

Search for displaced leptons in proton-proton collisions at
 $\sqrt{s} = 13 \text{ TeV}$

Dissertation

Presented in Partial Fulfillment of the Requirements for the Degree
Doctor of Philosophy in the Graduate School of The Ohio State
University

By

Bryan Cardwell, B.S., M.S.

Graduate Program in Physics

The Ohio State University

2021

Dissertation Committee:

Christopher Hill, Advisor

Stanley Durkin

Linda Carpenter

Andrew Heckler



© Copyright by

Bryan Cardwell

2021

Abstract

A search is presented for new long-lived particles that propagate a measurable distance through the CMS detector before decaying to leptons. The search is performed in $113\text{--}118\text{fb}^{-1}$ of proton-proton collision data produced by the CERN LHC at a center-of-mass energy of 13 TeV and collected by the CMS detector in 2016, 2017, and 2018. Events are selected with two leptons (an electron and a muon, two electrons, or two muons) that both have transverse impact parameter values between 0.01 and 10 cm. Using transverse impact parameter as the discriminating variable allows for sensitivity to displaced decays without requiring that the leptons form a common vertex. The search is designed to be sensitive to a wide range of new physics models that produce displaced di-lepton final states. The observation is consistent with the background-only hypothesis, and limits are set on the product of the cross-section of top squark pair production and the branching fraction to a lepton and a b or d quark through an R-parity-violating vertex. For a proper decay length hypothesis of 2 cm, top squarks with masses up to 1500 GeV are excluded at the 95 % confidence level.

For Cristiana and Montague.

Acknowledgments

Chris, thank you for this opportunity. Your thoughts are the contours of this thesis, and your foresight has continually put me in the right place at the right time.

Jamie, thank you for answering every naive question and for the kindness and care with which you spoon-fed me achievable tasks as I inched towards consciousness. Andrew, thank you for always being one step ahead of me and so graciously showing me the way. Brian, thank you for all the times you happily entertained whatever ultra-specific issue I happened to be facing.

Juliette, thank you for everything, most of all the teamwork, the guidance, and being so unrelentingly reliable. It's strange to think that our day-to-day professional lives will diverge soon; I look forward to when they cross again.

Everyone in MTest, thank you for the immediate and effortless companionship and teamwork. Corrinne, thank you for the advice, perspective, and care. Susan, thank you for constantly working side by side. Andre, thank you for saving me when I dropped that tiny screw. Thank you all for all the sleepless nights and pushing so hard towards our common goal. Lorenzo, thank you for the all-day-every-day support, the laughter, and the constant willingness to explain.

PRB 3000, the first two years of grad school would have broken me if not for our shenanigans. Thank you for being my family, dragging me through core classes, and

helping me feel ok as an imposter. Jovan, your coffee fueled each day and knowing you were on the other side of the Walhalla helped make Clintonville feel like home.

Paul Bloom and David Horner, thank you for showing me physics, for putting up with me every time I insisted on making things harder than they had to be, and for giving me a second chance. Stuart Wick, thank you for introducing me to Minutemen and infinite-dimensional function spaces. Bob Stanek, thank you for showing me experimental particle physics and for always having my back.

Dad, Julie, Bob, Ann, Jerry, Patrizia, Mark, Jeanne, and all my family, thank you for the unconditional support and love I've felt all my life. Thank you for looking at the world with me and always wanting to talk about cool stuff. I am in a position to write this thesis because of your examples, efforts, and care.

Cristiana and Montague, you are on every page. Thank you for sharing this adventure with me.

Vita

2015	B.S. Physics, North Central College
2015-2017	Graduate Teaching Associate, The Ohio State University
2017	M.S. Physics, The Ohio State University
2017-present	Graduate Research Associate, The Ohio State University

Publications

Research Publications

Included in the author list of 78 publications from CMS.

Fields of Study

Major Field: Physics

Table of Contents

	Page
Abstract	ii
Dedication	iii
Acknowledgments	iv
Vita	vi
List of Tables	x
List of Figures	xiv
1. Introduction	1
1.1 The Standard Model	2
1.1.1 Quantum chromodynamics	4
1.1.2 The electroweak theory	4
1.1.3 The Higgs mechanism	5
1.1.4 Current status	7
1.2 Beyond the Standard Model	11
1.2.1 Supersymmetry	11
1.2.2 Long-lived particles	13
1.2.3 Displaced supersymmetry	16
2. The Large Hadron Collider and Compact Muon Solenoid experiment . .	19
2.1 The Large Hadron Collider	19
2.1.1 Injection chain	20
2.1.2 Main ring	21
2.2 The Compact Muon Solenoid experiment	27
2.2.1 Solenoid magnet	29

2.2.2	Tracker	29
2.2.3	Electromagnetic calorimeter	35
2.2.4	Hadronic calorimeter	37
2.2.5	Muon system	39
2.2.6	Trigger	41
2.2.7	Physics object reconstruction	43
3.	Search for displaced leptons	48
3.1	Overview	50
3.1.1	Long-lived particles at the LHC	50
3.1.2	Displaced leptons signature	52
3.1.3	Analysis strategy	55
3.2	Data and simulated samples	58
3.2.1	Experimental data	58
3.2.2	Simulated background events	58
3.2.3	Simulated signal events	59
3.3	Event selection	62
3.3.1	Triggers	62
3.3.2	Preselection	64
3.3.3	Prompt control region	74
3.3.4	Inclusive signal region	77
3.4	Corrections to simulation	80
3.4.1	Pileup	80
3.4.2	Lepton ID	80
3.4.3	Lepton d_0 resolution	80
3.4.4	Trigger efficiency	85
3.5	Background estimation	89
3.5.1	Background sources	89
3.5.2	Data-driven ABCD method	91
3.5.3	Closure tests in control regions	95
3.5.4	ABCD correction and systematic uncertainty	97
3.5.5	Testing full background estimation procedure	103
3.5.6	Additional background checks	104
3.6	Systematic uncertainties	112
3.6.1	Integrated luminosity	112
3.6.2	Pileup	112
3.6.3	Displaced tracking efficiency	114
3.6.4	Trigger efficiency	114
3.6.5	Lepton ID and isolation	114
3.6.6	Muon pixel hit efficiency	115
3.6.7	Lepton d_0 resolution	116

3.7	Results	117
3.7.1	Observed events	117
3.7.2	Limits	122
3.7.3	Additional likelihood tests	124
3.8	Outlook	130
4.	Conclusion	132
Appendices		134
A.	Impact of APV25 saturation on displaced tracking	134
B.	Poorly measured lepton $ d_0 $ at large $ \eta $	138
C.	Displaced tracking efficiency	140
D.	Study of R-hadron treatment in simulation	148
Bibliography		155

List of Tables

Table	Page
2.1 Luminosity parameters used in Eq. (2.1).	22
3.1 The top squark production cross sections used when generating the simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}\bar{b}$ and $\tilde{t}\tilde{t} \rightarrow \bar{l}d\bar{l}\bar{d}$ events. Cross sections are calculated at NLO + NLL by the LHC SUSY Cross Section Working Group. . .	61
3.2 The $e\mu$ preselection criteria. The electron and muon p_T thresholds increase in 2017 in accordance with the increased HLT electron and muon p_T thresholds.	65
3.3 The ee preselection criteria. The electron p_T threshold increases in 2017 and 2018 in accordance with the increased HLT electron p_T threshold.	66
3.4 The $\mu\mu$ preselection criteria. The muon p_T threshold increases in 2017 and 2018 in accordance with the increased HLT muon p_T threshold.	67
3.5 The electron tight ID requirements, which are identical to the tight cut-based ID from the CMS EGamma Physics Object Group with the d_0 and d_z requirements removed. Electron ID quantity definitions are available in Ref. [73].	68
3.6 The muon tight ID requirements, which are identical to the tight cut-based ID from the CMS Muon Physics Object Group with the requirements on d_0 and d_z removed.	69
3.7 The cumulative efficiency for simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}\bar{b}$ signal events to pass the 2018 inclusive signal region selection, for several choices of \tilde{t} mass and $c\tau$. The corrections described in Section 3.4 are applied.	79

3.8	The average σ_{align} for electrons and muons, for the 2017 and 2018 analyses.	83
3.9	The unscaled p_T^{miss} triggers used to create an orthogonal data sample for the trigger efficiency calculation.	86
3.10	The trigger efficiency scale factors in each channel and year.	87
3.11	The p_T boundaries between the low- and high- p_T bins of SR I in each channel.	94
3.12	Closure test results in data and background simulation in the 100–500 μm region. The average extrapolated ratios and their statistical uncertainties are given. For each test, the A, B, C, and D regions are defined as follows: A is 20–30 μm in prompt lepton $ d_0 $ and 20–100 μm in displaced lepton $ d_0 $, B is 20–30 μm in prompt lepton $ d_0 $ and 100–500 μm in displaced lepton $ d_0 $, C is always 20–100 μm in displaced lepton $ d_0 $, D (the test region) is always 100–500 μm in displaced lepton $ d_0 $, and we perform repeated tests while simultaneously varying the C and D prompt lepton $ d_0 $ within the 30–100 μm range.	96
3.13	Closure test results in data and background simulation in the 500 μm –10 cm region. The ratios of the actual to the estimated yield and their statistical uncertainties are given. The A, B, C, and D regions are defined as follows: A is 20–30 μm in prompt lepton $ d_0 $ and 20–100 μm in displaced lepton $ d_0 $, B is 20–30 μm in prompt lepton $ d_0 $ and 500 μm –10 cm in displaced lepton $ d_0 $, C is 30–100 μm in prompt lepton $ d_0 $ and 20–100 μm in displaced lepton $ d_0 $, and D (the test region) is 30–100 μm in prompt lepton $ d_0 $ and 500 μm –10 cm in displaced lepton $ d_0 $	98
3.14	The correction factors and the uncorrected and corrected background estimates in SR I. The correction factor uncertainties include both the uncertainty in the average and the additional uncertainty obtained from varying the fit extrapolation point. The total uncertainty (statistical plus systematic) is given for the corrected background estimates.	103
3.15	The systematic uncertainty and the background estimates in SRs II, III, and IV. The total uncertainty (statistical plus systematic) is given for each estimate.	104

3.16	Closure test results in background simulation in the SRs with all background estimate corrections and uncertainties applied. The estimated number of events, the actual number of events, and their total uncertainties (statistical plus systematic) are given. In cases where the actual number of events is zero, the uncertainty is given by the product of the average background simulation event weight and the upper bound of the 68% Poisson interval given by a single observation of zero events.	105
3.17	Some properties of the seven events found in data with the material interactions selection inverted.	106
3.18	Background estimates in data while applying the 2018 $\mu\mu$ preselection and the additional requirement of at least one b -tagged jet. The estimates with at least one b -tagged jet are about an order of magnitude below the nominal prediction.	108
3.19	A closure test of the ABCD method in 2018 QCD simulation in the $\mu\mu$ channel with the muon isolation criterion inverted. The estimates from the ABCD method, the actual yields in simulation, and the ratios of the actual to the estimated yields are shown.	109
3.20	Systematic uncertainties in the signal efficiency for all years and channels. The mean is provided in cases where the uncertainty varies by signal sample. Uncertainties in the same row are treated as correlated among the years of data taking, except for the displaced tracking and muon pixel hit efficiencies, where the 2016 uncertainty is treated as uncorrelated with the 2017 and 2018 uncertainties.	113
3.21	The number of estimated background and observed events in each channel and SR. For each estimate, the total uncertainty is given. The p_T boundaries that separate the low- and high- p_T bins of SR I are listed in Table 3.11.	118
3.22	The pre- and post-fit predictions for each signal region bin.	127
C.1	Cosmic-ray muon arrival time requirements used when measuring displaced tracking efficiency in data and simulation.	143

C.2	Mean measured efficiency to reconstruct both lepton tracks in simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}$ events in data and simulation and the resulting systematic uncertainty. The top squark mass and proper decay length are assumed to be 1800 GeV and 100 cm.	147
-----	---	-----

List of Figures

Figure	Page
1.1 The Standard Model particle content.	3
1.2 Summary of Standard Model production cross section measurements from the CMS experiment [12].	8
1.3 The Higgs boson mass corrected by a fermion loop (left) and a fermion mass corrected by photon loop (right).	10
1.4 Corrections to the Higgs boson mass from the top quark (left) and top squark (right) cancel in exact SUSY. The top quark and top squark contributions are enhanced by the large coupling between the Higgs boson and the top quark.	12
1.5 Mass limits at 95 % CL for a simplified model of gluino pair production with gluino decays to pairs of top quarks and the LSP (left) and top squark pair production with squark decays to a top quark and the LSP (right) from several CMS analyses [19].	13
1.6 Masses and proper decay lengths of many Standard Model particles. Particles with proper decay lengths above approximately 10^{-4} m will be noticeably long-lived in collider detectors such as CMS [20].	14
1.7 Long-lived decays of the neutron (left) and muon (right).	15
1.8 Pair-produced top squarks that each decay to a bottom quark and a lepton through an R-parity-violating LQD vertex.	18
2.1 Layout of the LHC experiments [27].	20

2.2	A diagram of the CERN accelerator complex. The analysis presented in Chapter 3 utilizes protons accelerated by LINAC 2, BOOSTER (also known as PSB), PS, SPS, and finally the LHC before their ultimate collision inside CMS [32].	21
2.3	The peak instantaneous (top) and integrated (bottom) luminosity delivered by the LHC during proton operation between 2011 and 2018 [33].	24
2.4	Diagram of an LHC dipole magnet in cross-section [34].	25
2.5	The CMS detector [37].	28
2.6	The stored-energy-over-mass ratio, E/M , for several detector magnets (left), and a cross-sectional view of the four-layer winding of reinforced conductor in the CMS superconducting solenoid (right) [39].	30
2.7	Measured single hit efficiency per layer as a function of the instantaneous luminosity (left) and pileup interactions (right) in data taken with the original CMS pixel detector in 2016 [41].	33
2.8	Comparison of the original and Phase-1 CMS pixel detector layouts in y - z plane [40].	34
2.9	Layout of the CMS silicon tracker. TIB, TOB, TID, and TEC refer to subdetectors of the strip detector while PIXEL refers to the original pixel detector. The Phase-1 pixel detector is contained within the same volume [39].	35
2.10	The observed local p-value for the five decay modes and the overall combination as a function of the SM Higgs boson mass showing the importance of the ECAL mass resolution in the discovery of the Higgs boson by the CMS experiment. The dashed line shows the expected local p-values for a SM Higgs boson with a mass m_H [8].	36
2.11	Layout of the hadron calorimeter barrel (HB), outer (HO), endcap (HE), and forward (HF) subdetectors [39].	38
2.12	Muon momentum resolution as a function of momentum when using the CMS muon system, the CMS inner tracker, and the combination of the two subdetectors in two different η ranges [38].	40

2.13	Sketch of a CMS muon system drift cell showing drift lines and isochrones. The plates at the top and bottom of the cell are at ground potential while the voltages applied to the electrodes are +3600 V for wires, +1800 V for strips, and −1200 V for cathodes [39].	41
2.14	A sketch of a transverse slice of the CMS detector showing representative particle interactions used to identify and reconstruct particles with the CMS Particle Flow algorithm [48].	47
3.1	Illustration of several possible experimental signatures of long-lived particles [53].	51
3.2	Illustration of the displaced leptons signature showing the definition of d_0 in a transverse view of the CMS detector. X denotes a new long-lived particle, ℓ denotes an electron or muon, and Y denotes any other decay products of the new long-lived particle. When interpreting the results of the Displaced Leptons analysis with the Displaced SUSY model, X refers to a top squark and Y refers to a b or d quark. . . .	53
3.3	Distribution of data (colors) and simulated Displaced SUSY events (black boxes) in the plane defined by electron $ d_0 $ and muon $ d_0 $. The size of the black boxes are proportional to the bin content, and the bins along the x and y axes include underflow. All events are required to pass the $e\mu$ preselection defined in Section 3.3.	56
3.4	Comparison of the the muon isolation pileup correction term in the standard muon isolation and the modified muon isolation in simulated $t\bar{t}$ events that pass the 2018 $e\mu$ preselection. Muon $ d_0 $ is constrained to be less than 100 μm in the plot on the left and between 500 and 1000 μm in the plot on the right.	70
3.5	The distribution of $t\bar{t}$ simulated events in the plane defined by electron $ d_0 $ and muon $ d_0 $. The standard isolation is applied in the plot on the left, and the modified isolation is applied in the plot on the right, and the events in both plots are required to pass the remaining 2018 $e\mu$ preselection criteria and the additional constraint that the parent of at least one lepton is a heavy-flavor meson.	71
3.6	The muon modified isolation distribution for simulated $t\bar{t}$ background and $t\bar{t} \rightarrow \bar{l}b\bar{l}b$ signal events that pass the 2018 $e\mu$ preselection with no isolation criterion applied.	71

3.7	The electron η - ϕ distribution in a prompt-muon, displaced-electron control region in 2017 (left) and 2018 (right) data before vetoing the regions affected by pixel power-supply issues.	72
3.8	The di-muon $\cos(\alpha)$ distribution of 2018 NoBPTX data and simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events (left) and the di-muon Δt distribution in simulated cosmic-ray muon events in 2016 conditions (center) and $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ (right) events in 2018 conditions. In each case, all $\mu\mu$ preselection criteria except those relating to cosmic rejection are applied.	73
3.9	The electron (left) and muon (right) $ d_0 $ distributions for 2018 simulation events that pass the $e\mu$ preselection criteria. The upper two plots show $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ simulation for a single \tilde{t} mass and four different proper decay lengths; each histogram is normalized to unity. The lower two plots show the background simulation normalized to the integrated luminosity. In all of the histograms, the last bin includes the overflow. All of the corrections from Section 3.4 are applied.	75
3.10	The cumulative number of simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events that pass each criterion in the 2018 $e\mu$ (top), ee (middle), and $\mu\mu$ (bottom) preselection. Several \tilde{t} proper decay lengths are shown. The jet criteria do not exclude any events and are simply an artifact of the analysis framework.	76
3.11	The electron (top) and muon (bottom) p_T (left), η (center), and $ d_0 $ (right) distributions in the $e\mu$ prompt control region for 2016 data and simulated background events. The rightmost bin in each plot contains the overflow entries.	77
3.12	The electron p_T (left), η (center), and $ d_0 $ (right) distributions in the ee prompt control region for 2016 data and simulated background events. The rightmost bin in each plot contains the overflow entries.	78
3.13	The muon p_T (left), η (center), and $ d_0 $ (right) distributions in the $\mu\mu$ prompt control region for 2016 data and simulated background events. The rightmost bin in each plot contains the overflow entries.	78
3.14	The uncorrected lepton $ d_0 $ distributions in the $e\mu$ prompt control region, for electrons (left) and muons (right), for 2017 data and simulation (top), and 2018 data and simulation (bottom). The rightmost bin in each plot contains the overflow entries.	81

3.15	The average lepton $ d_0 $ as a function of ϕ in the $e\mu$ prompt control region, for electrons (left) and muons (right), for 2017 data and simulation.	82
3.16	The lepton d_0 distributions with Gaussian fits in the 2017 $e\mu$ prompt control region for (working from left to right) electrons in data, electrons in background simulation, muons in data, and muons in background simulation.	83
3.17	The lepton d_0 distributions with Gaussian fits in the 2018 $e\mu$ prompt control region for (working from left to right) electrons in data, electrons in background simulation, muons in data, and muons in background simulation.	83
3.18	The corrected lepton $ d_0 $ distributions in the $e\mu$ prompt control region, for electrons (left) and muons (right), for 2017 data and simulation (top), and 2018 data and simulation (bottom). The rightmost bin in each plot contains the overflow entries.	84
3.19	Trigger efficiency as a function of leading electron (left) or leading muon (right) p_T in the $e\mu$ channel in 2016 (top), 2017 (middle), and 2018 (bottom) in data and simulated background $t\bar{t}$ events.	88
3.20	The fraction of electrons (left) and muons (right) from different parents as a function of lepton $ d_0 $ in simulated $t\bar{t}$ events that pass the 2018 $e\mu$ channel preselection. Note that the vast majority of leptons whose parent is a W boson are produced in prompt decays.	90
3.21	The fraction of muons from different parents as a function of muon $ d_0 $ in simulated Drell-Yan events that pass the 2018 $\mu\mu$ channel preselection and the additional constraint that the leading and subleading muon both come from the same type of parent particle.	92
3.22	A diagram of the ABCD method overlaid on simulated background events passing the 2018 $e\mu$ preselection. A, B, and C are control regions, and D corresponds to the inclusive SR, which includes SRs I, II, III, and IV. Underflow events are included in the bins along the left and bottom edges. When performing the background estimate, regions B and C are further subdivided to coincide with the SR for which the estimate is being performed.	93

3.23	Background estimation closure tests in data, in 100–500 μm subregions of regions B (left) and C (right) in the $e\mu$ channel. Each plot shows the ratio of the actual to the estimated number of events as a function of the prompt lepton $ d_0 $ in 2016 (top) and 2017–18 (bottom). A linear fit is shown in black along with the 68 % and 95 % confidence intervals of the extrapolated fit value.	100
3.24	Background estimation closure tests in data, in 100–500 μm subregions of regions B (left) and C (right) in the ee channel. Each plot shows the ratio of the actual to the estimated number of events as a function of the prompt lepton $ d_0 $ in 2016 (top) and 2017–18 (bottom). A linear fit is shown in black along with the 68 % and 95 % confidence intervals of the extrapolated fit value.	101
3.25	Background estimation closure tests in data, in 100–500 μm subregions of regions B (left) and C (right) in the $\mu\mu$ channel. Each plot shows the ratio of the actual to the estimated number of events as a function of the prompt lepton $ d_0 $ in 2016 (top) and 2017–18 (bottom). A linear fit is shown in black along with the 68 % and 95 % confidence intervals of the extrapolated fit value.	102
3.26	The di-muon invariant mass distribution in data and QCD multijet simulation events that pass the 2018 $\mu\mu$ preselection with the muon isolation criterion inverted.	109
3.27	The di-muon invariant mass distribution in data events that pass the 2018 $\mu\mu$ prompt control region (black), SR I (blue), and SR IV (red) criteria with the muon isolation and di-muon ΔR criteria inverted. The equivalent distribution in the prompt control region is also shown in green.	110
3.28	The pixel hit efficiency as a function of muon $ d_0 $ in simulated cosmic ray events and NoBPTX data in 2016 (left), 2017 (center), and 2018 (right) conditions.	116
3.29	The number of estimated background and observed events in each channel and SR, with a representative signal overlaid, for 2016 (top) and 2017–2018 (bottom). For each background estimate and signal yield, the total uncertainty is given.	119

3.30	The number of estimated background and observed events in each channel and SR, with a representative signal overlaid, for 2016–2018. For each background estimate and signal yield, the total uncertainty is given. Statistical uncertainties are given on observed yields.	120
3.31	Two-dimensional distributions of $ d_0^a $ and $ d_0^b $, for the events in data that pass the $e\mu$ (left), ee (center), and $\mu\mu$ (right) preselection. The bins along the x and y axes contain underflow. The inclusive signal region covers the region between 100 μm and 10 cm in each $ d_0 $ variable shown.	121
3.32	Two-dimensional distributions of $ d_0^a $ and $ d_0^b $, for data events in the inclusive SR in the $e\mu$ (left), ee (center), and $\mu\mu$ (right) channels. . .	121
3.33	Muon ϕ and di-muon invariant mass distributions for data and composite background simulation in the $\mu\mu$ channel in the inclusive SR. The background simulation is normalized to the ABCD method background prediction.	123
3.34	The 95% CL upper limits on the top squark mass ($m_{\tilde{t}}$) as a function of its lifetime ($c\tau$), for the $e\mu$, ee , and $\mu\mu$ channels. The $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ (top) and $\tilde{t}\tilde{t} \rightarrow \bar{l}d\bar{l}d$ (bottom) processes are shown.	125
3.35	The 95% CL upper limits on the top squark mass ($m_{\tilde{t}}$) as a function of its lifetime ($c\tau$). The colors indicate the observed 95% CL upper limit on the cross section. The $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ (left) and $\tilde{t}\tilde{t} \rightarrow \bar{l}d\bar{l}d$ (right) processes are shown.	126
3.36	The distribution of pulls for each signal region bin, where pulls are calculated as the difference between the post-fit background yield and the pre-fit background yield divided by the pre-fit background uncertainty (left). The distribution of pulls for each background nuisance parameter, where pulls are calculated as the difference between the post-fit value and the pre-fit value divided by the pre-fit uncertainty (right). .	128
3.37	The observed asymptotic significances for the $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ process as a function of \tilde{t} mass and lifetime using the combined results.	129
3.38	The observed asymptotic significances for the $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ process as a function of \tilde{t} mass and lifetime in the ee (left), $e\mu$ (center), $\mu\mu$ channels (right).	129

A.1	Transverse decay length distribution of reconstructed $K_S^0 \rightarrow \pi^+\pi^-$ candidates for various runs in 2016 (top left), 2017 (top right) and 2018 (bottom). The peak instantaneous luminosity of each run is indicated in the legend, and each distribution is normalized to the integrated luminosity of the run. In the 2016 plot, runs taken before (after) the APV25 saturation effect was mitigated are shown by solid (dashed) lines. In the 2017 plot, the run with lowest (highest) instantaneous luminosity examined in each data-taking era is plotted with a dashed (solid) line. Towards the end of 2017 (orange and black lines) the LHC ran with approximately 20% fewer proton bunches, meaning that the instantaneous luminosity per bunch crossing was higher than one might naively expect from the instantaneous luminosities shown in the legend.	137
B.1	The standard deviation of the leading muon d_0 as a function of the leading muon η for simulated background events (left). To ensure that the variation in width is purely due to d_0 resolution effects, we use a sample of simulated Drell-Yan events from which the $Z \rightarrow \tau\tau \rightarrow ll$ events have been removed. Muon η distribution for simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}$ events (right). The $\mu\mu$ preselection with a loosened $ \eta $ requirement is applied in both plots.	139
B.2	Electron η distribution for simulated background events in which the electron parent particles are required to be SM mesons (left). Electron η distribution for simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}$ events (right). The $e\mu$ preselection with a loosened η requirement is applied in both plots.	139
C.1	Distribution of the arrival time of cosmic rays at their point of closest approach to the beamline as measured by the muon system in 2016 (left), 2017 (center), and 2018 (right) in data and simulation. Only cosmic-ray muons with $ d_0 < 8$ cm and $ z_0 < 20$ cm are considered.	144
C.2	Measured downward tracking efficiency versus cosmic ray muon arrival time in 2016 (left), 2017 (center), and 2018 (right) in data and simulation. Only cosmic-ray muons with $ d_0 < 8$ cm and $ z_0 < 20$ cm are considered.	144
C.3	Measured downward tracking efficiency in 2016 (top), 2017 (middle), and 2018 (bottom) versus $ d_0 $ (left) and $ d_z $ (right) in data and simulation. $ d_z $ ($ d_0 $) impact parameter is constrained to less than 20 cm (1.2 cm) when plotting against $ d_0 $ ($ d_z $).	146

D.1	Inclusive signal region cutflows for four signal samples in the ee (top) and $\mu\mu$ (bottom) channels. The top squark mass and proper decay length are 1000 GeV and 1 cm. In the sample corresponding to the black (red) curves, R-hadron material interactions are not modeled, but the top squark decay is performed in PYTHIA and the resulting b (d) quark produces a jet. In the samples corresponding to the blue and green curves, the R-hadron material interactions and decay are modeled with GEANT4. In the sample corresponding to the blue curves, the R-hadron decays to a lepton and a neutral pion, and in the sample corresponding to the green curves, the R-hadron decays to a lepton and a non-physical final-state quark.	151
D.2	Kinematic distributions of electrons (top), muons (middle), and jets (bottom) before any selection is applied for four signal samples in which the top squark mass and proper decay length are 1000 GeV and 1 cm. In the sample corresponding to the black (red) curves, R-hadron material interactions are not modeled, but the top squark decay is performed in PYTHIA and the resulting b (d) quark produces a jet. In the samples corresponding to the blue and green curves, the R-hadron material interactions and decay are modeled with GEANT4. In the sample corresponding to the blue curves, the R-hadron decays to a lepton and a neutral pion, and in the sample corresponding to the green curves, the R-hadron decays to a lepton and a non-physical final-state quark. In each plot, the rightmost bin contains overflow entries.	152
D.3	Kinematic distributions of electrons in the ee signal region (top), muons in the $\mu\mu$ signal region (middle), and jets in the $\mu\mu$ signal region (bottom) for four signal samples in which the top squark mass and proper decay length are 1000 GeV and 1 cm. In the sample corresponding to the black (red) curves, R-hadron material interactions are not modeled, but the top squark decay is performed in PYTHIA and the resulting b (d) quark produces a jet. In the samples corresponding to the blue and green curves, the R-hadron material interactions and decay are modeled with GEANT4. In the sample corresponding to the blue curves, the R-hadron decays to a lepton and a neutral pion, and in the sample corresponding to the green curves, the R-hadron decays to a lepton and a non-physical final-state quark. In each plot, the rightmost bin contains overflow entries.	153

Chapter 1: Introduction

Particle physicists seek to understand the fundamental constituents of nature and their interactions. At the present moment, the Standard Model of particle physics (along with the general theory of relativity) represents the most complete understanding humanity has yet achieved. At the same time, the Standard Model cannot explain gravity or several observed but not yet understood aspects of nature and therefore must give way to a more complete theory at a higher energy scale. One approach to understanding what might be hiding beneath the Standard Model is to use high-energy particle collisions to probe nature at ever smaller length scales. This thesis presents a search for beyond the Standard Model physics in 13 TeV proton-proton collision data collected by the Compact Muon Solenoid experiment at the Large Hadron Collider, which is the highest-energy particle collider ever constructed.

The search presented in this thesis targets new long-lived particles that produce displaced leptons, a unique signature that could evade many existing searches for new physics. In this chapter, I present theoretical context in the form of a brief overview of the Standard Model and targeted discussion of beyond the Standard Model physics and long-lived particles. I then present the experimental context with an overview of the Large Hadron Collider and the Compact Muon Solenoid detector in Chapter 2, present the search itself in Chapter 3, and conclude in Chapter 4.

1.1 The Standard Model

The Standard Model of particle physics (SM) describes all known particles and their non-gravitational interactions. Developed and experimentally verified over the past six decades, the SM posits the existence of twelve spin- $\frac{1}{2}$ particles, the fermions, that make up all observed matter; twelve spin-1 particles, the gauge bosons, that communicate the electromagnetic, weak, and strong forces; and one fundamental scalar, the Higgs boson, which breaks electroweak symmetry, giving mass to the gauge bosons and fermions.

The fermions and gauge bosons can be classified according to the forces with which they interact. The fermions are divided into six quarks, which carry color and interact via the strong force, and six leptons, which do not. Furthermore, all six quarks and three of the leptons carry electric charge and interact electromagnetically. The charged leptons include the electron, muon, and tau, and the neutral leptons are called neutrinos. All fermions interact via the weak force. The gauge bosons include the photon, which communicates the electromagnetic force; the W^+ , W^- , and Z bosons, which communicate the weak force; and eight gluons that communicate the strong force. Of these, only the W^+ and W^- bosons are electrically charged and only the gluons carry color charge. Finally, the fermions are grouped into three generations, each with two quarks, one charged lepton, and one neutral lepton. Figure 1.1 diagrams the grouping of the SM particles and lists some of their properties.

In the SM, the interactions between particles are governed by two theories: quantum chromodynamics, which describes the strong force, and the electroweak theory, which describes the electromagnetic and weak forces. The following sections provide a brief overview of these theories.

Standard Model of Elementary Particles

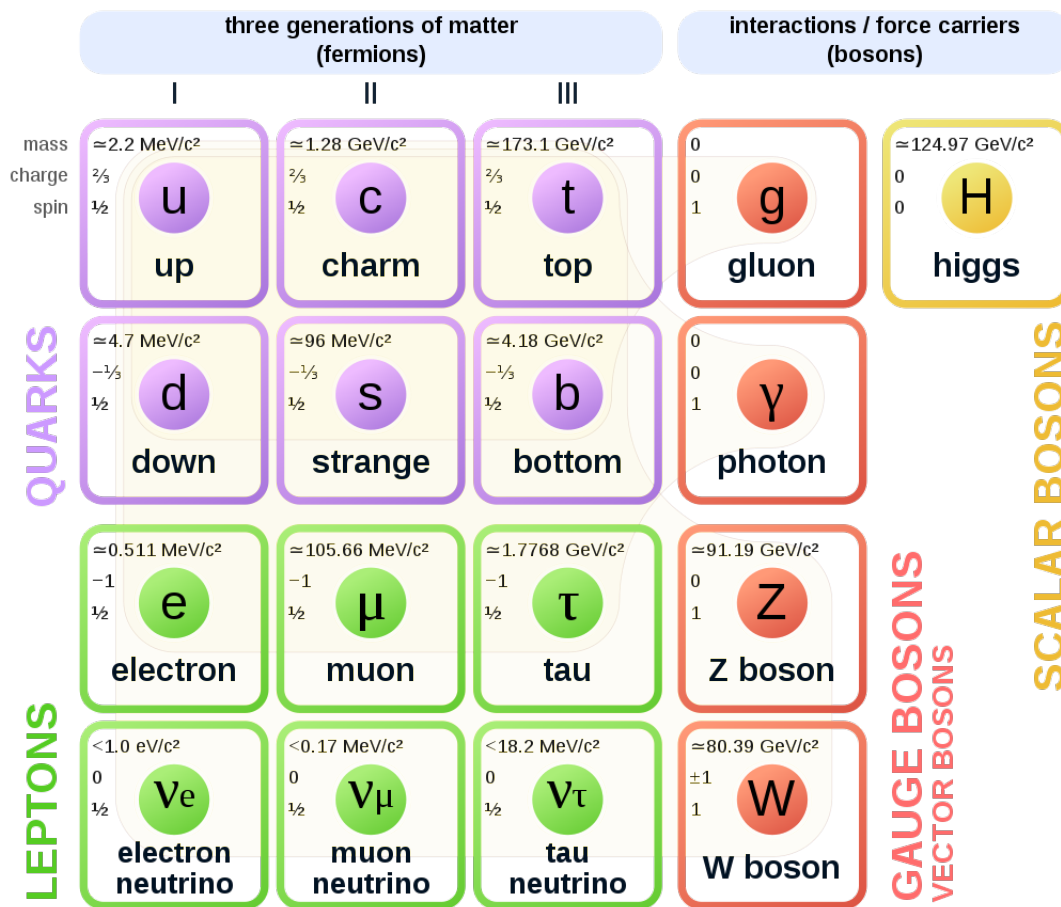


Figure 1.1: The Standard Model particle content.

1.1.1 Quantum chromodynamics

Quantum chromodynamics (QCD) describes the strong interactions between quarks and gluons and is based on the $SU(3)_c$ symmetry group, where the subscript c refers to color charge. In QCD, all quarks and gluons carry color charge, which allows interactions between two quarks of the same generation and a gluon, three gluons, or four gluons. QCD is responsible for the formation of all hadrons (such as protons and neutrons), and leads to two unique phenomena: confinement and asymptotic freedom [1, 2].

Confinement accounts for the experimental fact that an isolated particle with color charge has never been directly observed. Composite particles composed of quarks and gluons are always neutral under color, and attempts to separate the constituent particles will only produce new hadrons. This phenomenon is the result of the unique running of the strong coupling constant, which increases with decreasing energy (and therefore increasing distance).

Asymptotic freedom is the other side of the coin: if the strong coupling constant decreases as the interaction energy increases, then the strong interaction becomes more and more feeble at high energies. In high-energy interactions (such as those at the Large Hadron Collider), the strong coupling constant is in fact small enough to render the quarks nearly free. In this regime, perturbative calculations become possible.

1.1.2 The electroweak theory

The electroweak theory unifies the electromagnetic and weak interactions and is based on the $SU(2)_L \otimes U(1)_Y$ symmetry group. It posits two new charges: weak

isospin, which has three components, and hypercharge, Y . The third component of weak isospin, T_3 , is $\pm\frac{1}{2}$ for all left-handed fermions and 0 otherwise, while Y varies according to $Q = T_3 + \frac{1}{2}Y$, where Q is the familiar electric charge. Each generation of left-handed quarks or leptons forms an $SU(2)$ doublet. The first-generation doublets, for example, are:

$$\begin{pmatrix} \nu_e \\ e^- \end{pmatrix}_L, \begin{pmatrix} u \\ d \end{pmatrix}_L \quad (1.1)$$

where, as in $SU(2)_L$, the L denotes left-handed chiral states. The three generators of $SU(2)_L$ result in three massless spin-1 bosons: W^1 , W^2 , and W^3 , while $U(1)_Y$ gives rise to one massless spin-1 boson, B^0 . After electroweak symmetry is broken by the mechanism explained in the following subsection, the physical W^\pm bosons are identified as superpositions of W^1 and W^2 while the Z boson and the photon are identified as superpositions of W^3 and B^0 [3].

Terms in the electroweak Lagrangian involve either two left-handed fermions and a W^\pm or Z boson, two electrically charged particles and a photon, or charge-conserving combinations of W^\pm bosons, Z bosons, and photons that include three or four particles. Conspicuously missing, however, are mass terms for the electroweak gauge bosons or fermions.

1.1.3 The Higgs mechanism

As shown in Fig. 1.1, the fermions and W^\pm and Z bosons all have nonzero mass. Accounting for this fact within the context of the electroweak theory is difficult because explicit mass terms violate the gauge and chiral symmetry of $SU(2)_L \otimes U(1)_Y$. For example, a term such as

$$\frac{1}{2}m_A^2 A^\mu A_\mu, \quad (1.2)$$

which assigns mass m_A to gauge boson A , becomes

$$\frac{1}{2}m^2(A^\mu - \partial^\mu\alpha)(A_\mu - \partial_\mu\alpha) \neq \frac{1}{2}m^2 A^\mu A_\mu \quad (1.3)$$

under a $U(1)_Y$ gauge transformation, and a term such as

$$m_f \bar{f} f = m_f (\bar{f}_R f_L + \bar{f}_L f_R), \quad (1.4)$$

which assigns mass m_f to fermion f , breaks chiral symmetry by coupling the right- and left-handed components of the fermion.

If the gauge and chiral symmetries of $SU(3)_c \otimes SU(2)_L \otimes U(1)_Y$ truly are symmetries of nature and the fermions and W^\pm and Z bosons truly have nonzero mass, then another mechanism must be at work. Spontaneous symmetry breaking, which occurs when the vacuum state does not exhibit all of the symmetries of the underlying theory, fits the bill. In such a situation, each spontaneously broken continuous symmetry gives rise to a massless scalar particle [4]. In the case of spontaneously broken continuous *gauge* symmetries, however, there exists a mechanism by which the massless bosons are removed and some of the gauge bosons associated with the generators of the symmetries acquire mass [5, 6, 7]. In the SM, this mechanism, known as the Higgs mechanism, breaks electroweak symmetry, gives mass to the fermions and W^\pm and Z bosons, and results in one massive scalar particle, the Higgs boson.

The Higgs mechanism adds the scalar doublet

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.5)$$

whose potential is given by

$$V(\Phi^\dagger\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2, \quad (1.6)$$

to the SM. If $\mu^2 < 0$ and $\lambda > 0$, then $\Phi^\dagger\Phi = -\frac{\mu^2}{2\lambda}$ defines a circle of minima in the $\phi^+-\phi^0$ plane. Even though the potential remains invariant under $SU(2)_L \otimes U(1)_Y$, nature must spontaneously choose a vacuum state somewhere along this circle. Because the vacuum state does not respect $SU(2)_L \otimes U(1)_Y$, the symmetry is said to be spontaneously broken.

This procedure has three significant consequences. First, three of the four degrees of freedom originally associated with Φ are now associated with the longitudinal components of the W^\pm and Z bosons, which causes them to acquire mass while the photon remains massless. Second, Φ 's remaining degree of freedom adds a single massive scalar, the Higgs boson, to the theory. Third, the interaction between the fermions and the nonzero vacuum state of the scalar field produces fermion mass terms that obey chiral symmetry.

1.1.4 Current status

The SM is remarkably successful. It describes all known particles and their non-gravitational interactions, and it has passed countless experimental tests over the last several decades. Figure 1.2 gives an idea of the scale of this success by comparing theoretical predictions of SM production cross sections with measurements performed by the CMS experiment: theory and experiment agree across 41 different SM processes at 7, 8, and 13 TeV. In 2012, the CMS and ATLAS experiments independently discovered an approximately 125 GeV scalar particle with properties consistent with the SM Higgs boson [8, 9]. Further measurements in 7, 8, and 13 TeV proton-proton collisions at the Large Hadron Collider continue to agree with SM predictions of the Higgs boson properties [10, 11]. We finally have meaningful evidence as to the origin

facts of nature without explanation, but the history of science implies that a deeper understanding is likely hiding beneath the surface. Finally, the observed value of the Higgs boson mass is not only unexplained, it is unnatural. This final issue, which is explained in the following paragraphs, is a powerful motivation to search for new physics at currently accessible energy scales.

Naturalness

The naturalness criterion states that an effective theory such as the SM must not be overly sensitive to the details of the underlying higher-energy theory. For this independence of scales to hold, each dimensionless parameter in the effective theory must either be order one or protected by a custodial symmetry [15]. Such a criterion may or may not be respected by nature, but history and simple probability are on its side.

In the case of the SM, the dimensionless parameter in question is the ratio of the mass of the Higgs boson to the energy scale at which the higher-energy theory takes over. All SM parameters are affected by interactions with virtual particles through loop diagrams such as those shown in Fig. 1.3, but the Higgs boson mass is particularly sensitive. As a fundamental scalar, the Higgs boson lacks the chiral and gauge symmetries enjoyed by the fermions and gauge bosons. These symmetries, known as custodial symmetries, protect the fermion and gauge boson masses by guaranteeing that all corrections are proportional to the bare masses themselves. If the SM is indeed valid up to M_P , then matching the observed mass of the Higgs boson requires the bare mass of the Higgs boson to be coincidentally equal and opposite to the sum

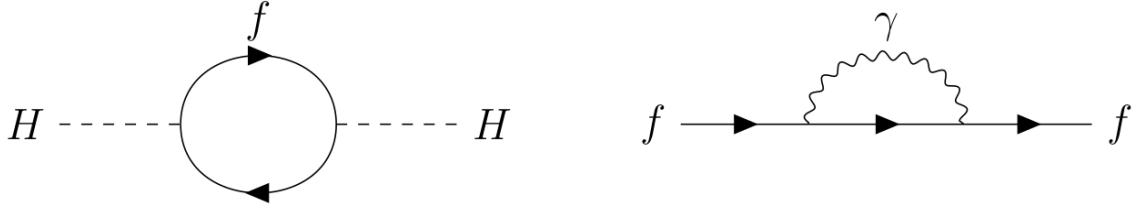


Figure 1.3: The Higgs boson mass corrected by a fermion loop (left) and a fermion mass corrected by photon loop (right).

of the terms that correct it, to approximately one part in 10^{32} [14]. While such miraculous fine tuning is technically possible, it could also be evidence of a deeper physical mechanism at work.

1.2 Beyond the Standard Model

Many beyond the SM (BSM) theories have been proposed to address the issues discussed in Section 1.1.4. Theories such as large extra dimensions address the unnatural Higgs boson mass by allowing gravity to spread across more than three spatial dimensions, which lowers M_P and therefore the size of the Higgs boson mass corrections [16]. Other theories posit new symmetries that protect the Higgs boson mass from large corrections. The following section explains supersymmetry, which protects the Higgs boson mass with a new symmetry between bosons and fermions.

1.2.1 Supersymmetry

Supersymmetry (SUSY) introduces a new symmetry in which every SM particle fits into a larger multiplet with an inherent symmetry between bosons and fermions. In its simplest form, SUSY predicts one new boson for every SM fermion, one new fermion for every SM boson, and one new Higgs doublet. The increase in particle multiplicity necessitates a new naming convention: the spin-0 superpartners of the SM fermions are called sfermions (e.g. sleptons or squarks) while the spin- $\frac{1}{2}$ superpartners to the SM bosons add “ino” to the end of their SM counterpart (e.g. Higgsino or wino).

When calculating contributions to the Higgs boson mass from loop diagrams, one finds that fermion loops differ in sign from boson loops, which means that in SUSY every bosonic correction to the Higgs boson mass is cancelled by a fermionic correction and vice versa. If SUSY were an exact symmetry of nature, the cancellation would be perfect, and the observed Higgs boson mass would match the bare Higgs boson mass exactly [17]. Figure 1.4 shows a sample leading-order cancellation.

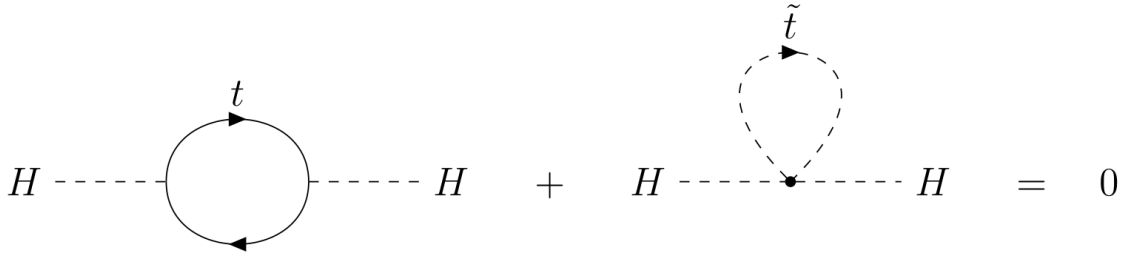


Figure 1.4: Corrections to the Higgs boson mass from the top quark (left) and top squark (right) cancel in exact SUSY. The top quark and top squark contributions are enhanced by the large coupling between the Higgs boson and the top quark.

Exact SUSY also requires that SUSY particles have the same mass as their SM counterparts, so the uniformly null results in collider searches imply that if SUSY exists, it must be a broken symmetry. In broken SUSY, the diagrams in Fig. 1.4 no longer exactly cancel. Instead, the resulting correction is proportional to the mass of the top squark [18], which means that broken SUSY can still resolve the naturalness problem if the masses of the SUSY particles that correct the Higgs boson mass are themselves approximately at the weak scale (say $\mathcal{O}(1 \text{ TeV})$). Many natural SUSY scenarios are therefore excluded as collider experiments set increasingly large lower bounds on SUSY particle masses. Figure 1.5, for example, shows that the exclusion bounds on top squark and gluino masses extend above 1 TeV in several recent CMS analyses.

As the available natural SUSY parameter space is further constricted, it is important to investigate signatures of new physics that conventional analyses may be missing. One possibility, new long-lived particles, is presented in the following section.

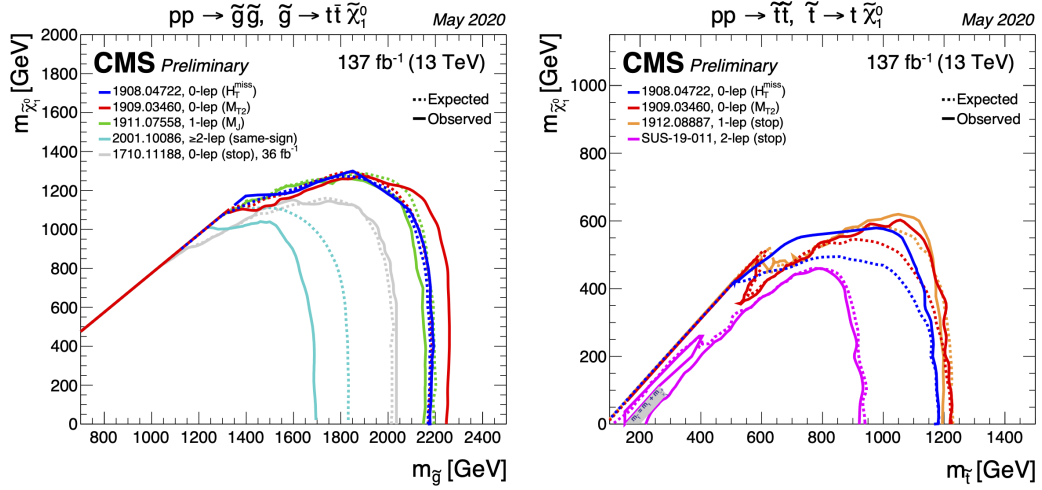


Figure 1.5: Mass limits at 95 % CL for a simplified model of gluino pair production with gluino decays to pairs of top quarks and the LSP (left) and top squark pair production with squark decays to a top quark and the LSP (right) from several CMS analyses [19].

1.2.2 Long-lived particles

In the context of collider physics, long-lived particles (LLPs) are particles whose lifetimes are such that they decay a measurable distance from the collision point. This category includes everything from particles that decay less than 1 mm away from the collision to particles that propagate through the entire detector. As shown in Fig 1.6, LLPs are common in the SM.

Long-lived SM particles arise from several mechanisms. First, symmetries such as charge and baryon number conservation ensure that particles such as electrons and protons are absolutely stable. Second, small coupling constants and highly virtual intermediate states decrease the decay rate of particles such as muons, whose 2.2 μs lifetime is the product of a weak decay through a virtual W boson (the W boson mass is about 760 times that of the muon). Finally, limited decay phase space increases

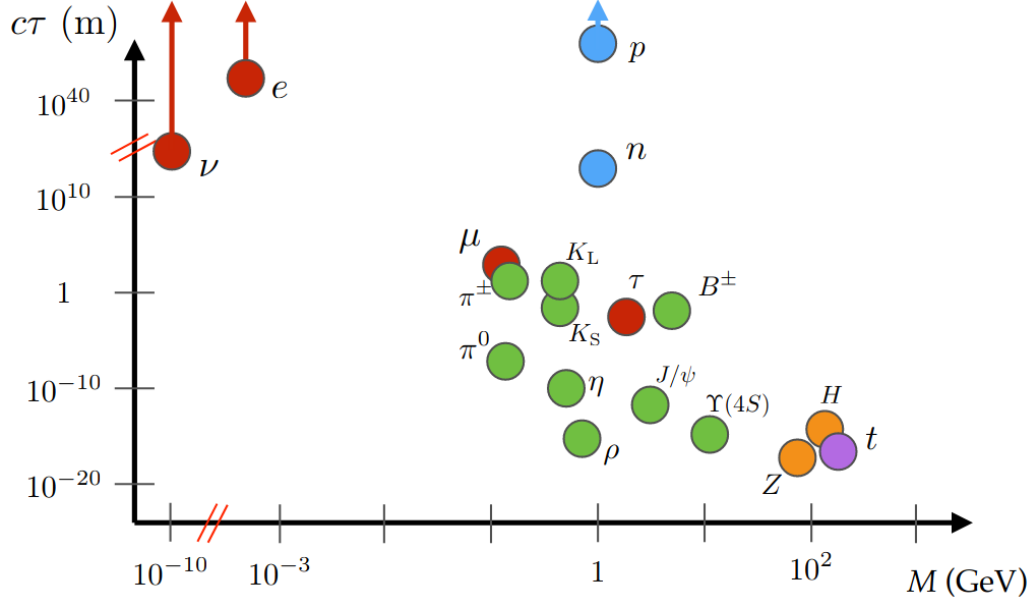


Figure 1.6: Masses and proper decay lengths of many Standard Model particles. Particles with proper decay lengths above approximately 10^{-4} m will be noticeably long-lived in collider detectors such as CMS [20].

the lifetime of particles such as the neutron, whose decay into a proton, an electron, and an electron neutrino is slowed by the near mass degeneracy of the neutron and the proton. The mass difference between the neutron and its decay products is less than 1 MeV. The muon and neutron decays are diagrammed in Fig. 1.7.

In BSM physics, the same fundamental mechanisms could very well produce new LLPs. Many SUSY scenarios, for example, introduce a symmetry known as R-parity that prevents proton decay. In models with exact R-parity conservation, SM particles are assigned R-parity 1 and SUSY particles are assigned R-parity -1 . Conserving the product of R-parity at each vertex has two phenomenological consequences: SUSY particles must be produced in pairs, and the lightest SUSY particle (LSP) must be absolutely stable. A neutral, weakly interacting LSP would therefore pass through

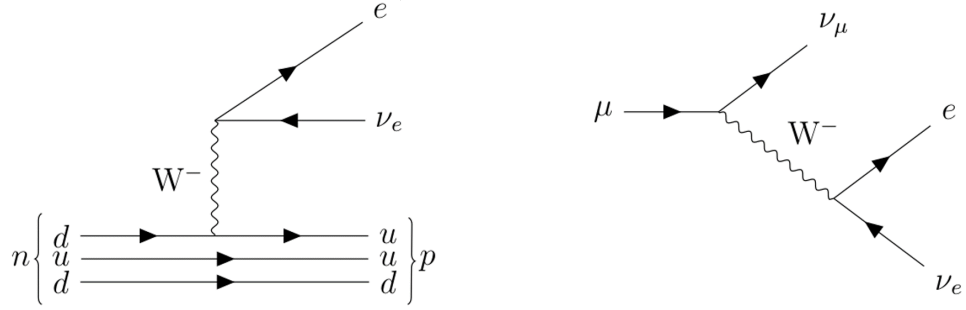


Figure 1.7: Long-lived decays of the neutron (left) and muon (right).

most detectors (including the CMS detector described in Section 2.2) without interacting. The resulting momentum imbalance is a standard signature in many SUSY searches [21].

On the other hand, SUSY models with weakly coupled R-parity-violating (RPV) terms produce long-lived but not perfectly stable LSPs. The following section will give a detailed overview of one such model. A similar situation arises in gauge-mediated SUSY breaking models where the gravitino is the LSP. The strength of the coupling between the next-to-LSP (NLSP) and the gravitino is inversely proportional to the energy scale at which SUSY is broken. A high SUSY breaking scale therefore suppresses the NLSP decay rate, making it long lived [22].

LLPs also arise from particular SUSY mass spectra. Models in the Split SUSY paradigm, for example, propose that the spin-0 SUSY particles are significantly more massive than the spin- $\frac{1}{2}$ SUSY particles. In these models, the gluino becomes long lived when its decay to two quarks and a neutral spin- $\frac{1}{2}$ SUSY particle is suppressed by a highly virtual intermediate squark [23]. Other SUSY models produce LLPs by limiting decay phase space. Some anomaly-mediated SUSY breaking models, for

example, predict that the NLSP and LSP are nearly degenerate in mass. Just like the neutron decaying into a proton, the lack of available phase space suppresses the decay and produces a long-lived NLSP [24].

In summary, LLPs are a general feature of the SM, and it is reasonable to assume that the same mechanisms that produce SM LLPs will also manifest in BSM physics. The following subsection gives an overview of the phenomenology of a SUSY model relevant to the analysis presented in Chapter 3, while the discussion of experimental details of this model and LLP searches at the Large Hadron Collider will be saved until after describing the Large Hadron Collider and the CMS experiment in Chapter 2.

1.2.3 Displaced supersymmetry

As mentioned in Section 1.2.1, weak-scale SUSY has the potential to explain the seemingly unnatural observed value of the Higgs boson mass. With this appealing outcome in mind, experimental physicists have been searching for signs of SUSY in high-energy particle collisions for the last few decades. In particular, searches at the Large Hadron Collider, where the 7, 8, and 13 TeV proton-proton collisions could potentially be producing SUSY particles with masses above 1 TeV, are actively excluding large swaths of the natural SUSY parameter space. By examining the common assumptions behind a majority of these searches, the proponents of the Displaced SUSY model find an approach that avoids proton decay constraints while naturally producing long-lived SUSY particles that would be undetected by most collider searches [25].

Displaced SUSY is an RPV SUSY model that respects proton lifetime constraints by allowing terms that violate lepton number but not baryon number. The Minimal

Supersymmetric Standard Model, which is a simple supersymmetric extension of the SM, allows the following baryon and lepton violating operators:

$$\frac{1}{2}\lambda''_{ijk}U_iD_jD_k, \frac{1}{2}\lambda_{ijk}L_iL_jE_k, \lambda'_{ijk}L_iQ_jD_k, \epsilon_iL_iH_u \quad (1.7)$$

where the first term violates baryon number and the remaining terms violate lepton number. The i , j , and k indices run over the three generations of fermions and U , D , L , E , Q , and H refer to the SUSY multiplets whose SM components are right-handed up-type quarks, right-handed down-type quarks, left-handed leptons, right-handed leptons, left-handed quarks, and Higgs bosons, respectively [17]. Most SUSY models introduce R-parity to forbid all of these terms and therefore disallow proton decay, but this approach may be overkill: separately conserving either lepton number or baryon number is sufficient to prevent proton decay.

Such a situation can arise naturally if a gauge-unifying, R-parity-conserving $SU(5)$ theory exists at high energies but is broken at lower energies [26]. In this scenario, the baryon-number-violating terms are suppressed and mixing between L and H becomes possible. The final term in expression (1.7) then dominates, and we are left with the following lepton-number-violating terms after rotating to the mass basis:

$$\epsilon_i y_{jk}^d L_i Q_j D_k, \epsilon_i y_{jk}^e L_i L_j E_k \quad (1.8)$$

where the ϵ factors are lepton-Higgs mixing angles and the y factors are the Yukawa coupling constants associated with the interactions between the fermions and the SM Higgs field. Because these y factors are proportional to fermion mass, the lepton-violating processes will favor third-generation fermions.

The dimensionless lepton-Higgs mixing angles are protected by the custodial R-parity conservation of the higher-energy theory and will therefore naturally be small.

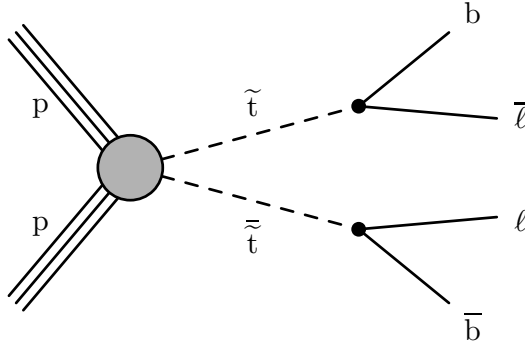


Figure 1.8: Pair-produced top squarks that each decay to a bottom quark and a lepton through an R-parity-violating LQD vertex.

Just as in the SM, these small coupling constants can lead to macroscopic SUSY particle lifetimes that would evade most collider searches. In particular, long-lived LSP squarks can decay to a quark and charged lepton through a displaced vertex, as in Fig. 1.8.

The distinct experimental signature of processes such as the one shown in Fig. 1.8 is a major motivating factor behind the analysis presented in Chapter 3. The particular experimental consequences of such processes will be further explored in that context.

Chapter 2: The Large Hadron Collider and Compact Muon Solenoid experiment

2.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is the highest-energy particle collider ever constructed. Located at the European Organization for Nuclear Research (CERN) and housed in a 27 km ring approximately 100 m below the French/Swiss countryside, the LHC is designed to accelerate two counter-rotating beams of protons to 7 TeV (and sometimes beams of heavy ions to 2.8 TeV) and collide them at four points around the ring. Each collision point is instrumented with a dedicated detector: the ATLAS (A Toroidal LHC ApparatuS) and CMS (Compact Muon Solenoid) experiments are general purpose detectors designed to reconstruct the remnants of proton-proton collisions at the highest collision rates offered by the LHC, LHCb (LHC beauty) studies b-quark decays from proton-proton collisions produced at lower collision rates, and ALICE (A Large Ion Collider Experiment) studies heavy-ion collisions [27]. Figure 2.1 shows the location of each experiment around the LHC ring. The analysis presented in Chapter 3 utilizes proton-proton collision data collected by the CMS experiment, and the following discussion is focused accordingly.

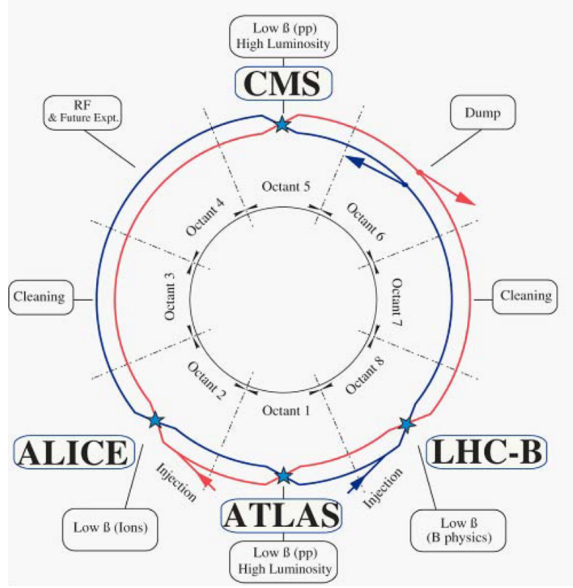


Figure 2.1: Layout of the LHC experiments [27].

2.1.1 Injection chain

The protons ultimately collided by the LHC must first travel through much of the CERN accelerator complex, which is diagrammed in Fig. 2.2. The 6.5 TeV proton beams relevant to this thesis start their journey as the nuclei of hydrogen atoms in a bottle of hydrogen gas. After having their electrons stripped away with an electric field, the protons are accelerated to an energy of 50 MeV with the Linac 2 linear accelerator [28]. The Proton Synchrotron Booster (PSB) next accelerates the protons to an energy of 1.4 GeV before injecting them into the Proton Synchrotron (PS) [29]. The PS was the first synchrotron constructed at CERN and was the highest-energy particle accelerator in the world at the time of its first operation [30]. Today, it accelerates protons to an energy of 25 GeV before passing them along to the Super Proton Synchrotron (SPS), which is the 7 km proton-antiproton collider at which the

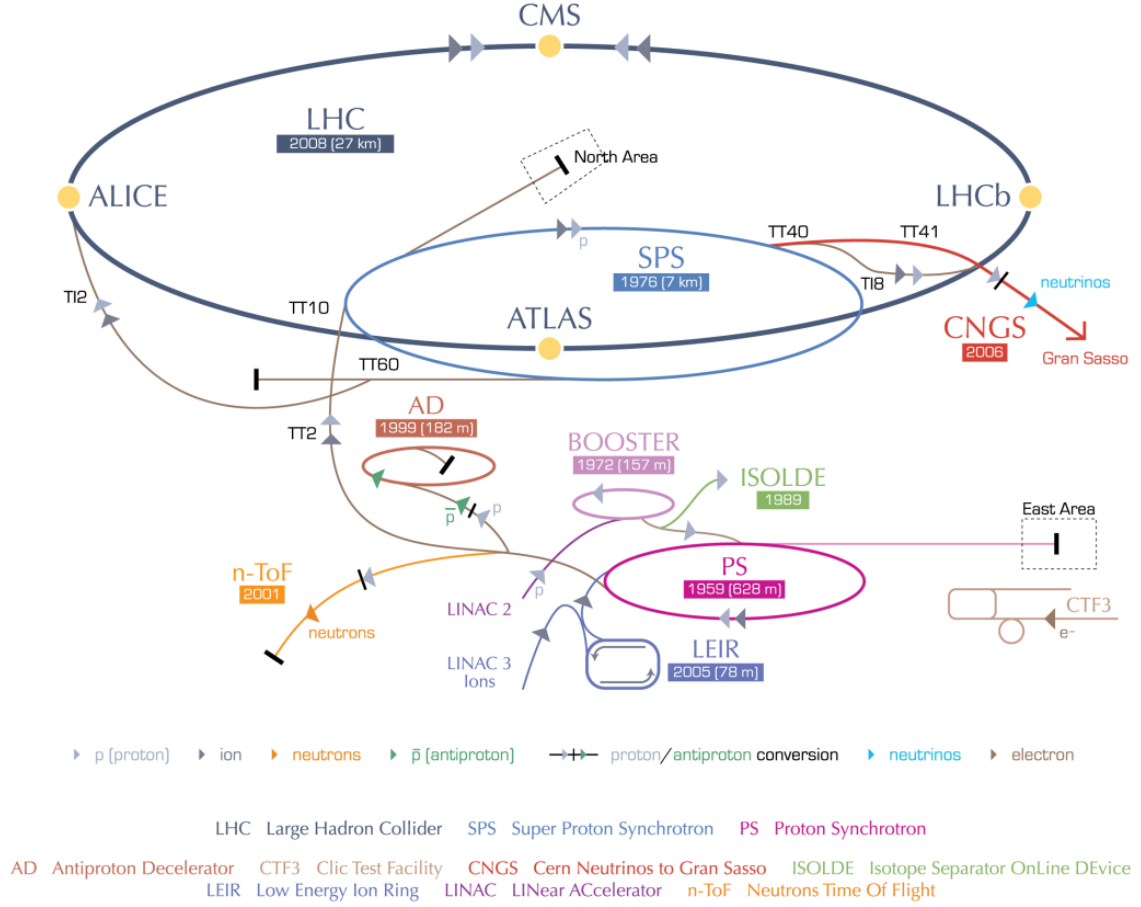


Figure 2.2: A diagram of the CERN accelerator complex. The analysis presented in Chapter 3 utilizes protons accelerated by LINAC 2, BOOSTER (also known as PSB), PS, SPS, and finally the LHC before their ultimate collision inside CMS [32].

W^\pm and Z bosons were discovered in 1983 [31]. As the final step in the LHC injection chain, SPS accelerates protons to an energy of 450 GeV before injecting them into the LHC [27].

2.1.2 Main ring

Maximizing the physics potential of the LHC requires simultaneously maximizing the collision energy and the number of interesting collisions per unit time. The main

Table 2.1: Luminosity parameters used in Eq. (2.1).

Parameter	Description
N_b	Number of particles per bunch
n_b	Number of bunches per beam
f_{rev}	Revolution frequency
γ_r	Relativistic gamma factor
ϵ_n	Normalized transverse beam emittance
β^*	Beta function at collision point
F	Geometric luminosity reduction factor

LHC ring is therefore designed to accelerate the 450 GeV protons it receives from SPS to 7 GeV and collimate them into intense beams to be collided at high rates. The number of interesting collisions per unit time is ultimately the product of the total cross section of the processes one deems interesting, σ , and the instantaneous luminosity, L , which is given by:

$$L = \frac{N_b^2 n_b f_{rev} \gamma_r}{4\pi \epsilon_n \beta^*} F \quad (2.1)$$

where the parameters are defined as in Table 2.1 [27].

The LHC is designed to deliver a maximum instantaneous luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ to ATLAS and CMS, but operational improvements, most notably a reduction in ϵ_n and β^* , have allowed the LHC to exceed this goal by up to a factor of approximately two in the 2016, 2017, and 2018 data-taking periods [33]. As shown in Fig. 2.3, the total integrated luminosity delivered during this period is approximately a factor of five times greater than that of the 2011–2012 period.

Increasing the number of interesting collisions per unit time comes with the unfortunate side effect of increasing the total number of collisions in each bunch crossing. Due to storage and processing limitations, the CMS and ATLAS detectors only record

data from the small subset of bunch crossings whose properties imply that they may contain interesting processes (see Section 2.2.6). In a given bunch crossing, usually only one collision is considered interesting. The remaining collisions, referred to as “pileup”, add a layer of difficulty to collecting and analyzing the data while increasing the total radiation dose that the detectors must withstand.

The ability to produce high-energy proton collisions at such high rates depends on several impressive technological feats, notably the superconducting magnets that steer and shape the beams and the superconducting radio-frequency (RF) cavities that accelerate the protons and determine their bunch structure.

Superconducting magnets

The LHC magnet system relies on superconducting NbTi magnets that are cooled to below 2 K with superfluid helium and are capable of producing fields in excess of 8 T. The design of the main dipole magnets that are responsible for keeping the beams in a circular trajectory is heavily influenced by the size of the LHC tunnel, which originally housed the Large Electron-Positron Collider (LEP). Unlike LEP, which collided particles and antiparticles, the LHC requires two separate beam pipes, each with its own dipole magnetic field. This requirement, along with the limited tunnel cross section, motivates the “two-in-one” magnet design in which both superconducting magnets share a common cold mass and cryostat, as shown in Fig. 2.4 [27].

In addition to the main dipole magnets, the LHC also employs quadrupole magnets for beam focusing and sextupole, octupole, and decapole magnets to correct the field at the edges of the dipoles [27].

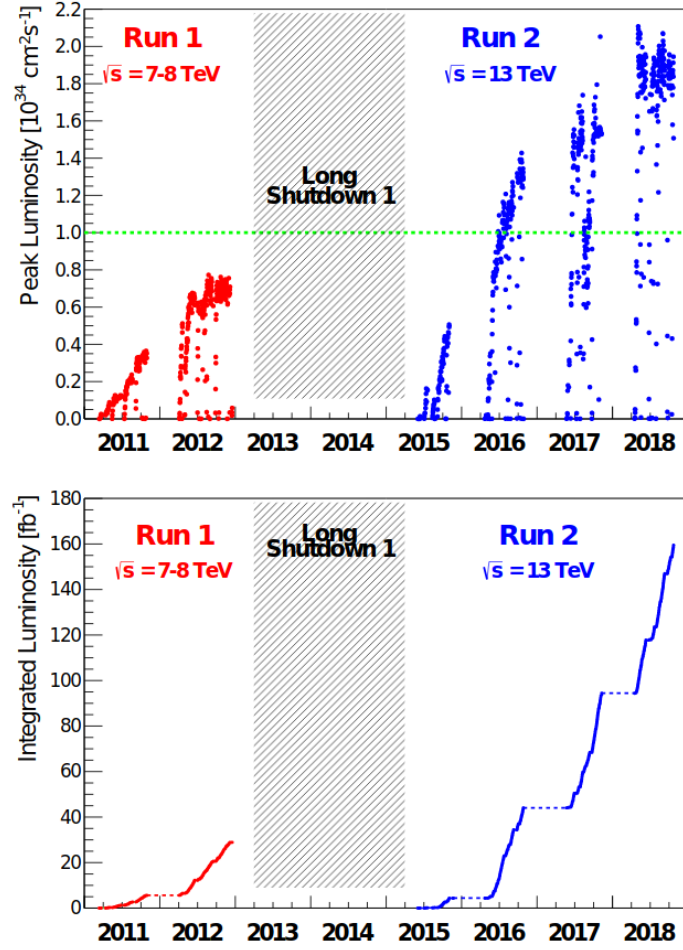


Figure 2.3: The peak instantaneous (top) and integrated (bottom) luminosity delivered by the LHC during proton operation between 2011 and 2018 [33].

LHC DIPOLE : STANDARD CROSS-SECTION

CERN AC/DI/MM - HE107 - 30 04 1999

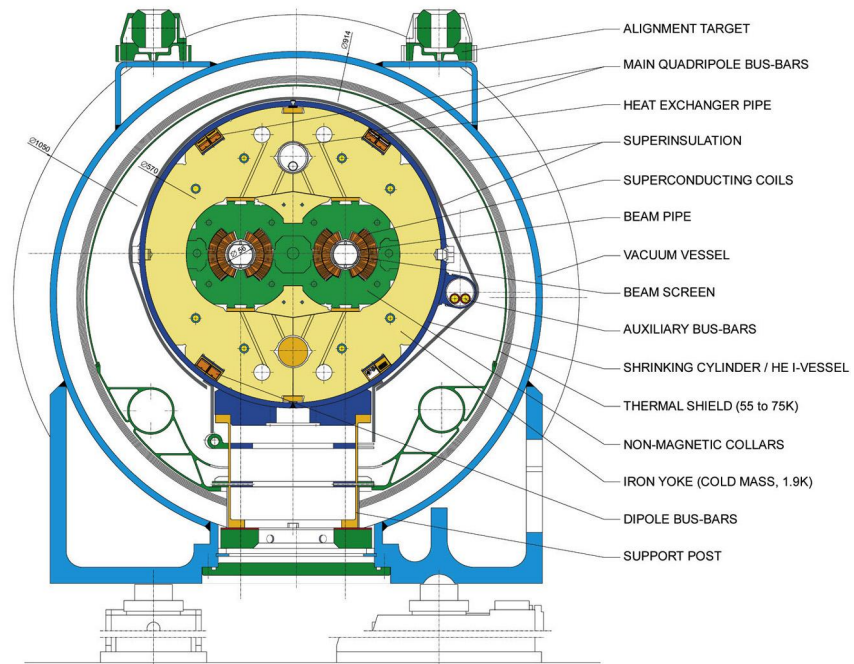


Figure 2.4: Diagram of an LHC dipole magnet in cross-section [34].

Radio-frequency cavities

An RF superconducting cavity system is responsible for capturing the 450 GeV protons injected into the LHC from SPS, accelerating them to the full collision energy, defining their bunch structure, and storing them. The main RF system operates at 400 MHz and is located in Octant 4 (see Fig.2.1). Each RF cavity contains an oscillating electromagnetic field whose phase is synchronized with the arrival of the proton bunches such that the protons passing through the cavity always feel a force in the direction of their motion. The applied force naturally varies for protons that are slightly out of phase in such a way as to keep the protons tightly bunched in the longitudinal direction [27].

2.2 The Compact Muon Solenoid experiment

The Compact Muon Solenoid experiment (CMS) is designed for the general study of the highest-energy, highest-luminosity proton-proton (and heavy ion) collisions the LHC can provide. The detector design is driven by the particular goals of exploring physics at the TeV scale and discovering the origin of electroweak symmetry breaking [35]. More than 4000 collaborators from institutions across more than forty countries work together to collect and analyze the data using a global computing grid [36].

To reconstruct the variety of particles that emerge from high-energy proton-proton collisions, CMS uses several complimentary subdetectors nested radially about the collision point. A 4 T superconducting solenoid magnet provides a powerful magnetic field to bend the trajectories of charged particles, thus enabling the determination of their momenta. Working from the center out, the CMS detector consists of an all-silicon tracker, a lead-tungstate scintillating crystal electromagnetic calorimeter, a sampling hadronic calorimeter composed of brass absorber and plastic scintillator tiles, the superconducting solenoid magnet, and a muon system with three varieties of gaseous detectors. Figure 2.5 shows the detector layout, and the remainder of this chapter is devoted to a brief overview of each subsystem as well as the triggering and reconstruction strategies employed by CMS.

CMS uses a right-handed coordinate system centered on the nominal collision point with positive x direction pointing towards the center of the LHC ring and the positive y direction pointing vertically upward. The azimuthal angle in the x - y plane, denoted ϕ , is measured from the positive x axis, and the polar angle θ is measured from the positive z axis. The angle from the z axis is more commonly described

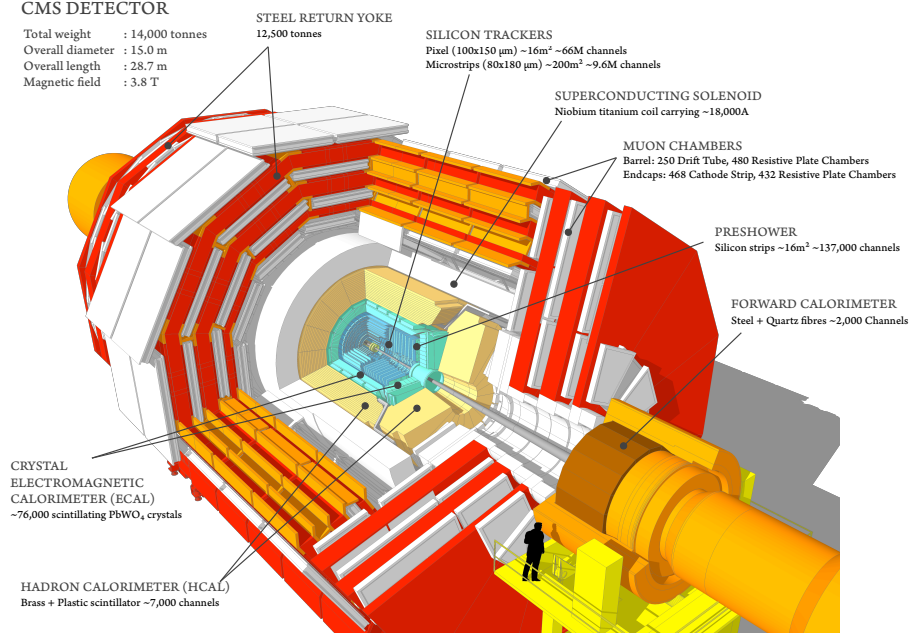


Figure 2.5: The CMS detector [37].

in terms of the pseudorapidity $\eta = -\ln \tan(\theta/2)$. Distances in the η - ϕ plane are commonly referred to as $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. The component of the momentum that is transverse to the beam direction (i.e., in the x - y plane) is denoted p_T , and the magnitude of the negative vector sum of the p_T of all the reconstructed particles in an event is denoted p_T^{miss} . The magnitude of p_T is also sometimes referred to as transverse energy or E_T [38, 39].

The CMS detector has undergone several upgrades since its initial construction. The description presented here will focus on the detector conditions relevant to the analysis presented in Section 3.

2.2.1 Solenoid magnet

The superconducting solenoid is designed to produce a 4 T magnetic field throughout the 6.3 m diameter, 12.5 m long cylindrical volume that contains the tracker and calorimeters. The magnetic field is produced by running 19 kA through 2168 turns of NbTi superconducting cable that are cooled with liquid helium. The flux returns through an iron yoke that also houses the muon system [39].

The strong magnetic field is critical to CMS's ability to unambiguously distinguish muons and anti-muons with transverse momenta up to 1 TeV [38], and much of the overall detector design is guided by the choice of a large superconducting solenoid. The uniform magnetic field alters the trajectories of charged particles immediately upon their production, which results in significant bending power within a relatively small radius and therefore enables a compact detector.

The CMS solenoid stores a uniquely large amount of energy in its magnetic field when compared to other collider detector magnets, especially when viewed relative to its mass. To avoid deformations from the strong magnetic field, the superconducting coils are reinforced with an aluminium alloy so that the coil layers themselves handle 70% of the magnetic hoop stress. Figure 2.6 compares several collider detector magnets in the energy/mass vs energy plane and shows a cross-sectional view of the reinforced conductor coils in cross section. This approach allows for a relatively thin solenoid that is less likely to scatter muons before they reach the muon system [39].

2.2.2 Tracker

In the region closest to the proton collisions, CMS employs a silicon tracker to reconstruct particle trajectories along with primary and secondary vertices. Efficiently

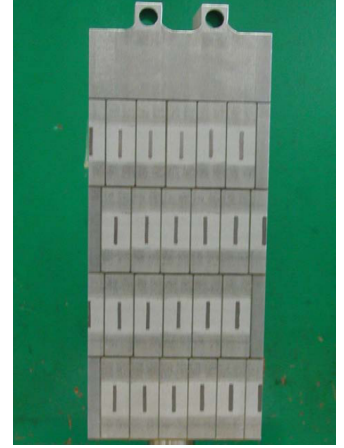
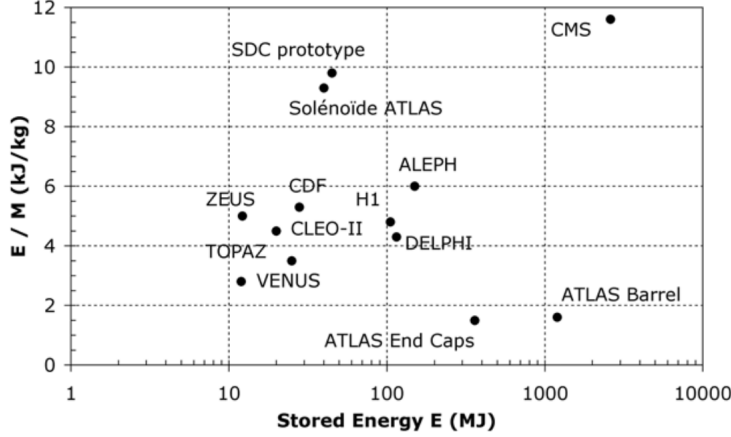


Figure 2.6: The stored-energy-over-mass ratio, E/M , for several detector magnets (left), and a cross-sectional view of the four-layer winding of reinforced conductor in the CMS superconducting solenoid (right) [39].

and accurately performing these tasks allows CMS to measure charged particle momenta, distinguish between particles from the primary and pileup vertices, and identify heavy-flavor decays. The high particle flux necessitates a highly granular, fast, and radiation hard detector that introduces the smallest possible amount of material to the region inside the calorimeters. The resulting detector contains over 200 m^2 of active silicon, making it the largest silicon tracker ever built. The tracker has a length of 5.8 m, a diameter of 2.5 m, and is divided into two subdetectors. Inside a radius of 20 cm, the particle flux demands the use of silicon pixel detectors, while silicon micro-strip detectors suffice in the region beyond 20 cm [39].

The pixel and strip detectors both utilize the same fundamental detection mechanism: the energy depositions of incident charged particles form electron-hole pairs when passing through the depletion region of a reverse-biased p-n junction, and the resulting induced current signifies the presence of a charged particle. The reverse bias

voltage sweeps away charge carriers to reduce thermal noise and maximize the sensitive volume. The fine two-dimensional segmentation of each pixel detector layer, along with its minimal thickness and known radius, provides a precise three-dimensional position measurement. Each strip detector layer, in contrast, takes advantage of its minimal thickness, known radius, and fine one-dimensional segmentation to provide a precise two-dimensional position measurement. In both cases, the detector planes are tilted to allow the charge liberated by a single incident particle to spread across multiple pixels or strips. This scheme significantly improves the position resolution by combining the charge measurements of multiple adjacent pixels or strips.

The original pixel detector was replaced between the 2016 and 2017 data-taking periods in preparation for higher luminosities [40]. As the analysis presented in Section 3 uses data collected in 2016–2018 and is particularly dependent on tracker measurements, the original pixel detector, 2017–2018 (Phase-1) pixel detector, and strip detector are described separately below.

Original pixel detector

The original CMS pixel detector covers the $|\eta| < 2.5$ region and is composed of three cylindrical barrel layers at $r = 4.4, 7.3,$ and 10.2 cm and four endcap disks 34.5 and 46.5 cm up and down the beamline from the nominal collision point. Each layer or disk is instrumented with several pixel modules that are each composed of a silicon sensor bump bonded to custom ASIC read-out chips (ROCs). Each sensor is $285\text{ }\mu\text{m}$ thick and typically comprises $66\ 560\ 100 \times 150\text{ }\mu\text{m}$ pixels. The nearly square pixel shape enables approximately $15\text{ }\mu\text{m}$ hit resolution in both the r - ϕ and z directions [38, 39].

The silicon pixel sensors are composed of highly doped n-type silicon implants in an n-type silicon substrate with a p-n junction on the opposite side. These “n-on-n” sensors are more expensive than traditional “p-on-n” sensors but are chosen for their radiation hardness. Specifically, the increased electron contribution to the signal charge lessens the impact of the radiation-damage-induced reduction in carrier lifetime [39, 21]. Signals whose analog pulse height exceeds a pre-determined value (set individually for each pixel) are read out and amplified by the bump-bonded ROCs. The resulting hit location and pulse height information are read out at each L1 trigger (see Section 2.2.6).

The original pixel detector was designed for a maximum instantaneous luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, which corresponds to approximately 25 pileup collisions per bunch crossing in CMS. As shown in Fig. 2.3, the LHC first exceeded this instantaneous luminosity in 2016. Figure 2.7 shows the resulting hit efficiency loss in the innermost layer of the pixel detector as a function of instantaneous luminosity and pileup. While the higher-radius pixel layers are less affected, important quantities such as the transverse impact parameter resolution are highly dependent on the measurements of the layer closest to the interaction. To mitigate this effect, CMS replaced the original pixel detector between the 2016 and 2017 data-taking periods.

Phase-1 pixel detector

The Phase-1 pixel detector represents an incremental improvement over the original CMS pixel detector: the same fundamental technology fills the same physical footprint and reuses many of the existing services but nevertheless achieves higher rate capabilities, improved radiation tolerance, and more robust tracking [40]. This

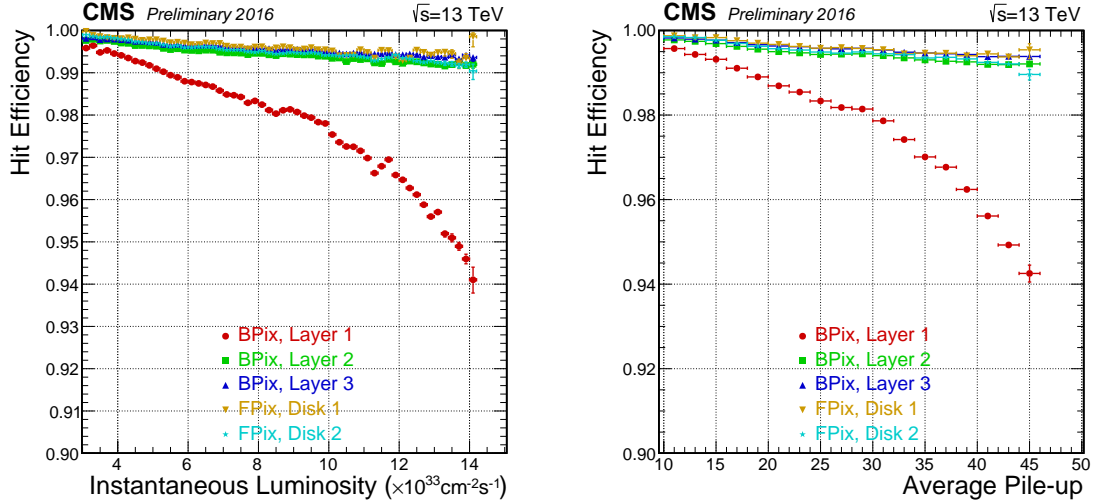


Figure 2.7: Measured single hit efficiency per layer as a function of the instantaneous luminosity (left) and pileup interactions (right) in data taken with the original CMS pixel detector in 2016 [41].

is accomplished by adding one additional layer to the barrel and each endcap, decreasing the radius of the innermost barrel layer to 2.9 cm, upgrading the ROCs, and reducing the material budget of the cooling system and mechanical structure [40, 42]. Figure 2.8 compares the geometries of the original and Phase-1 pixel detectors.

The loss of hit efficiency at high instantaneous luminosity observed in the original pixel detector is significantly reduced in the Phase-1 pixel detector. Despite the higher particle flux that accompanies the shift to a smaller radius, the Phase-1 innermost pixel layer maintains a single hit efficiency well over 98 % when operating at an instantaneous luminosity of $1.4 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ during the 2017 data-taking period [43]. Looking at Fig. 2.7, the equivalent quantity for the original pixel detector is approximately 94 %.

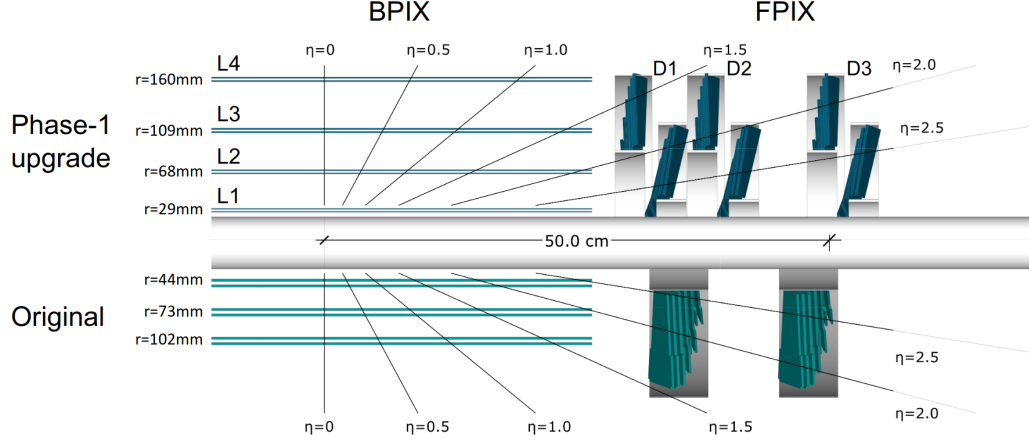


Figure 2.8: Comparison of the original and Phase-1 CMS pixel detector layouts in y - z plane [40].

Strip tracker

The strip tracker surrounds the pixel detector with silicon micro-strip sensors in 10 cylindrical barrel layers between $r = 20$ cm and $r = 110$ cm and 12 disks on each side of the barrel detector that extend to $|z| = 282$ cm and cover up to $|\eta| < 2.5$. The strip pitch generally increases with radius and results in hit resolutions that vary from 23 to 530 μm [39]. Figure 2.9 shows the layout of the entire silicon tracker.

Sitting farther from the interaction point allows the strip tracker to utilize more typical p-on-n type sensors that are composed of highly doped p-type silicon strips embedded in an n-type silicon substrate with a layer of highly doped n-type silicon on the side opposite the strips. As the n-type silicon substrate is slowly converted to p-type silicon by bulk radiation damage, the p-n junction moves from the strips to the back side of the sensor. Custom integrated circuits amplify, shape, and store the analog signals from each strip sensor. At each L1 trigger, the analog signals are

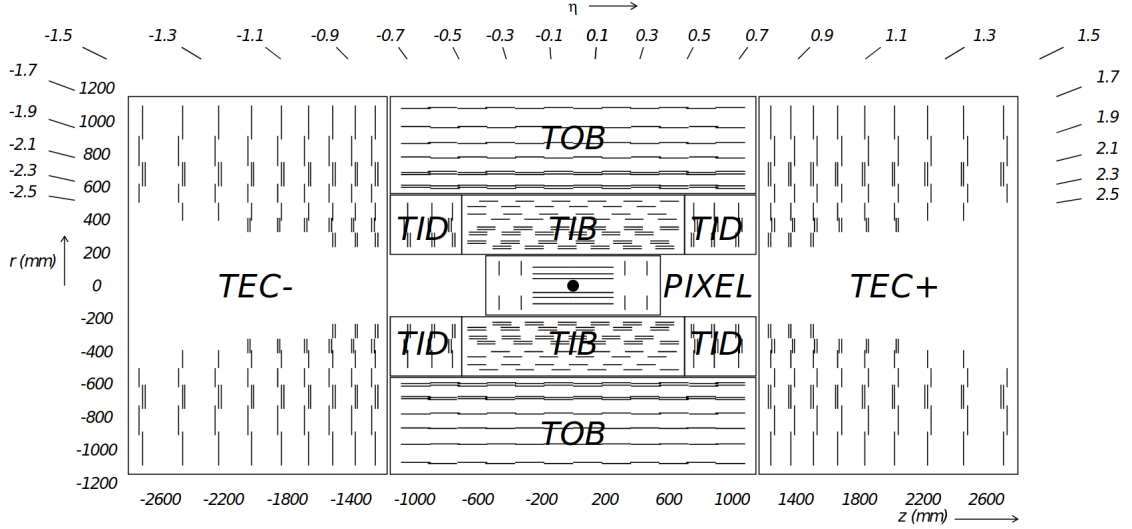


Figure 2.9: Layout of the CMS silicon tracker. TIB, TOB, TID, and TEC refer to subdetectors of the strip detector while PIXEL refers to the original pixel detector. The Phase-1 pixel detector is contained within the same volume [39].

transferred off-detector by optical fibers to be processed. Reading out the full analog information reduces the tracker material budget and need for custom radiation-hard electronics by performing analog-to-digital conversion, pedestal subtraction, and other processing outside the detector volume [39].

2.2.3 Electromagnetic calorimeter

After traversing the inner tracker, particles next encounter the electromagnetic calorimeter (ECAL). As a homogeneous scintillation calorimeter, ECAL uses 61 200 lead tungstate crystals in the barrel and 7324 in each endcap to reconstruct the energy deposited during electromagnetic showers. Lead tungstate crystals allow for a fast (80 % of light emitted within 25 ns), compact (radiation length = 0.89 cm),

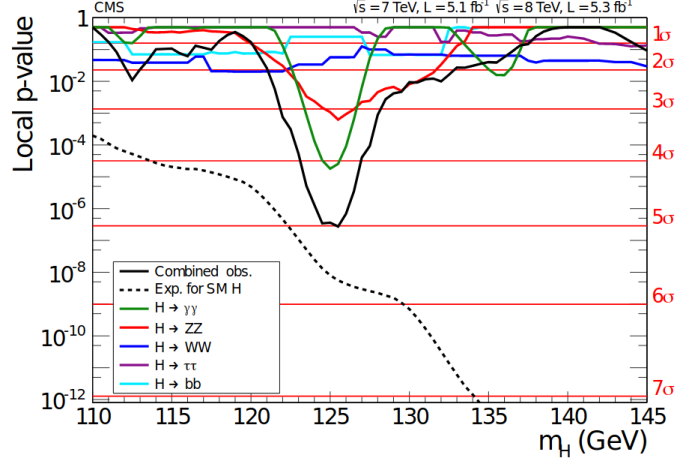


Figure 2.10: The observed local p-value for the five decay modes and the overall combination as a function of the SM Higgs boson mass showing the importance of the ECAL mass resolution in the discovery of the Higgs boson by the CMS experiment. The dashed line shows the expected local p-values for a SM Higgs boson with a mass m_H [8].

fine-grained (Molière radius = 2.2 cm), and radiation hard (up to 10 Mrad) calorimeter. The main drawback is the relatively low light yield (30 photon/MeV), which necessitates photodetectors with intrinsic gain that work in magnetic fields [39, 38].

The ECAL performance requirements were heavily influenced by the possibility to reconstruct the decay of a Higgs boson to two photons [35]. Despite the small branching fraction and irreducible background, this decay channel provides a clean signature of a narrow mass peak on top of a smoothly falling background. Thanks in large part to the excellent ECAL energy resolution, the di-photon channel provided the largest significance and best mass resolution (approximately 1 GeV resolution at 125 GeV) in the CMS Higgs boson discovery in 2012, as shown in Fig. 2.10 [8].

The barrel section extends radially from 129 to 177 cm and covers up to $|\eta| < 1.479$. The crystals are tapered to approximately project back to the nominal collision point

but not so perfectly that likely particle trajectories align with cracks. Each crystal is approximately one Molière radius wide and 25 radiation lengths deep. The crystals in each endcap section are arranged in an x - y grid that starts at $|z| = 315$ cm and covers $1.479 < |\eta| < 3.0$.

2.2.4 Hadronic calorimeter

Particles that survive the ECAL will next encounter the hadronic calorimeter (HCAL). As the ECAL constitutes approximately 25 radiation lengths but only one interaction length, all but the particles that decay through the strong force will be filtered out before reaching the HCAL. In addition to reconstructing the decays of hadrons, the HCAL plays a particularly important role in measuring p_T^{miss} . By maximizing the coverage in η and overall amount of material in terms of interaction lengths, HCAL ensures that nearly all particles (other than muons, neutrinos, and hypothetical BSM particles) decay and deposit all their energy before reaching the solenoid. Muon momentum is reconstructed with the tracker and muon system, so only neutrinos and hypothetical BSM particles will contribute to p_T^{miss} . Reliable p_T^{miss} measurements are particularly important when searching for new weakly interacting particles with large lifetimes such as potential dark matter candidates or R-parity-conserving LSPs that can pass through the detector without interacting.

With these goals and the constraint of fitting within the solenoid volume in mind, HCAL is designed as a sampling calorimeter that uses 3.7 mm thick plates of plastic scintillator interspersed within approximately 5 cm thick brass absorber plates to

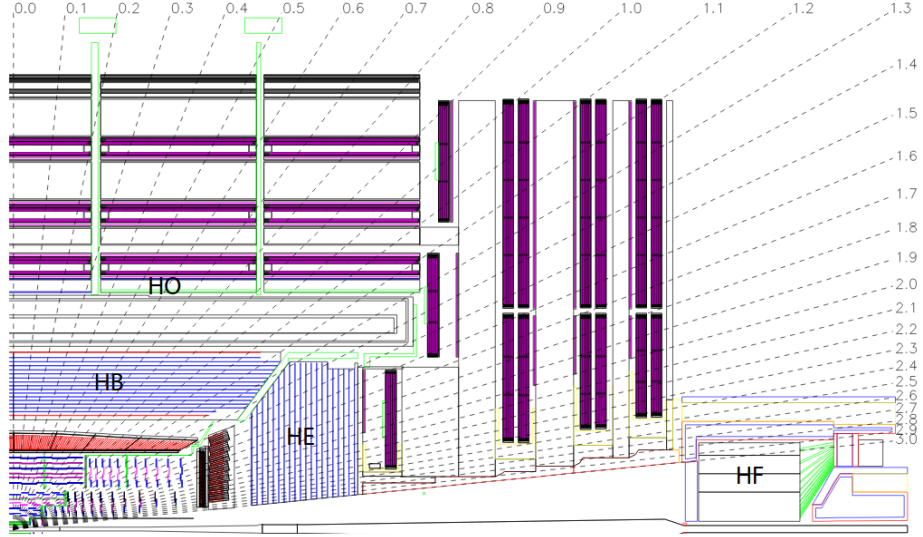


Figure 2.11: Layout of the hadron calorimeter barrel (HB), outer (HO), endcap (HE), and forward (HF) subdetectors [39].

reconstruct the energy deposited during hadronic showers. Embedded wavelength-shifting fibers capture the scintillation light and transfer it to clear fibers to be read out by hybrid photodiodes.

The barrel section ($|\eta| < 1.4$) is segmented into 32 towers in η and 64 in ϕ that each contain 17 active scintillator layers. In the $|\eta| < 1.26$ range, an extra layer of scintillator tiles (or two at $\eta = 0$) sits just outside the solenoid and increases the minimum effective HCAL interaction length to greater than 11.8. Each endcap spans a pseudorapidity range of 1.3 to 3.0 with 14 towers in η and 5 to 10° ϕ segmentation. Also, a steel and quartz fiber forward calorimeter (HF) sits 11.2 m from the interaction point and covers $3 < |\eta| < 5$. In HF, particles produce Cherenkov light when traversing the quartz fibers that run parallel to the beamline. Figure 2.11 shows the layout of the barrel, endcap, and forward HCAL subdetectors.

2.2.5 Muon system

The CMS muon system is composed of three varieties of gaseous detectors embedded in the iron return yoke outside the superconducting solenoid. In the central region ($|\eta| < 1.2$), the low muon and neutron rates along with the lower magnetic field allow the use of drift tube (DT) chambers. At higher η ($0.9 \leq |\eta| < 2.4$), cathode strip chambers (CSCs) are required to handle the higher rates and larger magnetic field. Finally, resistive plate chambers (RPCs), which provide more accurate time measurements and worse spatial resolution than the DTs and CSCs, complement the other detectors out to $|\eta| < 1.9$ [38, 44].

As shown in Fig. 2.12, the muon momentum resolution of the inner tracker is about an order of magnitude better than that of the muon system for low- p_T muons. The muon system is critical, however, for maintaining the $< 10\%$ momentum resolution that is necessary to unambiguously differentiate muons and anti-muons up to 1 TeV. In addition to improving muon reconstruction, the muon system provides information to the L1 trigger (see Section 2.2.6) and is capable of triggering on muons with good efficiency, high background rejection, and about 15–25% p_T resolution without input from the rest of the detector.

The DTs are organized into four stations, each of which contain up to 70 DT chambers that each measure the muon hit position in either the r - ϕ or z direction. Each chamber is composed of two or three collections of four-layer groupings of 13 by 42 mm drift cells. As diagrammed in Fig. 2.13, a 2–4 m anode wire runs down the center of each drift cell while electrode and cathode strips line the top, bottom, and walls of the cell. The cells are filled with an Ar/CO₂ gas mixture that is ionized

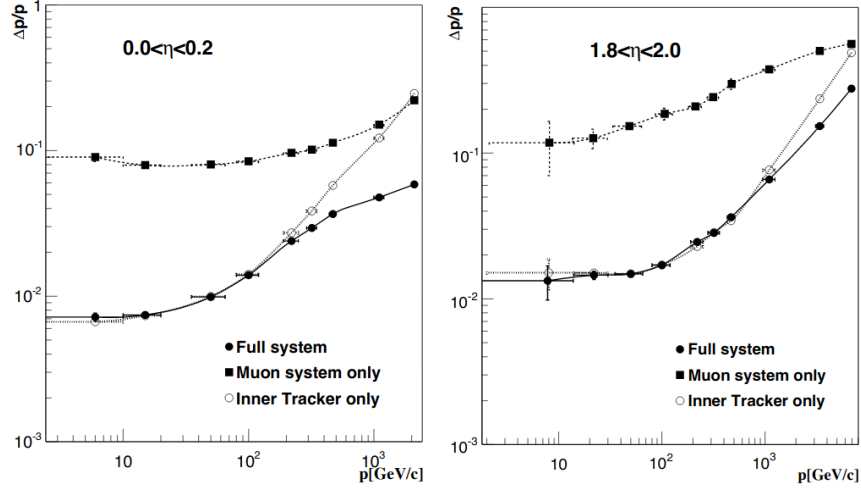


Figure 2.12: Muon momentum resolution as a function of momentum when using the CMS muon system, the CMS inner tracker, and the combination of the two subdetectors in two different η ranges [38].

by charged particles traversing the cell. The liberated electrons cause avalanches in the large electric fields before being collected by the anode wire.

Each endcap contains four CSC stations, each with six layers of CSCs whose cathode strips run radially outward to provide muon hit position measurements in the r - ϕ plane while the anode wires run in the azimuthal direction to provide measurements in η . As in the DTs, charged particles ionize a gas mixture inside each chamber (this time an Ar/CO₂/CF₄ mixture), which leads to an avalanche of electrons that are collected by an anode wire. In the CSCs, however, several anode wires share the same chamber and a two-coordinate position measurement is obtained by also reading out the induced charge on the perpendicular cathode strips.

The barrel and endcap regions also contain several RPC planes that ensure muons are associated with the correct bunch crossing and serve as a complementary muon

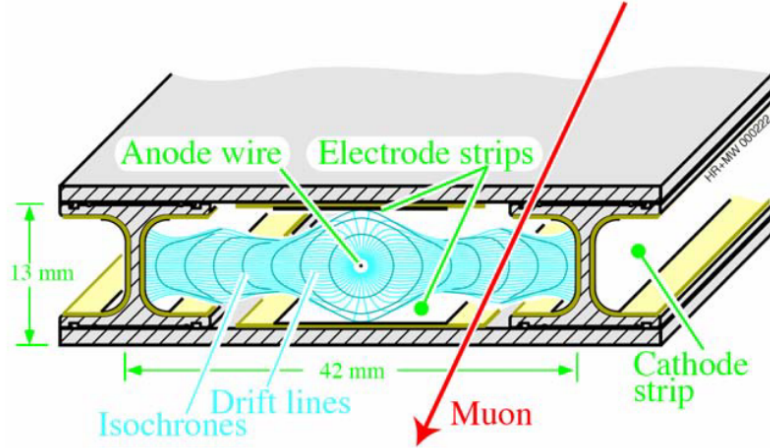


Figure 2.13: Sketch of a CMS muon system drift cell showing drift lines and isochrones. The plates at the top and bottom of the cell are at ground potential while the voltages applied to the electrodes are +3600 V for wires, +1800 V for strips, and −1200 V for cathodes [39].

trigger system. Each RPC plane is composed of parallel resistive plates separated by two 2 mm gas gaps. Charged particles ionize the gas, a voltage differential leads to an avalanche, and the total induced signal is read out as the sum of the induced signal in each gap. The double-gap approach allows for high-efficiency operation with lower voltage differentials than would be required for a single gap. At lower voltages, the RPCs operate in “avalanche mode”, which allows for stable performance and time resolution on the order of 2 ns despite particle fluxes up to hundreds of Hz cm^{-2} [39, 21, 45].

2.2.6 Trigger

The LHC provides bunch crossings to CMS at a nominal rate of 40 MHz. It is not possible to read out and store full event information at anywhere near this rate, so

the CMS trigger system is designed to identify the most interesting events to be read out at approximately 1 kHz. The rate reduction happens in two stages: the Level-1 (L1) trigger and High-Level Trigger (HLT). The L1 trigger uses custom electronics to analyze input from ECAL, HCAL, and the muon system to analyze each event in less than $3.8\,\mu\text{s}$ and select potentially interesting events at a rate of approximately 100 kHz. The HLT, which is implemented in software on a dedicated processor farm, uses information from all subdetectors to further reduce the rate to approximately 1 kHz [39, 46].

The L1 trigger processes calorimeter and muon system information in two separate trigger systems before combining the two outputs in a global trigger system that decides whether to pass a given event along to the HLT. The L1 algorithms are implemented on field programmable gate arrays housed on boards that communicate via optical link at $10\,\text{GBs}^{-1}$. The L1 Calorimeter Trigger uses coarsely segmented information from ECAL and HCAL to reconstruct electron, photon, tau lepton, and jet candidates, and the L1 Muon Trigger combines information from the DTs, CSCs, and RPCs to reconstruct muon candidates [47].

The HLT utilizes more than 400 different algorithms (called “paths”) that run on a farm of over 26 000 commercial processors. Each path runs a sequence of modules that either reconstruct or filter physics objects. The modules are ordered such that the fastest algorithms filter the object collections early, and the most resource-intensive algorithms only run when necessary. HLT paths generally mimic the selection and reconstruction algorithms used in CMS analyses, including the Particle Flow concept described in the following section, but are more aggressively optimized to satisfy the significant time and computing constraints imposed by the high event rate [46].

2.2.7 Physics object reconstruction

CMS uses a particle-flow (PF) algorithm to reconstruct the properties of individual particles from the combination of all subdetector measurements. Starting from charged particle tracks from the tracker and muon system and clusters of energy deposited in the ECAL and HCAL, CMS’s PF algorithm aims to reconstruct all final-state electrons, muons, photons, and charged and neutral hadrons in a given event. In this section, I first describe the reconstruction of tracks and energy clusters before moving on to the individual particle identification and reconstruction. A complete description of the CMS PF algorithm is available in Ref. [48].

Charged particle tracks

Charged particle tracks are reconstructed with an iterative procedure in which the highest-purity algorithms are run first, and the hits associated with each reconstructed track are masked before moving on to the next step. Despite the middling reconstruction efficiency of each individual step, this approach results in higher efficiency than could be achieved with any single tracking algorithm and without meaningfully increasing the rate of misreconstruction. This general principle applies to all charged particle tracks, but the tracks associated with candidate electrons and muons receive special consideration.

To better handle electron trajectories affected by radiative energy loss, CMS employs a special iterative tracking procedure that includes a Gaussian-sum filter (GSF) [49]. This approach improves the overall reconstruction efficiency, allows reconstruction of lower- p_T electrons, and helps identify electrons from photon conversions and distinguish electrons from charged hadrons.

Muon track reconstruction benefits from measurements in the tracker and the muon system. Candidate muon tracks are placed in one of three categories depending on which subdetectors are used in their reconstruction: standalone muon tracks only use muon system hits; tracker muon tracks only use tracker hits and the requirement of at least one consistent muon system hit; and global muon tracks are reconstructed from a global fit of tracker and muon system hits.

Calorimeter energy clusters

Energy deposits in the calorimeters are clustered separately in ECAL and HCAL with a Gaussian-mixture model that assumes the energy deposits arise from an arbitrary number of Gaussian energy deposits whose amplitude and location are allowed to vary while the width is determined by the calorimeter properties. The clusters are first seeded by cells with energy above some threshold and greater than the energy of the surrounding cells. Nearby clusters are then merged before being fed to the Gaussian-mixture algorithm. Finally, several corrections are applied to the cluster energies to ensure accurate responses to photons and hadrons.

Particle-flow reconstruction

The tracks and clusters are then identified with and used to reconstruct all individual particles in an event. The first step is to link tracks and clusters together into groups that correspond one or a few particles. Tracker tracks are extrapolated outwards and linked with the nearest ECAL and HCAL clusters that are within a set radius in the η - ϕ plane. In the case of candidate electron tracks, tracker tracks and ECAL deposits consistent with electron radiative losses are also linked with the candidate electron track. ECAL and HCAL clusters are similarly linked together by

proximity in the η - ϕ plane. Due to the high granularity of CMS subdetectors, the number of tracks and clusters in a linked group is largely independent of the total number of particles in an event.

Each group of linked tracks and clusters is then processed by the PF particle identification and reconstruction algorithm. As in track reconstruction, particle reconstruction is an iterative process in which the tracks and clusters are masked after being associated with a particle. Figure 2.14 diagrams the basic concept used to identify muons, electrons, photons, and charged and neutral hadrons. Each step of the PF algorithm is summarized below.

Muons are reconstructed first from isolated global muon candidates, then non-isolated global muon candidates, and finally tracker muon (standalone muon) candidate tracks that are particularly well measured and consistent with hits in the muon system (tracker). Muon momentum is taken from the tracker track when $p_T < 200$ GeV and from the combination of tracker and muon system hits that yields the best fit otherwise.

Electron and isolated photon reconstruction, which occur together after muon reconstruction, are necessarily interrelated by the high probability that an electron radiates a photon or a photon pair-produces electrons when interacting with tracker material. Electrons are identified from GSF tracks with a corresponding ECAL cluster while isolated photons are identified from isolated ECAL clusters. The total electron energy accounts for radiative losses that show up as ECAL clusters, and both electrons and isolated photons require a high ratio of ECAL cluster energy to nearby HCAL cluster energy.

Next, nonisolated photons and charged and neutral hadrons are reconstructed from the remaining tracks and clusters. Within the tracker acceptance ($|\eta| < 2.5$), ECAL (HCAL) clusters without associated tracks are identified as photons (neutral hadrons). At higher η , nearby ECAL and HCAL clusters are assumed to arise from the same hadron shower and ECAL clusters without nearby HCAL clusters are identified as photons. Discrepancies between track momenta and associated HCAL cluster energy are also used to identify neutral hadrons and muons. Finally, a post-processing step corrects for rare failure modes that can potentially produce inaccurately large $p_{\text{T}}^{\text{miss}}$ measurements.

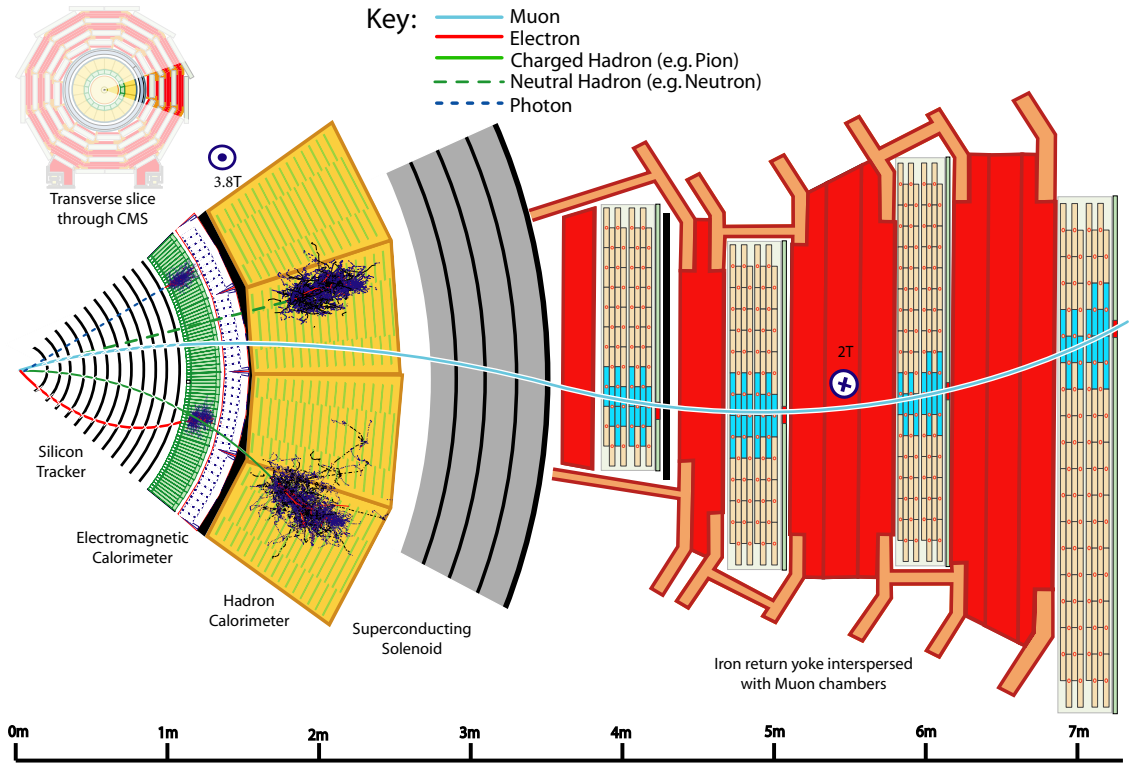


Figure 2.14: A sketch of a transverse slice of the CMS detector showing representative particle interactions used to identify and reconstruct particles with the CMS Particle Flow algorithm [48].

Chapter 3: Search for displaced leptons

In this chapter, I present a search for new long-lived particles (LLPs) that could be produced in proton-proton collisions at a center-of-mass energy of 13 TeV at the LHC and detected by the CMS experiment. The search targets the unique signature of “displaced leptons” that could be produced when LLPs decay to leptons after propagating a measurable distance from the location of the proton-proton collision. The candidate signal events include at least two leptons (one electron and one muon, two electrons, or two muons) whose transverse impact parameters are between 0.01 and 10 cm. Choosing transverse impact parameter as the main discriminating variable allows us to target pairs of displaced leptons without requiring that they form a common vertex. We apply an otherwise minimal event selection to retain sensitivity to a wide range of new physics models.

The Ohio State CMS group has previously performed two related searches for displaced leptons in the electron-muon final state: one in 19.7 fb^{-1} of 8 TeV data and another in 2.6 fb^{-1} of 13 TeV data [50, 51]. The search presented here is the most sensitive to date. Some of the most significant improvements include:

- adding sensitivity to the electron-electron and muon-muon final states
- simplifying the event selection to reduce model dependence

- introducing a custom lepton isolation definition to significantly reduce the background from heavy-flavor meson decays
- implementing a new, fully data-driven background estimation procedure
- expanding the signal region to include leptons with smaller transverse impact parameters
- adding a second signal interpretation
- analyzing more than a factor of 40 times more data than the previous 13 TeV analysis [51]

The ATLAS collaboration has also recently submitted a paper in which they target displaced leptons with transverse impact parameters between 0.3 and 30 cm. They observe no significant excess and set upper limits on the production cross section of long-lived sleptons in a gravity-mediated SUSY breaking model [52].

The remainder of this chapter is organized as follows: Section 3.1 introduces the displaced leptons experimental signature in the context of LLPs at the LHC and gives an overview of the analysis strategy, Section 3.2 defines the data and simulated SM and signal samples used in the analysis, Section 3.3 describes the event selection criteria and defines the signal and control regions, Section 3.4 describes the various corrections applied to the SM and signal simulation, Section 3.5 investigates the sources of background and defines the procedure for estimating their contribution to the signal region, Section 3.6 describes the sources of systematic uncertainty in the efficiency for simulated signal events to pass the signal region selection, and Section 3.7 presents the results.

3.1 Overview

3.1.1 Long-lived particles at the LHC

As discussed in Section 1.2.2, LLPs are common in the SM and naturally arise in many BSM scenarios. When produced in proton-proton collisions at the LHC, new LLPs have the potential to produce striking experimental signatures that differentiate them from SM backgrounds. As diagrammed in Fig. 3.1, new LLPs that decay within the detector volume can produce displaced vertices, physics objects whose trajectories do not point back to the location of the proton-proton collision, or particle tracks that vanish before reaching the outer edge of the tracker. Neutral LLPs that decay farther out can pass undetected through the tracker before producing signals in the calorimeters or muon system, and heavy, charged LLPs that are stable on detector length scales can be identified by their unusually large charge depositions.

The identification of any of these LLP signatures above the expected background rates would be a clear sign of new physics, but their atypical nature adds an inherent layer of difficulty to their study. In fact, standard reconstruction algorithms and event selections frequently discard such unusual signatures and render the majority of LHC analyses insensitive to many BSM scenarios that include LLPs [54]. As a solution to the naturalness problem may require new physics at the LHC energy scale, it is critical that physicists face the challenges posed by LLP analyses and look everywhere BSM physics may be hiding.

Interest in LLPs is growing as evidence of BSM physics continues to evade LHC physicists. Recent CMS searches in 13 TeV proton-proton collisions target disappearing tracks [55] or use the ECAL timing capabilities to target photons or jets whose

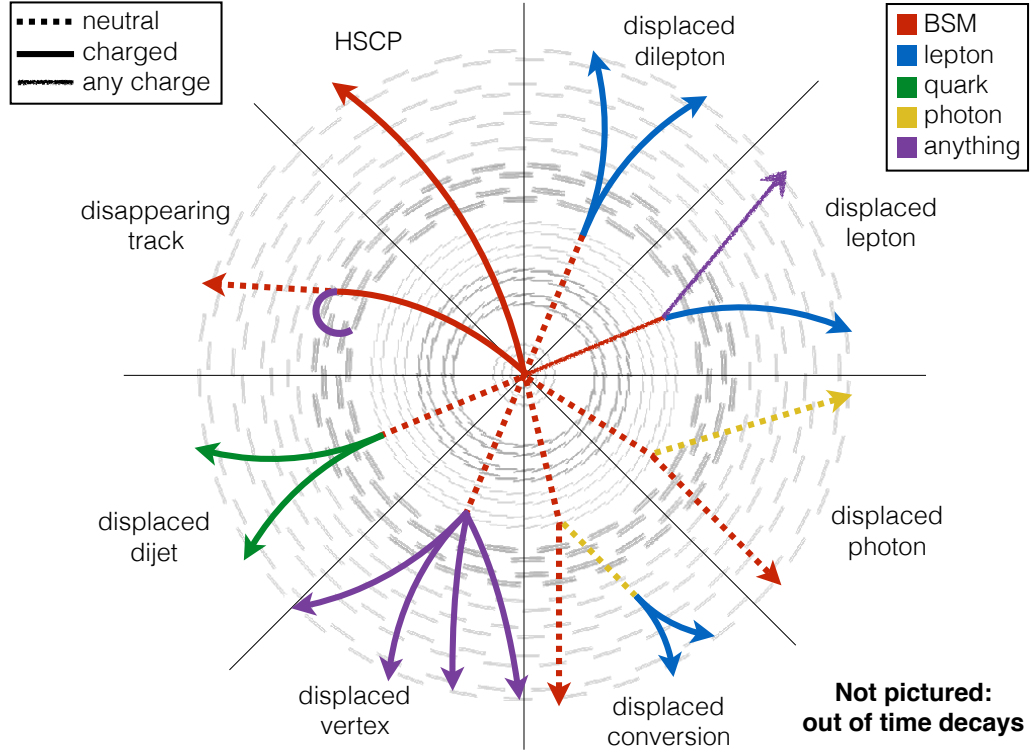


Figure 3.1: Illustration of several possible experimental signatures of long-lived particles [53].

production at a displaced vertex delays the time of their detection [56, 57]. All observations thus far agree with SM predictions, but these and other LLP analyses are probing regions of BSM parameter space that are untested by conventional analyses.

3.1.2 Displaced leptons signature

The Displaced Leptons analysis targets electrons and muons produced in the decays of new LLPs. Inspired by models such as Displaced SUSY (see Section 1.2.3), we take care to maintain sensitivity to leptons that are produced in separate long-lived decays as well as those that share a common displaced vertex. This goal is achieved by selecting pairs of leptons with large transverse impact parameters and setting no constraints on the presence or absence of displaced vertices. Figure 3.2 shows the benefit of such an approach: the pair-produced new LLPs, labeled X , each decay to a single lepton and an unspecified second particle. The lepton transverse impact parameter, which is labeled d_0 in the figure and defined explicitly below, is then used to identify the displaced nature of the lepton decays without explicitly constraining the other decay products in any way. Note also that the same strategy would successfully identify two leptons from a single long-lived particle decay.

Lepton transverse impact parameter, d_0 , is defined as the distance of closest approach in the transverse plane of the helical trajectory of the lepton track to the CMS beamspot, which is the center of the region in which the proton bunches cross. d_0 is commonly measured with respect to the primary vertex, but in the case of leptons produced in displaced decays, the association between a given primary vertex and the resulting leptons is unreliable. We determine d_0 from measured properties of the

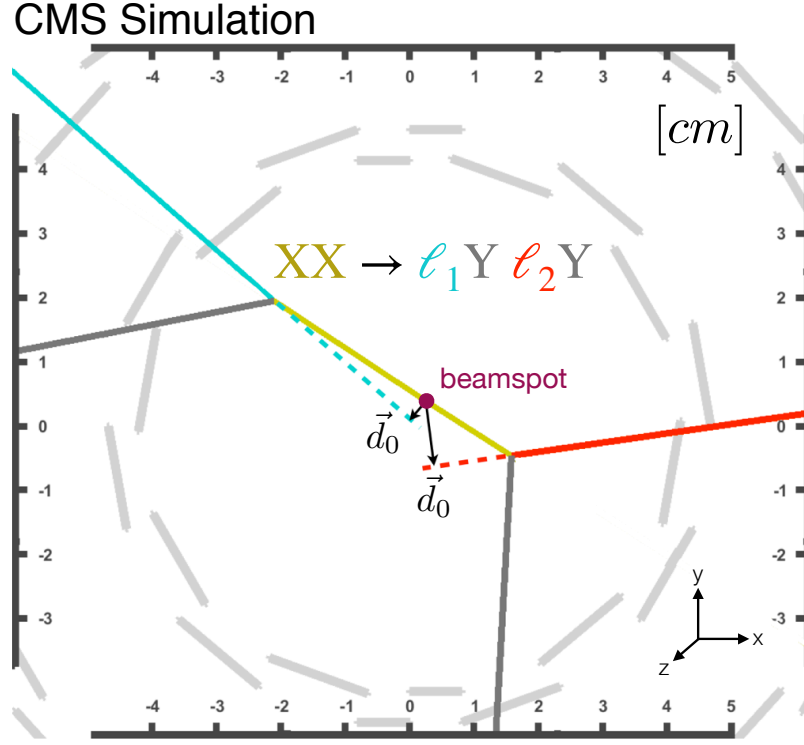


Figure 3.2: Illustration of the displaced leptons signature showing the definition of d_0 in a transverse view of the CMS detector. X denotes a new long-lived particle, ℓ denotes an electron or muon, and Y denotes any other decay products of the new long-lived particle. When interpreting the results of the Displaced Leptons analysis with the Displaced SUSY model, X refers to a top squark and Y refers to a b or d quark.

lepton track using the following equation:

$$d_0 = \frac{(v_x - x_0)p_y - (v_y - y_0)p_x}{p_T} \quad (3.1)$$

where v_x and v_y refer to the x and y coordinates of the lepton track reference point, which is usually chosen to be the point of closest approach to the center of CMS, x_0 and y_0 refer to the x and y coordinates of the beamspot, and p_x , p_y , and p_T refer to the magnitudes of the x , y , and transverse components of the lepton's momentum. $|d_0|$ is commonly used throughout the Displaced Leptons analysis because we generally care about the magnitude of d_0 but not its sign.

If we are to use lepton $|d_0|$ as the discriminating variable in an LLP search, we must ensure it scales appropriately with the parent LLP lifetime. Typically, a particle with lifetime τ will travel a distance $d = \beta\gamma c\tau$, where β and γ are relativistic factors and c is the speed of light, so d and τ are directly correlated. Figure 3.2 shows that lepton $|d_0|$ is determined by the distance travelled by the new LLP in the transverse plane and the angle between the transverse momenta of the new LLP and the lepton. Unless this angle is pathologically constrained to zero, $|d_0|$ and τ will be directly correlated as well. In fact, the maximum value of lepton $|d_0|$, which occurs when the angle between the transverse momenta of the new LLP and the lepton is $\frac{\pi}{2}$, is exactly equal to the transverse distance between the beamspot and location of the LLP decay.

Figure 3.3 shows the distribution of data and simulated Displaced SUSY events in the plane defined by electron and muon $|d_0|$. As expected, the data events, which are dominated by leptons from promptly decaying parent particles, are concentrated at low $|d_0|$ values while the Displaced SUSY events are spread across the entire plane. This figure shows the power of $|d_0|$ as a discriminating variable: requiring two leptons

with $|d_0| > 100\,\mu\text{m}$ eliminates nearly all the SM background without requiring that the leptons form a common vertex.

Another possible discriminating variable could be $|d_0|/\sigma_{|d_0|}$, where $\sigma_{|d_0|}$ is the uncertainty in $|d_0|$. Such a discriminating variable could potentially reduce the background from leptons with poorly measured $|d_0|$, but we choose to use $|d_0|$ because of its straightforward correspondence to the parent particle lifetime. We also find that $\sigma_{|d_0|}$ is often underestimated, which reduces the potential benefit of using $|d_0|/\sigma_{|d_0|}$.

3.1.3 Analysis strategy

Having seen that $|d_0|$ can be used to identify leptons from long-lived particle decays without requiring that the leptons form a common vertex, we now define a strategy to target such a signature. The basic analysis strategy is outlined here and described in detail in the following sections.

In addition to maximizing our sensitivity to models such as Displaced SUSY, we also strive to develop an analysis that is model independent, signature based, and easy to reinterpret. With these goals in mind, we perform a relatively simple cut-and-count analysis in which our main event selection sets no constraints on any non-lepton physics object. Unlike previous displaced leptons analyses [50, 51], we allow final states with more than two leptons and set no constraints on the lepton charge. As explained in Section 3.3, we select events in three analysis channels: electron-electron (ee), electron-muon ($e\mu$), and muon-muon ($\mu\mu$). We then divide the events into different regions of the plane defined by the $|d_0|$ of the two leptons that define the analysis channel. Figure 3.3, for example, shows this plane in the

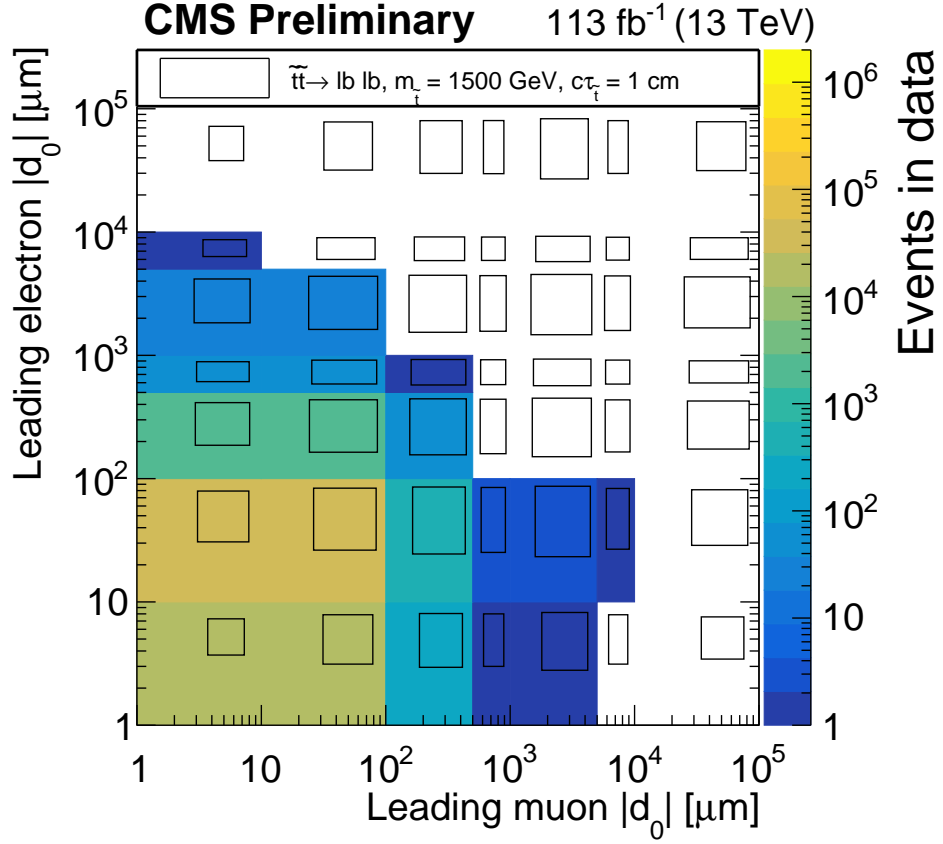


Figure 3.3: Distribution of data (colors) and simulated Displaced SUSY events (black boxes) in the plane defined by electron $|d_0|$ and muon $|d_0|$. The size of the black boxes are proportional to the bin content, and the bins along the x and y axes include underflow. All events are required to pass the $e\mu$ preselection defined in Section 3.3.

$e\mu$ channel. The signal region is defined as the region in which both leptons have $|d_0| > 100\,\mu\text{m}$.

Following the procedure defined in Section 3.5, we use the data in the non-signal regions to estimate the SM background in the signal region. Finally, we compare the background estimates and data yields in the signal region. In the absence of a significant excess, we use simulated signal events to constrain the available parameter space of the Displaced SUSY model. The results will also be used to constrain the production cross section of long-lived sleptons as well as the Higgs boson branching fraction to pairs of long-lived scalars, but these additional interpretations will unfortunately not be available in time to appear in this thesis. To avoid biasing the result, we blind ourselves to the data in the signal region and wait to observe the signal-region data until after we define and test the entire analysis procedure and receive pre-approval from the CMS Exotica Physics Analysis Group.

3.2 Data and simulated samples

3.2.1 Experimental data

This analysis uses proton-proton collision data taken in 2016, 2017, and 2018 at a center-of-mass energy of 13 TeV. In 2016, we use only the last two run periods due to lower displaced tracking efficiency caused by an analog pipeline voltage saturation problem in the silicon strip detector during the earlier run periods (see Appendix A). In 2017, we use all run periods in the ee channel and all but the earliest run period in the $e\mu$ and $\mu\mu$ channels because the $e\mu$ and $\mu\mu$ triggers are not available in the earliest run period. In 2018, we use all available run periods in all three channels. Ultimately, this analysis uses an integrated luminosity of $16.1 \pm 0.4 \text{ fb}^{-1}$ from 2016 in all three channels, $41.5 \pm 1.0 \text{ fb}^{-1}$ ($36.7 \pm 0.8 \text{ fb}^{-1}$) from 2017 in the ee channel ($e\mu$ and $\mu\mu$ channels), and $59.7 \pm 1.5 \text{ fb}^{-1}$ from 2018 in all three channels.

The search is performed in the `MuonEG`, `DoubleEG` (in 2016–2017), `EGamma` (in 2018), and `DoubleMu` primary datasets. We also use the `MET` dataset to study the trigger efficiency and the `Cosmics` and `NoBPTX` datasets to study the displaced tracking efficiency. All data are reconstructed in the `07Aug17`, `31Mar2018`, `17Sep2018` reprocessing campaigns with software versions `CMSSW_8_0_31`, `CMSSW_9_4_8`, and `CMSSW_10_2_0`, respectively. The sole exception is data-taking era D of the 2018 `EGamma` dataset, which is reconstructed in the `22Jan2019` campaign. In all cases, we use the CMS `MiniAOD` event format.

3.2.2 Simulated background events

This analysis employs a fully data-driven background estimation technique that does not rely on simulated SM events. We do, however, use simulated SM events

to study possible sources of background and verify the validity of the background estimation technique. The samples corresponding to 2016, 2017, and 2018 data conditions are from the `PdmVMCcampaignRunIISummer16`, `PdmVMCcampaignRunIIFall17`, and `PdmVMCcampaignRunIIAutumn18` production campaigns and were reconstructed in `CMSSW_8_0_31`, `CMSSW_9_4_8`, `CMSSW_10_2_0` with the MiniAODSIM event format. The samples simulating Z +jets, W +jets, and $t\bar{t}$ production are generated using MADGRAPH5_aMC@NLO [58, 59, 60], while the samples simulating diboson (WW , WZ , and ZZ with leptonic and semi-leptonic decays) and single-top-quark production are simulated with POWHEG v2 [61, 62, 63, 64, 65]. PYTHIA 8.2 [66] is used to simulate the parton showering and hadronization for all processes. The modeling of the underlying event is generated using the CUETP8M1 [67] and CP5 tunes [68] for simulated samples corresponding to the 2016 and 2017–18 data sets, respectively.

3.2.3 Simulated signal events

We use simulated signal events to guide the analysis strategy and interpret our results. Samples of simulated $pp \rightarrow t\bar{t}$ events in which the top squarks decay to a lepton and a b quark or d quark are produced at leading order using PYTHIA 8.2. For simplicity, all lepton flavors are generated with equal branching fractions. The top squarks can form strongly-produced hadronic states called R-hadrons, which are generated with PYTHIA but whose interactions with matter are not modeled with GEANT4. Such interactions are not expected to have a significant impact on the analysis because the R-hadrons do not traverse a significant number of interaction lengths before decaying. As described in Appendix D, we nevertheless study the impact of the R-hadron interactions using the “cloud model,” which assumes that the

top squark is surrounded by a cloud of colored, light constituents that interact during scattering [69, 70], and find the effect on the signal efficiency to be insignificant. To generate the samples, we start with a SUSY Les Houches Accord file [71] corresponding to Snowmass Points and Slopes point 1a [72] and modify the mass and width of the top squark according to the sample being produced. We generate samples with top squark masses from 100 to 1800 GeV at 100 GeV intervals and with top squark lifetimes at each decade from 0.1 mm to 1 m. After producing these samples, we also employ a lifetime reweighting technique to effectively produce eight additional lifetime points between each pair of adjacent lifetimes. In the case of the 1 m samples, we also use an equivalent technique to effectively produce nine additional lifetime points between 1 and 10 m. The production cross sections for each top squark mass hypothesis are taken from the website of the LHC SUSY Cross Section Working Group and reproduced in Table 3.1. The signal samples are reconstructed in the same campaigns and with the same conditions as the SM background samples described in 3.2.2.

Table 3.1: The top squark production cross sections used when generating the simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}\bar{b}$ and $\tilde{t}\tilde{t} \rightarrow \bar{l}d\bar{l}\bar{d}$ events. Cross sections are calculated at NLO + NLL by the LHC SUSY Cross Section Working Group.

\tilde{t} mass [GeV]	Cross section [pb]	Cross section uncertainty
100	1.52×10^3	15 %
200	6.45×10^1	14 %
300	8.52×10^0	14 %
400	1.84×10^0	14 %
500	5.18×10^{-1}	13 %
600	1.75×10^{-1}	13 %
700	6.70×10^{-2}	13 %
800	2.83×10^{-2}	14 %
900	1.29×10^{-2}	15 %
1000	6.15×10^{-3}	16 %
1100	3.07×10^{-3}	17 %
1200	1.60×10^{-3}	19 %
1300	8.50×10^{-4}	20 %
1400	4.61×10^{-4}	22 %
1500	2.56×10^{-4}	24 %
1600	1.41×10^{-4}	27 %
1700	8.08×10^{-5}	29 %
1800	4.67×10^{-5}	31 %

3.3 Event selection

3.3.1 Triggers

The events are required to pass different triggers in each channel. Standard CMS electron and muon triggers are not designed for displaced objects, so we use non-standard triggers for both electrons and muons. For muons, we remove all trigger requirements relating to the muon d_0 , longitudinal impact parameter (d_z), or the vertex from which the muon originates. For electrons, we actually use photon triggers, which collect events with electrons as well as photons but do not rely on any tracking information. See Section 2.2.6 for a brief overview of the CMS trigger system.

In the $e\mu$ channel, 2016 data and corresponding simulated events are required to pass the logical OR of two HLT paths (`HLT_Mu38NoFiltersNoVtx_Photon38_CaloIdL_v*` OR `HLT_Mu28NoFiltersNoVtxDisplaced_Photon28_CaloIdL_v*`) that were both originally designed for the 2015 CMS displaced leptons analysis [51]. The first trigger requires at least one muon with $p_T > 38$ GeV and places no constraints on the vertex, d_0 , or d_z . The second trigger requires at least one muon with $p_T > 28$ GeV and $|d_0| > 0.01$ cm. Each of these two triggers also requires at least one photon that passes a loose calorimeter-based identification. The first (second) trigger requires that the photon E_T is greater than 38 GeV (28 GeV). The signal efficiency with these dedicated triggers is significantly higher than that of standard muon-photon HLT paths.

In 2017 and 2018, data and corresponding simulated events in the $e\mu$ channel are required to pass `HLT_Mu43NoFiltersNoVtx_Photon43_CaloIdL_v*`. The muon p_T and photon E_T thresholds are raised with respect to 2016 due to increased pileup. A

version of the 2016 trigger that requires displaced muons is not available in 2017 and 2018.

In the ee channel, 2016 data and corresponding simulated events are required to pass the logical OR of two HLT paths (`HLT_Diphoton30_18_R9Id_OR_IsoCaloId_AND_HE_R9Id_Mass90_v*` OR `HLT_DoublePhoton60_v*`). The first requires a leading photon with $E_T > 30$ GeV and a subleading photon with $E_T > 18$ GeV. Photons must pass calorimeter identification criteria involving isolation, the ratio of HCAL to ECAL energy, and shower shape, and the di-photon invariant mass must be > 90 GeV. This path is highly efficient at low top squark mass. The second trigger simply requires at least two photons with $E_T > 60$ GeV. This path is highly efficient at large top squark mass and lifetime.

In 2017 and 2018, data and corresponding simulated events in the ee channel are required to pass `HLT_Diphoton30_22_R9Id_OR_IsoCaloId_AND_HE_R9Id_Mass90_v*` OR `HLT_DoublePhoton70_v*`. The photon E_T thresholds are raised with respect to 2016 due to increased pileup.

In the $\mu\mu$ channel, 2016 data and corresponding simulated events are required to pass the logical OR of two HLT paths (`HLT_DoubleMu33NoFiltersNoVtx_v*` OR `HLT_DoubleMu23NoFiltersNoVtxDisplaced_v*`) that were both designed for this analysis. The first trigger requires at least two muons with $p_T > 33$ GeV and sets no constraints on the vertex, d_0 , or d_z . The second trigger requires at least two muons with $p_T > 23$ GeV and $|d_0| > 0.01$ cm. The signal efficiency with these dedicated triggers is significantly higher than that of standard di-muon HLT paths.

In 2017 and 2018, data and corresponding simulated events in the $\mu\mu$ channel are required to pass `HLT_DoubleMu43NoFiltersNoVtx_v*`. The muon p_T threshold is

raised with respect to 2016 due to increased pileup. A version of the 2016 trigger that requires displaced muons is not available in 2017 and 2018.

3.3.2 Preselection

Starting from the events collected with the triggers described above, we next apply a set of preselection criteria to select the events to be analyzed. The preselection criteria vary by channel and year, but the fundamental goal is always to select events with at least one good reconstructed lepton of each flavor required by the channel.

Specifically, the $e\mu$ preselection selects events with at least one PF electron and at least one global PF muon, the ee preselection selects events with at least two PF electrons, and the $\mu\mu$ preselection selects events with at least two global PF muons (see Section 2.2.7 for a discussion of the PF algorithm). We set requirements on these electrons and muons as shown in Tables 3.2, 3.3, and 3.4. The electron and muon p_T requirements are chosen such that the trigger efficiency is independent of lepton p_T (see Section 3.4.4), while electron and muon $|\eta|$ requirements are chosen to reduce the fraction of leptons with poorly measured d_0 (see Appendix B). Electrons that traverse the gap between the endcap and barrel detectors are also rejected due to the known decrease in reconstruction performance in this region.

We use a tight cut-based identification (ID) on the electrons and muons to select well-reconstructed leptons, but unlike the standard ID definitions used in many CMS analyses, we do not place any requirements on d_0 or d_z . In all other respects, we follow the cut-based ID recommendations of the CMS EGamma and Muon Physics Object Groups. The electron ID corresponds to `egmGsfElectronIDs:cutBasedElectronID-Summer16-80X-V1-tight` in 2016, `egmGsfElectronIDs:cutBasedElectron`

Table 3.2: The $e\mu$ preselection criteria. The electron and muon p_T thresholds increase in 2017 in accordance with the increased HLT electron and muon p_T thresholds.

Selection variable		Object-level selections	
Number	Electron	Muon	
	≥ 1	≥ 1	
p_T	$> 42 \text{ GeV}$ (2016)	$> 40 \text{ GeV}$ (2016)	
	$> 45 \text{ GeV}$ (2017 and 2018)	$> 45 \text{ GeV}$ (2017 and 2018)	
$ \eta $	< 1.5	< 1.5	
	not in ECAL gap	-	
$\eta - \phi$	veto ($1.0 < \eta < 1.5$ and $\phi > 2.7$) (2017)	veto ($1.0 < \eta < 1.5$ and $\phi > 2.7$) (2017)	
	veto ($0.3 < \eta < 1.2$ and $0.4 < \phi < 0.8$) (2018)	veto ($0.3 < \eta < 1.2$ and $0.4 < \phi < 0.8$) (2018)	
ID	Tight (cut-based)	Tight (cut-based)	
Custom isolation	Tight	Tight	
Event-level selections			
Zero $\mu\mu$ pairs with $\cos\alpha < -0.99$			
Reject $\Delta t < -20$, if both timing ndof > 7			
At least one $e\mu$ pair with $\Delta R(e, \mu) > 0.2$			
Reject events with candidate leptons form a displaced vertex in the tracker material			

Table 3.3: The ee preselection criteria. The electron p_T threshold increases in 2017 and 2018 in accordance with the increased HLT electron p_T threshold.

Selection variable	Object-level selections
Number	Electron ≥ 2
p_T	$> 65 \text{ GeV}$ (2016) $> 75 \text{ GeV}$ (2017 and 2018)
$ \eta $	< 1.5 not in ECAL gap
$\eta - \phi$ (pixel power supply issue)	veto ($1.0 < \eta < 1.5$ and $\phi > 2.7$) (2017) veto ($0.3 < \eta < 1.2$ and $0.4 < \phi < 0.8$) (2018)
ID	Tight (cut-based)
Custom isolation	Tight
Event-level selections	
At least one ee pair with $\Delta R(e, e) > 0.2$	
Reject events with candidate leptons form a displaced vertex in the tracker material	
Reject events with displaced muons in the $e\mu$ channel inclusive signal region	

Table 3.4: The $\mu\mu$ preselection criteria. The muon p_T threshold increases in 2017 and 2018 in accordance with the increased HLT muon p_T threshold.

Selection variable	Object-level selections
Number	Muon ≥ 2
p_T	$> 35 \text{ GeV}$ (2016) $> 45 \text{ GeV}$ (2017 and 2018)
$ \eta $	< 1.5
$\eta - \phi$ (pixel power supply issue)	veto ($1.0 < \eta < 1.5$ and $\phi > 2.7$) (2017) veto ($0.3 < \eta < 1.2$ and $0.4 < \phi < 0.8$) (2018)
ID	Tight (cut-based)
Custom isolation	Tight
Event-level selections	
Zero $\mu\mu$ pairs with $\cos \alpha < -0.99$	
Reject $\Delta t < -20$, if both timing ndof > 7	
At least one $\mu\mu$ pair with $\Delta R(\mu, \mu) > 0.2$	
Reject events with candidate leptons form a displaced vertex in the tracker material	
Reject events with displaced electrons in the $e\mu$ channel inclusive signal region	

Table 3.5: The electron tight ID requirements, which are identical to the tight cut-based ID from the CMS EGamma Physics Object Group with the d_0 and d_z requirements removed. Electron ID quantity definitions are available in Ref. [73].

Electron ID requirements	
full5x5 $\sigma_{i\eta i\eta} <$	0.0104 (2018, 2017)
	0.00998 (2016)
	0.00255 (2018)
$ \delta\eta_{\text{Seed}} <$	0.00353 (2017)
	0.00308 (2016)
	0.022 (2018)
$ \delta\phi_{\text{In}} <$	0.0499 (2017)
	0.0816 (2016)
	0.026 + 1.15/E + 0.0324 ρ /E (2018)
H/E <	0.026 + 1.12/E + 0.0368 ρ /E (2017)
	0.0414 (2016)
	0.0287 + 0.506/ p_{T} (2018)
PF isolation <	0.0361 (2017)
	0.0588 (2016)
	0.159 (2018)
$ 1/E - 1/p <$	0.0278 (2017)
	0.0129 (2016)
expected missing inner hits <=	1
pass conversion veto	yes

ID-Fall17-94X-V1-tight in 2017, and `egmGsfElectronIDs:cutBasedElectronID-Fall17-94X-V2-tight` in 2018. The electron and muon tight ID requirements are summarized in Tables 3.5 and 3.6.

We also require that electrons and muons are isolated. Specifically, we use a modified isolation definition that accounts for the fact that displaced leptons may be associated with the wrong primary vertex. The standard PF isolation assumes all energy from primary vertices other than the leading primary vertex is due to

Table 3.6: The muon tight ID requirements, which are identical to the tight cut-based ID from the CMS Muon Physics Object Group with the requirements on d_0 and d_z removed.

Muon ID requirements
Is a global muon
Is a PF muon
$\chi^2/\text{n}_{\text{dof}}$ of the global-muon track fit is < 10
At least one muon-chamber hit included in the global-muon track fit
Muon segments in at least two muon stations
At least 1 valid pixel hit
At least 6 tracker layers with hits

pileup, which is not true when the primary vertex ordering is altered by an incorrectly associated lepton. We have therefore modified the isolation definition to be agnostic to the primary vertex ordering by allowing PF candidates from any primary vertex to contribute to the isolation sum and by using a simple ρ -based pileup correction, where ρ is the total transverse energy of all the PF candidates in an event divided by the total detector area. The modified isolation is calculated as:

$$\text{relative isolation} = \frac{\max\left(0, p_{\text{T}}^{h^{\pm}} + E_{\text{T}}^{h^0} + E_{\text{T}}^{\gamma} - \rho\pi R^2\right)}{p_{\text{T}}^{\ell}} \quad (3.2)$$

where R is the radius of a cone in the η - ϕ plane that is centered on the lepton, $p_{\text{T}}^{h^{\pm}}$ is the total p_{T} of charged hadrons in the cone, $E_{\text{T}}^{h^0}$ is the total E_{T} of neutral hadrons in the cone, E_{T}^{γ} is the total contribution of photons in the cone, ρ is defined as above, and p_{T}^{ℓ} is the lepton p_{T} . R is set to 0.3 for electrons and 0.4 for muons.

Figure 3.4 shows how the size of the pileup correction term depends on lepton displacement in the standard isolation but not in the modified isolation described here. We use the modified isolation definition for both electrons and muons while keeping the original tight working point for electrons and slightly tightening the tight working

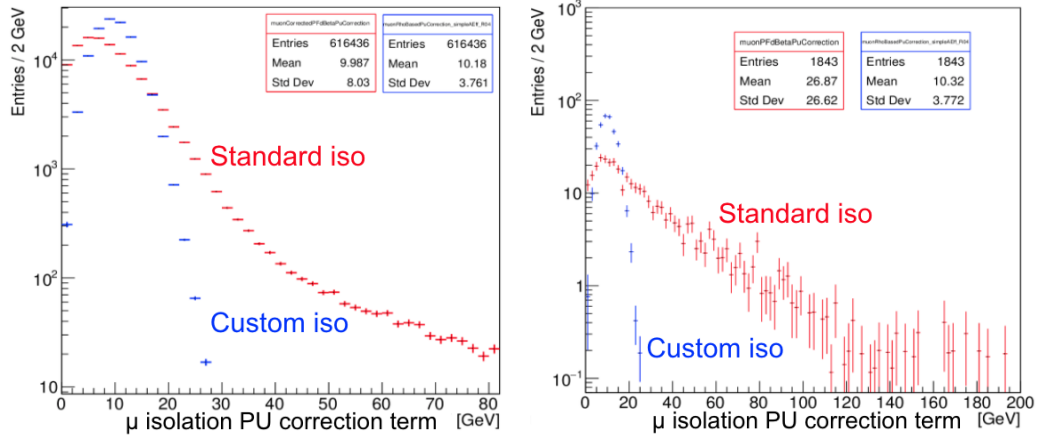


Figure 3.4: Comparison of the the muon isolation pileup correction term in the standard muon isolation and the modified muon isolation in simulated $t\bar{t}$ events that pass the 2018 $e\mu$ preselection. Muon $|d_0|$ is constrained to be less than 100 μm in the plot on the left and between 500 and 1000 μm in the plot on the right.

point for muons. In the end, we require that the relative isolation is < 0.10 for muons and < 0.0588 for electrons in 2016 and $< 0.0287 + 0.506/p_T$ for electrons in 2017 and 2018. As shown in Figs. 3.5 and 3.6, this modified PF isolation rejects substantially more background when the leptons are displaced but does not significantly alter the signal yield. We note, however, that there may still be some minor dependence on the primary vertex selection in the PF muon requirement because the PF muon selection includes some loose isolation requirements where the charged hadron component is constrained to the selected primary vertex.

We also reject electrons and muons in certain regions of the η - ϕ plane where lepton d_0 is more likely to be mismeasured. We identify these regions as highly populated bins in the electron η - ϕ distribution in a prompt-muon, displaced-electron control region in 2017 and 2018 data (see Fig.3.7). No such bins are present in 2016 data. The identified regions coincide with regions found by a previous CMS analysis [55] to

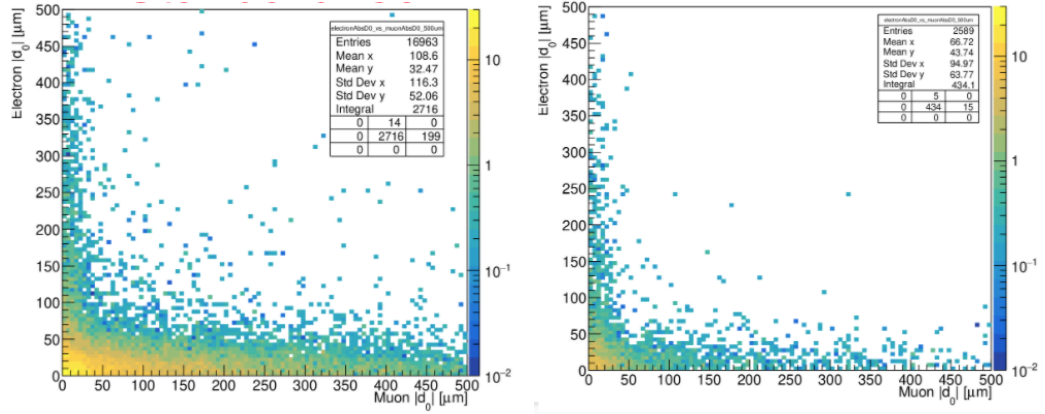


Figure 3.5: The distribution of $t\bar{t}$ simulated events in the plane defined by electron $|d_0|$ and muon $|d_0|$. The standard isolation is applied in the plot on the left, and the modified isolation is applied in the plot on the right, and the events in both plots are required to pass the remaining 2018 $e\mu$ preselection criteria and the additional constraint that the parent of at least one lepton is a heavy-flavor meson.

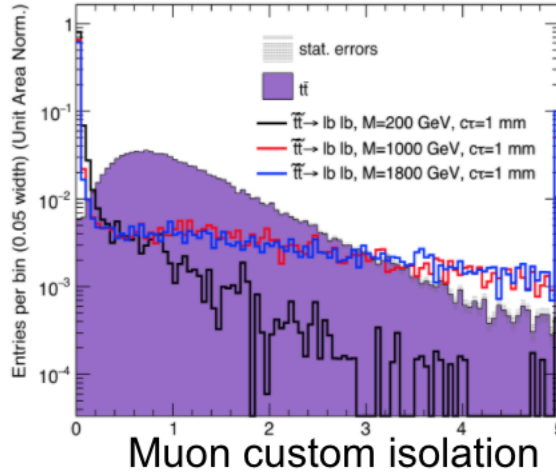


Figure 3.6: The muon modified isolation distribution for simulated $t\bar{t}$ background and $t\bar{t} \rightarrow \bar{l}b \, l\bar{b}$ signal events that pass the 2018 $e\mu$ preselection with no isolation criterion applied.

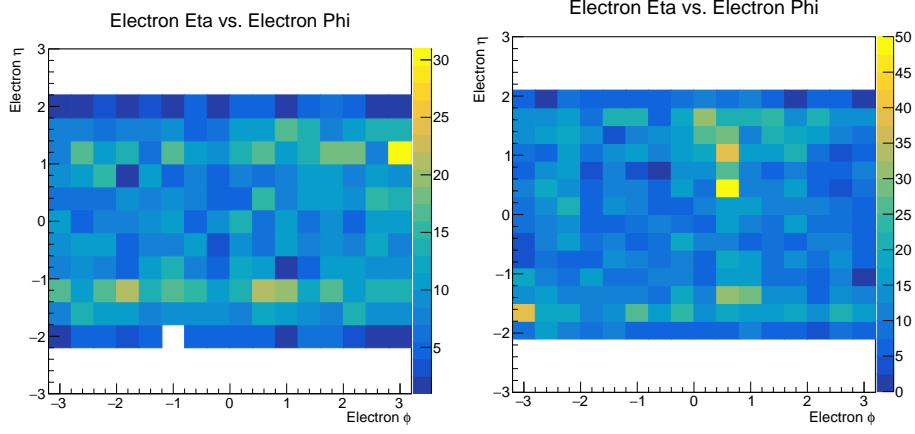


Figure 3.7: The electron η - ϕ distribution in a prompt-muon, displaced-electron control region in 2017 (left) and 2018 (right) data before vetoing the regions affected by pixel power-supply issues.

be affected by power supply issues in the pixel detector. The η - ϕ variation is more apparent for displaced electrons than displaced muons, so we use data in a prompt muon ($|d_0| < 40 \mu\text{m}$), displaced electron ($100 < |d_0| < 500 \mu\text{m}$) control region to define the regions used for both electrons and muons. In 2017, the rejected region is $1.0 < \eta < 1.5$ AND $\phi > 2.7$, and in 2018 the rejected region is $0.3 < \eta < 1.2$ AND $0.4 < \phi < 0.8$.

In addition to these object-level selections, we also impose a few event-level selections designed to remove potential backgrounds from cosmic rays, material interactions, and displaced decays of SM hadrons. To remove cosmic-ray muons in the $\mu\mu$ and $e\mu$ channels, we require there be zero pairs of muons with $\cos \alpha < -0.99$, where α is the 3D angle between the muons, and that the relative time between the leading two muons is inconsistent with the timing of cosmic-ray muons. To do this, we look at the muon time measured by the DTs and CSCs, which assume that the

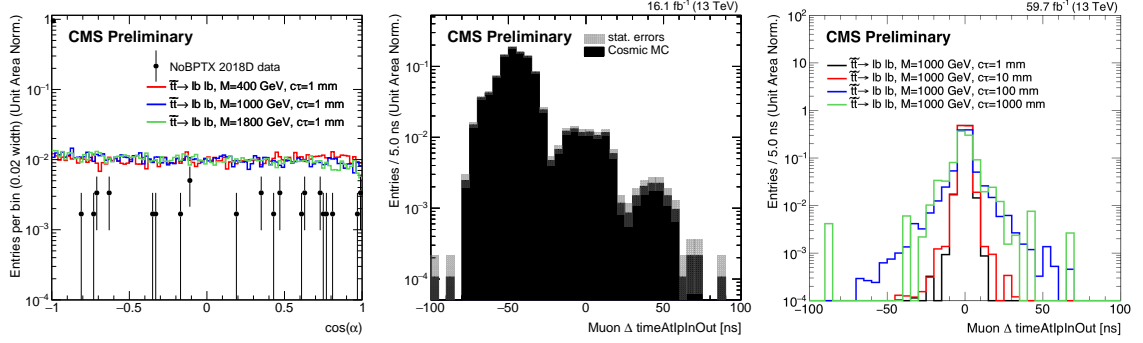


Figure 3.8: The di-muon $\cos(\alpha)$ distribution of 2018 NoBPTX data and simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events (left) and the di-muon Δt distribution in simulated cosmic-ray muon events in 2016 conditions (center) and $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ (right) events in 2018 conditions. In each case, all $\mu\mu$ preselection criteria except those relating to cosmic rejection are applied.

muons are traveling outwards from the center of the detector. We then use the muon ϕ measurements to determine which muon is above the other and find Δt , the time of the lower muon subtracted from the time of the upper muon. We reject events with $\Delta t < -20$ if the number of degrees of freedom of the timing measurements for both muons is greater than seven. Figure 3.8 shows the $\cos \alpha$ and Δt distributions of NoBPTX data and simulated cosmic and $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events.

To remove leptons from decays of SM hadrons, we require that the candidate leptons not be too close together in the η - ϕ plane. Specifically, we find that requiring $\Delta R > 0.2$ significantly reduces the contribution from SM hadrons without noticeably affecting the signal acceptance. To remove leptons from material interactions, we reject events in which the candidate leptons form a good displaced vertex that overlaps with the tracker material. The vertices are reconstructed with the Kalman Vertex Fitter, and a “good” vertex is one with $\chi^2/n_{\text{dof}} < 20$. The tracker material map is

obtained from the tracker material budget measurements [74, 75]. See Section 3.5.6 for tests in data that involve inverting the criteria described in this paragraph.

Finally, to ensure that the signal regions of all three channels are orthogonal to one another, we reject events in the ee ($\mu\mu$) channel with at least one muon (electron) that passes the $e\mu$ channel preselection and has $|d_0| > 100\,\mu\text{m}$.

In contrast to previous displaced leptons analyses [50, 51], we allow for the possibility of more than one lepton of each type in a given channel and set no requirements on the charge product of the lepton pair. These changes were made at the request of several theorists, including the authors of Ref. [76].

Figure 3.9 shows the electron and muon $|d_0|$ distributions in simulated signal and background events that pass the 2018 $e\mu$ preselection, and Fig. 3.10 shows the cumulative number of simulated signal events that pass each 2018 preselection criterion in all three channels for several top squark lifetime hypotheses.

3.3.3 Prompt control region

In order to verify the implementation of our selection and corrections to simulation (see Section 3.4), we define a prompt control region that is dominated by SM background events. Events in each channel’s prompt control region are selected by requiring that they pass all of the criteria defined in Section 3.3.2 as well as the requirement that the candidate leptons have $|d_0| < 50\,\mu\text{m}$. We define this region in each channel in order to check for reasonable agreement between simulated SM events and data after applying the corrections described in Section 3.4. Some examples are shown in Figs. 3.11, 3.12, and 3.13, which show the p_T , η , and $|d_0|$ distributions of the leptons in the $e\mu$, ee , and $\mu\mu$ prompt control regions, respectively, for 2016 data and

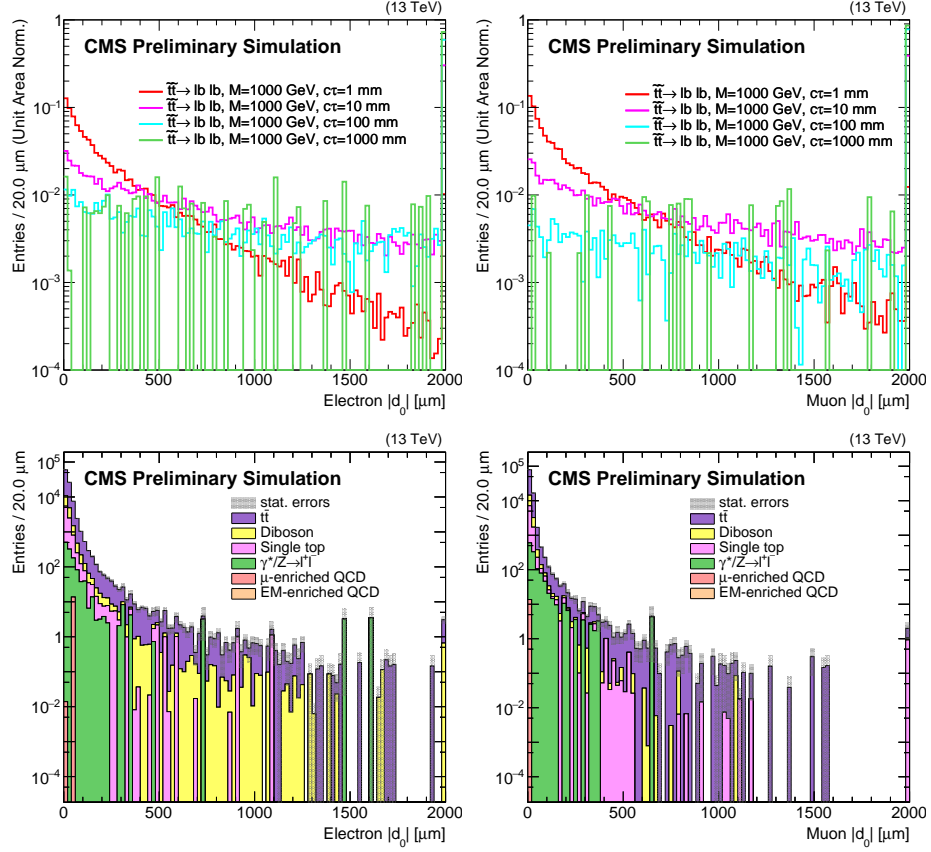


Figure 3.9: The electron (left) and muon (right) $|d_0|$ distributions for 2018 simulation events that pass the $e\mu$ preselection criteria. The upper two plots show $t\bar{t} \rightarrow \bar{l}b l\bar{b}$ simulation for a single \tilde{t} mass and four different proper decay lengths; each histogram is normalized to unity. The lower two plots show the background simulation normalized to the integrated luminosity. In all of the histograms, the last bin includes the overflow. All of the corrections from Section 3.4 are applied.

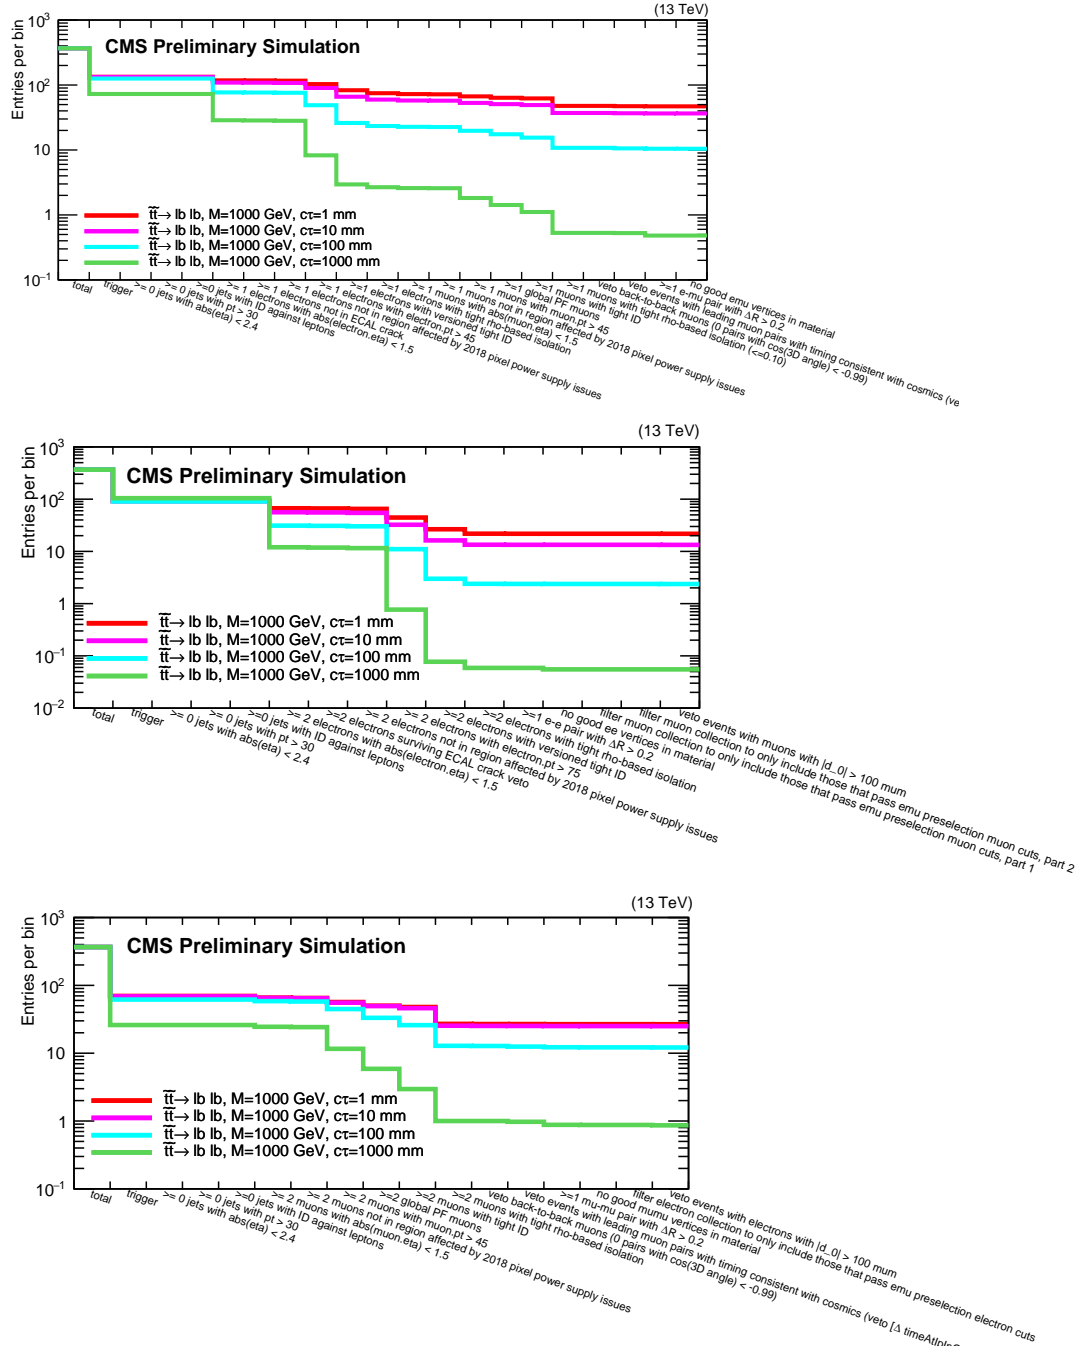


Figure 3.10: The cumulative number of simulated $t\bar{t} \rightarrow l\bar{b} l\bar{b}$ events that pass each criterion in the 2018 $e\mu$ (top), ee (middle), and $\mu\mu$ (bottom) preselection. Several \tilde{t} proper decay lengths are shown. The jet criteria do not exclude any events and are simply an artifact of the analysis framework.

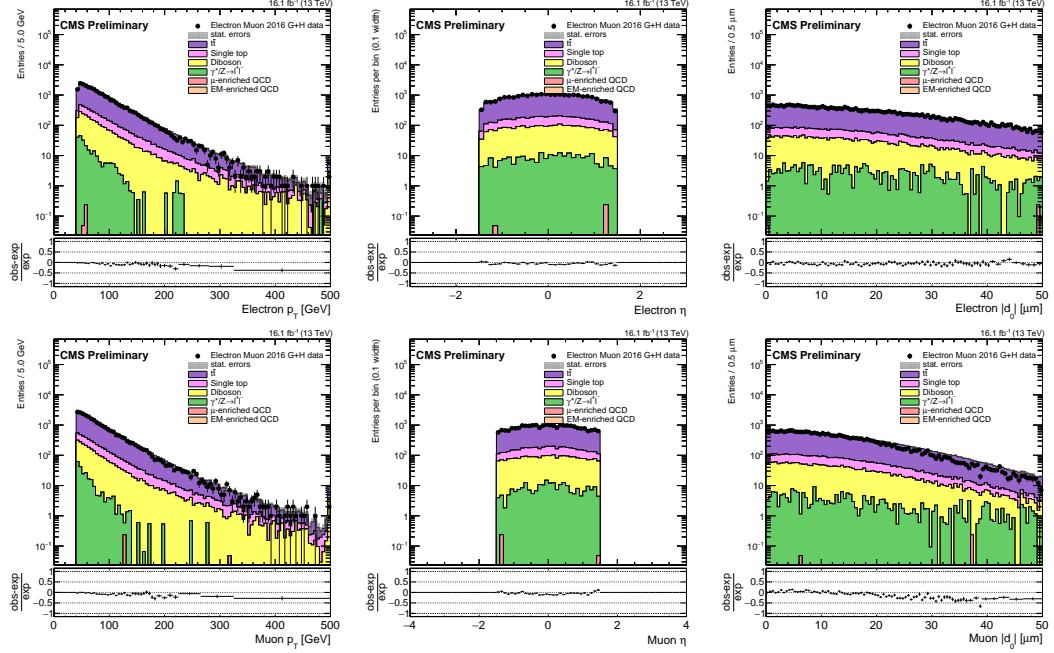


Figure 3.11: The electron (top) and muon (bottom) p_T (left), η (center), and $|d_0|$ (right) distributions in the $e\mu$ prompt control region for 2016 data and simulated background events. The rightmost bin in each plot contains the overflow entries.

background simulation. The data-driven background estimation technique employed in this analysis removes the need for exact agreement between data and simulation, but the absence of any significant discrepancies gives us confidence that we are accounting for the correct sources of prompt SM leptons and that our selection and corrections are functioning as intended.

3.3.4 Inclusive signal region

Finally, we define the region to which new physics may contribute significantly. The inclusive signal region is populated by events that pass all of the criteria defined in Section 3.3.2 as well as the requirement that the candidate leptons each have $100 \mu\text{m} < |d_0| < 10 \text{ cm}$. We do not select leptons with $|d_0| > 10 \text{ cm}$ because the tracking

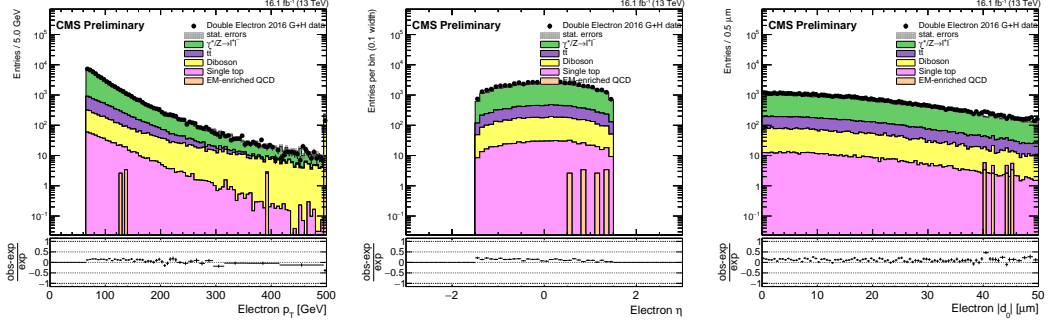


Figure 3.12: The electron p_T (left), η (center), and $|d_0|$ (right) distributions in the ee prompt control region for 2016 data and simulated background events. The rightmost bin in each plot contains the overflow entries.

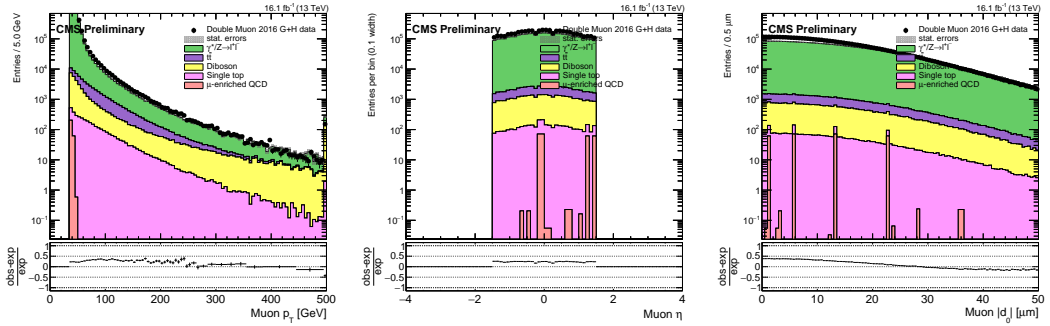


Figure 3.13: The muon p_T (left), η (center), and $|d_0|$ (right) distributions in the $\mu\mu$ prompt control region for 2016 data and simulated background events. The rightmost bin in each plot contains the overflow entries.

Table 3.7: The cumulative efficiency for simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ signal events to pass the 2018 inclusive signal region selection, for several choices of \tilde{t} mass and $c\tau$. The corrections described in Section 3.4 are applied.

$e\mu$ inclusive signal region			
	200 GeV	1000 GeV	1800 GeV
0.1 cm	2.1%	4.6%	4.6%
1 cm	3.5%	7.9%	8.8%
10 cm	1.0%	2.7%	3.4%
100 cm	0.047%	0.13%	0.16%

ee inclusive signal region			
	200 GeV	1000 GeV	1800 GeV
0.1 cm	0.47%	2.1%	2.1%
1 cm	0.60%	2.8%	3.3%
10 cm	0.11%	0.59%	0.76%
100 cm	0.003%	0.014%	0.022%

$\mu\mu$ inclusive signal region			
	200 GeV	1000 GeV	1800 GeV
0.1 cm	1.4%	2.6%	2.5%
1 cm	3.1%	5.6%	5.9%
10 cm	1.5%	3.2%	3.7%
100 cm	0.11%	0.23%	0.34%

efficiency drops sharply after this point, as shown in Section C. This requirement also ensures that the leptons originate within the pixel volume, which is effectively required by the pixel hit requirement of the tight lepton IDs. Table 3.7 lists the cumulative efficiency for $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events to pass the full 2018 inclusive signal region selection for several signal points. To ensure sensitivity to a wide range of new particle masses and lifetimes, we further subdivide the inclusive signal region into bins defined by the $|d_0|$ of each candidate lepton and the p_T of one candidate lepton. The exact binning is described in Section 3.5.2.

3.4 Corrections to simulation

We apply several corrections to the simulated background and signal events in order to account for known differences between simulation and data. Each correction is described individually in the following sections.

3.4.1 Pileup

The simulation is corrected so that its distribution of the number of pileup interactions matches that of 2016, 2017, and 2018 data. Each simulated sample is reweighted, event-by-event, by scale factors derived by dividing the pileup distribution in data by the pileup distribution in the given simulated sample.

3.4.2 Lepton ID

We apply scale factors provided by the CMS Physics Object Groups to correct for known differences in the lepton reconstruction and tight ID performance between data and simulation. Although our lepton ID differs from the standard tight ID in that we do not set requirements on d_0 or d_z , the standard scale factors are still sufficient because they are derived from leptons from Z boson decays that are representative of leptons in the bulk of the $|d_0|$ distribution, and we apply additional systematic uncertainties to account for possible differences at larger $|d_0|$ (see Section 3.6).

3.4.3 Lepton d_0 resolution

As shown in Fig. 3.14, the agreement between data and simulation in the 2017 and 2018 electron and muon $|d_0|$ distributions was initially poor. We found that the average muon and electron d_0 fluctuated periodically with respect to ϕ in 2017 and 2018 data but not in background simulation, as can be seen in Fig. 3.15. This periodic

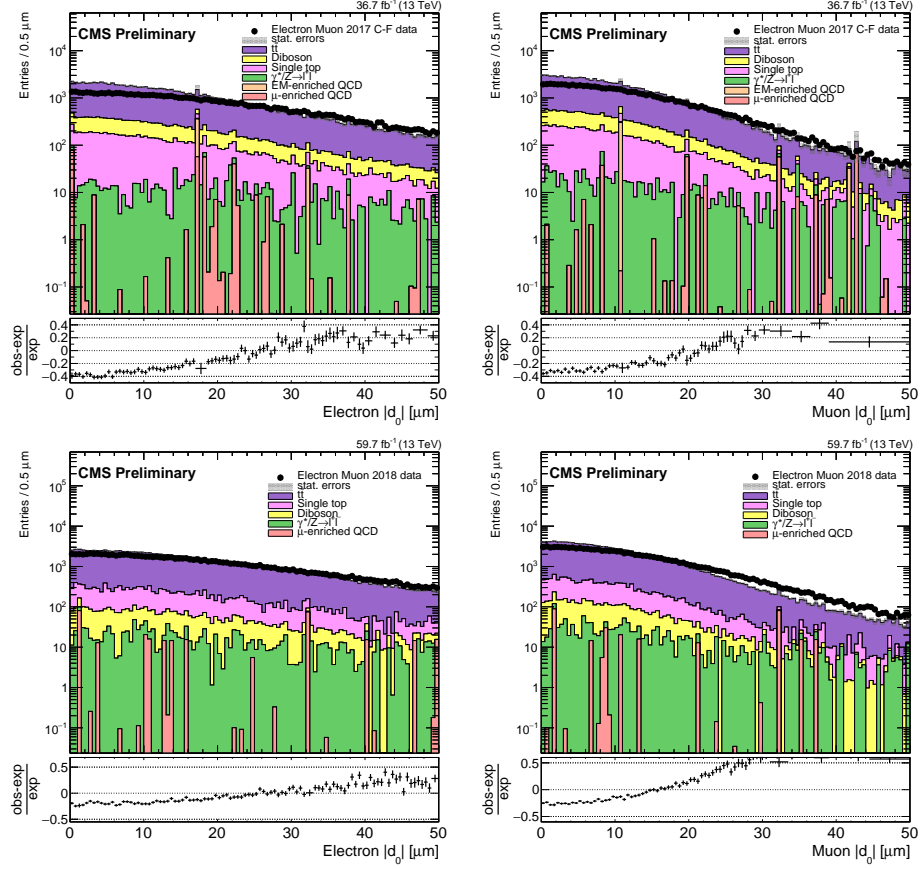


Figure 3.14: The uncorrected lepton $|d_0|$ distributions in the $e\mu$ prompt control region, for electrons (left) and muons (right), for 2017 data and simulation (top), and 2018 data and simulation (bottom). The rightmost bin in each plot contains the overflow entries.

fluctuation in data but not in simulation indicates that an overly optimistic simulated tracker alignment is responsible for the unrealistically narrow d_0 distribution in simulation.

To account for the overly optimistic alignment in simulation, we smear the electron and muon d_0 in 2017 and 2018 simulation in each channel's prompt control region to better model the d_0 distribution in data. To do this, we first fit the central regions of the background simulation and data d_0 distributions with Gaussian functions in each

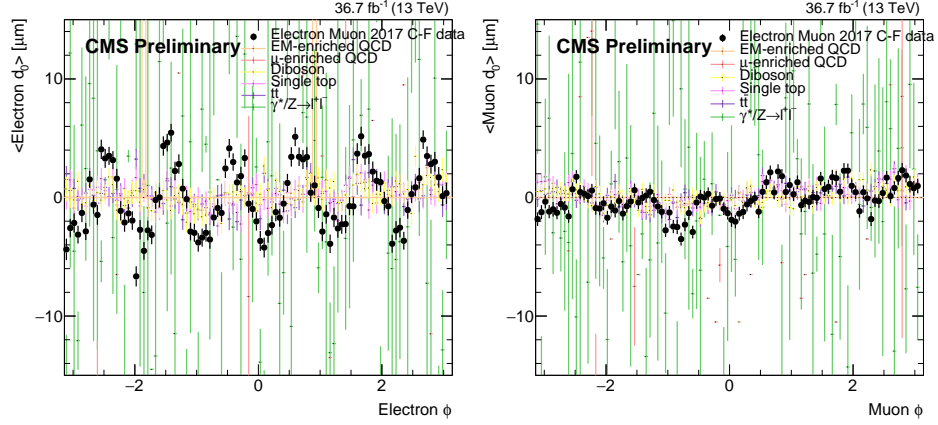


Figure 3.15: The average lepton $|d_0|$ as a function of ϕ in the $e\mu$ prompt control region, for electrons (left) and muons (right), for 2017 data and simulation.

channel's prompt control region and then compare the widths of the Gaussian fits. The fitted distributions are shown in Figs. 3.16 and 3.17 for the $e\mu$ channel. Assuming that the width of each Gaussian fit is mostly determined by the d_0 resolution, we define $\sigma_{data}^2 = \sigma_{bkg}^2 + \sigma_{align}^2$, where σ_{data} is the data Gaussian width, σ_{bkg} is the uncorrected background simulation Gaussian width, and σ_{align} is the additional component that is needed to make up the difference in d_0 resolution between background simulation and data. We find σ_{data} and σ_{bkg} from the fits and compute σ_{align} . The fit results are similar in the $e\mu$ channel shown here and in the same-flavor channels. We average the σ_{align} derived in the ee channel and the $e\mu$ channel for electrons, and in the $\mu\mu$ channel and the $e\mu$ channel for muons. The average σ_{align} is shown in Table 3.8. We then smear the simulation d_0 values with values drawn from a Gaussian distribution centered on zero and with a width of the average σ_{align} . The smearing is applied to both background and signal simulation. The corrected $|d_0|$ distributions are shown in Fig. 3.18.

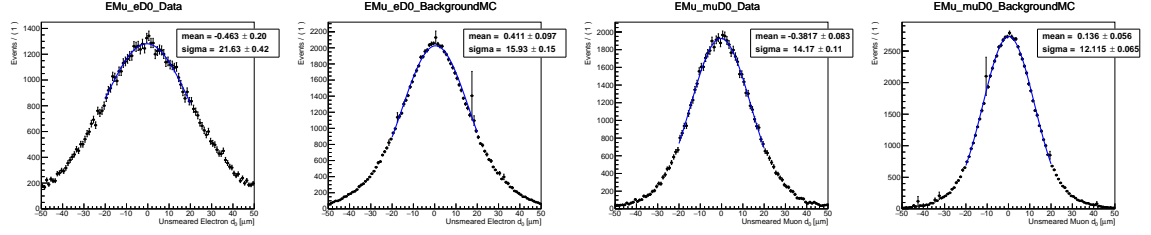


Figure 3.16: The lepton d_0 distributions with Gaussian fits in the 2017 $e\mu$ prompt control region for (working from left to right) electrons in data, electrons in background simulation, muons in data, and muons in background simulation.

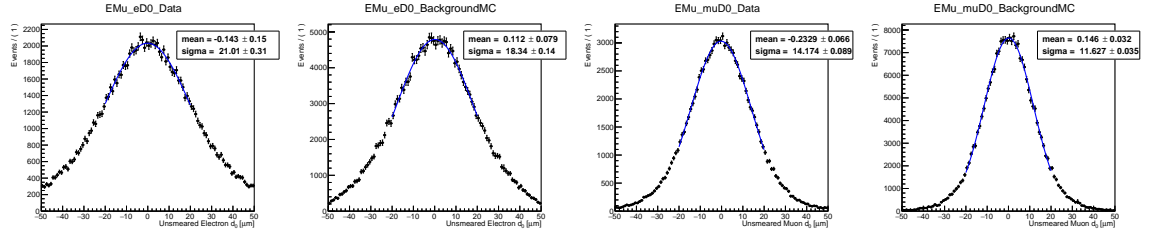


Figure 3.17: The lepton d_0 distributions with Gaussian fits in the 2018 $e\mu$ prompt control region for (working from left to right) electrons in data, electrons in background simulation, muons in data, and muons in background simulation.

Table 3.8: The average σ_{align} for electrons and muons, for the 2017 and 2018 analyses.

	2017	2018
Electrons	$14.75 \pm 0.36 \mu\text{m}$	$9.18 \pm 0.41 \mu\text{m}$
Muons	$7.57 \pm 0.12 \mu\text{m}$	$8.11 \pm 0.08 \mu\text{m}$

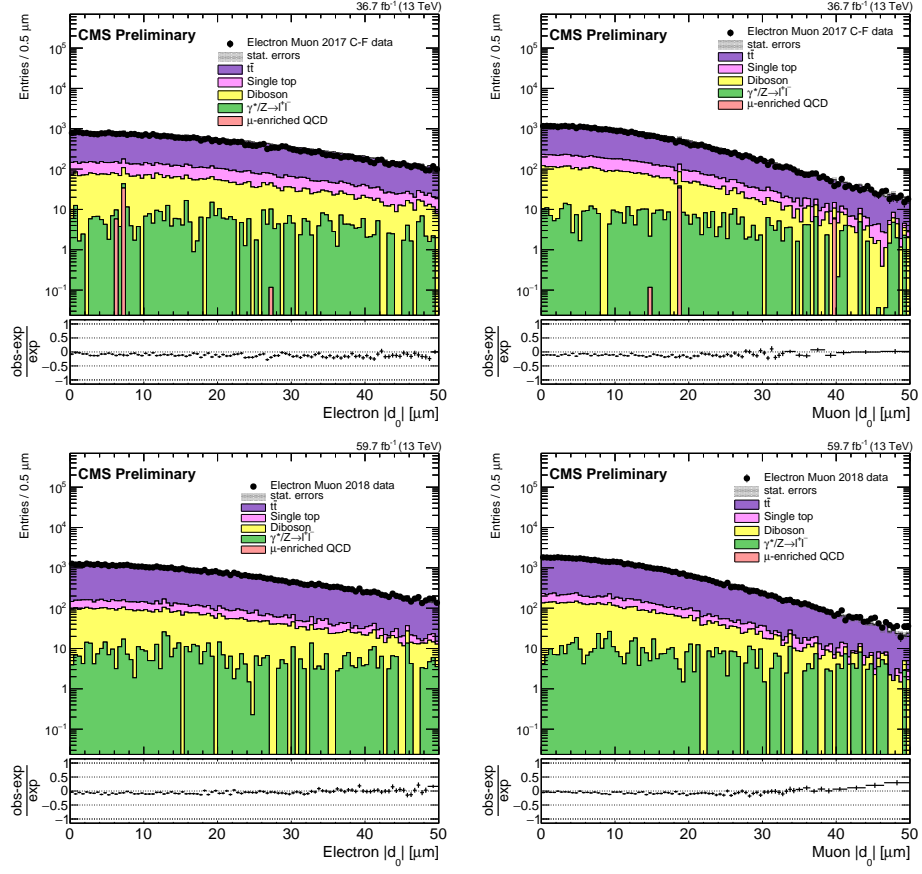


Figure 3.18: The corrected lepton $|d_0|$ distributions in the $e\mu$ prompt control region, for electrons (left) and muons (right), for 2017 data and simulation (top), and 2018 data and simulation (bottom). The rightmost bin in each plot contains the overflow entries.

This d_0 smearing has a minimal effect on the final result because the width of the Gaussian distribution from which the smearing values are drawn is small relative to the size of the signal region bins, but understanding the source of the poor agreement between data and simulation was important to validate our understanding of the SM background.

3.4.4 Trigger efficiency

We also apply scale factors to the simulated background and signal events to correct for differences in trigger efficiency between data and simulation. To measure the trigger efficiency, we first require that events pass an OR of several unprescaled p_T^{miss} triggers (see Table 3.9) and the preselection criteria with the lepton p_T requirement excluded. The p_T^{miss} triggers provide a sample of di-lepton events that is unbiased with respect to the main triggers used in the analysis, and excluding the lepton p_T requirement allows us to study the trigger efficiency as a function of lepton p_T . In the $e\mu$ channel, the electron (muon) p_T is required to be greater than 50 GeV when plotting against the muon (electron) p_T to disentangle the effect from the other leg of the muon-photon trigger. Data events are taken from the MET primary dataset (which contains events that pass p_T^{miss} triggers) and simulated background events are taken from $t\bar{t}$ simulation for the $e\mu$ channel and Drell-Yan simulation for the same-flavor channels.

To calculate the efficiency, we divide the lepton p_T distribution in events that pass the standard analysis triggers in addition to the OR of the p_T^{miss} triggers and the preselection by the lepton p_T distribution in events that pass the OR of the p_T^{miss} triggers and the preselection. We then compute the scale factor as the ratio of

Table 3.9: The unprescaled p_T^{miss} triggers used to create an orthogonal data sample for the trigger efficiency calculation.

2016

HLT_MET200
 HLT_MonoCentralPFJet80_PFMETNoMu110_PFMHTNoMu110_IDTight
 HLT_PFMET120_PFMHT120_IDTight
 HLT_PFMET170_HBHECleaned
 HLT_PFMET300
 HLT_PFMETNoMu120_PFMHTNoMu120_IDTight

2017

HLT_CaloMET350_HBHECleaned
 HLT_MonoCentralPFJet80_PFMETNoMu120_PFMHTNoMu120_IDTight
 HLT_PFMET120_PFMHT120_IDTight
 HLT_PFMET250_HBHECleaned
 HLT_PFMETNoMu120_PFMHTNoMu120_IDTight

2018

HLT_CaloMET350_HBHECleaned
 HLT_MonoCentralPFJet80_PFMETNoMu120_PFMHTNoMu120_IDTight
 HLT_PFMET120_PFMHT120_IDTight
 HLT_PFMET200_HBHE_BeamHaloCleaned
 HLT_PFMET250_HBHECleaned
 HLT_PFMETNoMu120_PFMHTNoMu120_IDTight

Table 3.10: The trigger efficiency scale factors in each channel and year.

	2016	2017	2018
$e\mu$ channel, electrons	0.974 ± 0.016	0.961 ± 0.013	0.965 ± 0.012
$e\mu$ channel, muons	0.975 ± 0.016	0.972 ± 0.014	0.967 ± 0.012
ee channel	1.000 ± 0.099	1.000 ± 0.131	1.000 ± 0.185
$\mu\mu$ channel	0.956 ± 0.012	0.930 ± 0.010	0.935 ± 0.011

the efficiency in data to the efficiency in simulation in the plateau of the efficiency distribution. The $e\mu$ channel efficiency distributions are shown in Fig. 3.19, and the resulting scale factors for all channels are listed in Table 3.10.

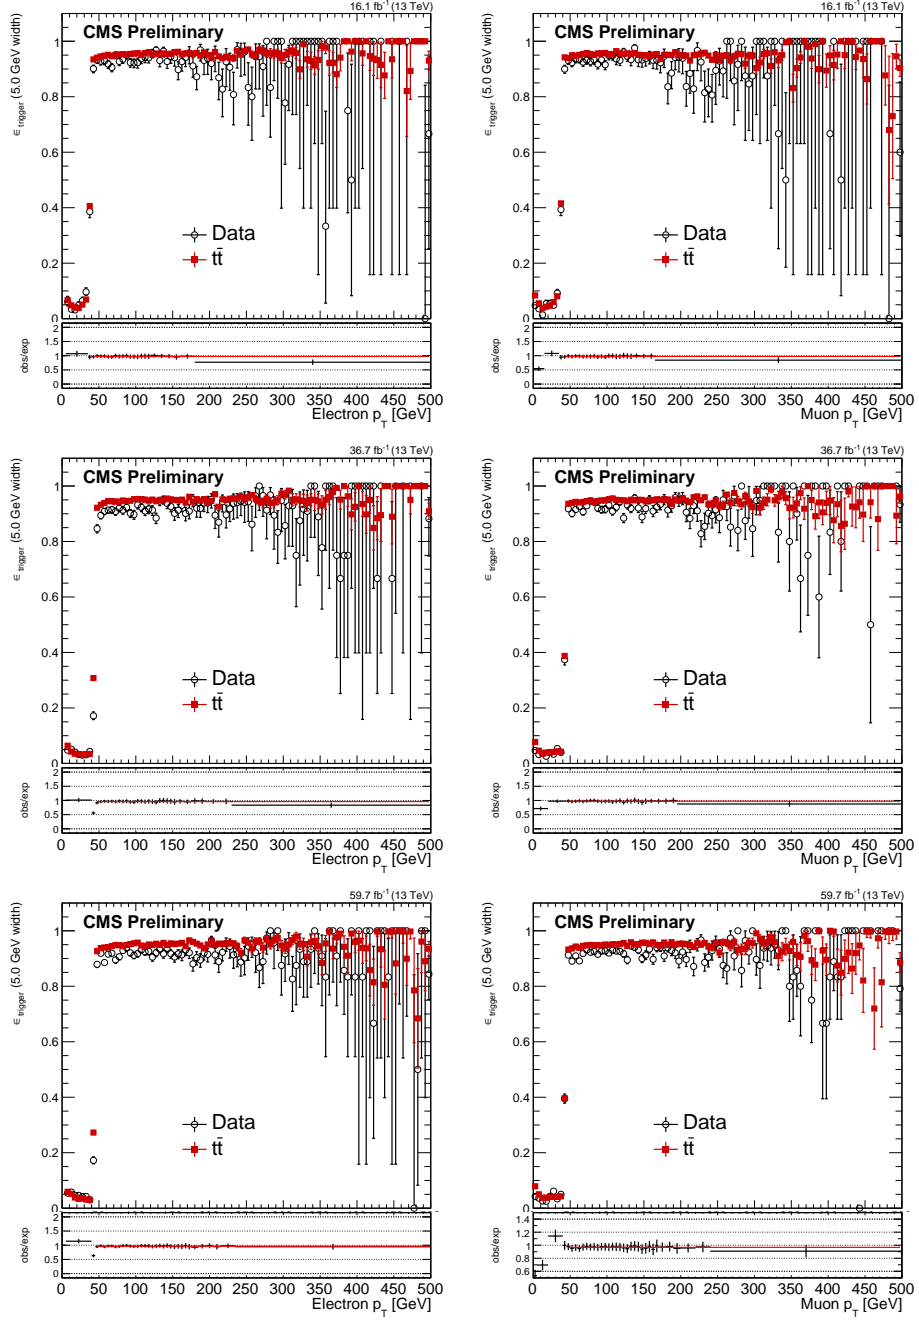


Figure 3.19: Trigger efficiency as a function of leading electron (left) or leading muon (right) p_T in the $e\mu$ channel in 2016 (top), 2017 (middle), and 2018 (bottom) in data and simulated background $t\bar{t}$ events.

3.5 Background estimation

3.5.1 Background sources

The vast majority of leptons from SM processes are prompt. There are, however, a few notable processes that produce leptons with large $|d_0|$ values that may pass the preselection: (1) leptons from prompt decays whose $|d_0|$ is poorly measured (“mismeasurements”), (2) leptons from decays of tau leptons (“taus”), and (3) leptons from decays of B or D mesons (“heavy flavor”). Note that the leptons from these processes generally do not share a common displaced vertex. We perform several cross checks to ensure that processes in which leptons share a common displaced vertex do not contribute significantly in the signal regions (SRs). Section 3.5.6 presents several additional studies that confirm that the SR contributions of leptons from material interactions, cosmic rays, and SM hadrons are either negligible or already accounted for by the background estimation procedure outlined in Section 3.5.2.

The $|d_0|$ distributions of leptons from the three main background sources vary according to lepton flavor and parent particle. Tau leptons, B mesons, and D mesons have proper decay lengths of $87\,\mu\text{m}$, approximately $500\,\mu\text{m}$, and approximately $100\,\mu\text{m}$, respectively, so leptons from taus will generally be less displaced than leptons from heavy flavor. Furthermore, mismeasurements are more common for electrons than for muons due to the superior muon d_0 resolution. Figure 3.20, which shows the relative contribution of each source of leptons as a function of $|d_0|$ in simulated $t\bar{t}$ events that pass the $e\mu$ channel preselection, shows how mismeasurements dominate at all $|d_0|$ values for electrons while taus and heavy flavor contribute meaningfully for muons with $|d_0| \gtrsim 100\,\mu\text{m}$.

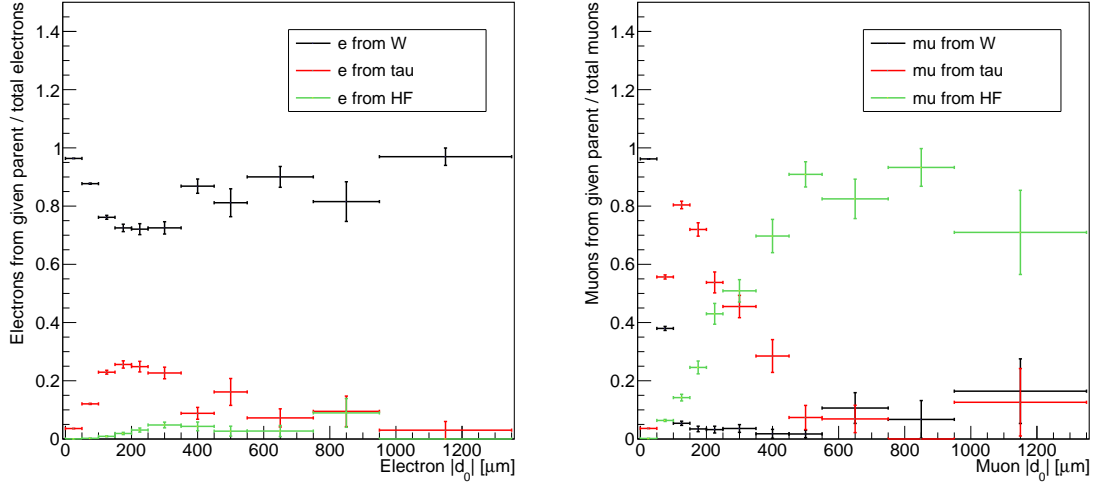


Figure 3.20: The fraction of electrons (left) and muons (right) from different parents as a function of lepton $|d_0|$ in simulated $t\bar{t}$ events that pass the 2018 $e\mu$ channel preselection. Note that the vast majority of leptons whose parent is a W boson are produced in prompt decays.

In addition to the $|d_0|$ behavior of individual leptons, we are also interested in understanding whether any background processes introduce correlation in $|d_0|$ between pairs of leptons. Having confirmed that the obvious potential source of $|d_0|$ - $|d_0|$ correlation, i.e. pairs of leptons that share a common displaced vertex, do not contribute meaningfully to the SR, we next investigate possible sources of correlation from mismeasurements, taus, and heavy flavor. As the preselection criteria ensure that we are operating in the regime in which lepton $|d_0|$ resolution does not depend meaningfully on p_T or η , and because the two candidate leptons in a given event are nearly always from prompt processes and/or independent decays, mismeasurements generally will not introduce $|d_0|$ - $|d_0|$ correlation.

Taus and heavy flavor, on the other hand, could potentially introduce $|d_0|$ - $|d_0|$ correlation through correlation in parentage between the two candidate leptons. As will be shown in Section 3.5.3, we find that taus lead to exactly this type of correlation through processes such as $Z \rightarrow \tau\tau \rightarrow ll$ while heavy flavor does not. Figure 3.21 shows the fraction of muons from different SM parents as a function of $|d_0|$ in simulated Drell-Yan events that pass the $\mu\mu$ preselection with the added requirement that both muons come from the same type of parent particle. Here, we see the same characteristic $|d_0|$ behavior of leptons from taus and also that the isolation criteria effectively eliminates events in which both leptons come from a heavy-flavor meson. We therefore conclude that while leptons from mismeasurements, taus, and heavy flavor will all contribute in the SR, only leptons from taus will introduce meaningful $|d_0|$ - $|d_0|$ correlation. Furthermore, the characteristic $|d_0|$ behavior of leptons from taus seen in Figs. 3.20 and 3.21 implies that this correlation will specifically impact the 100 to 500 μm region.

3.5.2 Data-driven ABCD method

We estimate the SR background yields with a data-driven method in which the lepton $|d_0|$ distributions serve as composite models of all background processes. Specifically, we employ an ABCD method using the $|d_0|$ of two leptons. We label the two $|d_0|$ values in each channel as $|d_0^a|$ and $|d_0^b|$, which correspond to the leading electron and leading muon in the $e\mu$ channel, the leading and subleading electrons in the ee channel, and the leading and subleading muons in the $\mu\mu$ channel. As a first step, we categorize the events that pass the preselection criteria into four regions (A, B, C, and D) of the $|d_0^a|$ - $|d_0^b|$ plane, as shown in Fig. 3.22.

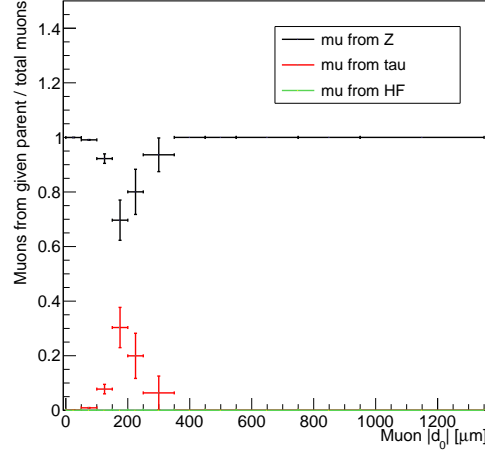


Figure 3.21: The fraction of muons from different parents as a function of muon $|d_0|$ in simulated Drell-Yan events that pass the 2018 $\mu\mu$ channel preselection and the additional constraint that the leading and subleading muon both come from the same type of parent particle.

In order to maximize sensitivity to a wide range of new particle lifetimes, we subdivide region D into four SRs:

- SR I: $100 \leq |d_0^a| < 500 \mu\text{m}$, $100 < |d_0^b| < 500 \mu\text{m}$
- SR II: $100 \leq |d_0^a| < 500 \mu\text{m}$, $500 \mu\text{m} < |d_0^b| < 10 \text{ cm}$
- SR III: $500 \mu\text{m} \leq |d_0^a| < 10 \text{ cm}$, $100 < |d_0^b| < 500 \mu\text{m}$
- SR IV: $500 \mu\text{m} \leq |d_0^a| < 10 \text{ cm}$, $500 \mu\text{m} < |d_0^b| < 10 \text{ cm}$

The exact boundaries between the four SRs are motivated by the expected contributions of the different background sources, as explained in 3.5.1. This approach also necessitates that the definitions of regions B and C vary in accordance with the SR for which a given estimate is performed (e.g. only the events in the $100 \leq |d_0^a| < 500 \mu\text{m}$ range of region B are considered when estimating the yields of SR I and II). Finally,

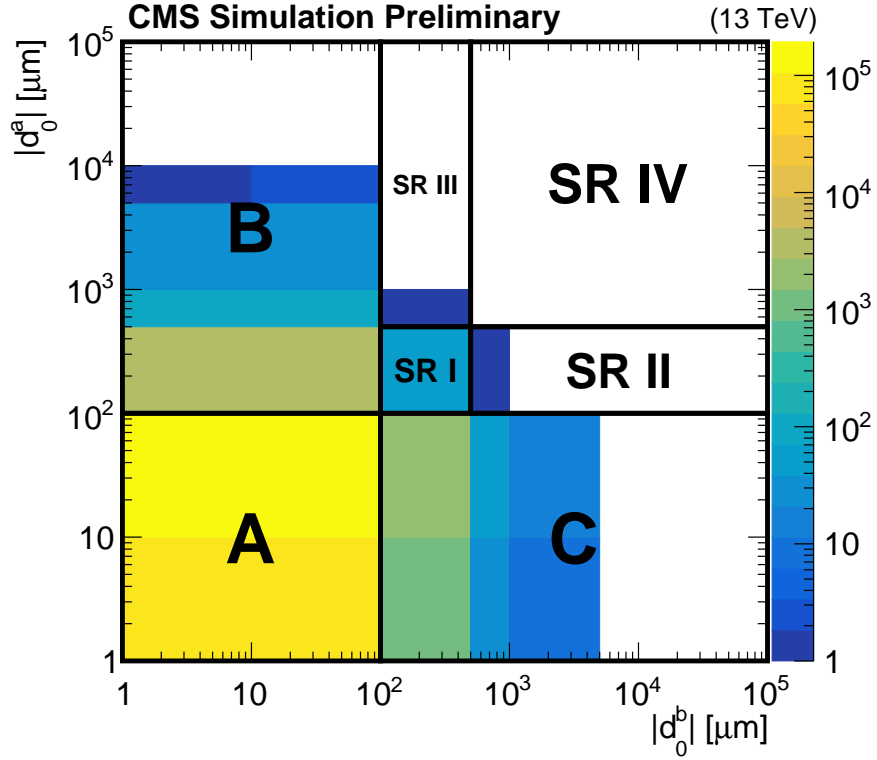


Figure 3.22: A diagram of the ABCD method overlaid on simulated background events passing the 2018 $e\mu$ preselection. A, B, and C are control regions, and D corresponds to the inclusive SR, which includes SRs I, II, III, and IV. Underflow events are included in the bins along the left and bottom edges. When performing the background estimate, regions B and C are further subdivided to coincide with the SR for which the estimate is being performed.

Table 3.11: The p_T boundaries between the low- and high- p_T bins of SR I in each channel.

	p_T boundary [GeV]
2016 $e\mu$	leading μ $p_T = 90$
2017+2018 $e\mu$	leading μ $p_T = 140$
2016 ee	leading e $p_T = 300$
2017+2018 ee	leading e $p_T = 400$
2016 $\mu\mu$	leading μ $p_T = 100$
2017+2018 $\mu\mu$	leading μ $p_T = 100$

we subdivide SR I into two bins using one lepton’s p_T to further increase sensitivity to high-mass, low-lifetime new physics. Table 3.11 lists the p_T boundary in each channel and year.

We then use the number of events in regions A, B, and C to estimate the expected background in each SR. The basic estimation procedure depends on the assumption that $|d_0^a|$ and $|d_0^b|$ are uncorrelated. If this assumption holds, then $N_B/N_A = N_D/N_C$ and the number of background events in D is equal to $N_B \times N_C/N_A$, where N_X is the number of background events in the given region. We find that $|d_0^a|$ and $|d_0^b|$ are indeed uncorrelated over much of the $|d_0^a|$ - $|d_0^b|$ plane, but the correlation discussed in Section 3.5.1 renders the simple ABCD method insufficient to estimate the background in SR I. After quantifying the degree of correlation in Section 3.5.3, we define a procedure to correct the simple SR I ABCD estimate in Section 3.5.4.

When performing the background estimate and closure tests, we treat the 2016 data and simulation separately from the 2017–2018 data and simulation to avoid any correlations between $|d_0^a|$ and $|d_0^b|$ that may arise from the differences between

the original and Phase-1 pixel detectors employed by CMS in 2016 and 2017–2018, respectively (see Section 2.2.2).

3.5.3 Closure tests in control regions

We perform several closure tests of the background estimation procedure in data and simulation to test the method and quantify the degree of $|d_0^a|$ - $|d_0^b|$ correlation from the processes discussed in 3.5.1. Two series of tests are performed, the first in the 100–500 μm subregions of regions B and C and the second in the 500 μm –10 cm subregions of regions B and C.

100–500 μm tests

We perform closure tests in subregions of regions B and C where one lepton is more prompt (20–100 μm) and the other is more displaced (100–500 μm). In these closure tests, we estimate the background yield using the simple ABCD method and use the ratio of the actual number of events to the estimated number of events as a measure of nonclosure (and therefore $|d_0^a|$ - $|d_0^b|$ correlation). With this approach, a ratio of 1 corresponds to closure and negligible correlation while ratios greater than 1 correspond to nonclosure and positive correlation. To understand how the degree of correlation varies with $|d_0|$, we perform several closure tests in each sideband while varying the $|d_0|$ range of the more-prompt lepton and fit the resulting ratios with a first-degree polynomial. We then extrapolate the resulting fit from the closure test regions to SR I to estimate the degree of nonclosure in SR I. We perform identical procedures in regions B and C and then average the resulting extrapolated ratios.

Table 3.12 shows the average extrapolated ratios for three rounds of closure tests: one in background simulation with the $Z \rightarrow \tau\tau \rightarrow ll$ events removed, one in the full

Table 3.12: Closure test results in data and background simulation in the 100–500 μm region. The average extrapolated ratios and their statistical uncertainties are given. For each test, the A, B, C, and D regions are defined as follows: A is 20–30 μm in prompt lepton $|d_0|$ and 20–100 μm in displaced lepton $|d_0|$, B is 20–30 μm in prompt lepton $|d_0|$ and 100–500 μm in displaced lepton $|d_0|$, C is always 20–100 μm in displaced lepton $|d_0|$, D (the test region) is always 100–500 μm in displaced lepton $|d_0|$, and we perform repeated tests while simultaneously varying the C and D prompt lepton $|d_0|$ within the 30–100 μm range.

	Bkg. simulation without $Z \rightarrow \tau\tau \rightarrow ll$	Full bkg. simulation	Data
2016 $e\mu$	0.9 ± 0.3	1.6 ± 0.6	0.9 ± 1.3
2017+2018 $e\mu$	1.1 ± 0.4	1.6 ± 0.7	3.1 ± 0.8
2016 ee	0.8 ± 0.5	0.8 ± 0.5	0.6 ± 0.6
2017+2018 ee	0.8 ± 1.0	1.6 ± 0.9	1.5 ± 0.4
2016 $\mu\mu$	1.1 ± 0.8	2.0 ± 0.8	2.5 ± 0.9
2017+2018 $\mu\mu$	2.6 ± 2.8	7.8 ± 3.7	4.2 ± 1.5
Average	1.2 ± 0.5	2.6 ± 0.7	2.1 ± 0.4

background simulation, and one in data. The average extrapolated ratios are always compatible with one in background simulation without $Z \rightarrow \tau\tau \rightarrow ll$ events, but they generally increase when the $Z \rightarrow \tau\tau \rightarrow ll$ events are included. Furthermore, the average extrapolated ratios from the full background simulation generally describe the average extrapolated ratios in data. From these results, we conclude that within our statistical uncertainties, $Z \rightarrow \tau\tau \rightarrow ll$ events are the only meaningful source of correlation and that the degree of correlation observed in data is modeled reasonably well in simulation. We also observe that the variation in the degree of correlation across channels matches our expectations: correlation increases with the number of muons in the final state and is greater in 2017–2018 than 2016 because of the improved d_0 resolution made possible by the Phase-1 tracker upgrade (described in Section 2.2.2).

500 μm –10 cm tests

We next perform closure tests in more-displaced subregions of regions B and C where one lepton is more prompt (20–100 μm) and the other is more displaced (500 μm –10 cm). We again use the ratio of the actual number of events to the estimated number of events as the measure of nonclosure, but in these tests we expect the ratio to be consistent with one because $Z \rightarrow \tau\tau \rightarrow ll$ events do not contribute meaningfully beyond 500 μm . Table 3.13 shows that this is indeed the case for background simulation (with and without $Z \rightarrow \tau\tau \rightarrow ll$ events) and for data. These results imply that $|d_0^a|$ and $|d_0^b|$ are uncorrelated beyond 500 μm , which means that a simple ABCD procedure will be adequate for estimating the background yields in SRs II, III, and IV.

3.5.4 ABCD correction and systematic uncertainty

The closure tests of Section 3.5.3 show that $|d_0^a|$ and $|d_0^b|$ are frequently positively correlated in the 100–500 μm region but are uncorrelated otherwise. To account for this correlation as well as other possible unforeseen sources of nonclosure, we define a procedure to correct the simple ABCD estimate in SR I and assign a systematic uncertainty to the simple ABCD estimate in all SRs.

100–500 μm correction and systematic uncertainty

Figures 3.23, 3.24, and 3.25 show the results of the data closure tests in the $e\mu$ ee , and $\mu\mu$ channels, respectively, in the subregions of regions B and C where one lepton is more prompt (20–100 μm) and the other is more displaced (100–500 μm). These plots show the ratio of the actual to the estimated number of events as a function of the prompt lepton $|d_0|$. In each plot, the binning of the prompt lepton axis is initially

Table 3.13: Closure test results in data and background simulation in the $500\,\mu\text{m}$ – $10\,\text{cm}$ region. The ratios of the actual to the estimated yield and their statistical uncertainties are given. The A, B, C, and D regions are defined as follows: A is 20 – $30\,\mu\text{m}$ in prompt lepton $|d_0|$ and 20 – $100\,\mu\text{m}$ in displaced lepton $|d_0|$, B is 20 – $30\,\mu\text{m}$ in prompt lepton $|d_0|$ and $500\,\mu\text{m}$ – $10\,\text{cm}$ in displaced lepton $|d_0|$, C is 30 – $100\,\mu\text{m}$ in prompt lepton $|d_0|$ and 20 – $100\,\mu\text{m}$ in displaced lepton $|d_0|$, and D (the test region) is 30 – $100\,\mu\text{m}$ in prompt lepton $|d_0|$ and $500\,\mu\text{m}$ – $10\,\text{cm}$ in displaced lepton $|d_0|$.

Region B			
	Bkg. simulation without $Z\rightarrow\tau\tau\rightarrow ll$	Bkg. simulation	Data
2016 $e\mu$	$1.1^{+0.3}_{-0.3}$	$1.1^{+0.3}_{-0.3}$	$0.4^{+1.0}_{-0.4}$
2017+2018 $e\mu$	$0.9^{+0.3}_{-0.2}$	$0.9^{+0.3}_{-0.2}$	$0.7^{+0.3}_{-0.3}$
2016 ee	$0.4^{+0.6}_{-0.3}$	$0.4^{+0.6}_{-0.3}$	$1.4^{+1.6}_{-0.9}$
2017+2018 ee	$0.5^{+0.8}_{-0.4}$	$0.3^{+0.4}_{-0.2}$	$1.0^{+0.3}_{-0.3}$
2016 $\mu\mu$	$0.7^{+0.3}_{-0.3}$	$0.7^{+0.3}_{-0.3}$	$0.8^{+0.3}_{-0.3}$
2017+2018 $\mu\mu$	$0.8^{+1.8}_{-0.7}$	$0.4^{+1.0}_{-0.4}$	$1.8^{+0.6}_{-0.7}$

Region C			
	Bkg. simulation without $Z\rightarrow\tau\tau\rightarrow ll$	Bkg. simulation	Data
2016 $e\mu$	$0.8^{+0.4}_{-0.3}$	$0.8^{+0.4}_{-0.3}$	1.0 (0 vs 0)
2017+2018 $e\mu$	$0.8^{+0.3}_{-0.2}$	$0.8^{+0.3}_{-0.2}$	$0.7^{+1.3}_{-0.7}$
2016 ee	$4.0^{+5.8}_{-3.1}$	$4.0^{+5.8}_{-3.1}$	$0.7^{+1.0}_{-0.6}$
2017+2018 ee	$3.5^{+2.6}_{-1.8}$	$2.1^{+2.6}_{-1.5}$	$1.0^{+0.3}_{-0.3}$
2016 $\mu\mu$	$1.2^{+0.5}_{-0.4}$	$1.3^{+0.6}_{-0.4}$	$0.6^{+0.4}_{-0.3}$
2017+2018 $\mu\mu$	$0.4^{+0.4}_{-0.3}$	$0.5^{+0.5}_{-0.3}$	$0.5^{+0.3}_{-0.2}$

10 μm wide. Starting from most-displaced bin, we test to see if any bin has fewer than 5 events, and if so, we combine it with whichever neighboring bin has fewer events, repeating until all bins have at least 5 events.

In each of the two sidebands, we then fit the resulting ratios with a straight line, where the slope and y-intercept are allowed to vary, and extrapolate the fit to 200 μm , which is where we expect the largest contribution from tau lepton decays (see Section 3.5.1). 200 μm also happens to be approximately the center-of-mass of the 100–500 μm bin in background simulation. We average the two extrapolated ratios and derive a correction and systematic uncertainty from this average extrapolated ratio.

If the average extrapolated ratio is > 1.0 , we take the central value as a multiplicative correction to the background estimate and the uncertainty in the average as a systematic uncertainty in the background estimate. In this case, we also vary the 200 μm extrapolation point by $\pm 50 \mu\text{m}$ (the approximate width of the peak in the tau lepton contribution as a function of $|d_0|$). We apply the difference from this variation in extrapolation point as an additional systematic uncertainty in the background estimate. If the average is ≤ 1.0 , we set the correction equal to 1.0 and use the uncertainty in the average as a symmetric systematic uncertainty about 1.0. Table 3.14 shows the resulting correction factors along with the uncorrected and corrected SR I background estimate.

500 μm –10 cm systematic uncertainty

In the 500 μm –10 cm region, we derive a systematic uncertainty in the background estimate from the data closure tests shown in Section 3.5.3. We take the largest deviation from 1.0 that occurs in the ratio of the actual to the estimated number

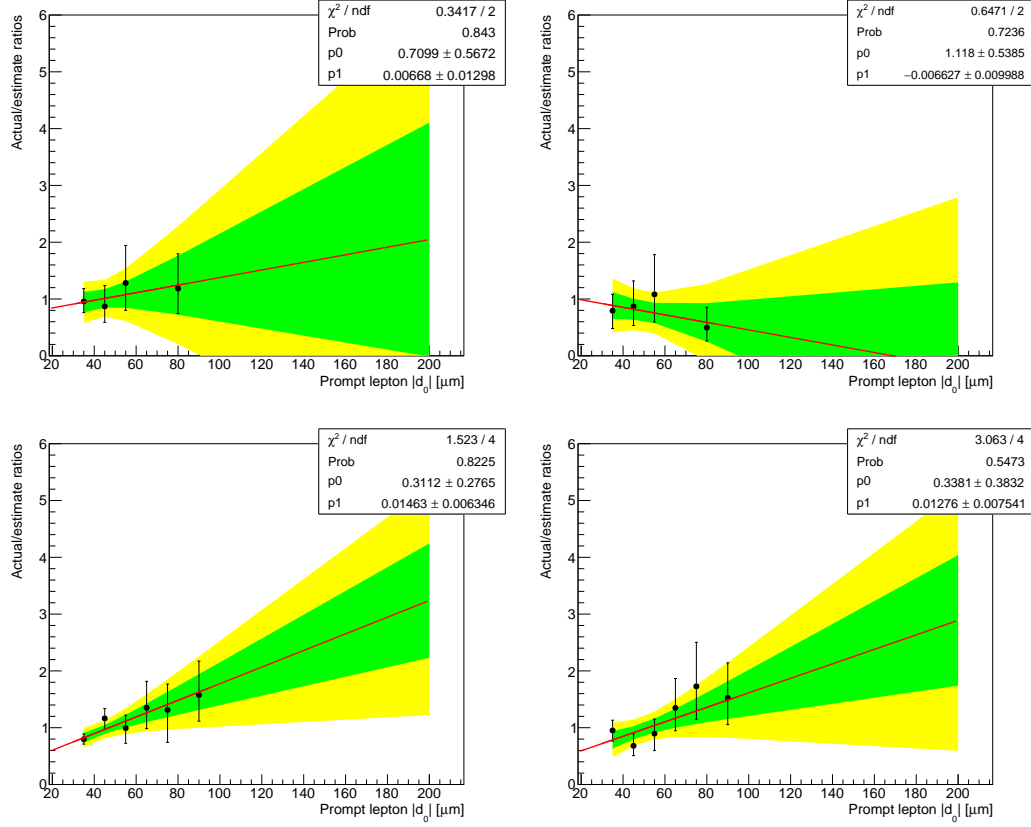


Figure 3.23: Background estimation closure tests in data, in 100–500 μm subregions of regions B (left) and C (right) in the $e\mu$ channel. Each plot shows the ratio of the actual to the estimated number of events as a function of the prompt lepton $|d_0|$ in 2016 (top) and 2017–18 (bottom). A linear fit is shown in black along with the 68 % and 95 % confidence intervals of the extrapolated fit value.

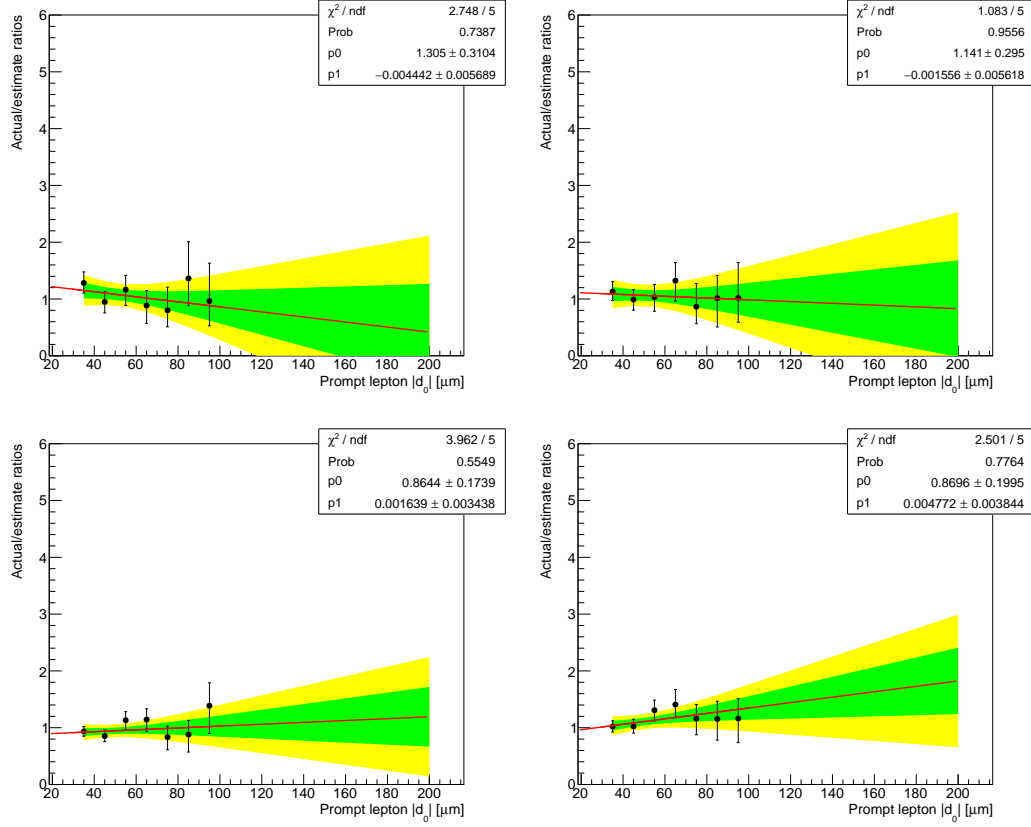


Figure 3.24: Background estimation closure tests in data, in 100–500 μm subregions of regions B (left) and C (right) in the ee channel. Each plot shows the ratio of the actual to the estimated number of events as a function of the prompt lepton $|d_0|$ in 2016 (top) and 2017–18 (bottom). A linear fit is shown in black along with the 68 % and 95 % confidence intervals of the extrapolated fit value.

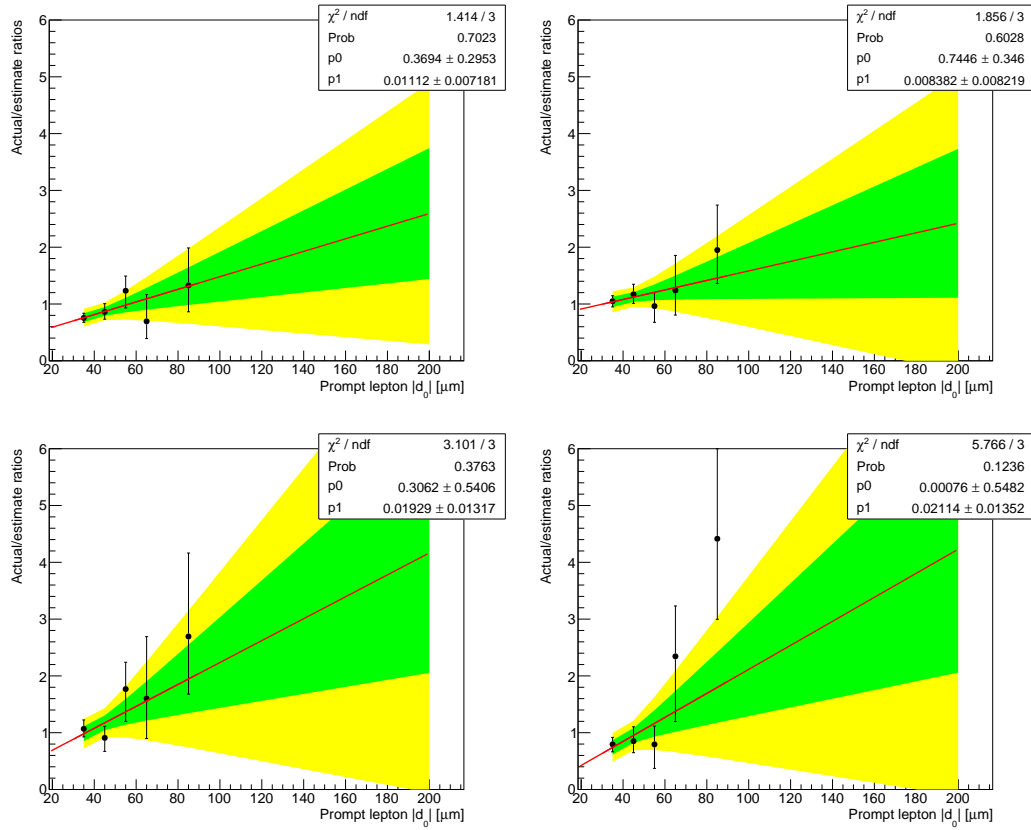


Figure 3.25: Background estimation closure tests in data, in 100–500 μm subregions of regions B (left) and C (right) in the $\mu\mu$ channel. Each plot shows the ratio of the actual to the estimated number of events as a function of the prompt lepton $|d_0|$ in 2016 (top) and 2017–18 (bottom). A linear fit is shown in black along with the 68 % and 95 % confidence intervals of the extrapolated fit value.

Table 3.14: The correction factors and the uncorrected and corrected background estimates in SR I . The correction factor uncertainties include both the uncertainty in the average and the additional uncertainty obtained from varying the fit extrapolation point. The total uncertainty (statistical plus systematic) is given for the corrected background estimates.

	Correction factor	Uncorrected estimate	Corrected estimate
2016 $e\mu$	$1.0^{+1.3}_{-1.0}$	$4.21^{+0.38}_{-0.40}$	$4.2^{+5.4}_{-4.2}$
2017+2018 $e\mu$	3.0 ± 1.0	$12.53^{+0.64}_{-0.61}$	38 ± 13
2016 ee	1.00 ± 0.60	$18.30^{+0.94}_{-0.91}$	18 ± 11
2017+2018 ee	$1.51^{+0.43}_{-0.42}$	41.6 ± 1.3	63^{+18}_{-17}
2016 $\mu\mu$	2.5 ± 1.0	3.07 ± 0.08	7.7 ± 3.1
2017+2018 $\mu\mu$	4.2 ± 1.8	1.00 ± 0.04	4.2 ± 1.8

of events plus its uncertainty, in either of the two closure tests that correspond to a given SR, as a systematic uncertainty. This is a conservative approach that produces a large systematic uncertainty in the small background yields that we predict in these regions. Table 3.15 shows the systematic uncertainty and the predicted number of events in SRs II, III, and IV.

3.5.5 Testing full background estimation procedure

Having defined the full background estimation procedure and seen that the $|d_0^a|$ - $|d_0^b|$ correlation observed in data is also present in simulated background events, we now perform a final closure test of the full background estimation method using simulated background events in SRs I–IV.

Table 3.15: The systematic uncertainty and the background estimates in SRs II, III, and IV. The total uncertainty (statistical plus systematic) is given for each estimate.

	Systematic uncertainty	SR II	SR III	SR IV
2016 $e\mu$	98%	0.15 ± 0.15	$0.09^{+0.12}_{-0.09}$	$0.003^{+0.004}_{-0.003}$
2017+2018 $e\mu$	106%	$0.71^{+0.76}_{-0.71}$	$0.23^{+0.27}_{-0.23}$	$0.01^{+0.02}_{-0.01}$
2016 ee	199%	$0.51^{+1.02}_{-0.51}$	$0.43^{+0.85}_{-0.43}$	$0.01^{+0.02}_{-0.01}$
2017+2018 ee	37%	3.6 ± 1.4	2.8 ± 1.1	$0.24^{+0.10}_{-0.09}$
2016 $\mu\mu$	64%	0.17 ± 0.11	0.19 ± 0.12	0.01 ± 0.01
2017+2018 $\mu\mu$	140%	$0.14^{+0.19}_{-0.14}$	$0.08^{+0.12}_{-0.08}$	$0.01^{+0.02}_{-0.01}$

Table 3.16 shows the estimated and actual number of simulated background events in SRs I–IV. The listed estimates include all corrections and statistical and systematic uncertainties as discussed in 3.5.4. The uncertainties in the actual values are purely statistical. The general agreement between estimated and actual yields leads us to conclude that the background estimation procedure is valid and the assigned systematic uncertainties are sufficient to cover any potential sources of nonclosure that we have not explicitly considered.

3.5.6 Additional background checks

We perform a few additional studies to check for other potential sources of background. We find that their SR contributions are either negligible or already covered by the background estimation method described above.

Table 3.16: Closure test results in background simulation in the SRs with all background estimate corrections and uncertainties applied. The estimated number of events, the actual number of events, and their total uncertainties (statistical plus systematic) are given. In cases where the actual number of events is zero, the uncertainty is given by the product of the average background simulation event weight and the upper bound of the 68% Poisson interval given by a single observation of zero events.

	SR I	SR II	SR III	SR IV
2016 $e\mu$ estimated	$7.4^{+4.8}_{-4.2}$	0.07 ± 0.07	$0.096^{+0.105}_{-0.096}$	0.001 ± 0.001
2016 $e\mu$ actual	$5.0^{+1.5}_{-1.2}$	$0.07^{+0.09}_{-0.05}$	$0.005^{+0.011}_{-0.004}$	$0.000^{+0.037}_{-0.000}$
2017+2018 $e\mu$ estimated	13.5 ± 6.4	$0.37^{+0.40}_{-0.37}$	$0.34^{+0.36}_{-0.34}$	0.02 ± 0.02
2017+2018 $e\mu$ actual	$19.1^{+11.4}_{-7.6}$	$0.52^{+0.41}_{-0.25}$	$0.00^{+0.24}_{-0.00}$	$0.00^{+0.24}_{-0.00}$
2016 ee estimated	9.3 ± 5.0	$0.12^{+0.23}_{-0.12}$	$0.14^{+0.28}_{-0.14}$	$0.002^{+0.004}_{-0.002}$
2016 ee actual	$13.4^{+3.4}_{-2.8}$	$0.15^{+0.19}_{-0.09}$	$1.03^{+1.36}_{-0.67}$	$0.000^{+0.550}_{-0.000}$
2017+2018 ee estimated	18 ± 11	$0.59^{+0.27}_{-0.26}$	$0.45^{+0.21}_{-0.20}$	0.02 ± 0.01
2017+2018 ee actual	$8.2^{+6.5}_{-3.9}$	$0.17^{+0.23}_{-0.11}$	$0.00^{+0.17}_{-0.00}$	$0.00^{+0.17}_{-0.00}$
2016 $\mu\mu$ estimated	1.3 ± 0.6	0.04 ± 0.04	0.03 ± 0.03	0.002 ± 0.002
2016 $\mu\mu$ actual	$3.3^{+1.8}_{-1.2}$	$0.11^{+0.14}_{-0.07}$	$0.06^{+0.14}_{-0.05}$	$0.000^{+0.110}_{-0.000}$
2017+2018 $\mu\mu$ estimated	2.7 ± 1.4	0.04 ± 0.04	0.02 ± 0.02	0.002 ± 0.002
2017+2018 $\mu\mu$ actual	$7.1^{+6.9}_{-3.8}$	$0.00^{+0.15}_{-0.00}$	$0.00^{+0.15}_{-0.00}$	$0.078^{+0.179}_{-0.064}$

Material interactions

In order to further study the material interactions, we invert the preselection criterion that rejects good vertices in the material. In data, we find seven events, across all channels and years, that pass the preselection with this inverted criterion. As shown in Table 3.17, three of these events are in the prompt control region and four are in region B or region C. The lepton vertices in these events coincide with the material as we expect: two are in the beampipe, one is in the pixel detector inner shield, and four are in the first layer of the pixel detector. Even with the material interaction veto inverted, we find no SR events resulting from material interactions and therefore conclude that material interactions are not a significant background after the full selection is applied.

Table 3.17: Some properties of the seven events found in data with the material interactions selection inverted.

Channel, Era	(d_0^a , d_0^b) [μm]	Region	Vertex (x, y, z) [cm]	Material
$e\mu$, 2017C	(-14, -10)	A	(-2.5, 1.4, 6.8)	pixel L1
$e\mu$, 2018D	(46, -14)	A	(0.9, 2.1, 0.1)	beampipe
ee , 2018D	(198, -34)	B	(-1.9, 0.5, 2.7)	beampipe
$\mu\mu$, 2016G	(407, -8)	B	(-1.4, 4.0, 6.3)	pixel L1
$\mu\mu$, 2016G	(-17, -2215)	C	(-2.6, 3.1, 6.6)	pixel L1
$\mu\mu$, 2016H	(2, 0)	A	(-1.6, -3.5, 12)	inner shield
$\mu\mu$, 2017F	(522, -13)	B	(-1.1, -3.0, -7.5)	pixel L1

Cosmic-ray muons

To estimate the SR contribution of cosmic-ray muons, we perform a study in which we invert the Δt and $\cos\alpha$ criteria in the $\mu\mu$ preselection and check how many events

are in the SRs. We find three data events with the criteria inverted (one event per year, all in SR IV). Next, we use NoBPTX data, which is dominated by cosmic-ray muon events, to estimate the efficiency for cosmic-ray muons to pass the Δt and $\cos(\alpha)$ criteria after passing the rest of the $\mu\mu$ preselection criteria. While 3736 NoBPTX data events pass the preselection criteria with the Δt and $\cos(\alpha)$ criteria removed, zero NoBPTX data events pass the full preselection. To conservatively estimate the efficiency, we fluctuate the number of passing events up to 1 and find an efficiency of $1/3736$. We therefore find the approximate upper bound on the SR contribution of cosmic-ray muons to be $3 \times \frac{1}{3736} = 0.0008$, which is negligible compared to the background estimation in each SR.

Heavy-flavor mesons

We perform two studies to estimate an upper limit on the SR contribution of leptons from heavy-flavor mesons. First, we estimate SR yields with a simple ABCD method in 2018 $\mu\mu$ preselection data while additionally requiring at least one medium CSVv2 b -tagged jet [77]. The test is performed in the $\mu\mu$ channel because it contains the smallest relative SR contribution from mismeasurements and should therefore be most sensitive to heavy flavor. As shown in Table 3.18, the background estimates are about an order of magnitude smaller than when no b -tagged jet is required in our usual preselection.

Next, we look at 2018 data and simulated QCD multijet events that pass the $\mu\mu$ preselection with the isolation criterion inverted. These samples are dominated by muons from B meson decays, and the QCD simulation describes the data well in the region outside of the Z boson peak, as shown in Fig. 3.26. We use this QCD multijet sample to test the heavy-flavor background in two ways. First, we perform a simple

Table 3.18: Background estimates in data while applying the 2018 $\mu\mu$ preselection and the additional requirement of at least one b -tagged jet. The estimates with at least one b -tagged jet are about an order of magnitude below the nominal prediction.

	SR I	SR II	SR III	SR IV
Preselection (corrected)	2.6 ± 1.0	$0.09^{+0.12}_{-0.09}$	$0.05^{+0.07}_{-0.05}$	$0.007^{+0.010}_{-0.007}$
Preselection + 1 b jet	0.19 ± 0.03	$0.008^{+0.007}_{-0.004}$	$0.005^{+0.004}_{-0.002}$	$0.0002^{+0.0002}_{-0.0001}$

ABCD estimate in the simulated QCD multijet events to check for $|d_0^a| - |d_0^b|$ correlation. As shown in Table 3.19, we find no evidence of correlation, which indicates that the background estimation already accounts for the heavy-flavor background. Second, we estimate the approximate heavy-flavor background in the SRs by taking the ratio of SR to prompt control region events in QCD multijet simulation from the anti-isolated region and the normalization from the number of simulated QCD multijet events that pass the $\mu\mu$ preselection. Using this approach, we estimate that the heavy-flavor background to be $0.06^{+0.13}_{-0.05}$ events in SR I and $0.0015^{+0.0034}_{-0.0012}$ events in SR IV, which is small relative to the nominal prediction shown in the first row of Table 3.18.

We therefore conclude that the heavy-flavor SR contribution is small and already accounted for in our background estimates.

Low-mass SM hadrons

To estimate an upper limit on the SR contribution of leptons from decays of low-mass SM hadrons, we examine 2018 data and QCD multijet simulation in the $\mu\mu$ channel with both the muon isolation and the ΔR requirements inverted. As shown in Fig. 3.27, this region is dominated by low-mass muon pairs, with clear J/ψ , ψ' ,

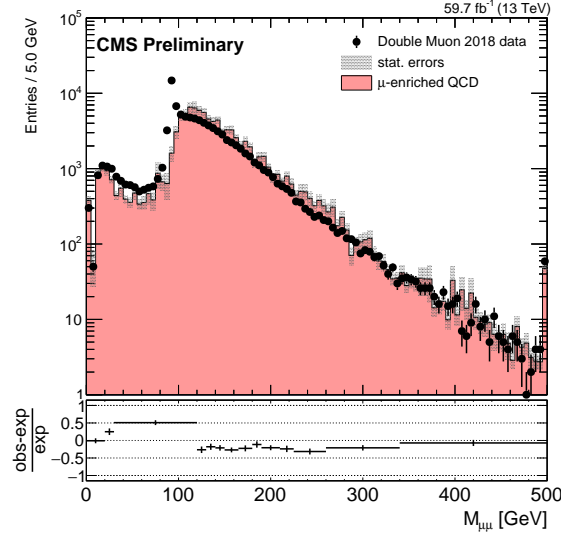


Figure 3.26: The di-muon invariant mass distribution in data and QCD multijet simulation events that pass the 2018 $\mu\mu$ preselection with the muon isolation criterion inverted.

Table 3.19: A closure test of the ABCD method in 2018 QCD simulation in the $\mu\mu$ channel with the muon isolation criterion inverted. The estimates from the ABCD method, the actual yields in simulation, and the ratios of the actual to the estimated yields are shown.

Region	Estimated yield	Actual yield	Ratio of actual to estimate
SR I	9500 ± 1100	11000 ± 1000	1.2 ± 0.2
SR II	1740^{+310}_{-280}	2200^{+330}_{-290}	1.3 ± 0.3
SR III	1450^{+280}_{-240}	1500^{+180}_{-160}	1.0 ± 0.2
SR IV	265^{+62}_{-54}	268^{+61}_{-50}	1.0 ± 0.3

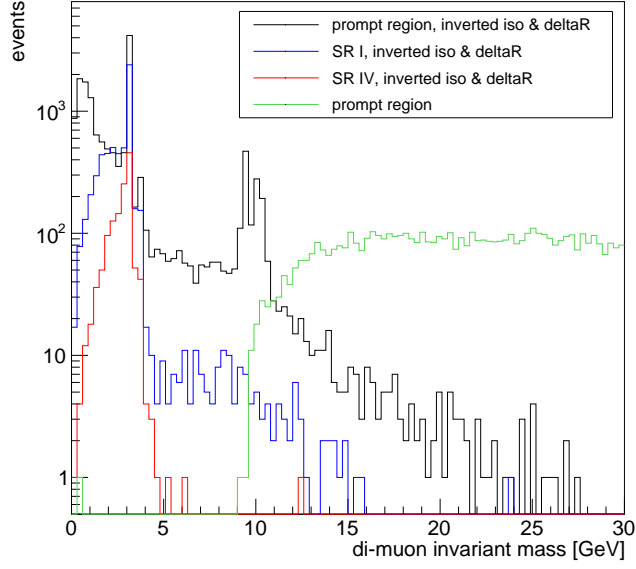


Figure 3.27: The di-muon invariant mass distribution in data events that pass the 2018 $\mu\mu$ prompt control region (black), SR I (blue), and SR IV (red) criteria with the muon isolation and di-muon ΔR criteria inverted. The equivalent distribution in the prompt control region is also shown in green.

and Υ mass peaks. Many of these leptons are displaced, especially those in the J/ψ mass range. To estimate the fraction of such leptons that will be displaced, we take the ratio of SR to prompt control region events of SM hadrons that decay to leptons from data in this region.

Even though the inverted-isolation region is dominated by low-mass muon pairs, the only QCD multijet simulation event that survives the 2018 $\mu\mu$ preselection has a di-muon invariant mass of approximately 300 GeV. Furthermore, the muons are not near each other in the η - ϕ plane ($\Delta R \approx 3$), which is inconsistent with the low-mass SM hadron events that dominate the region with the inverted isolation and ΔR criteria. To find a normalization from which to estimate the low-mass SM hadron

SR contribution, we therefore turn to the inverted-isolation sample used above in the heavy-flavor meson cross check. In this sample, the ratio of events with $\Delta R < 0.5$ to events with $2.8 < \Delta R < 3.2$ is about 0.1. We find 0.2 QCD multijet simulated events that pass the nominal preselection, and so we estimate that of the events passing the 2018 $\mu\mu$ preselection, about 0.02 contain pairs of muons produced in low-mass SM hadron decays. We estimate the SR contributions using this preselection normalization and the ratio of SR to prompt control region events from the sample of SM hadrons that decay to leptons in data. We find this contribution is less than $0.006^{+0.013}_{-0.005}$ events in SR I and less than $0.001^{+0.002}_{-0.001}$ events in SR IV, which, if compared with the nominal prediction shown in the first row of Table 3.26, are respectively negligible and easily covered by the 140 % systematic uncertainty already applied to the background prediction in this region.

3.6 Systematic uncertainties

Several systematic uncertainties are applied to the simulated signal efficiency to account for uncertainty in the signal yields arising from possible mismodeling of the detector conditions and response. The following subsections describe each systematic uncertainties in turn, and Table 3.20 summarizes all systematic uncertainties applied to the simulated signal.

3.6.1 Integrated luminosity

The integrated luminosities of the 2016, 2017, and 2018 data-taking periods are individually known with uncertainties in the 2.3–2.5 % range [78, 79, 80], while the total Run 2 (2016–2018) integrated luminosity has an uncertainty of 1.8 %, the improvement in precision reflecting the uncorrelated time evolution of some systematic effects.

3.6.2 Pileup

The simulation of pileup events assumes a total inelastic proton-proton cross section of 69.2 mb with an associated uncertainty of 5 % [81]. The systematic uncertainty arising as a result of the modeling of pileup events is estimated by varying the cross section of the minimum bias events by 5 % when generating the target pileup distributions. The pileup weights are recomputed with these new distributions and applied to the simulated events to obtain the variation in the yields in the inclusive signal region. The average uncertainty is between 1 and 2 %. We treat these uncertainties as 100 % correlated across the three years of data taking.

Table 3.20: Systematic uncertainties in the signal efficiency for all years and channels. The mean is provided in cases where the uncertainty varies by signal sample. Uncertainties in the same row are treated as correlated among the years of data taking, except for the displaced tracking and muon pixel hit efficiencies, where the 2016 uncertainty is treated as uncorrelated with the 2017 and 2018 uncertainties.

Systematic uncertainty	2016	2017	2018
<i>Integrated luminosity</i>	1.8%	1.8%	1.8%
<i>Pileup</i>			
- $e\mu$ channel	0.5%	0.6%	0.5%
- ee channel	0.5%	0.9%	0.8%
- $\mu\mu$ channel	0.2%	0.1%	0.2%
<i>Displaced tracking efficiency</i>	14%	5.8%	2.4%
<i>Trigger efficiency</i>			
- $e\mu$ channel, electrons	1.6%	1.3%	1.2%
- $e\mu$ channel, muons	1.6%	1.4%	1.2%
- ee channel	10%	13%	19%
- $\mu\mu$ channel	1.2%	1.0%	1.1%
<i>Muon trigger efficiency at large d_0</i>			
- $e\mu$ channel, muons	20%	20%	20%
- $\mu\mu$ channel	20%	20%	20%
<i>Lepton identification and isolation</i>			
- $e\mu$ channel, electrons	1.2%	3.6%	3.5%
- $e\mu$ channel, muons	0.05%	0.07%	0.06%
- ee channel	2.4%	7.2%	7.0%
- $\mu\mu$ channel	0.10%	0.14%	0.12%
<i>Muon pixel hit efficiency</i>			
- $e\mu$ channel, muons	32%	12%	16%
- $\mu\mu$ channel	73%	23%	30%
<i>Lepton d_0 correction</i>			
- $e\mu$ channel, electrons	—	0.001%	0.001%
- $e\mu$ channel, muons	—	0.003%	0.001%
- ee channel	—	0.11%	0.11%
- $\mu\mu$ channel	—	0.11%	0.11%

3.6.3 Displaced tracking efficiency

The systematic uncertainty associated with the modeling of the displaced tracking efficiency is derived from a dedicated study using cosmic-ray muons. Following the results of the study presented in Appendix C, we assign 14.1 %, 5.8 %, and 2.4 % systematic uncertainties in 2016, 2017, and 2018, respectively. Because the CMS pixel detector was upgraded between 2016 and 2017 data taking, we treat the 2017 and 2018 systematic uncertainties as fully correlated and the 2016 systematic uncertainty as uncorrelated with 2017 and 2018.

3.6.4 Trigger efficiency

The trigger efficiency systematic uncertainty is given by the uncertainty in the measured trigger efficiency scale factors (see Section 3.4.4). These uncertainties are 1 % or less for the $e\mu$ and $\mu\mu$ channels and about 10 % for the ee channel. In addition, we have studied the trigger efficiency in signal as a function of $|d_0|$ for events in the trigger p_T plateau. To cover the change observed in the muon trigger efficiency over the full $|d_0|$ range, we assign an additional 20 % uncertainty.

We treat the trigger efficiency uncertainties as 100 % correlated across the three years of data taking.

3.6.5 Lepton ID and isolation

To find the systematic uncertainty associated with the corrections to the lepton ID and isolation, we fluctuate the lepton scale factors up and down by their uncertainty and observe the change in the event yields in the inclusive signal region. The average uncertainty for electrons is about 3 % in the $e\mu$ channel and about 7 % in

the ee channel, while the average uncertainty for muons is $< 1\%$. We treat these uncertainties as 100 % correlated across the three years of data taking.

3.6.6 Muon pixel hit efficiency

The requirement in the muon ID that muons have at least one pixel hit could in principle have some appreciable $|d_0|$ dependence, so we perform a dedicated study to ensure that we account for any differences in $|d_0|$ dependence between data and simulation. Figure 3.28 shows the efficiency of this requirement in cosmic simulation and NoBPTX data (described in appendix C) as a function of muon $|d_0|$. For events in the denominator of these plots, we require that at least 2 global, PF muons have $|\eta| < 1.0$, $p_T > 25\text{ GeV}$, no displaced vertices in the tracker material, $|d_z| < 15\text{ cm}$, and pass all the tight ID criteria except the pixel hit requirement. We also require the muons to be separated by $\Delta R > 0.2$. The events in the numerator must pass the same requirements in addition to the requirement that the muons have at least one pixel hit. From this plot, we calculate the mean efficiency to identify the muons in the simulated signal events using the same procedure as is used for the displaced tracking efficiency systematic uncertainty described in Appendix C. We repeat the procedure in both cosmic simulation and in NoBPTX data, and using the ratio of these two efficiencies, we derive the relative systematic uncertainty in the signal. The average uncertainty is about 16 % (32 %) in the $e\mu$ ($\mu\mu$) channel. As the pixel detector was upgraded after 2016, the 2017 and 2018 systematic uncertainties are treated as fully correlated while the 2016 uncertainty is treated as uncorrelated with the 2017 and 2018 uncertainties.

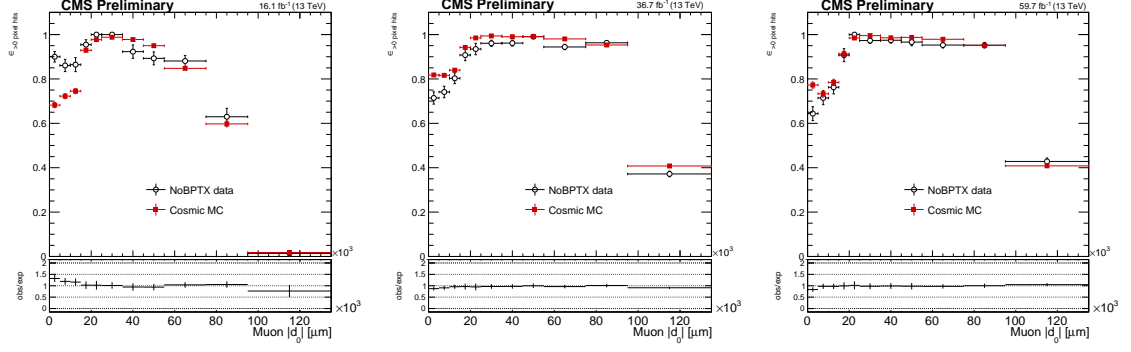


Figure 3.28: The pixel hit efficiency as a function of muon $|d_0|$ in simulated cosmic ray events and NoBPTX data in 2016 (left), 2017 (center), and 2018 (right) conditions.

3.6.7 Lepton d_0 resolution

To find the systematic uncertainty associated with the corrections to the lepton d_0 (see Section 3.4.3), we fluctuate the lepton d_0 corrections up and down by their uncertainty and observe the change in the event yields in the inclusive signal region. The average uncertainty is $< 1\%$. We treat these uncertainties as 100 % correlated in 2017 and 2018. No d_0 correction or systematic uncertainty is needed for 2016 simulation.

3.7 Results

After unblinding and comparing the observed SR yields with the predicted background yields, we find no significant excess. Table 3.21 shows the predicted number of background events and the observed yields in each SR, and Figs. 3.29 and 3.30 visually summarize this same information.

3.7.1 Observed events

In general, the observed SR events appear to be SM events from proton-proton collisions. Specifically, we see no evidence of leptons from cosmic rays, material interactions, or signal. Figure 3.31 shows two-dimensional $|d_0|$ distributions of data events that pass the preselection, and Fig. 3.32 shows the same but for data events in the inclusive SR. As expected for SM background, the observed events are concentrated in the low $|d_0|$ region.

We also use event display visualizations to examine each observed SR event individually and record our observations below.

In the $e\mu$ channel, the SR events tend to have several jets and often have $p_T^{\text{miss}} > 100 \text{ GeV}$. Many events have muon ϕ values such that the muon system hits are all near the edges of detector sections or muon η values such that the muon is near the barrel/endcap transition in the muon system. There are also a few events in which the electron and/or muon are associated with a different primary vertex than their associated track.

In the ee channel, the majority of SR events contain at least one electron with $|\eta| > 1.1$, which increases the probability that their d_0 is poorly measured. Across all three years, most events fall into one of three categories:

Table 3.21: The number of estimated background and observed events in each channel and SR. For each estimate, the total uncertainty is given. The p_T boundaries that separate the low- and high- p_T bins of SR I are listed in Table 3.11.

	SR I, low- p_T bin	SR I, high- p_T bin	SR II	SR III	SR IV
<i>2016 $e\mu$</i>					
- estimated	$3.8^{+4.8}_{-3.8}$	$0.41^{+0.53}_{-0.41}$	$0.09^{+0.12}_{-0.09}$	0.15 ± 0.15	$0.003^{+0.004}_{-0.003}$
- observed	8	1	0	0	0
<i>2017+2018 $e\mu$</i>					
- estimated	38 ± 13	$0.75^{+0.41}_{-0.34}$	$0.23^{+0.27}_{-0.23}$	$0.71^{+0.76}_{-0.71}$	$0.01^{+0.02}_{-0.01}$
- observed	28	3	0	1	0
<i>2016 ee</i>					
- estimated	18 ± 11	$0.22^{+0.17}_{-0.16}$	$0.51^{+1.02}_{-0.51}$	$0.43^{+0.85}_{-0.43}$	$0.01^{+0.02}_{-0.01}$
- observed	40	0	0	1	0
<i>2017+2018 ee</i>					
- estimated	62^{+18}_{-17}	$0.85^{+0.33}_{-0.35}$	2.8 ± 1.1	3.6 ± 1.4	$0.24^{+0.10}_{-0.09}$
- observed	48	0	1	4	0
<i>2016 $\mu\mu$</i>					
- estimated	7.4 ± 3.0	0.25 ± 0.11	0.17 ± 0.11	0.19 ± 0.12	0.01 ± 0.01
- observed	15	0	0	1	0
<i>2017+2018 $\mu\mu$</i>					
- estimated	3.5 ± 1.5	0.69 ± 0.31	$0.08^{+0.12}_{-0.08}$	$0.14^{+0.19}_{-0.14}$	$0.01^{+0.02}_{-0.01}$
- observed	1	1	1	1	0

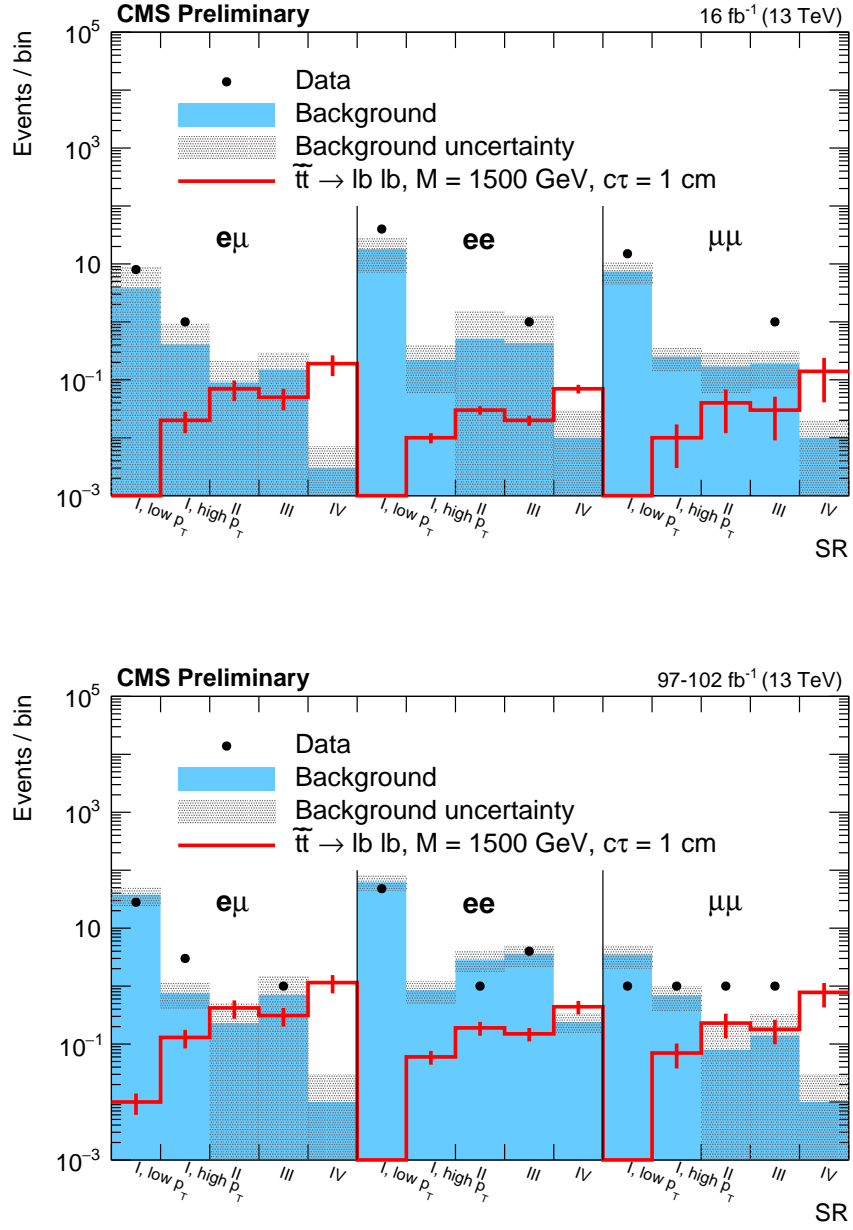


Figure 3.29: The number of estimated background and observed events in each channel and SR, with a representative signal overlaid, for 2016 (top) and 2017–2018 (bottom). For each background estimate and signal yield, the total uncertainty is given.

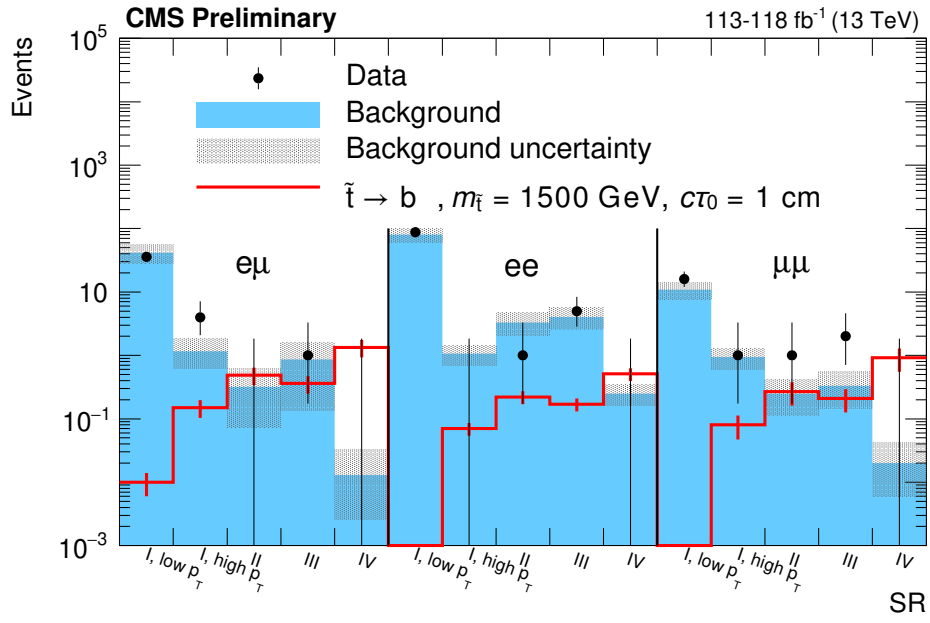


Figure 3.30: The number of estimated background and observed events in each channel and SR, with a representative signal overlaid, for 2016–2018. For each background estimate and signal yield, the total uncertainty is given. Statistical uncertainties are given on observed yields.

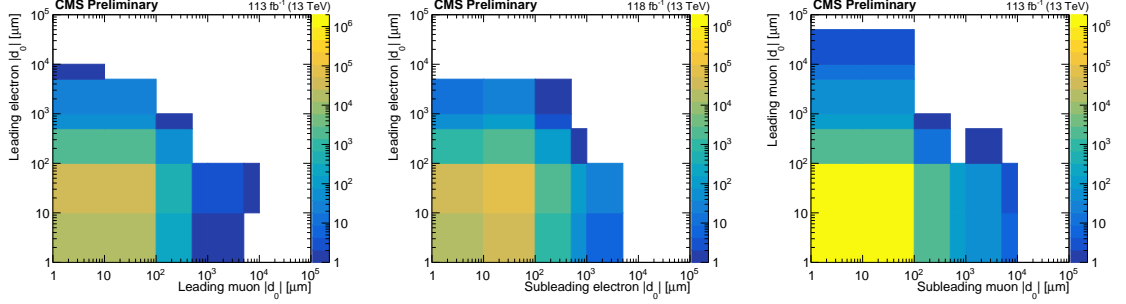


Figure 3.31: Two-dimensional distributions of $|d_0^a|$ and $|d_0^b|$, for the events in data that pass the $e\mu$ (left), ee (center), and $\mu\mu$ (right) preselection. The bins along the x and y axes contain underflow. The inclusive signal region covers the region between 100 μm and 10 cm in each $|d_0|$ variable shown.

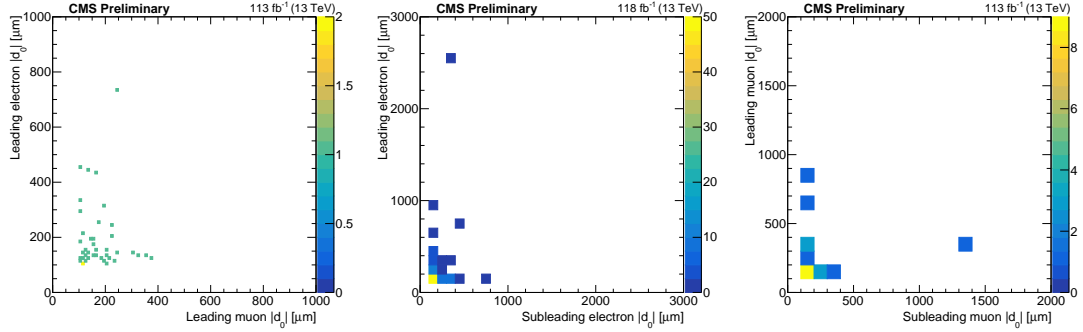


Figure 3.32: Two-dimensional distributions of $|d_0^a|$ and $|d_0^b|$, for data events in the inclusive SR in the $e\mu$ (left), ee (center), and $\mu\mu$ (right) channels.

1. Events with two electrons that appear to be from a boosted Z boson with an invariant mass between 80 and 100 GeV, opposite one or two jets
2. Events with two electrons approximately back-to-back in ϕ with an invariant mass greater than 100 GeV and p_T^{miss} usually between 10 and 40 GeV
3. Events that are similar to type 2 but with at least one jet and frequently p_T^{miss} between 70 and 110 GeV

In the $\mu\mu$ channel, many events have an invariant mass consistent with the mass of the Z boson and p_T^{miss} less than about 60 GeV. Most of the events found in 2017 and 2018 have an invariant mass higher than the Z boson mass and could be $t\bar{t}$ events. Eight of the sixteen SR events in 2016 have two muons with ϕ values of about $\pm\pi/2$ (i.e. in the y - z plane). Every muon pair in these eight events have an invariant mass consistent with a Z boson, and the $\cos(\alpha)$ and Δt distributions of these muons confirm that they are not from cosmic rays. As shown in Fig. 3.33, these features are clearly visible in the ϕ and di-muon invariant mass distributions of the SR data when compared with background simulation. Thirteen of the sixteen muons in these eight events have only one valid pixel hit, and event displays of these events show that the muon track often passes between or at the edge of pixel modules near the place where the two halves of the original pixel detector barrel are joined. We believe that this feature causes the muon d_0 values to be poorly measured.

3.7.2 Limits

The data show no significant excess over background, so we set upper limits on the product of the signal production cross section (σ) and branching fraction (\mathcal{B}) using the HybridNew statistical method of the “Combine” tool developed by the CMS

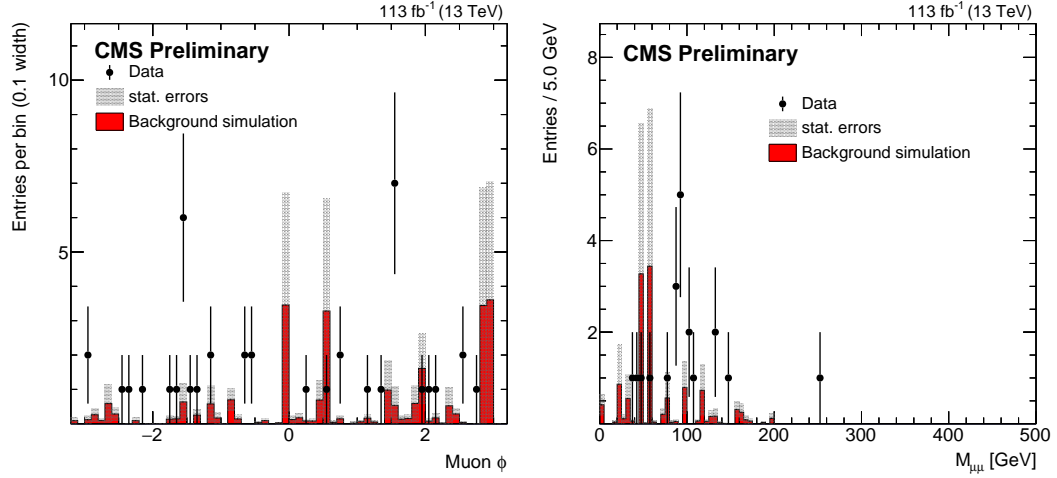


Figure 3.33: Muon ϕ and di-muon invariant mass distributions for data and composite background simulation in the $\mu\mu$ channel in the inclusive SR. The background simulation is normalized to the ABCD method background prediction.

Higgs working group [82, 83, 84, 85]. The ABCD estimate is performed in Combine, which means that any signal contamination in the control regions is accounted for automatically. We perform a simultaneous counting experiment in each signal region bin. Figure 3.34 shows the 95% confidence level (CL) upper limits on the top squark mass as a function of its lifetime.

The variation in the size and shape of the exclusion regions between the three channels is mostly explained by variation in signal yields between the three channels. Looking at the high- p_T bin of SR I, which is the most sensitive bin for top squarks with large masses and small lifetimes, we find that the simulated signal yield is highest in the $e\mu$. This difference between the $e\mu$ and same-flavor channels is a result of simple combinatorics: the two independent top squark decays will result in an $e\mu$ final state twice as often as an ee or $\mu\mu$ final state. In this bin, the ee - and $\mu\mu$ -channel signal yields are similar. SR IV drives the sensitivity for top squarks with large

lifetimes. Because CMS identifies muons with higher efficiency than it does electrons, the $\mu\mu$ channel has the largest simulated signal yield in SR IV when considering top squarks with lifetimes $\gtrsim 10$ cm. For this same reason, the ee channel has the smallest signal yield out of the three channels in SR IV when considering top squark lifetimes $\gtrsim 10$ cm. Taking all of these effects together, we find that the $e\mu$ channel is the most sensitive for lifetimes $\lesssim 10$ cm while the $\mu\mu$ channel is the most sensitive for lifetimes $\gtrsim 10$ cm.

Figure 3.35 shows the 95% CL upper limits for the combination of the three channels. The top squark limits assume $\mathcal{B}(\tilde{t} \rightarrow bl) = \mathcal{B}(\tilde{t} \rightarrow dl) = 100\%$, and each l has an equal probability of being an electron, a muon, or a tau lepton.

3.7.3 Additional likelihood tests

We also perform several statistical tests to help assess the significance of the differences between the observed and predicted SR yields and to ensure that the likelihood is handling the observed yields in a reasonable way.

We first compare the best-fit background yields under the background-only hypothesis while masking the signal regions with the best-fit background yields under the signal+background hypothesis using the full information from all signal and control regions. For simplicity, we refer to the first quantity as the pre-fit prediction and the second as the post-fit prediction. Table 3.22 lists the pre- and post-fit predictions for each channel and SR, and Fig. 3.36 (left) shows associated pull distribution. We find that the differences between the pre- and post-fit predictions are consistent with the variation one would expect from purely statistical effects.

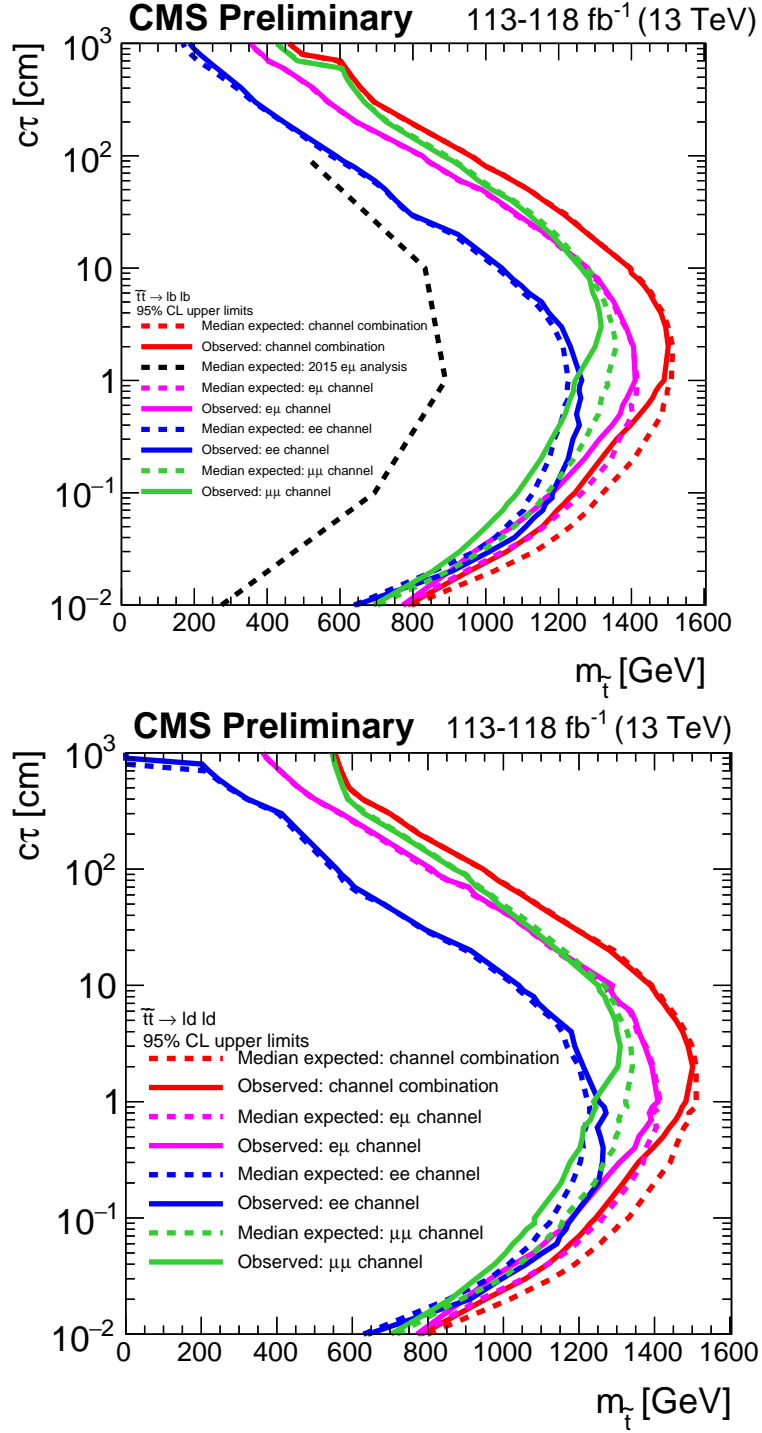


Figure 3.34: The 95% CL upper limits on the top squark mass ($m_{\tilde{t}}$) as a function of its lifetime ($c\tau$), for the $e\mu$, ee , and $\mu\mu$ channels. The $\tilde{t}\tilde{t} \rightarrow \bar{l}b \bar{l}b$ (top) and $\tilde{t}\tilde{t} \rightarrow \bar{l}d \bar{l}d$ (bottom) processes are shown.

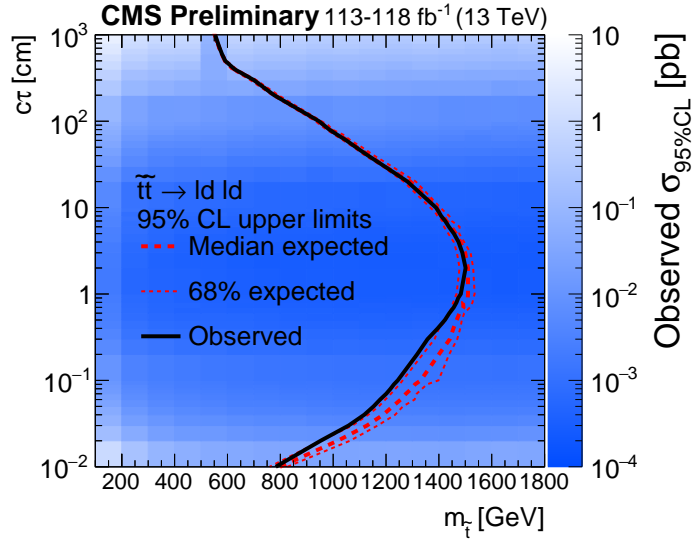
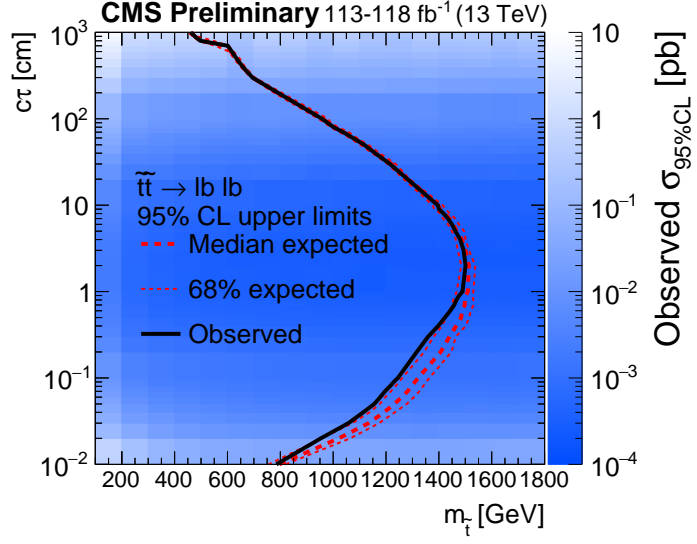


Figure 3.35: The 95% CL upper limits on the top squark mass ($m_{\tilde{t}}$) as a function of its lifetime ($c\tau$). The colors indicate the observed 95% CL upper limit on the cross section. The $\tilde{t}\tilde{t} \rightarrow \bar{l}b l\bar{b}$ (left) and $\tilde{t}\tilde{t} \rightarrow \bar{l}d l\bar{d}$ (right) processes are shown.

Table 3.22: The pre- and post-fit predictions for each signal region bin.

	SR I, low- p_T bin	SR I, high- p_T bin	SR II	SR III	SR IV
<i>2016 $e\mu$</i>					
- pre-fit	3.8 ± 3.9	0.40 ± 0.45	0.09 ± 0.11	0.15 ± 0.13	0.003 ± 0.003
- post-fit	7.1 ± 2.0	0.76 ± 0.31	0.08 ± 0.08	0.14 ± 0.14	0.003 ± 0.003
<i>2017+2018 $e\mu$</i>					
- pre-fit	38 ± 14	0.75 ± 0.40	0.23 ± 0.37	0.71 ± 0.90	0.01 ± 0.02
- post-fit	31 ± 5	0.68 ± 0.25	0.20 ± 0.17	0.65 ± 0.48	0.01 ± 0.01
<i>2016 ee</i>					
- pre-fit	18 ± 11	0.22 ± 0.17	0.51 ± 2.41	0.43 ± 2.06	0.01 ± 0.06
- post-fit	35 ± 5	0.40 ± 0.14	0.50 ± 0.75	0.44 ± 0.53	0.01 ± 0.02
<i>2017+2018 ee</i>					
- pre-fit	62 ± 17	0.85 ± 0.31	2.8 ± 0.9	3.6 ± 1.2	0.25 ± 0.09
- post-fit	50 ± 6	0.65 ± 0.19	2.5 ± 0.7	3.2 ± 0.9	0.22 ± 0.06
<i>2016 $\mu\mu$</i>					
- pre-fit	7.4 ± 3.3	0.25 ± 0.11	0.17 ± 0.11	0.19 ± 0.12	0.01 ± 0.01
- post-fit	11 ± 2	0.37 ± 0.10	0.19 ± 0.10	0.21 ± 0.12	0.01 ± 0.01
<i>2017+2018 $\mu\mu$</i>					
- pre-fit	3.4 ± 1.6	0.69 ± 0.32	0.08 ± 0.12	0.14 ± 0.18	0.01 ± 0.02
- post-fit	2.5 ± 1.1	0.51 ± 0.22	0.14 ± 0.36	0.23 ± 0.63	0.02 ± 0.05

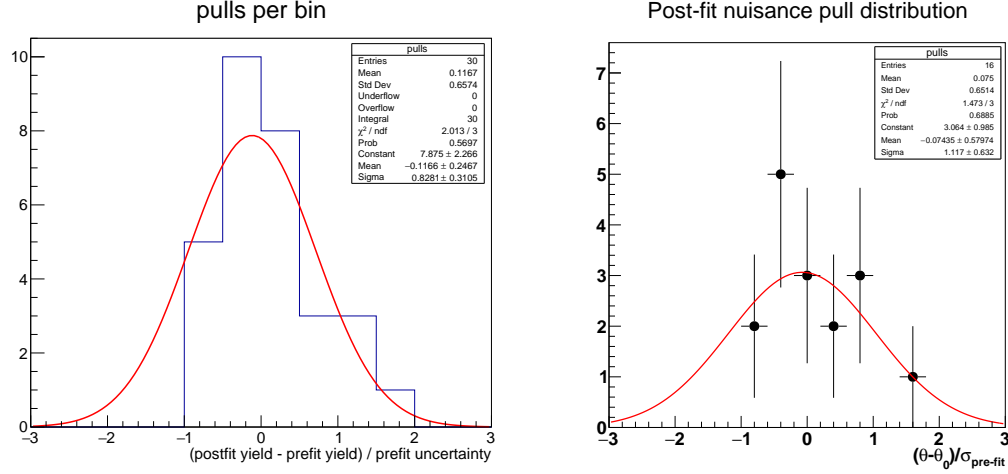


Figure 3.36: The distribution of pulls for each signal region bin, where pulls are calculated as the difference between the post-fit background yield and the pre-fit background yield divided by the pre-fit background uncertainty (left). The distribution of pulls for each background nuisance parameter, where pulls are calculated as the difference between the post-fit value and the pre-fit value divided by the pre-fit uncertainty (right).

Next, we examine the equivalent pull distribution for background yield nuisance parameters. Figure 3.36 (right) shows that the differences in nuisance parameter values before and after the fit are also consistent with the variation one would expect from purely statistical effects.

Finally, we check the observed asymptotic significance of the $\tilde{t}\tilde{t} \rightarrow \bar{l}b\, l\bar{b}$ signal model. As shown in Fig. 3.37, the observed significance is less than two for all signal points we consider when looking at the combination of all channels as well as each channel individually. We therefore conclude that the observed yields do not constitute a significant excess.

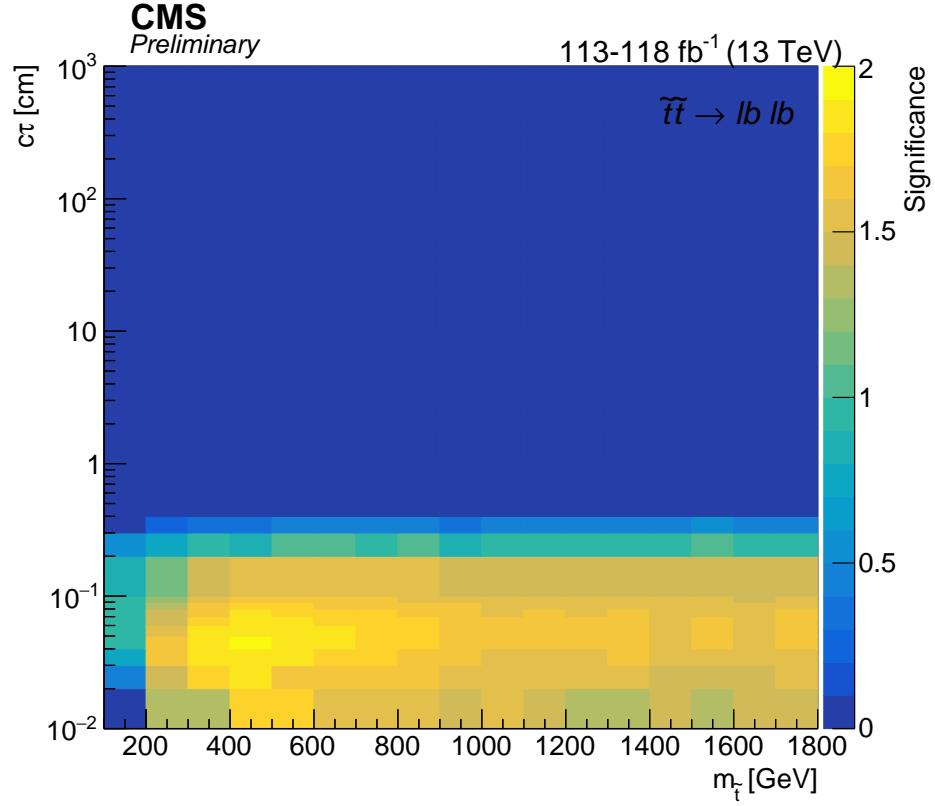


Figure 3.37: The observed asymptotic significances for the $\tilde{t}\tilde{t} \rightarrow \bar{l}b l\bar{b}$ process as a function of \tilde{t} mass and lifetime using the combined results.

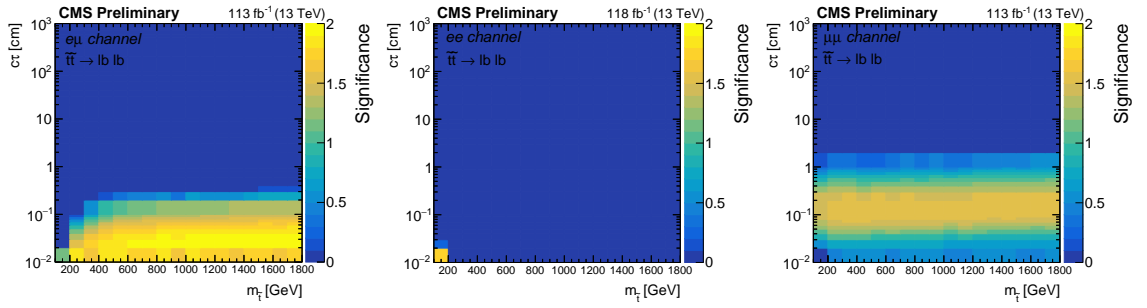


Figure 3.38: The observed asymptotic significances for the $\tilde{t}\tilde{t} \rightarrow \bar{l}b l\bar{b}$ process as a function of \tilde{t} mass and lifetime in the ee (left), $e\mu$ (center), $\mu\mu$ channels (right).

3.8 Outlook

The analysis presented here will likely stand for several years as the most sensitive search in CMS data for new long-lived particles that produce pairs of displaced leptons that are not required to form a common vertex. The LHC is expected to be the highest-energy particle collider in the world for at least another two decades, and it will take several years of data-taking to significantly improve upon the integrated luminosity of the 2016-2018 data-taking period. Assuming that Run 3 of the LHC provides approximately 300 fb^{-1} of 14 TeV proton-proton collisions by the end of 2024 [86, 87], for example, repeating this same analysis would likely only increase the reach in top squark mass reach by a few hundred GeV [88]. On the other hand, there are several potential changes to the analysis strategy that may be worth pursuing.

The most straightforward improvement would be to study the electron and muon identification requirements with an eye to improving the signal efficiency, especially at large $|d_0|$. In particular, the missing inner hit and pixel hit requirements applied to electrons and muons, respectively, in the current analysis effectively limit the maximum LLP decay length to the radius of the CMS pixel detector, which is 16 cm. Any gains in signal efficiency would of course have to be balanced against the likely increase in the mismeasurement background. Thinking along similar lines, it may be interesting to investigate the effects of relaxing the lepton isolation requirement.

A more challenging angle would be to explicitly consider tau leptons in the final state. The analysis presented here is sensitive to displaced taus that decay leptonically to electrons and muons, but a future analysis could likely expand this sensitivity by explicitly studying the $|d_0|$ behavior of displaced taus. Given the tau decay branching

fractions [21], the largest gain would likely come from considering hadronic tau decays, though this route would also likely represent a considerable challenge.

Finally, one could perform an analysis similar to the one presented here but specifically target new low-mass long-lived particles. The lepton p_T requirements imposed by the trigger limit the low-mass sensitivity of the current analysis. One possible approach would be to adopt a different triggering strategy in the next data-taking period, but it may be that CMS has already collected the ideal dataset in which to perform such a search. In 2018, CMS debuted a novel trigger strategy in which specialized triggers collected approximately ten billion unbiased B-hadron-decay events [89]. The triggers use a tag-and-probe strategy that actually requires the presence of at least one displaced muon whose p_T can be as low as 7 GeV. The trade-off is that most of the muons will be embedded in b -tagged jets, which will likely necessitate changes to the analysis strategy. Such a search could be an interesting way to cover new ground with existing data, and may be of particular interest given the growing tension between experiment and SM predictions of lepton universality [90] and the muon anomalous magnetic moment [91].

Chapter 4: Conclusion

A search has been presented for new long-lived particles that propagate a measurable distance before decaying to leptons inside the CMS detector. The resulting displaced lepton signature is targeted by selecting events with two leptons (an electron and a muon, two electrons, or two muons) whose transverse impact parameters are between 0.01 and 10 cm. The search is performed in $113\text{--}118\text{ fb}^{-1}$ of proton-proton collision data produced by the CERN LHC at a center-of-mass energy of 13 TeV and collected by the CMS detector in 2016, 2017, and 2018. This analysis is the first at CMS to target pairs of displaced electrons or muons without requiring that they form a common vertex. The observation is consistent with the background-only hypothesis, and limits are set on the product of the cross section of top squark pair production and the branching fraction to a lepton and a b or d quark through an R-parity-violating vertex. For a proper decay length hypothesis of 2 cm, top squarks with masses up to 1500 GeV are excluded at the 95 % confidence level.

Searches for BSM LLPs are critical to exploring the available new-physics parameter space and ultimately to understanding whether new physics exists at currently accessible energy scales. The analysis presented here explicitly constrains the natural parameter space of RPV SUSY models, but more importantly, it also constrains any not-yet-imagined new physics scenarios that could produce displaced leptons. There

are still many stones unturned, and the analysis presented in this thesis shines a light on one more region of this unexplored space.

Appendix A: Impact of APV25 saturation on displaced tracking

A portion of the data collected by the CMS detector in 2016 is affected by the APV saturation effect described in Ref. [92]. The APV25 saturation effect is a byproduct of the production of heavily ionising particles (HIPs) in inelastic interactions between hadrons and the nuclei of silicon sensors. The energy deposits of HIPs can be up to 1000 times greater than those of typical particles produced in LHC collisions. These large energy deposits can saturate the analog readout of the APV25 chips [93] that are used to read out the CMS silicon strip tracker, which is described in Section 2.2.2. Due to a feature of the APV25 powering scheme that normally helps stabilize the pulse-height baseline, a single saturated channel can inadvertently suppress the outputs of the 127 other APV25 channels for hundreds of nanoseconds.

Only around one in every 1000 incident hadrons will result in saturating HIP events, so the effect is only significant at high instantaneous luminosities. In 2016, the instantaneous luminosity increased to greater than the original LHC design goal, and the effect began influencing detector performance. Starting in run 278802, the tracker front-end electronics were reconfigured to substantially reduce their sensitivity to the APV25 saturation effect.

The deadtime associated with a saturating HIP event can cause some tracker hits to be lost. This effect can lead to reduced tracking efficiency, and it is reasonable to suspect that the loss of efficiency may be more significant for displaced particles that may have fewer tracker hits to begin with. To investigate the impact on displaced tracking, one of our collaborators, Ian Tomalin, performed a study with $K_S^0 \rightarrow \pi^+\pi^-$ decays. From this study, we conclude that the Displaced Leptons analysis should only use data taken after the APV25 saturation effect was mitigated. We therefore use only eras **G** and **H** in 2016 and all available data in 2017 and 2018.

Using data collected in 2016, 2017, and 2018 with the `HLT_ZeroBias` trigger, the reconstructed K_S^0 candidate decay lengths are compared among several different runs that correspond to a wide range of instantaneous luminosities and data-taking periods. The K_S^0 candidates used come from the `generalV0Candidates:Kshort` collection. Each candidate must have a pair of oppositely charged tracks consistent with the K_S^0 mass and coming from a common vertex that is at least 2 cm from the beam line. The tracks are required to have at least one pixel hit and $|\eta| < 2$. In 2016 (2017-18), the tracks are required to have $p_T > 0.7$ GeV (1.5 GeV).

Figure A.1 shows the reconstructed transverse decay length of the K_S^0 candidates for data from all three years. Each distribution is normalized to the integrated luminosity of the run from which it is taken. In the 2016 plot, the solid (dashed) lines correspond to runs taken before (after) the APV25 saturation effect was mitigated. In the pre-mitigation runs, the K_S^0 transverse decay length distribution falls rapidly with increasing instantaneous luminosity, but the dependence on instantaneous luminosity is significantly reduced in the post-mitigation runs and in all 2017 and 2018

runs. The narrower transverse decay length at higher luminosity in the 2016 pre-mitigation runs is interpreted as an instantaneous-luminosity-dependent decrease in tracking efficiency.

Given the size of this effect, we decide not to use 2016 data taken before run 278802 in the Displaced Leptons analysis. The instantaneous-luminosity-dependent displaced tracking efficiency would be difficult to quantify, which would lead to large systematic uncertainties. Furthermore, the signal yield would be suppressed by the lower displaced tracking efficiency. Finally, studies of displaced tracking efficiency with cosmic-ray data are insensitive to the APV25 saturation effect because the instantaneous luminosity during dedicated cosmic runs is zero.

A small dependence on instantaneous luminosity is also apparent in the 2017 and 2018 distributions shown in Fig. A.1. The runs with the lowest (highest) instantaneous luminosity examined in each data-taking period are shown with dashed (solid) lines. This may be due to a residual APV saturation effect or possibly another luminosity-related tracker inefficiency. We do not find it necessary to take any special measures to account for this small effect in 2017 and 2018.

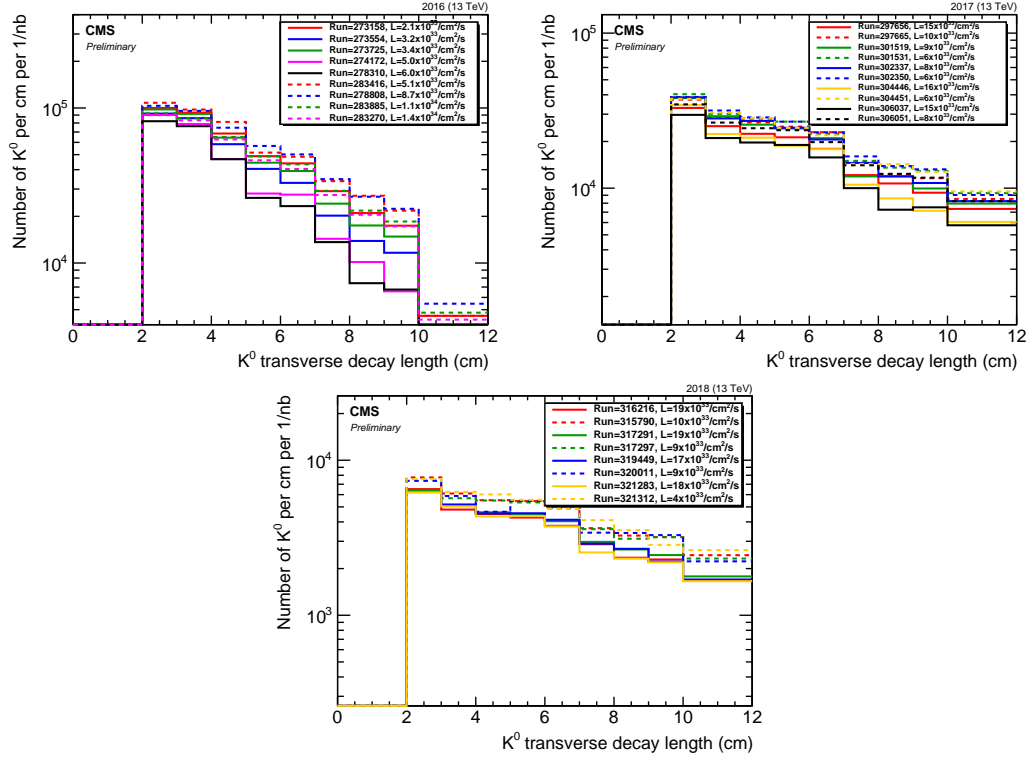


Figure A.1: Transverse decay length distribution of reconstructed $K_S^0 \rightarrow \pi^+\pi^-$ candidates for various runs in 2016 (top left), 2017 (top right) and 2018 (bottom). The peak instantaneous luminosity of each run is indicated in the legend, and each distribution is normalized to the integrated luminosity of the run. In the 2016 plot, runs taken before (after) the APV25 saturation effect was mitigated are shown by solid (dashed) lines. In the 2017 plot, the run with lowest (highest) instantaneous luminosity examined in each data-taking era is plotted with a dashed (solid) line. Towards the end of 2017 (orange and black lines) the LHC ran with approximately 20% fewer proton bunches, meaning that the instantaneous luminosity per bunch crossing was higher than one might naively expect from the instantaneous luminosities shown in the legend.

Appendix B: Poorly measured lepton $|d_0|$ at large $|\eta|$

We require muons to have $|\eta| < 1.5$ due to the observed increase in width of the muon d_0 distribution at large $|\eta|$ in Drell-Yan simulation with $Z \rightarrow \tau\tau \rightarrow ll$ events removed (see Fig. B.1 (left)). The width visibly increases at large $|\eta|$ in all three years but is less pronounced in 2017 and 2018 due to the improved performance of the Phase-1 pixel detector. The upgraded pixel detector is also responsible for the overall difference in d_0 width between years. Requiring muon $|\eta| < 1.5$ has two effects: (1) it dramatically reduces the mismeasurement background in 2016 data in the $\mu\mu$ channel, and (2) it removes a possible source of $|d_0^a|$ - $|d_0^b|$ correlation in which the correlation between muons in η leads to correlation between muons in $|d_0|$. As shown in Fig. B.1 (right), muons from $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events tend to have small $|\eta|$, so requiring muon $|\eta| < 1.5$ has a minimal effect on the signal acceptance.

Electron d_0 resolution also worsens at large $|\eta|$. Furthermore, Fig. B.2 (left) shows that electrons from SM mesons are particularly concentrated $|\eta| > 1.5$. As in the muon case, electrons from $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}b$ events tend to have $|\eta| < 1.5$ (see Fig. B.2), which implies that requiring electron $|\eta| < 1.5$ will reduce the mismeasurement and SM meson backgrounds without significantly reducing signal acceptance.

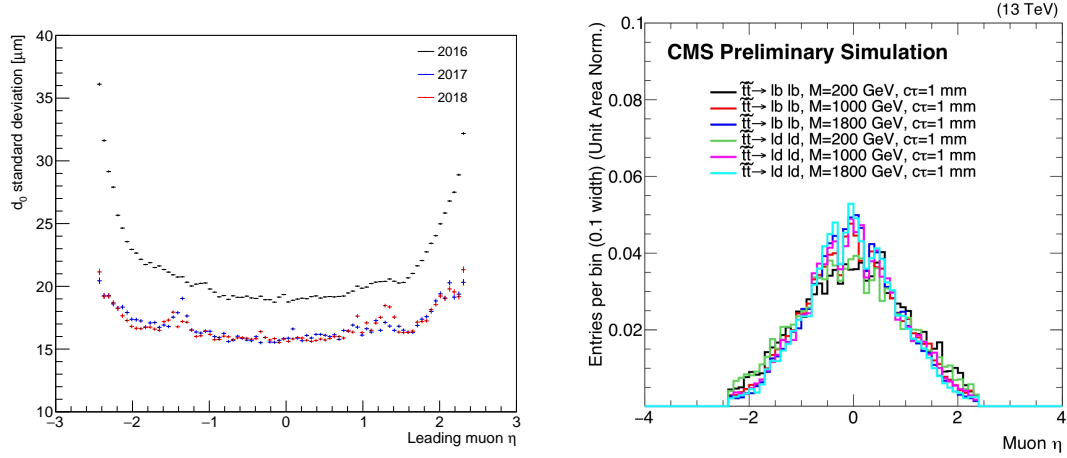


Figure B.1: The standard deviation of the leading muon d_0 as a function of the leading muon η for simulated background events (left). To ensure that the variation in width is purely due to d_0 resolution effects, we use a sample of simulated Drell-Yan events from which the $Z \rightarrow \tau\tau \rightarrow ll$ events have been removed. Muon η distribution for simulated $\bar{t}t \rightarrow \bar{l}b l\bar{b}$ events (right). The $\mu\mu$ preselection with a loosened $|\eta|$ requirement is applied in both plots.

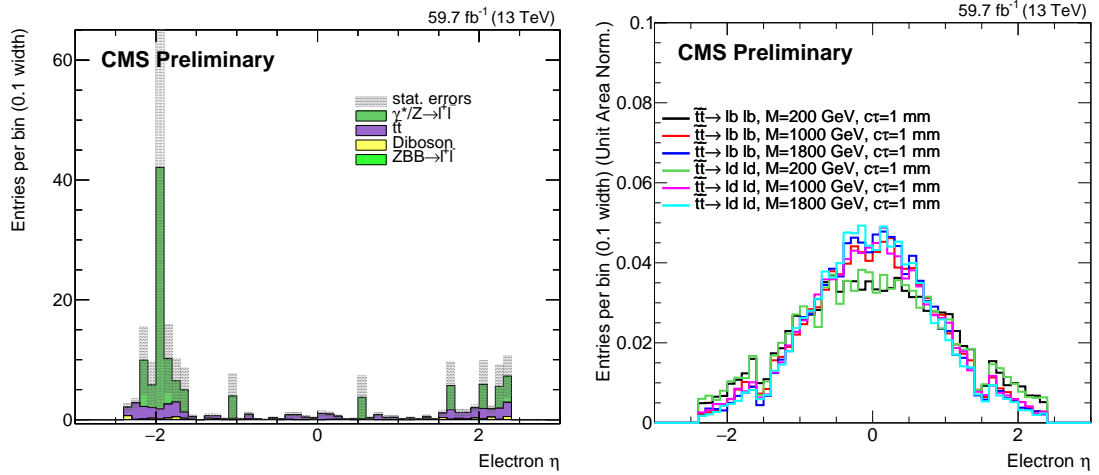


Figure B.2: Electron η distribution for simulated background events in which the electron parent particles are required to be SM mesons (left). Electron η distribution for simulated $\bar{t}t \rightarrow \bar{l}b l\bar{b}$ events (right). The $e\mu$ preselection with a loosened η requirement is applied in both plots.

Appendix C: Displaced tracking efficiency

To measure the efficiency to reconstruct displaced, isolated, high- p_T muons, our collaborator Ian Tomalin performed a study with cosmic rays. The basic idea is to approximate the displaced tracking efficiency in data and simulation from the fraction of cosmic-ray muons reconstructed in the muon system that have a corresponding track in the tracker. The results of this study are used to assign a systematic uncertainty in the signal efficiency (see Section 3.6) and define the upper bound of 10 cm on the inclusive signal region (see Section 3.3).

First, the cosmic-ray dataset is chosen. CMS collects cosmic-ray data in two types of runs: (1) dedicated cosmic runs in which cosmic-ray muons are reconstructed with dedicated reconstruction algorithms and (2) parasitic cosmic runs in which the triggers veto bunch-crossing events to collect cosmic-ray data in otherwise normal proton-proton running conditions. In the parasitic cosmic runs, the cosmic-ray muons are reconstructed with the standard reconstruction algorithms as well as some dedicated reconstruction algorithms. These two types of runs are then collected into two datasets: (1) **Cosmics**, which contains only the dedicated cosmic runs and (2) **NoBPTX**, which contains both the dedicated and parasitic cosmic runs. In each case, the strip tracker electronics operate in the same mode used for proton-proton collisions. The

NoBPTX datasets are used in this study because they include the same reconstruction algorithms as used in the Displaced Leptons analysis.

Events are collected with the `HLTL2Mu10NoVertexNoBPTX3BXv*` trigger, which vetoes events with proton-proton collisions and requires that a muon with $p_T > 10$ GeV is reconstructed in the muon system. As in the Displaced Leptons analysis, the trigger does not explicitly constrain the muon d_0 or d_z . Following the study presented in Appendix A, only eras G and H are used in 2016. The set of data-taking periods with reliable detector performance is identified with particular attention paid to the following properties: (1) suitable cosmic trigger timing configuration, (2) data quality assessed from reconstructed (as opposed to trigger-level) quantities, (3) trigger, tracker, muon system, and track reconstruction known to be functioning well, and (4) magnetic field value in normal range.

To compare the displaced tracking efficiency between data and simulation, simulated cosmic-ray events are produced using the CMSCGEN generator [94]. The simulated cosmic-ray muons have $p_T > 20$ GeV, $|d_0| < 40$ cm, $|d_z| < 80$ cm, and arrive within a 30 ns window centered on the time at which tracker readout efficiency is greatest. The detector response is modeled with GEANT [95].

In both data and simulated events, cosmic rays are reconstructed in the tracker using the same track reconstruction algorithm that is used during proton-proton collisions. This algorithm assumes particles propagate outwards from the center of the detector, so cosmic-ray muons are typically reconstructed as two tracks: one moving upward and one moving downwards. In the muon system, cosmic-ray muons are reconstructed with two dedicated algorithms. The first is a two-leg algorithm that reconstructs each cosmic-ray muon as two separate muons, one in the top half of

CMS and the other in the bottom. The second is a one-leg algorithm that reconstructs cosmic rays as a single muon that traverses the entire detector. The longer lever arm provided by the one-leg algorithm generally results in more accurate measurements of the muon curvature in the magnetic field.

A preliminary event selection is then applied to the data and simulation. The events are required to have one one-leg muon with $p_T > 20 \text{ GeV}$ and at least 50 hits in the muon system. While the p_T requirement of the muons selected in this study is lower than the Displaced Leptons muon p_T requirement, the tracking efficiency does not depend significantly on p_T in the relevant range. The one-leg muon is also required to be within 0.3 rad of two two-leg muons in ϕ in order to reject cosmic-ray muon candidates that do not traverse the entire detector.

The tracking efficiency is inferred from the fraction of selected one-leg muons that are associated with a tracker track that has $p_T > 15 \text{ GeV}$ and is within 0.2 rad in ϕ of the selected one-leg muon. To mimic the Displaced Leptons event selection, tracker tracks are also required to have at least one pixel hit. Because each cosmic-ray muon generally results in two tracker tracks, two separate efficiencies are measured. These efficiencies are referred to as “upward” and “downward” according to the direction of the relevant tracker track. The downward efficiency is expected to be more reliable because the tracker readout electronics assume that particles propagate outward from the center of the detector. The upward and downward efficiencies measured in this study agree within a few percent, but the downward efficiencies are used for the definitive measurement.

The muon system also measures the cosmic-ray muon arrival time. The average of the two two-leg muon arrival times provides the most precise measurement of the

Table C.1: Cosmic-ray muon arrival time requirements used when measuring displaced tracking efficiency in data and simulation.

	2016–2017	2018
Data	$-13 \text{ ns} < t_{\text{muon}} < -7 \text{ ns}$	$-8 \text{ ns} < t_{\text{muon}} < -2 \text{ ns}$
Simulation	$-38 \text{ ns} < t_{\text{muon}} < -32 \text{ ns}$	$-40 \text{ ns} < t_{\text{muon}} < -34 \text{ ns}$

cosmic-ray muon arrival time. This approach provides a time resolution on the order of 5 ns, whereas the time resolution provided by one-leg muons can be an order of magnitude greater.

The selected one-leg muons with $|d_0| < 8 \text{ cm}$ and $|d_z| < 20 \text{ cm}$ are used to study the tracking efficiency dependence on cosmic-ray muon arrival time. Figure C.1 shows the distribution of the measured arrival time while Fig. C.2 shows the measured downward tracking efficiency as a function of measured arrival time. The disagreement between data and simulation in measured arrival time is simply an artifact of the specified simulation time window. Despite this offset, the efficiency shows a clear peak with a width of approximately 15 ns in both data and simulation. Based on the results of Fig. C.2, muons used in the downward tracking efficiency measurement must also pass the requirements listed in Table C.1. Using the upward efficiency produces similar results that are shifted by approximately 5 ns. The same sample of one-leg muons is also used to measure the downward tracking efficiency as a function of run number. No meaningful dependence is identified.

The downward tracking efficiency is then measured using all selected one-leg muons that meet the timing requirements specified in Table C.1. The downward tracking efficiency as a function of $|d_0|$ and $|d_z|$ is shown for all three years in Fig. C.3. The

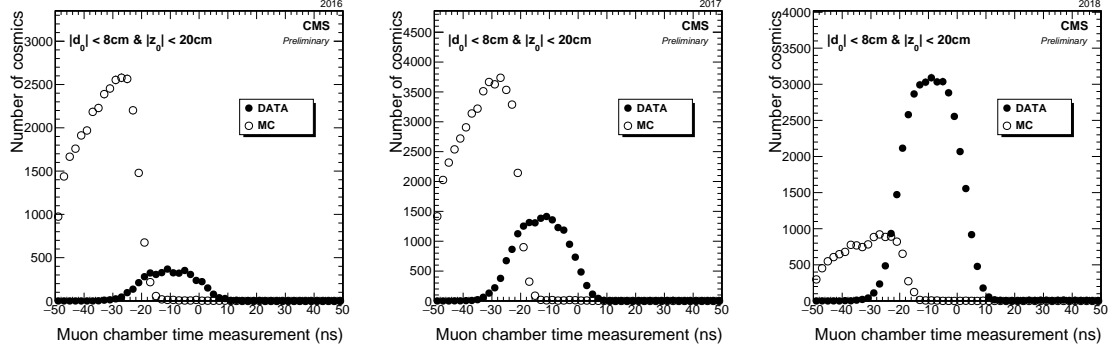


Figure C.1: Distribution of the arrival time of cosmic rays at their point of closest approach to the beamline as measured by the muon system in 2016 (left), 2017 (center), and 2018 (right) in data and simulation. Only cosmic-ray muons with $|d_0| < 8$ cm and $|z_0| < 20$ cm are considered.

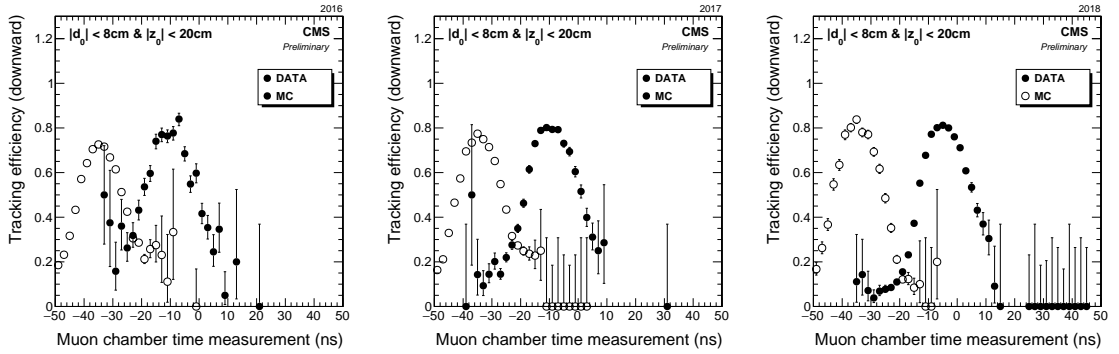


Figure C.2: Measured downward tracking efficiency versus cosmic ray muon arrival time in 2016 (left), 2017 (center), and 2018 (right) in data and simulation. Only cosmic-ray muons with $|d_0| < 8$ cm and $|z_0| < 20$ cm are considered.

tracking efficiency is found to be nonzero out to at least $|d_0| \leq 10$ cm and $|d_z| \leq 30$ cm in all three years. We also note that removing the pixel-hit requirement increases this range to approximately $|d_0| \leq 30$ cm and $|d_z| \leq 50$ cm.

The results of Fig. C.3 are used to estimate a systematic uncertainty in the simulated signal efficiency arising from the displaced tracking efficiency. A simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b \bar{l}b$ sample with a top squark mass of 1800 GeV and proper decay length $c\tau = 100$ cm is considered as it produces leptons with the largest impact parameters and therefore represents the most challenging scenario for the displaced track reconstruction. To accommodate the pixel hit requirements, only those events in which both top squarks decay within the volume of the pixel detector are considered. First, the $|d_0|$ and $|d_z|$ of both leptons in this subset of signal events are noted. Next, a two-dimensional plot of tracking efficiency as a function of $|d_0|$ and $|d_z|$, $\epsilon(|d_0|, |z_0|)$, is produced from the cosmic-ray muons in data and simulated cosmic-ray events. Using this plot, the mean efficiency to reconstruct both lepton tracks in the simulated signal events is evaluated as:

$$\frac{1}{N} \sum_i \epsilon(|d_0|_i^{(1)}, |z_0|_i^{(1)}) \epsilon(|d_0|_i^{(2)}, |z_0|_i^{(2)}) \quad (\text{C.1})$$

where the sum extends over the N events in the signal sample and the superscripts (1) and (2) denote the two leptons in each event. For each year, the relative systematic uncertainty in the efficiency to reconstruct both lepton tracks is then taken from the ratio of the efficiencies in data and simulation. The resulting efficiencies and systematic uncertainties are listed in Table C.2.

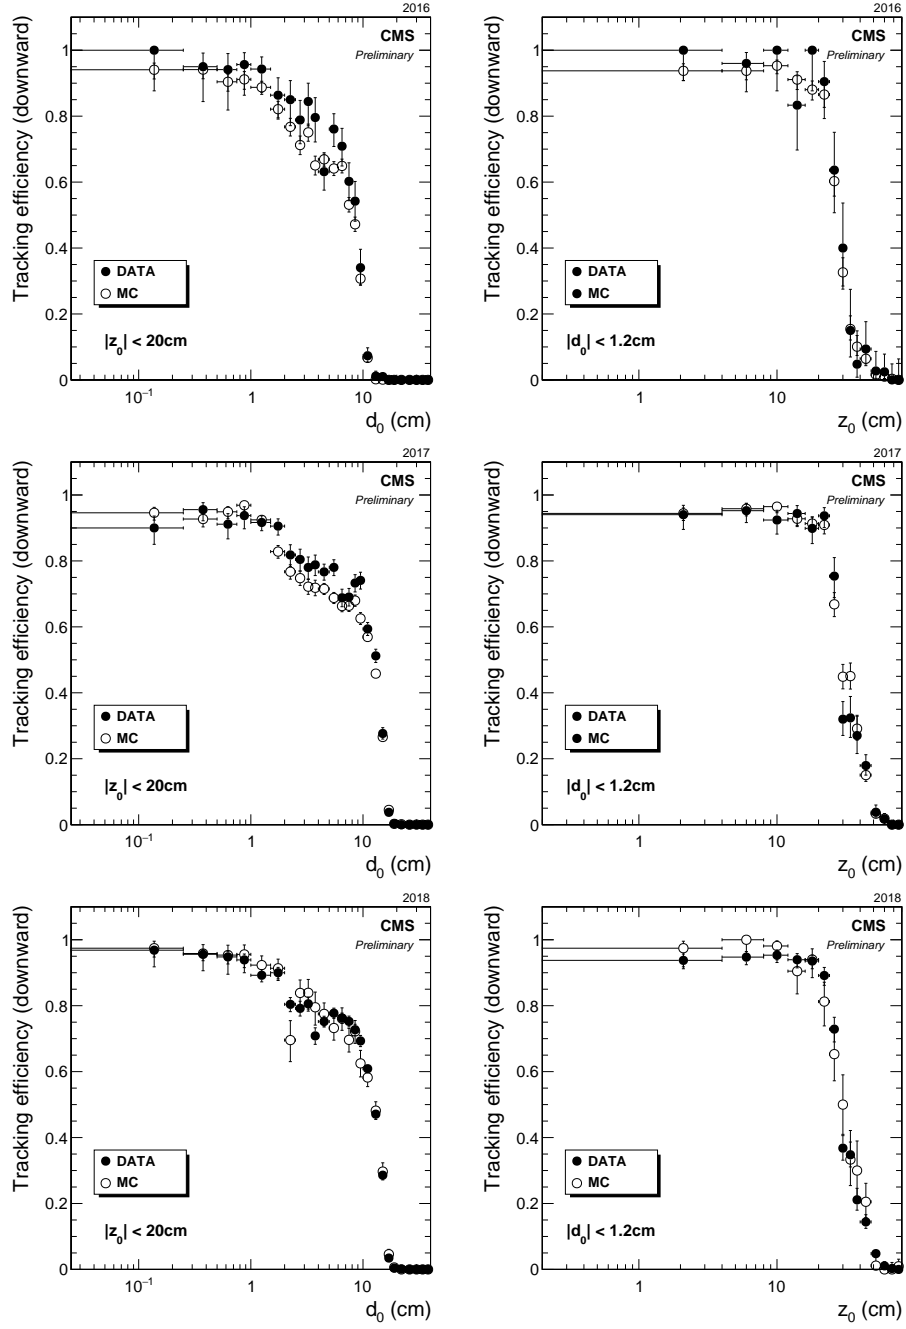


Figure C.3: Measured downward tracking efficiency in 2016 (top), 2017 (middle), and 2018 (bottom) versus $|d_0|$ (left) and $|d_z|$ (right) in data and simulation. $|d_z|$ ($|d_0|$) impact parameter is constrained to less than 20 cm (1.2 cm) when plotting against $|d_0|$ ($|d_z|$).

Table C.2: Mean measured efficiency to reconstruct both lepton tracks in simulated $\tilde{t}\tilde{t} \rightarrow \bar{l}b\bar{l}\bar{b}$ events in data and simulation and the resulting systematic uncertainty. The top squark mass and proper decay length are assumed to be 1800 GeV and 100 cm.

	2016	2017	2018
Efficiency in data	$57.5 \pm 2.1\%$	$55.3 \pm 1.1\%$	$56.1 \pm 0.7\%$
Efficiency in simulation	$50.3 \pm 1.0\%$	$52.3 \pm 0.7\%$	$57.5 \pm 1.1\%$
Systematic uncertainty	$14.1 \pm 4.3\%$	$5.8 \pm 2.3\%$	$2.4 \pm 2.2\%$

Appendix D: Study of R-hadron treatment in simulation

In the $pp \rightarrow \tilde{t}\tilde{t}$ signal processes, the top squarks can form strongly produced hadronic states called “R-hadrons.” In the $\tilde{t} \rightarrow bl$ and $\tilde{t} \rightarrow dl$ signal samples used in this analysis, R-hadron formation and decay are modeled with PYTHIA while the R-hadron interactions with material are not modeled. Given that the lepton identification criteria effectively limit the number of interaction lengths an R-hadron can traverse before decaying, we expect the impact of such interactions to be negligible. To test this hypothesis, we generate alternative signal samples in which the R-hadron material interactions and decay are modeled in GEANT4 and compare the resulting signal efficiency to that of the nominal samples. We find differences in signal efficiency on the order of 20% that are completely attributable to differences in the treatment of jet formation between the PYTHIA and GEANT4 implementations and therefore conclude that the R-hadron material interactions do not meaningfully impact the analysis.

In the alternative samples, the R-hadron interactions are modeled following the “cloud model”, which assumes that the top squark is surrounded by a cloud of colored, light constituents that interact during scattering [69, 70]. Such interactions can potentially alter the rate of R-hadron energy loss or cause the R-hadron charge value to change mid-flight. In the context of CMSSW, simultaneously enabling R-hadron

interactions with material and R-hadron decay within the detector volume constrains us to perform both the R-hadron propagation and decay with GEANT4. We generate two additional sets of signal samples for this study. In the first, the R-hadron decays to a lepton and bottom quark. Because GEANT4 does not perform the parton shower and hadronization that lead to jet formation, the bottom quark is in this case unrealistically treated as a final-state particle. To help study the effect of this missing jet formation, we generate a second sample in which the R-hadron decays to a lepton and a neutral pion, which then decays to two photons and initiates an electromagnetic shower.

Figure D.1 shows how modeling the R-hadron interactions and decay in GEANT4 affects the inclusive signal region cutflows in the ee and $\mu\mu$ channels. The alternative samples generally pass the full selection at a slightly higher rate than the existing signal samples. The overall magnitude of the effect is similar across channels and signal points, and a few clear trends are apparent in the cutflows. First, events from the neutral pion sample pass the ee trigger at a higher rate than other samples because of the increased number of high-momentum photons and electrons from the neutral pion decay, an effect that carries through the full selection. Second, $\tilde{t} \rightarrow bl$ events pass the $\mu\mu$ trigger at a higher rate due to the increased number of muons from b jets, but this effect is neutralized by the muon isolation criteria. Third, events in which the R-hadron decays to a lepton and a final-state b quark in GEANT4 pass the ee trigger at lower rates due to the reduced number of electrons from jets but pass the lepton quality and isolation criteria at higher rates because the absence of jets increases the probability that leptons from the top squark decay are well isolated. A similar effect

can be seen in the neutral pion sample in the $\mu\mu$ cutflow. The effects on electrons and muons identified in the same-flavor channels are also apparent in the $e\mu$ channel.

The stated causes of the effects described above are supported by the kinematic distributions of electrons, muons, and jets in each of these samples before and after the inclusive signal region selections are applied. For example, Fig. D.2 shows that before any selection is applied, events in which the R-hadron decays to a lepton and a neutral pion contain a clear relative excess of high- p_T , isolated electrons; $\tilde{t} \rightarrow bl$ events contain a clear relative excess of high- p_T , non-isolated muons; and events in which the R-hadron decays to a lepton and a non-physical final-state b quark (and to a lesser degree, events in which the R-hadron decays to a lepton and a neutral pion) show a distinct lack of central, high- p_T jets.

Figure D.3 shows how the effects described above carry through to the events that pass the inclusive signal region selections. As expected, the events in which the R-hadron decays to a lepton and a neutral pion have a relative excess of high- p_T , well-isolated electrons when the full ee signal region selection is applied. The muon isolation cut eliminates the differences between the two samples in which the top squark is decayed in PYTHIA when the $\mu\mu$ inclusive signal region selection is applied, but the relative excess of well-isolated muons from the two other samples becomes more pronounced. Finally, the stark differences in jet activity remain apparent.

In the end, modeling the R-hadron interactions and decay with GEANT4 has an approximately 20% effect on the signal efficiency that does not vary meaningfully with analysis channel, top squark mass, or top squark lifetime. We find that all observable differences are due to known, non-physical effects of the modeling of the R-hadron decay and therefore conclude that our nominal signal samples provide the

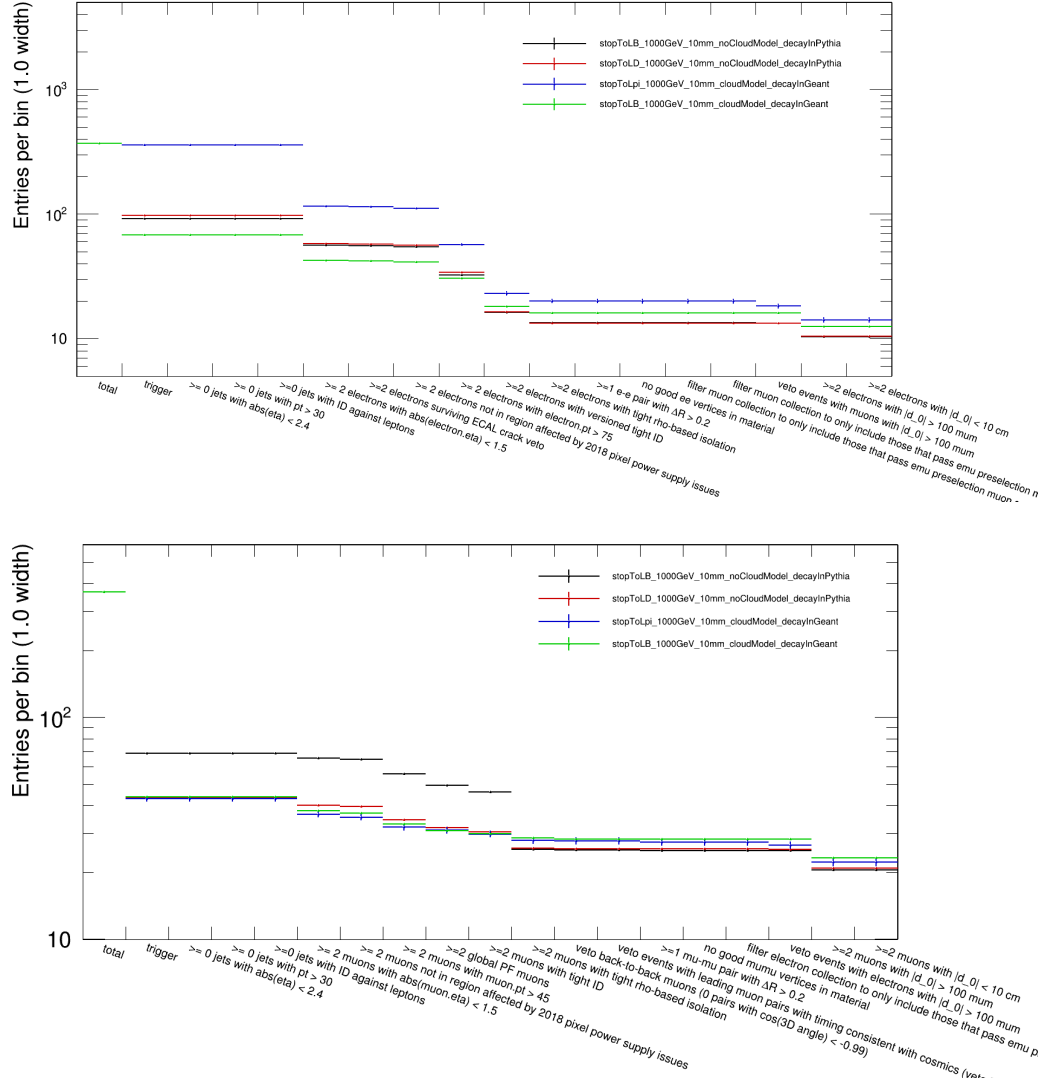


Figure D.1: Inclusive signal region cutflows for four signal samples in the ee (top) and $\mu\mu$ (bottom) channels. The top squark mass and proper decay length are 1000 GeV and 1 cm. In the sample corresponding to the black (red) curves, R-hadron material interactions are not modeled, but the top squark decay is performed in PYTHIA and the resulting b (d) quark produces a jet. In the samples corresponding to the blue and green curves, the R-hadron material interactions and decay are modeled with GEANT4. In the sample corresponding to the blue curves, the R-hadron decays to a lepton and a neutral pion, and in the sample corresponding to the green curves, the R-hadron decays to a lepton and a non-physical final-state quark.

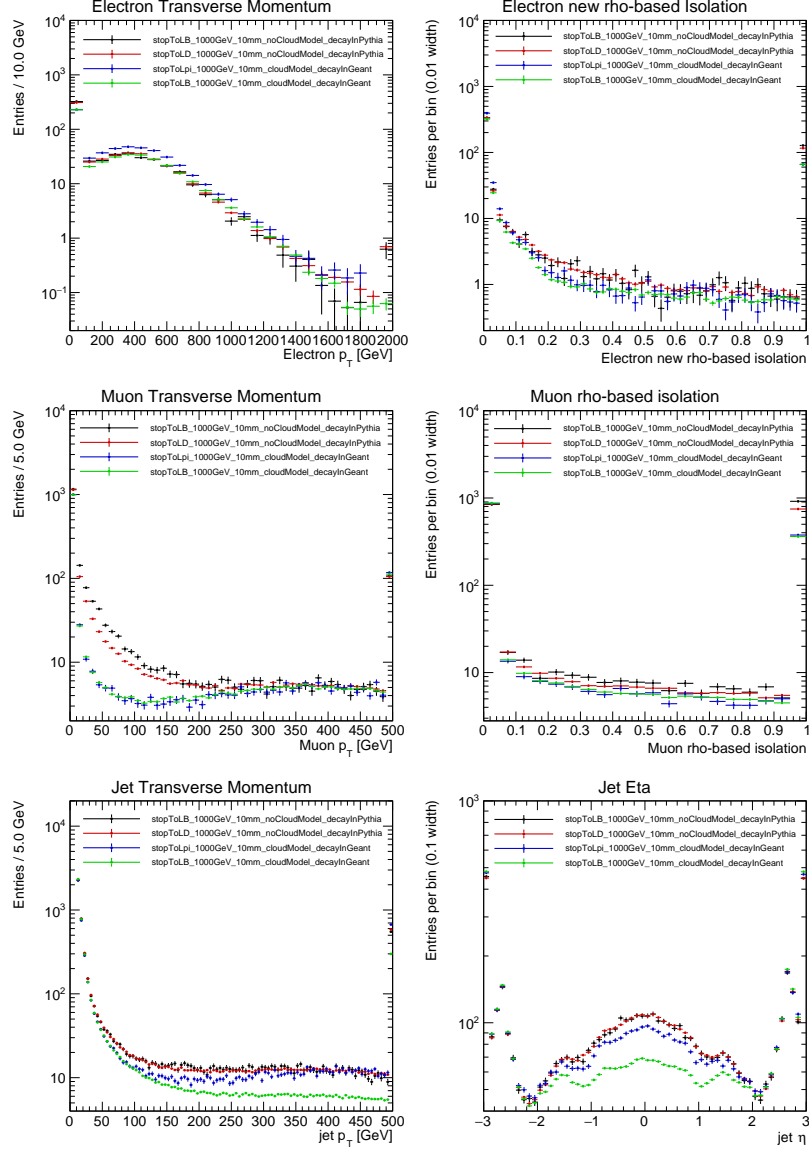


Figure D.2: Kinematic distributions of electrons (top), muons (middle), and jets (bottom) before any selection is applied for four signal samples in which the top squark mass and proper decay length are 1000 GeV and 1 cm. In the sample corresponding to the black (red) curves, R-hadron material interactions are not modeled, but the top squark decay is performed in PYTHIA and the resulting b (d) quark produces a jet. In the samples corresponding to the blue and green curves, the R-hadron material interactions and decay are modeled with GEANT4. In the sample corresponding to the blue curves, the R-hadron decays to a lepton and a neutral pion, and in the sample corresponding to the green curves, the R-hadron decays to a lepton and a non-physical final-state quark. In each plot, the rightmost bin contains overflow entries.

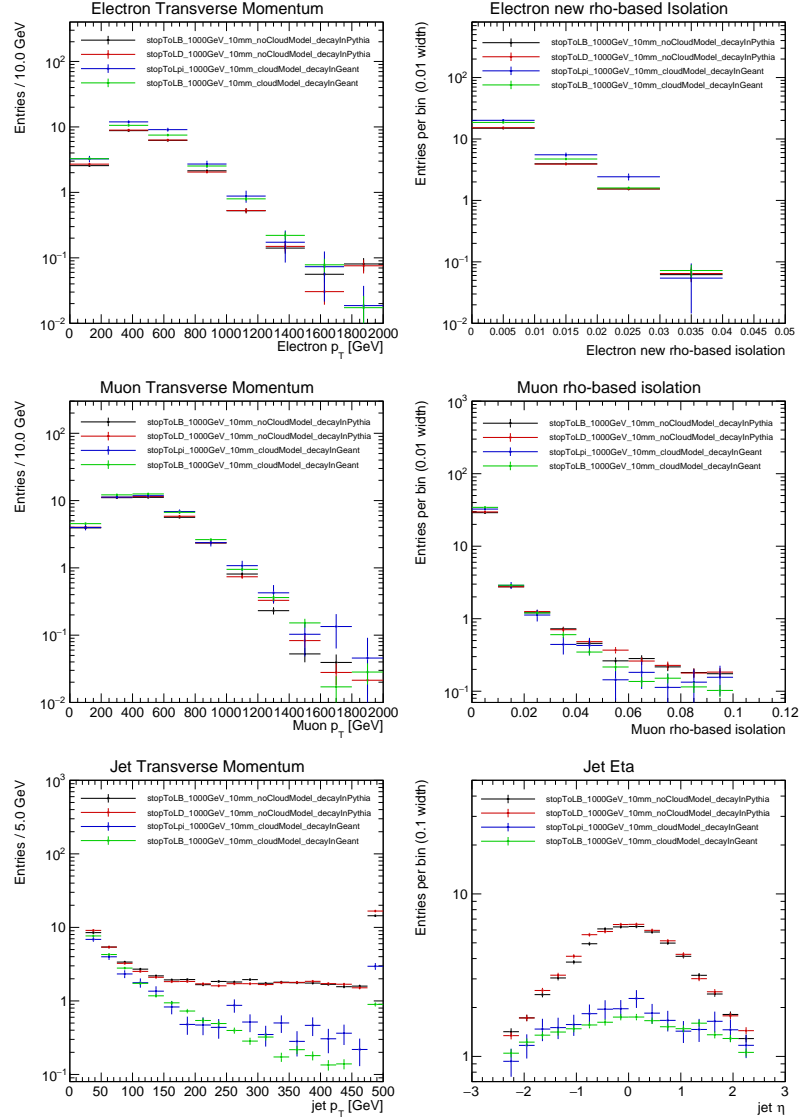


Figure D.3: Kinematic distributions of electrons in the ee signal region (top), muons in the $\mu\mu$ signal region (middle), and jets in the $\mu\mu$ signal region (bottom) for four signal samples in which the top squark mass and proper decay length are 1000 GeV and 1 cm. In the sample corresponding to the black (red) curves, R-hadron material interactions are not modeled, but the top squark decay is performed in PYTHIA and the resulting b (d) quark produces a jet. In the samples corresponding to the blue and green curves, the R-hadron material interactions and decay are modeled with GEANT4. In the sample corresponding to the blue curves, the R-hadron decays to a lepton and a neutral pion, and in the sample corresponding to the green curves, the R-hadron decays to a lepton and a non-physical final-state quark. In each plot, the rightmost bin contains overflow entries.

best estimate of the signal efficiency and any real effects from R-hadron interactions must be comparatively small and therefore insignificant to the analysis as a whole.

Bibliography

- [1] D.J. Gross and F. Wilczek, *Asymptotically Free Gauge Theories - I*, *Phys. Rev. D* **8** (1973) 3633.
- [2] H. Fritzsch, M. Gell-Mann and H. Leutwyler, *Advantages of the Color Octet Gluon Picture*, *Phys. Lett. B* **47** (1973) 365.
- [3] S. Weinberg, *A Model of Leptons*, *Phys. Rev. Lett.* **19** (1967) 1264.
- [4] J. Goldstone, A. Salam and S. Weinberg, *Broken Symmetries*, *Phys. Rev.* **127** (1962) 965.
- [5] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321.
- [6] P.W. Higgs, *Broken symmetries, massless particles and gauge fields*, *Phys. Lett.* **12** (1964) 132.
- [7] G.S. Guralnik, C.R. Hagen and T.W.B. Kibble, *Global Conservation Laws and Massless Particles*, *Phys. Rev. Lett.* **13** (1964) 585.
- [8] CMS collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *PLB* **716** (2012) [1207.7235].
- [9] ATLAS collaboration, *Observation of a new particle in the search for the standard model Higgs boson with the ATLAS detector at the LHC*, *PLB* **716** (2012) [1207.7214].
- [10] CMS collaboration, *Combined measurements of Higgs boson couplings in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) [1809.10733].
- [11] ATLAS collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, *Phys. Rev. D* **101** (2020) 012002.

- [12] “Summaries of CMS cross section measurements.”
`twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsCombined`.
- [13] M. Fukugita and P.J.E. Peebles, *The Cosmic energy inventory*, *Astrophys. J.* **616** (2004) 643 [[astro-ph/0406095](#)].
- [14] G.F. Giudice, *Naturally Speaking: The Naturalness Criterion and Physics at the LHC*, 0801.2562.
- [15] G. 't Hooft, *Naturalness, chiral symmetry, and spontaneous chiral symmetry breaking*, *NATO Sci. Ser. B* **59** (1980) 135.
- [16] N. Arkani-Hamed, S. Dimopoulos and G.R. Dvali, *The Hierarchy problem and new dimensions at a millimeter*, *Phys. Lett. B* **429** (1998) 263 [[hep-ph/9803315](#)].
- [17] S.P. Martin, *A Supersymmetry primer*, *Adv. Ser. Direct. High Energy Phys.* **18** (1998) 1 [[hep-ph/9709356](#)].
- [18] N. Craig, *The State of Supersymmetry after Run I of the LHC*, in *Beyond the Standard Model after the first run of the LHC*, Sep, 2013 [[1309.0528](#)].
- [19] “CMS Supersymmetry Physics Results.”
`twiki.cern.ch/twiki/bin/view/CMSPublic/PhysicsResultsSUS`.
- [20] B. Shuve, “Theory Overview of Long-Lived Particles at the LHC.”
`indico.cern.ch/event/607314/contributions/2542308`.
- [21] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *PTEP* **2020** (2020) 083C01.
- [22] Z. Liu and B. Tweedie, *The Fate of Long-Lived Superparticles with Hadronic Decays after LHC Run 1*, *JHEP* **06** (2015) 042 [[1503.05923](#)].
- [23] W. Kilian, T. Plehn, P. Richardson and E. Schmidt, *Split supersymmetry at colliders*, *Eur. Phys. J. C* **39** (2005) 229 [[hep-ph/0408088](#)].
- [24] F.E. Paige and J.D. Wells, *Anomaly mediated SUSY breaking at the LHC*, in *1st Les Houches Workshop on Physics at TeV Colliders*, 6, 1999 [[hep-ph/0001249](#)].
- [25] P.W. Graham, D.E. Kaplan, S. Rajendran and P. Saraswat, *Displaced Supersymmetry*, *JHEP* **07** (2012) [[1204.6038](#)].
- [26] L.J. Hall and M. Suzuki, *Explicit R-Parity Breaking in Supersymmetric Models*, *Nucl. Phys. B* **231** (1984) 419.
- [27] L. Evans and P. Bryant, eds., *LHC Machine*, *JINST* **3** (2008) S08001.

- [28] “Linear accelerator 2.” cds.cern.ch/record/1997427, Sep, 2012.
- [29] “The Proton Synchrotron Booster.” cds.cern.ch/record/1997372, Jul, 2012.
- [30] “Proton Synchrotron.” cds.cern.ch/record/1479637, 1959.
- [31] “The Super Proton Synchrotron.” cds.cern.ch/record/1997188, Jan, 2012.
- [32] E. Mobs, *The CERN accelerator complex - 2019*, Jul, 2019.
- [33] J. Wenninger, *Operation and Configuration of the LHC in Run 2*, Mar, 2019.
- [34] AC Team, “Diagram of an LHC dipole magnet.” cds.cern.ch/record/40524, Jun, 1999.
- [35] CMS collaboration, *CMS technical design report, volume II: Physics performance*, *J. Phys. G* **34** (2007) 995.
- [36] “CMS collaboration.” cms.cern/collaboration.
- [37] CMS collaboration, T. Sakuma, *3D SketchUp images of the CMS detector (120918)*, May, 2016.
- [38] CMS collaboration, G.L. Bayatian et al., *CMS physics: Technical design report volume I: Detector performance and software*, 2006.
- [39] CMS collaboration, *The CMS Experiment at the CERN LHC*, *JINST* **3** (2008) S08004.
- [40] The Tracker Group of the CMS collaboration, *The CMS Phase-1 Pixel Detector Upgrade*, *JINST* **16** (2021) [2012.14304].
- [41] “Run II Pixel Performance plots for data and simulation.” twiki.cern.ch/twiki/bin/view/CMSPublic/PixelOfflinePlots2016.
- [42] CMS collaboration, D.A. Matzner Dominguez et al., eds., *CMS Technical Design Report for the Pixel Detector Upgrade*, .
- [43] CMS collaboration, *Performance and operation of the CMS Phase 1 pixel detector*, *Nucl. Instrum. Meth. A* **936** (2019) 688.
- [44] CMS collaboration, *Performance of the CMS muon detector and muon reconstruction with proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JINST* **13** (2018) P06015 [1804.04528].
- [45] CMS collaboration, *The CMS RPC Detector Performance and Stability during LHC RUN-2*, *JINST* **14** (2019) C11012 [1808.10488].

- [46] CMS collaboration, *The CMS trigger in Run 2*, Tech. Rep. CMS-CR-2017-340, CERN, Geneva (Oct, 2017), DOI.
- [47] CMS collaboration, *The CMS Level-1 trigger system for LHC Run II*, *JINST* **12** (2017) C03021.
- [48] CMS collaboration, *Particle-flow reconstruction and global event description with the CMS detector*, *JINST* **12** (2017) [1706.04965].
- [49] W. Adam, R. Frühwirth, A. Strandlie and T. Todorov, *Reconstruction of electrons with the gaussian-sum filter in the CMS tracker at the LHC*, *Journal of Physics G: Nuclear and Particle Physics* **31** (2005) N9.
- [50] CMS collaboration, *Search for Displaced Supersymmetry in events with an electron and a muon with large impact parameters*, *Phys. Rev. Lett.* **114** (2015) 061801 [1409.4789].
- [51] CMS collaboration, *Search for displaced leptons in the e-mu channel*, Tech. Rep. CMS-PAS-EXO-16-022 (2016).
- [52] ATLAS collaboration, *Search for displaced leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, 2011.07812.
- [53] Image courtesy of Jamie Antonelli.
- [54] J. Alimena et al., *Searching for long-lived particles beyond the Standard Model at the Large Hadron Collider*, *J. Phys. G* **47** (2020) 090501 [1903.04497].
- [55] CMS collaboration, *Search for disappearing tracks in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Lett. B* **806** (2020) 135502 [2004.05153].
- [56] CMS collaboration, *Search for long-lived particles using delayed photons in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Rev. D* **100** (2019) 112003 [1909.06166].
- [57] CMS collaboration, *Search for long-lived particles using nonprompt jets and missing transverse momentum with proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Phys. Lett. B* **797** (2019) 134876 [1906.06441].
- [58] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) [1405.0301].
- [59] R. Frederix and S. Frixione, *Merging meets matching in MC@NLO*, *JHEP* **12** (2012) [1209.6215].

- [60] J. Alwall et al., *Comparative study of various algorithms for the merging of parton showers and matrix elements in hadronic collisions*, *Eur. Phys. J. C* **53** (2008) [0706.2569].
- [61] S. Frixione and B.R. Webber, *Matching NLO QCD computations and parton shower simulations*, *JHEP* **06** (2002) [hep-ph/0204244].
- [62] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, *JHEP* **11** (2004) [hep-ph/0409146].
- [63] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, *JHEP* **11** (2007) [0709.2092].
- [64] S. Alioli, P. Nason, C. Oleari and E. Re, *NLO vector-boson production matched with shower in POWHEG*, *JHEP* **07** (2008) [0805.4802].
- [65] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, *JHEP* **06** (2010) [1002.2581].
- [66] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) [1410.3012].
- [67] CMS collaboration, *Event generator tunes obtained from underlying event and multiparton scattering measurements*, *Eur. Phys. J. C* **76** (2016) [1512.00815].
- [68] CMS collaboration, *Extraction and validation of a new set of CMS PYTHIA8 tunes from underlying-event measurements*, *EPJC* **80** (2020) [1903.12179].
- [69] R. Mackeprang and A. Rizzi, *Interactions of coloured heavy stable particles in matter*, *Eur. Phys. J. C* **50** (2007) 353 [hep-ph/0612161].
- [70] R. Mackeprang and D. Milstead, *An updated description of heavy-hadron interactions in GEANT-4*, *Eur. Phys. J. C* **66** (2010) 493 [0908.1868].
- [71] J. Alwall et al., *A standard format for Les Houches event files*, *Comput. Phys. Commun.* **176** (2007) [hep-ph/0609017].
- [72] B. Allanach et al., *The Snowmass points and slopes: Benchmarks for SUSY searches*, *EPJC* **25** (2002) [hep-ph/0202233].
- [73] CMS collaboration, *Performance of Electron Reconstruction and Selection with the CMS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV*, *JINST* **10** (2015) P06005 [1502.02701].

- [74] CMS collaboration, *Precision measurement of the structure of the CMS inner tracking system using nuclear interactions*, *JINST* **13** (2018) P10034 [1807.03289].
- [75] CMS collaboration, *Precision measurement of the structure of the CMS inner tracking system using nuclear interactions with data collected in 2018*, Tech. Rep. CMS-DP-2019-001 (Feb, 2019).
- [76] J.A. Evans and J. Shelton, *Long-lived staus and displaced leptons at the LHC*, *JHEP* **04** (2016) [1601.01326].
- [77] CMS collaboration, *Identification of b quark jets at the CMS Experiment in the LHC Run 2*, Tech. Rep. CMS-PAS-BTV-15-001 (2016).
- [78] CMS collaboration, *CMS luminosity measurements for the 2016 data taking period*, CMS Physics Analysis Summary CMS-PAS-LUM-17-001 (2017).
- [79] CMS collaboration, *CMS luminosity measurements for the 2017 data-taking period at $\sqrt{s} = 13$ TeV*, CMS Physics Analysis Summary CMS-PAS-LUM-17-004 (2018).
- [80] CMS collaboration, *CMS luminosity measurements for the 2018 data-taking period at $\sqrt{s} = 13$ TeV*, CMS Physics Analysis Summary CMS-PAS-LUM-18-002 (2018).
- [81] CMS collaboration, *Measurement of the inelastic proton-proton cross section at $\sqrt{s} = 13$ TeV*, *JHEP* **07** (2018) [1802.02613].
- [82] T. Junk, *Confidence level computation for combining searches with small statistics*, *Nucl. Instrum. Meth. A* **434** (1999) [hep-ex/9902006].
- [83] A.L. Read, *Presentation of search results: the CL_s technique*, in *Durham IPPP Workshop: Advanced Statistical Techniques in Particle Physics*, (Durham, UK), Mar, 2002, DOI.
- [84] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) [1007.1727].
- [85] ATLAS collaboration, CMS collaboration, LHC Higgs Combination Group, *Procedure for the LHC Higgs boson search combination in Summer 2011*, Tech. Rep. CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11 (2011).
- [86] B.A. Petersen and C. Schwick, *Experiment Requests and Constraints for Run 3*, in *9th LHC Operations Evian Workshop*, 2019.

- [87] “Longer term LHC schedule.”
lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm.
- [88] G. Salam and A. Weiler, “Collider Reach.” collider-reach.web.cern.ch.
- [89] CMS collaboration, *Recording and reconstructing 10 billion unbiased b hadron decays in CMS*, *EPJ Web Conf.* **245** (2020) 01025.
- [90] LHCb collaboration, *Test of lepton universality in beauty-quark decays*,
2103.11769.
- [91] MUON G-2 collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm*, *Phys. Rev. Lett.* **126** (2021) 141801 [2104.03281].
- [92] W. Adam, W. Beaumont, T. Bergauer, O. Bouhali, B. Clerbaux, E. De Langhe et al., *The effect of highly ionising particles on the CMS silicon strip tracker*, *Nucl. Instrum. Methods Phys. Res., A* **543** (2005) 463.
- [93] M.J. French, L.L. Jones, Q.R. Morrissey, A. Neviani, R. Turchetta, J.R. Fulcher et al., *Design and results from the APV25, a deep sub-micron CMOS front-end chip for the CMS tracker*, *Nucl. Instrum. Methods Phys. Res., A* **466** (2001) 359.
- [94] P.A. Biallass, T. Hebbeker and K. Hoepfner, *Simulation of Cosmic Muons and Comparison with Data from the Cosmic Challenge using Drift Tube Chambers*, Tech. Rep. CMS-NOTE-2007-024 (9, 2007).
- [95] GEANT4 collaboration, *GEANT4—a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250.