 where the various types of nonprompt photons are also explained.

5.5 Trigger selection and efficiency measurement

Data collected by the CMS detector is streamed into different primary data sets (PD), based on whether events pass at least one of a set corresponding HLT requirements. Some of these PDs and their corresponding triggers are based on the presence of certain objects, e.g. triggers for **DoubleMuon** data sets require the presence of at least two muons with certain kinematic properties. Other PDs require for example a minimum amount of missing transverse energy, as is the case for the **MET** PD. Sometimes the trigger thresholds of a trigger are so low that the resulting rate of events passing it would not be possible to store. In this case it is possible to randomly save only $1/N$ of the events passing the trigger, in which case N is called the prescaling which can be used to estimate the true number of events from the recorded value. For this measurement however these triggers are avoided, and only unprescaled triggers are used.

Since all three final states (ee , $\mu\mu$ and $e\mu$) are used, overlap also needs to be avoided between them. This is achieved by first performing the lepton selection, so that the lepton final state an event will be categorized under is fixed, and then applying the corresponding trigger requirements. A separate technicality comes from the fact that triggers requiring either one or two leptons are used, and accordingly both single lepton and 2-leptons PDs are used. To avoid introducing overlap when combining these data sets events passing 2-lepton triggers are first taken, after which events from single lepton PDs are added only if they pass single-lepton triggers and fail 2-lepton triggers. For the $e\mu$ channel events from both the single- e and single- μ PDs need to be added to the $e\mu$ PD, so that a similar procedure needs to be used for each data set being added.

A list of the full set of triggers is of little use here, but the important point is that the p_T requirements of the triggers used are slightly lower than those used in the object and event selection. The kinematic features of reconstructed particles used by the HLTs are slightly different from those used at the analysis level due to the more complete offline reconstruction procedure and the application of corrections after the trigger step. Kinematic thresholds fully aligned with the HLT level requirements would therefore lead to large inefficiencies just above those thresholds.

Trigger efficiencies have been measured in both data and simulated $t\bar{t}$ samples, for events which pass the dilepton selection used in this analysis (two oppositely charged leptons, selected following the criteria discussed in section 5.3.1). The ratio of the trigger efficiencies in data and simulation, measured as a function of the p_T of the two leptons, provide data/simulation scale factors which are applied throughout the analysis. Similar to other CMS analyses which employ a $t\bar{t}$ dilepton topology [146–148], the trigger efficiencies in data are measured using the unbiased **MET** and **JetHT** primary data sets. For each year a set of triggers is selected which are uncorrelated with lepton triggers (i.e. only pure E_T^{miss} and H_T -related triggers which do not contain lepton vetoes) and are unprescaled throughout the whole data taking period. This means that the data set obtained by applying said triggers is unbiased when it comes to the performance of the lepton triggers, and can therefore be used to measure their performance.

The assumption of the lepton and selected **MET** and **JetHT** triggers being uncorrelated has been explicitly checked using the correlation ratio α defined [146] as:

$$\alpha = \frac{\epsilon_{\text{lep}} \times \epsilon_{\text{ref}}}{\epsilon_{\text{ref,lep}}} \quad (5.2)$$




FIGURE 5.5: The observed (points) and predicted (shaded histograms) signal and background yields as functions of the kinematic variables explained in the text. Distributions are shown after the event selection for the three lepton flavour channels combined, with all relevant corrections applied. The vertical bars on the points show the statistical uncertainty in the data, and the hatched bands the systematic uncertainty in the predictions. The lower panels show the ratio of the event yields in data to the sum of the predictions.

where ϵ_{lep} is the efficiency of the lepton triggers, ϵ_{ref} is the efficiency of the reference triggers and $\epsilon_{\text{ref,lep}}$ the efficiency of the logical AND combination of the lepton and reference triggers. If the reference and lepton triggers are perfectly uncorrelated, the correlation ratio α is equal to 1. The values for α have been measured for all channels in the $t\bar{t}$ and $t\bar{t}\gamma$ simulated samples, and its deviation from 1 is found to be small for all cases.

In the end the trigger efficiencies measured in the MET and JetHT data sets and a $t\bar{t}$ dilepton sample are used to calculate the actual scale factors used in the analysis. The averages of the correction factors obtained in the MET and JetHT data sets are used in the analysis, and the corresponding systematic




FIGURE 5.6: The observed (points) and predicted (shaded histograms) signal and background yields as functions of the kinematic variables explained in the text. Distributions are shown after the event selection for the three lepton flavour channels combined, with all relevant corrections applied. The vertical bars on the points show the statistical uncertainty in the data, and the hatched bands the systematic uncertainty in the predictions. The lower panels show the ratio of the event yields in data to the sum of the predictions.

uncertainty is estimated by comparing the total efficiency in each data set to their average. This comparison is performed for the different channels separately, the result of which is shown in table 5.3 for the 3 years. Of course these values also simply serve to illustrate the overall trigger efficiency within the lepton selection of the analysis. These differences are assigned as an extra uncertainty in the trigger scale factors, correlated between the channels. The statistical uncertainty in the correction factors is propagated as well, but uncorrelated between the channels.




FIGURE 5.7: The observed (points) and predicted (shaded histograms) signal and background yields as functions of the lepton channel and the photon p_T for events in which the selected photon is reconstructed in the endcap region of the ECAL ($|\eta| > 1.566$). All other requirements are the same as in the regular event selection. Compared to the signal region a significantly larger background coming from nonprompt photons is found.

TABLE 5.3: Integrated trigger efficiencies for the **MET** and **JetHT** data sets, and the difference between them which is used as an estimate of the systematic uncertainty in the trigger scale factors.

		ee	$e\mu$	$\mu\mu$
2016	JetHT	95.8%	96.5%	97.6%
	MET	94.1%	95.8%	96.8%
	average	94.9%	96.1%	97.2%
	difference	0.8%	0.4%	0.4%
2017	JetHT	96.0%	96.3%	96.9%
	MET	93.0%	94.8%	94.2%
	average	94.5%	95.5%	95.5%
	difference	1.5%	0.8%	1.4%
2018	JetHT	96.6%	96.7%	97.2%
	MET	94.1%	95.9%	96.9%
	average	95.3%	96.3%	97.1%
	difference	1.3%	0.4%	0.1%

5.6 Background prediction

In principle both the signal and the various backgrounds are predicted using MC simulation. Sometimes however it is necessary to correct these predictions, when for example a shortcoming of the sample used is known, or some discrepancy with the data is observed. This can of course not be checked in the signal region, but must be looked at in a region enriched in the relevant process. The distinction should also be made between backgrounds corresponding to specific processes, like $Z\gamma$ or single top production, and backgrounds defined based on the presence of a certain type of object. In this analysis the relevant object is that of the nonprompt photon, which can appear in any of the background processes, but is best considered as a category of its own.

First, I will describe the data based correction applied to the $Z\gamma$ background, is a priori predicted using simulations. I will then describe the nonprompt photon background, and the method that is used to estimate it directly from data. The reason for not simply taking this background from simulation is not the observation of a discrepancy between data and simulation, but the knowledge that the objects in this category are potentially not well simulated. While such data-driven estimations have their own shortcomings, which is reflected in the uncertainties assigned to them, the predictions they provide are more trustworthy.




FIGURE 5.8: The $m(l\bar{l}\gamma)$ and $m(l\bar{l})$ distributions in the SF channels of the signal region, with the requirements on these variables removed. Grey lines indicate range requirements applied to these variables in the signal region event selection. The $Z\gamma$ process peaks in both distributions around the Z boson mass (approximately 91 GeV). Most of the events falling outside of these peaks in either distribution correspond to the bulk of the peak in the other distribution, so that applying range requirements on both variables effectively removes $Z\gamma$ contributions.

5.6.1 $Z\gamma$ background

As could be seen in figure 5.4 on page 64 the $Z\gamma$ process, while being almost completely absent in the $e\mu$ channel, forms the dominant prompt photon background source in the SF channels. Figure 5.8 shows the $m(l\bar{l}\gamma)$ and $m(l\bar{l})$ distributions in the SF channels of the signal region with the requirements on these variables relaxed, illustrating how large this background would be if not removed in this way. A region highly enriched in $Z\gamma$ events can be created by inverting the $m(l\bar{l}\gamma)$ selection requirement used to define the signal region. By keeping $m(l\bar{l})$ off-Z like in the signal region the contribution from Drell-Yan events, which would enter the selection through nonprompt photons, is kept low. As can be seen in figure 5.10 the N_j/N_b shape is not well modelled for $Z\gamma$. This control region is therefore used to determine bin-by-bin correction factors for $Z\gamma$, which are then applied in the signal region to correct both the $Z\gamma$ shape and yield. This imperfect modelling is not surprising as the simulated $Z\gamma$ sample includes at most 1 jet in the ME calculation. When deriving these corrections the SF channels and the 3 years are combined in order to reduce the statistical uncertainty in the correction factors. These correction factors are also shown in figure 5.10, where the lower b jet multiplicity bins are included to illustrate that they mainly deviate from one for higher jet multiplicities. While there are nearly no $Z\gamma$ events with 3 jets in the signal region, the $\geq 3j, 2b$ and $\geq 3j, \geq 3b$ bins in this map are merged to limit the associated statistical uncertainty.

Various potential systematic effects on the validity of these scale factors were investigated. For example due to the orthogonality between the signal region and the $Z\gamma$ control region the ratio between t-channel and s-channel production of $Z\gamma$ events is expected to be somewhat different in the signal region compared to the control region, which can be confirmed using a generator level check of the photon origins. These channels could alternatively also be referred to by where the photon is produced: In the s-channel diagram the photon is produced as final state radiation (FSR), while in the t-channel it comes in the form of initial state radiation (ISR). Diagrams illustrating these production channels are shown in figure 5.9. A comparison between the various photon origin contributions is shown in figure 5.11. In the $Z\gamma$ control region there is almost exclusively s-channel production, whereas in the signal region about 70% s-channel and 30% t-channel production is found. As will be discussed later in the section on systematic uncertainties one would normally assign a normalization uncertainty of 5% to $Z\gamma$ [149], if the prediction were coming purely from MC simulation. On the other hand when the yield of a process is directly corrected using data one would typically drop this normalization uncertainty entirely, and instead assign an uncertainty related to the extrapolation from the control region to the signal region. Instead of following this approach here the $Z\gamma$ yield is simply not considered to be fully constrained, and the normalization uncertainty is not removed. Instead, 30% of the input normalization uncertainty is kept, and only the 70% corresponding to s-channel production is considered to be constrained, resulting in a remaining 1.5% uncertainty in the $Z\gamma$ yield.




FIGURE 5.9: Diagrams illustrating the two leading production mechanisms for $Z\gamma$.

Additional uncertainties are assigned to the $Z\gamma$ background to take into account the statistical uncertainties in these correction factors. To propagate these to the final result, taking into account that they are statistically independent, the correction factors are each varied independently within their uncertainty, leading to 5 separate nuisances (the 0 b jet bins are not relevant to the signal region). Additionally, the fact that there is a small amount of signal present in the $Z\gamma$ control region is taken into account. Within the $N_b > 0$ selection relevant to the signal region this accounts for 9.7% of the total yield. Given the theory uncertainty of 18.5% on the prediction of the signal yield (as will be explained in the results section), a flat uncertainty of 1.8% is assigned to the $Z\gamma$ yield in the signal region.⁷ This flat uncertainty is independent of the $Z\gamma$ normalization uncertainty mentioned earlier.

While the photon kinematics of these two fractions are significantly different the effect of the correction procedure on the overall kinematics in the signal region is limited. To illustrate this, the photon p_T spectrum with and without corrections is shown in the right-hand graph of figure 5.12, but the same was checked for various other variables. The corrections for the ee and $\mu\mu$ channel are found to be consistent as shown in the left-hand graph of figure 5.12, so that they can be combined without issue.

This control region will not be included in the final fit used to extract the signal strength. Doing so would be a viable alternative to the method described here relying on the derivation and application of correction factors. Since it is however the N_j/N_b shape of $Z\gamma$, and not just the total yield being corrected it would be necessary to provide separate histograms for each of the N_j/N_b bins to the fit, and this both for the $Z\gamma$ control region and the signal region. The correlations of systematic uncertainties between the signal and control region would however automatically be treated correctly, and would potentially allow the control region to constrain uncertainties which predominantly affect $Z\gamma$. The issue of the main production channels differing between the signal and control regions would however remain with this alternative method, so that shifting this complexity to the fitting step was deemed unnecessary, while also being impractical.

5.6.2 Nonprompt photon categories

While photons are generally simply referred to as being either prompt or nonprompt, the nonprompt contribution is made up of photons, both real and fake, from various sources. Since the nonprompt photon category does indeed contain real photons the distinction between prompt and nonprompt photons should also be made very clear. The categorization is based on a matching procedure at the generator level, where a reconstructed photon is matched to the nearest generator particle with a p_T within 50% of that of the reconstructed photon. The matching fails if no such generator particle is found within a cone of radius 0.3 around the photon. If a matched particle is found to be a generator

⁷Where a flat uncertainty means an uncertainty that has the same percentual size in all bins of any given distribution.




FIGURE 5.10: Data versus simulation comparison in the $Z\gamma$ control region as a function of the number of jets and b quark jets (left), and the correction factors derived from it (right).




FIGURE 5.11: Photon origins in the signal region (left) and the $Z\gamma$ control region (right) in a $Z\gamma$ sample. The “other” category consists mostly of quarks. Photons categorized as being radiated from Z bosons are in reality radiated by one of the leptons it decays into.

photon originating from either a lepton, quark, or boson it is labelled as a prompt photon. In all other cases the photon is nonprompt, in which case four subcategories can be distinguished as described below. The $\sigma_{i\eta i\eta}$ and charged isolation values referred to in these descriptions are illustrated in figure 5.13.

- Hadronic photons:

Photons of this kind are real, but are produced inside a jet and therefore not part of the signal definition. The photon is either matched to a generator photon with a hadronic origin, or the matching procedure failed but both a generator photon and neutral meson are found within the $\Delta R < 0.3$ cone around the photon. This second option catches cases where the matching procedure would fail based on the p_T requirement, and the photon would otherwise be mistakenly be put into the next category. In other analyses the term “fragmentation photon” is sometimes used. Since these are real photons they typically have low values of $\sigma_{i\eta i\eta}$, but feature high charged isolation values as they are surrounded by the charged particles that are part of the jet.

- Misidentified electron:

This category consists of electrons wrongly identified as photons, and is therefore based on the reconstructed photon being matched to a generator electron. This category is negligible in the dilepton channel, and is further reduced by applying the pixel seed veto. In terms of shower width ($\sigma_{i\eta i\eta}$) these objects are very similar to prompt photons, but typically feature higher charged isolation values.




FIGURE 5.12: Comparison between the correction factors derived in the ee and $\mu\mu$ channels (left). Comparison between the photon p_T spectra in the signal and control regions, with and without the corrections applied (right).

- Fake photon:

In this case the object identified as a photon is in reality a hadronic jet. This jet may contain fragmentation photons itself, but thanks to the p_T requirement in the matching procedure the reconstructed photon is not matched to them. A photon is labelled as fake if it does not fall into any of the aforementioned categories. These objects are not real photons, and since they are in reality hadronic jets their $\sigma_{i\eta i\eta}$ and charged isolation values tend to take on larger values.

- Pileup photon:

This category consists of photons with no nearby generator particles, which also means there is no generator particle to check the origin of. Because of this peculiar feature, seemingly implying they appear out of nowhere, they were initially referred to as “magic photons” among some of the people involved in the analysis. As will be illustrated here there is good reason to believe that these photons in reality originate from secondary interactions, or in other words correspond to pileup events. There is also no guarantee these photons are all real, as they can correspond to any prompt or nonprompt photon type, but originating from a different interaction vertex. Such wrong assignments are relatively common due to the fact that no track information is available for photons. In this case the matching procedure fails, and no generator particle is found within the $\Delta R < 0.3$ cone around the photon.⁸

These photons exhibit low charged isolation values, and seem to feature a mixture of both low and high values of the shower shape $\sigma_{i\eta i\eta}$. It however turns out that these photons are only isolated with respect to the primary vertex, from which they do not originate. This can be seen in the right graph of figure 5.14, where a variable is shown with a similar definition to the charged hadron isolation, with one key difference. Instead of simply calculating the charged isolation with respect to the primary vertex, it is calculated for every reconstructed vertex and the maximum is taken. This essentially makes it the maximal/“worst case” charged isolation of the photon and since pileup photons are likely to correspond to a hadronic or fake photon from a different vertex this variable takes on larger values for these objects than for other photon categories. As shown in the left graph of figure 5.14 they also feature the lowest transverse momentum values of any of the categories. The jets in these secondary interactions are presumably softer than those associated with the primary vertex, which explains why pileup photons have

⁸Typically in such a check only stable generator particles are considered. Here however all particles, even decayed ones, are considered as true pileup photons are so distinctly isolated from any other particles. This detail was however found to not make any noticeable difference.




FIGURE 5.13: The charged isolation (left) and $\sigma_{i\eta i\eta}$ (right) distributions for the various prompt and nonprompt photon categories, each normalized to unity. The identification criteria making up the photon ID are applied except for the requirement on the variable being shown each of the graphs.




FIGURE 5.14: The transverse momentum (left) and maximal charged isolation (right) distributions for the various prompt and nonprompt photon categories, each normalized to unity. The maximal charged isolation variable is explained in the text and can to some degree be used to distinguish pileup photons from prompt photons.

even lower momenta than those in the hadronic category. while this feature is too interesting not to mention, this additional variable was in the end not used or studied further.

Where these categories are found in the 2D parameters space is also illustrated in the left-hand diagram of figure 5.15. Finally, the relative yields of the nonprompt photon categories within the signal region are shown in the right-hand graph of figure 5.15, where the overwhelmingly dominant contribution of prompt photons is left out.

These definitions align with the distinction between $t\bar{t}\gamma$ and $t\bar{t}$ events. Thanks to the overlap removal procedure described earlier all signal events, meaning $t\bar{t}$ events with a prompt photon, are found to come from the $t\bar{t}\gamma$ samples. All events from $t\bar{t}$ samples left after overlap removal fall into one of the nonprompt photon categories. $t\bar{t}\gamma$ does however contribute a very small number of events to the nonprompt photon background as well, through events in which the prompt photon is missed and some other nonprompt photon is selected instead.

5.6.3 Estimation of nonprompt photons

The background formed by nonprompt photons from all processes accounts for only about 9% of the total yield in the signal region. Nevertheless, the various types of nonprompt photons are estimated




FIGURE 5.15: Left: Illustration showing where the (fake) photons from each category are distributed in the $\sigma_{in\eta}$ / charged isolation grid. The boxes show only where the bulk of the events fall, while in reality plenty of events fall outside them. Most notably the hadronic photon category has non-negligible numbers of events with large $\sigma_{in\eta}$ values, but is at the same time much less spread out than the fake photon category. Different drawing styles are merely used to make it clear where the boxes overlap, and the values indicated on the axes correspond to selection criteria which will become relevant in the next section. Right: Yields of the various nonprompt photon categories in the signal region, with both the $\sigma_{in\eta}$ and charged isolation requirements applied.

from data together using a measurement region obtained by inverting the charged isolation requirement on the photon. The $m(l\ell\gamma)$ and $m(l\ell)$ Z boson mass window requirements are kept in order to limit the number of prompt photons from $Z\gamma$ in the measurement region, and to reduce the fraction of nonprompt photons from DY. Reducing the DY fraction in the measurement region is not strictly necessary, but leads to a better closure due to the fact that nonprompt photons originate mostly from $t\bar{t}$ events in the signal region. In the measurement region the requirement on the charged hadron isolation of the photon is inverted with respect to the application region, but required to stay below 15 GeV in order not to deviate too far from the signal region selection. The 2 regions used in the estimation method can be summarized as follows:

Measurement region:

- $N_j \geq 1$ (SF channels) / no jet requirement ($e\mu$ channel)
- $1.141 \text{ GeV} < \text{charged hadron isolation of selected photon} < 15 \text{ GeV}$

Application region, of which the signal region is a subset:

- $N_b \geq 1$
- charged hadron isolation of selected photon $< 1.141 \text{ GeV}$ (as in the photon ID)

Any requirement not listed here is the same as in the signal region event selection.

The estimation is performed by measuring the ratio of the number of events passing the $\sigma_{in\eta}$ threshold from the photon ID ($\sigma_{in\eta} < 0.01015$) to the number falling into the sideband $\sigma_{in\eta} > 0.012$. This leaves a gap between the two $\sigma_{in\eta}$ thresholds, which reduces the number prompt photons in the sideband. Because of the four subregions used when performing an estimate in this way it is typically referred to as a parametrized ABCD method. A visualization of the method is shown in the left-hand image of figure 5.16. The C and D subregions are part of the measurement region, and the C/D ratio is applied to the B subregion of the application region to obtain an estimate of the nonprompt photon background in the subregion A, which corresponds to the signal region. Or written out mathematically, for the case of an estimation parametrized along two axes:

$$\text{prediction}(i, j) = B(i, j) \times C(i, j) / D(i, j) \quad (5.3)$$

Where i and j respectively refer to bins along the axis of the first and second parameter.




FIGURE 5.16: Visual summary of the ABCD subregions used to estimate the nonprompt background (left). Comparison of the $\sigma_{i\eta i\eta}$ distribution shape of nonprompt photons between the measurement and application region in 2016, using nonprompt photons from all simulated process samples (right). The subregions corresponding to passing the $\sigma_{i\eta i\eta}$ requirement, falling into the sideband, and falling into the gap between them are indicated using dotted lines.




FIGURE 5.17: Data-simulation comparison in the measurement and application regions. Dashed lines show the three subregions formed by the requirements on $\sigma_{i\eta i\eta}$.

The measurement and application regions are shown in figure 5.17 with lines indicating the $\sigma_{i\eta i\eta}$ threshold values.

As mentioned in the description of the event selection events with more than one photon passing the photon ID with the $\sigma_{i\eta i\eta}$ and charged hadron isolation requirements relaxed have already been vetoed. This removes only a small fraction of events but it ensures there is no ambiguity between which is the selected photon in the different regions of the ABCD method. To put it differently, not applying this veto would mean that a single event could end up in more than one of the ABCD subregions if it has multiple photons which could each fail of pass the requirements. The right-hand graph of figure 5.16 shows a comparison of the $\sigma_{i\eta i\eta}$ shape for nonprompt photons in the measurement and application regions.

The corrections to the predicted $Z\gamma$ yields described earlier are applied in the measurement and application regions when deriving and applying this estimation. Any prompt photon contribution has to be subtracted from data in order to perform the estimation, and the measurement region is therefore optimized to keep this fraction small. The fraction of prompt photons in the four subregions was already shown in the illustrative image pointed to earlier (5.16), but the detailed background composition is additionally shown in figure 5.19 where the four graphs are arranged in the same way as in the illustration. This estimate is performed in bins of the p_T and η of the selected photon. As this is a small background with limited statistics being estimated these bins are kept relatively wide. The bin




FIGURE 5.18: Estimation maps used in the $t\bar{t}$ MC simulation based closure tests (left) and the data-driven nonprompt photon background prediction (right). The uncertainties shown correspond to the statistical uncertainty in the ratios taken in the measurement region. In the right-hand graph this includes the statistical uncertainties from the prompt photon contributions that were subtracted using information from simulation before taking the ratio.

edges as a function of the $|\eta|$ of the photon are placed at $|\eta| = 0, 0.435, 0.783$ and 1.13 , corresponding to the edges of the different ECAL modules. The closure tests described later were repeated with various bin widths and using an unbinned approach and were found to be consistent with the final binning choice, which shows that the estimation is reasonably independent of the binning used. As the method is the same for the three years they are combined to produce a single map to use for the estimation. The maps for the data-driven prediction and a simulation-based closure test described next are shown in figure 5.18.

To verify the method a closure test is performed using simulated samples. The estimation maps are measured in simulated $t\bar{t}$ and DY samples since the measurement region is enriched in both processes. The resulting estimation to simulation comparison is shown for the 3 years of data combined in figure 5.20 to limit the statistical uncertainty in the closure test. The statistical uncertainty in the sideband events these factors are applied to is treated as a statistical uncertainty in the bins of the estimation, as with any other background. Finally, it must be ensured that this estimate also properly predicts the distributions considered later in for the differential measurements. The closure checks in simulation are therefore performed as a function of all these variables. A selected few of these are shown in figure 5.21 for the 3 years of data combined, but the check was performed for all relevant variables. Based on the closure tests presented a flat systematic uncertainty of 5% is assigned to the estimated nonprompt photon background, and an additional (uncorrelated) uncertainty of 50% is added for events with a photon p_T of above 80 GeV. These same closure tests were also performed using simulations for each data taking year individually, from which similar results were obtained.

Having described the background estimation method as it is implemented in the analysis, I should describe how this exact method was developed and what systematic effects were studied. The category of nonprompt photons is made up of several different contributions, as already described in 5.6.2. These contributions are significantly different in terms of kinematics and detector signatures, and therefore their relative contributions strongly depend on the selection applied. The contribution from fake photons in fact does not correspond to photons at all. At the same time these categories are similar enough for it to be nearly impossible to create regions enriched in just one of them. Early attempts to perform a background estimate using an ABCD method resulted in bad closure test results. A condition that needs to be fulfilled in order for an ABCD method to work is that the two variables used should be uncorrelated. Intuitively this seemed to be the case, as for example the shower shape of a hadronic photon should at most depend on its p_T and η values. Since the parametrization was performed as a function of these two variables this dependence has been taken into account, so that the method was expected to work. A simple calculation however revealed that the correlation between $\sigma_{i\eta i\eta}$ and the charged isolation was not close to zero, and in fact changed depending on other selection criteria. When the correlation was however calculated for the photon categories individually,




FIGURE 5.19: The four subregions used in the ABCD method for 2016, arranged in correspondence with figure 5.16.




FIGURE 5.20: Simulation based closure test of the nonprompt photon background estimate in the 3 channels (left), and as a function of the number of jets and b quark jets (right), for the 3 years of data combined. The error bars on the points in the ratio graph show the combination of the statistical uncertainty in the direct MC simulation prediction shown in red and the prediction from the sideband. The hashed band shows the sum of the flat and high- p_T uncertainties assigned to the prediction.




FIGURE 5.21: Simulation based closure test of the nonprompt photon background estimate as a function of some of the variables considered for unfolding for the 3 years of data combined. The hashed band shows the sum of the flat and high- p_T uncertainties assigned to the prediction.

correlation values of close to zero were found. The takeaway message from this was that while the two variables are more or less uncorrelated within each nonprompt photon category individually, it is the presence of multiple categories which introduces correlation in the global picture. This is further illustrated in the caption of figure 5.22.

In the end the chosen method relies on the event selection of the sideband. It was optimized so that the $\sigma_{i\eta i\eta}$ distribution would match the one in the application region, and would produce an estimate that is not very dependent on the binning used. Aside from the subtraction of prompt photon contributions the estimation method is itself independent of the MC simulation. The central claim is therefore not that the fractions of nonprompt photon sources match the ones in data, but only that the $\sigma_{i\eta i\eta}$ distribution matches between the measurement and application regions.




FIGURE 5.22: Simplified scatter graph of two photon categories as a function of $\sigma_{i\eta i\eta}$ and the charged isolation, for the purpose of illustrating how correlation arises when combining different nonprompt photon categories. If a 2D scatter graph is made for a set of events which are randomly distributed following a normal distribution in both variables, then the result looks somewhat like an ellipse. If the two variables are uncorrelated this ellipse is oriented horizontally or vertically (the shape itself is arbitrarily dependent on the chosen length of the axes). For correlated variables these ellipses are tilted. In the example shown here the two categories circled in red and blue correspond to events for which $\sigma_{i\eta i\eta}$ and the charged isolation are uncorrelated. When taken together however correlation is found, despite the lack of correlation within each of the categories.

5.6.4 Other backgrounds

Several other processes contribute to the background through prompt photons. Single top processes with a prompt photon are categorized under $t\gamma$ in figures. This includes single top production through the s and t-channel (top quark + some other quark final state), and tW production, but the category is completely dominated by the tW process in the signal region due to the extra lepton in its final state. All other minor backgrounds like diboson and triboson events with a photon, $W\gamma$, $t\bar{t}Z+\gamma$, $tZq+\gamma$,... are included in the Other+ γ category. Any contribution of these processes through nonprompt photons is included in the nonprompt photon contribution estimated from data.

5.7 Systematic uncertainties

Below the various sources of systematic effects in this measurement are discussed. These consist of uncertainties in the signal acceptance and background yields related to corrections and calibrations, limitations of the theoretical predictions, uncertainties in values measured elsewhere like the integrated luminosity, and shortcomings of data-driven predictions. Unless stated otherwise these systematic uncertainties are implemented as shape variations, meaning that an alternative set of histograms is provided to the fitting software for every up and down variation. This means that contrary to a “flat” systematic uncertainty, through which every bin of a distribution is assigned the same percentual uncertainty, every bin of a distribution can be affected more or less by the same uncertainty source. It is for example typical for corrections applied to a certain physics object to be dependent on the transverse momentum, with the uncertainties related to these calibrations becoming larger towards the high momentum tails due to the lack of statistics in this region in the measurement where these corrections were derived. Usually the magnitude and shape of uncertainties can be assessed in a straightforward, if at times technically complex, way by varying model parameters or examining how corrections are derived. Sometimes however this is not the case, so that a more intuitive assessment needs to be made, and conservatively large flat uncertainties are typically applied.

Corrections and their variations are applied to simulated events with the exception of the electron/photon scale corrections, for which the corrections are applied on data. The corresponding systematic uncertainties can however simply be assessed by applying the relevant variations to the simulated events. Up and down variations in data respectively lead to the same effect as the down and up variations in simulation, so that overall the same variation is described. The usage of the

term “shape” reflects the way it is used in the fitting tool described later, which means that these shape variations can modify both the shape and the total yield of distributions, but the change in the yield is sometimes intentionally removed through normalization to the nominal yield. All systematic uncertainties listed below are propagated to the fit as nuisance parameters, and are included in the systematic uncertainty bands shown on some of the figures in this work.⁹ The uncertainties associated with the efficiencies of reconstruction procedures and identification criteria all stem from the fact that these efficiencies differ slightly between data and simulation. Because of this correction factors need to be applied to simulated events in order to correct for this discrepancy. The measurement of these correction factors is in itself subject to statistical and systematic uncertainties, and it is these uncertainties which need to be taken into account here. More generally each of the experimental uncertainties is due to the fact that any correction or calibration involves a comparison between data and simulation, and therefore a measurement, with all the same limitations of the measurement of a physics process applying. Theoretical uncertainties are in part due to the inherent limitations of limited order ME level calculations, but also stem from the fact that some steps in the event simulation are not well understood. In this case the only option is essentially to see how the prediction is affected by switching between available models for these steps, or to vary the parameters of the model within some reasonable range.

- The uncertainties associated with the measured values of the integrated **luminosity** add up to 1.2%, 2.3%, 2.5% for 2016, 2017, and 2017 data respectively [150–152]. The corresponding nuisance parameters are however implemented broken down into components that are either fully correlated or fully uncorrelated between the data sets from different years. These uncertainties essentially account for the fact that the total number of collisions each data set corresponds to is in itself a measured value.
- By default, the **pileup** profile in simulated events does not match the one observed in data. Samples are generated using a predicted pileup profile and the events are later reweighted based on the number of true interactions to correct the pileup profile. To adjust the simulated pileup activity to data, a reweighting based on a total inelastic cross section is performed and a 4.6% up and down uncertainty in this cross section is applied, which is propagated to the reweighting histogram. The true pileup profile in data is calculated based on the instantaneous luminosity (per period of 23 seconds), and assuming a total inelastic cross section of 69.2 mb. These pileup profiles are compared to those found in the simulated samples to produce weights that can then correct the simulations to agree with the data.
- **Trigger efficiency corrections** are applied as described in section 5.5. The trigger scale factors are parametrized as a function of both the p_T of the leading and subleading leptons and derived for each channel separately. The uncertainty consists of a statistical component given by the statistical uncertainty of the bin in the scale factor map of each channel, and a systematic uncertainty determined by comparing the efficiencies measured in two separate primary data sets to their average. This systematic uncertainty is determined per channel, but correlated between them since there is overlap when it comes to which triggers feed into each channel.
- Uncertainties for the **efficiency of lepton reconstruction, identification and isolation** are propagated by varying the relevant scale factors. For the electron reconstruction the relevant uncertainties were provided along with the scale factors described in section 5.3.1. For muons no scale factors need to be applied for the reconstruction, and there are therefore no corresponding uncertainties to evaluate. The uncertainties associated with the lepton ID are split into their systematic and statistical parts as mentioned in section 5.3.1, of which the former is by far the dominant component. These uncertainties cover both the identification and isolation selection criteria imposed by the chosen requirement on the discriminator output value.
- Uncertainties associated with the medium **photon ID** and the **pixel seed veto** are applied, both parametrized as a function of the p_T and η of the photon. Uncertainties for very high energy photons were additionally applied.

⁹The total systematic uncertainty band is only drawn on figures specifically showing the signal region, where this information is directly relevant to the final results. While most systematic uncertainties are relevant within any selection, their effect is typically of secondary importance when looking at e.g. a control region used in a data-driven estimation procedure. An exception is the $Z\gamma$ control region, where it was deemed necessary for the systematic uncertainties to be taken into account individually when extracting correction factors to be applied in the signal region.

- For both electrons and photons **momentum scale and resolution smearing** corrections are applied to data and simulation both to make the inefficiencies in the measurement of the electron/photon energy match and to correct the measured energy to the true particle energy on average. The smearing is applied in simulated events to match the resolution observed in data. The resolution in simulations is intentionally made slightly better than in data to make the correction work in this way, since the inverse of smearing is not possible. The uncertainties associated with the scale and resolution corrections are propagated through the analysis by varying the electron and photon momenta.
- The **jet energy scale** corrections (JEC) are associated with a wide range of uncertainty sources, which are grouped based on their origin and their correlation between the data sets from different years. This makes it so that the resulting grouped uncertainties are either fully correlated or fully uncorrelated between data sets which, in combination with the fact that the grouping lowers the total number of systematic uncertainties, reduces the complexity of the implementation. These variations modify the jet momenta, the effect of which is propagated through the analysis.
- **Jet energy resolution** corrections are applied which modify the jet momenta, and the corresponding uncertainties are propagated to all relevant observables. These corrections are similar to the smearing corrections applied to photons and electrons, and are evaluated by varying the degree to which the reconstructed jet energy values deviate from those of their generator level counterparts.
- The simulated efficiencies for the **identification of b quark jets** and for the misidentification of c quark, light-quark or gluon jets are corrected for using scale factors and their respective uncertainties [153], provided as a function of the jet p_T and η . The uncertainties are assessed by varying the scale factors for heavy quarks jets (bottom and charm) and light jets (from lighter quarks or gluons) up and down separately. For the heavy flavours 2 uncertainty components are provided: one fully correlated between the data sets from different years, and one fully uncorrelated between them.
- The **statistical uncertainties in the predictions** are also taken into account. For each histogram used in the likelihood fit an uncorrelated nuisance parameter is introduced for every bin of the distribution [154].
- Uncertainties in the **parton distribution functions (PDF)** are taken into account by taking the root mean square (RMS) of the deviations produced by 100 PDF variations called replicas. Only the shape dependence of this uncertainty is taken into account for the signal, given that the overall normalization of $t\bar{t}\gamma$ will be measured. This is done by normalizing the weights such that the sum of ME level weights in each sample stays constant between these variations. This sum is determined for the total sample independent of any selection or skimming, as keeping the sum of the weights after skimming or the total yield constant would ignore migration effects included in the PDF weight variations.
- Systematic uncertainties related to **factorization and renormalization scales** at the matrix-element level are evaluated by considering the envelope of 6 variations: changing the factorization up (factor 2.0) and down (factor 0.5), changing the renormalization scale up and down, and changing both scales simultaneously. In the same way as was done for the PDF uncertainties, only effects on the shape and acceptance are taken into account when it comes to the signal.
- The uncertainty in the **initial state** and **final state** radiation is considered. These are evaluated using the default set of weight variations where up and down correspond to factors of 2 and 0.5 respectively.
- For $t\bar{t}\gamma$ dedicated simulated samples were generated to assess the systematic uncertainty due to **colour reconnection** and **underlying event** settings. These variations are also available for $t\bar{t}$ samples, but since the entire $t\bar{t}$ phase space is either removed in the overlap removal procedure or included in the estimated nonprompt background these variation samples are not needed. Ideally for the colour reconnection systematic uncertainty one would model the systematic uncertainty as the pre-fit envelope of the three available variations, with the resulting maximum deviation in each bin being added and subtracted to produce the up and down variations (since the variations

provided are one-sided). Because of the small statistics available in these samples this leads to an unphysically large uncertainty stemming from taking the envelope over large fluctuations in the bins used in the fit. The variations are therefore provided separately to the fit, with the variation being taken as the "up" variation, and its inverse w.r.t. the nominal situation as the "down" variation. Similarly the alternative underlying event samples were found to have insufficient statistics, washing out the systematic deviations these samples are supposed to model. The uncertainty for the underlying event systematic is therefore implemented as a flat uncertainty of 0.5% on the signal. This uncertainty was obtained by looking at the effect of the variations in a single bin for the 3 years of data combined, so that the effect from statistical fluctuations could be minimized.

- The uncertainty in the **nonprompt photon background estimation**, which was determined using a closure test in simulation as described in section 5.6.2, is propagated to the final distributions by varying the relevant estimation factors within the uncertainties. Two systematic uncertainties are assigned to this background: A flat uncertainty of 5% across the entire photon p_T range, and an additional separate uncertainty of 50% for events with a photon p_T of above 80 GeV. This background forms about 9% of the total yield in the signal region.
- The **normalizations of the backgrounds** are implemented as flat log-normally distributed nuisance parameters applied to each background category individually. The input uncertainty for the $t\gamma$ background is set to 10%. A priori the input uncertainty in $Z\gamma$ is 5%, but the correction described in section 5.6.1 means most of the $Z\gamma$ yield is already constrained using data. This leaves an input normalization uncertainty of 1.5% for $Z\gamma$. Finally, a 30% normalization uncertainty is assigned to the Other+ γ category. This 30% is a rough conservative estimate of the uncertainty used to account for any possible mismodelling effects and theory uncertainties, but since the contribution of this category is completely negligible in the signal region the exact input value for this uncertainty has no impact on the final result. Both DY and $t\bar{t}$ are part of the total nonprompt background estimated from data. Any DY/ $t\bar{t}$ events where the selected photon is prompt are removed in the overlap removal procedure, or to put it differently these events are part of $Z\gamma/t\bar{t}\gamma$. These two parts cover the entire samples, and hence no other yield remains to assign a normalization uncertainty to.
- A trigger inefficiency not accounted for elsewhere is present in the 2016 and 2017 data sets, usually referred to as the **EGamma prefiring** issue. The gradual loss in transparency of ECAL crystals over time together with the associated trigger electronics occasionally causes an event to be assigned to the preceding bunch crossing. The probability of this happening depends on the opacity of the crystals, and therefore this issue worsens over time. Practically this issue leads to these events being rejected by the L1 triggers. The problem was addressed in time for the 2018 data taking period, but needs to be retroactively corrected for in the 2016 and 2017 data sets. Corrections for this are applied through weights and the associated uncertainty is treated as a systematic uncertainty and propagated through the analysis. For 2018 this reweighting and its associated uncertainty are irrelevant.
- The uncertainty associated with the modelling of the momentum transfer from the b quark to the B hadron, referred to as the **b-fragmentation** uncertainty, can be evaluated by varying the Bowler-Lund parameter in the samples. To avoid having to generate new simulated samples for this the effect is implemented through reweighting using generator level information.
- The yield and kinematic distributions in $t\bar{t}\gamma$ are to a limited degree dependent on the **top quark mass** used in the simulations. This value is set to 172.5 GeV by default, and the current best associated uncertainty is around 0.3 GeV [31]. No top mass variation samples are available for the $t\bar{t}\gamma$ signal, and it is therefore difficult to properly assess the dependence on the top mass. For $t\bar{t}$ samples these variations are available, but the $t\bar{t} + \gamma$ events present in these samples only include photons added in the showering step. The effect of varying the top quark mass is therefore evaluated through a reweighting procedure, in which a weight is applied for each top quark in the event. These weights are obtained by taking the ratio of the Breit-Wigner (BW) distribution with the centre of the mass peak shifted up or down by 0.3 GeV with respect to the BW distribution used in the nominal samples. These weight functions for the ME level top quark masses in the signal events are then evaluated and applied for both top quarks simultaneously to obtain the up and down variations. A reweighting procedure is finally applied

to remove the normalization effect this has on the total sample yield, so that only the acceptance effect is left when looking at the signal region selection. The impact on the signal yield is found to be 0.4%, where it should be noted that this effect is nearly flat as a function of the photon p_T . Since this ad-hoc method of assessing the uncertainty has its limitations compared to using separately produced signal samples with different top quark mass settings this uncertainty is not included in the final fit.

- The correction of the $Z\gamma$ shape and yield explained in section 5.6.1 has 6 different uncertainties associated with it. Five of these are used to take into account the statistical uncertainties in the correction factors, and are implemented as shape uncertainties, while one is associated with the presence of signal in the control region, and takes the form of a flat 1.8% uncertainty in the $Z\gamma$ yield.

Systematic uncertainties specific to $t\bar{t}$ are not listed and used, since after the overlap removal between $t\bar{t}\gamma$ and $t\bar{t}$ no $t\bar{t}$ events with a prompt photon remain. All remaining $t\bar{t}$ events are part of the non-prompt background which is estimated from data.

The $Z\gamma$ corrections described in section 5.6.1 are re-derived for every systematic variation. This means that the control region histograms used to derive the correction factors are produced for every variation, after which these corrections are applied when producing histograms corresponding to the signal region with this variation. Doing this instead of applying the nominal $Z\gamma$ corrections everywhere means that the impact of the systematic on $Z\gamma$ is compensated for to the extent that the effect of the variation is the same in the signal and control regions. For systematic uncertainties which are anti-correlated between the signal and control regions the impact is however increased. The most intuitive example of this is the electron/photon resolution uncertainty: This variation makes the Z-peak wider or thinner, which causes events to migrate from the $Z\gamma$ control region to the signal region or vice versa.

The treatment of the correlations of systematic uncertainties between the data sets from each of the 3 years are summarized in table 5.4. This table also shows both pre-fit and post-fit magnitudes of the uncertainty sources, so that their relative importance can already be kept in mind moving forward. The correlation pattern for the luminosity uncertainty is based on the findings of Ref. [155]. In practice the partial correlations are implemented by splitting the systematic uncertainties up into correlated and uncorrelated parts in the fitting software, based on the method described in Ref. [156]. An example of how systematic uncertainties can modify not only the normalization of a prediction but also its shape is given in figure 5.23.




FIGURE 5.23: Impact of four dominant systematic uncertainties in the signal prediction as a function of the transverse momentum and absolute value of the rapidity of the selected photon. The impact is defined here as the fractional difference, per bin, of the predicted yield between a systematic variation and the nominal situation. Full lines correspond to the “up” variations, and dotted lines correspond to the “down” variations. This example is meant to illustrate how some systematic uncertainties affect the shape of distributions, while others only modify the overall normalization.

TABLE 5.4: Summary of the uncertainty sources in the $t\bar{t}\gamma$ cross section measurements. The first and second column list various uncertainty sources and the corresponding treatment of correlations between the three data-taking years, where $\propto$ means fully correlated, $\sim$ means partially correlated, and $\times$ means uncorrelated. For each uncertainty source the “pre-fit” magnitude is estimated using a cut-and-count analysis of the expected and observed event yields, and this per bin of $p_T(\gamma)$ and for the three data-taking years individually. The typical ranges obtained in this way are shown in column three. The last column shows the impact of each uncertainty source on the measured inclusive $t\bar{t}\gamma$ cross section after the fit to the data performed later on (“post-fit”).

	Source	Correlation	Uncertainty [%]	
			Pre-fit range	Post-fit
Experimental	Integrated luminosity	$\sim$	1.3–3.2	1.7
	Pileup	$\propto$	0.1–1.4	0.7
	Trigger efficiency	$\times$	0.6–1.7	0.6
	Electron selection efficiency	$\sim$	1.0–1.3	1.0
	Muon selection efficiency	$\sim$	0.3–0.5	0.5
	Photon selection efficiency	$\sim$	0.4–3.6	1.1
	Electron & photon energy	$\propto$	0.0–1.1	0.1
	Jet energy scale	$\sim$	0.1–1.3	0.5
	Jet energy resolution	$\propto$	0.0–0.6	<0.1
	b-tagging efficiency	$\sim$	0.9–1.4	1.1
	L1 prefiring	$\propto$	0.0–0.8	0.3
	Values of μ_F and μ_R	$\propto$	0.3–3.5	1.3
Theoretical	PDF choice	$\propto$	0.3–4.5	0.3
	Parton shower: ISR & FSR scale	$\propto$	0.3–3.5	1.3
	Parton shower: colour reconnection	$\propto$	0.0–8.4	0.2
	Parton shower: b-fragmentation	$\propto$	0.0–2.2	0.7
	Underlying event tune	$\propto$	0.5	0.5
Background	$Z\gamma$ correction & normalization	$\propto$	0.0–0.2	0.1
	$t\gamma$ normalization	$\propto$	0.0–0.9	0.8
	Other+ γ normalization	$\propto$	0.3–1.0	0.8
	Nonprompt γ normalization	$\propto$	0.0–1.8	0.7
	Size of background samples	$\times$	1.5–7.6	0.9
	Total systematic uncertainty			3.6
	Statistical uncertainty			1.4
	Total uncertainty			3.9

5.8 Fiducial region definition

As part of the inclusive measurement the measured fiducial cross section for $t\bar{t}\gamma$ production in the dilepton channel will be provided. For this value to have meaning the exact selection at the generator level defined as the signal needs to be defined, and the particle collections which make up the generator level need to be specified. Therefore, while the definition of the fiducial region will be most relevant for the section on unfolding, the definition of this region is already given here.

The particle level (PL) being unfolded corresponds to stable final state particles, where the simulation is performed up to the level of parton showering and hadronization but before any detector simulation is performed. Selection criteria on these particles and events still need to be defined, which are chosen to be similar to those applied at the reconstruction level. This ensures that no extrapolation is made to kinematic regions for which no information is available at the reconstruction level. This set of particle level selections form the fiducial phase space, in which direct comparison between the unfolded measurement and theory predictions is possible without unjustified extrapolation.

The fiducial region definition is summarized in table 5.5. Photons found within $\Delta R < 0.1$ of a lepton are added onto the lepton, meaning that their four-vectors are summed, merging them into a single higher energy lepton. Photons are also required to be isolated, meaning that no stable generator particles with $p_T > 5$ GeV are found within a cone of $\Delta R < 0.1$ (except for neutrinos). At the particle level b quark jets are identified using a ghost-matching procedure [157].¹⁰ Leptons and photons on the other hand can simply be identified using generator truth information. The particle level object definitions were implemented in line with those generally used in CMS analyses, but with settings adjusted to be more suitable for this measurement.

TABLE 5.5: Object and event selection applied at the particle level, which defines the fiducial phase space. The two lepton p_T requirements respectively apply to the highest and second-highest p_T lepton. For the photon “Isolated” implies no stable particles with $p_T > 5$ GeV are present within a cone of $\Delta R = 0.1$, excluding neutrinos. The parameters N_l , N_γ , and N_b refer to the numbers of leptons, photons, and b jets, respectively. The horizontal line separates the object level requirements from the event level requirements.

Leptons	Photons	Jets	b jets
$p_T > 25$ (15) GeV	$p_T > 20$ GeV	$p_T > 30$ GeV	$p_T > 30$ GeV
$ \eta < 2.4$	$ \eta < 1.44$	$ \eta < 2.4$	$ \eta < 2.4$
	$\Delta R(\gamma, l) > 0.4$	$\Delta R(\text{jet}, l) > 0.4$	$\Delta R(\text{jet}, l) > 0.4$
	isolated	$\Delta R(\text{jet}, \gamma) > 0.1$	$\Delta R(\text{jet}, \gamma) > 0.1$
			matched to b hadron
$N_l = 2$ (OC)	$N_\gamma = 1$		$N_b \geq 1$
$m(ll) > 20$ GeV			

In order to be consistent with the reconstruction level, the fiducial region selection is applied on top of the overlap removal procedure described in section 5.2.4. The overlap removal is essentially a selection applied at the parton level. While it needs to be included to be fully correct, the procedure mainly removes events from the $t\bar{t}$ samples, whereas in the signal samples the removed fraction is 1% at most. Tables showing the changes in yield or cross section when applying the fiducial and reconstruction level selection criteria one by one are given in tables 5.6 and 5.7. While the selections are the same for the 3 data taking years the detector configuration slightly differs between these periods, and the actual efficiencies are also not exactly the same. In this table the average values for the 3 years of data taking are therefore listed, weighted by the total luminosity each data set corresponds to. The total cross section after applying all fiducial level selection criteria listed in the table corresponds to the predicted fiducial cross section. The measured fiducial cross section quoted in the next section is obtained by multiplying this predicted cross section by the signal strength extracted in the fit.

¹⁰This is simply a trick used to identify the origin of a jet. During the showering and subsequent clustering step the information of which particle formed the jet is in principle lost. The trick is to add an additional particle with negligible energy to the jet early on, the type of which then serves as a label for the jet it is clustered into later.

TABLE 5.6: Table showing the effect of each requirement for the particle level and reconstruction level selection criteria, corresponding to the luminosity weighted average for the three data sets from different years. Although this is less intuitive the trigger requirements have to be applied after the reconstruction level lepton selection because of the channel-dependent trigger selection.

	requirement	cross section (fb)			
		dileptonic	semileptonic	hadronic	total
Particle level requirements	full sample	2186.92	7387.92	6058.66	15633.51
	overlap removal	2113.32	7149.21	5873.66	15136.19
	PL lepton selection	676.68	26.92	11.42	715.01
	PL photon selection	229.25	4.76	2.50	236.50
	PL $ \eta(\gamma) < 1.4442$	180.40	3.65	1.93	185.98
	PL $m(ll) > 20$	175.42	1.89	0.00	177.31
	PL N_j, N_b requirement	152.98	1.53	0.00	154.51
		cross section $\times$ selection efficiency (fb)			
		dileptonic	semileptonic	hadronic	total
Reconstruction level requirements	lepton selection	110.76	0.73	0.00	111.50
	pass triggers	107.55	0.69	0.00	108.24
	photon selection	50.66	0.04	0.00	50.70
	$ \eta(\gamma) < 1.4442$	50.65	0.04	0.00	50.69
	photon = prompt	50.59	0.04	0.00	50.63
	$m(ll) > 20$	50.54	0.04	0.00	50.58
	$m(ll)$ off-Z	45.06	0.04	0.00	45.09
	$m(ll\gamma)$ off-Z	41.77	0.03	0.00	41.80
	N_j, N_b requirement	33.72	0.03	0.00	33.75

5.9 Inclusive cross section measurement

The cross section measurement happens through a binned likelihood fit to the distribution of the p_T of the selected photon, resulting in a signal strength, pulls, and constraints on the nuisance parameters representing the systematic uncertainty sources [158]. In section 3.8 the general principles of such a procedure were already described, but the main ideas bear repeating: The parameter of interest in a normal fit is the signal strength r , corresponding to the factor by which the predicted signal must be scaled for it to optimally match with the observation. In the fit the test statistic $q(r) = -2 \ln L(r, \hat{\theta}_r) / L(\hat{r}, \hat{\theta})$ is used, which is zero for the global best fit value of r , and larger elsewhere. The likelihood function $L(r, \theta)$ used in this test statistic is the product of the probabilities representing each of the nuisances and the Poisson probabilities for each of the bins. In this notation $\hat{\theta}_r$ represents the values of the nuisance parameters for which the test statistic is minimized, and $\hat{r}$ and $\hat{\theta}$ are respectively the signal strength for which the test statistic is globally minimized and its corresponding set of optimized nuisance parameter values. Also for other values of the signal strength the nuisance parameters are optimized to minimize the test statistic, which is called profiling. The large number of nuisances make this a complex optimization procedure, so it is therefore simplified by using the asymptotic approximation [159, 160].

The $t\bar{t}\gamma$ signal forms the dominant contribution in all signal region bins of the three channels, but a higher purity is obtained in the $e\mu$ channel. The nonprompt contribution in the figures is the total nonprompt photon background estimated from data as described in section 5.6.2.

All systematic variations, as described in section 5.7, are added as nuisance parameters to the fit. For the factorization and renormalization scales and PDF systematic uncertainties an envelope / RMS of the available variations is first taken, after which the upper and lower edges of these envelopes can similarly be added as up and down variations. Normalization uncertainties for the different backgrounds are implemented as log-normally distributed flat uncertainties, with values based on the approximate uncertainties in the cross sections of the backgrounds.

The statistical uncertainties in the signal and background samples are uncorrelated between the bins, and are therefore treated as separate nuisance parameters. Systematic uncertainties are implemented as shapes, meaning that for every up and down variation a set of histograms is provided to the fit. For the implementation of the statistical uncertainties the autoMCstats [154] function of the Combine tool is used, which corresponds to a method sometimes referred to as the “Barlow-Beeston-lite

TABLE 5.7: Table showing the effect of each requirement for the reconstruction level selection criteria, corresponding to the luminosity weighted average for the three data sets from different years. Although this is less intuitive the trigger requirements have to be applied after the lepton selection because of the channel-dependent trigger selection.

requirement	cross section $\times$ selection efficiency (fb)			
	dileptonic	semileptonic	hadronic	total
full sample	2186.92	7387.92	6058.66	15633.51
overlap removal	2113.32	7149.21	5873.66	15136.19
lepton selection	452.75	28.50	2.53	483.78
pass triggers	434.56	24.85	1.87	461.28
photon selection	76.73	1.03	0.02	77.78
$ \eta(\gamma) < 1.4442$	67.14	0.88	0.02	68.03
photon = prompt	66.89	0.85	0.01	67.76
$m(ll) > 20$	65.39	0.79	0.00	66.18
$m(ll)$ off-Z	58.41	0.69	0.00	59.11
$m(ll\gamma)$ off-Z	54.32	0.62	0.00	54.94
N_j, N_b requirement	40.66	0.40	0.00	41.06

approach” [161]. This method assigns a single nuisance parameter to the total expected yield in each bin, which is constrained by its total statistical uncertainty. All other systematic uncertainties are implemented to be fully correlated among the signal regions. In the case of linear systematic uncertainties the type of uncertainty and its magnitude can simply be specified.

From the fit the following signal strength is extracted:

$$r = 1.134 \pm 0.016 \text{ (stat)} \pm 0.042 \text{ (syst)} \quad (5.4)$$

Multiplying this value by the expected fiducial cross section obtained earlier finally yields the measured or observed inclusive cross section for $t\bar{t}\gamma$ in a 2 lepton final state:

$$\sigma_{\text{fid}}(pp \rightarrow t\bar{t}\gamma) = 175.2 \pm 2.5 \text{ (stat)} \pm 6.3 \text{ (syst) fb} \quad (5.5)$$

which is compatible with the theoretical prediction of 155 ± 27 fb within the corresponding large uncertainties. As is apparent here the experimental uncertainties in the $t\bar{t}\gamma$ cross section that can be achieved using currently available data sets are much better than the theoretical uncertainty on the prediction obtained from the cross section calculation performed at NLO. The pulls and constraints on the nuisance parameters, the meaning of which is explained below, are shown in figure 5.24. In this figure the Δr ranges shown on the right correspond to the fractional change in the signal strength obtained if the nuisance is varied up or down within one standard deviation. Notice how for most uncertainties pulling the nuisance up reduces the signal strength, which is to be expected since usually this corresponds to an increase of the predicted yields. The points on the left of the figure indicate to what degree, and in what direction, the nuisances have been pulled in the fit. The error bars around these points show whether the uncertainty is constrained by the fit. If the half width of the error bar corresponds to a range of 1 on the axis the uncertainty is unconstrained. If it is however shorter when fitting to data, this means that the freedom offered to the nuisance by the statistical uncertainties is smaller than initially provided. This would in principle imply that the measurement could be used to measure the nuisance parameter more precisely than whatever source it was obtained from. In practice small constraints are simply an artifact of minor statistical fluctuations in the bins, and are no cause for concern. Large constraints would however mean that the uncertainty was overestimated. As shown in the figure only small pulls are extracted from the fit, and none of the nuisances are significantly constrained.

Scaling the signal and background predictions according to the extracted signal strength and nuisance parameters produces what is referred to as post-fit distributions. Figure 5.25 shows post-fit graphs for the distribution of the reconstructed photon p_T , for the 3 years of data combined but split up into the three lepton channels. The fitting procedure can be repeated for each of the lepton channels individually, the results of which are shown in figure 5.26. As expected the cross sections extracted in each of the channels are compatible within the statistical uncertainties, and the smallest statistical




FIGURE 5.24: Pulls and constraints for the dominant systematic uncertainties in the fit used to extract the signal strength. The meaning of these values is explained in the text.

and systematic uncertainties are obtained in the $e\mu$ channel. The figure also shows two scans of the test statistic as a function of the signal strength, one profiling over the nuisances, and one with the nuisances fixed to the values they take on in the global minimum.

The photon p_T distribution was chosen as input to the fit because it separates the signal events from the backgrounds to some degree, resulting in a reduced total uncertainty in the fit. The signal strength itself should however be reasonably independent of the variable being considered. This can simply be checked by repeating the fit using different input distributions, while keeping all other settings the same. Figure 5.27 shows the result of performing this check for a wide range of kinematic variables, and shows that the extracted signal strengths are compatible within the uncertainties.




FIGURE 5.25: The observed (points) and predicted (shaded histograms) event yields as a function of the reconstructed photon p_T after applying the event selection, for the three decay channels. Predictions are scaled in accordance with the signal strength and nuisances extracted from a fit to the data. The vertical bars on the points show the statistical uncertainties in data, and the hatched bands the systematic uncertainty in the predictions. The lower panels of each graph show the ratio of the event yields in data to the predictions.




FIGURE 5.26: Measured inclusive fiducial production cross sections for $t\bar{t}\gamma$ production decaying into two leptons (left). The measurement obtained using all three decay channels is compared to the results obtained by performing fits to each of the channels individually. Horizontal lines on the points respectively show the statistical and total uncertainties. The SM prediction is shown as a vertical line, and the corresponding theoretical uncertainty corresponds to the shaded grey band. Likelihood scans of the signal strength (right). The solid black line shows the profile likelihood scan, while the dashed red line instead has the nuisance parameters fixed to their best fit values.




FIGURE 5.27: Signal strengths extracted by repeating the entire fitting procedure using different input distributions. The statistical and total uncertainties in each of the signal strengths are indicated using the inner and outer error bars. The uncertainties in the signal strength obtained using the photon p_T distribution as in the main result are also drawn as coloured bands over the entire range for ease of comparison. As one would expect the results are compatible within the uncertainties.




FIGURE 5.28: A visual representation of the unfolding process. Information from simulated signal events at the particle and reconstruction level are used to construct a response matrix, which encodes the effects of the detector response and selection steps. By inverting this matrix and applying it to an observed signal distribution a distribution at the particle level is obtained. In this illustration the same binning is used on both axes of the response matrix. In reality a finer binning is typically used on the reconstructed variable axis, which improves the stability of the unfolding procedure.

5.10 Differential cross section measurement

5.10.1 Unfolding methods and tools

The way distributions of reconstructed variables are discussed may wrongly give the impression that they also directly correspond to kinematic features of the physics processes. This is however not the case, not only because the various selection steps affect the distributions, but also due to the fact that measurements are to some degree distorted by the response of the detector itself. In order to deliver measurements of observables which correspond more directly to the actual physics processes the effects of the detector and event selection need to be removed.¹¹ The response of the detector, the inefficiencies of the reconstruction, and the effects of the analysis selection steps can be modelled based on simulated events, so that these effects be removed from measured distributions using a procedure called unfolding. While the Combine software used to perform the fits can be used to this end the implementation provided by the TUnfold package [162] is used instead.

The general procedure for unfolding is as follows: A response matrix is produced using signal events passing the fiducial region selection criteria, and this is done for every distribution for which the unfolding is to be performed. This response takes on the form of a 2D matrix with a reconstructed variable on one axis, and its generator level counterpart on the other. The axis ranges are typically similar or the same, but a finer binning is used for the reconstruction level. The underflow bins along with the reconstructed variable axis are used to keep track of events which fail to get reconstructed correctly or do not pass the reconstruction level selection. TUnfold takes these into account when performing the unfolding, where this fraction corresponds to $1 -$ the selection efficiency, with the selection efficiency being the fraction of events in the fiducial region which pass the reconstruction level selection. A visual aid in understanding the unfolding method is provided in figure 5.28.

¹¹I write “more directly” as the unfolded distributions still very much depend on the fiducial region definition. The distortions caused by the detector and the reconstruction level selections are however removed, so that the unfolded results are conceptually much closer to the “true physics” kinematics in the events.

Various approaches to unfolding exist, each with their own advantages and drawbacks. While the unfolding framework used to obtain the results described in this section was developed from scratch on top of the TUnfold package, the method itself is very similar to the one using in Ref. [54] mainly authored by Illia Khvastunov and the rest of the UGent CMS team. The person who implemented the unfolding step in said analysis, Joscha Knolle, provided explanations, advice, and suggestions regarding my own implementation of these methods. The main feature of this approach is the fact that no information extracted from a fit is used. Instead, predicted backgrounds are subtracted from the observed reconstruction level distributions, so that a differential measurement of the signal is obtained that is independent of the MC simulation for the signal. Such an approach is only valid if the pulls and constraints seen in the inclusive fits are small, as is the case for this analysis. A signal strength deviating from unity is however not a problem, as subtracting the backgrounds from the data distributions results in a yield that is very close to the predicted yield scaled by the signal strength.

When filling the response matrix all relevant weights are applied as in the inclusive cross section measurement, but some subtleties need to be taken into account due to the fact that events failing the reconstruction level selection are also used here. For events which fail this selection, and are therefore sorted into the underflow bins, the reconstruction-related weights are undefined. These events get only the weights related to the simulation applied to them, like generator weights and the weights accounting for the luminosity. Events passing the reconstruction level selection are assigned a weight obtained by multiplying the generator-related weights with the weights related to fulfilling the reconstruction level selection criteria. These reconstruction-related weights account for data-simulation differences in the performance of each of these selections and are typically close to one. For each of these events an event should also be added to the corresponding underflow bin with a weight given by the product of the generator weights multiplied by $(1 - \text{the product of the reconstruction level weights})$. In other words the use of weights to account for the different performance of reconstruction level selections between data and simulation splits an event into a passing and failing event, with weights that add up to 1. In the inclusive cross section measurement only the fraction that represents passing the reconstruction level selection criteria is of interest, but in unfolding the failing fraction is needed as well. Of course, because the scale factors are typically close to one, these failing events end up with weights of close to zero. Nevertheless failing to take them into account still introduces a small inaccuracy.

Unfolding is only performed on the signal, and therefore the various background predictions are subtracted from the data. In addition to the background contribution the outside migration fraction, referred to as f_{out} , is subtracted. This outside migration fraction consists of events that pass the reconstruction level selection, and are therefore present in the inclusive cross section measurement histograms, but fail any of the fiducial level selection criteria listed in table 5.5. f_{out} relies on generator level information for its definition, and therefore needs to be taken from simulation. TUnfold is capable of keeping track of the statistical uncertainties in the data, predicted backgrounds, and outside migration fraction itself throughout the subtractions and unfolding. It was verified that manually evaluating the statistical uncertainties gives the same results as obtained when using the TUnfold functions to do so, and from that point on the provided functions for error propagation were used.

The full unfolding procedure is performed using pre-fit information. This is justified by the fact that the pulls and constraints obtained from the fits in section 5.9 are small with respect to the uncertainties. Additionally, the $Z\gamma$ and nonprompt photon backgrounds already contain data information since $Z\gamma$ is corrected based on data in a control region, and the nonprompt photon background is obtained using a data-driven background estimation.

The same systematic uncertainties that are considered in the inclusive cross section measurement need to be taken into account here, and aside from differences explicitly mentioned below their implementations and magnitudes are essentially the same. In order to propagate the uncertainties to the unfolded distributions the unfolding is repeated for every variation, after which the total systematic uncertainty can be obtained by summing the deviations from the nominal case in quadrature. Which of the inputs is varied and the exact procedure depends on the type of systematic uncertainty:

- The normalization uncertainties for the backgrounds can be evaluated by scaling each of the backgrounds up or down within their uncertainty before performing the background subtraction. These variations do not impact the response matrix so that the nominal one can be used. The

luminosity can be treated in a similar way by scaling all backgrounds together up or down by the uncertainty in the luminosity.¹²

- Uncertainties associated with the MC simulation affect both the backgrounds and the signal, with the exception of the nonprompt photon background estimated using data. Both the background histograms and the response matrices are therefore varied in each of the repeated unfolding procedures. For the PDF and factorization and normalization scale uncertainties this means the unfolding is repeated for each of the PDF sets and scale choices. The PDF uncertainty in each bin is then obtained by taking the RMS of the deviations of all obtained unfolded distributions, similar to how this uncertainty is evaluated at the reconstruction level. A similar approach is used for the factorization and normalization scale uncertainties, where the envelope of the obtained deviations is taken, once again using the same approach as in the inclusive cross section measurement.
- For reconstruction and selection-related uncertainties a practical consideration needs to be made. While it is possible to evaluate their effects by varying the backgrounds and the response matrix as it is done for the modelling uncertainties, the same results can be obtained by varying the data instead of the response matrix. This prevents statistical fluctuations from affecting the results, since when filling the response matrix the available simulated events are distributed over both reconstruction level and generator level bins. Since generator level variables are completely independent of these reconstruction level variations the exact same results are obtained with this alternative approach. In practice this means looking at the sum of the predicted backgrounds plus the predicted signal, comparing the variation and nominal case, and scaling the data accordingly. The unfolding is then performed using this data variation, the variation of the backgrounds, and the nominal response matrix.

TABLE 5.8: Definition of the observables used in the differential cross section measurement.

Symbol	Definition
$p_T(\gamma)$	Transverse momentum of the photon
$ \eta(\gamma) $	Absolute value of the pseudorapidity of the photon
$\min\Delta R(\gamma, l)$	ΔR between the photon and the closest lepton
$\Delta R(\gamma, l_1)$	ΔR between the photon and the leading lepton
$\Delta R(\gamma, l_2)$	ΔR between the photon and the subleading lepton
$\min\Delta R(\gamma, b)$	ΔR between the photon and the closest b jet
$ \Delta\eta(ll) $	Pseudorapidity difference between the two leptons
$\Delta\varphi(ll)$	Azimuthal angle difference between the two leptons
$p_T(ll)$	Transverse momentum of the dilepton system
$p_T(l_1) + p_T(l_2)$	Scalar sum of the transverse momenta of the two leptons
$\min\Delta R(l, j)$	Smallest ΔR between any of the selected leptons and jets
$p_T(j_1)$	Transverse momentum of the leading jet

All of the variables considered in the unfolding are defined in table 5.8, and two examples of corresponding response matrices are shown in figure 5.29. Distributions showing the performance of the selection (the outside migration fraction, efficiency, and purity) for the same two variables are shown in figure 5.30. For the purpose of making these performance figures the generator level binning is temporarily made finer to match the reconstruction level. As mentioned earlier the outside migration fraction f_{out} corresponds to the fraction of events in a certain reconstruction level bin, which do not pass the fiducial level selection criteria. The efficiency shown in these graphs is simply the fraction of events in generator level bin i that ends up in the corresponding reconstruction level bin $j = i$. It therefore parametrizes both the event selection efficiency, and the probability of an event being sorted into the correct reconstruction level bin. Finally, the purity is a variable representing what fraction of events in a certain reconstruction level bin also fell into the same bin range at the particle level.

¹²The nonprompt photon background is not scaled along with the other backgrounds, as it is estimated from data and therefore unaffected by the luminosity uncertainties.




FIGURE 5.29: Example response matrices for the transverse momentum of the selected photon (left) and the transverse momentum of the leading jet (right).

An overall efficiency of around 20% is found, depending on the distribution and its binning. While a higher efficiency is in principle better the actual value of the efficiency is highly dependent on the fiducial region definition. For the unfolding procedure using 3 years of data combined the response matrix is formed by stitching the response matrices of the 3 years together while summing the underflow bins. The diagonality of a response matrix essentially depends on whether the correct objects are reconstructed and selected and, if they are, whether their reconstructed kinematic properties match their the corresponding particle level values closely enough. Resolution effects cause events to migrate from one bin at the generator level to a neighbouring one at the reconstruction level, whereas the selection of a wrong object can also cause migration to a random bin at the reconstruction level.

The matrix inversion needed to perform the unfolding is non-trivial. It essentially corresponds to finding an estimator for the yield in each of the particle level bins as a function of the reconstruction level bin contents [163]. If the response matrix has large values in off-diagonal bins the problem quickly becomes ill-defined, leading to estimators with large associated variances. This in turn leads to unphysical fluctuations in the unfolded distributions. This is to a large degree already remedied by using a finer binning at the reconstruction level, and by choosing reasonably wide bins in general. Other solutions to these kinds of fluctuations fall under the umbrella term of regularization techniques. Many such methods are available, but all of them are in essence ways to smoothen the output distributions and reduce the correlations between bins. All of these methods do however introduce a bias into the results, so that it is preferable to avoid using them if possible. Checks were performed to verify that applying regularization did not significantly reduce the correlation between bins, and to confirm that the response matrix inversions are well-defined. Variables with worse resolutions inherently exhibit more bin-to-bin migration, so that for these variables coarser bins are used in order to reduce the yields in off-diagonal bins.

5.10.2 Differential cross sections

The unfolded data distributions for the three data years combined are given in figures 5.31 and 5.32. The figures also show the theory predictions corresponding to the signal samples after the fiducial level selection. The impacts of the renormalization and factorization scales and the choice of PDFs on the theory prediction are evaluated by varying the signal sample directly. For these the envelope and RMS of the different variations they include is once again taken. One difference with the reconstruction level is that here the weights used for the PDF and renormalization/factorization scale are not adjusted to keep their sum constant. These uncertainties are not meant to affect the measured cross section, and therefore should not impact the total yield of the samples, but when assessing the uncertainty in the theoretical prediction the full effect of the variations should be taken into account. In addition to the theory prediction obtained from the samples used in the rest of the analysis (MADGRAPH5_aMC@NLO




FIGURE 5.30: The efficiency, purity, and outside migration fraction as a function the transverse momentum of the selected photon (left) and the transverse momentum of the leading jet (right). While the fiducial region definition, the object reconstruction, and event selection affect these values they are also simply dependent on the binning used.

+ PYTHIA8 for the showering step) theory predictions obtained using Herwig7 for the showering simulation are also shown.

As the signal samples were generated at leading order, but their cross sections are corrected using NLO K-factors, NLO signal samples are used to assess the impact of the PDF and renormalization/factorization scale uncertainties instead. The available NLO signal samples do not include photons from top quark decay products, but this matter is unrelated to the effects of the PDF choice and scales so that using the NLO sample should be acceptable for this specific application. The bin-by-bin deviations observed in the NLO samples are projected onto the LO sample distributions to obtain the final uncertainty bands.

In addition to the unfolded results themselves, correlation matrices for the uncertainties in these distributions were calculated. Examples of such matrices are given in figure 5.33. In order to produce these first the covariance between the bins (i, j) for a variation due to a statistical or systematic uncertainty is calculated as

$$\text{Cov}(i, j) = \left(\sigma_i^{\text{variation}} - \sigma_i^{\text{nominal}} \right) \cdot \left(\sigma_j^{\text{variation}} - \sigma_j^{\text{nominal}} \right). \quad (5.6)$$

The “up” and “down” variations for each uncertainty source are evaluated separately, and the covariance is obtained by taking the average of the two variations. The total covariances are obtained by summing the individual contributions from each uncertainty source. From this, the correlation between the bins (i, j) is calculated as

$$\text{Corr}(i, j) = \text{Cov}(i, j) \cdot \left(\text{Cov}(i, i) \cdot \text{Cov}(j, j) \right)^{-1/2}. \quad (5.7)$$

These matrices are given separately for the statistical uncertainty in the data, and the systematic uncertainties in which the statistical uncertainties in the predicted backgrounds are included. As one would expect for the statistical uncertainties large correlation values are mostly limited to neighbouring bins. The systematic uncertainties on the other hand are very much correlated between bins, as most systematic uncertainties affect the overall acceptance significantly.

5.10.3 Normalized differential cross sections

Finally, it is interesting to look at the normalized versions of these distributions. In order to obtain these results the unfolded and theoretically predicted distributions are normalized to one, so that purely the shapes can be compared:

$$n_i = \sigma_i \cdot \left(\sum_j \sigma_j \right)^{-1}, \quad (5.8)$$

Where n_i and σ_i are the bin-values of the normalized and absolute distributions, respectively. As described earlier the unfolding is repeated for every systematic variation to assess the impact of




FIGURE 5.31: Differential $t\bar{t}\gamma$ production cross sections as functions of kinematic variables. The data are represented by points, with inner (outer) vertical bars indicating the statistical (total) uncertainties. The predictions obtained with the MADGRAPH5_aMC@NLO event generator interfaced with PYTHIA8 and HERWIG7 are shown as horizontal lines. The theoretical uncertainty bands and χ^2/dof values shown correspond to the PYTHIA8 model.




FIGURE 5.32: Differential $t\bar{t}\gamma$ production cross sections as functions of kinematic variables. The data are represented by points, with inner (outer) vertical bars indicating the statistical (total) uncertainties. The predictions obtained with the MADGRAPH5_amc@NLO event generator interfaced with PYTHIA8 and HERWIG7 are shown as horizontal lines. The theoretical uncertainty bands and χ^2/dof values shown correspond to the PYTHIA8 model.




FIGURE 5.33: Correlation matrices of the systematic uncertainty (top) and statistical uncertainty (bottom) in the measured differential cross section as a function of $p_T(\gamma)$ (left), and $p_T(j_1)$ (right).

these systematic uncertainties on the results. The results of this procedure are given in figures 5.34 and 5.35. The theory uncertainties mainly correspond to a (large) normalization effect, which is why the theory uncertainty band is significantly reduced in the normalized distributions. To evaluate the uncertainty in the normalized result, each of the variations is normalized according to Eq. (5.8) before the deviation from the nominal case is determined. These normalized distributions allow one to focus solely on the shape differences between measurements and predictions, which can be for example be useful in assessing which variables are not well modelled by the leading order $t\bar{t}\gamma$ simulation.

5.10.4 Comparison with theoretical predictions

To evaluate the agreement between the measured differential distributions and the theory predictions based on PYTHIA8 χ^2 values are calculated. The method follows the procedure described in Ref. [164]. In the χ^2 calculation, both the uncertainty in the measured distributions as well as in the prediction are taken into account.

For the absolute distributions and starting from the covariance matrix $\text{Cov}(i, j)$, calculated according to Eq. (5.6) and summed for the statistical, systematic, and theoretical contributions, the χ^2 value is calculated as:

$$\chi^2_{\text{absolute}} = \sum_{i=1}^N \sum_{j=1}^N \left(\sigma_i^{\text{meas.}} - \sigma_i^{\text{pred.}} \right) \cdot \left(\sigma_j^{\text{meas.}} - \sigma_j^{\text{pred.}} \right) \cdot \left[\text{Cov}^{\text{stat.}}(i, j) + \text{Cov}^{\text{syst.}}(i, j) + \text{Cov}^{\text{theo.}}(i, j) \right]^{-1}. \quad (5.9)$$




FIGURE 5.34: Normalized differential $t\bar{t}\gamma$ production cross sections as functions of kinematic variables. The data are represented by points, with inner (outer) vertical bars indicating the statistical (total) uncertainties. The predictions obtained with the MADGRAPH5_aMC@NLO event generator interfaced with PYTHIA8 and HERWIG7 are shown as horizontal lines. The theoretical uncertainty bands and χ^2/dof values shown correspond to the PYTHIA8 model.




FIGURE 5.35: Normalized differential $t\bar{t}\gamma$ production cross sections as functions of kinematic variables. The data are represented by points, with inner (outer) vertical bars indicating the statistical (total) uncertainties. The predictions obtained with the MADGRAPH5_AMC@NLO event generator interfaced with PYTHIA8 and HERWIG7 are shown as horizontal lines. The theoretical uncertainty bands and χ^2/dof values shown correspond to the PYTHIA8 model.

Here, N is the number of bins of the distribution, which corresponds to the number of degrees of freedom to which the χ^2 value is to be compared.

For the normalized distributions, the number of degrees of freedom (dof) is $N - 1$, since the normalization enforces that the cross section value of any bin can be obtained from calculating one minus the sum of all other bins. This normalization condition is however not respected by the variations calculated for the theory uncertainties with special prescriptions, where the envelope (scale variations) or RMS (PDF) of several variations are taken. The envelope/RMS for these uncertainties are calculated after applying the normalization, so that the integrals of the corresponding histograms deviate slightly from unity. If one were to use Eq. 5.9 to calculate a χ^2 value by dropping one arbitrary bin from the sums (as suggested in Ref. [164]), the result would have a bias from the choice of the dropped bin, and different choices would result in different χ^2 values.

Instead, a χ^2 value for the agreement of the normalized measured and predicted distribution is computed from the covariance matrices of the absolute distributions in the following way: The value χ^2_{absolute} calculated according to Eq. (5.9) describes both the agreement in shape and normalization. In addition to this a χ^2 value is calculated for just the normalization of the absolute distributions (i.e. for the denominator in Eq. (5.8) used for the calculation of the normalized distribution), or written out:

$$\chi^2_{\text{normalization}} = \left(\sum_{i=1}^N \sigma_i^{\text{meas.}} - \sum_{i=1}^N \sigma_i^{\text{pred.}} \right)^2 \cdot \left[\sum_{i=1}^N \sum_{j=1}^N \left[\text{Cov}^{\text{stat.}}(i, j) + \text{Cov}^{\text{syst.}}(i, j) + \text{Cov}^{\text{theo.}}(i, j) \right] \right]^{-1}, \quad (5.10)$$

which corresponds to one degree of freedom. Then, the χ^2 value for the agreement between the normalized measured and predicted distributions can be calculated as $\chi^2_{\text{normalized}} = \chi^2_{\text{absolute}} - \chi^2_{\text{normalization}}$. This value corresponds correctly to $N - 1$ degrees of freedom, and has no bias from a bin choice. The calculated χ^2 values and corresponding number of degrees of freedom are shown on each of the results.

For the purpose of assessing which distributions show good agreement between measurement and prediction the normalized differential distributions and their corresponding χ^2 values are the more straightforward results to scrutinize. The χ^2/dof values for the $p_T(\gamma)$ and $|\eta|(\gamma)$ distributions of respectively 5.2/5 and 3.3/6 indicate good agreement within the uncertainties. In terms of the χ^2/dof values the variables related to the angle between the photon and the decay products of the top quarks also show good agreement, but a consistent trend of finding smaller separations in the prediction than in the measurement can be noticed. As these variables are sensitive to the origin of the photon these trends could hint at a mismodelling of the photon origins in the simulated signal samples, which would not be surprising given that the samples are produced at LO. Kinematic variables not involving the photon offer sensitivity to the modelling of the top quark decay, with $|\Delta\eta(l)|$ and $\Delta\varphi(l)$ being the most interesting due to their sensitivity to $t\bar{t}$ spin correlations [144, 165]. Some small deviations with respect to the predictions are found, most notably for $|\Delta\eta(l)|$ with a χ^2/dof value of 11.3/6, but the measurements remain compatible with the prediction within the experimental uncertainties. Unsurprisingly the result showing the largest difference between the prediction using PYTHIA8 and HERWIG is the $p_T(j_1)$ distribution.

5.11 Effective field theory interpretation

An introduction to EFT (section 4.1) and its practical implementation (sections 4.2 and 4.4) was already given earlier, so that this information could be placed in a broader context. We can therefore now focus on the technicalities and results specific to the $t\bar{t}\gamma$ process.

As mentioned earlier the c_{tZ} and c_{tZ}^I Wilson coefficients are selected as being the most relevant to this process, and the Wilson coefficients of all other operators are set to zero. The effect of these modifications is probed in the measured distribution of the photon p_T at the reconstruction level, which is also used in the inclusive cross section measurement. This kinematic variable was found to be the most sensitive to new physics effects, in the sense that its shape changes significantly as a function of the Wilson coefficients, in addition to the normalization effects that affect any variable. It




FIGURE 5.36: Predicted photon p_T distribution of $t\bar{t}\gamma$ signal process compared for different values of the c_{tZ} and c_{tZ}^I Wilson coefficients.

also separates the signal from the remaining backgrounds to some degree, which further benefits the sensitivity. With a fixed EFT expansion parameter of $\Lambda = 1$ TeV, the expected EFT modifications for non-zero values of c_{tZ} and c_{tZ}^I are estimated at the particle level and used to calculate per-event weights corresponding to the ratio of the predicted SM and EFT cross sections in bins of the photon p_T . The nominal simulation is then reweighted after applying the full analysis selection criteria to the reconstructed events in order to estimate the expected EFT modifications at the reconstruction level. In the parametrization of the EFT effects both linear and quadratic terms in the Wilson coefficients are included. Figure 5.36 shows the effect of setting these Wilson coefficients to some (relatively large) non-zero values on the photon p_T distribution.

The similar measurement by CMS using final states with one lepton and jets (l+jets) presented in Ref. [6] probes an orthogonal fiducial phase space of the $t\bar{t}\gamma$ production process defined for the single-lepton decay channel of the $t\bar{t}$ system. While the analysis presented here manages to deliver the most precise inclusive cross section thanks to very small background contributions, the measurement in Ref. [6] has a significantly larger number of signal events at large photon p_T values in data, which increases the sensitivity to modifications described by the studied Wilson coefficients. To further improve the constraints on the Wilson coefficients, a combined EFT interpretation of both measurements is also performed.

The constraints on the Wilson coefficients are measured using a profile likelihood fit constructed in the same way as for the inclusive cross section measurement. Since here it is not the signal strength being extracted, the modelling uncertainties in the $t\bar{t}\gamma$ signal normalization are added to the set of systematic uncertainties. The minimized likelihood value obtained in a fit using the EFT-predicted photon p_T distribution is compared to the corresponding likelihood value using the SM prediction. In the combination with the l+jets measurement the correlations between the nuisance parameters applied in both analyses are treated appropriately.

One-dimensional scans of the negative log-likelihood value difference to the best fit value are shown in figure 5.37 for each Wilson coefficient, with the other Wilson coefficient set to zero in the fit. These results are shown for dilepton events alone and for the combination with the l+jets analysis, where in both cases the photon p_T distribution is used as input. The corresponding 68 and 95% CL intervals are listed in table 5.9. Fits are also performed where both Wilson coefficients are varied simultaneously in the fit. The result of the two-dimensional scans and the corresponding 68 and 95% CL contours are shown in figure 5.38. In the two-dimensional scans the best fit point, indicated using

red diamonds in figure 5.38, is found at $(c_{tZ}, c_{tZ}^I) = (-0.28, -0.02)$ for the dilepton-only interpretation, and at $(-0.22, -0.05)$ for the combined interpretation. The SM expectation of zero for both Wilson coefficients is 0.7 (1.2) standard deviations away from the best fit point for the dilepton-only (combined) interpretation. In other words the SM expectation is compatible with the 95% CL limits obtained in the fits, and this for both the one- and two-dimensional case of the dilepton-only and combined analyses.

Other measurements [53, 54, 88, 91, 95] and global EFT fits [99, 166–169] have also derived constraints on these Wilson coefficients. The 95% CL limits on c_{tZ} and c_{tZ}^I from a selection of these analyses are compared to the results of this analysis in figure 5.39. Back in section 4.2 the higher sensitivity of the $t\text{-}\gamma$ vertex to the c_{tZ} and c_{tZ}^I operators with respect to the $t\text{-}Z$ vertex was already illustrated. Together with the larger production cross section of $t\bar{t}\gamma$ and its relatively small backgrounds due to the presence the high energy photon in its signal definition this leads to stronger constraints on c_{tZ} and c_{tZ}^I than are obtainable from a $t\bar{t}Z$ measurement. The limits on c_{tZ} and c_{tZ}^I quoted in the CMS $t\bar{t}Z$ measurement using 77.5 fb^{-1} of data [54] are about three times less stringent than the combined limits obtained here, while the limit on c_{tZ} from a recent measurement of $t\bar{t}Z$ and tZq production by CMS using 138 fb^{-1} of data [91] is a factor of two less stringent. While the global fit analysis uses input from $t\bar{t}\gamma$ and $t\bar{t}Z$ measurements, among others, its sensitivity is further limited by the fact that only linear terms in the Wilson coefficients are considered as discussed back in section 4.8.

TABLE 5.9: Summary of the one-dimensional 68 and 95% CL intervals obtained for the Wilson coefficients c_{tZ} and c_{tZ}^I using the photon p_T distribution from this analysis or the combination of this analysis with the l+jets analysis from Ref. [6]. The profiled results correspond to the fits where the other Wilson coefficient is left free in the fit, otherwise it is set to zero.

			Dilepton result	
Wilson coefficient			68% CL interval [(Λ/TeV) ²]	95% CL interval [(Λ/TeV) ²]
Expected	c_{tZ}	$c_{tZ}^I = 0$	[-0.28, 0.35]	[-0.42, 0.49]
		profiled	[-0.28, 0.35]	[-0.42, 0.49]
	c_{tZ}^I	$c_{tZ} = 0$	[-0.33, 0.30]	[-0.47, 0.45]
		profiled	[-0.33, 0.30]	[-0.47, 0.45]
Observed	c_{tZ}	$c_{tZ}^I = 0$	[-0.43, -0.09]	[-0.53, 0.52]
		profiled	[-0.43, 0.17]	[-0.53, 0.51]
	c_{tZ}^I	$c_{tZ} = 0$	[-0.47, -0.03] $\cup$ [0.07, 0.38]	[-0.58, 0.52]
		profiled	[-0.43, 0.33]	[-0.56, 0.51]
Dilepton & l+jets combination				
Wilson coefficient			68% CL interval [(Λ/TeV) ²]	95% CL interval [(Λ/TeV) ²]
Expected	c_{tZ}	$c_{tZ}^I = 0$	[-0.15, 0.19]	[-0.25, 0.29]
		profiled	[-0.15, 0.19]	[-0.25, 0.29]
	c_{tZ}^I	$c_{tZ} = 0$	[-0.17, 0.18]	[-0.27, 0.27]
		profiled	[-0.18, 0.18]	[-0.27, 0.27]
Observed	c_{tZ}	$c_{tZ}^I = 0$	[-0.30, -0.13]	[-0.36, 0.31]
		profiled	[-0.30, 0.00]	[-0.36, 0.31]
	c_{tZ}^I	$c_{tZ} = 0$	[-0.47, -0.03] $\cup$ [0.07, 0.38]	[-0.58, 0.52]
		profiled	[-0.43, 0.33]	[-0.56, 0.51]




FIGURE 5.37: Distributions of the observed (solid line) and expected (dashed line) negative log-likelihood difference from the best fit value for the one-dimensional scans of the Wilson coefficients c_{tZ} (upper) and c_{tZ}^1 (lower), using the photon p_T distribution from this analysis (left) or the combination of this analysis with the l +jets analysis from Ref. [6] (right). In the scans the other Wilson coefficient is profiled. The green (orange) bands indicate the 68 (95)% CL limits on the Wilson coefficients.




FIGURE 5.38: Result from the two-dimensional scan of the Wilson coefficients c_{tZ} and c_{tZ}^I using the photon p_T distribution from this analysis (left) or the combination of this analysis with the $l+jets$ analysis from Ref. [6] (right). The shading quantified by the colour scale on the right reflects the negative log-likelihood difference with respect to the best fit value that is indicated by the red diamond. The 68% (dashed curve) and 95% (solid curve) CL contours are shown for the observed result. The orange circle indicates the SM prediction.




FIGURE 5.39: Comparison of observed 95% CL intervals for the Wilson coefficients c_{tZ} (upper panel) and c_{tZ}^I (lower panel). For the CMS $t\bar{t}Z$ [54], $t\bar{t}\gamma$ $l+jets$ [6] and $t\bar{t}\gamma$ dilepton results, the limits are shown for the case where all other considered Wilson coefficients are fixed to zero. The dashed lines indicate the corresponding result of the $t\bar{t}\gamma$ $l+jets$ and dilepton combination. For the CMS result based on $t\bar{t}Z$ and tZq events [91], as well as from a global fit to the LHC, LEP, and Tevatron data [99], the limits where all other considered Wilson coefficients are fixed to zero are shown with solid lines, and the marginalized limits from the full fits are shown with dashed-and-dotted lines.

Chapter 6

Conclusion and outlook

This thesis documents the measurement of $t\bar{t}\gamma$ production in the dilepton final state ($e^\pm\mu^\mp$, e^+e^- , and $\mu^+\mu^-$) using 138 fb^{-1} of proton-proton collision data recorded using the CMS detector. A fiducial phase space is defined at the particle level, and a reconstruction level selection with similar selection criteria aimed at reducing the background contributions is applied. The background coming from nonprompt photons is estimated from a sideband in data and the $Z\gamma$ prediction obtained from simulation is corrected using a control region. Within the fiducial phase space a cross section of $175.2 \pm 2.5\text{ (stat)} \pm 6.3\text{ (syst)}\text{ fb}$ is obtained using a fit to the p_T distribution of the selected photon, which is found to be compatible with the SM prediction of $155 \pm 27\text{ fb}$ within the uncertainties. The result benefits from the relatively small uncertainties in the lepton selection, a recently improved measurement of the 2016 period luminosity, and the fact that the e^+e^- and $\mu^+\mu^-$ final states were included despite being less pure than the $e^\pm\mu^\mp$ channel. Differential distributions are obtained in the same fiducial region obtained through an unregularized unfolding procedure. This includes results for various kinematic distributions related to the photon, leptons, and jets in the events. All are found to agree with the SM predictions within the uncertainties, with some deviations being attributed to the limited precision afforded by the leading order signal prediction. The results are interpreted within the framework of SMEFT, from which constraints on the c_{tZ} and c_{tZ}^I Wilson coefficients are extracted. A combined interpretation of these results and the similar measurement in the lepton+jets final state is also performed, leading to the most stringent limits on the aforementioned coefficients available so far.

Conceptually the most straightforward way to improve upon measurements in CMS is by increasing the centre-of-mass energy, increasing the total recorded luminosity, and reducing the impact of systematic uncertainties. In practice even small increases in the beam energy are technically difficult to achieve, and reducing systematic uncertainties requires the work of several experts over long periods of time. With the second long shutdown period of the LHC now having ended CMS and ATLAS are preparing for new data taking periods with an increased centre-of-mass energy of $\sqrt{s} = 13.6\text{ TeV}$, with a planned integrated luminosity of around twice that collected between 2016 and 2018. A potential large jump in centre of mass energy could be made with the construction of the so-called Future Circular Collider (FCC). The end goal for this machine with a circumference of 100 km would be proton-proton collisions at a centre-of-mass energy of $\sqrt{s} = 100\text{ TeV}$. Such an energy increase would increase the cross section for $t\bar{t}\gamma$ production by a factor of 5, as illustrated in figure 6.1. The figure also shows the dependence on centre-of-mass energy for other processes like $t\bar{t}Z$ production, for which the cross section benefits even more from the energy increase.

Similar to the LEP experiment which was located in the LHC tunnel before being replaced by the current hadron collider the FCC tunnel would potentially first be used to house an electron-positron collider (FCC-ee) with a centre-of-mass energy range between 90 and 400 GeV. While $t\bar{t}\gamma$ production is possible within this energy range such an experiment would allow us to study the $t\text{-}\gamma$ and $t\text{-}Z$ couplings in a very direct way using $e^+e^- \rightarrow Z/\gamma^* \rightarrow t\bar{t}$ events [171]. Lepton-lepton collisions inherently lead to very “clean” events, where nearly every event can be fully reconstructed with high precision, and for which no trigger system is needed. The $t\bar{t}$ production process could even be targeted specifically using a dedicated run at $\sqrt{s} = 365\text{ GeV}$, and contributions from $t\text{-}\gamma$ and $t\text{-}Z$ couplings disentangled by using polarized beams.

The $t\bar{t}\gamma$ process also gains an additional use when studied in electron-positron collisions, as would also be the case in the proposed International Linear Collider (ILC) and the Compact Linear Collider (CLIC). The most straightforward way to perform a top quark mass measurement using such a collider would be to perform a scan over the centre-of-mass energy around the production threshold of top quark pair production. Since leptons are not composite particles the energy entering the collisions is precisely known, such that the dependence of the $t\bar{t}$ production cross section on the invariant mass of the incoming particles can be measured directly. Running the colliders at various different beam




FIGURE 6.1: Cross sections for various $t\bar{t}X(X)$ production processes as functions of the centre-of-mass energy in proton-proton colliders [170].

energies can however be circumvented by looking at $t\bar{t}\gamma$ events instead. In $t\bar{t}\gamma$ events with a photon coming from initial state radiation the photon carries away a certain amount of energy, so that the energy available for the production of a $t\bar{t}$ pair is reduced [172]. Keeping the convention of writing the square of the invariant mass as s this remaining energy s' can be written as:

$$s' = s \left(1 - \frac{2E_\gamma}{\sqrt{s}} \right) \quad (6.1)$$

In lepton-lepton collision events the photon radiated in these events can reliably be reconstructed and selected, so that a differential cross section measurement as a function of $\sqrt{s'}$ is possible. The predicted and hypothetical measured spectra for such a measurement are shown in figure 6.2. The strong dependence of the distribution on the top quark mass is readily apparent, illustrating the sensitivity such a measurement can offer. The observable is furthermore only dependent on the photon energy and the invariant mass in the event, and is completely inclusive with regards to the top quark decay products. Using this method the top quark mass could be determined with a precision of 110 MeV in the initial stage of CLIC using an integrated luminosity of 1 ab^{-1} at $\sqrt{s} = 380 \text{ GeV}$, and with a precision of around 150 MeV at the ILC, with an integrated luminosity of 4 ab^{-1} at $\sqrt{s} = 500 \text{ GeV}$ [172].

As the structure of this text already suggested the main point of interest is at this point not necessarily $t\bar{t}\gamma$ production itself, but the EFT framework in which it can be interpreted in combination with other measurement sensitive to electroweak couplings. Effective field theory is at this point possibly one of the most exciting topics in experimental particles physics, as it allows us to cast the net wider than ever. Its hurdles come in the form of the large number of operators to be constrained, the sensitivities to be optimized within individual analyses, and the growing number of precision measurements to be interpreted simultaneously. That is why most of the outlook I can provide on standard model effective field theory has already been described in Chapter 4. Any ideas that were floating around at the time when I was working on my master thesis on the topic of EFT have since been made concrete in a first publication, or are at the very least being worked on. One other ambitious project being started in collaboration between CMS and ATLAS is however worth mentioning.

The main goal in the context of EFT is the simultaneous extraction of constraints on a wide range of operators, using an equally wide range of measurements. Such an endeavour performed by theorists was described in the introduction to EFT. The framework developed in the context of said publication was made publicly available, so that measurements can be added on to extend the interpretation. In the end the limiting factor in this approach is that only reinterpretations of published results are possible, in which correlations can at most be approximated.

Within CMS and ATLAS efforts are under way to make combined interpretations of different measurements performed by both experiments possible, taking into account the full sets of uncertainties and correlations. Such combinations are typically referred to as “likelihood based” combinations, and could be accomplished by relying on the fact that the fitting tools used in CMS and ATLAS produce similar file formats. These efforts are organized by the LHC TOP and LHC EFT working groups. Such a coordinated effort between multiple groups from two different experiments is organizationally challenging, so that a priori the only clear part of the plan is the end goal. For many measurements the likelihood information required to perform full combinations are no longer available, with the data being made public being limited to inclusive or differential measurements with the associated total uncertainties.

The current plans of the working groups are roughly as follows: Combined reinterpretations of CMS and ATLAS measurements are to be performed using publicly available results. The results obtained in this way will not be much different from what external theorists have been able to produce so far, but the combination exercise creates a context in which communication between the collaborations is promoted. Within this context decisions can then be made regarding common conventions, systematic correlations, fitting tools and methods, and the format of potential future publications. Once a common ground has been established people performing precision measurements in the future can provide all needed information to include their results in a full combination. In the context of top physics this should lead to the most stringent constraints on the electroweak couplings of the top quark achievable using the available measurements.

For $t\bar{t}\gamma$ a first step would be a combined interpretation of the CMS results obtained in the lepton+jets and dilepton channels, and the ATLAS measurement of $t\bar{t}\gamma$ and $tW\gamma$ production in the $e\mu$ channel [5]. In section 4.2, where the selection of operators most relevant to $t\bar{t}\gamma$ was performed, the close connection between the $t\text{-}\gamma$ and $t\text{-}Z$ couplings was illustrated. This link between the $t\text{-}\gamma$ and $t\text{-}Z$ couplings of course also implies a link between $t\bar{t}\gamma$ and $t\bar{t}Z$ processes. Potentially this synergy could be exploited by performing a combination in two directions: between CMS and ATLAS, and between the $t\bar{t}\gamma$ and $t\bar{t}Z$ processes.




FIGURE 6.2: The predicted differential cross section for $t\bar{t}$ production in association with a photon as a function of $\sqrt{s'}$ for two hypothetical top quark mass scenarios (top), and pseudo-data for CLIC and ILC illustrating what measurements of this differential cross section could look like with realistic integrated luminosity values [172].

*Ik vroeg eens aan een karper
 Of hij ooit had nagegaan
 Waarom de vijver donker wordt
 Bij elke nieuwe maan
 En de karper zei "maakt mij niks uit
 Wie 't licht heeft uitgedaan
 Zolang de modder maar vanonder zit
 En het water bovenaan"*

*Ik vroeg eens aan een knotwilg
 Of hij zich ooit had afgevraagd
 Waarom tijdens de zomer
 Het waterpeil verlaagt
 En de knotwilg zei "maakt mij niks uit
 Wie 't water heeft verjaagd
 Zolang mijn hart van hout is
 En de klei mijn wortels draagt"*

*Ik vroeg eens aan de regen
 Of hij ooit had stilgestaan
 Bij de vraag wat hij nog betekent
 Zonder zee of oceaan
 "Maakt mij niks uit", zo sprak de regen
 "Daarvoor laat ik geen traan
 Van zodra ik niet meer regen
 Zal de regenboog er staan"*

*En de regen liep
 Naar de vijver diep
 En niemand vroeg zich af
 Wie de wereld schiep
 Want de karper sliep
 En de knotwilg sliep*

*Toen vroeg ik aan mezelf
 En ik pijnigde m'n brein
 Of alles wat ik deed en liet
 Niet te groot was of te klein
 En ik dacht, maakt mij niks uit
 Want niks is vastomlijnd
 Kijk de zon hoeveel
 Ze ook naast de aarde schijnt*

Ik vroeg eens aan de regen - Urbanus

UNIVERSITEIT GENT

Summary

A measurement of $t\bar{t}\gamma$ production in the dilepton channel using 138 fb^{-1} of proton-proton collision data recorded by the CMS experiment at $\sqrt{s} = 13\text{ TeV}$ was presented. The signal is defined such that it includes $t\bar{t}$ events with an additional photon with $p_T > 20\text{ GeV}$, that can be radiated anywhere from initial state radiation to any of the decay products of the top quarks. This measurement is performed in a selection with two opposite-charge leptons, identified using a multivariate discriminator optimized to reject nonprompt leptons. The invariant mass of the two leptons $m(l\bar{l})$ is required to be $> 20\text{ GeV}$, and should additionally not be close to the Z boson mass in order to reject the Drell-Yann background. This corresponds to the requirement $|m(l\bar{l}) - m_Z| > 15\text{ GeV}$, but is applied only for ee and $\mu\mu$ events. The same requirement is imposed on the invariant mass of the two leptons and photon $|m(l\bar{l}\gamma)|$. Events are required to have exactly one photon in the barrel region of the detector passing a set of identification criteria, with no additional photons passing a looser version of these requirements.

Selected photons are categorized as prompt or nonprompt based on a matching procedure to the nearest generator particle with similar energy, and this promptness definition of the photon is what separates the signal from $t\bar{t}$ background events. The background originating from processes entering the selection through a nonprompt photon are estimated from data using a parametrized ABCD method for which a nonprompt photon enriched sideband is created by relaxing and inverting some of the photon selection criteria. $Z\gamma$ production forms the dominant background in the ee and $\mu\mu$ channels, and is predicted using Monte Carlo simulation, but with sideband based corrections for its N_j/N_b shape applied.

An inclusive cross section is obtained within the fiducial phase space using a maximum likelihood fit to the photon p_T distribution, with a value of $175.2 \pm 2.5\text{ (stat)} \pm 6.3\text{ (syst) fb}$. This value is compatible with the next-to-leading order standard model prediction of $153 \pm 27\text{ fb}$ within the uncertainties. Differential results are obtained through an unfolding procedure for variables ranging from photon, lepton, and jet kinematics, to the angles between these objects, and are compared to two theory predictions obtained using different parton shower models.

From the measurement of $t\bar{t}\gamma$ production constraints can be put on SMEFT operators that could modify electroweak dipole moments of the top quark. Using a particle level reweighting procedure the effect of the c_{tZ} and c_{tZ}^I wilson coefficients on the yield and $p_T(\gamma)$ shape are modelled, after which constraints on their magnitude are extracted from a likelihood fit. A combined interpretation of these results with those obtained in the lepton+jets channel is performed, leading to the most stringent constraint on these Wilson coefficients to date.

UNIVERSITEIT GENT

Samenvatting

Een meting van $t\bar{t}\gamma$ productie in het twee-lepton kanaal met 138 fb^{-1} aan proton-proton botsing data verzameld door het CMS experiment bij $\sqrt{s} = 13 \text{ TeV}$ werd beschreven. Het signaal is gedefinieerd als de productie van een paar top en antitop quarks met een extra foton met $p_T > 20 \text{ GeV}$, dat eerder waar uitgezonden kan worden vanaf de inkomende partonen tot de vervalproducten van de top quarks. Deze meting werd uitgevoerd in een selectie met twee leptonen met tegengestelde lading, geïdentificeerd met behulp van een multivariate discriminator die geoptimaliseerd werd om niet-prompte leptonen te verwerpen. Aan de invariante massa van de twee leptonen $m(ll)$ wordt de $> 20 \text{ GeV}$ voorwaarde opgelegd, en deze waarde mag bijkomend ook niet dicht liggen bij de invariante massa van het Z boson zodat Drell-Yann achtergrond verwijderd wordt. Dit komt overeen met de voorwaarde dat $|m(ll) - m_Z| > 15 \text{ GeV}$, die enkel word toegepast in ee en $\mu\mu$ events. Dezelfde eis wordt ook opgelegd aan de invariante massa van de twee leptonen en het foton $|m(ll\gamma)|$. Er wordt geëist dat exact één foton in de centrale regio van de detector wordt gevonden dat aan een aantal discrete identificatie criteria voldoet, met geen enkel bijkomstig foton dat aan een minder strenge set van criteria voldoet.

Geselecteerde fotonen worden onderverdeeld in prompt of niet-prompt op basis van een toekenningsprocedure aan het dichtsbijzijnde generator deeltje met een gelijkaardige energie, en het is deze definitie van “promptheid” die het onderscheid maakt tussen het signaal en de $t\bar{t}$ achtergrond. De achtergrond die voortkomt uit processen die doorheen de selectie geraken doormiddel van een niet-prompt foton wordt voorspeld met behulp van een zogenaamde ABCD methode waarvoor een selectie vrijrijkt in niet-prompte fotonen wordt gevormd door enkele van de foton selectie criteria minder streng te maken of om te keren. $Z\gamma$ productie vormt de dominante achtergrond in de ee en $\mu\mu$ kanalen, en wordt voorspeld met behulp van Monte Carlo simulatie, maar dan met correcties voor de N_j/N_b vorm toegepast.

De inclusieve crossectie wordt bekomen door het uitvoeren van een maximum likelihood fit met de p_T distributie als input, en een waarde van $175.2 \pm 2.5 \text{ (stat)} \pm 6.3 \text{ (syst)} \text{ fb}$ wordt gevonden. Deze waarde is compatibel met de op-een-na hoogste orde standaard model voorspelling van $153 \pm 27 \text{ fb}$ binnen de onzekerheden. Differentiële resultaten worden bekomen doormiddel van een soort van deconvolutie operatie. Zo worden resultaten verkregen voor de kinematische variabelen van het foton, de leptonen, jets, en de hoeken tussen deze deeltjes. Deze resultaten worden vergeleken met die verkregen uit 2 theoretische voorspellingen met verschillende parton stralingsmodellen.

Uit de meting van $t\bar{t}\gamma$ productie kunnen limieten op SMEFT operatoren gehaald worden die de elektrozwakke dipool momenta van de top quarks zouden kunnen aanpassen. Met behulp van een simulatie-niveau herschalingsprocedure kan het effect van de c_{tZ} en c_{tZ}^I Wilson coëfficiënten op het $p_T(\gamma)$ profiel gemodelleerd worden, waarna grenzen op de waarden van deze coëfficiënten kunnen gevonden worden doormiddel van een maximum likelihood fit. Een gecombineerde interpretatie van deze resultaten met die gevonden in het lepton+jets kanaal werd ook uitgevoerd, waaruit de tot noch toe sterkste grenzen op deze twee Wilson coefficienten bekomen werden.

Appendix A

Charge misidentification prediction for a $t\bar{t}t\bar{t}$ search

In addition to the $t\bar{t}\gamma$ analysis, in which the dominant nonprompt photon background is estimated using control regions in data, I also performed a similar estimate of the background coming from electrons with misidentified charges in the context of a search for 4 top quark production being performed by the UGent CMS team. This background is known to not be well modelled in MC simulation due to the strong dependence on the physical features of the detector. I will briefly describe the signal process and how charge misidentification forms a significant background in its measurement, before going into detail on how this misidentification occurs and how the background arising from it can be predicted from data.

A.1 Production and measurement of 4 top quarks

The simultaneous production of four top quarks ($t\bar{t}t\bar{t}$) is an as of yet undiscovered process predicted by the standard model, and one of the rarest processes currently being searched for by CMS and ATLAS.¹ The production of four top quarks can occur in a multitude of ways, some examples of which are illustrated in the form of Feynman diagrams as shown in figure A.1. Both gg and $q\bar{q}$ initial states contribute, and the dominant contributions to the production involve only QCD interactions including triple and quartic gluon couplings. Diagrams involving intermediate Z and H bosons also appear at leading order, but only lead to subdominant contributions.




FIGURE A.1: Example LO Feynman diagrams for $t\bar{t}t\bar{t}$ production. The process can occur through both gg and $q\bar{q}$ initial states and can involve triple and quartic gluon couplings, electroweak couplings, and Higgs couplings.

The decay modes of the four top quarks give rise to final states with zero to four prompt leptons and high to very high jet multiplicities. Typically analyses focus on just one or a few of these final states at a time. Many such searches have been performed by both CMS and ATLAS [173–180]. The latest result by ATLAS [180] quotes an expected significance of 2.6 standard deviations, and an observed significance of 4.7 standard deviations, obtained using collision data collected in 2016, 2017, and 2018, and considering the 1, 2, and 3 lepton final states.

¹There is in principle little doubt that this process exists, but in order for a process to be “discovered” it needs to be observed with a significance of at least 5 standard deviations.

The analysis being performed in the UGent CMS group to which I only contributed the measurement of the charge misidentification rate is highly complex due to the inclusion of 2, 3, and 4 lepton final states, multiple control regions, and the application of multivariate analysis techniques. A full description is therefore beyond the scope of this thesis. I therefore focus only on the final state in which charge misidentification is relevant, namely the dilepton channel, and only mention object selection requirements to the degree they are relevant to the charge misidentification background.

Leptons are selected using a loose working point of an updated version of the MVA discriminator used in the $t\bar{t}\gamma$ analysis. In the search for four top production in the dilepton final state only events with two leptons with the same charge are considered for the signal region. Oppositely charged dilepton pairs do form a final state of $t\bar{t}t\bar{t}$, but this channel is highly dominated by $t\bar{t}$ events. Even with this selection applied, as well as additional requirements on the jet multiplicity, b jet multiplicity, and the total transverse momentum from jets in the event, the signal region still suffers from large backgrounds coming from for example $t\bar{t}Z$ and $t\bar{t}W$ production and nonprompt leptons. The top graph of figure A.2 shows the predicted yield of various contributions in the signal region targeting final states with two leptons with the same charge. The category labelled as “ChargeMisID” contains events from any process in which at least one of the leptons has a reconstructed charge opposite from the generator particle it is matched to. The charge of an electron is determined by looking at the curvature of its track and the position of the associated ECAL supercluster with respect to the track. To reduce this background a charge consistency requirement is applied for electrons, which means that the charges reconstructed using pixel tracker information, full track information, and the position of the supercluster with respect to the track are required to be consistent. Despite this extra requirement the contribution coming from this background is not negligible, and since the trust in the modelling of this background is limited the goal is now to estimate this background from data.

A.2 Charge misidentification of electrons and muons

Applying a charge consistency requirement already reduces the fraction of electrons for which the charge is misidentified significantly with respect to what one would find using only identification requirements aimed at removing nonprompt leptons. Nevertheless, this charge misidentification still occurs with a low rate. This causes events with oppositely charged leptons, mostly from $t\bar{t}$ production, to contribute in signal regions where two selected leptons with the same charge are required. This is mainly caused by electrons radiating bremsstrahlung photons as they pass through the detector material. These photons can then convert into electron-positron pairs, whose tracks are hard to distinguish from the one of the electron which radiated the photon. Due to their larger mass this bremsstrahlung is much less of a problem for muons which, in addition to the better track reconstruction offered by the muon systems, leads to a charge misidentification rate that is about 2 orders of magnitude smaller than for electrons.² The charge misidentification rate is therefore important in the ee and $e\mu$ final states, but negligible for the $\mu\mu$ final state.

The rate at which electron charges are misidentified shows a strong dependence on the pseudorapidity and transverse momentum of the particle. The former dependence can be understood as a consequence of the fact that a particle moves more parallel to the magnetic field as its pseudorapidity increases, while also having to pass through a larger amount of detector material. The dependence on the transverse momentum is more easily understood, as higher momenta directly lead to less strongly curved tracks.

A.3 Charge misidentification background estimate

The charge misidentification rate is measured in simulated DY samples within a selection requiring exactly 2 selected electrons, forming an invariant mass within 15 GeV of the mass of the Z boson. Both electrons are also required to be prompt using generator information. This promptness requirement does not remove all electrons from photon conversions, e.g. if a photon produced as initial state radiation converts into electrons, these electrons are considered to be prompt. A charge misidentified

²The power radiated by accelerated charged particles is described by the Larmor formula, in which a mass dependence of the order of m^{-4} to m^{-6} appears depending on the angle between the acceleration and the velocity vectors of the particle.




FIGURE A.2: Top: Signal and background yields in the same-charge dilepton signal region as predicted by MC simulation, as a function of the flavour of the selected leptons. Bottom: Signal and background yields in the same-charge dilepton signal region of the $t\bar{t}t\bar{t}$ analysis, as a function of the leading lepton p_T . The left graph shows the charge misidentification background, labelled as “ChargeMisID”, as predicted by MC simulation, while in the right graph this contribution is estimated from data using the method described here. These are preliminary results kindly provided by Niels Van den Bossche.

electron is therefore defined as a reconstructed electron for which the charge is opposite to that of the generator particle to which it is matched. Cases in which the electron is matched to a neutral generator particle, which is exactly what occurs for prompt electrons from photon conversions, are not counted as misidentified charges.

The charge misidentification rate is defined as the number of electrons for which the charge is misidentified, as defined above, divided by the number of electrons for which this is not the case. It should be noted that while the preceding sentence accurately describes the approach used, simply including all electrons in the denominator, or leaving out electrons from conversions in the denominator, has a negligible impact on the obtained result. This is true because the rate of charge misidentification and photon conversions (within the relevant object and event selection) is so small. The measured charge misidentification rates are given in figure A.3 as a function of the p_T and $|\eta|$ of the leptons. The measurement is performed for each data taking year separately, as the rates are highly dependent on the detector configuration and performance. One significant change in between years comes from the

fact that an additional layer was added to the pixel detector during the short shutdown period before the data taking period of 2017 started, leading to a reduction in the charge misidentification rate with respect to 2016. Since the rate is so low statistics for electrons with a misidentified charge are inherently limited. Therefore the binning is chosen to be coarse enough to keep the statistics in each bin high, while still capturing the p_T and $|\eta|$ dependence sufficiently. In order to reduce statistical uncertainties several DY samples are also combined while performing this measurement, as only the ratios are of interest and the total yield does not need to agree with data.

When using these rates to perform an estimate or closure test they are applied as weights for each of the electrons in an opposite-charge selection. If this is done in the $e\mu$ channel only one weight needs to be applied. When this is performed in the ee channel, as will be the case in the closure tests, either or both of the two electrons can have a misidentified charge. The total weight to be applied to the event is therefore given by:

$$w = P(e^-)(1 - P(e^+)) + P(e^+)(1 - P(e^-)) \approx P(e^-) + P(e^+) \quad (\text{A.1})$$

Where P is the (measured) probability of an electron/positron having a misidentified charge. The approximation in the equation relies on the fact that the probability is around 1% at most, with values down to as low as $\mathcal{O}(0.001\%)$.

In order to check that the chosen binning is sufficiently fine to model the kinematics in the prediction a pseudo-closure test can be performed by applying the estimation weights to simulated DY and $t\bar{t}$ samples in an opposite-charge selection, and comparing the resulting distributions to those obtained directly using a same-charge selection. The results of this are shown in figures A.4, A.5, and A.6. While closure in terms of the total yield will always be obtained in a check like this, the good agreement as a function of the lepton kinematic variables shows that the estimation maps can model these sufficiently well. The shift in the invariant mass distribution is expected, as electrons with a misidentified charge are biased towards having lost more energy due to bremsstrahlung.

A true closure test in data can now be performed by applying the measured weights to opposite-charge events with the same selection mentioned above (without the promptness requirement, which uses generator information):

- exactly 2 selected electrons
- $|m(ee) - m(Z)| < 15 \text{ GeV}$

An underprediction of the order of 50% is found for 2017 and 2018, while 2016 seems to be well estimated, consistent with what was found in Ref. [181]. Such differences are to be expected, as the charge misidentification rate is highly dependent on the material budget, so that even a small difference between the simulated detector and the real one will produce such a discrepancy. From the difference in the total observed and predicted yields flat correction factors for the charge misidentification rates measured in each data set are derived, to be applied when performing the actual prediction. The changes made to the pixel tracker between 2016 and 2017 are the most obvious cause for this change, which aligns with the fact that the 2017 and 2018 correction factors match almost perfectly. The result of these closure tests, with the correction factors applied, are shown in figures A.7, A.8, and A.9. An uncertainty band of 15% is drawn on the ratio pads to illustrate that a flat systematic uncertainty of this size is sufficient to cover the uncertainty on this prediction in the final fit. The correction factors themselves are given in table A.1. The right graph of figure A.2 shows the dilepton same-charge signal region with the charge misidentification contribution estimated from data using the method described here, compared to the prediction purely from MC simulation on the left.

The dependence of the charge misidentification rate on the pseudorapidity of the electron is very large, so that one could wonder if an additional parametrization as a function of the p_T is necessary. Conversely, one could assume that the $|\eta|$ distributions of leptons in different processes are not so different, so that only binning as a function of the p_T might seem reasonable. Below the flaws in this line of thought are explained.

The 2D charge misidentification maps show that while the $|\eta|$ dependence of the rate is very large, the dependence on the p_T is also large (typically a factor 2 for low vs high p_T electrons). Figure A.10 shows the p_T and $|\eta|$ distributions of the leading and subleading leptons in DY and $t\bar{t}$ events overlaid

TABLE A.1: Ratio of the observed yield in the same charge final state (in a 2 electron, DY enriched selection) to the prediction obtained using the MC based estimation weights. These factors are applied on top of the estimation weights in the closure test in data, and are to be applied when performing the actual charge misidentification background estimation in the $t\bar{t}t\bar{t}$ analysis.

year	correction factor
2016	0.966
2017	1.509
2018	1.515

and normalized to unity. This is shown within a selection requiring exactly 2 prompt electrons, with an invariant mass close to the Z boson mass. For the $t\bar{t}$ sample this invariant mass requirement has no physical basis, but if anything this requirement would bias the distributions to look more alike, which is the opposite of the point being made here.

Figure A.11 then shows the charge misidentification rates measured in DY and $t\bar{t}$, as a function of the p_T and $|\eta|$ of the electrons separately. The observed differences are the result of the situation described above, and underline why a 2D parametrization is needed. The argument can be summarized as follows:

- The $|\eta|$ distributions are only moderately different between processes, but the $|\eta|$ dependence of the charge misidentification rate is very large.
- The dependence of the charge misidentification rate on the p_T of the electrons is much smaller, but the p_T spectra of the leptons are very different between processes.

In the end this means that opting for a parametrization using both p_T and $|\eta|$ is the lesser of two evils, and a coarser binning needs to be used for each of the parameters than if only one were used, in order to limit the statistical uncertainty in each bin. The main purpose of the systematic uncertainty of 15% assigned to the estimated background is intended to cover the inability of the coarsely binned estimation map to fully model the kinematics of this background.




FIGURE A.3: Electron charge misidentification rates based on truth information measured in 2016, 2017 and 2018 DY simulation samples.




FIGURE A.4: Pseudo-closure test of the charge misidentification estimate performed in DY and $t\bar{t}$ simulation samples for 2016. The grey band in the ratio graph visualizes the 15% uncertainty assigned to this estimate.




FIGURE A.5: Pseudo-closure test of the charge misidentification estimate performed in DY and $t\bar{t}$ simulation samples for 2017. The grey band in the ratio graph visualizes the 15% uncertainty assigned to this estimate.




FIGURE A.6: Pseudo-closure test of the charge misidentification estimate performed in DY and $t\bar{t}$ simulation samples for 2018. The grey band in the ratio graph visualizes the 15% uncertainty assigned to this estimate.




FIGURE A.7: Closure test of the charge misidentification estimate performed in data within a DY enriched selection, for 2016. The grey band in the ratio graph visualizes the 15% uncertainty assigned to this estimate.




FIGURE A.8: Closure test of the charge misidentification estimate performed in data within a DY enriched selection, for 2017. The grey band in the ratio graph visualizes the 15% uncertainty assigned to this estimate.




FIGURE A.9: Closure test of the charge misidentification estimate performed in data within a DY enriched selection, for 2018. The grey band in the ratio graph visualizes the 15% uncertainty assigned to this estimate.




FIGURE A.10: Pseudorapidity and transverse momentum of the leading and subleading leptons (electrons) compared between DY and $t\bar{t}$ within the same selection.




FIGURE A.11: Charge misidentification rates measured in DY and $t\bar{t}$ as a function of the lepton pseudorapidity (left) and transverse momentum (right).

Bibliography

- [1] T. Aaltonen et al. “Evidence for $t\bar{t}\gamma$ Production and Measurement of $\sigma_{t\bar{t}\gamma}/\sigma_{t\bar{t}}$ ”. In: *Phys. Rev. D* 84 (2011), p. 031104. DOI: [10.1103/PhysRevD.84.031104](https://doi.org/10.1103/PhysRevD.84.031104). arXiv: [1106.3970](https://arxiv.org/abs/1106.3970) [hep-ex].
- [2] G. Aad et al. “Observation of top-quark pair production in association with a photon and measurement of the $t\bar{t}\gamma$ production cross section in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector”. In: *Phys. Rev. D* 91.7 (2015), p. 072007. DOI: [10.1103/PhysRevD.91.072007](https://doi.org/10.1103/PhysRevD.91.072007). arXiv: [1502.00586](https://arxiv.org/abs/1502.00586) [hep-ex].
- [3] A. Sirunyan et al. “Measurement of the semileptonic $t\bar{t} + \gamma$ production cross section in pp collisions at $\sqrt{s} = 8$ TeV”. In: *JHEP* 10 (2017), p. 006. DOI: [10.1007/JHEP10\(2017\)006](https://doi.org/10.1007/JHEP10(2017)006). arXiv: [1706.08128](https://arxiv.org/abs/1706.08128) [hep-ex].
- [4] M. Aaboud et al. “Measurements of inclusive and differential fiducial cross-sections of $t\bar{t}\gamma$ production in leptonic final states at $\sqrt{s} = 13$ TeV in ATLAS”. In: *Submitted to: Eur. Phys. J.* (2018). arXiv: [1812.01697](https://arxiv.org/abs/1812.01697) [hep-ex].
- [5] G. Aad et al. “Measurements of inclusive and differential cross-sections of combined $t\bar{t}\gamma$ and $tW\gamma$ production in the $e\mu$ channel at 13 TeV with the ATLAS detector”. In: *JHEP* 09 (2020), p. 049. DOI: [10.1007/JHEP09\(2020\)049](https://doi.org/10.1007/JHEP09(2020)049). arXiv: [2007.06946](https://arxiv.org/abs/2007.06946) [hep-ex].
- [6] A. Tumasyan et al. “Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the single-lepton channel and EFT interpretation at $\sqrt{s} = 13$ TeV”. In: *JHEP* 12 (2021), p. 180. DOI: [10.1007/JHEP12\(2021\)180](https://doi.org/10.1007/JHEP12(2021)180). arXiv: [2107.01508](https://arxiv.org/abs/2107.01508) [hep-ex].
- [7] A. Tumasyan et al. “Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the dilepton channel and effective field theory interpretation in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: (Jan. 2022). arXiv: [2201.07301](https://arxiv.org/abs/2201.07301) [hep-ex].
- [8] W. Buchmüller and D. Wyler. “Effective Lagrangian Analysis of New Interactions and Flavour Conservation”. In: *Nucl. Phys. B* 268 (1986), p. 621. DOI: [10.1016/0550-3213\(86\)90262-2](https://doi.org/10.1016/0550-3213(86)90262-2).
- [9] B. Grzadkowski et al. “Dimension-Six Terms in the Standard Model Lagrangian”. In: *JHEP* 10 (2010), p. 085. DOI: [10.1007/JHEP10\(2010\)085](https://doi.org/10.1007/JHEP10(2010)085). arXiv: [1008.4884](https://arxiv.org/abs/1008.4884) [hep-ph].
- [10] CMS Collaboration. *Measurement of the inclusive and differential $t\bar{t}\gamma$ cross sections in the dilepton channel and effective field theory interpretation in proton-proton collisions at $\sqrt{s} = 13$ TeV*. HEPData (collection). <https://doi.org/10.17182/hepdata.113657>. 2022.
- [11] D. Dobur et al. “Photon radiation effects in kinematic reconstruction of top quarks”. In: *JINST* 16.12 (2021), P12001. DOI: [10.1088/1748-0221/16/12/P12001](https://doi.org/10.1088/1748-0221/16/12/P12001). arXiv: [2107.07586](https://arxiv.org/abs/2107.07586) [hep-ex].
- [12] G. Mestdach. “Inclusive and differential $t\bar{t}\gamma$ measurement in the dilepton channel and effective field theory interpretation”. In: *14th International Workshop on Top Quark Physics*. Dec. 2021. arXiv: [2112.12541](https://arxiv.org/abs/2112.12541) [hep-ex].
- [13] Wikimedia Commons. *Standard Model of Elementary Particles*. 2018. URL: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles_Anti.svg.
- [14] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (9 1964), pp. 321–323. DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.321>.
- [15] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (16 1964), pp. 508–509. DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.13.508>.
- [16] K. Kumericki. *Feynman Diagrams for Beginners*. 2016. arXiv: [1602.04182](https://arxiv.org/abs/1602.04182) [physics.ed-ph].
- [17] Wikimedia Commons. *All Standard Model Vertices*. 2020. URL: https://commons.wikimedia.org/wiki/File:All_Standard_Model_Vertices.pdf.
- [18] CMS Collaboration. *Summaries of CMS cross section measurements*. URL: <https://twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsCombined>.

- [19] T. et al Aaltonen. “High-precision measurement of the W boson mass with the CDF II detector”. In: *Science* 376.6589 (2022), pp. 170–176. DOI: [10.1126/science.abk1781](https://doi.org/10.1126/science.abk1781). eprint: <https://www.science.org/doi/pdf/10.1126/science.abk1781>. URL: <https://www.science.org/doi/abs/10.1126/science.abk1781>.
- [20] S. Chatrchyan et al. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021). arXiv: [1207.7235 \[hep-ex\]](https://arxiv.org/abs/1207.7235).
- [21] G. Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020). arXiv: [1207.7214 \[hep-ex\]](https://arxiv.org/abs/1207.7214).
- [22] R.L. Workman et al. “Review of Particle Physics”. In: (). to be published (2022).
- [23] W. Greiner et al. *Structure and Dynamics of Elementary Matter*. Jan. 2004. ISBN: 978-1-4020-2446-7. DOI: [10.1007/978-1-4020-2705-5](https://doi.org/10.1007/978-1-4020-2705-5).
- [24] E. Corbelli and P. Salucci. “The Extended Rotation Curve and the Dark Matter Halo of M33”. In: *Mon. Not. Roy. Astron. Soc.* 311 (2000), pp. 441–447. DOI: [10.1046/j.1365-8711.2000.03075.x](https://doi.org/10.1046/j.1365-8711.2000.03075.x). arXiv: [astro-ph/9909252](https://arxiv.org/abs/astro-ph/9909252).
- [25] X. P. Wu et al. “A comparison of different cluster mass estimates: consistency or discrepancy ?” In: *Mon. Not. Roy. Astron. Soc.* 301 (1998), p. 861. DOI: [10.1046/j.1365-8711.1998.02055.x](https://doi.org/10.1046/j.1365-8711.1998.02055.x). arXiv: [astro-ph/9808179](https://arxiv.org/abs/astro-ph/9808179).
- [26] D. Clowe et al. “A direct empirical proof of the existence of dark matter”. In: *Astrophys. J. Lett.* 648 (2006), pp. L109–L113. DOI: [10.1086/508162](https://doi.org/10.1086/508162). arXiv: [astro-ph/0608407](https://arxiv.org/abs/astro-ph/0608407).
- [27] D. Vannerom. “Dark Matter searches with CMS”. In: *PoS DIS2019* (2019), p. 111. DOI: [10.22323/1.352.0111](https://doi.org/10.22323/1.352.0111).
- [28] B. T. Carlson. “Dark Matter searches with the ATLAS detector”. In: (2020). URL: <https://cds.cern.ch/record/2727868>.
- [29] T. Marrodán Undagoitia and L. Rauch. “Dark matter direct-detection experiments”. In: *J. Phys. G* 43.1 (2016), p. 013001. DOI: [10.1088/0954-3899/43/1/013001](https://doi.org/10.1088/0954-3899/43/1/013001). arXiv: [1509.08767 \[physics.ins-det\]](https://arxiv.org/abs/1509.08767).
- [30] A. Buras. “CP violation in electroweak interactions”. In: *Scholarpedia* 10.8 (2015). revision #151973, p. 11418. DOI: [10.4249/scholarpedia.11418](https://doi.org/10.4249/scholarpedia.11418).
- [31] P.A. Zyla et al. “Review of Particle Physics”. In: *PTEP* 2020.8 (2020), p. 083C01. DOI: [10.1093/ptep/ptaa104](https://doi.org/10.1093/ptep/ptaa104).
- [32] M. Kobayashi and T. Maskawa. “CP Violation in the Renormalizable Theory of Weak Interaction”. In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: [10.1143/PTP.49.652](https://doi.org/10.1143/PTP.49.652).
- [33] W. R. Innes. “Observation of a Dimuon Resonance Near 9.5-GeV in 400-GeV Proton-Nucleus Collisions”. In: *5th SLAC Summer Institute on Particle Physics : Quark Spectroscopy and Hadron Dynamics (Topical Conference last 3 days) (SSI 77)*. July 1977.
- [34] F. Abe et al. “Observation of top quark production in $\bar{p}p$ collisions”. In: *Phys. Rev. Lett.* 74 (1995), pp. 2626–2631. DOI: [10.1103/PhysRevLett.74.2626](https://doi.org/10.1103/PhysRevLett.74.2626). arXiv: [hep-ex/9503002](https://arxiv.org/abs/hep-ex/9503002).
- [35] F. Abe et al. “Evidence for top quark production in $\bar{p}p$ collisions at $\sqrt{s} = 1.8$ TeV”. In: *Phys. Rev. Lett.* 73 (1994), pp. 225–231. DOI: [10.1103/PhysRevLett.73.225](https://doi.org/10.1103/PhysRevLett.73.225). arXiv: [hep-ex/9405005](https://arxiv.org/abs/hep-ex/9405005).
- [36] I. I. Y. Bigi et al. “Production and Decay Properties of Ultraheavy Quarks”. In: *Phys. Lett. B* 181 (1986), pp. 157–163. DOI: [10.1016/0370-2693\(86\)91275-X](https://doi.org/10.1016/0370-2693(86)91275-X).
- [37] “A profile likelihood approach to measure the top quark mass in the lepton+jets channel at $\sqrt{s} = 13$ TeV”. In: (2022).
- [38] A. Sirunyan et al. “Observation of $t\bar{t}H$ production”. In: *Phys. Rev. Lett.* 120.23 (2018), p. 231801. DOI: [10.1103/PhysRevLett.120.231801](https://doi.org/10.1103/PhysRevLett.120.231801). arXiv: [1804.02610 \[hep-ex\]](https://arxiv.org/abs/1804.02610).
- [39] M. Aaboud et al. “Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector”. In: *Phys. Lett. B* 784 (2018), pp. 173–191. DOI: [10.1016/j.physletb.2018.07.035](https://doi.org/10.1016/j.physletb.2018.07.035). arXiv: [1806.00425 \[hep-ex\]](https://arxiv.org/abs/1806.00425).

- [40] A. Sirunyan et al. “Measurement of the top quark Yukawa coupling from $t\bar{t}$ kinematic distributions in the dilepton final state in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 102.9 (2020), p. 092013. DOI: [10.1103/PhysRevD.102.092013](https://doi.org/10.1103/PhysRevD.102.092013). arXiv: [2009.07123](https://arxiv.org/abs/2009.07123) [hep-ex].
- [41] G.K. Krintiras. “Study of the top quark production in complementary phase space regions and impact on PDFs in CMS”. In: *29th Rencontres de Blois on Particle Physics and Cosmology*. Dec. 2017. arXiv: [1712.06103](https://arxiv.org/abs/1712.06103) [hep-ex].
- [42] A. Giammanco. “Single top quark production at the LHC”. In: *Rev. Phys.* 1 (2016), pp. 1–12. DOI: [10.1016/j.revip.2015.12.001](https://doi.org/10.1016/j.revip.2015.12.001). arXiv: [1511.06748](https://arxiv.org/abs/1511.06748) [hep-ex].
- [43] M. Czakon and A. Mitov. “Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders”. In: *Comput. Phys. Commun.* 185 (2014), p. 2930. DOI: [10.1016/j.cpc.2014.06.021](https://doi.org/10.1016/j.cpc.2014.06.021). arXiv: [1112.5675](https://arxiv.org/abs/1112.5675) [hep-ph].
- [44] A. Kulesza et al. “Associated top quark pair production with a heavy boson: differential cross sections at NLO+NNLL accuracy”. In: *Eur. Phys. J. C* 80.5 (2020), p. 428. DOI: [10.1140/epjc/s10052-020-7987-6](https://doi.org/10.1140/epjc/s10052-020-7987-6). arXiv: [2001.03031](https://arxiv.org/abs/2001.03031) [hep-ph].
- [45] R. Frederix, D. Pagani, and M. Zaro. “Large NLO corrections in $t\bar{t}W^\pm$ and $t\bar{t}t\bar{t}$ hadroproduction from supposedly subleading EW contributions”. In: *JHEP* 02 (2018), p. 031. DOI: [10.1007/JHEP02\(2018\)031](https://doi.org/10.1007/JHEP02(2018)031). arXiv: [1711.02116](https://arxiv.org/abs/1711.02116) [hep-ph].
- [46] M. Aliev et al. “HATHOR: HAdronic Top and Heavy quarks crOSS section calculatoR”. In: *Comput. Phys. Commun.* 182 (2011), pp. 1034–1046. DOI: [10.1016/j.cpc.2010.12.040](https://doi.org/10.1016/j.cpc.2010.12.040). arXiv: [1007.1327](https://arxiv.org/abs/1007.1327) [hep-ph].
- [47] P. Kant et al. “HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions”. In: *Comput. Phys. Commun.* 191 (2015), pp. 74–89. DOI: [10.1016/j.cpc.2015.02.001](https://doi.org/10.1016/j.cpc.2015.02.001). arXiv: [1406.4403](https://arxiv.org/abs/1406.4403) [hep-ph].
- [48] *Observation of single-top-quark production in association with a photon at the ATLAS detector*. Tech. rep. Geneva: CERN, 2022. URL: <http://cds.cern.ch/record/2805217>.
- [49] A. Sirunyan et al. “Observation of Single Top Quark Production in Association with a Z Boson in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. Lett.* 122.13 (2019), p. 132003. DOI: [10.1103/PhysRevLett.122.132003](https://doi.org/10.1103/PhysRevLett.122.132003). arXiv: [1812.05900](https://arxiv.org/abs/1812.05900) [hep-ex].
- [50] G. Aad et al. “Observation of the associated production of a top quark and a Z boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *JHEP* 07 (2020), p. 124. DOI: [10.1007/JHEP07\(2020\)124](https://doi.org/10.1007/JHEP07(2020)124). arXiv: [2002.07546](https://arxiv.org/abs/2002.07546) [hep-ex].
- [51] S. Chatrchyan et al. “Observation of the associated production of a single top quark and a W boson in pp collisions at $\sqrt{s} = 8$ TeV”. In: *Phys. Rev. Lett.* 112.23 (2014), p. 231802. DOI: [10.1103/PhysRevLett.112.231802](https://doi.org/10.1103/PhysRevLett.112.231802). arXiv: [1401.2942](https://arxiv.org/abs/1401.2942) [hep-ex].
- [52] G. Aad et al. “Measurement of the production cross-section of a single top quark in association with a W boson at 8 TeV with the ATLAS experiment”. In: *JHEP* 01 (2016), p. 064. DOI: [10.1007/JHEP01\(2016\)064](https://doi.org/10.1007/JHEP01(2016)064). arXiv: [1510.03752](https://arxiv.org/abs/1510.03752) [hep-ex].
- [53] A. Sirunyan et al. “Measurement of the cross section for top quark pair production in association with a W or Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* 08 (2018), p. 011. DOI: [10.1007/JHEP08\(2018\)011](https://doi.org/10.1007/JHEP08(2018)011). arXiv: [1711.02547](https://arxiv.org/abs/1711.02547) [hep-ex].
- [54] A. Sirunyan et al. “Measurement of top quark pair production in association with a Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* 03 (2020), p. 056. DOI: [10.1007/JHEP03\(2020\)056](https://doi.org/10.1007/JHEP03(2020)056). arXiv: [1907.11270](https://arxiv.org/abs/1907.11270) [hep-ex].
- [55] *Measurement of the cross section of top quark-antiquark pair production in association with a W boson in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Tech. rep. Geneva: CERN, 2022. URL: <https://cds.cern.ch/record/2804479>.
- [56] A. Tumasyan et al. “Inclusive and differential cross section measurements of single top quark production in association with a Z boson in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* 02 (2022), p. 107. DOI: [10.1007/JHEP02\(2022\)107](https://doi.org/10.1007/JHEP02(2022)107). arXiv: [2111.02860](https://arxiv.org/abs/2111.02860) [hep-ex].
- [57] A. Habjia. “Synthesis on Heavy Higgs from the Standard Model to 2HDM and Beyond”. In: *JOURNAL OF ADVANCES IN PHYSICS* 18 (Oct. 2020), pp. 110–142. DOI: [10.24297/jap.v18i.8881](https://doi.org/10.24297/jap.v18i.8881).

- [58] Encyclopedia Britannica. *gamma ray*. URL: <https://www.britannica.com/summary/subatomic-particle#/media/1/570533/73050>.
- [59] K. R. Lang. *Invisible gamma ray photon*. URL: https://ase.tufts.edu/cosmos/view_picture.asp?id=1357.
- [60] Wikimedia Commons. *Cern accelerator complex*. 2011. URL: <https://commons.wikimedia.org/wiki/File:Cern-accelerator-complex.svg>.
- [61] I. Neutelings. *CMS coordinate system*. URL: https://tikz.net/axis3d_cms/.
- [62] T. Sakuma and T. McCauley. “Detector and Event Visualization with SketchUp at the CMS Experiment”. In: *J. Phys. Conf. Ser.* 513 (2014). Ed. by D. L. Groep and D. Bonacorsi, p. 022032. DOI: [10.1088/1742-6596/513/2/022032](https://doi.org/10.1088/1742-6596/513/2/022032). arXiv: [1311.4942](https://arxiv.org/abs/1311.4942) [physics.ins-det].
- [63] A. Benaglia. *The CMS ECAL performance with examples*. Tech. rep. Geneva: CERN, 2013. DOI: [10.1088/1748-0221/9/02/C02008](https://doi.org/10.1088/1748-0221/9/02/C02008). URL: <https://cds.cern.ch/record/1632384>.
- [64] CMS. *Particle-flow event reconstruction in CMS and performance for jets, taus, and missing transverse energy*. CMS Physics Analysis Summary CMS-PAS-PFT-09-001. CERN, 2009. URL: <http://cdsweb.cern.ch/record/1194487>.
- [65] P. Billoir and S. Qian. “Simultaneous pattern recognition and track fitting by the Kalman filtering method”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 294.1 (1990), pp. 219–228. ISSN: 0168-9002. DOI: [https://doi.org/10.1016/0168-9002\(90\)91835-Y](https://doi.org/10.1016/0168-9002(90)91835-Y). URL: <https://www.sciencedirect.com/science/article/pii/016890029091835Y>.
- [66] W. Adam et al. *Reconstruction of Electrons with the Gaussian-Sum Filter in the CMS Tracker at the LHC*. Tech. rep. Geneva: CERN, 2005. URL: <https://cds.cern.ch/record/815410>.
- [67] S. Chatrchyan et al. “Description and performance of track and primary-vertex reconstruction with the CMS tracker”. In: *JINST* 9.10 (2014), P10009. DOI: [10.1088/1748-0221/9/10/P10009](https://doi.org/10.1088/1748-0221/9/10/P10009). arXiv: [1405.6569](https://arxiv.org/abs/1405.6569) [physics.ins-det].
- [68] CMS Collaboration and T. Mc Cauley. “Collisions recorded by the CMS detector on 14 Oct 2016 during the high pile-up fill”. CMS Collection. 2016. URL: <https://cds.cern.ch/record/2231915>.
- [69] A. Sirunyan et al. “Electron and photon reconstruction and identification with the CMS experiment at the CERN LHC”. In: (Dec. 2020). arXiv: [2012.06888](https://arxiv.org/abs/2012.06888) [hep-ex].
- [70] V. Khachatryan et al. “Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV”. In: *JINST* 10.06 (2015), P06005. DOI: [10.1088/1748-0221/10/06/P06005](https://doi.org/10.1088/1748-0221/10/06/P06005). arXiv: [1502.02701](https://arxiv.org/abs/1502.02701) [physics.ins-det].
- [71] M. Cacciari, G. P. Salam, and G. Soyez. “The anti- k_t jet clustering algorithm”. In: *JHEP* 04 (2008), p. 063. DOI: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063). arXiv: [0802.1189](https://arxiv.org/abs/0802.1189) [hep-ex].
- [72] A. Sirunyan et al. “Performance of missing transverse momentum reconstruction in proton-proton collisions at $\sqrt{s} = 13$ TeV using the CMS detector”. In: *JINST* 14.07 (2019), P07004. DOI: [10.1088/1748-0221/14/07/P07004](https://doi.org/10.1088/1748-0221/14/07/P07004). arXiv: [1903.06078](https://arxiv.org/abs/1903.06078) [hep-ex].
- [73] V. Khachatryan et al. “The CMS trigger system”. In: *JINST* 12.01 (2017), P01020. DOI: [10.1088/1748-0221/12/01/P01020](https://doi.org/10.1088/1748-0221/12/01/P01020). arXiv: [1609.02366](https://arxiv.org/abs/1609.02366) [physics.ins-det].
- [74] J. C. Collins, D. E. Soper, and G. F. Sterman. “Factorization of Hard Processes in QCD”. In: *Adv. Ser. Direct. High Energy Phys.* 5 (1989), pp. 1–91. DOI: [10.1142/9789814503266_0001](https://doi.org/10.1142/9789814503266_0001). arXiv: [hep-ph/0409313](https://arxiv.org/abs/hep-ph/0409313).
- [75] N. Metropolis and S. Ulam. “The Monte Carlo Method”. In: *Journal of the American Statistical Association* 44.247 (1949), pp. 335–341. ISSN: 01621459. URL: <http://www.jstor.org/stable/2280232>.
- [76] S. Höche. “Introduction to parton-shower event generators”. In: *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders*. 2015, pp. 235–295. DOI: [10.1142/9789814678766_0005](https://doi.org/10.1142/9789814678766_0005). arXiv: [1411.4085](https://arxiv.org/abs/1411.4085) [hep-ph].
- [77] A. Canepa. “Searches for supersymmetry at the Large Hadron Collider”. In: *Reviews in Physics* 4 (2019), p. 100033. ISSN: 2405-4283. DOI: <https://doi.org/10.1016/j.revip.2019.100033>. URL: <https://www.sciencedirect.com/science/article/pii/S2405428318300091>.

- [78] A. Sirunyan et al. “Search for resonant and nonresonant new phenomena in high-mass dilepton final states at $\sqrt{s} = 13$ TeV”. In: *JHEP* 07 (2021), p. 208. DOI: [10.1007/JHEP07\(2021\)208](https://doi.org/10.1007/JHEP07(2021)208). arXiv: [2103.02708](https://arxiv.org/abs/2103.02708) [hep-ex].
- [79] D. J. Griffiths. *Introduction to elementary particles; 2nd rev. version*. Physics textbook. New York, NY: Wiley, 2008. URL: <https://cds.cern.ch/record/111880>.
- [80] J. A. Aguilar-Saavedra et al. *Interpreting top-quark LHC measurements in the standard-model effective field theory*. LHC TOP WG note CERN-LPCC-2018-01. 2018. arXiv: [1802.07237](https://arxiv.org/abs/1802.07237). URL: <http://cds.cern.ch/record/2305783>.
- [81] G. Aad et al. “Combination of the W boson polarization measurements in top quark decays using ATLAS and CMS data at $\sqrt{s} = 8$ TeV”. In: *JHEP* 08.08 (2020), p. 051. DOI: [10.1007/JHEP08\(2020\)051](https://doi.org/10.1007/JHEP08(2020)051). arXiv: [2005.03799](https://arxiv.org/abs/2005.03799) [hep-ex].
- [82] C. Degrande et al. “Effective Field Theory: A Modern Approach to Anomalous Couplings”. In: *Annals Phys.* 335 (2013), pp. 21–32. DOI: [10.1016/j.aop.2013.04.016](https://doi.org/10.1016/j.aop.2013.04.016). arXiv: [1205.4231](https://arxiv.org/abs/1205.4231) [hep-ph].
- [83] A. Bouzas and F. Larios. “Electromagnetic dipole moments of the Top quark”. In: *Phys. Rev. D* 87 (2013), p. 074015. DOI: [10.1103/PhysRevD.87.074015](https://doi.org/10.1103/PhysRevD.87.074015). arXiv: [1212.6575](https://arxiv.org/abs/1212.6575) [hep-ph].
- [84] T. Aaltonen et al. “Exclusion of exotic top-like quarks with $-4/3$ electric charge using jet-charge tagging in single-lepton $t\bar{t}$ events at CDF”. In: *Phys. Rev. D* 88.3 (2013), p. 032003. DOI: [10.1103/PhysRevD.88.032003](https://doi.org/10.1103/PhysRevD.88.032003). arXiv: [1304.4141](https://arxiv.org/abs/1304.4141) [hep-ex].
- [85] G. Aad et al. “Measurement of the top quark charge in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. In: *JHEP* 11 (2013), p. 031. DOI: [10.1007/JHEP11\(2013\)031](https://doi.org/10.1007/JHEP11(2013)031). arXiv: [1307.4568](https://arxiv.org/abs/1307.4568) [hep-ex].
- [86] V. M. Abazov et al. “Measurement of the Electric Charge of the Top Quark in $t\bar{t}$ Events”. In: *Phys. Rev. D* 90.5 (2014). [Erratum: *Phys. Rev. D* 90, 079904 (2014)], p. 051101. DOI: [10.1103/PhysRevD.90.051101](https://doi.org/10.1103/PhysRevD.90.051101). arXiv: [1407.4837](https://arxiv.org/abs/1407.4837) [hep-ex].
- [87] J. A. Aguilar-Saavedra, M. C. N. Fiolhais, and A. Onofre. “Top Effective Operators at the ILC”. In: *JHEP* 07 (2012), p. 180. DOI: [10.1007/JHEP07\(2012\)180](https://doi.org/10.1007/JHEP07(2012)180). arXiv: [1206.1033](https://arxiv.org/abs/1206.1033) [hep-ph].
- [88] A. Sirunyan et al. “Search for new physics in top quark production with additional leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV using effective field theory”. In: *JHEP* 03 (2021), p. 095. DOI: [10.1007/JHEP03\(2021\)095](https://doi.org/10.1007/JHEP03(2021)095). arXiv: [2012.04120](https://arxiv.org/abs/2012.04120) [hep-ex].
- [89] *The dim6top model, a complete top-quark EFT implementation*. <https://feynrules.irmp.ucl.ac.be/wiki/dim6top>.
- [90] C. Degrande et al. “Automated one-loop computations in the standard model effective field theory”. In: *Phys. Rev. D* 103.9 (2021), p. 096024. DOI: [10.1103/PhysRevD.103.096024](https://doi.org/10.1103/PhysRevD.103.096024). arXiv: [2008.11743](https://arxiv.org/abs/2008.11743) [hep-ph].
- [91] N. Tonon et al. “Probing effective field theory operators in the associated production of top quarks with a Z boson in multilepton final states at $\sqrt{s} = 13$ TeV”. In: *JHEP* 12 (2021), p. 083. DOI: [10.1007/JHEP12\(2021\)083](https://doi.org/10.1007/JHEP12(2021)083). arXiv: [2107.13896](https://arxiv.org/abs/2107.13896) [hep-ex].
- [92] V. Khachatryan et al. “Observation of top quark pairs produced in association with a vector boson in pp collisions at $\sqrt{s} = 8$ TeV”. In: *JHEP* 01 (2016), p. 096. DOI: [10.1007/JHEP01\(2016\)096](https://doi.org/10.1007/JHEP01(2016)096). arXiv: [1510.01131](https://arxiv.org/abs/1510.01131) [hep-ex].
- [93] A. Alloul, B. Fuks, and V. Sanz. “Phenomenology of the Higgs Effective Lagrangian via FEYN-RULES”. In: *JHEP* 04 (2014), p. 110. DOI: [10.1007/JHEP04\(2014\)110](https://doi.org/10.1007/JHEP04(2014)110). arXiv: [1310.5150](https://arxiv.org/abs/1310.5150) [hep-ph].
- [94] R. Contino et al. “Effective Lagrangian for a light Higgs-like scalar”. In: *JHEP* 07 (2013), p. 035. DOI: [10.1007/JHEP07\(2013\)035](https://doi.org/10.1007/JHEP07(2013)035). arXiv: [1303.3876](https://arxiv.org/abs/1303.3876) [hep-ph].
- [95] M. Aaboud et al. “Measurement of the $t\bar{t}Z$ and $t\bar{t}W$ cross sections in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 99 (2019), p. 072009. DOI: [10.1103/PhysRevD.99.072009](https://doi.org/10.1103/PhysRevD.99.072009). arXiv: [1901.03584](https://arxiv.org/abs/1901.03584) [hep-ex].
- [96] G. Aad et al. “Measurements of the inclusive and differential production cross sections of a top-quark antiquark pair in association with a Z boson at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Eur. Phys. J. C* 81.8 (2021), p. 737. DOI: [10.1140/epjc/s10052-021-09439-4](https://doi.org/10.1140/epjc/s10052-021-09439-4). arXiv: [2103.12603](https://arxiv.org/abs/2103.12603) [hep-ex].

- [97] G. Aad et al. “Measurement of the energy asymmetry in $t\bar{t}j$ production at 13 TeV with the ATLAS experiment and interpretation in the SMEFT framework”. In: (Oct. 2021). arXiv: [2110.05453 \[hep-ex\]](#).
- [98] E. Maguire, L. Heinrich, and G. Watt. “HEPData: a repository for high energy physics data”. In: *J. Phys. Conf. Ser.* 898.10 (2017). Ed. by Richard Mount and Craig Tull, p. 102006. DOI: [10.1088/1742-6596/898/10/102006](#). arXiv: [1704.05473 \[hep-ex\]](#).
- [99] J. Ellis et al. “Top, Higgs, Diboson and Electroweak Fit to the Standard Model Effective Field Theory”. In: *JHEP* 04 (2021), p. 279. DOI: [10.1007/JHEP04\(2021\)279](#). arXiv: [2012.02779 \[hep-ph\]](#).
- [100] *LHCTopWG Summary Plots: Top Quark Mass*. https://twiki.cern.ch/twiki/bin/view/LHCPhysics/LHCTopWGSummaryPlots#Top_Quark_Mass.
- [101] P.-F. Duan et al. “QCD corrections to associated production of $t\bar{t}\gamma$ at hadron colliders”. In: *Phys. Rev. D* 80 (2009), p. 014022. DOI: [10.1103/PhysRevD.80.014022](#). arXiv: [0907.1324 \[hep-ph\]](#).
- [102] K. Melnikov, M. Schulze, and A. Scharf. “QCD corrections to top quark pair production in association with a photon at hadron colliders”. In: *Phys. Rev. D* 83 (2011), p. 074013. DOI: [10.1103/PhysRevD.83.074013](#). arXiv: [1102.1967 \[hep-ph\]](#).
- [103] P.-F. Duan et al. “Next-to-leading order QCD corrections to $t\bar{t}\gamma$ production at the 7 TeV LHC”. In: *Chin. Phys. Lett.* 28 (2011), p. 111401. DOI: [10.1088/0256-307X/28/11/111401](#). arXiv: [1110.2315 \[hep-ph\]](#).
- [104] A. Kardos and Z. Trócsányi. “Hadroproduction of $t\bar{t}$ pair in association with an isolated photon at NLO accuracy matched with parton shower”. In: *JHEP* 05 (2015), p. 090. DOI: [10.1007/JHEP05\(2015\)090](#). arXiv: [1406.2324 \[hep-ph\]](#).
- [105] F. Maltoni, D. Pagani, and I. Tsirikos. “Associated production of a top-quark pair with vector bosons at NLO in QCD: impact on $t\bar{t}H$ searches at the LHC”. In: *JHEP* 02 (2016), p. 113. DOI: [10.1007/JHEP02\(2016\)113](#). arXiv: [1507.05640 \[hep-ph\]](#).
- [106] P.-F. Duan et al. “Electroweak corrections to top quark pair production in association with a hard photon at hadron colliders”. In: *Phys. Lett. B* 766 (2017), p. 102. DOI: [10.1016/j.physletb.2016.12.061](#). arXiv: [1612.00248 \[hep-ph\]](#).
- [107] D. Pagani et al. “Automated EW corrections with isolated photons: $t\bar{t}\gamma$, $t\bar{t}\gamma\gamma$ and $t\gamma j$ as case studies”. In: *JHEP* 09 (2021), p. 155. DOI: [10.1007/JHEP09\(2021\)155](#). arXiv: [2106.02059 \[hep-ph\]](#).
- [108] G. Bevilacqua et al. “Hard Photons in Hadroproduction of Top Quarks with Realistic Final States”. In: *JHEP* 10 (2018), p. 158. DOI: [10.1007/JHEP10\(2018\)158](#). arXiv: [1803.09916 \[hep-ph\]](#).
- [109] G. Bevilacqua et al. “Precise predictions for $t\bar{t}\gamma/t\bar{t}$ cross section ratios at the LHC”. In: *JHEP* 01 (2019), p. 188. DOI: [10.1007/JHEP01\(2019\)188](#). arXiv: [1809.08562 \[hep-ph\]](#).
- [110] G. Bevilacqua et al. “Off-shell vs on-shell modelling of top quarks in photon associated production”. In: *JHEP* 03 (2020), p. 154. DOI: [10.1007/JHEP03\(2020\)154](#). arXiv: [1912.09999 \[hep-ph\]](#).
- [111] U. Baur et al. “Probing electroweak top quark couplings at hadron colliders”. In: *Phys. Rev. D* 71 (2005), p. 054013. DOI: [10.1103/PhysRevD.71.054013](#). arXiv: [hep-ph/0412021](#).
- [112] M. Schulze and Y. Soreq. “Pinning down electroweak dipole operators of the top quark”. In: *Eur. Phys. J. C* 76 (2016), p. 466. DOI: [10.1140/epjc/s10052-016-4263-x](#). arXiv: [1603.08911 \[hep-ph\]](#).
- [113] O. Bessidskaia Bylund et al. “Probing top quark neutral couplings in the Standard Model Effective Field Theory at NLO in QCD”. In: *JHEP* 05 (2016), p. 052. DOI: [10.1007/JHEP05\(2016\)052](#). arXiv: [1601.08193 \[hep-ph\]](#).
- [114] J. A. Aguilar-Saavedra et al. “Shedding light on the $t\bar{t}$ asymmetry: the photon handle”. In: *JHEP* 04 (2014), p. 188. DOI: [10.1007/JHEP04\(2014\)188](#). arXiv: [1402.3598 \[hep-ph\]](#).
- [115] J. Bergner and M. Schulze. “The top quark charge asymmetry in $t\bar{t}\gamma$ production at the LHC”. In: *Eur. Phys. J. C* 79 (2019), p. 189. DOI: [10.1140/epjc/s10052-019-6707-6](#). arXiv: [1812.10535 \[hep-ph\]](#).

- [116] 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”. In: *JHEP* 07 (2014), p. 079. DOI: [10.1007/JHEP07\(2014\)079](https://doi.org/10.1007/JHEP07(2014)079). arXiv: [1405.0301](https://arxiv.org/abs/1405.0301) [hep-ph].
- [117] R. D. Ball et al. “Parton distributions from high-precision collider data”. In: *Eur. Phys. J. C* 77 (2017), p. 663. DOI: [10.1140/epjc/s10052-017-5199-5](https://doi.org/10.1140/epjc/s10052-017-5199-5). arXiv: [1706.00428](https://arxiv.org/abs/1706.00428) [hep-ph].
- [118] P. Nason. “A New method for combining NLO QCD with shower Monte Carlo algorithms”. In: *JHEP* 11 (2004), p. 040. DOI: [10.1088/1126-6708/2004/11/040](https://doi.org/10.1088/1126-6708/2004/11/040). arXiv: [hep-ph/0409146](https://arxiv.org/abs/hep-ph/0409146).
- [119] S. Frixione, G. Ridolfi, and P. Nason. “A Positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction”. In: *JHEP* 09 (2007), p. 126. DOI: [10.1088/1126-6708/2007/09/126](https://doi.org/10.1088/1126-6708/2007/09/126). arXiv: [0707.3088](https://arxiv.org/abs/0707.3088) [hep-ph].
- [120] S. Frixione, P. Nason, and C. Oleari. “Matching NLO QCD computations with Parton Shower simulations: the POWHEG method”. In: *JHEP* 11 (2007), p. 070. DOI: [10.1088/1126-6708/2007/11/070](https://doi.org/10.1088/1126-6708/2007/11/070). arXiv: [0709.2092](https://arxiv.org/abs/0709.2092) [hep-ph].
- [121] S. Alioli et al. “NLO single-top production matched with shower in POWHEG: s - and t -channel contributions”. In: *JHEP* 09 (2009), p. 111. DOI: [10.1088/1126-6708/2009/09/111](https://doi.org/10.1088/1126-6708/2009/09/111). arXiv: [0907.4076](https://arxiv.org/abs/0907.4076) [hep-ph].
- [122] S. Alioli et al. “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”. In: *JHEP* 06 (2010), p. 043. DOI: [10.1007/JHEP06\(2010\)043](https://doi.org/10.1007/JHEP06(2010)043). arXiv: [1002.2581](https://arxiv.org/abs/1002.2581) [hep-ph].
- [123] E. Re. “Single-top Wt -channel production matched with parton showers using the POWHEG method”. In: *Eur. Phys. J. C* 71 (2011), p. 1547. DOI: [10.1140/epjc/s10052-011-1547-z](https://doi.org/10.1140/epjc/s10052-011-1547-z). arXiv: [1009.2450](https://arxiv.org/abs/1009.2450) [hep-ph].
- [124] J. M. Campbell et al. “Top-Pair Production and Decay at NLO Matched with Parton Showers”. In: *JHEP* 04 (2015), p. 114. DOI: [10.1007/JHEP04\(2015\)114](https://doi.org/10.1007/JHEP04(2015)114). arXiv: [1412.1828](https://arxiv.org/abs/1412.1828) [hep-ph].
- [125] H. Hartanto et al. “Higgs boson production in association with top quarks in the POWHEG BOX”. In: *Phys. Rev. D* 91 (2015), p. 094003. DOI: [10.1103/PhysRevD.91.094003](https://doi.org/10.1103/PhysRevD.91.094003). arXiv: [1501.04498](https://arxiv.org/abs/1501.04498) [hep-ph].
- [126] J. M. Campbell and R. K. Ellis. “MCFM for the Tevatron and the LHC”. In: *Nucl. Phys. B Proc. Suppl.* 205-206 (2010), p. 10. DOI: [10.1016/j.nuclphysbps.2010.08.011](https://doi.org/10.1016/j.nuclphysbps.2010.08.011). arXiv: [1007.3492](https://arxiv.org/abs/1007.3492) [hep-ph].
- [127] J. M. Campbell, R. K. Ellis, and C. Williams. “Bounding the Higgs Width at the LHC Using Full Analytic Results for $gg \rightarrow e^-e^+\mu^-\mu^+$ ”. In: *JHEP* 04 (2014), p. 060. DOI: [10.1007/JHEP04\(2014\)060](https://doi.org/10.1007/JHEP04(2014)060). arXiv: [1311.3589](https://arxiv.org/abs/1311.3589) [hep-ph].
- [128] R. D. Ball et al. “Parton distributions for the LHC Run II”. In: *JHEP* 04 (2015), p. 040. DOI: [10.1007/JHEP04\(2015\)040](https://doi.org/10.1007/JHEP04(2015)040). arXiv: [1410.8849](https://arxiv.org/abs/1410.8849) [hep-ph].
- [129] T. Sjöstrand et al. “An introduction to PYTHIA8.2”. In: *Comput. Phys. Commun.* 191 (2015), p. 159. DOI: [10.1016/j.cpc.2015.01.024](https://doi.org/10.1016/j.cpc.2015.01.024). arXiv: [1410.3012](https://arxiv.org/abs/1410.3012) [hep-ph].
- [130] A. Sirunyan et al. “Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements”. In: *Eur. Phys. J. C* 80 (2020), p. 4. DOI: [10.1140/epjc/s10052-019-7499-4](https://doi.org/10.1140/epjc/s10052-019-7499-4). arXiv: [1903.12179](https://arxiv.org/abs/1903.12179) [hep-ex].
- [131] P. Skands, S. Carrazza, and J. Rojo. “Tuning PYTHIA8.1: the Monash 2013 Tune”. In: *Eur. Phys. J. C* 74 (2014), p. 3024. DOI: [10.1140/epjc/s10052-014-3024-y](https://doi.org/10.1140/epjc/s10052-014-3024-y). arXiv: [1404.5630](https://arxiv.org/abs/1404.5630) [hep-ph].
- [132] V. Khachatryan et al. “Event generator tunes obtained from underlying event and multiparton scattering measurements”. In: *Eur. Phys. J. C* 76 (2016), p. 155. DOI: [10.1140/epjc/s10052-016-3988-x](https://doi.org/10.1140/epjc/s10052-016-3988-x). arXiv: [1512.00815](https://arxiv.org/abs/1512.00815) [hep-ex].
- [133] CMS Collaboration. *Investigations of the impact of the parton shower tuning in PYTHIA8 in the modelling of $t\bar{t}$ at $\sqrt{s} = 8$ and 13 TeV*. CMS Physics Analysis Summary CMS-PAS-TOP-16-021. 2016. URL: <http://cds.cern.ch/record/2235192>.
- [134] GEANT4 Collaboration. “GEANT4 - a simulation toolkit. GEANT4. A Simulation toolkit”. In: *Nucl. Instrum. Methods Phys. Res., A* 506 (2002), 250–303. 54 p. DOI: [10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8). URL: <https://cds.cern.ch/record/602040>.

- [135] CMS Collaboration. *Object definitions for top quark analyses at the particle level*. CMS Note CMS-NOTE-2017-004. 2017. URL: <https://cds.cern.ch/record/2267573>.
- [136] A. Sirunyan et al. “Observation of Single Top Quark Production in Association with a Z Boson in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. Lett.* 122 (13 2019), p. 132003. DOI: [10.1103/PhysRevLett.122.132003](https://doi.org/10.1103/PhysRevLett.122.132003). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.122.132003>.
- [137] T. Cornelis et al. *Search for tZq in trilepton events*.
- [138] K. Skovpen et al. *TOP LeptonMVA*. 2020/06/03.
- [139] *ECAL scale and resolution corrections*. <https://twiki.cern.ch/twiki/bin/view/CMS/EgammaMiniAODV2>.
- [140] *Rochester Corrections for muons*. <https://twiki.cern.ch/twiki/bin/view/CMS/RochcorMuon>.
- [141] *Jet identification criteria for Run II*. <https://twiki.cern.ch/twiki/bin/view/CMS/JetID>.
- [142] A. Sirunyan et al. “Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV”. In: *JINST* 13.05 (2018), P05011. DOI: [10.1088/1748-0221/13/05/P05011](https://doi.org/10.1088/1748-0221/13/05/P05011). arXiv: [1712.07158](https://arxiv.org/abs/1712.07158) [physics.ins-det].
- [143] *Cut Based Photon ID for Run 2*. <https://twiki.cern.ch/twiki/bin/viewauth/CMS/CutBasedPhotonIdentificationRun2>.
- [144] A. Sirunyan et al. “Measurement of the top quark polarization and $t\bar{t}$ spin correlations using dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 100 (2019), p. 072002. DOI: [10.1103/PhysRevD.100.072002](https://doi.org/10.1103/PhysRevD.100.072002). arXiv: [1907.03729](https://arxiv.org/abs/1907.03729) [hep-ex].
- [145] M. Aaboud et al. “Measurements of top-quark pair spin correlations in the $e\mu$ channel at $\sqrt{s} = 13$ TeV using pp collisions in the ATLAS detector”. In: *Eur. Phys. J. C* 80.8 (2020), p. 754. DOI: [10.1140/epjc/s10052-020-8181-6](https://doi.org/10.1140/epjc/s10052-020-8181-6). arXiv: [1903.07570](https://arxiv.org/abs/1903.07570) [hep-ex].
- [146] T. Arndt, C. Diez Pardos, and A. Meyer. *Measurement of the 2016 Trigger Efficiencies for a dilepton selection for a $t\bar{t}$ analysis*. CMS Note 2016/392. CERN, 2017.
- [147] R. Schoefbeck et al. *Search for scalar top quark pair production in the dilepton final state at $\sqrt{s} = 13$ TeV in the 2016 dataset*. CMS Note 2017/024. CERN, 2017.
- [148] I. Khvastunov et al. *Measurement of the top pair-production in association with a W or Z boson in pp collisions at 13 TeV (2016 and 2017 dataset)*. CMS Note 2018/025. CERN, 2018.
- [149] G. Aad et al. “Measurements of $Z\gamma$ and $Z\gamma\gamma$ production in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 93.11 (2016), p. 112002. DOI: [10.1103/PhysRevD.93.112002](https://doi.org/10.1103/PhysRevD.93.112002). arXiv: [1604.05232](https://arxiv.org/abs/1604.05232) [hep-ex].
- [150] A. Sirunyan et al. “Precision luminosity measurement in proton-proton collisions at $\sqrt{s} = 13$ TeV in 2015 and 2016 at CMS”. Submitted to *Eur. Phys. J. C*. 2021.
- [151] *CMS luminosity measurement for the 2017 data-taking period at $\sqrt{s} = 13$ TeV*. Tech. rep. CMS-PAS-LUM-17-004. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2621960>.
- [152] *CMS luminosity measurement for the 2018 data-taking period at $\sqrt{s} = 13$ TeV*. Tech. rep. CMS-PAS-LUM-18-002. Geneva: CERN, 2019. URL: <https://cds.cern.ch/record/2676164>.
- [153] CMS. *Identification of b quark jets at the CMS Experiment in the LHC Run 2*. CMS Physics Analysis Summary CMS-PAS-BTV-15-001. CERN, 2016. URL: <http://cds.cern.ch/record/2138504>.
- [154] J. S. Conway. *Incorporating Nuisance Parameters in Likelihoods for Multisource Spectra*. 2011. arXiv: [1103.0354](https://arxiv.org/abs/1103.0354) [physics.data-an].
- [155] O. Hindrichs. *Measurement of differential cross sections for the production of top quark pairs in lepton+jets events from pp collisions at $\sqrt{s} = 13$ TeV*.
- [156] E. Ranken et al. *Measuring the top quark Yukawa coupling using kinematic distributions of $t\bar{t}$ dilepton decays from the full run 2 dataset*. CMS Note 2019/236. CERN, 2019. URL: <http://cms.cern.ch/iCMS/user/NOTEINFO?CMSNOTEID=CMS%20AN-2019/236>.
- [157] Matteo Cacciari and Gavin P. Salam. “Pileup subtraction using jet areas”. In: *Physics Letters B* 659.1-2 (2008), pp. 119–126. ISSN: 0370-2693. DOI: [10.1016/j.physletb.2007.09.077](https://doi.org/10.1016/j.physletb.2007.09.077). URL: <http://dx.doi.org/10.1016/j.physletb.2007.09.077>.

- [158] G. Aad et al. “Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV”. In: *JHEP* 08 (2016), p. 045. DOI: [10.1007/JHEP08\(2016\)045](https://doi.org/10.1007/JHEP08(2016)045). arXiv: [1606.02266](https://arxiv.org/abs/1606.02266) [hep-ex].
- [159] ATLAS Collaboration, CMS Collaboration, and LHC Higgs Combination Group. *Procedure for the LHC Higgs boson search combination in Summer 2011*. Tech. rep. CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-011. 2011. URL: <https://cds.cern.ch/record/1379837>.
- [160] G. Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *Eur. Phys. J. C* 71 (2011), p. 1554. DOI: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0). arXiv: [1007.1727](https://arxiv.org/abs/1007.1727) [physics.data-an].
- [161] R. Barlow and C. Beeston. “Fitting using finite Monte Carlo samples”. In: *Computer Physics Communications* 77.2 (1993), pp. 219–228. ISSN: 0010-4655. DOI: [https://doi.org/10.1016/0010-4655\(93\)90005-W](https://doi.org/10.1016/0010-4655(93)90005-W). URL: <http://www.sciencedirect.com/science/article/pii/001046559390005W>.
- [162] S. Schmitt. “TUnfold, an algorithm for correcting migration effects in high energy physics”. In: *Journal of Instrumentation* 7.10 (2012), T10003–T10003. DOI: [10.1088/1748-0221/7/10/t10003](https://doi.org/10.1088/1748-0221/7/10/t10003). URL: <https://doi.org/10.1088/1748-0221/7/10/t10003>.
- [163] G. Cowan. “A survey of unfolding methods for particle physics”. In: *Conf. Proc. C* 0203181 (2002). Ed. by M. R. Whalley and L. Lyons, pp. 248–257.
- [164] A. Sirunyan et al. “Measurement of $t\bar{t}$ normalised multi-differential cross sections in pp collisions at $\sqrt{s} = 13$ TeV, and simultaneous determination of the strong coupling strength, top quark pole mass, and parton distribution functions”. In: *Eur. Phys. J. C* 80.7 (2020), p. 658. DOI: [10.1140/epjc/s10052-020-7917-7](https://doi.org/10.1140/epjc/s10052-020-7917-7). arXiv: [1904.05237](https://arxiv.org/abs/1904.05237) [hep-ex].
- [165] A. Sirunyan et al. “Measurements of $t\bar{t}$ differential cross sections in proton-proton collisions at $\sqrt{s} = 13$ TeV using events containing two leptons”. In: *JHEP* 02 (2019), p. 149. DOI: [10.1007/JHEP02\(2019\)149](https://doi.org/10.1007/JHEP02(2019)149). arXiv: [1811.06625](https://arxiv.org/abs/1811.06625) [hep-ex].
- [166] N. Hartland et al. “A Monte Carlo global analysis of the Standard Model Effective Field Theory: the top quark sector”. In: *JHEP* 04 (2019), p. 100. DOI: [10.1007/JHEP04\(2019\)100](https://doi.org/10.1007/JHEP04(2019)100). arXiv: [1901.05965](https://arxiv.org/abs/1901.05965) [hep-ph].
- [167] I. Brivio et al. “O new physics, where art thou? A global search in the top sector”. In: *JHEP* 02 (2020), p. 131. DOI: [10.1007/JHEP02\(2020\)131](https://doi.org/10.1007/JHEP02(2020)131). arXiv: [1910.03606](https://arxiv.org/abs/1910.03606) [hep-ph].
- [168] S. Bickmann et al. “Top and Beauty synergies in SMEFT-fits at present and future colliders”. In: *JHEP* 06 (2021), p. 010. DOI: [10.1007/JHEP06\(2021\)010](https://doi.org/10.1007/JHEP06(2021)010). arXiv: [2012.10456](https://arxiv.org/abs/2012.10456) [hep-ph].
- [169] J. Ethier et al. “Combined SMEFT interpretation of Higgs, diboson, and top quark data from the LHC”. In: *JHEP* 11 (2021), p. 089. DOI: [10.1007/JHEP11\(2021\)089](https://doi.org/10.1007/JHEP11(2021)089). arXiv: [2105.00006](https://arxiv.org/abs/2105.00006) [hep-ph].
- [170] P. Torrielli. “Rare Standard Model processes for present and future hadronic colliders”. In: (July 2014). arXiv: [1407.1623](https://arxiv.org/abs/1407.1623) [hep-ph].
- [171] P. Janot. “Top-quark electroweak couplings at the FCC-ee”. In: *JHEP* 04 (2015), p. 182. DOI: [10.1007/JHEP04\(2015\)182](https://doi.org/10.1007/JHEP04(2015)182). arXiv: [1503.01325](https://arxiv.org/abs/1503.01325) [hep-ph].
- [172] M. Boronat et al. “Top quark mass measurement in radiative events at electron-positron colliders”. In: *Phys. Lett. B* 804 (2020), p. 135353. DOI: [10.1016/j.physletb.2020.135353](https://doi.org/10.1016/j.physletb.2020.135353). arXiv: [1912.01275](https://arxiv.org/abs/1912.01275) [hep-ph].
- [173] A. Sirunyan et al. “Search for physics beyond the standard model in events with two leptons of same sign, missing transverse momentum, and jets in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 77 (2017), p. 578. DOI: [10.1140/epjc/s10052-017-5079-z](https://doi.org/10.1140/epjc/s10052-017-5079-z). arXiv: [1704.07323](https://arxiv.org/abs/1704.07323) [hep-ex].
- [174] A. Sirunyan et al. “Search for standard model production of four top quarks with same-sign and multilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 78 (2018), p. 140. DOI: [10.1140/epjc/s10052-018-5607-5](https://doi.org/10.1140/epjc/s10052-018-5607-5). arXiv: [1710.10614](https://arxiv.org/abs/1710.10614) [hep-ex].
- [175] M. Aaboud et al. “Search for new phenomena in events with same-charge leptons and b jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *JHEP* 12 (2018), p. 039. DOI: [10.1007/JHEP12\(2018\)039](https://doi.org/10.1007/JHEP12(2018)039). arXiv: [1807.11883](https://arxiv.org/abs/1807.11883) [hep-ex].

- [176] M. Aaboud et al. “Search for four-top-quark production in the single-lepton and opposite-sign dilepton final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 99 (2019), p. 052009. DOI: [10.1103/PhysRevD.99.052009](https://doi.org/10.1103/PhysRevD.99.052009). arXiv: [1811.02305](https://arxiv.org/abs/1811.02305) [[hep-ex](#)].
- [177] A. Sirunyan et al. “Search for the production of four top quarks in the single-lepton and opposite-sign dilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV ”. In: *JHEP* 11 (2019), p. 082. DOI: [10.1007/JHEP11\(2019\)082](https://doi.org/10.1007/JHEP11(2019)082). arXiv: [1906.02805](https://arxiv.org/abs/1906.02805) [[hep-ex](#)].
- [178] A. Sirunyan et al. “Search for production of four top quarks in final states with same-sign or multiple leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV ”. In: *Eur. Phys. J. C* 80 (2020), p. 75. DOI: [10.1140/epjc/s10052-019-7593-7](https://doi.org/10.1140/epjc/s10052-019-7593-7). arXiv: [1908.06463](https://arxiv.org/abs/1908.06463) [[hep-ex](#)].
- [179] G. Aad et al. “Evidence for $t\bar{t}t\bar{t}$ production in the multilepton final state in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Eur. Phys. J. C* 80 (2020), p. 1085. DOI: [10.1140/epjc/s10052-020-08509-3](https://doi.org/10.1140/epjc/s10052-020-08509-3). arXiv: [2007.14858](https://arxiv.org/abs/2007.14858) [[hep-ex](#)].
- [180] G. Aad et al. “Measurement of the $t\bar{t}t\bar{t}$ production cross section in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *JHEP* 11 (2021), p. 118. DOI: [10.1007/JHEP11\(2021\)118](https://doi.org/10.1007/JHEP11(2021)118). arXiv: [2106.11683](https://arxiv.org/abs/2106.11683) [[hep-ex](#)].
- [181] T. T. et al Tran. *Measurements of the $pp \rightarrow p$ process at $\sqrt{s} = 13$ TeV using the full Run-2 data sample*.