

The LHCb GPU high level trigger and measurements of neutral pion and photon production with the LHCb detector

by

Thomas J. Boettcher

B.S., Indiana University (2015)

Submitted to the Department of Physics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2021

© Massachusetts Institute of Technology 2021. All rights reserved.

Author
.....

Department of Physics
January 4, 2021

Certified by
.....

Michael Williams
Associate Professor
Thesis Supervisor

Accepted by
.....

Deepto Chakrabarty
Associate Department Head



The LHCb GPU high level trigger and measurements of neutral pion and photon production with the LHCb detector

by

Thomas J. Boettcher

Submitted to the Department of Physics
on January 4, 2021, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

In Run 3 of the LHC, the LHCb experiment will collect data without a hardware trigger. As a result, the LHCb software trigger must process the full LHC collision rate. This thesis presents Allen, an implementation of the LHCb high level trigger on GPUs. The Allen framework and reconstruction algorithms are described and the physics and computing performance are presented.

This thesis also presents studies of neutral pion and photon production in proton-lead collisions. The LHCb detector's forward geometry allows it to study these processes in a previously unexplored kinematic regime. Measurements of the π^0 nuclear modification factor at forward and backward rapidities are presented. These results provide constraints on the cold nuclear matter effects that modify the partonic structure of nuclei. Additionally, measurements of the direct photon excess ratio in proton-proton and proton-lead collisions are presented.

Thesis Supervisor: Michael Williams

Title: Associate Professor

Acknowledgments

First, I want to thank my advisor, Prof. Michael Williams. Mike treated me as a colleague and collaborator from my first day at MIT. He gave me the freedom to pursue my own ideas, and he was always ready with guidance and encouragement when I needed it. I learned as much from dropping into his office unannounced as I did from any class in grad school.

I feel lucky to have been part of the MIT LHCb group for the last five years. This group is full of kind and talented people working on an astounding range of physics. I want to thank Phil Ilten for his admirable patience and guidance as I found my footing during my first couple of years of grad school. I also want to thank Dan Craik, who was a constant reassuring presence during the late nights and early mornings we spent working on Allen. I'm also thankful for my colleagues in the LHCb collaboration. I'd especially like to thank Roel Aaij, Dorothea vom Bruch, Daniel Hugo Cámpora Pérez, and Vladimir Gligorov for bringing me into the Allen project. I'm extremely proud of what we've accomplished together.

I'm grateful for the friendship of my fellow grad students in the MIT LHCb and GlueX groups. I want to thank Constantin Weisser for helping me face the challenges of a year at CERN. I want to thank Yunjie Yang, who was a constant source of positivity on long days. I want to thank John Hardin, who may have graduated a year earlier had I not so frequently distracted him. I also want to thank the past and present students of the Building 26 Penthouse. Their company kept me going as we struggled through courses, prepared for quals, and placed prop bets on departmental colloquia. Finally, I want to thank the students of LNS more broadly. Because of them, LNS has been a great place to work for the last five years.

I'm lucky to have made too many friends throughout grad school to thank by name here, but I do want to thank my roommate, Graeme Sutcliffe. Grad school is much easier when you have a friend who's always ready to go to karaoke, make pierogies, or play a game of unihoc. When we were assigned as roommates first year of grad school, I had no idea he would become a lifelong friend.

I likely wouldn't have chosen to pursue physics in grad school without the support of my advisors and friends in undergrad. I want to thank Prof. Matthew Shepherd for introducing me to particle physics research at Indiana University. I also want to thank Diego Tonelli, who advised me as a summer student at CERN. I'm also grateful for my friends Josh Foster and Tom Dauer, who got me through far too many all-night homework sessions and made me a better physicist in the process.

I wouldn't be where I am today without my parents, Tom and Julie, who have supported my curiosity while providing endless love and encouragement. They are always ready to provide wisdom while letting me make just enough mistakes to develop some wisdom of my own. I couldn't ask for better parents.

Finally, I want to thank my girlfriend, Michelle. Since we first met in undergrad, her work ethic has constantly inspired me. She has somehow supported me and pushed me to be a better person while navigating medical school a thousand miles away. She is the love of my life, and I can't imagine tackling the challenges of the last five years with anyone else.

Contents

1	Introduction	21
2	Theoretical Background	23
2.1	Some Useful Concepts	23
2.1.1	Natural Units	23
2.1.2	The Standard Model of Particle Physics	23
2.1.3	Feynman Diagrams	24
2.2	QCD in Hadron Collisions	28
2.2.1	Quark-Gluon Plasma	30
2.2.2	Parton Distributions	30
2.2.3	Nuclear Modification of PDFs	33
2.2.4	Parton Saturation	33
2.3	Photon Production in Hadron Collisions	34
2.3.1	Direct Production	34
2.3.2	Thermal Production	37
2.3.3	Production in Hadron Decays	37
2.3.4	Experimental Status	37
3	The LHCb Experiment	39
3.1	The Large Hadron Collider	39
3.2	How to Study B -Physics at a Hadron Collider	41
3.3	The LHCb Detector	42
3.3.1	Tracking	43
3.3.2	Particle Identification	47
3.4	The LHCb Trigger	51
3.4.1	L0	51
3.4.2	HLT1	53
3.4.3	HLT2	54
4	The LHCb GPU High Level Trigger	57
4.1	The LHCb Upgrade	58
4.1.1	The Upgrade Trigger Strategy	58
4.1.2	VELO	60
4.1.3	Upstream Tracker	62
4.1.4	Scintillating Fibre Tracker	63

4.1.5	Muon System	63
4.1.6	HLT1 Reconstruction Requirements	63
4.2	The Allen Project	64
4.2.1	The Basics of GPUs	64
4.2.2	The Allen Framework	66
4.2.3	Integration	68
4.3	Trigger Algorithms For GPUs	68
4.3.1	The Allen Trigger Sequence	68
4.3.2	Global Event Cut	69
4.3.3	Track and Primary Vertex Reconstruction	69
4.3.4	Kalman Filter and Secondary Vertex Fitter	73
4.3.5	Trigger Lines, Selections, and Output	75
4.4	Performance	78
4.4.1	Track Reconstruction	78
4.4.2	Kalman Filter	83
4.4.3	Selections	85
4.4.4	Throughput	87
4.5	Conclusion	89
5	Intrinsic Charm and Forward $Z + c$ Production	93
5.1	Intrinsic Charm	93
5.2	$Z + c$ Production at LHCb	95
5.3	Bayesian Reweighting	97
5.4	Conclusion	99
6	Neutral Pion and Photon Production in Proton-Lead Collisions	103
6.1	Observables	103
6.1.1	Proton-Lead Collisions at LHCb	103
6.1.2	The Direct Photon Excess Ratio	104
6.1.3	Nuclear Modification Factors	105
6.1.4	Forward-Backward Ratio	106
6.2	Data and Simulation	106
6.2.1	Data	106
6.2.2	Simulation	106
6.3	Reconstruction and Selection	107
6.3.1	Converted Photons and Neutral Pions	107
6.3.2	Proxy Photons	108
6.4	Diphoton Mass Fit	110
6.5	$R_{p\text{Pb}}^{\pi^0}$ Measurement	111
6.5.1	π^0 Yields	111
6.5.2	Efficiency Corrections	113
6.5.3	Luminosity Calibration	117
6.5.4	Systematic Uncertainties	121
6.5.5	Results	122
6.6	R^γ Measurement	122

6.6.1	Photon and π^0 Yields	122
6.6.2	Efficiency and Acceptance Corrections	126
6.6.3	π^0 Decay Photon Fraction	133
6.6.4	Resolution Corrections	136
6.6.5	Systematic Uncertainties	143
6.6.6	Results	145
6.7	Conclusion	145
7	Conclusion	151
A	Kalman Filter Parameterizations	153
A.1	Propagation in the Velo, Ut, and SciFi	153
A.2	Propagation in the Velo-Ut Region	154
A.3	Propagation in the Ut-SciFi Region	154
A.4	Multiple Scattering Noise	155
B	Diphoton Mass Fits	157
B.1	Fits for $R_{p\text{Pb}}^{\pi^0}$	157
B.2	Fits for R^{γ}	178

List of Figures

2-1	Standard Model fermions and gauge bosons. Quarks are shown in blue, leptons in red, and gauge bosons in orange. The connections demonstrate couplings between the gauge bosons and fermions.	25
2-2	Divergent loop diagram contributing to $e^+e^- \rightarrow \mu^+\mu^-$ at NLO.	28
2-3	The QCD fine structure constant α_s as a function of energy scale Q . Calculated using the Wilson package [16].	29
2-4	Parton distribution functions at low and high Q^2 . Generated using LHAPDF [44] and the CT18NNLO [41] parton distribution functions.	32
2-5	Kinematic coverage of various experiments used in PDF fits.	32
2-6	The gluon nuclear modification factor in ^{208}Pb from the EPPS16 nPDF set [47].	34
2-7	Diagrams contributing to direct photon production at leading order including (a) $q\bar{q}$ scattering and (b) QCD compton scattering.	35
2-8	Example of a NLO diagram contributing to the fragmentation component of direct photon production.	35
2-9	Fractional contributions of prompt and fragmentation photons to the total direct photon cross section in 13 TeV pp collisions in LHCb acceptance. Calculated using the JETPHOX generator [65–67] using the CTEQ6 PDF set [68] and the BFG II photon fragmentation functions [69].	36
2-10	Relative contributions of the five most common sources of photons from hadron decays. Generated with CIMBA [78] using interpolation grids produced with PYTHIA8 [79].	38
3-1	Pseudorapidity of $b\bar{b}$ pairs produced via gluon fusion in pp collisions at $\sqrt{s} = 13$ TeV. The LHCb detector acceptance is outlined in black. Generated using PYTHIA8 [79] and the NNPDF2.3 LO PDF set [111].	42
3-2	Schematic of the LHCb detector in the $y - z$ plane.	43
3-3	Schematic of the VELO detector in the closed (3-3a) and open (3-3b) configurations. The ϕ strips are shown in blue and the r strips are shown in red. Strip pitches are scaled up for legibility. The beamline is denoted by black circles. Adapted from Ref. [113].	44
3-4	Schematic of the first TT plane. Each square represents a sensor consisting of 512 vertical silicon strips. The beamline is denoted by a black circle. Adapted from Ref. [114].	45

3-5	Schematic of an IT layer. Each shaded rectangle represents a sensor containing 384 vertical strips each. The IT consists of two sensor types of slightly different thicknesses. The thicker sensors are shown in red and the thinner sensors are shown in blue. Adapted from Ref. [116]. . .	46
3-6	Schematic of an OT x layer. F -type modules are shown in blue and S -type modules are shown in red. The beamline is denoted by a black circle. Adapted from Ref. [117].	47
3-7	Illustration of different LHCb track types in the $x - z$ plane. Adapted from Ref. [103].	48
3-8	Schematic of the upper right quadrants of the ECAL (a) and HCAL (b). The PS and SPD share the same segmentation as the ECAL but are 0.5% and 1% smaller, respectively. In both diagrams blue squares contain a single calorimeter cell and red squares contain 2x2 arrays of cells. The orange squares in a contain 3x3 arrays of cells. The beamline is denoted by the black circles.	49
3-9	Schematic of the upper right quadrant of a Muon station. The R4 region is shown in blue, R3 in red, R2 in orange, and R1 in green. The dimensions Δx and Δy vary for each Muon station and are described in the text. The beamline is denoted by a black circle.	50
3-10	Data acquisition and trigger sequences for Runs 1 (a) and 2 (b). The trigger sequence evolved throughout Run 1, and a shows the trigger as it operated in 2012.	52
4-1	Production rates for various particles at the Run 3 LHC design energy and luminosity. Cross sections are calculated using MADGRAPH5_AMC@NLO [130,131] and the NNPDF3.1 NLO PDF set [132].	58
4-2	Efficiency as a function of rate for a simulated hardware trigger. Passing events must have at least one generated stable hadron in LHCb acceptance satisfying a p_T requirement. The trigger rate is adjusted by changing the p_T threshold. Passing signal events must have at least one final state hadron satisfying the p_T requirement. Efficiencies are calculated with respect to reconstructible signal decays. The line thickness shows the statistical uncertainty.	59
4-3	Diagram of the LHCb data acquisition system for Run 3. In Run 3, the hardware level L0 trigger will be removed and HLT1 will be run on GPUs hosted in the Event Builder nodes.	60
4-4	Schematic of the upgraded Velo detector. The open and closed detector configurations are shown in a and b, respectively. Each blue square represents a 256×256 grid of silicon pixels. The relative position of the layers along the z axis is shown in c. Adapted from Ref. [134]. . .	61

4-5	Schematic of the upstream UT detector layer in the x configuration. Each blue square represents a sensor with 512 silicon strips. The red squares represent sensors with half pitch strips, and orange rectangles represent sensors with half pitch and half length. The downstream layers share the same layout, but have two additional staves of sensors. The beampipe is shown as a white circle at the center of the diagram. Adapted from Ref. [135].	62
4-6	Schematic of a SciFi detector layer in the x configuration. Each rectangle represents a single module consisting of upper and lower fiber mats made up of 2.4 m scintillating fibers. The beampipe is shown as a white circle in the center of the detector. Adapted from [135].	63
4-7	Comparison of theoretical maximum performance of a selection of GPUs and CPUs.	65
4-8	Illustration of the memory and thread layout of a GPU grid.	67
4-9	The Allen HLT1 sequence of algorithms. Input is raw data packaged into events by the event builder nodes and copied to the GPU. This is shown in blue. For selected events raw banks summarizing the passing decisions and the objects that triggered them are copied back to the CPU. This is shown in red.	69
4-10	Illustration of the mask clustering algorithm. Pixel borders are shown as thin lines and super pixel borders are shown in bold. Activated pixels that are not included in the cluster are shown in gray. Pixels included in the cluster at each step are shown in blue. The masks are shown in red. Adapted from Ref. [145].	70
4-11	Illustration of the <i>Search by Triplet</i> algorithm. Triplets are formed (a) and extrapolated to the next layer, where additional hits are added to the track (b). Hits in tracks with at least four hits are flagged (shown in red) and the triplet formation is repeated (c). Tracks are then forwarded to the next layer (d) and the steps are repeated until the entire detector has been processed. Adapted from Ref. [145].	71
4-12	Demonstration of the UT search window determination. The black grid demarcates UT sectors. Activated strips are shown in blue and the extrapolated VELO track position is shown as a black circle. Search windows are shown in red. Adapted from Ref. [146].	72
4-13	VELO track reconstruction efficiency as a function of (a) p and (b) p_T . The efficiency is calculated with respect to generated particles reconstructible as tracks in the VELO. The shaded regions show the generated distributions.	79
4-14	VELO-UT track reconstruction efficiency as a function of (a) p and (b) p_T . The efficiency is calculated with respect to generated particles reconstructible as long tracks. The shaded regions show the generated distributions.	79
4-15	Forward track reconstruction efficiency as a function of (a) p and (b) p_T . The efficiency is calculated with respect to particles reconstructible as long tracks. The shaded regions show the generated distributions.	80

4-16 Primary vertex reconstruction as a function of (a) the number of tracks originating from the generated primary vertex and (b) the z position of the generated primary vertex. The shaded regions show the generated distributions.	80
4-17 Position resolution of reconstructed primary vertices in (a) x , (b) y , and (c) z as a function of the number of tracks originating from the generated primary vertex. The shaded regions show the generated distributions.	81
4-18 Pull distributions of reconstructed primary vertex positions.	82
4-19 Resolution of the momentum estimates produced by the forward tracking algorithm and the <i>modified</i> parameterized Kalman filter. The shaded region shows the generated distribution.	83
4-20 The χ^2 per degree of freedom cumulative distributions for real and ghost tracks. Results are shown for the (a) <i>modified</i> and (b) <i>simplified</i> parameterized Kalman filters. The widths of the lines shows the statistical uncertainty.	84
4-21 One dimensional impact parameter resolution in (a) x and (b) y determined using various fitting strategies.	84
4-22 The χ^2_{IP} distribution for the various Allen track fits.	85
4-23 Pulls of the track state closest to the beamline determined using the simplified Kalman filter.	86
4-24 Trigger efficiencies for various simulated signal samples. The combined efficiency for all prototype trigger selections is shown in black, while the two most efficient selections are shown in blue and red. Efficiencies are calculated with respect to events with reconstructible signal decays as defined in the text.	88
4-25 Allen throughput as a function of theoretical maximum 32-bit TFLOPS.	89
4-26 Allen throughput as a function of the size of the raw data produced by the SciFi. This measurement is performed using a NVIDIA Tesla V100 16GB GPU. The distribution of SciFi event sizes in simulated minimum bias events is shown in the shaded region.	90
4-27 The relative contribution of each set of algorithms to the Allen processing time. Here “other” includes secondary vertex fitting, output preparation, selection evaluation, and Kalman filtering, which contribute 1.4%, 0.5%, 0.2%, and 0.2%, respectively.	90
5-1 Leading order Feynman diagrams contributing to $Z + c$ production.	95
5-2 Intrinsic charm PDFs based on variations of the CT14 NNLO PDF set [160]. The gray band shows the CT14 NNLO PDF uncertainty.	95
5-3 Comparison of NNPDF3.0 [161] and NNPDF3.1 [42] NNLO charm PDFs. The high- x structure in the NNPDF3.1 indicates that the global analysis favors a small valence-like intrinsic contribution to the charm quark PDF.	96
5-4 Expected measurements of $\sigma(Z + c)/\sigma(Z + j)$ as a function of $y(Z)$ for (a) Run 2 and (b) Run 3 calculated using various intrinsic charm models.	98

5-5	Comparison raw and weighted χ^2/ndf distributions calculated using the NNPDF3.1 replicas and hypothetical measurements assuming (a) BHPS1, (b) BHPS2, (c) SEA1, and (d) SEA2 intrinsic charm models.	100
5-6	Comparison of charm PDFs at $Q^2 = 10^4 \text{ GeV}^2$ with and without hypothetical $Z + c$ measurements assuming (a) BHPS1, (b) BHPS2, (c) SEA1, and (d) SEA2 intrinsic charm models.	101
5-7	Comparison of charm PDF uncertainties with and without hypothetical $Z + c$ measurements assuming (a) BHPS1, (b) BHPS2, (c) SEA1, and (d) SEA2 intrinsic charm models.	102
6-1	Converted photon (OS) and SS mass spectra after the full selection.	109
6-2	Example fits to the π^0 signal in 8.16 TeV $p\text{Pb}$ simulation.	111
6-3	Example comparisons between combinatoric background in 8.16 TeV $p\text{Pb}$ simulation and the proxy diphoton background model.	112
6-4	Example fits background to 8.16 TeV $p\text{Pb}$ simulation. The combinatoric background is shown in light gray, while the bremsstrahlung background is shown in dark gray. The blue line shows the full background fit.	112
6-5	Uncorrected π^0 yields in the forward rapidity window.	114
6-6	Uncorrected π^0 yields in the backward rapidity window.	115
6-7	Efficiency maps used to correct measured π^0 yields. The maps shown in a and b are used to correct both 13 TeV and 5.02 TeV pp measurements.	116
6-8	Fits to diphoton mass spectra for $\gamma^{\text{cav}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cav}}\gamma^{\text{cav}}$ candidates (b) in 13 TeV pp data.	117
6-9	Fits to diphoton mass spectra for $\gamma^{\text{cav}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cav}}\gamma^{\text{cav}}$ candidates (b) in 5.02 TeV pp data.	118
6-10	Fits to diphoton mass spectra for $\gamma^{\text{cav}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cav}}\gamma^{\text{cav}}$ candidates (b) in 8.16 TeV $p\text{Pb}$ data.	118
6-11	Fits to diphoton mass spectra for $\gamma^{\text{cav}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cav}}\gamma^{\text{cav}}$ candidates (b) in 8.16 TeV Pbp data.	119
6-12	Measured 5.02 and 13 TeV pp π^0 production cross sections and interpolation results.	122
6-13	Neutral pion nuclear modification factors measured in $p\text{Pb}$ (a) and Pbp (b) collisions at $\sqrt{s_{\text{NN}}} = 8.16 \text{ TeV}$ with the LHCb detector. The statistical uncertainties are illustrated using error bars and the systematic uncertainties are shown as boxes. NLO pQCD calculations are shown in blue. The shaded regions show the nPDF uncertainties.	123
6-14	Uncorrected inclusive converted photon yields used to measure R^γ	124
6-15	Uncorrected π^0 yields used to measure R^γ	125
6-16	Efficiency corrected π^0 yields from data. Hagedorn fit results are shown as blue curves.	127
6-17	Efficiency corrected π^0 yields from simulation. Hagedorn fit results are shown as blue curves.	128
6-18	Calculated π^0 acceptances from data.	129
6-19	Calculated π^0 acceptances from simulation.	130

6-20 Reconstruction and total π^0 efficiencies determined for each data sample.	131
6-21 Reconstruction and total π^0 efficiencies determined for each simulation sample.	132
6-22 Ratios of measured to true acceptance in simulation.	133
6-23 True and fully corrected measured π^0 yields in each simulation sample.	134
6-24 Ratio of fully corrected measured π^0 yields to the true yields in simulation.	135
6-25 Ratio of π^0 to η acceptance from PYTHIA8.	136
6-26 Measured η/π^0 ratios in data. Fit results are illustrated by the blue curves.	137
6-27 Measured η/π^0 ratios in simulation. Fit results are illustrated by the blue curves.	138
6-28 Measured $f_{\text{data}}^{\pi^0}/f_{\text{ref}}^{\pi^0}$ for each combination of data and reference sample. The data points show the measured values, while the curves show the fit results.	139
6-29 Unfolding maps obtained from each simulation sample.	140
6-30 Raw and unfolded converted photon yields from data.	141
6-31 Raw and unfolded π^0 yields from data.	142
6-32 Summary of systematic uncertainties for each R^γ measurement.	143
6-33 R^γ measurements performed with each reference sample. The 13 TeV pp data reference sample is omitted from a, as this measurement is trivially unity.	146
6-34 R^γ measured in pp collisions at $\sqrt{s} = 13$ TeV.	147
6-35 R^γ measured in pp collisions at $\sqrt{s} = 5.02$ TeV.	147
6-36 R^γ measured in $p\text{Pb}$ collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV.	148
6-37 R^γ measured in $\text{Pb}p$ collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV.	148
B-1 π^0 mass fits in 13 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$	158
B-1 π^0 mass fits in 13 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).	159
B-2 π^0 mass fits in 13 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$	160
B-2 π^0 mass fits in 13 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$ (cont.).	161
B-3 π^0 mass fits in 13 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$	162
B-3 π^0 mass fits in 13 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$ (cont.).	163
B-4 π^0 mass fits in 5.02 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$	164
B-4 π^0 mass fits in 5.02 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).	165
B-5 π^0 mass fits in 5.02 TeV pp data as a function of $p_{\text{T}}(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$	166

B-5 π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$ (cont.).	167
B-6 π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$	168
B-6 π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$ (cont.).	169
B-7 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$	170
B-7 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).	171
B-8 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$	172
B-8 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < y^*(\pi^0) < 3.5$ (cont.).	173
B-9 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$	174
B-9 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).	175
B-10 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$	176
B-10 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\pi^0)$ for $-4.0 < y^*(\pi^0) < -3.0$ (cont.).	177
B-11 π^0 mass fits in 13 TeV pp data as a function of $p_T(\gamma^{\text{cnv}})$	178
B-12 π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\gamma^{\text{cnv}})$	179
B-13 π^0 mass fits in 8.16 TeV pPb data as a function of $p_T(\gamma^{\text{cnv}})$	180
B-14 π^0 mass fits in 8.16 TeV Pbp data as a function of $p_T(\gamma^{\text{cnv}})$	181
B-15 η mass fits in 13 TeV pp data as a function of $p_T(\gamma^{\text{cnv}})$	182
B-16 η mass fits in 5.02 TeV pp data as a function of $p_T(\gamma^{\text{cnv}})$	183
B-17 η mass fits in 8.16 TeV pPb data as a function of $p_T(\gamma^{\text{cnv}})$	184
B-18 η mass fits in 8.16 TeV Pbp data as a function of $p_T(\gamma^{\text{cnv}})$	185

List of Tables

2.1	The generations of Standard Model fermions and their masses. Electric charges are shown for particles. Antiparticles will have opposite sign charges.	25
2.2	The Feynman diagram elements for Standard Model propagators.	26
2.3	Allowed Standard Model Feynman vertices, excluding the Higgs sector.	27
4.1	Requirements for one-track trigger selections implemented in Allen.	77
4.2	Requirements for two-track secondary vertex trigger selections implemented in Allen.	77
4.3	Event selection rates for the Allen prototype trigger selections.	87
5.1	$\langle x \rangle_{\text{IC}}$ values for each intrinsic charm model considered.	96
6.1	Summary of datasets used for the analyses presented in this chapter. The listed pp luminosities are the luminosities measured in Section 6.5.3, while the $p\text{Pb}$ and Pbp are the luminosities delivered to LHCb by the LHC during the 2016 proton-lead runs.	107
6.2	Summary of selection criteria for converted and ECAL photons.	110
6.3	Center-of-mass rapidity regions used for measuring $R_{p\text{Pb}}^{\pi^0}$	113
6.4	Efficiency corrections to the measured π^0 yields.	117
6.5	Luminosities for each sample used to calculate $R_{p\text{Pb}}^{\pi^0}$. The reported μ value for the $p\text{Pb}$ and Pbp samples is the visible minimum bias interaction rate μ_{MB}	120
6.6	Summary of the systematic uncertainties in the $p\text{Pb}$ (forward) and Pbp (backward) directions as explained in the text.	121

Chapter 1

Introduction

One of the main goals of high-energy collider physics is to produce *something* in the collisions and study its properties. This is arguably the primary goal of the Large Hadron Collider (LHC). For example a collider can produce a previously undiscovered particle. The LHC is most famous for producing a particle with a mass of about 125 GeV and showing that its properties are consistent with the Standard Model Higgs boson [1, 2]. The LHC also collides heavy ions with the goal of creating temperatures high enough to “melt” matter into its constituent quarks and gluons, creating a medium called Quark-Gluon Plasma (QGP) [3–5]. Particles created in these collisions are used to determine the properties of the QGP, such as its temperature and shear viscosity [6, 7].

The second goal is to use particle collisions to study the structure of the colliding particles. For example, physicists scatter electrons off of protons to determine the proton radius [8]. At higher energies, electron-proton collisions are used to study the distributions of quarks and gluons within the proton [9]. Analogously, proton-ion and electron-ion collisions are used to study the structure of the nucleus.

At hadron colliders like the LHC, these two goals are never completely separate. In order to measure the properties of the Higgs boson, physicists must also understand the structure of the protons they collide in order to produce the Higgs. Understanding the properties of the QGP requires disentangling effects caused by the medium from those caused by the complicated structure of the nucleus. Furthermore, the underlying structure of protons and nuclei means high energy collisions of these particles are often extremely messy, producing dozens or hundreds of particles that serve as backgrounds to the experimental signatures of interesting processes. This thesis presents three studies that reflect the inevitable entanglement of studying hadron collisions and studying hadrons themselves:

- Documentation of the Allen project.
- A phenomenological study of LHCb’s sensitivity to the intrinsic charm content of the proton through Z boson production in association with a charm jet ($Z + c$).

- Studies of neutral pion and direct photon production in proton-proton and proton-lead collisions using the LHCb detector.

This thesis begins by presenting some useful theoretical and experimental background information in Chapter 2. Chapter 3 presents a description of the LHCb experimental setup. This includes the LHC accelerator, the LHCb detector, and the LHCb trigger system.

Chapter 4 presents Allen a GPU implementation of LHCb’s first level software trigger. In Run 3, LHCb will collect data with no hardware trigger. This provides increased flexibility in data selection strategies, allowing LHCb to better separate signatures of interesting physical processes from large backgrounds produced in high energy hadron collisions. The flexibility of triggerless data acquisition brings comes with additional computational challenges. LHCb’s trigger system must process the full LHC collision rate in software. Chapter 4 shows how Allen accomplishes this task.

Chapter 5 presents a phenomenological study of $Z + c$ production at LHCb. $Z + c$ production is sensitive to the charm quark content of the proton. Studying this process at LHCb could provide evidence that the proton contains intrinsic charm. Such evidence would invalidate common assumptions about the structure of the proton, affecting theoretical calculations of a broad range of physical processes.

Finally, Chapter 6 presents a study of direct photon production in proton-proton and proton-lead collisions using the LHCb detector. Direct photon production is sensitive to the gluon content of the proton and nucleus. LHCb has the ability to study direct photon production in a previously unexplored kinematic regime. Additionally, direct photon production can be enhanced by thermal photons radiated by the QGP. Observation of thermal photons in proton-proton or proton-lead collisions would provide evidence of QGP formation in these collision systems. As a result, direct photons can be used as both a tool for probing nuclear structure and a tool for studying the medium created in high energy nuclear collisions.

Chapter 2

Theoretical Background

This chapter presents theoretical concepts that will be useful for understanding the physics presented in the rest of this thesis. This begins with a brief conceptual introduction to the Standard Model, followed by a more detailed discussion of QCD in hadron collisions. This includes an introduction to parton distribution functions in protons and nuclei. Finally, I will discuss how photons are produced in hadron collisions and how studying them reveals information about the partonic structure of protons and nuclei.

2.1 Some Useful Concepts

2.1.1 Natural Units

This thesis uses a natural unit system defined by the relation

$$c = \hbar = k_B = 1, \quad (2.1)$$

where c is the speed of light in vacuum, \hbar is the reduced Planck constant, and k_B is the Boltzmann constant. As a result, all dimensionful quantities used to describe microscopic physical phenomena in this thesis will have dimension [energy] n , where n is referred to as the quantity’s “mass dimension” or simply “dimension”. Mass, energy, momentum, and temperature all have mass dimension 1, while length and time all have mass dimension -1 . For descriptions of the LHC and the LHCb experimental setup, SI units are used for convenience.

2.1.2 The Standard Model of Particle Physics

The Standard Model (for a detailed description see Refs. [10] and [11]) describes three of the four fundamental forces of nature. This includes the strong, weak, and electromagnetic forces. The Standard Model consists of a collection of elementary fermions, which interact by exchanging gauge bosons.

The fundamental units of matter in the Standard Model are 12 fermions. These fermions are divided into quarks and leptons. These are further divided into three

generations with increasing masses. The Standard Model fermions are listed in Table 2.1. The leptons include the electron (e), muon (μ), and tau (τ), which all have charge -1 . The leptons also include corresponding chargeless neutrinos (ν_e , ν_μ , and ν_τ). The quarks include the up (u), down (d), strange (s), charm (c), bottom (b), and top (t) quarks. The u , c , and t quarks are referred to as “up-type” (q_u) quarks and have charge $+2/3$, while the d , s , and b quarks are called “down-type” (q_d) quarks and have charge $-1/3$. For each fundamental fermion, the Standard Model also includes an antifermion. Antifermions will have the opposite sign charge of the corresponding fermion.

These fermions interact by exchanging spin-1 particles called gauge bosons. The electromagnetic force acts on electrically charged fermions through photon exchange. Quarks are held together in hadrons through the exchange of gluons. All of the Standard Model fermions interact via the weak force by exchanging W^\pm and Z^0 bosons. Figure 2-1 illustrates these interactions. All of the massive Standard Model particles (except for possibly the neutrinos) become massive by interacting with the spin-0 Higgs boson (H). The strength of each force is characterized by a coupling constant¹

$$\alpha_s \sim 1, \tag{2.2}$$

$$\alpha_{\text{QED}} \sim 1/137, \tag{2.3}$$

$$\alpha_w \sim 10^{-7}, \tag{2.4}$$

where α_s , α_{QED} , and α_w characterize the strengths of the strong, electromagnetic, and weak interactions, respectively.

The Standard Model clearly does not describe all physical phenomena. Most glaringly, the Standard Model does not provide a theory of gravity, which is successfully described by General Relativity. The Standard Model also does not contain a satisfactory dark matter candidate or explain dark energy, which collectively make up almost all of the energy in the universe. Standard Model neutrinos were thought to be massless until the observation of neutrino oscillations [12, 13], and the origin of neutrino masses is still unknown. Despite these deficiencies, the Standard Model has been wildly successful, notwithstanding decades of rigorous testing. That said, finding the cracks in the Standard Model is one of the primary goals of modern particle physics.

2.1.3 Feynman Diagrams

The Standard Model is a quantum field theory (QFT). This means that particles are excited quantum states of underlying fields. When a force acts on one of these quantum states, it can cause a transition to another state. The transition rate Γ can

¹The coupling constants are approximated here at $Q^2 = 0$. The scale Q^2 is defined in the next section.

generation	q_u ($C = +2/3$)	q_d ($C = -1/3$)	l^- ($C = -1$)	ν_l ($C = 0$)
1	$\begin{pmatrix} u \\ \sim 2 \text{ MeV} \end{pmatrix}$	$\begin{pmatrix} d \\ \sim 5 \text{ MeV} \end{pmatrix}$	$\begin{pmatrix} e \\ 0.511 \text{ MeV} \end{pmatrix}$	$\begin{pmatrix} \nu_e \\ \sim 0 \text{ MeV} \end{pmatrix}$
2	$\begin{pmatrix} c \\ 1.3 \text{ GeV} \end{pmatrix}$	$\begin{pmatrix} s \\ 95 \text{ MeV} \end{pmatrix}$	$\begin{pmatrix} \mu \\ 106 \text{ MeV} \end{pmatrix}$	$\begin{pmatrix} \nu_\mu \\ \sim 0 \text{ MeV} \end{pmatrix}$
3	$\begin{pmatrix} t \\ 173 \text{ GeV} \end{pmatrix}$	$\begin{pmatrix} b \\ 4.2 \text{ GeV} \end{pmatrix}$	$\begin{pmatrix} \tau \\ 1.78 \text{ GeV} \end{pmatrix}$	$\begin{pmatrix} \nu_\tau \\ \sim 0 \text{ MeV} \end{pmatrix}$

Table 2.1: The generations of Standard Model fermions and their masses. Electric charges are shown for particles. Antiparticles will have opposite sign charges.

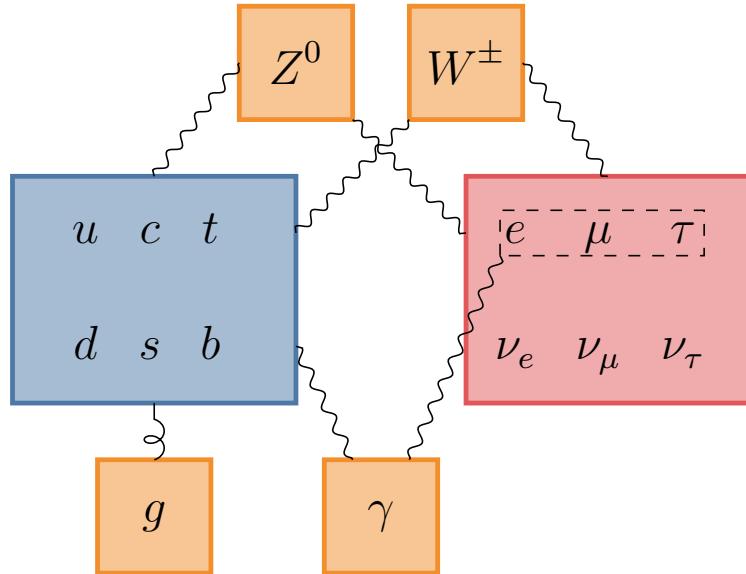


Figure 2-1: Standard Model fermions and gauge bosons. Quarks are shown in blue, leptons in red, and gauge bosons in orange. The connections demonstrate couplings between the gauge bosons and fermions.

particle	propagator
q, l	
W^\pm, Z^0, γ	
g	
H	

Table 2.2: The Feynman diagram elements for Standard Model propagators.

be calculated using Fermi’s Golden Rule,

$$\Gamma_{i \rightarrow f} = 2\pi |\mathcal{M}_{i \rightarrow f}|^2 \rho_f, \quad (2.5)$$

where i and f denote initial and final states, respectively, and ρ_f is the density of available states for a final state f . If the force causing the transition is sufficiently weak, the transition (or scattering) amplitude \mathcal{M} can be calculated using perturbation theory. Terms in this perturbative series are represented by Feynman diagrams. Detailed instructions for using Feynman diagrams to calculate scattering amplitudes can be found in many introductory texts on QFT [14] and particle physics [10, 11]. Here I will try to provide some intuition for the basic qualities of processes illustrated by Feynman diagrams.

Feynman diagrams have a time axis and a space axis. Here, the time axis is horizontal and the space axis is vertical. Propagating particles are represented as lines, which connect at vertices. A line that has at least one free end is an external line, while a line that connects two vertices is an internal line. Table 2.2 shows the lines used to represent each of the Standard Model particles. External lines represent initial and final state particles with definite momenta and quantum numbers. Internal lines represent virtual particles exchanged in interactions. Virtual particles can have energies and momenta that do not obey $M^2 = E^2 - p^2$. Vertices represent particle interactions. Table 2.3 shows all of the allowed Standard Model vertices. Quantities conserved by an interaction, such as momentum and conserved quantum numbers, must be conserved at vertices. In general, vertices will contribute a factor proportional to the coupling constant of the associated force (α_s , α_{QED} , or α_w) to the modulus squared of the scattering amplitude, $|\mathcal{M}_{i \rightarrow f}|^2$.

All Feynman diagrams with the same initial and final state particles contribute to the total scattering amplitude. The diagrams with the fewest vertices contributing to a process are called Leading Order (LO) diagrams. LO diagrams typically make the largest contributions to the scattering amplitude and are sometimes sufficient for predicting experimental results. Diagrams with additional vertices are referred to as Next-to-Leading Order (NLO), Next-to-Next-to-Leading Order (NNLO), and so on. In addition, diagrams will contribute depending on the relevant coupling constant. So a diagram proceeding through gluon exchange, for example will contribute more to the scattering amplitude than an identical diagram proceeding through γ or Z^0 exchange due to the relative sizes of α_s , α_{QED} , and α_w .

The rules of thumb presented above do not always hold. Consider the NLO scat-

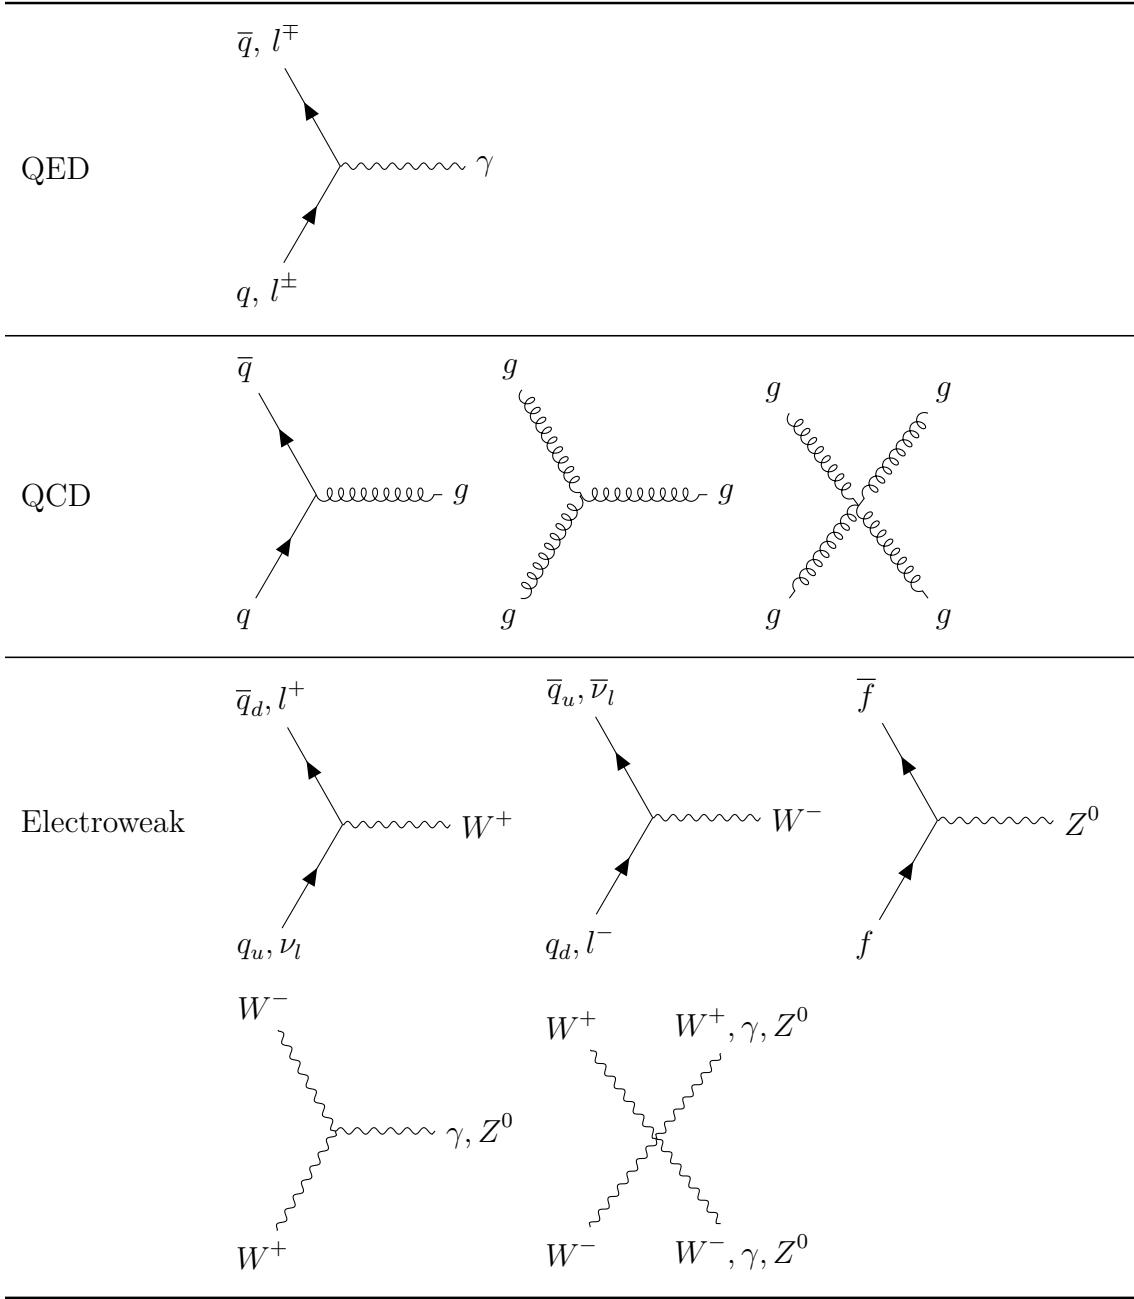


Table 2.3: Allowed Standard Model Feynman vertices, excluding the Higgs sector.

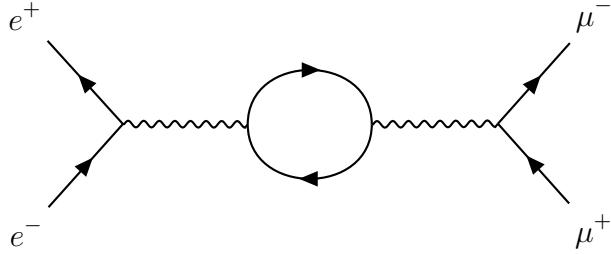


Figure 2-2: Divergent loop diagram contributing to $e^+e^- \rightarrow \mu^+\mu^-$ at NLO.

tering diagram of Figure 2-2 [10]. The qualitative rules of the previous section would indicate that this diagram has a smaller contribution to the scattering amplitude than the corresponding LO diagram. However, calculating this diagram requires integrating over all possible energies and momenta of the virtual particles in the loop. In the case of the diagram in Figure 2-2, this integral is divergent.

Of course the $e^+e^- \rightarrow \mu^+\mu^-$ cross section is not infinite. This is because the effects of these loop diagrams are offset by the effects of loop diagrams on the physical constants of the theory, such as particle masses and coupling constants. This means that when observables are measured in particle physics experiments, the effects of these diagrams are also measured. The importance of these diagrams will vary with the energy scale (or equivalently the length or time scale) of the process being probed. As a result, physical constants will vary with energy scale as well.

The process of “absorbing” these divergent loop diagrams into physical constants is called renormalization. The evolution of physical observables with the probed energy scale is given by a renormalization group equation [14]. Important examples of scale-dependent observables include the coupling constants α themselves, as well as the momentum distributions of quarks and gluons in the proton and nucleus.

The relevant energy scale of an interaction is often quantified by the momentum transfer of an interaction. If two particles interact by exchanging a third particle with four-momentum q , the energy scale relevant to the interaction is $Q^2 \equiv -q^2$. When discussing experiments involving proton scattering, such as deep inelastic lepton-proton scattering or high energy proton-proton collisions, it is useful to think of Q as the frequency of a probe interacting with the proton. As Q increases, the proton is probed on shorter time scales and the effects of short-lived fluctuations increase. On the other hand, Q is then inversely proportional to the wavelength of the probe. So probing at higher Q^2 reveals the effects of structures at smaller length scales.

2.2 QCD in Hadron Collisions

Quantum Chromodynamics (QCD) is a quantum field theory describing the strong force. The strong interaction is mediated by gluons and provides the attractive force that holds quarks together in hadrons. Particles that interact via QCD carry a charge that comes in three “colors”: red, blue, and green. Color-charged particles interact by exchanging gluons, the mediators of the strong force.

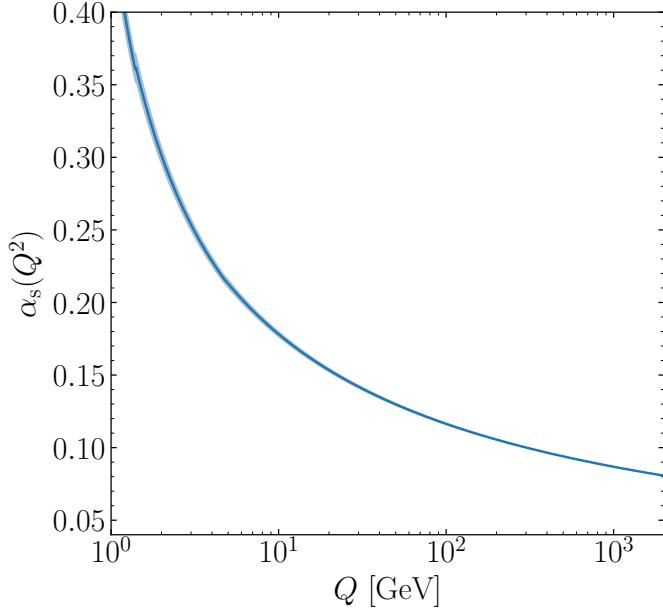


Figure 2-3: The QCD fine structure constant α_s as a function of energy scale Q . Calculated using the Wilson package [16].

The allowed QCD Feynman diagram vertices are shown in Table 2.3. Unlike photons, which are uncharged under the electromagnetic force, gluons carry color charge. As a result, gluons interact strongly with other gluons. These processes affect the range of the strong force. Unlike the electromagnetic force, which becomes stronger at short distances, the strong force becomes stronger as distance increases. Equivalently, the QCD fine structure constant α_s decreases as the energy scale Q increases. This so-called “running coupling” is shown in Figure 2-3. This quality of QCD is known as asymptotic freedom. At high energies and short distances, the QCD coupling is small and QCD calculations can be performed using perturbation theory. Conversely, at long distances and low energies α_s becomes large. This means that quarks do not propagate freely, but are confined to hadrons with no net color charge. Quarks and gluons with enough energy to escape a hadron will rapidly hadronize, producing a shower of particles called a jet. The scale at which the transition between asymptotic freedom and confinement occurs is given by Λ_{QCD} and has a value of about 200 MeV [15]. At energies much larger than Λ_{QCD} , the strong coupling constant is small enough for perturbation theory to apply. At scales much less than Λ_{QCD} , quarks and gluons are confined to hadrons.

The LHC collides protons and ions, which are collections of quarks and gluons held together by QCD. Quarks and gluons produced in these collisions hadronize via QCD. As a result, making sense of collisions at the LHC is impossible without understanding QCD.

2.2.1 Quark-Gluon Plasma

Relativistic heavy ion collisions produce extremely high-temperature matter. At these temperatures, quarks and gluons can have thermal energies larger than Λ_{QCD} [15]. As a result, heavy ion collisions can produce a medium of deconfined quarks and gluons. This state of matter is called the Quark-Gluon Plasma (QGP). The QGP was first conclusively observed at the Relativistic Heavy Ion Collider (RHIC) [17–20] and has subsequently been studied extensively at both RHIC and the LHC (see Refs. [21–26], for example).

Data from heavy ion collisions are well described by fluid dynamic models [15]. These models predict azimuthal anisotropies in the angular distributions of particles created in heavy ion collisions. This phenomenon has been observed at both RHIC and the LHC. Complicating this picture is the observation of azimuthal anisotropies in smaller collision systems such as $d\text{Au}$ collisions at RHIC and $p\text{Pb}$ and pp collisions at the LHC [27–31]. Similar collectivity was observed in photon-nucleus interactions in ultra-peripheral PbPb collisions [32], but not in e^+e^- [33,34] or e^-p [35] collisions.

Explaining collective behavior in small collision systems, as well as other signatures of QGP production observed in pp collisions [36], is an outstanding problem in relativistic heavy ion physics. Explanations include the production of “QGP droplets” in small collision systems [27], initial state effects described by the Color-Glass Condensate effective field theory [37], and final state effects described by models of color string interactions [38]. Searches for additional QGP signatures could help clarify the interpretation of data from small collision systems.

2.2.2 Parton Distributions

The structure of the proton has been studied for decades by scattering other particles off of it. When the proton is probed by very low energy particles, it appears to be a point-like particle. As the energy increases, the probe particle can begin to resolve the proton’s small but non-zero radius. At high enough energies the probes do not scatter off of the entire proton, but instead interact with the point-like partons that make up the proton. This was first observed in deep inelastic scattering experiments at the Standford Linear Accelerator Center in 1968 [39], confirming that the proton is a collection of what we now know are quarks and gluons.

Because protons are collections of quarks and gluons, collisions at the LHC are collisions of these partons. Each parton carries some fraction x of the total proton momentum. As a result, calculating cross sections of hard processes in proton-proton collisions with center of mass energy \sqrt{s} requires calculating a collection of parton-parton cross sections at a center of mass energy of $\sqrt{x_1 x_2 s}$. A generic cross section for a process that occurs in an inelastic proton-proton collision with center-of-mass energy \sqrt{s} and scale Q^2 is given by

$$\sigma_{pp} = \sum_{a,b} \int_0^1 \int_0^1 f_a(x_1, Q^2) f_b(x_2, Q^2) \sigma_{ab} dx_1 dx_2, \quad (2.6)$$

where a and b denote parton flavors and σ_{ab} is the parton-parton cross section. The functions $f_a(x, Q^2)$ and $f_b(x, Q^2)$ are the parton momentum fraction (x) distributions in the proton as seen by a virtual particle probing the proton at energy scale Q^2 . These momentum distributions are called parton distribution functions (PDFs). The PDFs contain the non-perturbative QCD of the proton structure, factorizing it from the high-energy parton-parton process whose cross section can be calculated using perturbation theory [40].

Because the proton structure is described by QCD in the non-perturbative regime, the PDFs must be determined from data. The cleanest way to measure the proton PDFs is through deep inelastic electron-proton scattering in which the proton is probed by a high-energy virtual photon. At high Q^2 , this is referred to as deep inelastic scattering (DIS). DIS from the HERA accelerator at DESY [9] is the largest source of data used in the determination of the proton PDFs. Additional data is provided by fixed target experiments, as well as experiments at hadron colliders such as the LHC and Tevatron. Data from hadron collider experiments are particularly important for constraining the gluon PDF, which is only accessible at NLO in DIS. In practice, the proton parton distribution functions are determined by global fits to all available data [41, 42].

The light quark PDFs are parameterized at an initial energy scale Q_0^2 . As Q^2 increases, QCD radiation produces additional soft quarks and gluons, including c and b quarks. This evolution can be calculated using the renormalization group equations for the PDFs, called the DGLAP equations [43]. Because quarks radiate gluons and gluons split into quarks, the DGLAP equations couple the quark and gluon PDFs and have the form

$$Q^2 \frac{\partial}{\partial \ln Q^2} \begin{pmatrix} q_i(x, Q^2) \\ g(x, Q^2) \end{pmatrix} = \frac{\alpha_s(Q^2)}{2\pi} \sum_{q_j, \bar{q}_j} \int_x^1 \frac{d\xi}{\xi} \begin{pmatrix} P_{q_j \rightarrow q_i} & P_{q_j \rightarrow g} \\ P_{g \rightarrow q_i} & P_{g \rightarrow g} \end{pmatrix} \begin{pmatrix} q_j(\xi, Q^2) \\ g(\xi, Q^2) \end{pmatrix} \quad (2.7)$$

where i and j are quark flavors. The splitting kernels $P_{a \rightarrow b}$ are functions of $z = \xi/x$ and α_s . These give the probability of particle a radiating a particle b carrying momentum fraction z of the original parton. DGLAP evolution in Q^2 is a fundamental assumption of all PDF fits.

The DGLAP equations reveal some important qualitative properties of the PDFs. As the energy scale Q^2 increases, the proton is probed at shorter time and distance scales. As a result, the probe becomes more sensitive to the short-lived sea quarks and gluons produced by QCD radiation. This is demonstrated in Figure 2-4. At large Q^2 , c and b quarks carry a significant portion of the proton's momentum. These heavy quark PDFs are assumed to be generated entirely through QCD radiation and are initialized at 0 at Q_0^2 in most PDF fits. This assumption has recently been reexamined to account for the possibility of non-perturbative, or intrinsic, contributions to the heavy quark PDFs. Evidence for intrinsic contributions to the s -, c -, and b -quark PDFs is inconclusive [42].

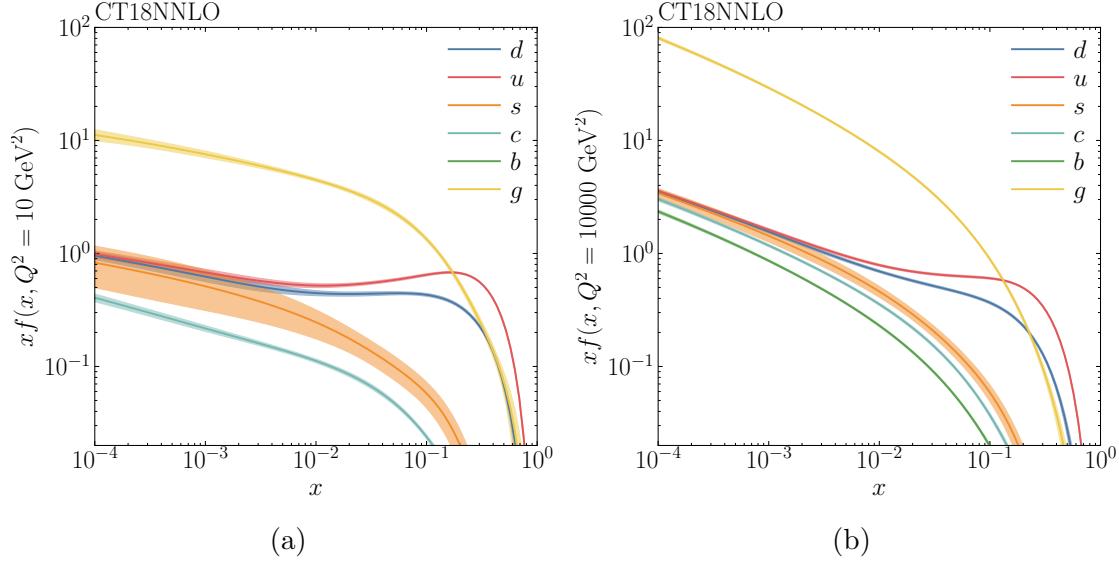


Figure 2-4: Parton distribution functions at low and high Q^2 . Generated using LHAPDF [44] and the CT18NNLO [41] parton distribution functions.

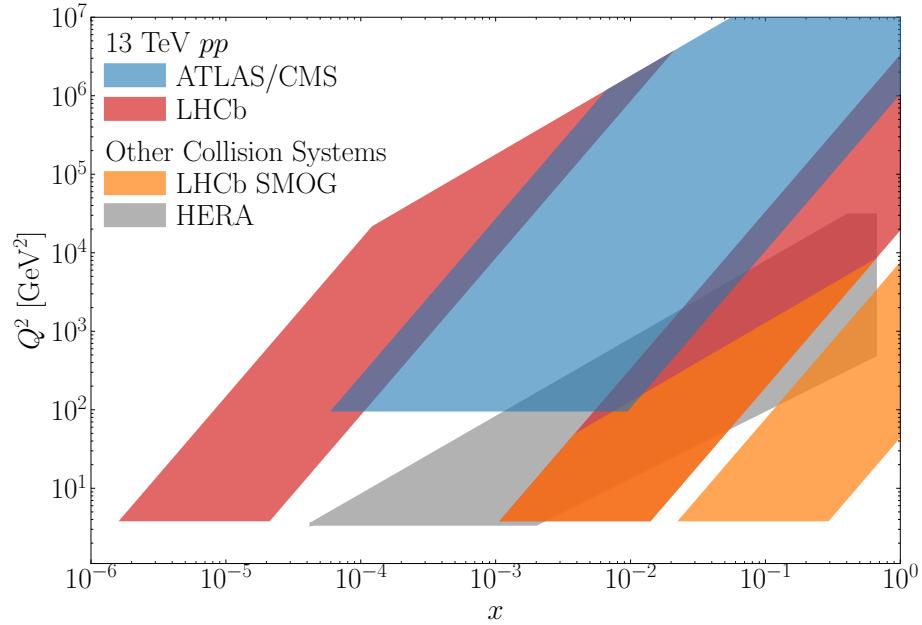


Figure 2-5: Kinematic coverage of various experiments used in PDF fits.

2.2.3 Nuclear Modification of PDFs

One would naïvely expect nuclei to behave as a collection of nucleons. In this case, cross-sections would scale with the number of nucleons in the nucleus. However, experiments show that DIS cross sections do not simply scale with the number of nucleons [45]. This indicates that a nucleus is more complicated than a collection of free nucleons.

Nuclear effects occur at the scale of typical nuclear binding energies of a few MeV per nucleon and consequently cannot be calculated using perturbative QCD. As a result, high-energy scattering data are described by absorbing nuclear effects into the initial PDF parameterizations using a nuclear modification factor R_i^A for a nucleus A and parton flavor i [46]. Perturbative evolution in Q^2 is then governed by the DGLAP equations. This means that nuclear effects are important at low energies but are slowly overtaken by perturbative evolution as Q^2 increases. The resulting PDFs are referred to as nuclear PDFs or nPDFs.

The gluon modification factor in ^{208}Pb from the EPPS16 nPDF set [47] is shown in Figure 2-6. At low x , the nuclear modification factor is less than one due to an effect called shadowing. This has been interpreted as a multiple scattering effect within the nucleus [48, 49]. An incoming parton scatters elastically off of one nucleon, decreasing the probability that it will scatter inelastically off of another nucleon. Because the PDFs must remain normalized, shadowing at low x implies anti-shadowing at high x . This is visible at $x \approx 0.1$ in Figure 2-6. At higher x , the modification factor dips below one again. This is called the EMC effect after the European Muon Collaboration that discovered it [45]. The cause of the EMC effect is unknown, though experimental evidence indicates that it may be associated with short range correlations between nucleons [50]. As x approaches one, the modification factor becomes large due to the Fermi motion of nucleons within the nucleus.

Nuclear PDFs are far less constrained by experimental data than the proton PDFs. Most of the data used is from fixed target DIS and Drell-Yan lepton pair production, as well as inclusive pion production from RHIC [46, 47, 51]. The gluon PDF is determined largely from inclusive pion production and is particularly poorly constrained. Only one nPDF set, EPPS16 [47], currently uses LHC data in its global fit.

2.2.4 Parton Saturation

In addition to non-perturbative effects that modify the nPDFs, novel perturbative effects may also modify parton distributions in nuclei. For example, as x decreases, the number density of gluons rapidly increases. At large number densities, gluon recombination will compete with gluon splitting, causing the nucleus to “saturate” with gluons [52]. It is important to note that parton saturation does not modify the nucleon PDF in the same way as shadowing, for example. Shadowing is a non-perturbative effect that can be absorbed into the initial PDF parameterization, while parton saturation is a perturbative effect that leads to non-DGLAP evolution of parton densities.

Saturation effects should be noticeable in interactions where the wavelength of the

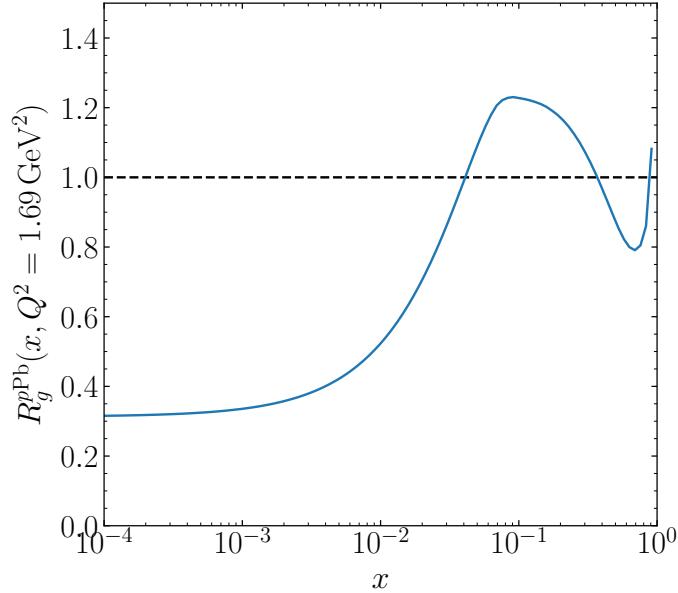


Figure 2-6: The gluon nuclear modification factor in ^{208}Pb from the EPPS16 nPDF set [47].

probe particle is comparable to the average distance between partons. This allows for the definition of a characteristic saturation scale Q_s^2 given by [52]

$$Q_s^2 \sim \frac{\alpha_s}{\pi R^2} x g(x, Q_s^2), \quad (2.8)$$

where R is the radius of the nucleon or nucleus and g is the gluon distribution. In a nucleus, the number of gluons will scale with the number of nucleons A . The volume of the nucleus will also scale with the number of nucleons, so the radius will scale as $A^{1/3}$. As a result, Q_s^2 will scale as $A^{1/3}$. This means that parton saturation will affect a larger kinematic region in collisions involving heavy nuclei. The saturation regime is typically described using an effective field theory called the Color Glass Condensate (CGC) [53].

2.3 Photon Production in Hadron Collisions

2.3.1 Direct Production

Direct photons are photons that are not produced in hadron decays. This includes photons produced in hard QCD processes. Feynman diagrams for the leading order processes are shown in Figure 2-7. Photons produced via hard QCD scattering processes are referred to as “prompt”. At high energies, such as those achieved at the LHC, the QCD compton process shown in Figure 2-7b is the dominant leading order contribution to the direct photon cross section. As a result, direct photon pro-

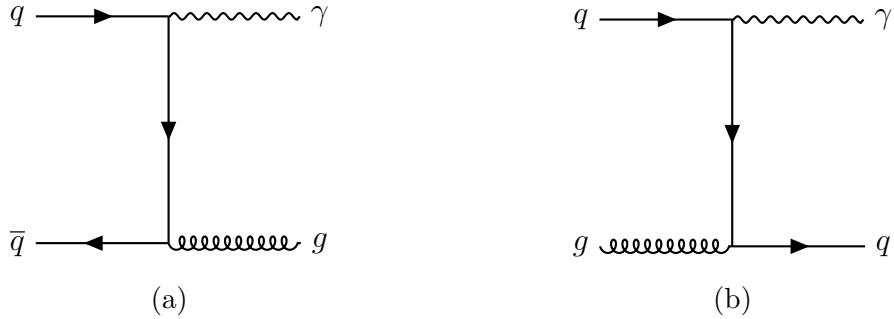


Figure 2-7: Diagrams contributing to direct photon production at leading order including (a) $q\bar{q}$ scattering and (b) QCD compton scattering.

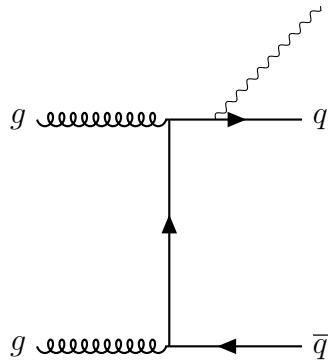


Figure 2-8: Example of a NLO diagram contributing to the fragmentation component of direct photon production.

duction has been proposed as a channel for constraining the gluon density in the proton [54, 55] and the nucleus [56–59]. As indicated in Figure 2-5, measuring direct photon production in proton-proton and proton-lead collisions at LHCb would constrain the gluon PDF at high and low x because of the LHCb detector’s forward acceptance. Because of their sensitivity to the low x gluon density, prompt photon measurements have also been proposed as probes of saturation physics and tests of the CGC effective theory [52, 60–62].

The use of direct photons as a probe of the gluon PDF and nPDF is complicated by NLO contributions to direct photon production. An example NLO diagram is shown in Figure 2-8. These photons are referred to as fragmentation photons. Although this is an NLO process, its contribution to the scattering matrix is comparable to the LO processes shown in Figure 2-7 [43]. At low p_T , the fragmentation contribution is much larger than the prompt contribution [58]. Figure 2-9 shows predictions for the fractional contributions of prompt and fragmentation photons to the total direct photon cross section in 13 TeV pp collisions in LHCb acceptance. In this low- p_T region, the fragmentation contribution is about four times larger than the prompt contribution. Fragmentation photons are accompanied by other particles produced in the parton shower. As a result, the fragmentation contribution may be suppressed by requiring that the photon be isolated from surrounding event activity [63, 64].

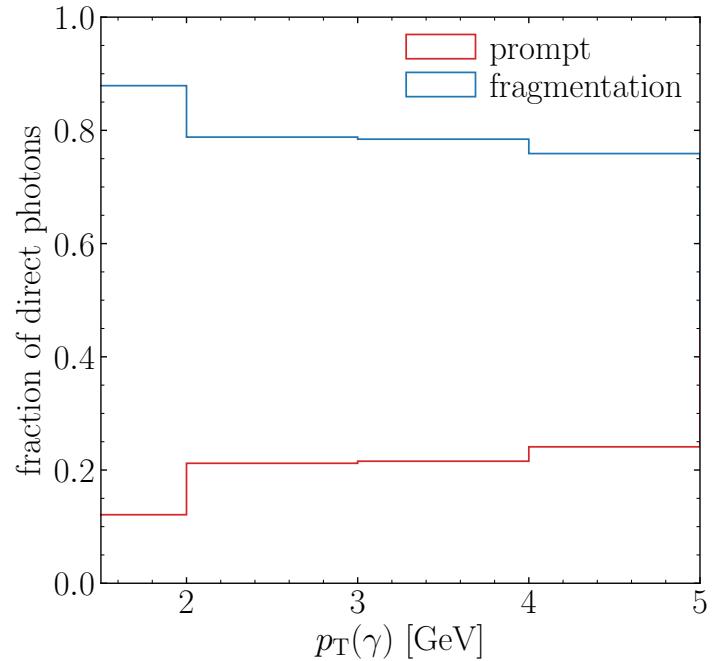


Figure 2-9: Fractional contributions of prompt and fragmentation photons to the total direct photon cross section in 13 TeV pp collisions in LHCb acceptance. Calculated using the JETPHOX generator [65–67] using the CTEQ6 PDF set [68] and the BFG II photon fragmentation functions [69].

2.3.2 Thermal Production

In addition to direct production in hard QCD processes, photons may be radiated by thermalized matter produced in heavy ion collisions, including the QGP and the subsequent hot hadron gas (HHG) [70–75]. Thermal photons are not the product of hadron decays, so they are considered an additional source of direct photons. In heavy ion collisions at low p_T , the thermal photon contribution is expected to be much larger than the prompt and fragmentation contributions [75]. Thermal photons have also been proposed as a potential signal of QGP production in pp collisions at LHC energies [76]. The experimental signature is a direct photon excess at low p_T (less than around 3 GeV), where the direct and fragmentation contributions are expected to be small.

2.3.3 Production in Hadron Decays

Most photons produced in high energy hadron collisions are the result of neutral meson decays. This idea was first proposed by J. R. Oppenheimer to explain the large numbers of low-energy photons in high energy cosmic ray showers [77]. Figure 2-10 shows predictions for the relative contributions for various sources of decay photons. Most decay photons originate from decays of π^0 and η mesons, which are produced in large numbers in high energy hadron collisions.

Although photons from meson decays are backgrounds in direct photon measurements, neutral meson production provides useful information about the nuclear gluon distribution as well [58]. This is especially true at forward rapidities, such as those accessible at LHCb, where low- p_T π^0 production is sensitive to x on the order of 10^{-5} .

2.3.4 Experimental Status

Direct photon measurements come in two main varieties: high- p_T (typically greater than 10 GeV) isolated photons, and low- p_T (typically less than 5 – 10 GeV) photons with backgrounds from decay photons subtracted. Isolated photons have been measured in fixed target, pp , and $p\bar{p}$ experiments at numerous energies. These measurements are summarized in Refs. [54] and [55]. These data are well described by perturbative QCD calculations and constrain the gluon PDF in the range $10^{-3} < x < 0.1$.

Low- p_T direct photons were first observed in high energy heavy ion collisions by the WA98 experiment in PbPb collisions at $\sqrt{s_{NN}} = 158$ GeV [80]. Most subsequent studies of direct photons at low p_T have focused on searches for thermal photons in central heavy ion collisions. The PHENIX collaboration observed an excess of direct photons with $p_T < 3$ GeV in central AuAu collisions with $\sqrt{s_{NN}} = 200$ GeV using internal photon conversions [81, 82]. They interpreted this as evidence for thermal photon production by the QGP. The STAR experiment has also observed a direct photon excess in AuAu collisions, although not as large as that seen by PHENIX [83]. More recently, PHENIX has observed similar excesses in AuAu collisions at $\sqrt{s_{NN}} = 39$ and 62.4 GeV [84], as well as in CuCu collisions at $\sqrt{s_{NN}} = 200$ GeV [85]. At the LHC, the ALICE experiment has observed a direct photon excess in central

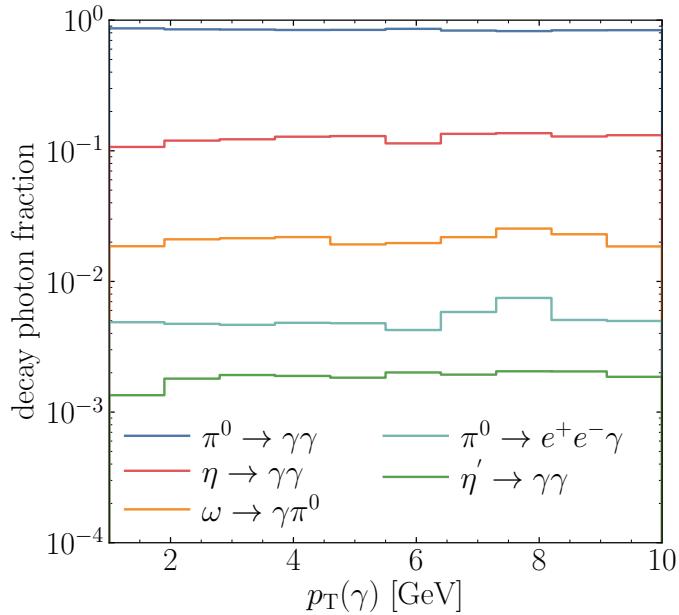


Figure 2-10: Relative contributions of the five most common sources of photons from hadron decays. Generated with CIMBA [78] using interpolation grids produced with PYTHIA8 [79].

PbPb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV [86], which has also been interpreted as evidence for thermal photon production.

In addition to ion-ion collisions, both the ALICE and PHENIX collaborations have studied direct photon production in smaller collision systems. PHENIX observed a direct photon signal in $\sqrt{s_{\text{NN}}} = 200$ GeV consistent with scaled pp measurements [87]. PHENIX has also produced preliminary measurements of direct photon production in high multiplicity dAu collisions, which show a hint of a direct photon excess consistent with thermal photon predictions [88]. The ALICE collaboration has studied direct photon production in pp collisions at $\sqrt{s} = 2.76$ and 8 TeV [89] and produced preliminary measurements in $p\text{Pb}$ collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV [90]. ALICE observes no significant direct photon signal with $p_T < 5$ GeV in any of these measurements.

In summary, there is a robust body of experimental evidence for excess direct photon production at $p_T < 3$ GeV in heavy ion collisions with a variety of beam energies and species. The interpretation of this excess, however, is uncertain. Both PHENIX and ALICE have measured elliptic flow of direct photons [91, 92]. Theoretical models have trouble replicating the observed large yields and elliptic flow simultaneously [93]. Furthermore, STAR and PHENIX direct photon measurements in AuAu collisions demonstrate a dramatic disagreement which has remained unresolved for more than a decade.

Chapter 3

The LHCb Experiment

The LHCb detector studies high energy hadron collisions at one of the LHC’s four interaction points. LHCb is designed specifically to study decays of hadrons containing c and b quarks. The design choices that allow LHCb to study heavy flavor decays also allow it to study more general hadron collider physics in kinematic regions that are inaccessible to the other LHC experiments. In this chapter I will discuss the design of the LHCb experiment and how it is used for both heavy flavor and general purpose hadron collider physics. This begins with the LHC, which is described in Section 3.1. Section 3.2 motivates studying heavy flavor physics at hadron colliders and presents the experimental requirements that motivate the LHCb detector’s design. The detector itself is described in Section 3.3. Finally the LHCb trigger, which selects which collisions to save, is described in Section 3.4.

3.1 The Large Hadron Collider

The Large Hadron Collider is a superconducting synchrotron occupying the 27 km ring created for the Large Electron Positron Collider. A detailed description can be found in Ref. [94]. The LHC is the largest particle accelerator in the world and is designed to accelerate two beams made up of bunches of protons in opposite directions. These beams can reach maximum energies of 7 TeV each, producing collisions with a center-of-mass energy of 14 TeV.

The LHC consists of eight 2.8 km arcs and eight 528 m straight sections. The arcs host the superconducting magnets used to bend and focus the LHC beams, and the straight sections host the accelerator’s access points. Four of the straight sections also house interaction points where the beams are brought together to produce collisions. These collisions are studied by detectors located at each interaction point. ATLAS [95] and CMS [96] are located at opposite sides of the LHC ring and are designed to collect data at the highest luminosity the LHC can provide in order to search for physics beyond the Standard Model at high energies. These are sometimes referred to as “general purpose detectors”. ALICE [97] and LHCb [98] are specialized detectors and are located at the interaction points on either side of the ATLAS experiment. The ALICE detector is designed to study heavy ion collisions while LHCb is built to study

heavy flavor decays.

The primary goal of the LHC is to search for physics beyond the Standard Model at high energies. Because no physics processes beyond the Standard Model have been observed at the LHC, if these processes occur they must be rare. As a result the LHC must maximize the rate of collisions as well as the center-of-mass energy of those collisions. The collision rate is defined by the accelerator's luminosity. For two identical bunched beams like those at the LHC, this is given by

$$\mathcal{L} = \frac{N^2 f}{4\pi\sigma_{\text{eff}}^2}, \quad (3.1)$$

where N is the number of protons in each bunch, f is the bunch crossing rate, and σ_{eff} is the effective RMS radius of the beams. The effective radius σ_{eff} includes a correction accounting for the affect of nonzero beam crossing angles and is given by

$$\sigma_{\text{eff}} = \sigma \sqrt{1 + \left(\frac{\sigma_z \phi}{2\sigma} \right)^2}, \quad (3.2)$$

where σ_z is the RMS beam width in the longitudinal direction and ϕ is the crossing angle between the two beams.

The LHC is designed to achieve a peak luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$. The need to maximize luminosity resulted in the choice to collide two beams of protons instead of beams of protons and antiprotons due to the higher intensities achievable with proton beams [99]. This is not the only possible choice, as a pair of proton-antiproton beams can share a vacuum and magnet system. This would be a much more convenient configuration for the narrow LEP tunnel than the two separate systems required by the LHC. Additionally the integrated luminosity delivered over a long period of LHC operations is limited by beam losses from collisions. As a result only four of the eight possible interaction points are used, and only ATLAS and CMS collect data at the highest possible luminosity [94].

The high beam intensities at the LHC result in multiple proton-proton collisions in each bunch crossing. The presence of multiple collisions in each bunch crossing is referred to as “pile-up”. Pile-up is also used to refer to the number of collisions per bunch crossing. The ATLAS and CMS detectors collected data with a peak pile-up of 37 in Run 1 [100] and up to 70 in Run 2 [101]. The performance of the LHCb detector decreases at high pile-up. In order to maintain data quality, the LHCb detector collects data with decreased luminosity. Lower luminosity is achieved using a transverse offset between the two beams at the interaction point [102]. The beam offset at LHCb was used to limit pile-up to an average of about 1.7 in Run 1 [103] and 1.1 in Run 2 [104].

The LHC began its first run of data taking (Run 1) in March of 2010 at a center-of-mass energy of 7 TeV. Data taking continued in 2011, and in April 2012 the center-of-mass energy was increased to 8 TeV. During Run 1 the LHC achieved a peak luminosity of $7.7 \times 10^{33} \text{ cm}^2\text{s}^{-2}$ [100]. ATLAS and CMS collected 28 fb^{-1} of data, while LHCb collected 3 fb^{-1} . Run 2 of the LHC began in 2015 at a center-of-

mass energy of 13 TeV and lasted until 2018. During Run 2 the LHC delivered about 160 fb^{-1} to ATLAS and CMS and about 6 fb^{-1} to LHCb [101].

In addition to regular proton-proton collisions, the LHC also produces heavy ion collisions. This includes lead-lead collisions at a center-of-mass energy per nucleon 2.76 TeV and 5.13 TeV, as well as proton-lead collisions at 5.02 TeV and 8.16 TeV. All four of the major LHC experiments study these collisions [6, 7]. In 2017, the LHC also produced xenon-xenon collisions at a center-of-mass energy per nucleon of 5.44 TeV. The LHCb detector also features the System for Measuring Overlap with Gas (SMOG), which is used for precision luminosity measurements [105]. SMOG can also be used to inject noble gases into the interaction region, allowing LHCb to study fixed-target proton-nucleus collisions.

3.2 How to Study B -Physics at a Hadron Collider

The primary purpose of the LHCb experiment is to study decays of hadrons containing bottom and charm quarks. The goal of this so-called heavy flavor physics is to precisely study decays of hadrons with masses of a few GeV in order to indirectly probe physics at scales of many TeV. Over the past two decades, successful heavy flavor physics programs have been carried out at electron-positron colliders, often referred to as B -factories after the heavy B mesons they produced [106]. These colliders collide electrons and positrons at a center-of-mass energy tuned to the $\Upsilon(4S)$ resonance. The $\Upsilon(4S)$ decays to pairs of B mesons, creating a sample of B meson pairs that can be studied with minimal background. Despite their success, the B -factories have limitations [107]. The current generation of e^+e^- heavy flavor experiments can only study B mesons containing a b quark accompanied by either a u or d , (called simply the B meson). The B -factories cannot produce heavier B mesons, such as the B_c meson, which contains b and c quarks¹. They also cannot produce baryons containing b quarks. All of these particles are produced in large amounts in high energy proton-proton collisions. Additionally, b and c quarks are produced at a much higher rate in proton-proton collisions than at the B -factories. A larger production rate could allow a heavy flavor experiment at a hadron collider to study rarer processes than those studied at electron-positron colliders. The primary obstacle to studying heavy flavor decays at a proton-proton collider is the messy collision environment. This means that a detector designed for studying heavy flavor decays at the LHC will look very different from the B -factory detectors, as well as from other LHC experiments.

The LHC produces $b\bar{b}$ quark pairs via hard QCD scattering, mostly between relatively low momentum gluons [108]. This results in $b\bar{b}$ pairs produced predominantly at angles close to the beamline. The pseudorapidity distributions of $b\bar{b}$ pairs produced in LHC collisions is shown in Figure 3-1. As a result, LHCb is constructed as a single arm forward spectrometer. This has the additional benefit simplifying the reconstruction of tracks left by charged particles. Instead of tracking curving tracks

¹The B -factories have collected data the $\Upsilon(5S)$ resonance, which can decay to pairs of B_s mesons containing b and s quarks. However, the B_s meson physics program at the B -factories is limited for reasons discussed in Ref. [106]

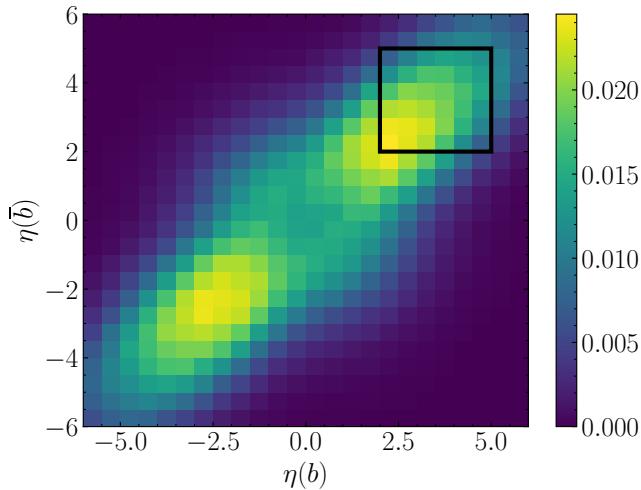


Figure 3-1: Pseudorapidity of $b\bar{b}$ pairs produced via gluon fusion in pp collisions at $\sqrt{s} = 13$ TeV. The LHCb detector acceptance is outlined in black. Generated using PYTHIA8 [79] and the NNPDF2.3 LO PDF set [111].

through a magnetic field, as in other LHC experiments [109, 110], tracks in LHCb can be roughly described as combinations of two straight track segments, one upstream of the magnet and the other downstream. Fast track reconstruction is particularly advantageous in the software trigger, which faces the difficult task of separating B decays from large backgrounds. This is discussed in detail in Chapter 4.

Ground state heavy flavor hadrons decay via the weak interaction. As a result, ground state b - and c - hadrons tend to have long lifetimes. Heavy flavor hadrons produced at the LHC will be highly boosted and will travel almost macroscopic distances before decaying. This means that the primary signature of a heavy flavor decay is a set of charged tracks meeting at a vertex displaced from a proton-proton collision point [107]². In order to identify these decays in the messy LHC environment, the LHCb detector must have excellent tracking near the beamline as close to the interaction point as possible. Finally, heavy flavor hadrons decay to a wide variety of final state particles. These particles must be correctly identified in order to accurately reconstruct the parent particle. This means the LHCb detector must be able to correctly identify all species of long lived charged hadrons and leptons.

3.3 The LHCb Detector

The LHCb detector is a single arm forward spectrometer. A description of the full detector can be found in Ref. [98]. The full detector has a nominal angular acceptance of $2 < \eta < 5$. LHCb's instrumented area includes less than 5% of the solid angle

²Although only ground state heavy flavor hadrons have long lifetimes, excited states will first decay to the ground state, which will then decay weakly and produce a secondary vertex.

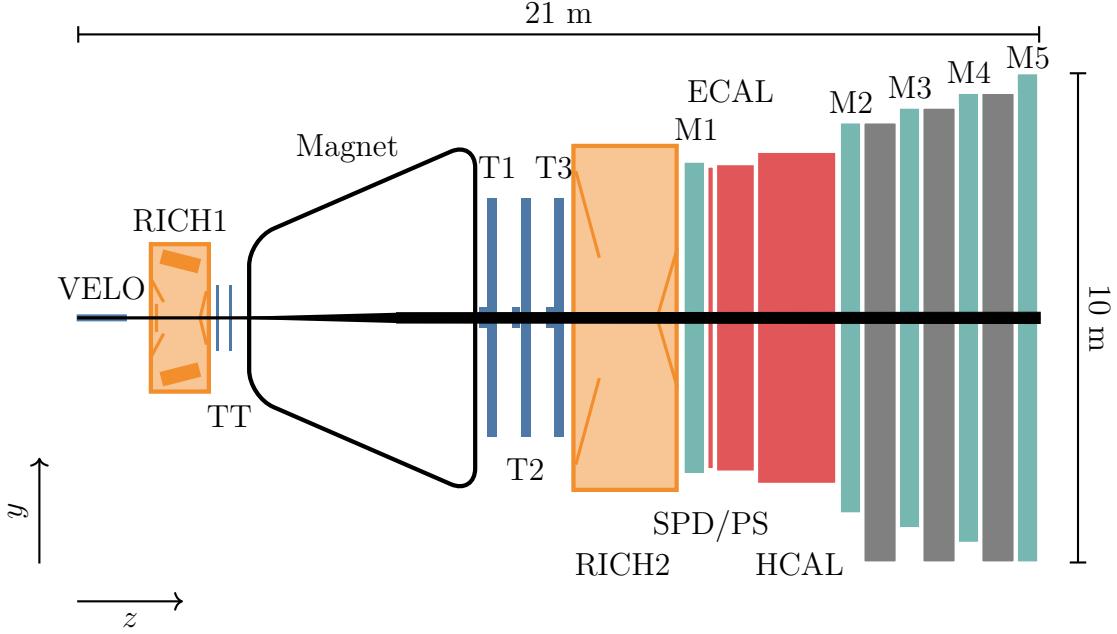


Figure 3-2: Schematic of the LHCb detector in the $y - z$ plane.

surrounding the interaction point but accepts about 25% of all produced $b\bar{b}$ pairs. Figure 3-2 shows a diagram of the full LHCb detector.

A tracking system made up of the Vertex Locator (VELO), Tracker Turicensis (TT), Inner Tracker (IT) and Outer Tracker (OT) measures the paths of charged particles and together with a dipole magnet measures their momenta. An electromagnetic calorimeter (ECAL) and hadronic calorimeter (HCAL) are used to measure particle energies. The calorimeters are used along with two ring imaging cherenkov detectors (RICH1 and RICH2) and five muon stations (M1-M5) to identify reconstructed particles.

The z axis of the LHCb coordinate system points from its origin at the center of the interaction region along the beamline towards the spectrometer. The y axis is perpendicular to the z axis and the plane of the LHC ring, pointing upward. The x axis is perpendicular to both of these, lying in the plane of the LHC ring and creating a right-handed coordinate system.

3.3.1 Tracking

Vertex Locator

Studying heavy flavor physics at the LHC requires distinguishing particles produced in decays of long lived b and c hadrons from those produced in the underlying pp collision. Additionally, many LHCb measurements require precise measurements of particle lifetimes (for example see Ref. [112]). As a result, the LHCb detector must have excellent primary and secondary vertex resolution and must precisely measure charged particle positions near the beamline. LHCb accomplishes this with the vertex

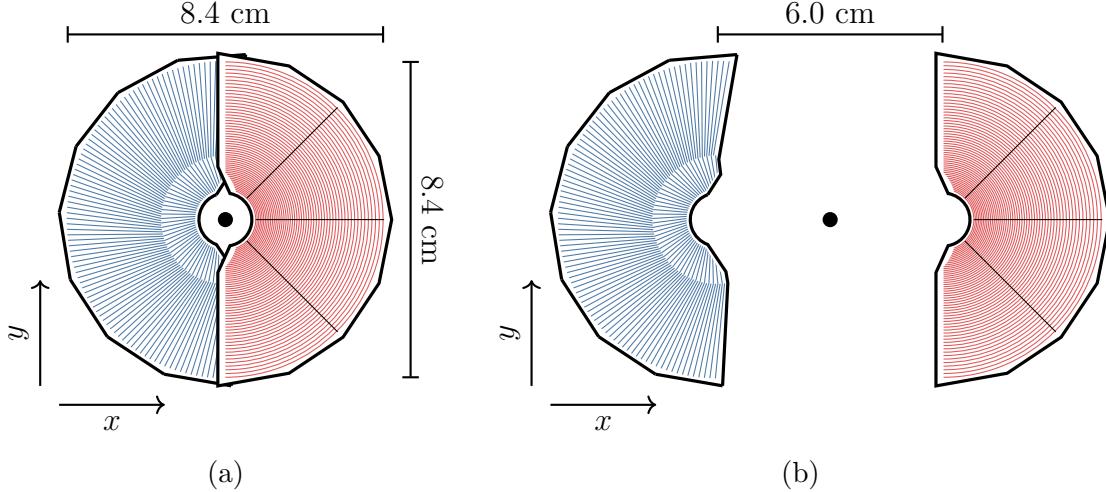


Figure 3-3: Schematic of the VELO detector in the closed (3-3a) and open (3-3b) configurations. The ϕ strips are shown in blue and the r strips are shown in red. Strip pitches are scaled up for legibility. The beamline is denoted by black circles. Adapted from Ref. [113].

locator (VELO).

The VELO is a silicon strip detector surrounding the interaction point [113]. LHCb's secondary vertex and track position resolution near the beamline increase with the distance between the beamline and the first track measurement. During stable beams the VELO is "closed" and the minimum distance between the beamline and the VELO's sensitive region is 8 mm. This is smaller than the transverse size of the LHC beams during beam injection, so the VELO can be "opened" to increase this distance to about 3 cm in order to avoid damage.

The VELO consists of two types of semicircular sensors. The r -sensors measure the radial coordinate r and ϕ -sensors measure the azimuthal angle ϕ . Figure 3-3 shows a $x - y$ projection of the r and ϕ sensors in the closed and open configurations. The r -sensors consist of four 45° segments with 512 strips each. These strips have a pitch of $38 \mu\text{m}$ closest to the beamline, increasing linearly to a maximum of $101.6 \mu\text{m}$. In order to reduce the strip pitch at the outer edge of the sensors, the ϕ sensors are divided into inner and outer regions. The strips in the inner has a strip pitch of $35.5 \mu\text{m}$ closest to the beamline, which increases to $78.3 \mu\text{m}$ as the strips run radially outward. The outer region begins at a radius of 17.25 mm and has a strip pitch of $39.3 \mu\text{m}$ at its innermost point, about half the pitch of the adjacent inner region. The strips of the ϕ sensors are also skewed relative to the radial in order to improve pattern recognition.

The VELO consists of 42 modules, each consisting of back-to-back r - and ϕ -sensors. The modules are arranged such that charged particles with $1.6 < \eta < 4.9$ originating from $|z| < 10.6$ cm will traverse the sensitive regions of at least three modules. The VELO also includes four pileup veto modules in the backward direction made up of only r -sensors.

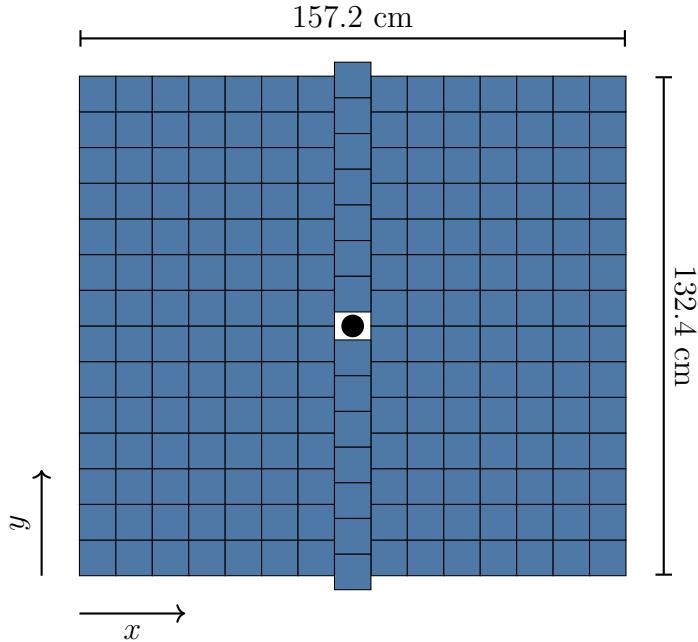


Figure 3-4: Schematic of the first TT plane. Each square represents a sensor consisting of 512 vertical silicon strips. The beamline is denoted by a black circle. Adapted from Ref. [114].

Tracker Turicensis

The Tracker Turicensis (TT) is a silicon strip detector located upstream of the LHCb magnet [114,115]. It is used in the LHCb software trigger to provide an initial momentum and charge determination for tracks, which allows for faster pattern recognition in the tracking stations downstream of the magnet. The TT consists of four detector layers grouped into two stations. The strips of the first and fourth layers are oriented vertically (known as x layers), while the second and third layers are rotated by $+5^\circ$ (u layer) and -5° (v layer), respectively. This maximizes precision in the bending plane of the magnet. The stereo angle provides sensitivity to the track's y position, which helps eliminate fake tracks.

Figure 3-4 shows the layout of the first TT layer. The TT layers are made up of half-modules. Each of these half-modules consists of 7 sensors arranged vertically with 512 vertical strips per sensor. The first two upstream layers are made up of 30 half-modules each, with 15 above and 15 below the $x - z$ plane. The two center half-modules are separated in y to accommodate the beampipe. The third and fourth layers have four additional half modules, extending further outward in the x direction.

Forward Trackers

The LHCb tracking system downstream of the magnet consists of an Inner Tracker (IT) [116] and Outer Tracker (OT) [117]. The IT is a silicon strip detector located close to the beamline. The OT is made up of straw drift tubes. These detectors are

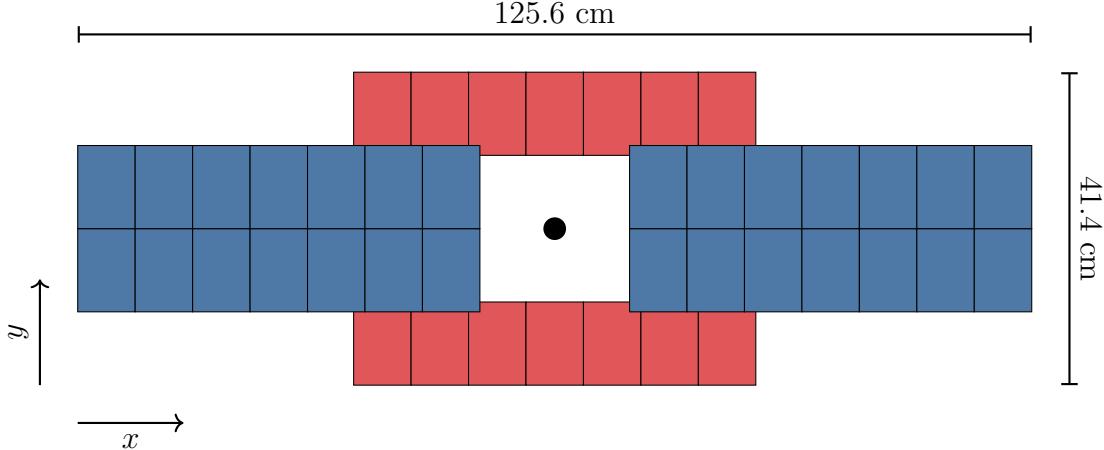


Figure 3-5: Schematic of an IT layer. Each shaded rectangle represents a sensor containing 384 vertical strips each. The IT consists of two sensor types of slightly different thicknesses. The thicker sensors are shown in red and the thinner sensors are shown in blue. Adapted from Ref. [116].

used in conjunction with the VELO, TT, and dipole magnet to precisely measure the momenta of charged particles. Both detector systems are arranged in three tracking stations (T stations) with four detector layers each. Each tracking station uses the same $x - u - v - x$ orientation used in the TT.

An IT x layer is shown in Figure 3-5. Each IT layer consists of 28 modules. The modules to the left and right of the beampipe consist of two sensors each, while the modules above and below consist of one sensor each. Each sensor has 384 vertical silicon strips with a pitch $198 \mu\text{m}$. The IT covers $4.5 < \eta < 4.9$ in the y direction and $3.4 < \eta < 5$ in the x direction.

Figure 3-6 shows an OT layer. The OT is made up of full size F modules and shorter S modules. Each module consists of two monolayers of 128 straw tubes each. The F module straw tubes are 4850 mm long, while the S module tubes are about half that. The straw tubes have an inner diameter of 4.9 mm and use a gas mixture of 70% argon and 30% carbon dioxide. This results in a drift time of less than 50 ns and a position resolution of about $200 \mu\text{m}$ in the bending plane.

Track Types

The LHCb detector can be used to reconstruct a number of different types of tracks [103]. *Long* tracks are the most commonly used tracks in analyses and consist of hits in the VELO, TT, and T stations. *Downstream* tracks are made up of TT and T stations hits but contain no hits in the VELO. These are often the products of long lived hadron decays, such as K_s^0 decays. Downstream tracks can also be used to reconstruct photons that convert in the VELO material, as in the analysis in Chapter 6. *Upstream* tracks are made up of hits in the VELO and TT, but contain no hits in the T stations. These are often low momentum tracks that are deflected out of detector acceptance in the magnetic field. *VELO* tracks consist entirely of hits in the

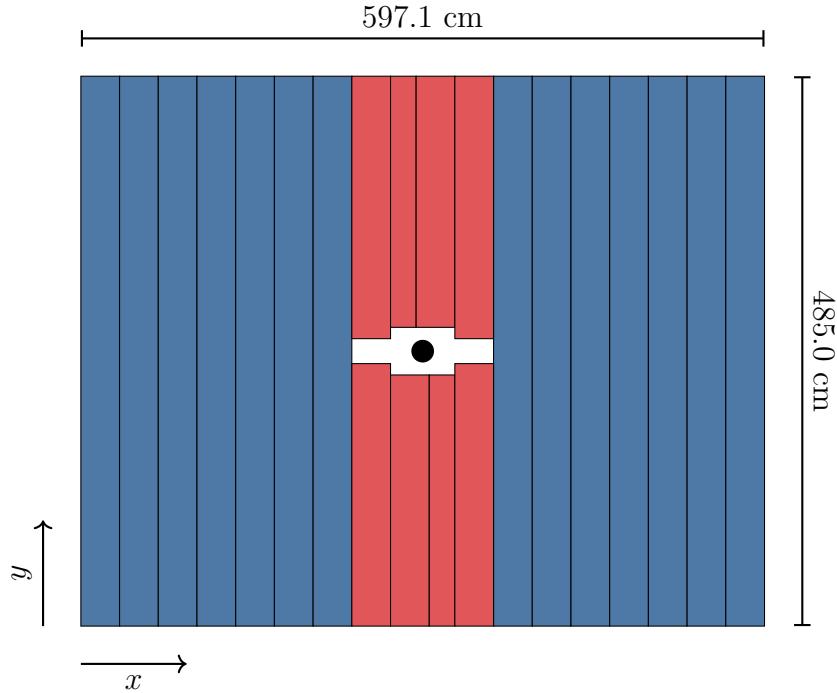


Figure 3-6: Schematic of an OT x layer. F -type modules are shown in blue and S -type modules are shown in red. The beamline is denoted by a black circle. Adapted from Ref. [117].

VELO. VELO tracks can be both forward and backward and are used to reconstruct primary vertices. T tracks contain only hits in the T stations. Figure 3-7 illustrates these different track types in the tracking system.

3.3.2 Particle Identification

Calorimeters

The LHCb detector includes both electromagnetic (ECAL) and hadronic (HCAL) calorimeters [118]. These calorimeters provide energy measurements for LHCb's hardware level trigger (L0) and provide energy measurements and particle identification information. Figure 3-8 shows the cell layout of the LHCb calorimeters.

The face of the LHCb ECAL is located about 12 m from the interaction point. The ECAL is longitudinally segmented into a Scintillating Pad Detector (SPD), Preshower Detector (PS), and main ECAL volume. The SPD is the first longitudinal segment. A 15 mm (about 2.5 radiation lengths) lead absorber separates the SPD and PS, which is in turn followed by the ECAL. All three segments use the same cell layout, but the cells are scaled to provide a one-to-one-to-one projective correspondence between cells in each detector.

Both the SPD and PS consist of 15 mm polystyrene pads read out by wavelength shifting fibers. The ECAL is a sampling calorimeter made up of alternating 2 mm layers of lead and 4 mm layers of scintillator. The ECAL has a total depth of 42 cm,

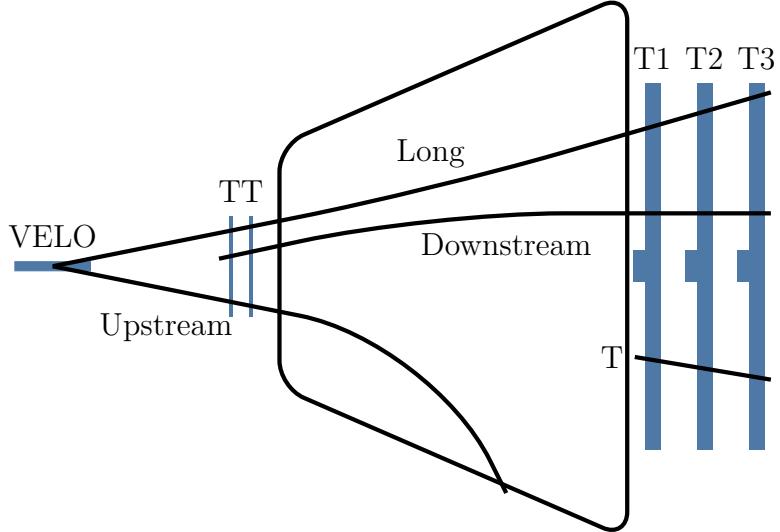


Figure 3-7: Illustration of different LHCb track types in the $x - z$ plane. Adapted from Ref. [103].

corresponding to 25 radiation lengths. These are read out by wavelength shifting fibers penetrating the scintillators and absorbers and running perpendicular to the detector face in a shashlik design.

The ECAL system must differentiate photons from electrons, and electrons from other charged particles. The SPD has a depth of 0.28 radiation lengths, so most photons will not produce a SPD signal. They then shower in the lead absorber and leave an electromagnetic shower in the PS and ECAL. Charged particles will leave a signal in the SPD, PS and ECAL. The PS and ECAL signals provide information about the shape of the electromagnetic shower, which can be used to distinguish electrons from other charged particles. The ECAL is also used to measure photon energies. Optimal energy resolution is vital for many LHCb measurements, such as those involving radiative B hadron decays. The LHCb ECAL provides an energy resolution of $\sigma_E/E = 10\%/\sqrt{E/\text{GeV}} \oplus 1\%$ [119]. The ECAL's photon energy measurement and electron identification capabilities are both vital for the measurement presented in Chapter 6.

The HCAL is a sampling calorimeter made up of tiles of iron and scintillator. The tiles are oriented parallel to the beamline in the $y - z$ plane. Each tile is $197 \times 256 \text{ mm}^2$ or $197 \times 127 \text{ mm}^2$ and 1 cm thick. These tiles are arranged in alternating layers in both the transverse and longitudinal direction. The HCAL is used primarily in L0 hadronic trigger algorithms, which select events with high- E_T hadrons. This binary decision requires only modest energy resolution, so the HCAL depth is only 5.6 hadronic interaction lengths. The HCAL provides an energy resolution of about $\sigma_E/E = 70\%/\sqrt{E/\text{GeV}} \oplus 10\%$ [98].

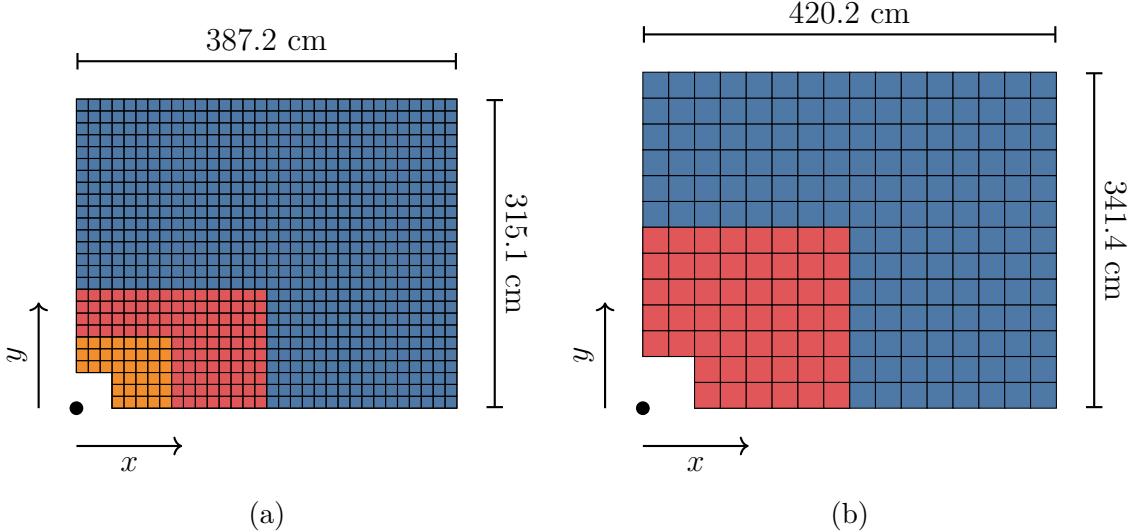


Figure 3-8: Schematic of the upper right quadrants of the ECAL (a) and HCAL (b). The PS and SPD share the same segmentation as the ECAL but are 0.5% and 1% smaller, respectively. In both diagrams blue squares contain a single calorimeter cell and red squares contain 2x2 arrays of cells. The orange squares in a contain 3x3 arrays of cells. The beamline is denoted by the black circles.

Muon System

Most long-lived charged particles will deposit all of their energy in either the ECAL or HCAL. Muons, however, do not lose large amounts of energy through radiation in the ECAL and do not interact with the HCAL via hadronic interactions. As a result muons consistently penetrate both the ECAL and HCAL. The LHCb Muon System [120] is made up of five tracking stations, four of which are located downstream of the HCAL and interspersed with iron absorbers. The Muon System identifies charged particles as muons based on their penetration depth.

Figure 3-9 shows the layout of the upper right quadrant of a muon station. The stations are denoted M1-M5. Station M1 is located in front of the calorimeter systems, while stations M2-M5 are located downstream of the calorimeters. Each muon station consists of four regions denoted R1-R4. Region R1 of station M1 is made up of gas electron multiplier (GEM) chambers, while the rest of the detector is made up of multi-wire proportional chambers (MWPCs).

The Muon System is used in the L0 trigger to identify high- p_T muons [121]. Station M1 is located in a region with a small magnetic field. This allows for muon momentum measurements using only the Muon System and provides a momentum resolution of about $\sigma_{p_T}/p_T = 20\%$. To facilitate this measurement, stations M1-M3 are more highly segmented than stations M4 and M5, which are used primarily for particle identification and don't require precise position measurements.

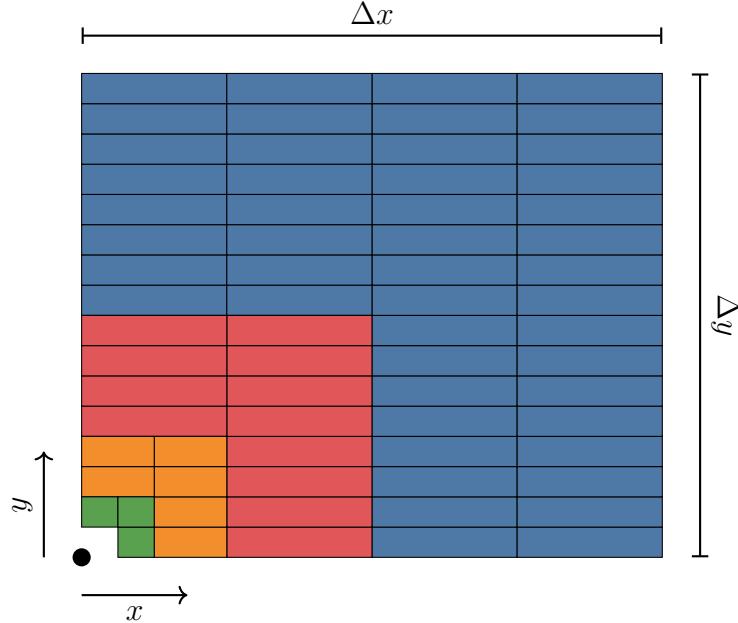


Figure 3-9: Schematic of the upper right quadrant of a Muon station. The R4 region is shown in blue, R3 in red, R2 in orange, and R1 in green. The dimensions Δx and Δy vary for each Muon station and are described in the text. The beamline is denoted by a black circle.

Ring Imaging Cherenkov Detectors

Heavy flavor hadrons frequently decay to hadronic final states containing charged pions and kaons. Successfully studying these decays requires correctly identifying these hadrons. LHCb accomplishes this using Ring Imaging Cherenkov (RICH) detectors [122]. A charged particle emits cherenkov light when its speed exceeds the phase velocity of light in the medium. The light is emitted in a cone around the charged particle with a velocity-dependent opening angle θ given by

$$\theta = \arccos \left(\frac{1}{n\beta} \right), \quad (3.3)$$

where n is the medium's index of refraction and $\beta = v/c$, where v is the particle's velocity. As a result, measuring the opening angle of the Cherenkov cone provides a measurement of the particle's velocity. Combined with a momentum measurement in the tracking system, this can be used to identify the charged particle.

The LHCb RICH detectors consist of media with $n > 1$, along with photodetectors used to measure the emitted cherenkov light. To study $K - \pi$ separation power, it is useful to calculate the difference in opening angles produced by kaons and pions of the same momentum. This is given by

$$\Delta\theta_{K\pi} \approx \frac{1}{\sqrt{1 - \frac{1}{n^2}}} \left(\frac{1}{\gamma_K^2} - \frac{1}{\gamma_\pi^2} \right), \quad (3.4)$$

where γ_K and γ_π are the relativistic γ factors for the kaon and pion mass hypotheses, respectively. While this separation increases as n approaches unity, photodetectors measure a limited range of opening angles. As a result, the index of refraction is tuned to identify particles in a certain momentum range. Smaller indices of refraction are used to measure higher momentum particles. Because of this, LHCb uses two RICH detectors with different radiating media to identify particles over a wide momentum range.

RICH1 is located between the VELO and TT and covers the full LHCb acceptance. RICH1 uses a C_4F_{10} radiator with $n = 1.0014$ at a wavelength of 400 nm. Until 2012, RICH1 also used an aerogel radiator with $n = 1.03$ at a wavelength of 400 nm. RICH1 is used to identify particles with momenta between 1 and 60 GeV. RICH2 is located between the T stations and the SPD. RICH2 uses a CF_4 radiator with an index of refraction of $n = 1.0005$ at 400 nm and is designed to identify particles with momenta greater than 60 GeV. Most high-momentum particles are produced at angles close to the beamline. As a result, RICH2 has a limited acceptance of $15 < \theta < 120$ mrad.

3.4 The LHCb Trigger

The LHC produces a nominal bunch crossing rate of 40 MHz when running with 2808 bunches per beam. In Run 1, the LHC ran with 1296 bunches [100]. This corresponded to a bunch crossing rate of 20 MHz. The number of bunches increased to about 2600 by the end of Run 2, approaching the LHC’s nominal bunch crossing rate [101].

During Runs 1 and 2, LHCb used a hardware trigger utilizing the calorimeters and Muon System to reduce the event rate to about 1 MHz [121]. Passing events were then sent to a CPU farm, which ran a two stage software trigger. The first stage of the software trigger (HLT1) used a partial event reconstruction to further reduce the event rate to about 30 – 50 kHz in Run 1 and about 100 kHz in Run 2. Finally the second stage (HLT2) performed a more complete reconstruction and selected events to write to storage. During Run 1 HLT2 selected events at a rate of about 5 kHz. This more than doubled in Run 2 to 12.5 kHz [104]. Figure 3-10 shows the trigger sequences for Runs 1 and 2. The work presented in Chapter 4 builds off of the Run 2 trigger model. Furthermore the analysis presented in Chapter 6 uses data collected during Run 2. Because of this I will focus on the Run 2 trigger design here. This is also described in full in Ref. [104].

3.4.1 L0

The L0 trigger uses hardware algorithms that select events based on information from the ECAL, HCAL, or Muon System. This includes 6 algorithms, or “lines.”

- **L0Hadron** uses the HCAL to select events with high-energy hadrons. Events are selected based on the 2x2 HCAL cluster with the highest E_T in the event,

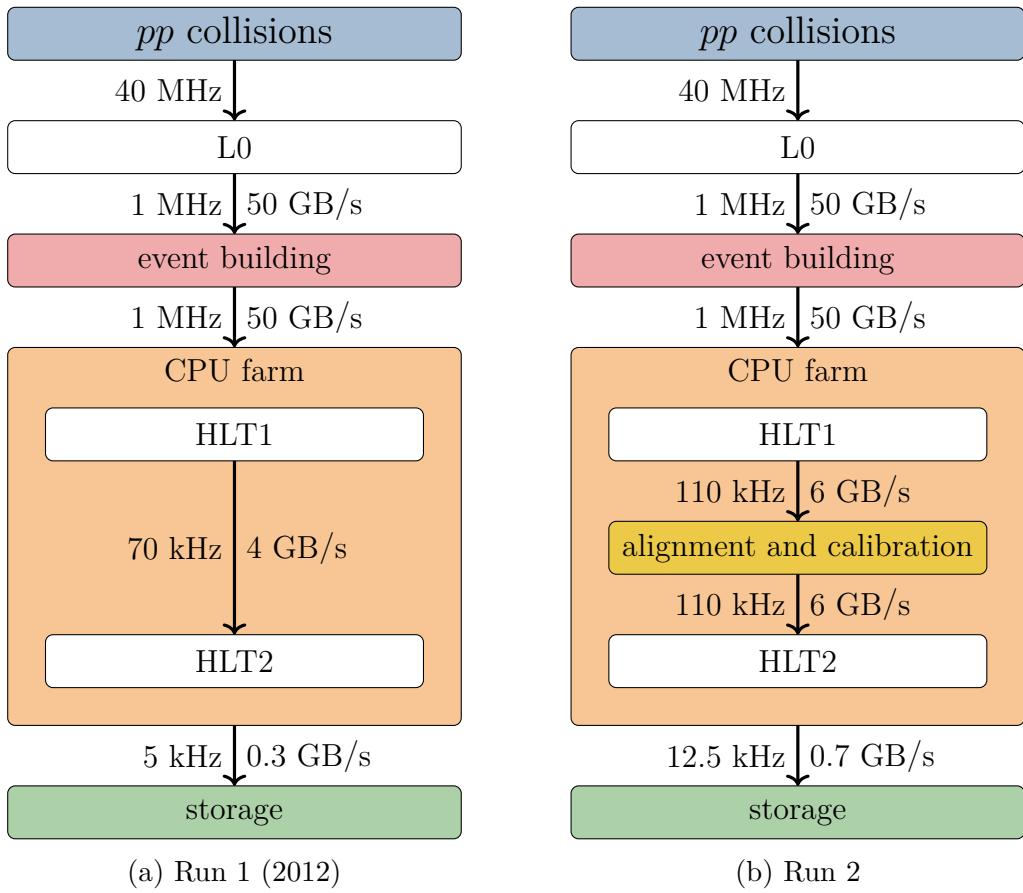


Figure 3-10: Data acquisition and trigger sequences for Runs 1 (a) and 2 (b). The trigger sequence evolved throughout Run 1, and a shows the trigger as it operated in 2012.

where E_T is given by

$$E_T = \sum_{i=1}^4 E_i \sin \theta_i, \quad (3.5)$$

where i denotes the HCAL cell, E_i is the measured energy in each cell, and θ_i is the angle between the z axis and a line connecting the cell and the LHCb origin. If the highest E_T ECAL cluster in the event is located in front of the highest E_T HCAL cluster, the ECAL cluster energy is included in the E_T calculation.

- **L0Photon** and **L0Electron** use the ECAL to select events with a high-energy electron or photon. These algorithms use the E_T of 2x2 ECAL clusters. Clusters are identified as electrons or photons using signals from the SPD and PS.
- **L0Muon**, **L0MuonHighPt**, and **L0DiMuon** select events with muons based on hits in the Muon System. Muon tracks are formed by looking for hits in each muon station forming a straight line. The track momentum is estimated by assuming the muon originates from the LHCb origin and is deflected by a single kick in the magnetic field. **L0Muon** selects events based on the p_T of the highest p_T muon in the event, while **L0DiMuon** uses the product of the transverse momenta of the two highest p_T muons in the event. **L0MuonHighPt** was added in Run 2 for electroweak physics and selects events with a higher p_T requirement than **L0Muon**.

Each Run 1 L0 line used the number of SPD hits in the event to exclude high-multiplicity. This Global Event Cut (GEC) speeds up event reconstruction by removing events that take a large portion of reconstruction time and suffer from reduced reconstruction performance. In Run 2, **L0MuonHighPt** did not use a GEC in order to reduce systematic uncertainties in electroweak measurements from estimating its efficiency. Additionally, **L0DiMuon** used a much looser GEC than the other L0 lines. The p_T and E_T requirements for the L0 lines are used to tune the L0 bandwidth division. In Run 2 the bandwidth division was tuned via an automated procedure that optimized signal efficiencies using various interesting decay channels.

3.4.2 HLT1

HLT1 performs a partial reconstruction of charged tracks to further reduce the event rate. Hits in the VELO are combined to form VELO tracks. These are used to reconstruct primary vertices (PVs) and are propagated to the TT to form upstream tracks. An initial momentum estimate is made and low momentum tracks are rejected. The track is then extrapolated to the T stations. The search region for T station hits is limited by requiring long tracks to have $p_T > 500$ MeV. The resulting long tracks are fit using a Kalman filter, and the fit quality is used to reject fake tracks.

Long tracks are extrapolated to the Muon System and matched to hits in order to identify muons. Only tracks with $p_T > 500$ MeV are reconstructed as long tracks, but all muons with $p > 3$ GeV can potentially reach the Muon System and can be successfully identified. To take advantage of this, upstream tracks are extrapolated

directly to the Muon System and matched to hits in the muon station. If the upstream track is identified as a muon, it is extrapolated to the T stations, where unused hits are added to the track. This allows for muon identification down to $p_T = 80$ MeV.

The primary aim of HLT1 is to select events with particles that decay far from a primary vertex. It uses two inclusive trigger lines to do this, one selecting events based on single tracks and the other based on two-track secondary vertices. `Hlt1TrackMVA` selects events with high- p_T displaced tracks based on a hyperbolic boundary in the p_T -displacement plane. `Hlt1TwoTrackMVA` uses a MatrixNet classifier that takes as input the vertex fit quality, vertex displacement, the individual track p_T , and individual track displacement. Optimization of these lines is described in Ref. [123].

HLT1 also selects events with muonic heavy flavor and electroweak boson decays using five muon trigger lines. Two of these select events based on single muon candidates. `Hlt1TrackMuonMVA` is a looser version of `Hlt1TrackMVA`, using a similar hyperbolic requirement in the p_T -displacement plane. `Hlt1SingleHighPtMuon` selects high p_T muons with no displacement requirement and is used for electroweak physics. The three remaining lines select events based on dimuon secondary vertices. `Hlt1DisplacedDimuon` selects displaced dimuon vertices, and `Hlt1HighMassDimuon` selects dimuons with masses around the J/ψ mass and higher with no displacement requirement. Finally, `Hlt1LowPtDimuon` selects soft dimuon candidates with a very tight displacement requirement. These lines only run on events passing the L0 muon lines except `Hlt1LowPtDimuon`, which runs on all events passing L0.

The combined inclusive and muon lines cover most of LHCb’s physics program. In addition special lines that select low multiplicity events are used to study central exclusive production, and exclusive lines are used to select two-body heavy flavor decays without biasing their lifetime measurement with displacement requirements. HLT1 also has dielectron lines, which are used for both heavy flavor physics and searches for exotic particles decaying to e^+e^- pairs. HLT1 also has various calibration lines that provide data for aligning the LHCb subdetectors.

3.4.3 HLT2

Events passing HLT1 are then processed by HLT2. HLT2 performs a full event reconstruction, beginning by repeating the HLT1 charged track and PV reconstruction. After the HLT1 reconstruction is repeated, remaining VELO tracks are extrapolated to the T stations to form long tracks with no p_T threshold. After this, unused hits in the T stations are combined into T tracks and extrapolated to the VELO to create additional long tracks. T tracks not matched to VELO tracks are extrapolated to the TT to form downstream tracks. As in HLT1, these tracks are fit using a Kalman filter and fake tracks are rejected using the fit quality. In addition to charged tracks, HLT2 uses information from the ECAL and HCAL to reconstruct neutral particles. HLT2 also uses the RICH detectors and calorimeters to perform particle identification on all reconstructed particles.

The algorithms used in HLT2 in Run 2 are of the same quality as those used in offline analyses. In Run 2, a 10 PB buffer was added between HLT1 and HLT2, allowing for real time calibration and alignment of the LHCb detector. This in turn allows

for offline-quality reconstruction in HLT2. As a result, HLT2 candidates themselves can be used in analyses. This has led to the division of HLT2 lines into FULL and TURBO streams. Events selected by lines in the FULL stream are saved to disk in their entirety, while only a subset of reconstructed objects are saved for events selected by TURBO lines [124]. This allows HLT2 to support a broader physics program while saving disk space.

The LHCb B physics program is driven by a few inclusive topological trigger lines [125]. These lines use a MatrixNet classifier to select events with displaced two-, three-, and four-track secondary vertices. The topological trigger is trained to select events with reconstructible b hadron decays [123]. They are trained to reject b hadrons with decay products outside of detector acceptance, charm hadrons, and particles originating from the primary vertex. Loosened thresholds are used when constituent tracks are identified as electrons or muons.

HLT2 also features a wide range of muon and dimuon lines used for heavy flavor physics, strange physics, electroweak physics, and searches for exotic particles. Additionally HLT2 selects events using hundreds of exclusive lines designed to select fully reconstructed decays. These form the basis of LHCb's charm physics program. Most of these lines are in the TURBO stream and only the trigger candidate itself is saved.

Chapter 4

The LHCb GPU High Level Trigger

The LHC is scheduled to resume collisions in 2022 after a long shutdown and upgrade. The upgraded accelerator will produce proton-proton collisions at a center of mass energy of 14 TeV. The LHCb detector will be upgraded as well, featuring an almost entirely new tracking system. The improved tracking system will allow the LHCb detector to collect data at luminosities five times larger than in Runs 1 and 2. In addition, LHCb’s front end electronics will be upgraded to allow the full detector to be read out at the LHC’s 40 MHz bunch crossing rate. As a result, the upgraded LHCb detector will collect data without a hardware trigger and will require a software trigger capable of reconstructing and selecting events at the full inelastic event rate of 30 MHz.

This chapter presents Allen, a software trigger developed for graphics processing units (GPUs). Allen has been designed to meet the triggering needs of the LHCb physics program and will be used as LHCb’s first level software trigger in Run 3. Section 4.1 begins by introducing the problem of triggering on heavy flavor decays at the LHC and includes a brief description of the upgraded LHCb detector. Section 4.2 introduces the Allen project, and section 4.3 details the algorithms used in the Allen trigger sequence. Finally section 4.4 presents Allen’s computing and physics performance.

Throughout this chapter, I will cover in detail the areas of the Allen project in which I played a central role. This includes the Kalman filter [126], secondary vertex fitter, and physics selection framework, as well as studies of the physics performance of the entire reconstruction and selection sequence. I will briefly summarize the other aspects of the Allen framework, which reflect the efforts of dozens of developers. The full Allen framework is summarized in Ref. [127] and described in detail in Ref. [128].

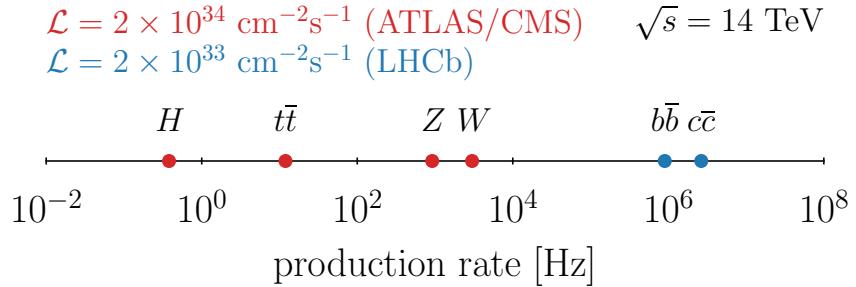


Figure 4-1: Production rates for various particles at the Run 3 LHC design energy and luminosity. Cross sections are calculated using `MADGRAPH5_AMC@NLO` [130,131] and the NNPDF3.1 NLO PDF set [132].

4.1 The LHCb Upgrade

4.1.1 The Upgrade Trigger Strategy

One of the primary motivations for studying heavy flavor physics at the LHC is the huge production rate of b and c hadrons in high-energy hadron collisions. Figure 4-1 shows the production rates of processes of interest to the various LHC experiments. General purpose detectors at the LHC study processes such as production of Higgs bosons, electroweak bosons, and top quarks. These occur at rates that vary from once every few hundred seconds to a few thousand times per second at the LHC’s Run 3 luminosity of $\mathcal{L} = 2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This is orders of magnitude below the production rate of b and c hadrons, which will exceed 1 MHz in LHCb acceptance [129]. Saving every event with a heavy flavor hadron at LHCb would produce about 200 GB/s. This amount of data is impossible to store using available resources.

This means that the goal of the LHCb trigger is inherently different from that of triggers at general purpose hadron collider experiments. Instead of picking out small signals from large backgrounds, the LHCb trigger must distinguish between similar signals and preferentially select more interesting decay channels. To this end, the two stages of the upgraded LHCb trigger roughly correspond to two different tasks:

- *HLT1* is the first stage of the software trigger. It performs a partial reconstruction of charged tracks and uses these for initial background reduction. HLT1 uses inclusive selections designed to select all heavy flavor decays in LHCb acceptance.
- *HLT2* performs a full reconstruction, including the calorimeters and RICH detector systems. HLT2 uses this offline-quality reconstruction to perform inclusive and exclusive selections of desirable b and c hadron decays.

Throughout this chapter, I will focus on the design and implementation of the first stage of this software trigger.

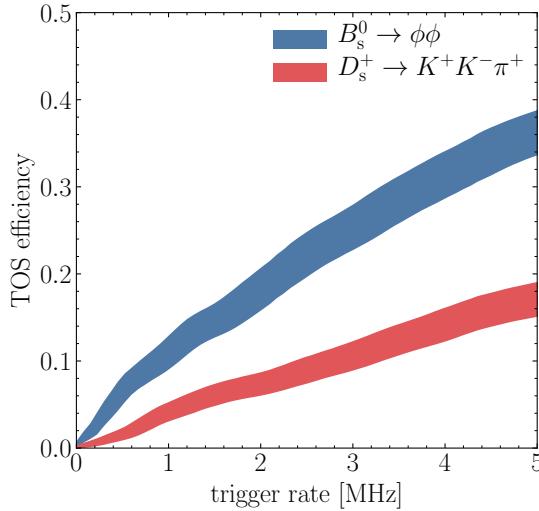


Figure 4-2: Efficiency as a function of rate for a simulated hardware trigger. Passing events must have at least one generated stable hadron in LHCb acceptance satisfying a p_T requirement. The trigger rate is adjusted by changing the p_T threshold. Passing signal events must have at least one final state hadron satisfying the p_T requirement. Efficiencies are calculated with respect to reconstructible signal decays. The line thickness shows the statistical uncertainty.

Ground state heavy flavor hadrons decay via the weak interaction and have relatively long lifetimes. Because of the high energies of LHC collisions, these hadrons will have large momenta and consequently larger lifetimes in the detector’s reference frame. Thus b and c hadrons will fly a distance of about a centimeter before decaying. The resulting final states usually include charged leptons and long-lived hadrons with momenta comparable to those of particles produced in the soft QCD processes that occur in most LHC collisions. These low-momentum particles make up the primary backgrounds in HLT1.

The similarity between signal and background makes triggering on heavy flavor decays in hadron collisions challenging. In Runs 1 and 2, LHCb’s hardware triggers relied on information from single subdetectors, such as the muon system or calorimeters. To show how this will perform in Run 3, we can study the performance of a hypothetical trigger line that selects events with a stable hadron passing a p_T threshold. This is comparable to the hadronic triggers used in LHCb’s L0 trigger in Runs 1 and 2. These lines selected events based on the highest E_T 2x2 HCAL cluster in the event [133]. Figure 4-2 shows the performance of a L0-like algorithm in simulated Run 3 collisions for two benchmark hadronic decay modes. In Runs 1 and 2, LHCb’s L0 operated at a trigger rate of 1 MHz. Figure 4-2 shows that this is expected to result in efficiencies of about 10% for hadronic B decays and less than 5% for hadronic charm decays.

Instead of relying only on high- p_T signatures, LHCb can successfully trigger on heavy flavor decays by selecting events with charged tracks originating far from a pp collision. This strategy requires information from the entire tracking system and is

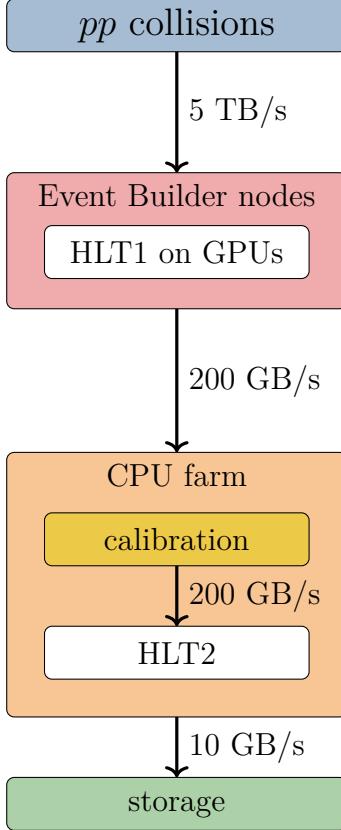


Figure 4-3: Diagram of the LHCb data acquisition system for Run 3. In Run 3, the hardware level L0 trigger will be removed and HLT1 will be run on GPUs hosted in the Event Builder nodes.

consequently difficult to implement in a hardware trigger. As a result, the LHCb experiment will run without a hardware trigger in Run 3 and will process the full 30 MHz inelastic event rate in software. Figure 4-3 outlines LHCb data acquisition sequence for Run 3.

4.1.2 VELO

The upgraded VELO detector [134] is similar in design to the original, but uses silicon pixels instead of microstrips. Figure 4-4 shows a schematic of the upgraded VELO detector. The detector will consist of 26 layers of pixel detectors made up of two modules each. Each module is made up of twelve square sensors consisting of 256×256 grids of $55 \mu\text{m}$ square pixels. These sensors are arranged in an “L” shape as shown in Figures 4-4a and 4-4b. As in Runs 1 and 2, the VELO is the primary source of information for determining the locations of pp collisions and particle decays, which are the basis of the HLT1 trigger strategy. The minimum distance from the beamline to the active region of the detector will be reduced from 8.2 mm in Runs 1 and 2 to 5.1 mm, resulting in improvements in impact parameter resolution over the original VELO.

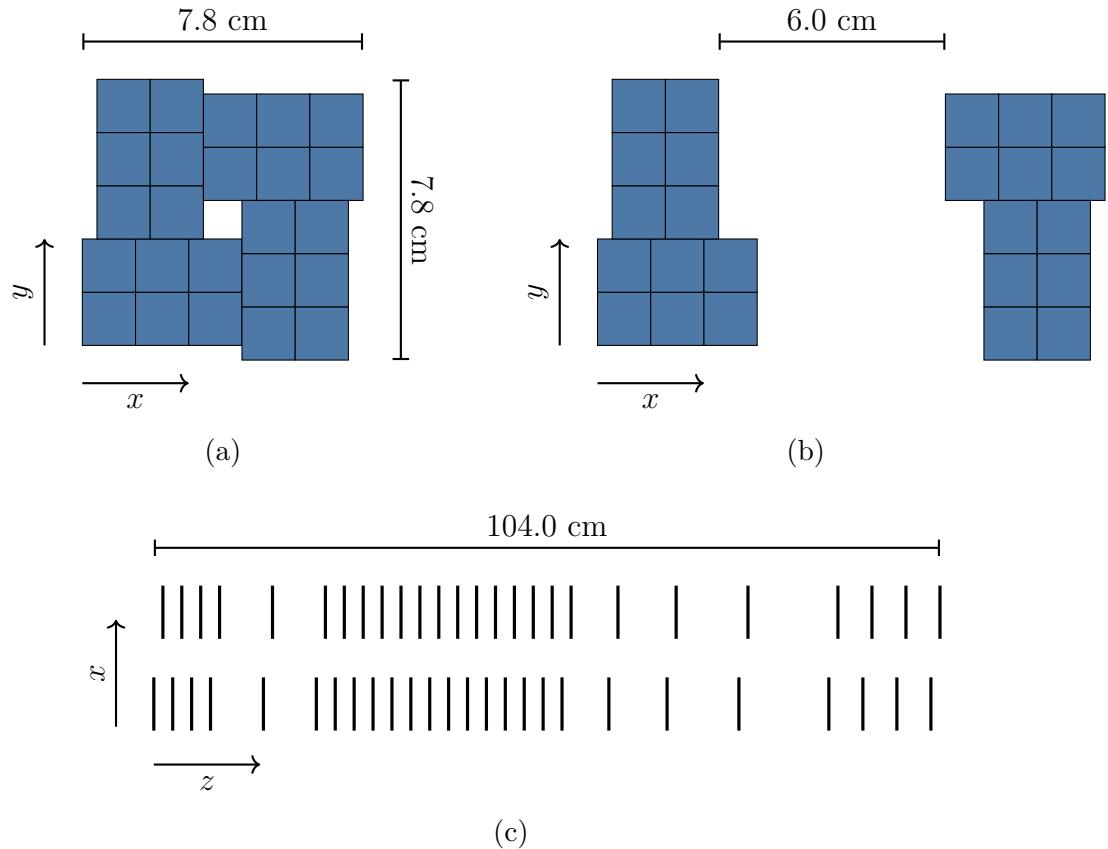


Figure 4-4: Schematic of the upgraded VELO detector. The open and closed detector configurations are shown in a and b, respectively. Each blue square represents a 256×256 grid of silicon pixels. The relative position of the layers along the z axis is shown in c. Adapted from Ref. [134].

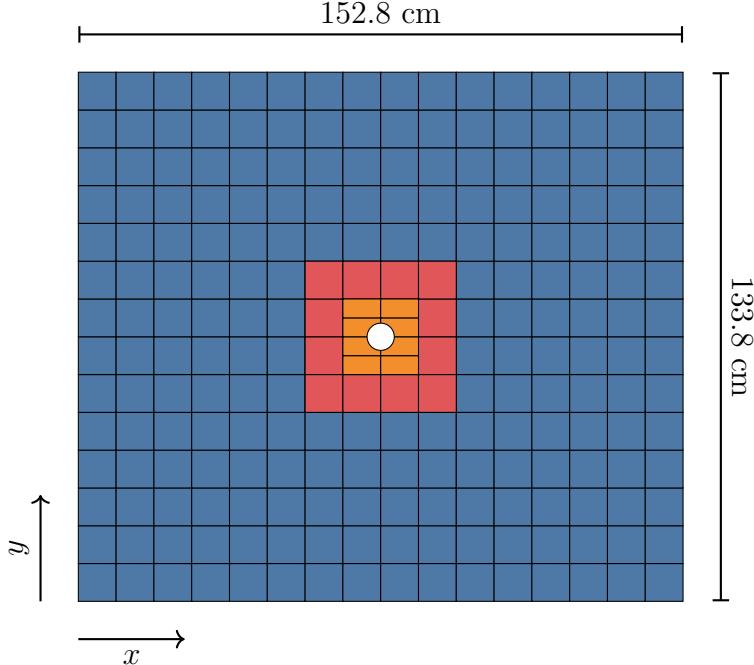


Figure 4-5: Schematic of the upstream UT detector layer in the x configuration. Each blue square represents a sensor with 512 silicon strips. The red squares represent sensors with half pitch strips, and orange rectangles represent sensors with half pitch and half length. The downstream layers share the same layout, but have two additional staves of sensors. The beampipe is shown as a white circle at the center of the diagram. Adapted from Ref. [135].

4.1.3 Upstream Tracker

The Upstream Tracker (UT) [135] is a silicon strip detector built to replace the TT, which is positioned upstream of the magnet. Figure 4-5 shows a diagram of a UT detector plane. The full detector consists of four planes of silicon strips, each made up of 16 or 18 vertical staves. Each stave consists of 14 square sensors. Most of these sensors measure 98.88 mm per side and contain 512 strips with a 190 μm pitch. Sensors near the center of the detector may be half width and half pitch, and those closest to the beampipe are shaped to form a circular aperture around the beam. The first and fourth UT layers are oriented vertically (known as x layers), while the second and third layers are oriented at $+5^\circ$ and -5° , respectively (u and v layers). The UT is used in HLT1 to provide a first momentum estimate and charge determination for tracks, which is used to determine the search region for tracks at the forward tracker. Furthermore, requiring tracks to have UT hits greatly reduces the fake track rate. The UT's angular acceptance is only about 95% that of the SciFi, so this requirement results in a decrease in tracking efficiency.

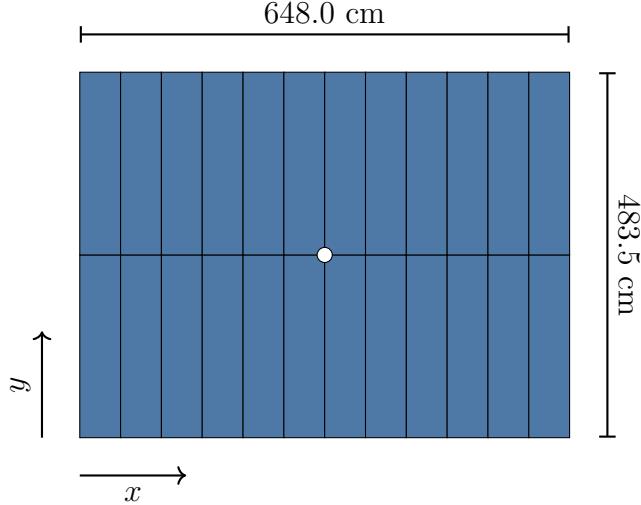


Figure 4-6: Schematic of a SciFi detector layer in the x configuration. Each rectangle represents a single module consisting of upper and lower fiber mats made up of 2.4 m scintillating fibers. The beampipe is shown as a white circle in the center of the detector. Adapted from [135].

4.1.4 Scitillating Fibre Tracker

The Scintillating Fibre Tracker (SciFi) [135] will serve as a forward tracker downstream of the beamline. The SciFi consists of three tracking stations with four detector layers each. The layers in each station are arranged in the same $x - u - v - x$ orientation as the UT. The detector layers are made up of 2.4 m long scintillating fibers with a diameter of 250 μm . The combination of VELO, UT, and SciFi tracks provides precise momentum measurements for charged particles.

4.1.5 Muon System

Muon identification is vital for triggering on leptonic and semileptonic heavy flavor decays, as well as decays of electroweak bosons and possible exotic particles such as dark photons [136]. Because the muon stations are well shielded, minimal changes are necessary to handle the increased luminosity of Run 3 [137]. Additional shielding will be placed in front of the second muon station (M2) where the particle flux is highest. In addition, the first muon station (M1) was primarily used for momentum measurements in L0. This will no longer be necessary in Run 3, so M1 will be removed in order to reduce the amount of material upstream of the calorimeters.

4.1.6 HLT1 Reconstruction Requirements

The basic requirements of the upgraded LHCb trigger system are described in Ref. [133]. In order to effectively trigger on heavy flavor decays, HLT1 must reconstruct two basic types of objects.

- *Charged tracks*: The trigger must read out data from the entire tracking system, cluster the signals into hits, and use those hits to reconstruct tracks.
- *Vertices*: This includes primary and secondary vertices. A primary vertex (PV) is a reconstructed collision point, while a secondary vertex (SV) is a reconstructed decay candidate.

The trigger can then select events based on the presence of tracks and SVs displaced from any primary vertex. Displacement is determined by measuring the track's minimum impact parameter with respect to each primary vertex in the event (IP). The significance of this impact parameter (χ_{IP}^2) is the primary discriminator used to separate prompt and displaced tracks. This requires precise reconstruction of tracks and PVs, as well as their covariance matrices.

Primary vertex reconstruction is performed using VELO tracks. All VELO tracks contribute to PV reconstruction, including those outside of the full detector acceptance. Backwards tracks, for example, can be left by particles travelling away from the LHCb spectrometer and can be used to locate PVs. As a result, HLT1 should reconstruct every track in the VELO acceptance and the reconstruction efficiency should be maximized over the VELO's entire acceptance.

Track reconstruction uses information from the VELO, UT, and SciFi detectors. HLT1 must be able to reconstruct long tracks above a certain kinematic threshold. Requiring a minimum momentum for reconstructed tracks restricts the search area for tracks in the SciFi, speeding up the track reconstruction. LHCb successfully carried out its b and c physics programs in Run 2 requiring $p_{\text{T}} > 500$ MeV in HLT1. Loosening this threshold, however, would allow for new trigger strategies and would benefit other areas of LHCb's physics program. For example LHCb studies rare strange hadron decays to final states containing muons. Muons must have $p > 3$ GeV in order to reach the muon stations, making this a more natural tracking threshold.

HLT1 must select candidates that are interesting and useful for offline analysis. This means that the quality of track reconstruction in HLT1 cannot be significantly worse than the full reconstruction used in HLT2 and offline. In HLT2, the track momentum resolution is around 0.5%. This means that the HLT1 momentum resolution should not be much larger than about 1% for charged particles from heavy flavor decays. This also means that the HLT1 covariance matrix quality of both tracks and primary vertices should be comparable to those in HLT2 to ensure similar χ_{IP}^2 calculations.

4.2 The Allen Project

4.2.1 The Basics of GPUs

The challenge of running LHCb with a triggerless readout prompted research and development into the use of heterogeneous computing architectures in HLT1, including GPUs. GPUs are designed to display graphics, which requires calculating the color

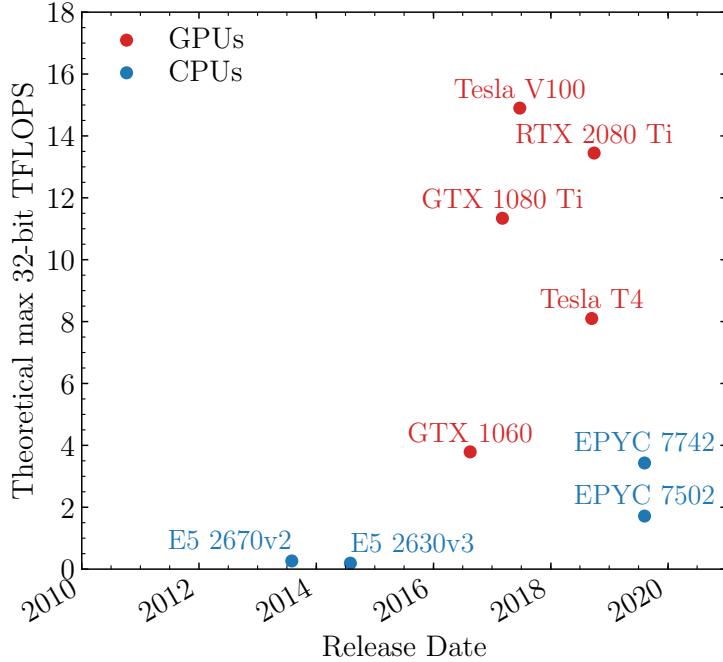


Figure 4-7: Comparison of theoretical maximum performance of a selection of GPUs and CPUs.

of millions of pixels at once. Modern GPUs use thousands of cores to do this, allowing them to perform trillions of single-precision floating point operations per second (FLOPS). Figure 4-7 demonstrates this, showing how the parallelism of GPUs allow them to achieve greater theoretical maximum performance than high-end CPUs. This includes consumer GPUs such as the NVIDIA RTX 2080 Ti.

Over the last two decades, GPUs have seen use in more general computing applications, including in high energy particle physics. The ALICE experiment uses GPUs in its trigger to reconstruct tracks left by charged particles in its time projection chamber [138], and the CMS collaboration is pursuing research and development to perform track reconstruction in their software trigger using GPUs [139]. In addition to LHC experiments, the Mu3e [140] and CBM [141] experiments have proposed using GPUs to perform both reconstruction and selection of single physics signatures in their triggers.

Figure 4-8 shows a diagram of the structure of GPU threads and memory. The thousands of GPU cores are spread across multiple streaming multiprocessors (SMs). Within a single SM, threads are physically grouped together into 32-thread *warps*. Warps are further combined into *blocks*. Blocks are combined into a *grid*, which may span multiple SMs. While the warps are physical groups of threads, blocks and grids are logical constructs and have programmable dimensions. GPU instructions are executed as a *kernel*, which is an algorithm running on a grid.

GPUs operate based on a same-instruction-multiple-thread (SIMT) computing

model [142]. This means the same instructions are executed in lock step across multiple threads on different data. Threads within a warp must execute instructions in lock step, making warps the basic processing unit of the SIMT model. Threads within a block also operate in lock step. This means that ideal GPU algorithms should minimize branching within a block. An `if-else` statement, for example, will result in both sets of instructions being executed on every thread, leading to wasted computations.

GPU memory is divided into several different categories based on its scope and storage location [143].

- *Global memory* is the largest and slowest memory resource on the GPU. Global memory is typically a few gigabytes and is stored separately from the SM. It is visible to all threads running an application.
- *Constant memory* is similar to global memory, but is read-only during kernel execution. Constant memory is slightly faster than global memory.
- *Shared memory* is stored on a SM and is visible to each thread in a block. Each SM has a few dozen kilobytes of shared memory. Shared memory access is much faster than global memory access and is the primary way threads within a block can share information.
- *Register memory* is also stored on the SM and is accessible only to the thread that wrote it. This is the fastest GPU memory.

In total, GPUs have only a few gigabytes of memory, while modern CPUs may feature over 100 GB. Furthermore, GPU memory access is much more restrictive than CPU memory access. Global memory, the vast majority of GPU memory, must be allocated and freed for all threads at once. In practice this means that most GPU memory cannot be allocated during kernel execution.

A GPU accesses global memory in chunks of 32, 64, or 128 bytes [143]. If threads within a warp access data that is stored consecutively, the memory can be coalesced into a single access. If the data is not stored consecutively in memory, many global memory accesses must be performed sequentially, leading to decreases in performance. As a result, it is often advantageous to store data in a structure-of-arrays (SoA) format rather than an arrays-of-structures (AoS) format.

This thread and memory layout reveals the challenges of developing applications for GPUs. GPUs offer immense computing power, but taking advantage of this power requires designing algorithms for a SIMT computing model that optimally utilize the GPUs limited memory. An ideal GPU application uses a small amount of memory, minimizes memory accesses, and minimizes branching such that individual threads always perform useful operations.

4.2.2 The Allen Framework

LHCb's GPU research and development has resulted in the Allen project. Allen is named after Frances E. Allen and is a framework for executing complex GPU

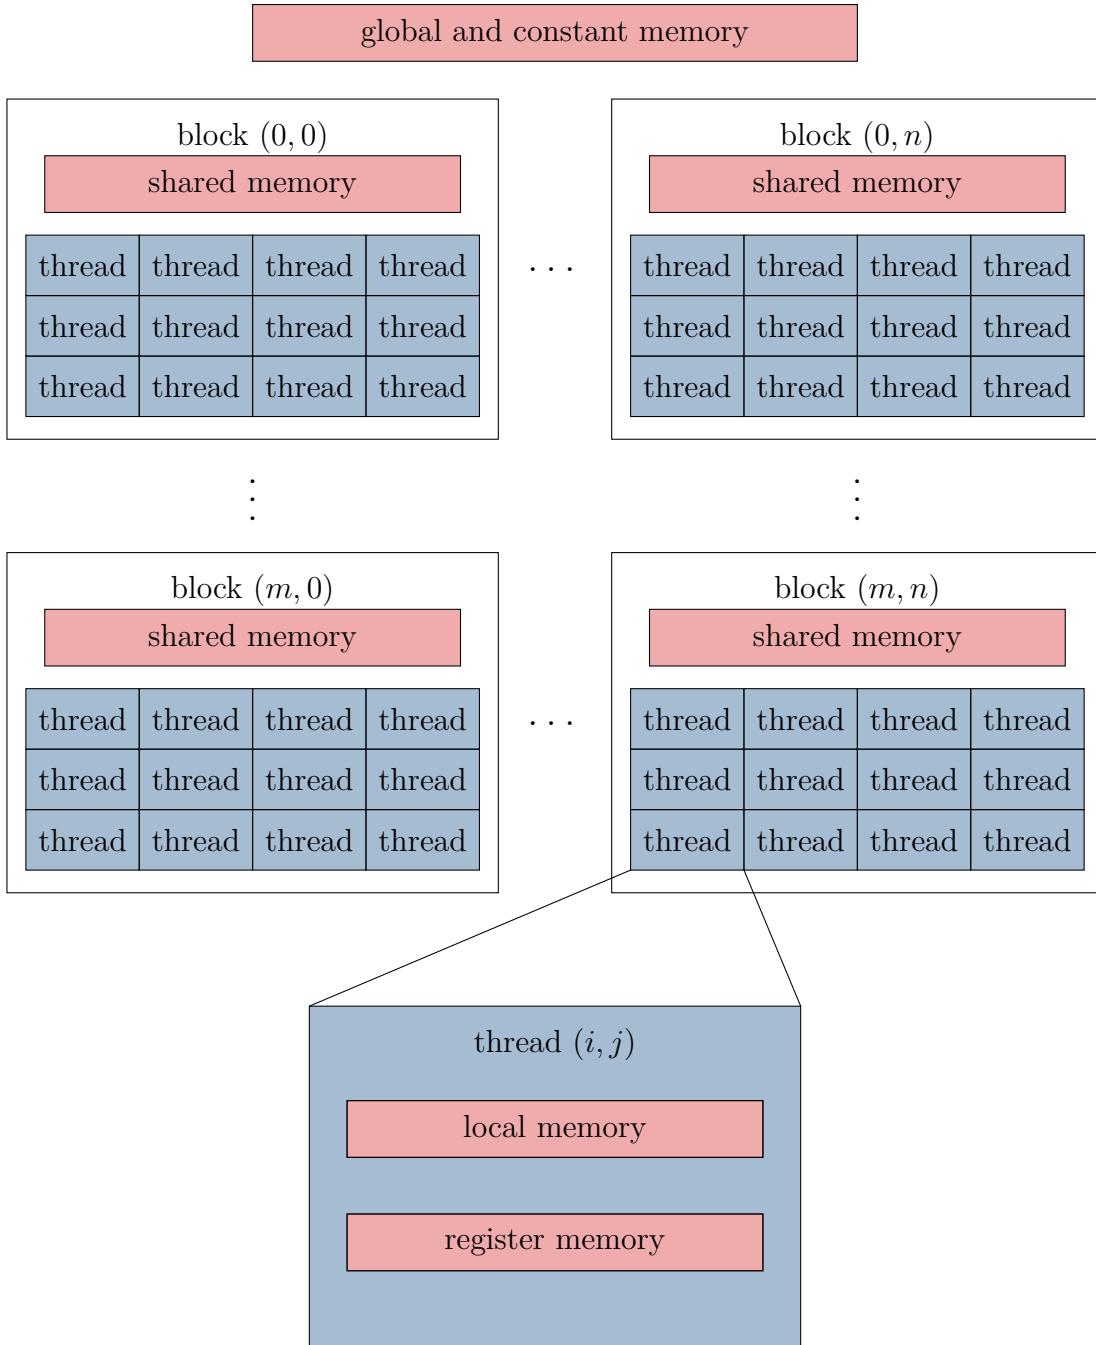


Figure 4-8: Illustration of the memory and thread layout of a GPU grid.

workflows. Allen also includes a collection of algorithms implementing LHCb’s HLT1. Allen is written primarily in C++ and CUDA [144]. While Allen can be executed as part of the LHCb software stack, it can also be used as a standalone application. Allen requires only CUDA v10.2 and a C++17 compatible compiler. It can be compiled and run on both GPUs and CPUs. It can also be compiled using HIP, allowing it to run on AMD GPUs as well.

Allen uses a configurable number of CPU threads, each steering a single sequence of GPU kernels called a *stream*. Each GPU stream processes a configurable number of events in parallel, executing the HLT1 reconstruction and selection sequence. Events are copied to and from the GPU in groups, effectively hiding the latency of CPU-GPU communication. The memory needed for each stream is allocated before execution. A memory manager handles allocating and freeing specific buffers between kernel executions within this preallocated memory. A scheduler checks the input and output of each kernel to determine when to allocate and free memory.

This framework effectively hides many of the difficult aspects of working with GPUs, such as allocating and freeing global memory during stream execution. This allows new developers to quickly begin writing algorithms. Because Allen can be compiled for CPUs, a GPU is not required to begin contributing. This accessibility is vital, as most LHCb collaboration members have little or no GPU programming experience.

4.2.3 Integration

In order for Allen to be a viable trigger solution, GPUs must fit into the LHCb data acquisition sequence. The sequence begins with the LHCb detector, which produces a peak data rate of 40 Tbit/s. The data is assembled into events by 170 event builder nodes. In the CPU-only trigger scenario, 40 Tbit/s of data is then transferred to the Event Filter Farm, which hosts CPUs that execute HLT1 and HLT2. Accepted events are then written to storage.

Figure 4-3 shows the planned Run 3 LHCb data acquisition sequence. The Allen project proposes to add GPUs to the event builder nodes. Each of these nodes has three PCI express ports, allowing it to host three GPUs. Performing this event filtering in GPUs will decrease the peak data rate between the EB nodes and the EFF from 40 Tbit/s to about 1 Tbit/s, allowing for significant cost savings on networking cables. As a result, adding GPUs to the EB nodes will be cheaper than the CPU-only trigger running in the EFF. Allen has undergone extensive integration tests to insure its long-term stability. These are documented in Ref. [128].

4.3 Trigger Algorithms For GPUs

4.3.1 The Allen Trigger Sequence

Figure 4-9 shows the Allen HLT1 sequence of algorithms. Large portions of the HLT1 reconstruction sequence are inherently parallel. Each event, for example, can

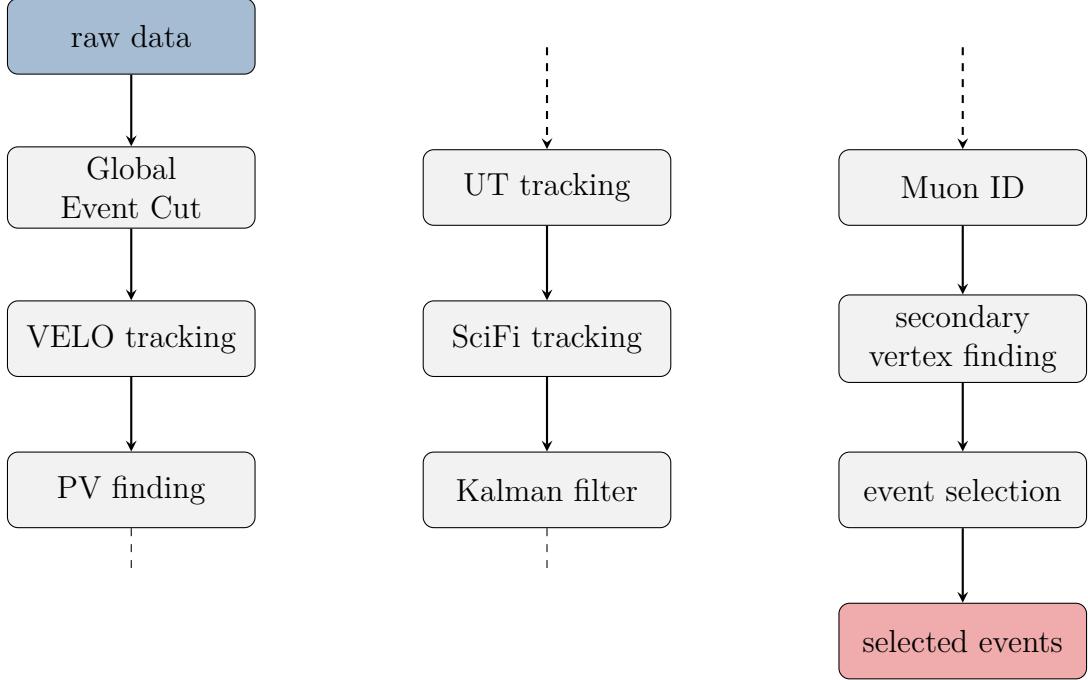


Figure 4-9: The Allen HLT1 sequence of algorithms. Input is raw data packaged into events by the event builder nodes and copied to the GPU. This is shown in blue. For selected events raw banks summarizing the passing decisions and the objects that triggered them are copied back to the CPU. This is shown in red.

be reconstructed independently of all other events. Furthermore individual track candidates can be reconstructed, combined into SVs and evaluated by trigger selections in parallel. This makes GPUs an attractive technology for HLT1, but it also requires rewriting or reoptimizing most of the HLT1 sequence to take advantage of the GPU’s massive parallelism.

4.3.2 Global Event Cut

High occupancy events use a disproportionately large amount of LHCb’s reconstruction time. Furthermore these events suffer from poor track reconstruction performance. As a result, a Global Event Cut (GEC) removes the highest occupancy events. The GEC selects events based on the size of the raw data produced by the UT and SciFi and is tuned to reject 7% of events. The GEC results in a 25% increase in throughput. Studies of the feasibility of removing the GEC are currently underway.

4.3.3 Track and Primary Vertex Reconstruction

Track reconstruction in HLT1 begins with the VELO. This is described in detail in Ref. [145]. VELO reconstruction begins with clustering detector data into *hits*. Figure 4-10 demonstrates the VELO clustering algorithm. VELO data is decoded in grids of 2x4 pixels called super pixels. VELO clusters are typically small, containing only

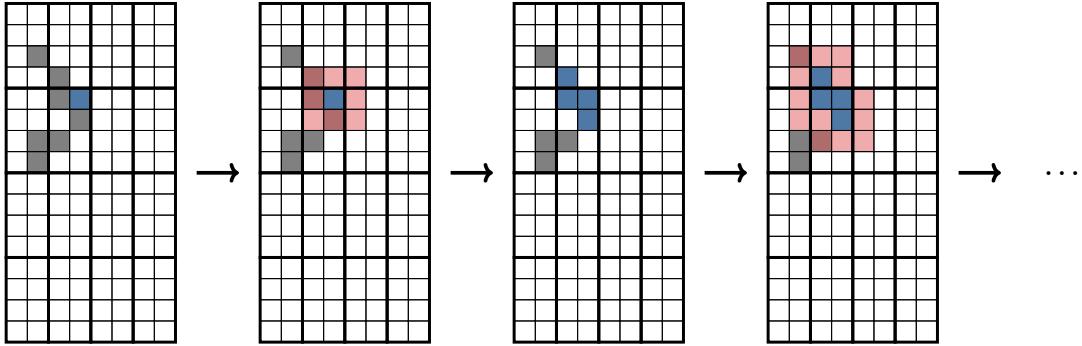


Figure 4-10: Illustration of the mask clustering algorithm. Pixel borders are shown as thin lines and super pixel borders are shown in bold. Activated pixels that are not included in the cluster are shown in gray. Pixels included in the cluster at each step are shown in blue. The masks are shown in red. Adapted from Ref. [145].

a few pixels each. In addition, a very small fraction of VELO pixels are activated in a single event, so VELO clusters rarely overlap. Seed pixels are chosen by requiring the absence of adjacent activated pixels in the upper right hemisphere of surrounding pixels. This seeding is highly parallelizable and provides an upper bound on the number of clusters in each module. Masks are used to extend clusters to pixels adjacent to the seed pixel. The masking is repeated until no more pixels can be added to the cluster. The vast majority of VELO clusters can be fully constructed in fewer than ten iterations of this algorithm. Each cluster is processed in parallel, allowing the clustering to occur in constant time.

VELO tracks are reconstructed using the *Search by Triplet* algorithm. Because the magnetic field in the VELO region is negligible, VELO tracks consist of hits lying along straight lines. Hits are sorted in the azimuthal angle ϕ , taking advantage of the fact that tracks originating from the beamline will have constant ϕ . Track seeds are created using three consecutive VELO modules. Threads are assigned to clusters in the middle module. All possible triplet combinations in a ϕ search window around the middle cluster are fit to a straight line and the triplet with the smallest χ^2 below a certain maximum χ^2 is kept as a track seed. This begins with the modules furthest from the collision point and proceeds inwards. Track candidates are extrapolated to the next module using the slope defined by their previous two hits. The closest cluster to the extrapolated track within a certain minimum distance is appended to the track. All hits used in tracks are flagged, and the triplet seeding procedure is repeated at the next module with the remaining unflagged hits. This alternating process of seeding and extrapolating tracks to the next module is repeated for every module. Figure 4-11 demonstrates the *Search by Triplet* algorithm.

VELO tracks are then extrapolated to the UT. The *CompassUT* [146] algorithm uses the extrapolated position as a seed to search for candidate UT tracks with three or four hits. The UT is divided into sectors and search windows are opened on a range of sectors surrounding the extrapolated VELO track. The search can be configured to consider three or five sectors. Figure 4-12 shows the three-sector configuration. The windows are limited in the x direction by the sector widths and in the y direction by

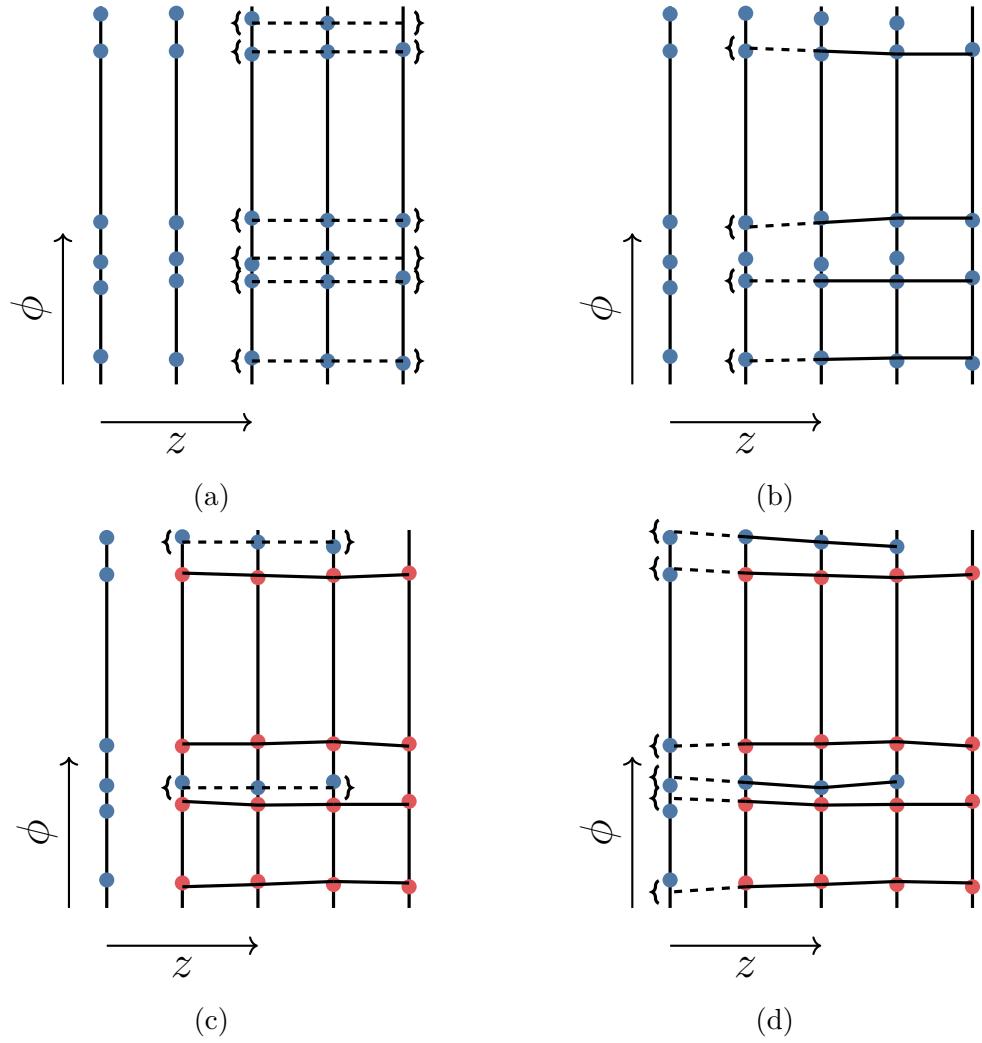


Figure 4-11: Illustration of the *Search by Triplet* algorithm. Triplets are formed (a) and extrapolated to the next layer, where additional hits are added to the track (b). Hits in tracks with at least four hits are flagged (shown in red) and the triplet formation is repeated (c). Tracks are then forwarded to the next layer (d) and the steps are repeated until the entire detector has been processed. Adapted from Ref. [145].

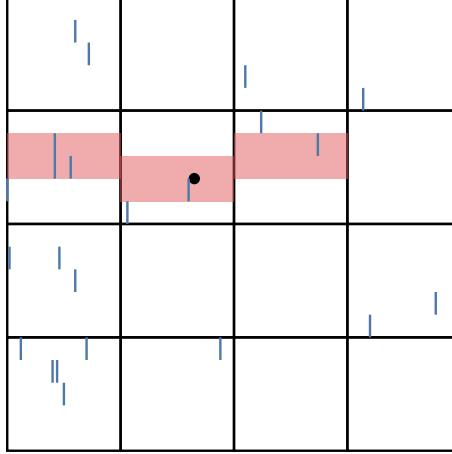


Figure 4-12: Demonstration of the UT search window determination. The black grid demarcates UT sectors. Activated strips are shown in blue and the extrapolated VELO track position is shown as a black circle. Search windows are shown in red. Adapted from Ref. [146].

the VELO track extrapolation uncertainty. The search is configured to only consider tracks with $p > 3$ GeV. The algorithm begins by searching for doublets in the first and third layer, then the second a fourth. The slopes of the doublet and the velo track are used to search for hits in the other layers. The small magnetic field in the UT region allows for the use of a single-kick approximation to determine the charge of the particle and obtain a first momentum estimate.

A parameterized extrapolation is used to forward tracks from the UT through the magnetic field to the SciFi. Track seeds are formed from triplets of hits in the x layers of each of the three SciFi tracking stations. The magnetic field within the SciFi is parameterized as

$$B_y(x, y, z) = B_0(x, y) + B_1(x, y)z \quad (4.1)$$

where B_0/B_1 is constant. This parameterization is used to extrapolate the track and add hits from the remaining x , u , and v layers. Using the seed triplet, the positions of the remaining hits can be predicted to within a few millimeters, allowing the search windows for these hits to be extremely small. This in turn allows Allen to relax the forward tracking kinematic threshold. Instead of requiring $p_T > 500$ MeV as in Run 2, Allen is able to reconstruct tracks with $p > 3$ GeV and no p_T requirement.

Reconstructed long tracks are then extrapolated to the Muon system. Tracks are matched to hits in the Muon system using the *IsMuon* algorithm. This algorithm has been used by LHCb since Run 1 and is documented in Ref. [147].

Reconstructed VELO tracks can be used to reconstruct PVs. VELO tracks are extrapolated to their point closest to the beamline and a histogram is filled with the resulting z positions. A fast velo-only Kalman filter is used to determine the covariance matrix of each track, and the histogram is filled using an approximate Gaussian kernel based on the z position uncertainty. This extrapolation and histogramming is performed in parallel for each track. Peaks in this distribution are used as PV seeds,

and a χ^2 -fit is used to determine the position of the PV in three dimensions.

4.3.4 Kalman Filter and Secondary Vertex Fitter

A Kalman filter is used in HLT1 to provide an accurate covariance matrix for each track extrapolated to its point closest to the beamline. Track states in LHCb are described by five-dimensional vectors $(x, y, t_x, t_y, q/p)$, where x and y are state positions, t_x and t_y are track slopes, and q/p is the track charge over momentum. Starting with a seed state, the Kalman filter alternates predicting a subsequent state and updating the state using detector information.

During the prediction step the state vector \vec{x} and covariance matrix C are propagated from step $k-1$ to step k :

$$\vec{x}_k^{\text{pred}} = f_{k-1}(\vec{x}_{k-1}), \quad (4.2)$$

$$C_k^{\text{pred}} = F_{k-1} C_{k-1} F_{k-1}^T + Q_{k-1}, \quad (4.3)$$

where f_{k-1} is the propagation function from step $k-1$ to step k , and F_{k-1} is the Jacobian matrix of this function. In LHCb, these steps correspond to layers of the tracking detectors. The matrix Q_{k-1} is the noise contribution to the covariance matrix. In the context of tracking at LHCb this noise is due to multiple scattering in the detector material.

After the prediction step, a measurement is used to update the state and covariance matrix. The measurement at step k provides a measurement vector \vec{m}_k and measurement covariance V_k . The measurement \vec{m}_k is related to the true state at step k by

$$\vec{m}_k = H_k \vec{x}_k^{\text{true}} + \vec{\epsilon}_k, \quad (4.4)$$

where H_k is referred to as the measurement matrix and $\vec{\epsilon}_k$ is the measurement error, which is assumed to be drawn from a Gaussian with a mean of zero and covariance V_k . Then the measurement residual \vec{r}_k is given by

$$\vec{r}_k = \vec{m}_k - H_k \vec{x}_k^{\text{pred}}, \quad (4.5)$$

and the residual covariance C_k^{res} is given by

$$C_k^{\text{res}} = V_k + H_k C_k^{\text{pred}} H_k^T. \quad (4.6)$$

This can then be used to calculate the Kalman gain K_k :

$$K_k = C_k^{\text{pred}} H_k^T C_k^{\text{res}}. \quad (4.7)$$

Finally, the post-fit state and covariance are given by

$$\vec{x}_k = \vec{x}_k^{\text{pred}} + K_k \vec{r}_k, \quad (4.8)$$

$$C_k = (I - K_k H_k) C_k^{\text{pred}}. \quad (4.9)$$

Because each step of the Kalman filter requires information from previous steps, Kalman filters are inherently sequential. As a result, the Allen Kalman filter is optimized for GPUs by minimizing memory usage and accesses, allowing it to quickly fit every track in an event in parallel.

The Allen Kalman filter is based on a parameterized Kalman filter developed for the CPU. The nominal LHCb Kalman filter uses a magnetic field map and Runge-Kutta extrapolator to predict track states and a detailed material map to determine noise from multiple scattering. The parameterized Kalman filter replaces these calculations with simple parameterizations. These parameterizations are described in Appendix A.

Kalman filters in LHCb historically follow the same basic fitting procedure:

1. Perform a VELO-to-SciFi (forward) fit.
2. Perform a SciFi-to-VELO (backward) fit.
3. Combine the states from the forward and backward fits using a smoothing procedure.
4. Use the smoothed states to remove outliers from the fit.
5. Repeat steps 1-4 until no outliers remain or a maximum number of outliers is removed.

Smoothing and outlier removal requires storing up to 42 (one for each detector layer) state vectors and covariance matrices for each track, dramatically increasing the required number of global memory accesses. Experience with Kalman filtering in LHCb has demonstrated that smoothing and outlier removal only slightly increase fit quality, so these steps are omitted in the GPU implementation of the parameterized Kalman filter.

Without smoothing, the forward and backward fits are independent, so Allen could in theory only perform a backward fit. However the UT-SciFi extrapolation is linearized around a trajectory originating from the origin. As a result, a forward fit is necessary for a precise momentum estimate. In order to determine the track state and covariance matrix closest to the beamline, the product of the track prediction Jacobian matrices between the end of the VELO and the end of the SciFi is calculated,

$$F = \prod_{k=m}^n F_k, \quad (4.10)$$

where m is number of hits in the VELO, and n is the total number of hits associated to the track. This is then used to propagate the final state of the forward track fit back to the end of the VELO

$$\vec{x}_0^{\text{back}} = F^{-1} \vec{x}_n^{\text{fwd}}, \quad (4.11)$$

$$C_0^{\text{back}} = F^{-1} C_n^{\text{fwd}} (F^{-1})^T, \quad (4.12)$$

where fwd and back denote states and covariance matrices used for the forward and backward fits, respectively. The resulting state and covariance matrix can be used to perform a VELO-only backward fit in order to acquire the state and covariance matrix closest to the beamline. This *modified* kalman filter provides a fit quality, track state, and covariance matrix similar to the full forward and backward fit and results in a 50% speedup.

As the Allen sequence has become faster, even this *modified* Kalman filter serves as a performance bottleneck, occupying around 30% of the sequence execution time. A *simplified* Kalman filter has been developed to provide an additional speedup. The IP and χ^2_{IP} are determined almost exclusively using information from the VELO. In addition, Allen's forward tracking algorithm already provides momentum resolution of around 1%. This means a single backward VELO fit will provide enough information to adequately describe HLT1 tracks. This VELO-only fit is based on the Kalman filter used for Allen's PV reconstruction. This separates the Kalman fit into two independent fits in (x, t_x) and (y, t_y) . This fit is improved by using the momentum estimate from the forward tracking and the noise parameterization from the full parameterized Kalman filter to more accurately account for multiple scattering. This fit is an order of magnitude faster than the *modified* GPU Kalman filter. This *simplified* fit is the baseline Kalman filter used in both the Allen and CPU implementations of HLT1. The *simplified* Kalman filter occupies less than 1% of the total Allen execution time.

Kalman filtered tracks are associated to the primary vertex which minimizes χ^2_{IP} and used to form two-track secondary vertices. In order to limit the number of secondary vertex candidates, only tracks with $\chi^2_{\text{IP}} > 9$ and associated to the same PV are combined. Tracks with $\chi^2_{\text{IP}} > 16$ are considered to be unassociated. Unassociated tracks may be combined with any other track to form a secondary vertex. In the case of dimuons, it is important that the lifetime of selected candidates is unbiased. As a result, these χ^2_{IP} and primary vertex requirements are ignored when both tracks are identified as muons. The secondary vertex fitter uses the track states determined by the Kalman filter to calculate the point of closest approach between each pair of tracks. This is used as the vertex seed position. The vertex fit then updates this seed position to minimize the vertex χ^2 determined by the vertex position, track positions, and track covariance matrices. The fit provides an estimate of the secondary vertex position and covariance matrix.

4.3.5 Trigger Lines, Selections, and Output

The HLT1 sequence is defined by a collection of *trigger lines*, which are sequences of reconstruction and selection algorithms that result in trigger decisions. In practice, HLT1 lines share large portions of their reconstruction. So while HLT1 consists of dozens of different lines, this corresponds to only a few distinct reconstruction sequences. This includes a main reconstruction sequence that runs on all events passing the GEC and is used for the majority of LHCb's physics program. In the past, a VELO-Muon reconstruction sequence has matched VELO tracks to hits in the Muon system. The VELO-Muon sequence was used in addition to the main reconstruction sequence in order to reconstruct muons that do not pass the forward

tracking threshold. Allen’s loose tracking threshold reduces HLT1’s reliance on these special reconstructions. For example, Allen’s forward tracking algorithm will already reconstruct all tracks for which muon identification is possible, removing the need for a special low-momentum muon sequence. As a result, all Allen trigger lines discussed here currently use the same reconstruction sequence. However, if the GEC cannot be removed entirely, an additional high- p_T muon reconstruction sequence will still be necessary to bypass the GEC for electroweak physics measurements, which suffer from large systematic uncertainties from the GEC.

Trigger lines used for heavy flavor physics analyses select events based on either single tracks or two-track secondary vertices. Allen consists of six prototype trigger selections that cover most of LHCb’s physics program. These selections correspond to trigger lines that used about 95% of LHCb’s HLT1 trigger rate in Run 2 and are tuned to produce an output rate of about 1 MHz. Table shows the criteria for one-track selections and Table shows the criteria for two-track selections. One-track selections are based on the track’s Kalman fit quality, p_T , χ^2_{IP} , and muon identification. Selections on secondary vertices are based on the features of the constituent tracks as well as features of the vertex itself. This includes the vertex fit quality (χ^2_{SV}), η of the PV–SV vector ($\eta(\text{PV} - \text{SV})$), the distance of closest approach of the two constituent tracks (DOCA), and the corrected mass (M_{corr}). The corrected mass is the mass of the SV calculated using the constituent track four-momenta and corrected for missing momenta using the angle between the SV momentum and the flight direction given by the vector connecting the PV and SV. This is given by

$$M_{\text{corr}} = \sqrt{m^2 + (p \sin \theta)^2} + |p \sin \theta|, \quad (4.13)$$

where m and p are the measured SV mass and momentum, respectively, and θ is the angle between the SV momentum and the PV-SV vector. The `OneTrackMVA` and `TwoTrackMVA` lines are arguably the most important. These two lines select the majority of hadronic decays studied at LHCb and use the vast majority of the HLT1 output trigger rate [148]. In addition to one- and two-track lines, studies are underway to implement three- and four-body lines. The addition of three- and four-body vertices would allow for selections that more specifically target the multi-body heavy flavor decays that make up most of LHCb’s physics program, providing increased trigger efficiencies at reduced rates.

The selection rate of each line can be decreased using a postscaler. This rejects passing events using a randomly generated number seeded using information in the raw event header. This means that the postscaler is deterministic and will always return the same decision for each line in each event. Allen evaluates trigger decisions for candidates in parallel, resulting in a selection stage that uses a negligible amount of the total sequence time. To test the scalability of the selection framework, 100 random one- and two-track selections were generated and added to the Allen sequence. This resulted in a slowdown of about 5%, indicating that Allen can handle the dozens of trigger lines used in Run 2.

After making trigger decisions, Allen produces decision reports (`DecReports`) with information about which trigger lines were accepted. Allen also produces selection

TrackMVA	$\chi^2/\text{NDF} < 2.5$ $p_T > 2 \text{ GeV}$ $\ln(\chi^2_{\text{IP}}) > \frac{1}{(p_T/\text{GeV}-1)^2} + \frac{1.248}{26(26-p_T/\text{GeV})} + \ln(7.4) \text{ if } p_T < 26 \text{ GeV}$ $\chi^2_{\text{IP}} > 7.4 \text{ if } p_T \geq 26 \text{ GeV}$
TrackMuonMVA	IsMuon $\chi^2/\text{NDF} < 100$ $p_T > 2 \text{ GeV}$ $\ln(\chi^2_{\text{IP}}) > \frac{1}{(p_T/\text{GeV}-1)^2} + \frac{1.248}{26(26-p_T/\text{GeV})} + \ln(7.4) \text{ if } p_T < 26 \text{ GeV}$ $\chi^2_{\text{IP}} > 7.4 \text{ if } p_T \geq 26 \text{ GeV}$
SingleHighPtMuon	IsMuon $\chi^2/\text{NDF} < 100$ $p_T > 6 \text{ GeV}$

Table 4.1: Requirements for one-track trigger selections implemented in Allen.

TwoTrackMVA	$\chi^2_{\text{SV}} < 25$ minimum track $p_T > 700 \text{ MeV}$ minimum track $\chi^2_{\text{IP}} > 12$ $\sum p_T > 2 \text{ GeV}$ $M_{\text{corr}} > 1 \text{ GeV}$ $2 < \eta(\text{PV} - \text{SV}) < 5$ $N(\text{tracks with } \chi^2_{\text{IP}} < 16) < 1$
DiMuonHighMass	both tracks IsMuon $\chi^2_{\text{SV}} < 25$ minimum track $p_T > 300 \text{ MeV}$ minimum track $p > 6 \text{ GeV}$ $M > 2.7 \text{ GeV}$ $\text{DOCA} < 0.2 \text{ mm}$
DiMuonLowMass	both tracks IsMuon $\chi^2_{\text{SV}} < 25$ minimum track $p_T > 500 \text{ MeV}$ minimum track $p > 3 \text{ GeV}$ minimum track $\chi^2_{\text{IP}} > 4$ $\text{DOCA} < 0.2 \text{ mm}$

Table 4.2: Requirements for two-track secondary vertex trigger selections implemented in Allen.

reports (`SelReports`), which contain summaries of the trigger candidates that caused the trigger lines to fire. This includes identification codes for each detector channel involved in the trigger decision. These reports are the primary output produced by Allen and are appended to the raw event data for use in subsequent trigger and analysis steps.

4.4 Performance

While Allen includes its own performance checkers, this section uses results obtained from tools developed for checking the CPU HLT1 performance. This allows for fair comparison between the two technologies. All result shown were produced using Allen v0.8. Reconstruction performance is determined using simulated minimum bias pp events with $\sqrt{s} = 14$ TeV.

4.4.1 Track Reconstruction

Reconstructed tracks are matched to generated charged particles. Matching requires 70% of the hits used to reconstruct the track match the trajectory of the generated particle. This is required for each subdetector in which the track leaves at least two hits except the UT, which is ignored for matching. Reconstructed tracks that are not matched to generated particles are considered fake or “ghost” tracks.

The Allen VELO, UT, and Forward tracking efficiencies are shown in Figures 4-13, 4-14, and 4-15, respectively. The forward tracking efficiency is much poorer for electrons, so these are shown separately. The UT and Forward reconstruction efficiencies are cumulative and include losses from the previous reconstruction steps. Efficiencies are calculated with respect to reconstructible generated particles. Particles must leave at least three hits in the VELO to be reconstructible as VELO tracks and at least three hits in both the VELO and SciFi to be reconstructible as long tracks.

The VELO efficiency plateaus at around 98%, which allows for efficient primary vertex reconstruction. The primary vertex reconstruction efficiency is shown in Figure 4-16. The position resolutions are shown in Figure 4-17, and pulls are shown in Figure 4-18. This shows that the PV position resolution is similar to that quoted in Ref. [134] and that the covariance matrix is well-calibrated.

The forward tracking efficiency plateaus at around 90% for non-electrons. Most of this efficiency loss occurs at the VELO-UT step due to the UT’s smaller acceptance. This is only slightly lower than the efficiency achieved in offline reconstruction in Run 1 [103]. Because a signal decay can produce a positive trigger decision when only a subset of its final state particles are reconstructed, the inefficiency in the plateau region contributes minimally to the overall trigger inefficiency. This is instead driven by tracks with momentum below the selected 3 GeV tracking threshold.

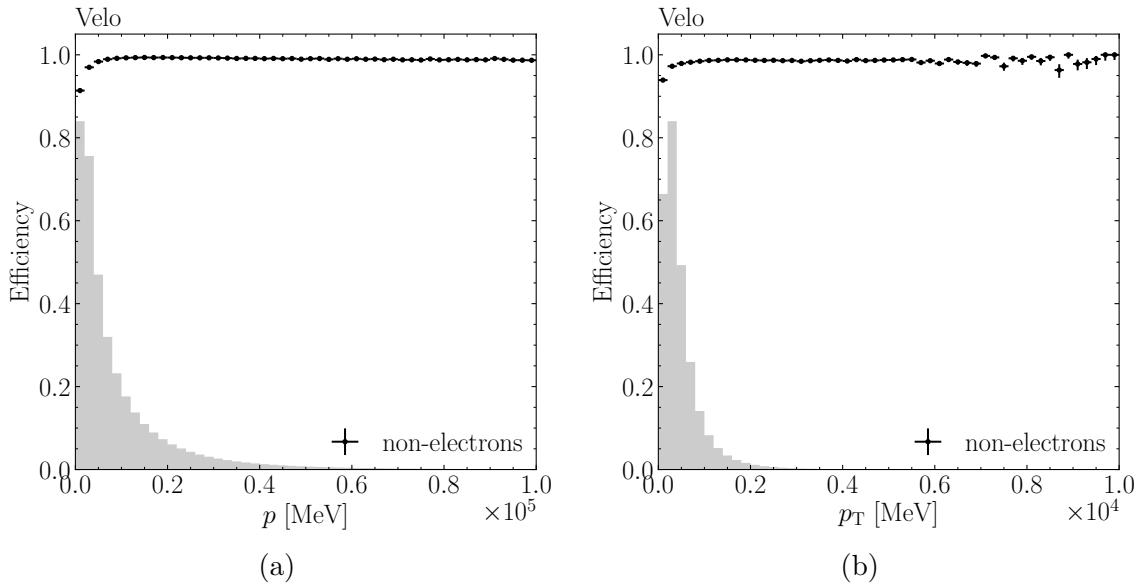


Figure 4-13: VELO track reconstruction efficiency as a function of (a) p and (b) p_T . The efficiency is calculated with respect to generated particles reconstructible as tracks in the VELO. The shaded regions show the generated distributions.

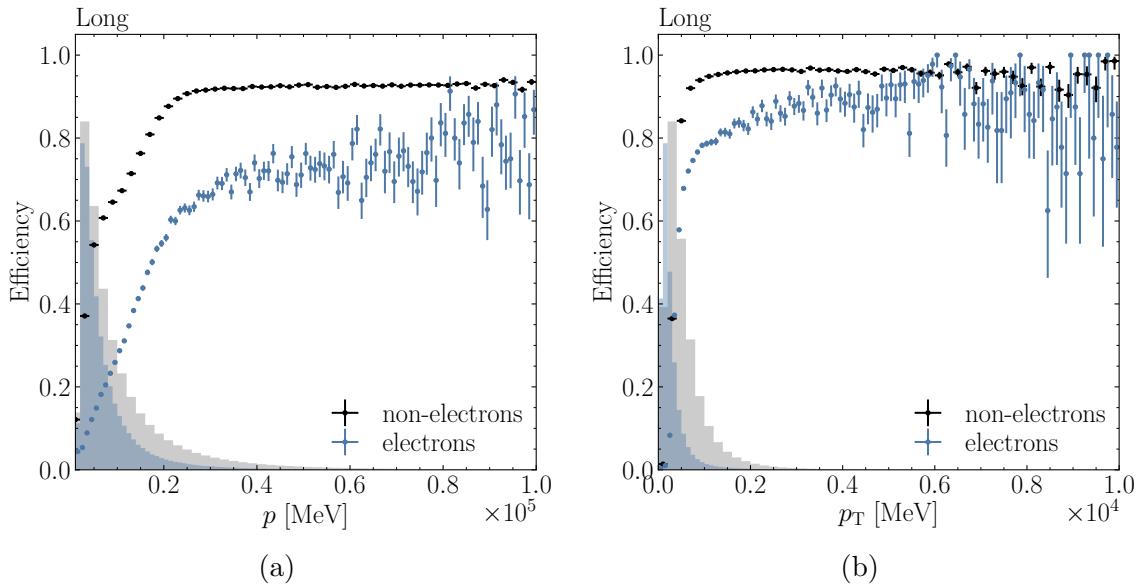


Figure 4-14: VELO-UT track reconstruction efficiency as a function of (a) p and (b) p_T . The efficiency is calculated with respect to generated particles reconstructible as long tracks. The shaded regions show the generated distributions.

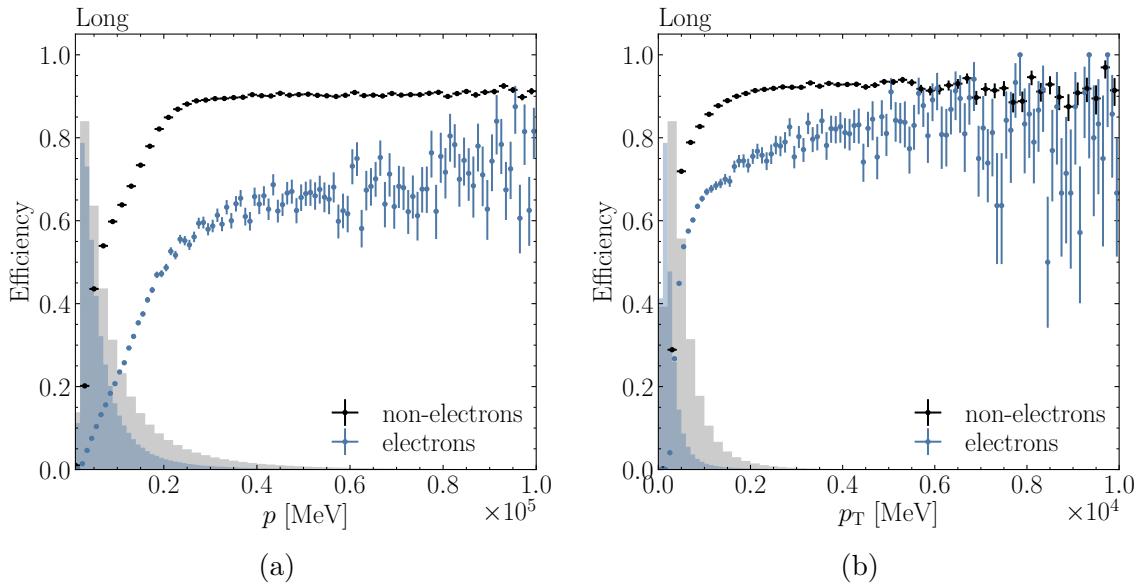


Figure 4-15: Forward track reconstruction efficiency as a function of (a) p and (b) p_T . The efficiency is calculated with respect to particles reconstructible as long tracks. The shaded regions show the generated distributions.

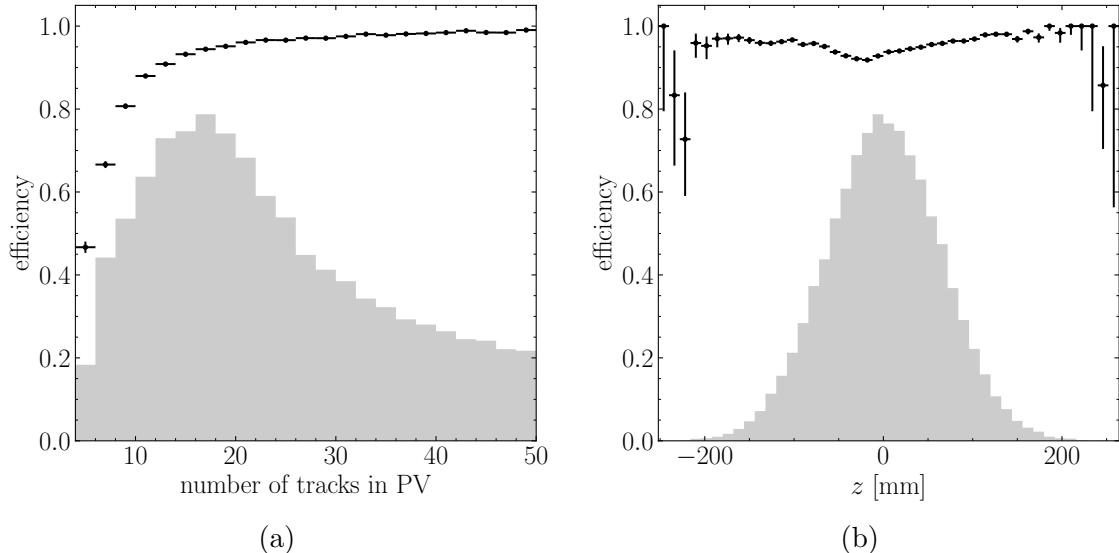


Figure 4-16: Primary vertex reconstruction as a function of (a) the number of tracks originating from the generated primary vertex and (b) the z position of the generated primary vertex. The shaded regions show the generated distributions.

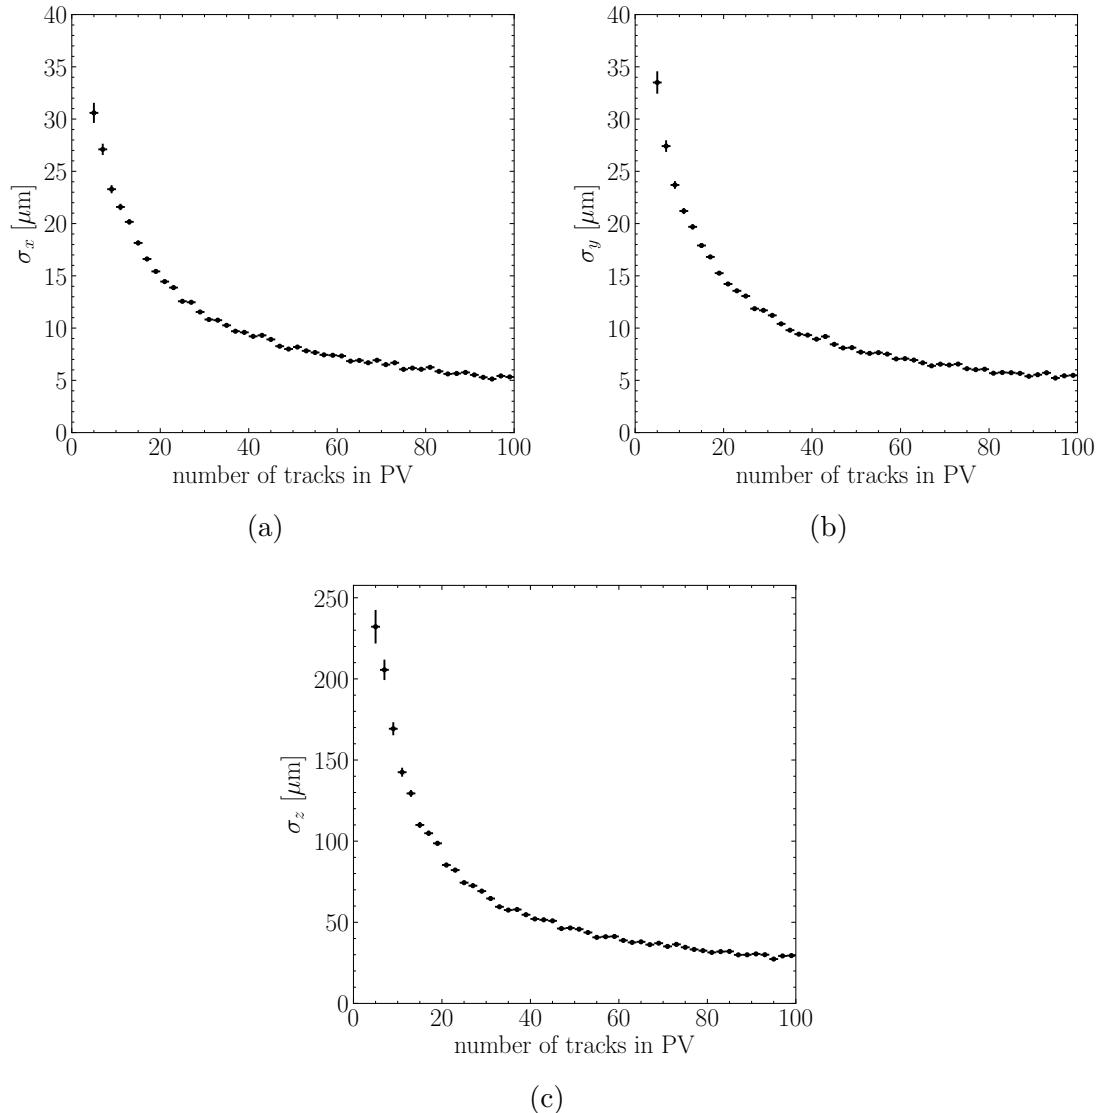


Figure 4-17: Position resolution of reconstructed primary vertices in (a) x , (b) y , and (c) z as a function of the number of tracks originating from the generated primary vertex. The shaded regions show the generated distributions.

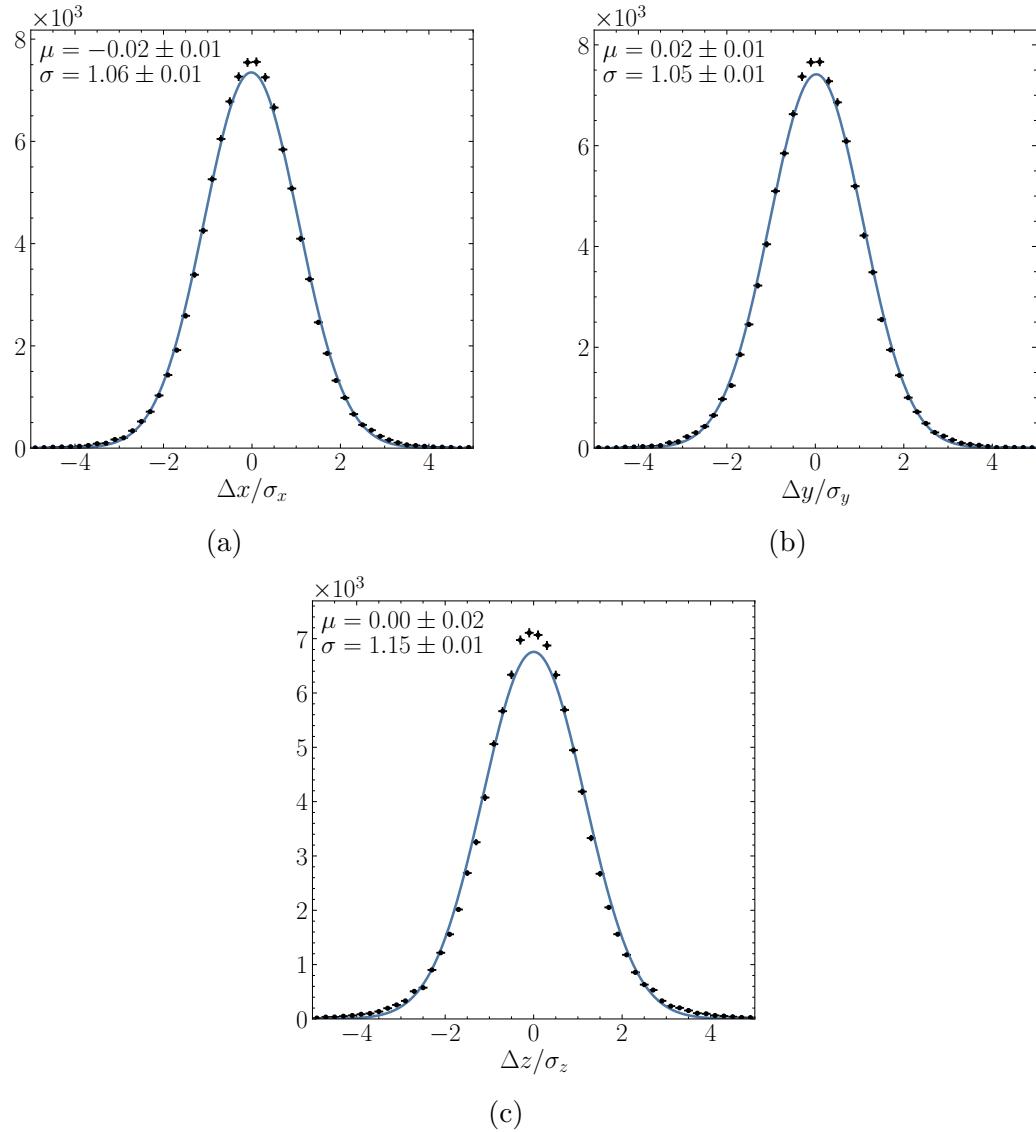


Figure 4-18: Pull distributions of reconstructed primary vertex positions.

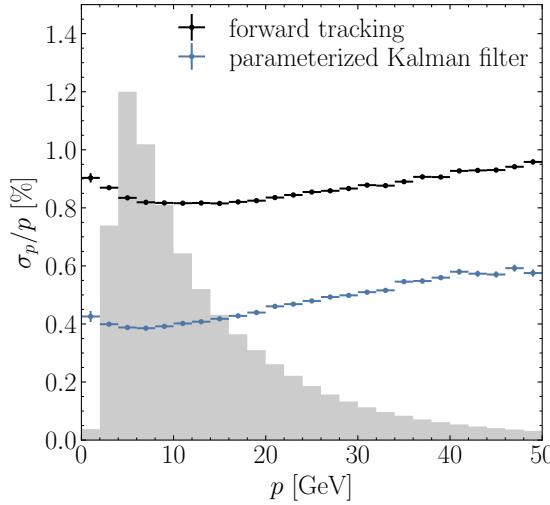


Figure 4-19: Resolution of the momentum estimates produced by the forward tracking algorithm and the *modified* parameterized Kalman filter. The shaded region shows the generated distribution.

4.4.2 Kalman Filter

The full *modified* parameterized Kalman filter can be used to obtain an improved momentum estimate for charged tracks. The momentum resolution is shown in Figure 4-19. This improved resolution is similar to results achieved in LHCb’s offline reconstruction [103].

Ghost tracks will tend to have poorer Kalman fit quality than real tracks, so the Kalman filter fit quality is used to eliminate ghost tracks. The cumulative χ^2 per degree of freedom distributions are shown for the *modified* and *simplified* parameterized Kalman filter configurations in Figure 4-20. Most ghost tracks are incorrect combinations of real VELO and SciFi tracks. As a result, the VELO-only *simplified* fit will provide less discrimination power than the full VELO-UT-SciFi *modified* fit. This effect is illustrated in Figure 4-20b. However, due to the momentum dependence of multiple scattering the fit quality of the VELO-only Kalman fit is correlated with momentum. Consequently, the *simplified* fit quality could still be a useful input to a multivariate ghost rejection algorithm.

The *simplified* Kalman filter provides an improved estimate of the track state and covariance matrix closest to the beamline. The improved estimate results in better impact parameter resolution and consequently better discrimination between tracks produced by prompt and displaced particles. The impact parameter resolution is shown in Figure 4-21, which demonstrates that the *modified* and *simplified* Kalman filters provide similar track descriptions near the beamline. This is in contrast to the VELO fit used in the primary vertex reconstruction. The fit used in PV reconstruction is tuned using the average momentum of all tracks in the VELO acceptance, resulting in worse performance at low and high momenta.

Track displacement in LHCb is most commonly expressed in terms of the impact

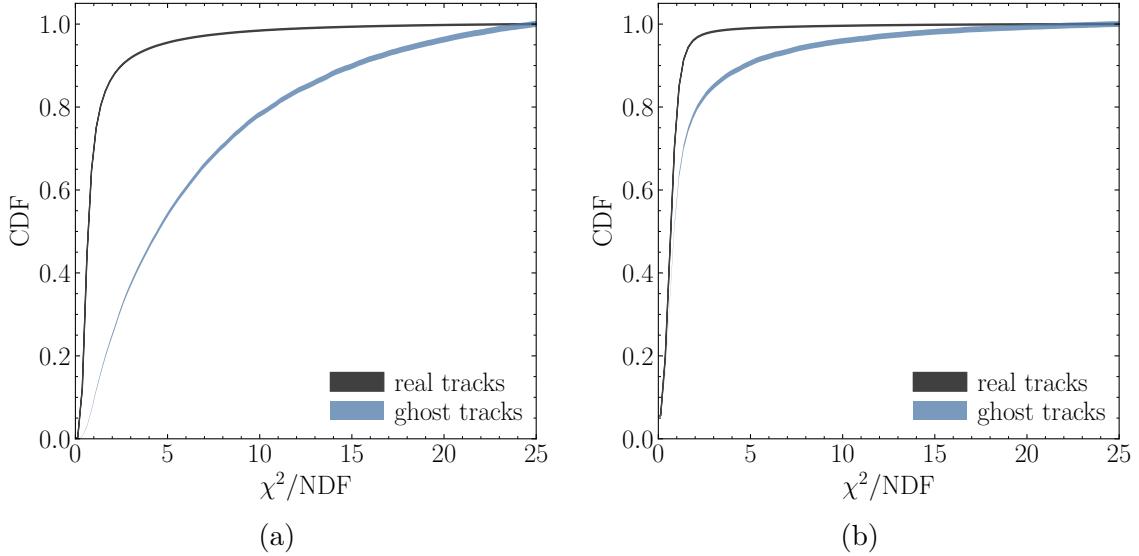


Figure 4-20: The χ^2 per degree of freedom cumulative distributions for real and ghost tracks. Results are shown for the (a) *modified* and (b) *simplified* parameterized Kalman filters. The widths of the lines shows the statistical uncertainty.

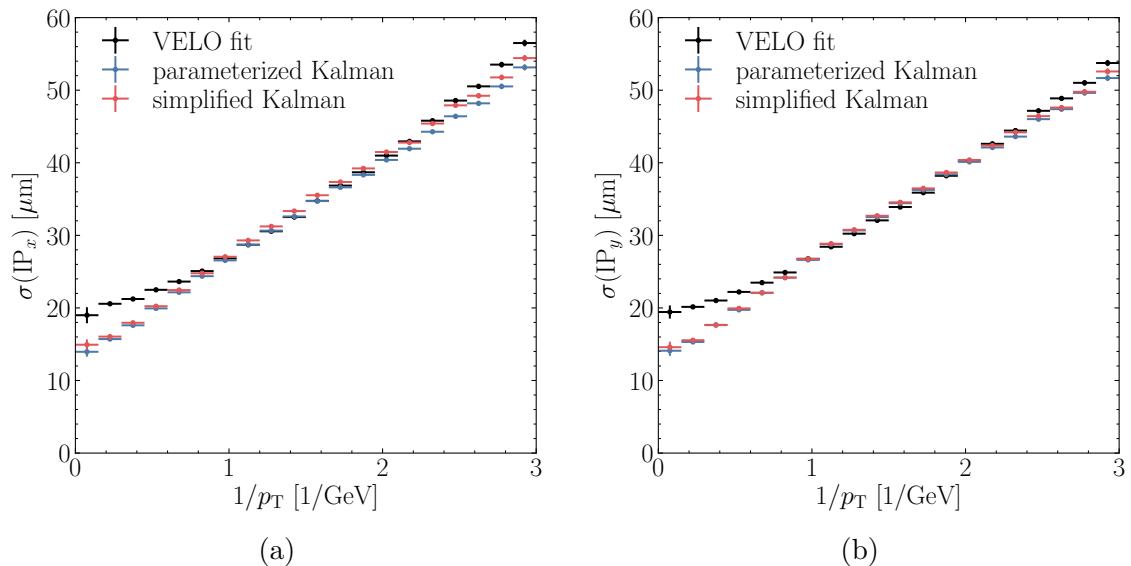


Figure 4-21: One dimensional impact parameter resolution in (a) x and (b) y determined using various fitting strategies.

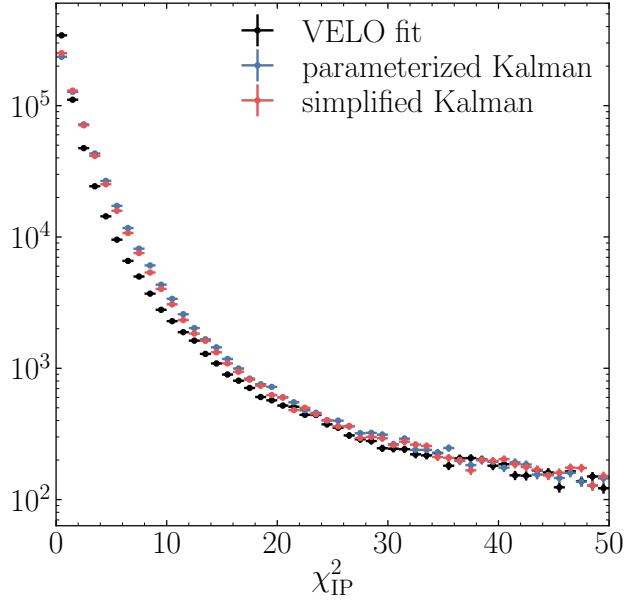


Figure 4-22: The χ^2_{IP} distribution for the various Allen track fits.

parameter significance, or χ^2_{IP} . A comparison of the χ^2_{IP} calculated using various Kalman filters is shown in Figure 4-22. This also shows similar results between the parameterized and simplified fits. Additionally, the pulls of the track state closest to the beamline determined using the *simplified* Kalman filter are shown in Figure 4-23, which indicates that the covariance matrix is well calibrated.

4.4.3 Selections

The trigger selection rate is determined using simulated minimum bias events. The event selection rate of individual selections is shown in Table 4.3. Allen selection efficiencies are determined using simulated signal samples of the decays

- $B_s^0 \rightarrow \phi\phi$
- $D_s^+ \rightarrow K^+ K^- \pi^+$
- $B^0 \rightarrow K^{*0} e^+ e^-$
- $B^0 \rightarrow K^{*0} \mu^+ \mu^-$
- $J/\psi \rightarrow \mu^+ \mu^-$
- $Z \rightarrow \mu^+ \mu^-$

Trigger efficiencies are shown in Figure 4-24. Trigger efficiencies are defined with respect to signal decays that are potentially reconstructible. The parent particles must have $p_T > 2 \text{ GeV}$. The short lived Z and J/ψ have no decay time requirement,

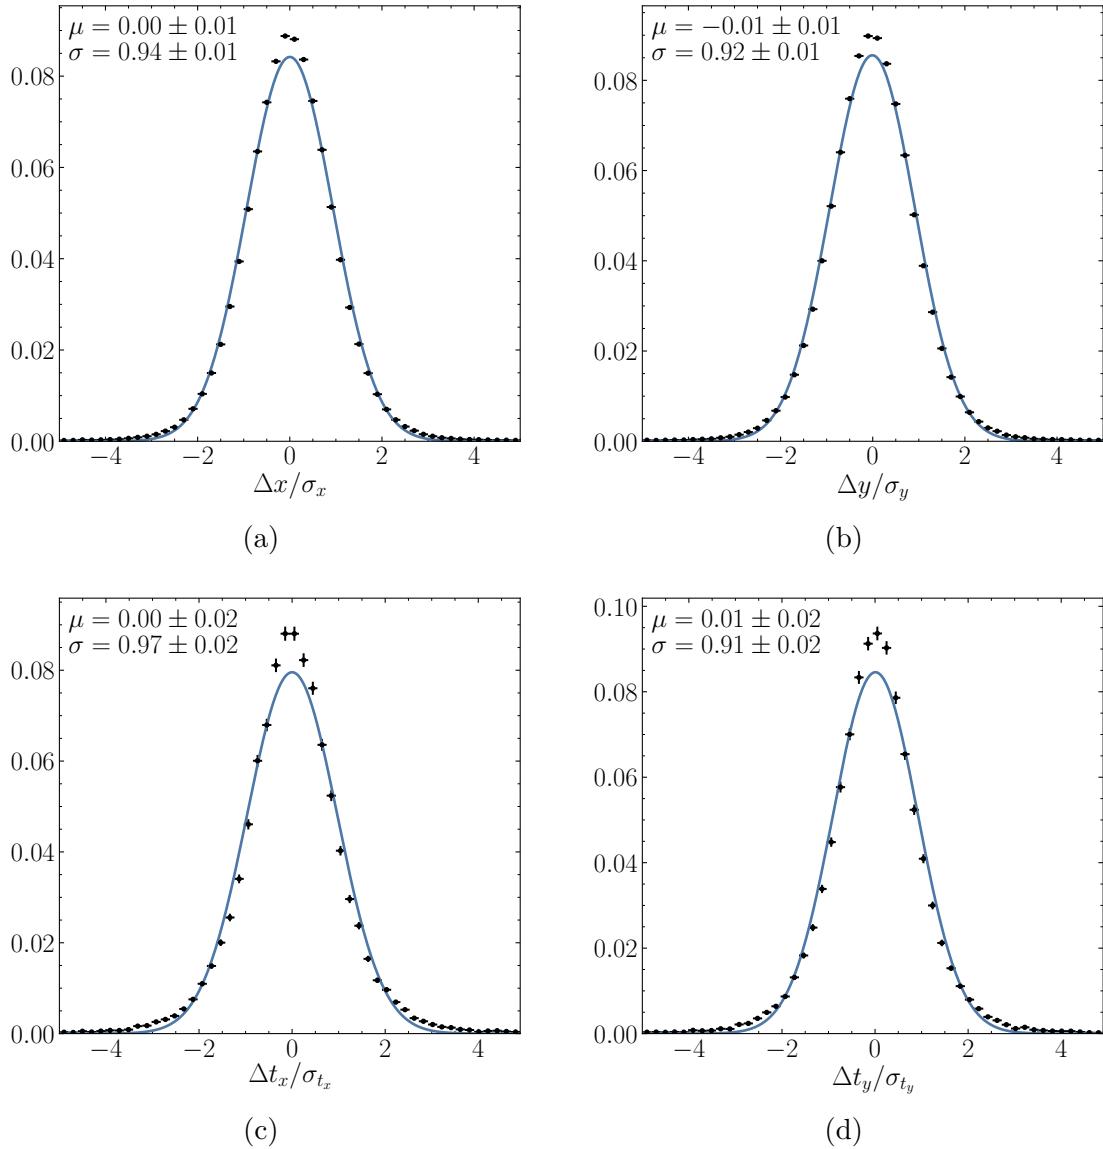


Figure 4-23: Pulls of the track state closest to the beamline determined using the simplified Kalman filter.

Table 4.3: Event selection rates for the Allen prototype trigger selections.

selection	rate [kHz]
TrackMVA	379 ± 11
TrackMuonMVA	24 ± 3
SingleHighPtMuon	5 ± 1
TwoTrackMVA	498 ± 12
DiMuonHighMass	125 ± 6
DiMuonLowMass	156 ± 7
Total	1013 ± 17

while all other parent particles must have a decay time greater than 0.2 ps. Final state charged particles must have $p_T > 250$ MeV and be within LHCb’s angular acceptance. The efficiency defition used here is preliminary and other definitions will be studied as Run 3 approaches. For example, efficiencies may be determined with respect to only those signal decays which can be fully reconstructed in LHCb’s offline reconstruction sequence and using only reconstructed trigger candidates matched to true simulated signal particles.

The selections presented here are based on simple rectangular cuts on the track and SV features. Experience from Runs 1 and 2 has shown that machine learning methods provide improved performance over simple cuts. As a result, the selections used in the Run 3 trigger will eventually be much more sophisticated than those shown here. Even these suboptimal selections demonstrate the power of LHCb’s Run 3 triggerless readout model. Comparing the efficiencies in 4-24 to those in 4-2 shows that Channels relying on hadronic signals, such as $B_s^0 \rightarrow \phi\phi$ and $D_s^+ \rightarrow K^+K^-\pi^+$, already show dramatic increases in efficiency relative to the hardware algorithms used in LHCb’s L0 trigger in Runs 1 and 2. Furthermore, these efficiencies are comparable to those predicted for a CPU implementation of HLT1 [148].

4.4.4 Throughput

The GPUs running Allen will be hosted by LHCb’s 170 event builder nodes. Each of these nodes can host three GPUs. Consequently, Allen must process the LHC’s 30 MHz event rate using fewer than 510 GPUs, resulting in a minimum throughput of about 60 kHz per card.

Allen’s throughput is tested on various GPUs. For each throughput measurement, Allen is run using 12 thread-stream pairs processing 1000 events each, repeating 100 times. Each measurement is repeated five times with a different sample of 1000 events, and the standard deviation is taken as the statistical uncertainty. Figure 4-25 shows Allen’s throughput on each card as a function of its maximum theoretical 32-bit TFLOPS. For each measurement, the statistical uncertainty is smaller than the plotted data point. Because of the attractive pricing of consumer GPUs, the NVIDIA RTX 2080 Ti is used as the primary benchmark card for throughput testing. Allen currently achieves throughputs of around 150 kHz on minimum bias pp collisions

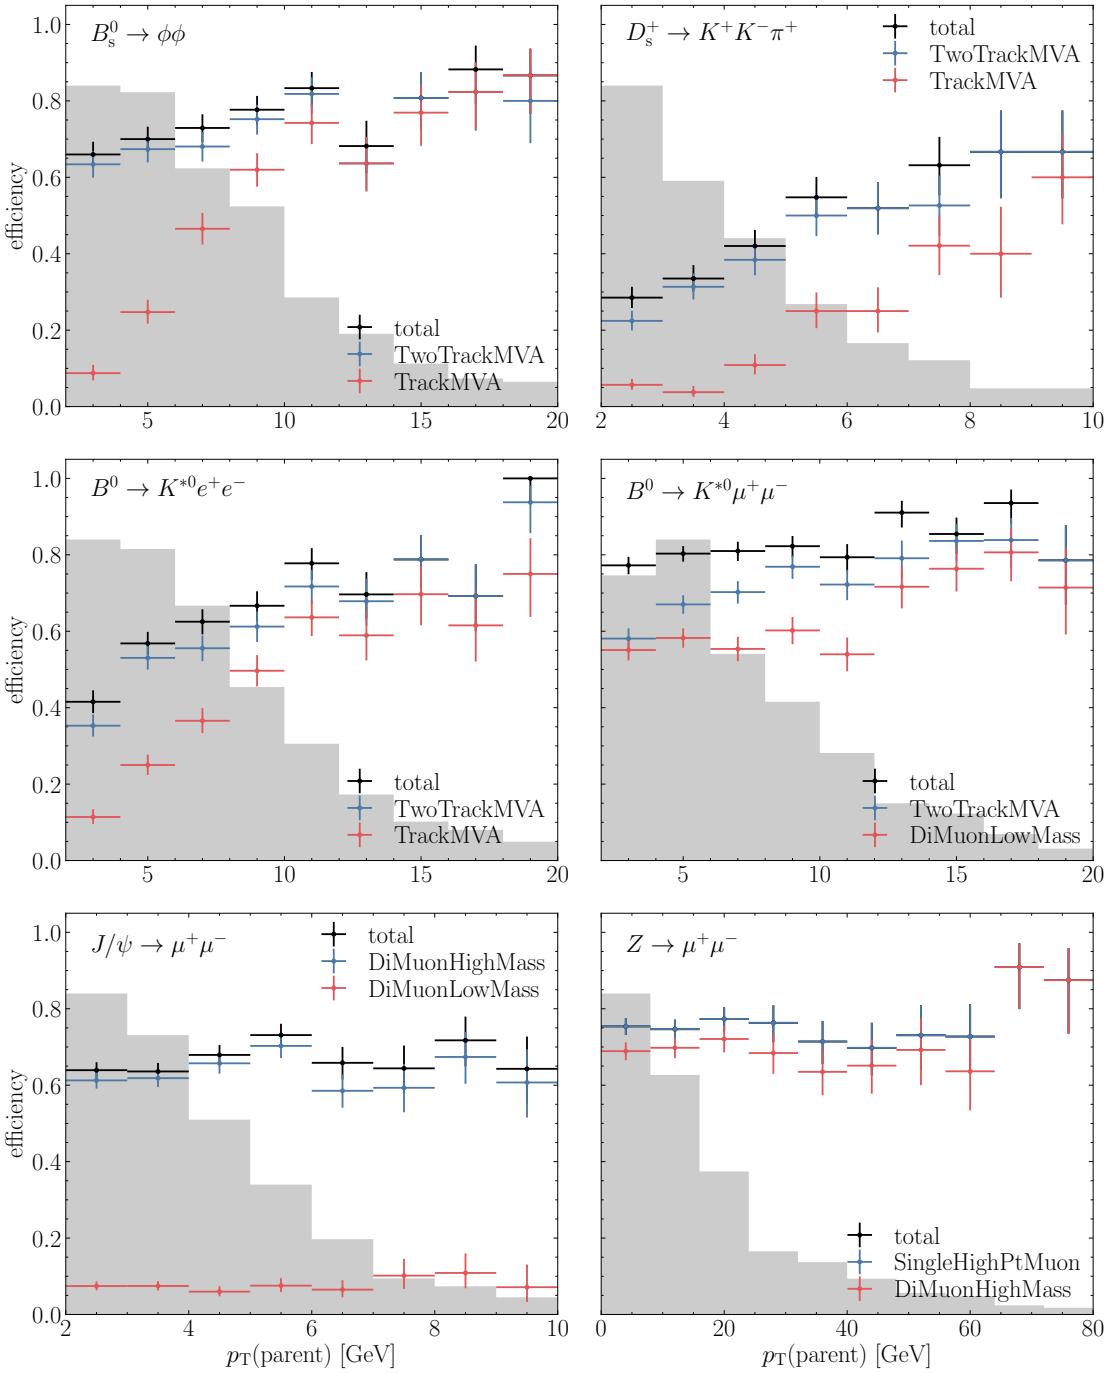


Figure 4-24: Trigger efficiencies for various simulated signal samples. The combined efficiency for all prototype trigger selections is shown in black, while the two most efficient selections are shown in blue and red. Efficiencies are calculated with respect to events with reconstructible signal decays as defined in the text.

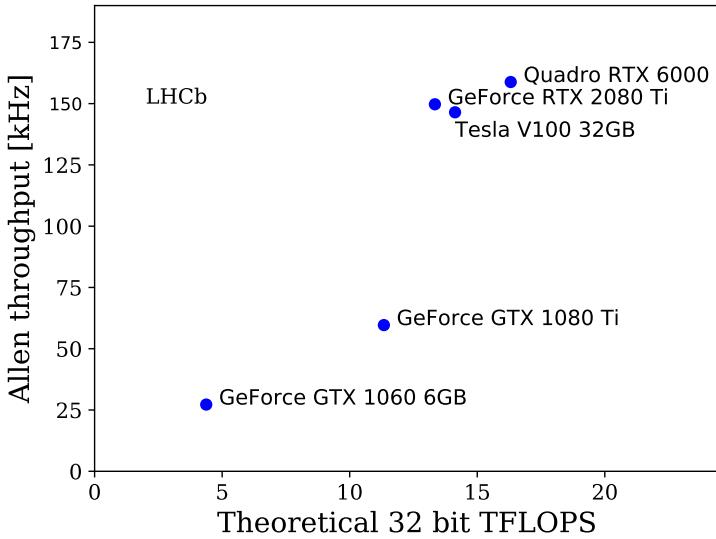


Figure 4-25: Allen throughput as a function of theoretical maximum 32-bit TFLOPS.

using this card. This is more than twice the throughput required to process the full LHC inelastic event rate. Figure 4-27 shows a breakdown of the contributions to Allen’s processing time.

Allen has been developed and tested using simulated events. Experience has shown that we can expect the occupancy of real collision events to be larger than in simulation. To test this, Allen’s throughput is measured using events with different SciFi data volumes. This gives Allen’s throughput as a function of detector occupancy. Figure 4-26 shows the results. This indicates that Allen can maintain a viable throughput at occupancies up to about 40% larger than those of simulated events.

Figure 4-25 demonstrates that Allen’s performance scales well with theoretical GPU performance. This indicates that Allen will be able to take advantage of incremental improvements in GPU performance. As a result, Allen’s throughput is expected to improve before the beginning of Run 3.

4.5 Conclusion

In this chapter I have presented the Allen project. Allen will be the first full software trigger executed on GPUs in a high energy physics experiment. Throughout this chapter I have demonstrated how Allen meets the minimum requirements set out in the LHCb Trigger Upgrade Technical Design Report [128]. Furthermore, Allen exceeds these goals with throughput headroom to spare. This will allow LHCb to pursue new trigger strategies in Run 3. Allen’s improvements in tracking threshold, for example, will increase the reconstruction efficiency of three- and four-track secondary vertices. This could allow LHCb to implement trigger lines similar to the HLT2 topological trigger [123,125] in HLT1. This could provide improved trigger efficiencies

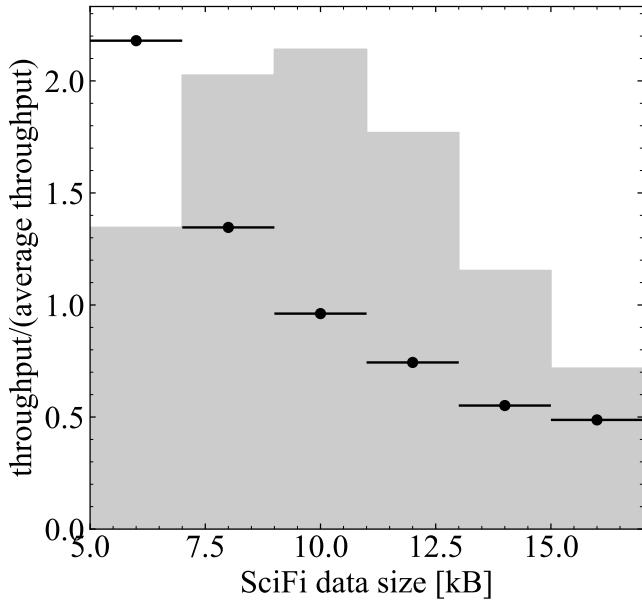


Figure 4-26: Allen throughput as a function of the size of the raw data produced by the SciFi. This measurement is performed using a NVIDIA Tesla V100 16GB GPU. The distribution of SciFi event sizes in simulated minimum bias events is shown in the shaded region.

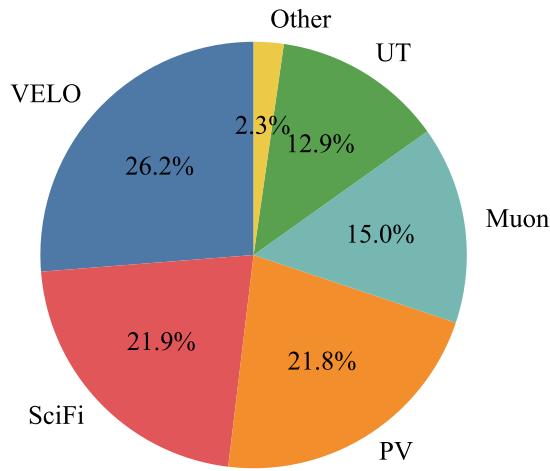


Figure 4-27: The relative contribution of each set of algorithms to the Allen processing time. Here “other” includes secondary vertex fitting, output preparation, selection evaluation, and Kalman filtering, which contribute 1.4%, 0.5%, 0.2%, and 0.2%, respectively.

at lower trigger rates.

Allen's throughput headroom also opens the possibility of reconstructing previously inaccessible physics signatures. In particular, this could allow for decoding and reconstruction of calorimeter clusters. This could allow for electron identification in HLT1. This has historically not been possible, and as a result LHCb's high profile analyses of final states containing electrons have relied on hadronic trigger signatures (see Ref. [149], for example). In these cases, dedicated electron triggers in HLT1 could result in increased trigger efficiencies. This includes the analysis of converted direct photon production presented in this thesis, which currently relies on relatively small samples of minimum bias events.

At its inception in 2017, the goals of the Allen project were extremely ambitious. They required rewriting an entire trigger system from the ground up for a new architecture. On paper, the results are impressive gains in throughput at an attractive cost with no sacrifice in physics performance. The less obvious result is that Allen will allow LHCb to expand its physics program in Run 3. This justification led LHCb to adopt Allen as its official baseline trigger system for Run 3 in June 2020 during the writing of this chapter.

Chapter 5

Intrinsic Charm and Forward $Z + c$ Production

This chapter presents a phenomenological study of LHCb’s ability to probe the intrinsic charm content of the proton by studying Z boson production in association with a charm quark. This contains and expands on work published in Ref. [150]. Since publication, Run 2 of the LHC has concluded. The measurement proposed in this chapter is nearing completion and should be published in early 2021. This chapter describes intrinsic charm and how $Z + c$ production can be used to look for it. It also demonstrates LHCb’s sensitivity to this phenomenon by presenting projected measurements assuming various intrinsic charm models. Finally, Bayesian reweighting is used to estimate the impact of these hypothetical measurements on global PDF fits.

5.1 Intrinsic Charm

Until recently, PDF fits have assumed that the heavy quark content of the proton is generated entirely perturbatively [151]. This means that at scales below some Q_0^2 , the heavy quark PDFs are zero. As Q^2 increases, heavy quarks are generated via DGLAP evolution as gluons split into sea c and b quarks. It then follows that the heavy quark PDFs have similar shapes to the sea quark PDFs.

This assumption was cast into doubt when the EMC experiment observed unexpectedly large charm hadron production rates in deep inelastic scattering [152]. To explain this excess, the authors of Ref. [153] proposed a non-perturbative contribution to the charm PDF. They proposed that the Fock state of the proton contains a $|uudc\bar{c}\rangle$ component. Because c quarks are heavy, it follows that they must then carry most of the proton’s momentum. This results in a valence-like momentum distribution.

The presence of a valence-like contribution to the charm quark PDF would have significant impacts on many experiments involving hadron collisions. This includes cross sections at the LHC, which would be affected by both a change in the charm quark PDF and the implied changes in the other PDFs from momentum sum rules. Intrinsic charm would also have a large effect on the cross section of potential Higgs-

mediated dark matter interactions with protons [154]. Intrinsic charm would also affect the rate of forward charm production by cosmic rays. When these charm hadrons decay semileptonically, they produce neutrinos that make up an important background in astrophysical neutrino searches [155, 156].

Intrinsic charm is typically quantified by the average momentum fraction carried by charm quarks in the proton. This is given by

$$\langle x \rangle_{\text{IC}} = \int_0^1 x c(x, Q^2 = Q_0^2) dx. \quad (5.1)$$

Global PDF analyses disagree on the allowed values of $\langle x \rangle_{\text{IC}}$, with some claiming that intrinsic charm is excluded at levels much less than the roughly 1% favored by DIS results [157]. This has sparked considerable debate in the theory community [158, 159]. Most of the data sensitive to the charm PDF at high x comes from fixed target DIS experiments, which typically have $Q < 10 \text{ GeV}$. Charm production at lower energy fixed target experiments is also sensitive to the charm PDF at high x , but at low Q^2 where hadronic and nuclear effects are important. As a result, this data is often excluded from PDF fits.

Because of the predicted valence-like momentum distribution of intrinsic charm, the ideal observable for studying intrinsic charm should be sensitive to the charm PDF at large x . In addition, it should ideally probe the proton at Q^2 large enough for hadronic and nuclear effects to be negligible in order to avoid the ambiguities that complicate interpretation of low energy fixed target experiments. Forward Z boson production in association with a charm quark satisfies these criteria. $Z + c$ production can proceed via gluon-charm scattering, as shown in Figure 5-1, providing direct sensitivity to the charm quark PDF. At forward rapidities, such as those covered by LHCb acceptance, one of these partons must have high x and the other low x . Furthermore, this must occur with a large momentum transfer due to the large mass of the Z boson. As a result, $Z + c$ production is an ideal probe of intrinsic charm. In practice, the cross section ratio $\sigma(Z + c)/\sigma(Z + j)$ is measured in order to cancel most systematic uncertainties.

Two classes of intrinsic charm models are considered here.

- *The BHPs model* has a valence-like momentum distribution with a large contribution at high x [153].
- *The SEA model* assumes the charm PDF has the same shape as the sea quark PDFs generated by DGLAP evolution, but is nonzero at the initial scale Q_0^2 .

These models were used in variations of the CT14 global analysis and are shown in Figure 5-2. For each model, two variations of the fit were produced using different values of $\langle x \rangle_{\text{IC}}$. The smaller of these values roughly corresponds to the best fit value of $\langle x \rangle_{\text{IC}}$, while the larger value is the largest value consistent with the default CT14 fit [160]. The low- $\langle x \rangle_{\text{IC}}$ models are denoted BHPs1 and SEA1, and the high- $\langle x \rangle_{\text{IC}}$ models are denoted BHPs2 and SEA2.

Since the publication of [150], the NNPDF collaboration released the first general purpose PDF set with a fitted intrinsic charm component [42]. The NNPDF3.1

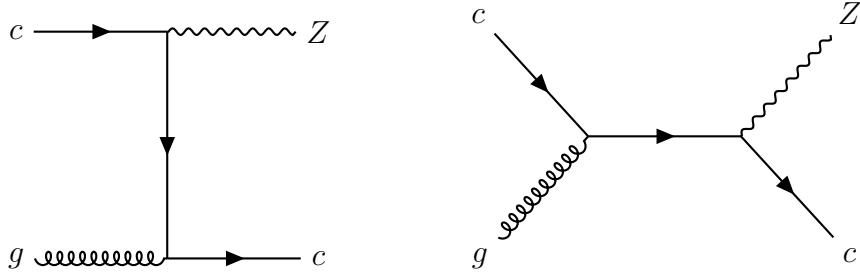


Figure 5-1: Leading order Feynman diagrams contributing to $Z + c$ production.

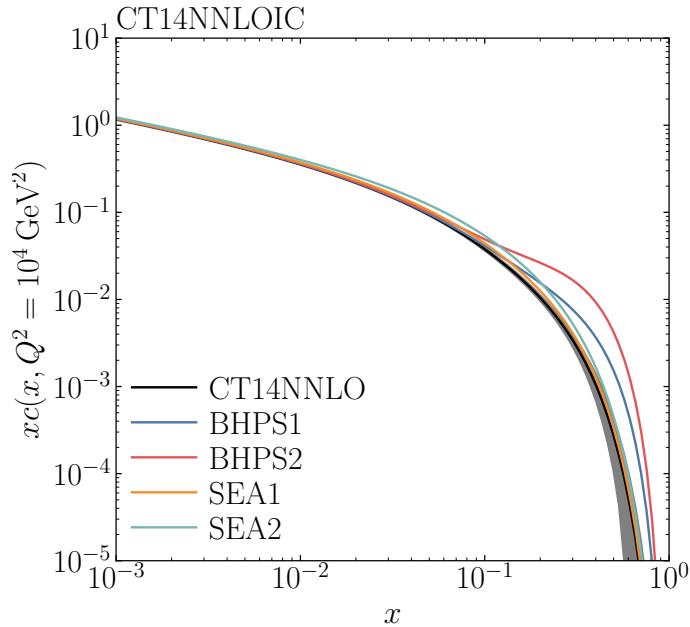


Figure 5-2: Intrinsic charm PDFs based on variations of the CT14 NNLO PDF set [160]. The gray band shows the CT14 NNLO PDF uncertainty.

charm PDF is parameterized with the flexibility to handle a valence-like intrinsic charm component. The NNPDF3.1 charm PDF is shown in Figure 5-3 compared to the charm PDF from NNPDF3.0 [161], which uses sea-like parameterization. The added flexibility in the NNPDF3.1 parameterization results in much larger PDF uncertainties at high x . The NNPDF3.1 fit favors smaller values of $\langle x \rangle_{\text{IC}}$ than those given by the CT14 analysis, but still favors a small intrinsic charm component at the level of about one- σ [42].

5.2 $Z + c$ Production at LHCb

The $Z + c$ and $Z + j$ cross sections are calculated for Run 2 ($\sqrt{s} = 13$ TeV) and Run 3 ($\sqrt{s} = 14$ TeV) using NLO $Z + j$ scattering matrix elements from POWHEGBOX [162]. These are combined with the CT14 NNLO PDF set [160]. Partons are showered

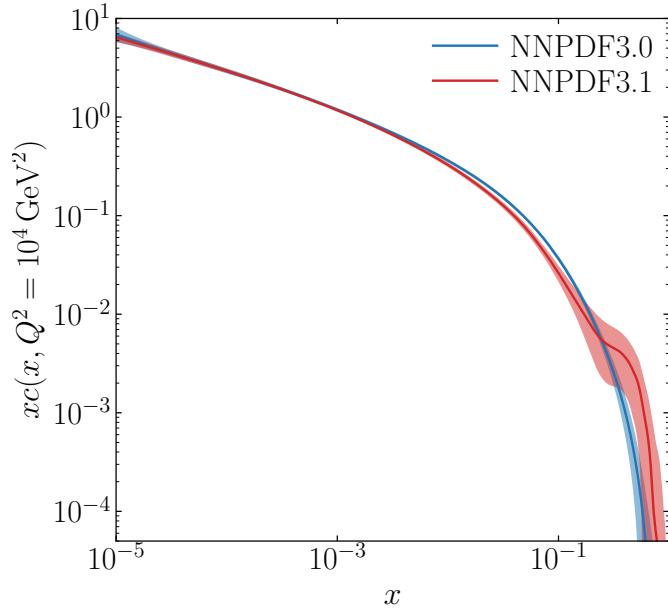


Figure 5-3: Comparison of NNPDF3.0 [161] and NNPDF3.1 [42] NNLO charm PDFs. The high- x structure in the NNPDF3.1 indicates that the global analysis favors a small valence-like intrinsic contribution to the charm quark PDF.

model	$\langle x \rangle_{\text{IC}}$
BHPS1	0.6%
BHPS2	2.0%
SEA1	0.6%
SEA2	1.5%

Table 5.1: $\langle x \rangle_{\text{IC}}$ values for each intrinsic charm model considered.

using PYTHIA8 [79] with the POWHEG method [163]. Hadronization is performed with PYTHIA, and hadron decays are performed by EvtGEN [164] interfaced with PHOTOS [165]. As a cross-check, the simulation is repeated using aMC@NLO matrix elements [130] and the FxFx parton shower method [166]. Visible final state particles with $2 < \eta < 5$ are clustered into jets with FASTJET [167] using the anti- k_t algorithm [168] with radius parameter $R = 0.5$.

Only decays of $Z \rightarrow \mu^+ \mu^-$ are considered, and Z bosons are required to have $60 < m_{\mu\mu} < 120$ GeV. Muons must have $p_T > 20$ GeV and $2.0 < \eta < 4.5$ for Run 2 and $2.0 < \eta < 5.0$ for Run 3. Jets must have $p_T > 20$ GeV and $2.2 < \eta < 4.2$. This results in nearly uniform c -jet tagging efficiency over the entire fiducial region. Only the highest- p_T jet in each event is considered. Jets are classified as charm jets if they contain a long-lived, promptly-produced c -hadron with $p_T > 2$ GeV. Statistical uncertainties are calculated by assuming an integrated luminosity of 5 fb^{-1} in Run 2 and 15 fb^{-1} in Run 3. Calculated yields are modified by an expected jet efficiency of 96% and a global event cut efficiency of 90% [103]. A charm jet tagging efficiency of about 25% was achieved in Run 1 [169] and is assumed for both Runs 2 and 3. The charm jet fake rate in Run 1 was found to be about 0.3% and is assumed to be negligible here.

The dominant systematic uncertainty in this measurement will be due to the c -jet tagging efficiency. In Run 1, the b - and c -jet tagging efficiencies were measured in data simultaneously and are strongly anticorrelated. This resulted in a c -jet tagging efficiency uncertainty of about 10%. In the future, this can be constrained by measuring $\sigma(c\bar{c})/\sigma(b\bar{b})$. With this in mind, a 5% relative systematic uncertainty is assumed.

Results are shown in Figure 5-4. The BHPs models have maximum impact at high rapidity, corresponding to high x . Meanwhile, the SEA models produce larger effects at low $y(Z)$. Figure 5-4 indicates that LHCb will be sensitive to BHPs-like intrinsic charm at the level of $\langle x \rangle_{\text{IC}} > 0.3\%$ and sea-like intrinsic charm at the level of $\langle x \rangle_{\text{IC}} > 1\%$.

5.3 Bayesian Reweighting

Reweighting techniques can be used to study the impact of new data on PDF fits without repeating the full global analysis. Here, Bayesian reweighting [170, 171] is used to study the effects of the hypothetical $\sigma(Z + c)/\sigma(Z + j)$ measurements from the previous section on the NNPDF3.1 fit.

The NNPDF3.1 PDF fit [42] uses an artificial neural network to parameterize the PDFs at the initial state Q^2 . The experimental data used in the fit is used to generate pseudodata sampled from Gaussian distributions centered at each data point. The PDF fit is then repeated with each pseudodata set, creating an ensemble of replica PDF sets. The central value of the PDF is then given by the mean of the PDF replica sets, and the uncertainty is given by the standard deviation. To study the effects of new data on the PDF fit, each replica is given a weight based on its compatibility with the new data. The PDF weights can be calculated using Bayes' theorem, as

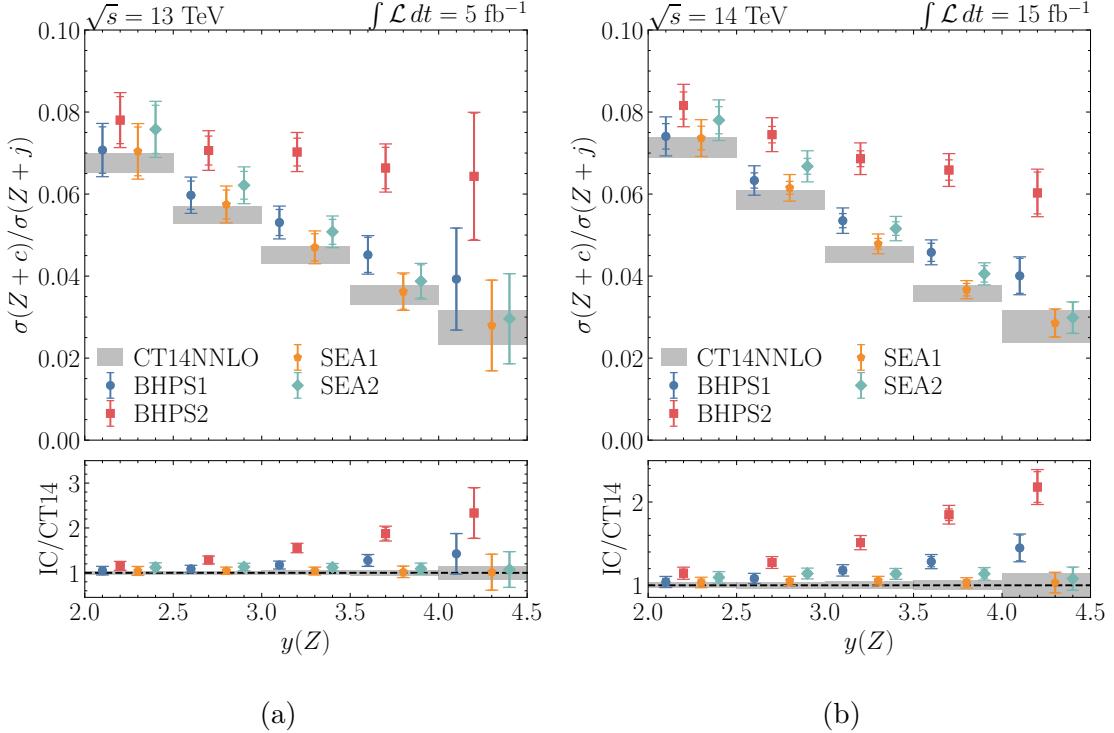


Figure 5-4: Expected measurements of $\sigma(Z + c)/\sigma(Z + j)$ as a function of $y(Z)$ for (a) Run 2 and (b) Run 3 calculated using various intrinsic charm models.

shown in Ref. [171]. Replica predictions are produced using each replica PDF and a χ^2 is given by

$$\chi^2 = (\vec{y} - \vec{y}_f) \sigma^{-1} (\vec{y} - \vec{y}_f), \quad (5.2)$$

where \vec{y} is the new measurement, \vec{y}_f is the prediction calculated using PDF replica f , and σ is the measurement covariance matrix. The weight of each PDF is then given by

$$w_k = \chi_k^{n-1} \exp\left(-\frac{1}{2}\chi_k^2\right), \quad (5.3)$$

where k denotes the PDF replica and n is the number of data points in the new measurement.

Here the NNPDF3.1 replicas are reweighted using the hypothetical Run 2 measurements from the previous section. The 5% uncertainty from the c -tagging efficiency measurements is assumed to be fully correlated between each data point. The raw and weighted χ^2/ndf distributions are shown in Figure 5-5. In each case, the reweighted distribution peaks near one, indicating that the reweighted PDF set describes the hypothetical $\sigma(Z + c)/\sigma(Z + j)$ measurements reasonably well.

After reweighting, some PDF replicas will have very small weights and no longer contribute significantly to ensemble averages. This reflects information loss due to the reweighting procedure [170]. The reweighted PDF ensemble will then have the

same accuracy as a refit with N_{eff} replicas, where N_{eff} is given by

$$N_{\text{eff}} = \exp \left(\frac{1}{N} \sum_{k=1}^N w_k \ln(N/w_k) \right), \quad (5.4)$$

where N is the total number of PDF replicas. If N_{eff} is small, then the new data either provides a large amount of information or is incompatible with the original fit [170]. Here $N = 1000$, and $N_{\text{eff}} > 400$ for each hypothetical measurement. This indicates that reweighted PDF sets produced using LHCb's eventual $Z + c$ measurement are likely to be accurate.

Comparisons of the NNPDF3.1 charm PDF before and after reweighting are shown in Figure 5-6 for $Q^2 = 10^4 \text{ GeV}^2$. The resulting charm PDFs all lie within one- σ of the default NNPDF3.1 charm PDF over all values of x , indicating that the hypothetical measurements presented in this chapter would all be consistent with the world's data in this kinematic region. As expected, the BHPS1 measurement results in almost no change to the central PDF value, reflecting the fact the the NNPDF3.1 fit favors a small valence-like intrinsic charm component. The BHPS2 measurement results in a large increase in the charm PDF at high x , which is disfavored at the one-*sigma* level for $x > 0.3$. Both sea-like models result in decreases in the charm PDF at high x . This again reflects small valence-like intrinsic charm component favored by the default fit that is absent in the sea-like PDFs.

Regardless of the intrinsic charm content of the proton, a measurement of $\sigma(Z + c)/\sigma(Z + j)$ at LHCb will provide constraints on the charm quark PDF at high x . Figure 5-7 shows the relative charm quark PDF uncertainties with and without the hypothetical measurements from this chapter. This shows that BHPS1-, SEA1-, and SEA2-like measurements would slightly decrease PDF uncertainties for $x > 0.2$, while a BHPS2-like measurement would provide powerful constraints for $x > 0.1$. This demonstrates that LHCb could potentially provide definitive proof of intrinsic charm in the proton in the next couple of months.

5.4 Conclusion

This chapter has demonstrated LHCb's ability to probe the intrinsic charm quark content of the proton by studying Z production in association with a charm quark. This chapter has also shown how the impact of this measurement on global PDF fits can be assessed. This measurement is currently in collaboration-wide review. The reweighting methods presented in this chapter can easily be applied to the actual measurement when it is available.

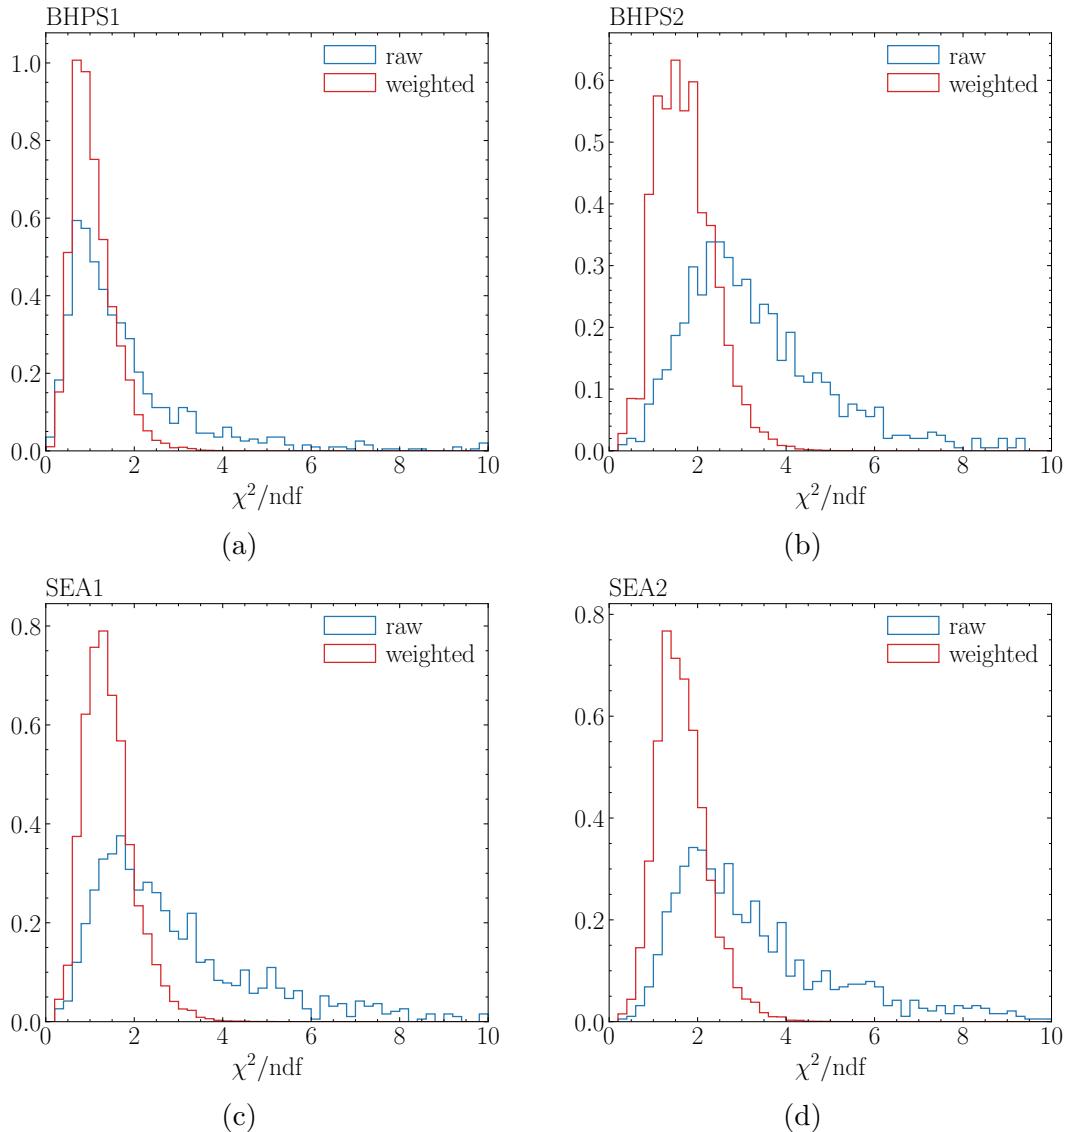


Figure 5-5: Comparison raw and weighted χ^2/ndf distributions calculated using the NNPDF3.1 replicas and hypothetical measurements assuming (a) BHPS1, (b) BHPS2, (c) SEA1, and (d) SEA2 intrinsic charm models.

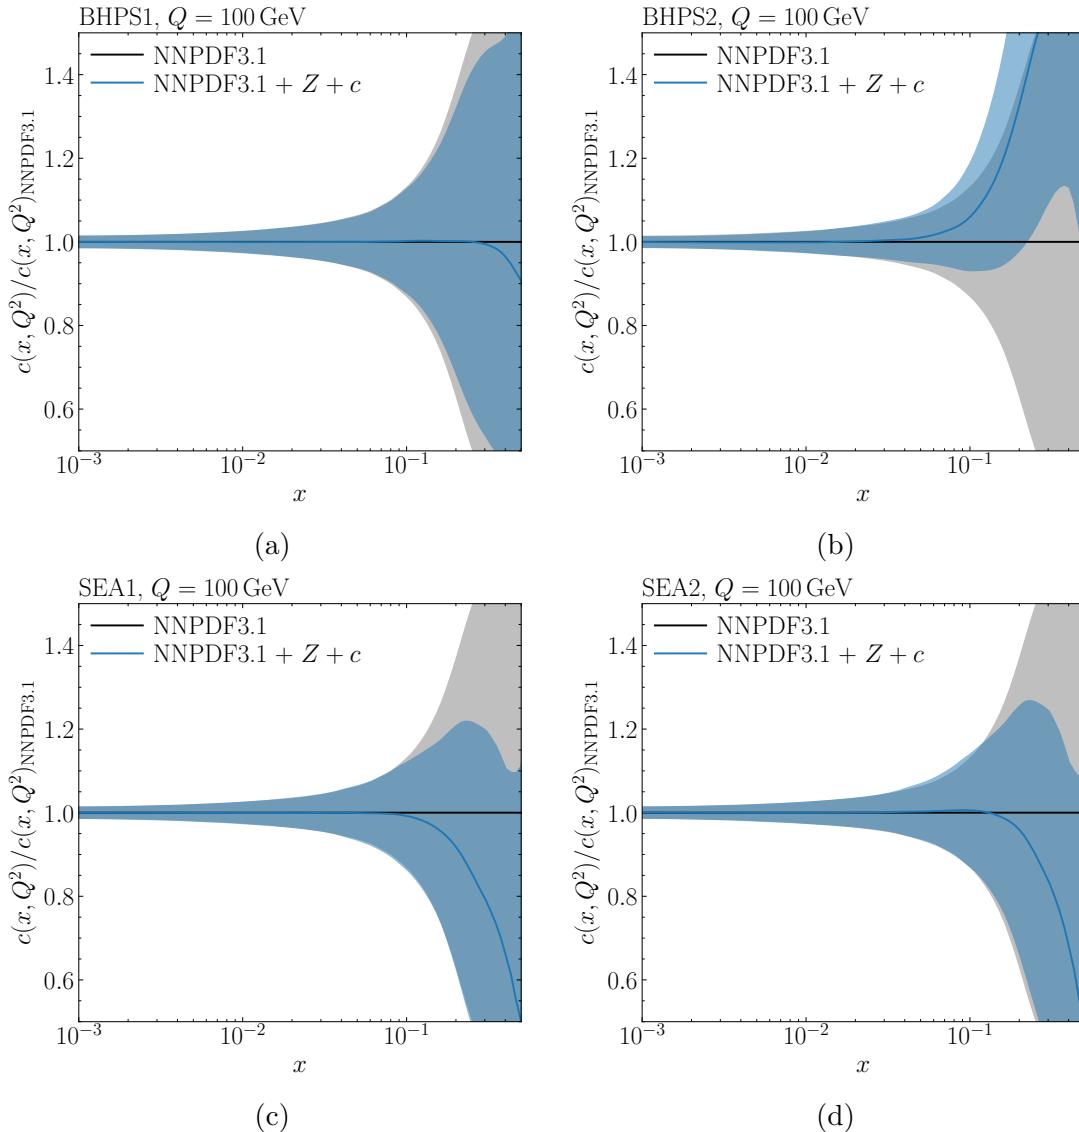


Figure 5-6: Comparison of charm PDFs at $Q^2 = 10^4 \text{ GeV}^2$ with and without hypothetical $Z + c$ measurements assuming (a) BHPs1, (b) BHPs2, (c) SEA1, and (d) SEA2 intrinsic charm models.

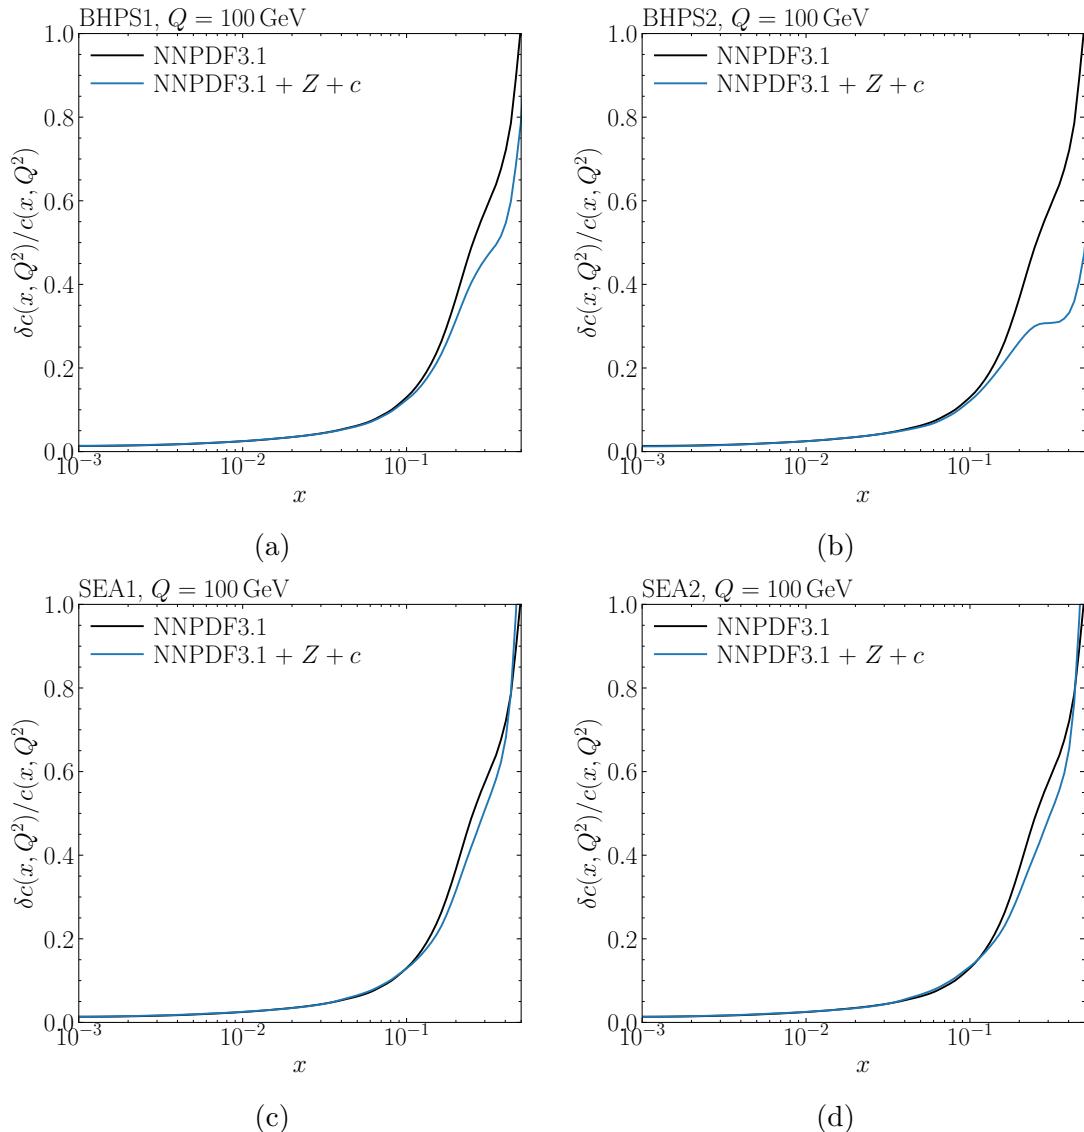


Figure 5-7: Comparison of charm PDF uncertainties with and without hypothetical $Z + c$ measurements assuming (a) BHPS1, (b) BHPS2, (c) SEA1, and (d) SEA2 intrinsic charm models.

Chapter 6

Neutral Pion and Photon Production in Proton-Lead Collisions

The LHCb detector was designed to study heavy flavor decays in relatively low multiplicity proton-proton collisions with no plans to record heavy ion collisions. However, the LHCb detector’s design also provides unique opportunities for studying heavy ion physics at the LHC. The LHCb detector is fully instrumented in the forward region, providing sensitivity to parton distributions at high and low x . Studying proton-lead collisions in these kinematic regions provides information about the cold nuclear matter effects that modify the structure of the nucleus.

This chapter presents measurements of direct photon production in proton-proton and proton-lead collisions using the LHCb detector. Direct photon production is sensitive to the nuclear gluon distribution and could potentially provide evidence of thermal photon production. Measuring direct photon production requires understanding backgrounds from meson decays. As a result, measurements of π^0 production are also presented in this chapter.

Section 6.1 explains the observables measured and outlines the strategies for measuring them, and Section 6.2 presents the data and simulation samples used to perform these measurements. Section 6.3 then explains how particles of interest are reconstructed and selected, and Section 6.4 presents the fit strategy used to extract neutral pion yields. Finally, Sections 6.5 and 6.6 present measurements of the π^0 nuclear modification factor and the direct photon excess ratio, respectively.

6.1 Observables

6.1.1 Proton-Lead Collisions at LHCb

LHCb recorded proton-lead collisions in 2013 at $\sqrt{s_{\text{NN}}} = 5.02 \text{ TeV}$ and in 2016 at $\sqrt{s_{\text{NN}}} = 8.16 \text{ TeV}$. Because of the detector’s single arm design, proton-lead collisions

were recorded in two different configurations:

- $p\text{Pb}$: The proton beam travels in $+z$ direction in the LHCb coordinate system, moving towards the spectrometer. In the center-of-mass frame, rapidities in the LHCb acceptance are positive in the $p\text{Pb}$ configuration.
- Pbp : The proton beam travels in the $-z$ direction, while the Pb beam travels in the $+z$ direction. Center-of-mass frame rapidities in LHCb acceptance in the Pbp configuration are negative.

Throughout this chapter, “proton-lead” will refer to collisions involving a proton and a lead ion in general, while “ $p\text{Pb}$ ” and “ Pbp ” will refer to collisions using the designated beam configuration.

In proton-ion collisions at the LHC, the ion beam energy per nucleon E_{NN}^A is given by

$$E_{NN}^A = \frac{ZE_p}{A}, \quad (6.1)$$

where E_p is the proton beam energy, Z is the atomic number of the nucleus, and A is the number of nucleons in the nucleus. As a result, the beam energies in proton-lead collisions are asymmetric in the laboratory frame. This energy asymmetry results in a rapidity shift of -0.465 in the center of mass frame relative to the laboratory frame in both $p\text{Pb}$ and Pbp collisions.

6.1.2 The Direct Photon Excess Ratio

Direct photon production at the LHC is accompanied by copious photon production from hadron decays. As a result, the direct photon contribution is quantified using the direct photon excess ratio given by

$$R^\gamma = \frac{\sigma_{\text{inc}}^\gamma}{\sigma_{\text{dec}}^\gamma}, \quad (6.2)$$

where $\sigma_{\text{inc}}^\gamma$ and $\sigma_{\text{dec}}^\gamma$ are the inclusive photon and hadron decay photon production cross sections, respectively. R^γ can be expressed as a double ratio given by

$$R^\gamma = \frac{\left(\sigma_{\text{inc}}^\gamma / \sigma^{\pi^0}\right)_{\text{data}}}{\left(\sigma_{\text{dec}}^\gamma / \sigma^{\pi^0}\right)_{\text{ref}}}, \quad (6.3)$$

where σ^{π^0} is the production cross section of photons from π^0 decays. The photon and π^0 yields are measured in both data and a reference sample, denoted “ref” here. The reference sample may be a Monte Carlo sample or a data sample with no direct photon contribution. Many systematic uncertainties conveniently cancel in Eq. 6.3,

which can be expressed in terms of experimentally measurable quantities as

$$R^\gamma = \frac{\left(N_{\text{inc}}^\gamma / N^{\pi^0}\right)_{\text{data}}}{\left(N_{\text{dec}}^\gamma / N^{\pi^0}\right)_{\text{ref}}} \times \frac{\epsilon_{\text{ref}}^\gamma}{\epsilon_{\text{data}}^\gamma} \times \frac{\epsilon_{\text{data}}^{\pi^0}}{\epsilon_{\text{ref}}^{\pi^0}} \times \frac{f_{\text{data}}^{\pi^0}}{f_{\text{ref}}^{\pi^0}}, \quad (6.4)$$

where N^{γ, π^0} is the specified particle yield, ϵ^{γ, π^0} is the reconstruction efficiency, and f^{π^0} is the fraction of decay photons originating from π^0 decays.

This chapter presents measurements of R^γ in pp , $p\text{Pb}$, and $\text{Pb}p$ collisions. The inclusive photon yield is measured using photons that convert to electron-positron pairs in the detector material upstream of the LHCb dipole magnet. Converted photons are combined with unconverted photons (also called “ECAL photons”) to reconstruct π^0 candidates. The converted photon efficiency cancels in Eq. 6.4, which reduces to

$$R^\gamma = \frac{\left(N_{\text{inc}}^\gamma / N^{\pi^0}\right)_{\text{data}}}{\left(N_{\text{dec}}^\gamma / N^{\pi^0}\right)_{\text{ref}}} \times \frac{\epsilon_{\text{ref}}^{\gamma, \text{ECAL}}}{\epsilon_{\text{data}}^{\gamma, \text{ECAL}}} \times \frac{f_{\text{data}}^{\pi^0}}{f_{\text{ref}}^{\pi^0}}. \quad (6.5)$$

Reference samples are constructed from simulation. The direct photon signal is expected to be negligible in 13 TeV pp , so the measurement is also performed using 13 TeV pp data as a reference sample as a cross-check.

6.1.3 Nuclear Modification Factors

One of the primary goals of heavy ion physics is to study how heavy ion collisions differ from superpositions of many nucleon-nucleon collisions. Modifications are quantified by the nuclear modification factor, denoted R_{pA} for proton-ion collisions and R_{AA} for ion-ion collisions. Here A is the species of the nucleus. Nuclear modification factors are often calculated by rescaling per-event yields in pp collisions by the average number of individual nucleon-nucleon collisions, N_{coll} . N_{coll} can be calculated using the Glauber model of heavy ion collisions [172]. In ion-ion collisions, N_{coll} can be calculated as a function of the impact parameter of the ions, allowing for measurements of nuclear modification factors in different centrality regions.

This chapter presents measurements of $R_{p\text{Pb}}^{\pi^0}$ in minimum bias collisions. Instead of expressing the nuclear modification factor in terms of N_{coll} , it is given as a ratio of measured cross sections as [58]

$$R_{p\text{Pb}}(p_{\text{T}}) = \frac{1}{A} \frac{d\sigma_{p\text{Pb}}/dp_{\text{T}}}{d\sigma_{pp}/dp_{\text{T}}}, \quad (6.6)$$

where $\sigma_{p\text{Pb}}$ and σ_{pp} are the cross sections of a particular process in proton-lead and proton-proton collisions, respectively. Here A is the number of nucleons in the nucleus. For proton-lead collisions at the LHC, $A = 208$.

6.1.4 Forward-Backward Ratio

Because LHCb records collisions in both $p\text{Pb}$ and $\text{Pb}p$ configurations, measurements at forward and backward rapidities can be compared. This is quantified using the forward-backward ratio R_{FB} , given by

$$R_{\text{FB}}(p_{\text{T}}) = \frac{d\sigma_{p\text{Pb}}/dp_{\text{T}}}{d\sigma_{\text{Pb}p}/dp_{\text{T}}}, \quad (6.7)$$

where $\sigma_{p\text{Pb}}$ and $\sigma_{\text{Pb}p}$ are cross-sections measured in the $p\text{Pb}$ and $\text{Pb}p$ configurations, respectively. This ratio is typically constructed in rapidity regions defined such that $y_1 < y_{p\text{Pb}}^* < y_2$ and $-y_2 < y_{\text{Pb}p}^* < -y_1$, where y^* is the rapidity in the center-of-mass frame and y_1 and y_2 define the measurement's fiducial region. However, this is not necessary, and ratios of cross sections in other pairs of rapidity windows can also be constructed. The forward-backward ratio is particularly useful for studying proton-lead collisions when no pp reference sample is available [58].

6.2 Data and Simulation

6.2.1 Data

The measurements presented in this chapter use four datasets summarized in Table 6.1. All of the measurements in this chapter use data selected with no (no bias) or minimal (minimum bias) trigger selections. No bias data is collected by triggering on random bunch crossings at a given rate. The no bias rate at LHCb is typically kept to a few hertz. Minimum bias events are similar to no bias events, but require a single reconstructed VELO track.

LHCb collected large samples of no bias pp data at $\sqrt{s} = 5 \text{ TeV}$ and 13 TeV in 2015. These data samples are made up of collisions between bunches that lead a bunch train. Leading bunch collisions are free from detector signals left over from previous bunch crossings. LHCb also collected large samples of $p\text{Pb}$ and $\text{Pb}p$ data at $\sqrt{s_{\text{NN}}} = 8.16 \text{ TeV}$ in 2016. These data were collected using a minimum bias trigger selection. In addition, LHCb collected minimum and no bias data in pp collisions at $\sqrt{s} = 7$ and 8 TeV and in $p\text{Pb}$ and $\text{Pb}p$ collisions at $\sqrt{s_{\text{NN}}} = 5 \text{ TeV}$. However, these samples are small and are not used for the measurements presented in this chapter.

6.2.2 Simulation

Simulated data samples were produced using the GAUSS software package [173]. Proton-proton collisions were generated using PYTHIA8 [79], while proton-lead collisions were generated using EPOS [174]. Unstable hadron decays were performed using EVTGEN [164]. Final state radiation was simulated using PHOTOS [165], and the LHCb detector was simulated using GEANT4 [175]. The simulated detector signals are then digitized using BOOLE [173].

year	beams	$\sqrt{s_{(\text{NN})}}$	$\mathcal{L} [\text{nb}^{-1}]$
2015	pp	13 TeV	9.7
2015	pp	5.02 TeV	3.4
2016	$p\text{Pb}$	8.16 TeV	12.5
2016	$\text{Pb}p$	8.16 TeV	17.4

Table 6.1: Summary of datasets used for the analyses presented in this chapter. The listed pp luminosities are the luminosities measured in Section 6.5.3, while the $p\text{Pb}$ and $\text{Pb}p$ are the luminosities delivered to LHCb by the LHC during the 2016 proton-lead runs.

The measurements in this chapter use a simulated no bias pp collisions with $\sqrt{s} = 13 \text{ TeV}$. No large sample of simulated no bias 8.16 TeV $p\text{Pb}$ or $\text{Pb}p$ collisions exists. Instead, simulated $p\text{Pb}$ and $\text{Pb}p$ collisions with b hadron decays are used. The b hadron decay products are removed in order to mimic minimum bias collisions.

6.3 Reconstruction and Selection

6.3.1 Converted Photons and Neutral Pions

Most photons in hadron collisions are the products of $\pi^0 \rightarrow \gamma\gamma$ decays. At p_{T} above about 2 GeV, the opening angle between π^0 decay photons in LHCb acceptance is small enough that the photons are reconstructed as a single cluster in the ECAL. This serves as a background to photons. To avoid this background, R^{γ} is measured using converted photons (γ^{cav}). Photons that convert upstream of the LHCb dipole magnet are reconstructed as pairs of oppositely charged tracks. Because of LHCb’s excellent tracking, converted photons have better momentum resolution than ECAL photons (γ^{cal}) [176]. However, the probability of a photon converting before the LHCb magnet is only about 20%, resulting in far fewer reconstructed converted photon candidates than ECAL photons.

Photons that convert in the VELO material are typically reconstructed as pairs of long tracks (LL photons), while photons that convert too late to form VELO tracks are reconstructed as pairs of downstream tracks (DD photons). The electrons from LL photons traverse more of the detector than those from DD photons. As a result, LL photons lose more energy via bremsstrahlung and have worse momentum resolution than DD photons. The analyses in this chapter primarily use DD photons.

The converted photon reconstruction and selection used here is similar to that used in Ref. [176]. Only photons with $2.5 < \eta < 4.0$ are considered in order to avoid regions of decreased acceptance near the edges of the detector. Converted photons are reconstructed from tracks with $p_{\text{T}} > 200 \text{ MeV}$. Fake (or ghost) tracks are removed using a neural network. The response of this neural network is referred to as `GhostProb` [177]. Tracks from converted photon candidates are required to be within ECAL acceptance. Because the photon is massless, dielectrons are produced in photon conversions with almost no opening angle. As a result, the tracks from

a converted photon candidate are required to have a difference in y position at the ECAL face of less than 3σ , where σ is the track position resolution. For pairs of downstream tracks, the opening angle resolution is dependent on the z position of the reconstructed vertex. As a result, the measured converted photon mass must be less than a z -dependent maximum value.

Tracks are identified as electrons using a combination of information from the RICH and calorimeter systems. Likelihoods for the electron and charged pion hypotheses are calculated for each subdetector and combined. Electrons are selected based on the difference in log likelihood between the electron and pion hypotheses ($\text{DLL}_{e-\pi}$). Electron identification is driven primarily by information from the ECAL. A track is considered “electron-like” if its extrapolated position at the ECAL face is within or near an ECAL cluster and has p/E near unity, where p is the measured track momentum and E is the ECAL cluster energy. Tracks from converted photon candidates must have $\text{DLL}_{e-\pi} > 0$.

Reconstructed converted photons are required to have an impact parameter significance $\chi_{\text{IP}}^2 < 200$ to insure that they originate from a primary vertex. In addition, the converted photon momentum must point back to the PV. This is enforced using the feature δ given by

$$\delta = \frac{|\vec{p} \times \hat{d}|}{p_{\text{T}}}, \quad (6.8)$$

where \vec{p} and p_{T} are the converted photon momentum vector and transverse momentum, respectively, and \hat{d} is the unit vector pointing between the PV and SV. Selected converted photons must have $\delta < 0.03$, their SVs must be downstream of the PV.

The purities of the selected samples are studied both in simulation and using same-charge dielectrons (also called “same-sign” or “SS”) to model combinatoric background. Selected converted photon and SS mass spectra are shown in Figure 6-1. The SS samples show that all data and simulation samples have combinatoric background contaminations of 2–2.5%. The background fraction is also determined in simulation by matching reconstructed converted photon candidates to generated photons. The resulting background contributions are about 0.5%. This discrepancy arises because the combinatoric background is made up mostly of converted photon candidates that share a track with another candidate in the event. These clone candidates are very close to parallel and may be matched to the same generated photon in simulation.

Converted photons are combined with ECAL photons to form π^0 candidates. ECAL photons must not be matched to conversion electrons and must have $p_{\text{T}} > 400 \text{ MeV}$ and $2.0 < \eta < 4.5$. Selection criteria for converted and ECAL photons are summarized in Table 6.2.

6.3.2 Proxy Photons

Neutral pions and charged pions have similar production mechanisms in hadron collisions. Throughout this chapter, charged tracks are used as proxies for neutral pions to simulate decay photon production in real collisions. Proxy π^0 candidates are long tracks with $\text{GhostProb} < 0.3$ and $\chi_{\text{IP}}^2 < 16$. Tracks are identified using a combination

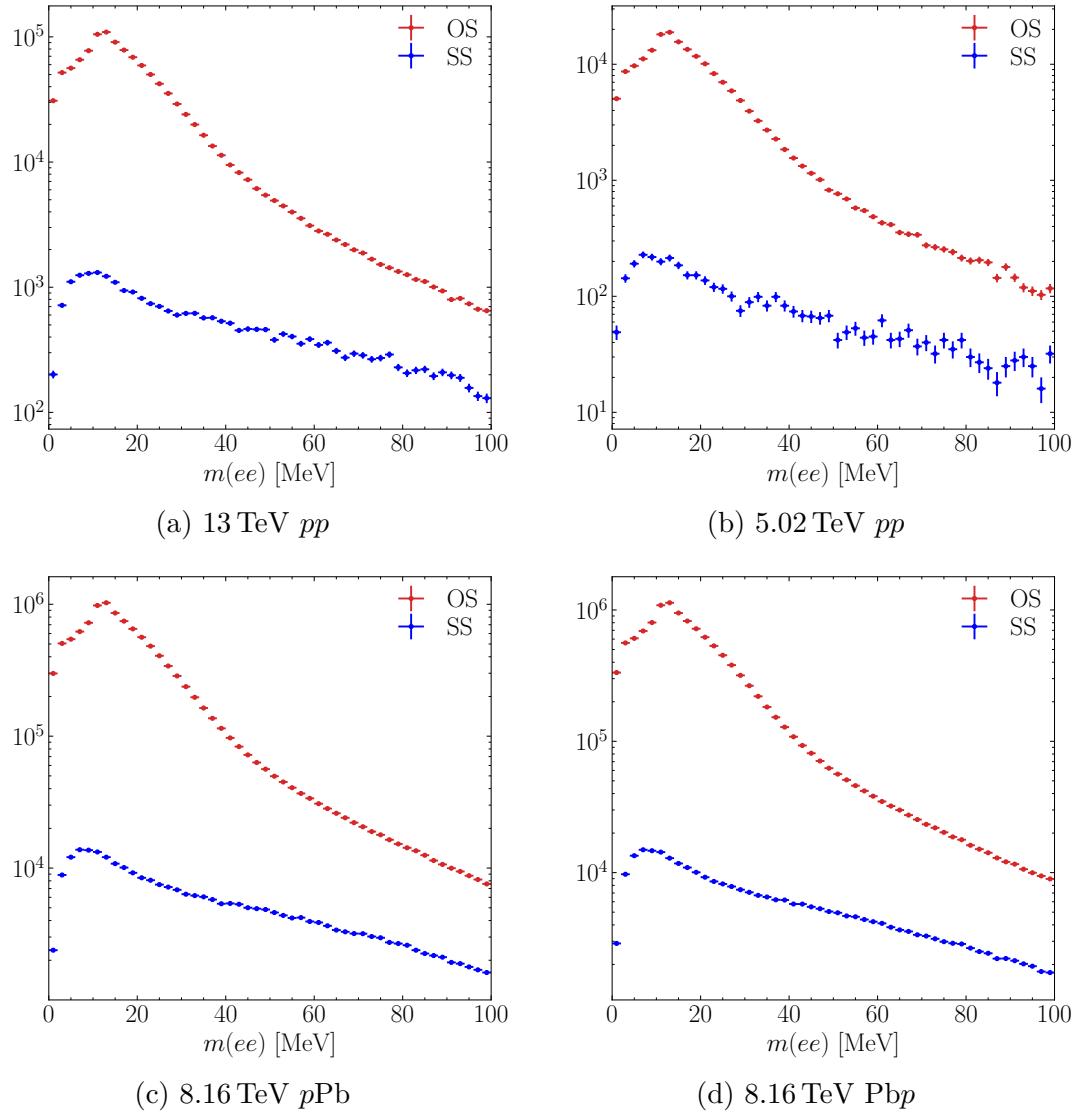


Figure 6-1: Converted photon (OS) and SS mass spectra after the full selection.

e^\pm	$p_T > 200 \text{ MeV}$ $\text{DLL}_{e-\pi} > 0$ In ECAL acceptance
γ^{cav}	$2.5 < \eta < 4.0$ $\chi_{\text{IP}}^2 < 200$ $\Delta y(e^+e^-) < 3\sigma$ $\Delta z(\text{PV, SV}) > 0$ $M < (20 + 0.04z_{\text{vtx}}/\text{mm})/\text{MeV}$ $\delta < 0.03$
γ^{cal}	$2.0 < \eta < 4.5$ $p_T > 400 \text{ MeV}$

Table 6.2: Summary of selection criteria for converted and ECAL photons.

of information from the different particle identification systems, including the RICH, calorimeters, and muon systems. Information from the different detectors is combined using a neural network. Different neural networks are trained for each particle species. The output of the π^\pm identification network is referred to as PNN_π . Proxy π^0 's must have $\text{PNN}_\pi > 0.1$. Each proxy candidate is assigned the π^0 mass hypothesis and its momentum is used to simulate a decay to two photons. One photon is randomly chosen as the converted photon candidate. These will henceforth be referred to as proxy photons.

6.4 Diphoton Mass Fit

The π^0 yields are extracted using binned maximum likelihood fits to the diphoton mass spectra. In order to improve cancellation of systematic uncertainties associated with the choice of fit model, the same fitting procedure is used to extract π^0 yields in the data and reference samples. The same fit model is used to extract the yields used in both the $R_{p\text{Pb}}^{\pi^0}$ and R^γ measurements.

The converted photons used to reconstruct the π^0 candidate may lose energy via Bremsstrahlung radiation, and ECAL photon energies may be overestimated due to electronic noise in the ECAL. As a result, the π^0 peak has a non-Gaussian shape. The π^0 signal peak is modelled using a two-sided Crystal Ball function [178], which consists of a Gaussian core with power-law tails. The tail parameters are fixed using fits to the mass spectra of simulated π^0 candidates. Example signal fits in simulation are shown in Figure 6-2. The mean and width of the Gaussian core are left free in the fit.

The combinatoric background is modeled using proxy photons. Proxy photons are combined with reconstructed ECAL photons. The mass spectra of the resulting proxy diphotons provide an accurate description of the combinatoric background in simulation. Example comparisons of the proxy diphoton background and combinatoric background in simulation are shown in Figure 6-3. An additional background compo-

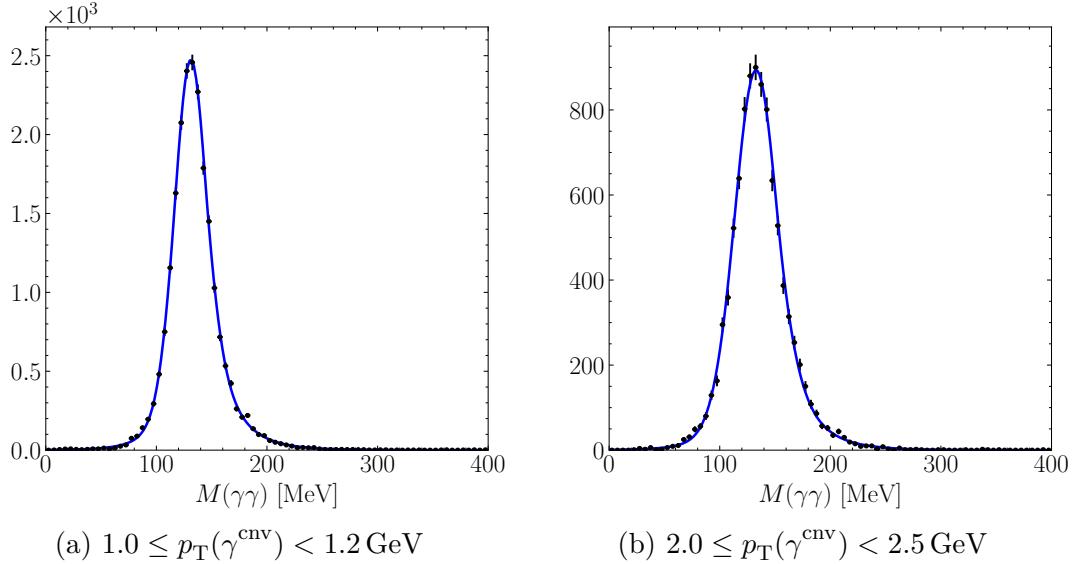


Figure 6-2: Example fits to the π^0 signal in 8.16 TeV $p\text{Pb}$ simulation.

ment arises from converted photons combined with bremsstrahlung photons radiated by the conversion electrons, resulting in a low-mass peak in the diphoton mass spectrum. The bremsstrahlung background is modelled by convolving the reconstructed converted photon mass distribution with the sum of two half-Gaussian distributions. One of the half-Gaussian widths is fixed using fits to background in simulation while the other is left free. Both background components are multiplied by a phase space factor that insures that the PDF vanishes at $m = 0$. The phase space factor is given by

$$f_{\text{PHSP}}(m, m_0) = \begin{cases} \sqrt{1 - \frac{4m_0^2}{m^2}}, & m > 2m_0 \\ 0, & m \leq 2m_0 \end{cases} \quad (6.9)$$

where m is the diphoton mass and m_0 is a fit parameter. Figure 6-4 demonstrates that the total background fit model provides a good description of backgrounds in simulation.

6.5 $R_{p\text{Pb}}^{\pi^0}$ Measurement

6.5.1 π^0 Yields

Neutral pion yields are extracted using the fit model described in Section 6.4. Pions reconstructed with converted photons with $p_T > 1 \text{ GeV}$ and ECAL photons with $p_T > 400 \text{ MeV}$ are considered. The fit is performed in bins of the $\pi^0 p_T$.

In order to measure a nuclear modification factor, the π^0 yields must be measured in proton-proton and proton-lead collisions in the same rapidity window in the center of mass frame. The rapidity shift in $p\text{Pb}$ and $\text{Pb}p$ collisions means that the resulting rapidity windows must be smaller than the full fiducial region and are different for

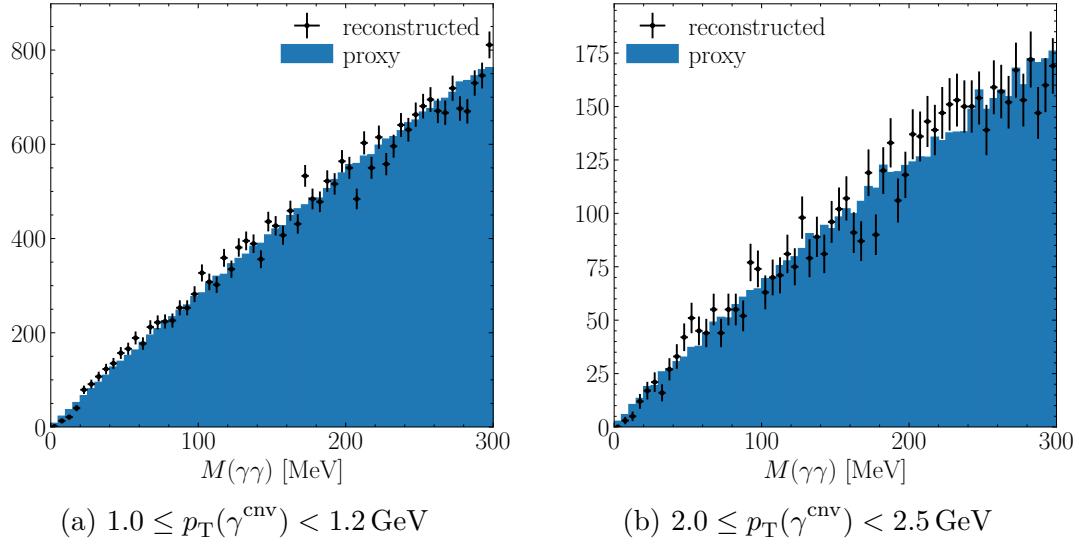


Figure 6-3: Example comparisons between combinatoric background in 8.16 TeV $p\text{Pb}$ simulation and the proxy diphoton background model.

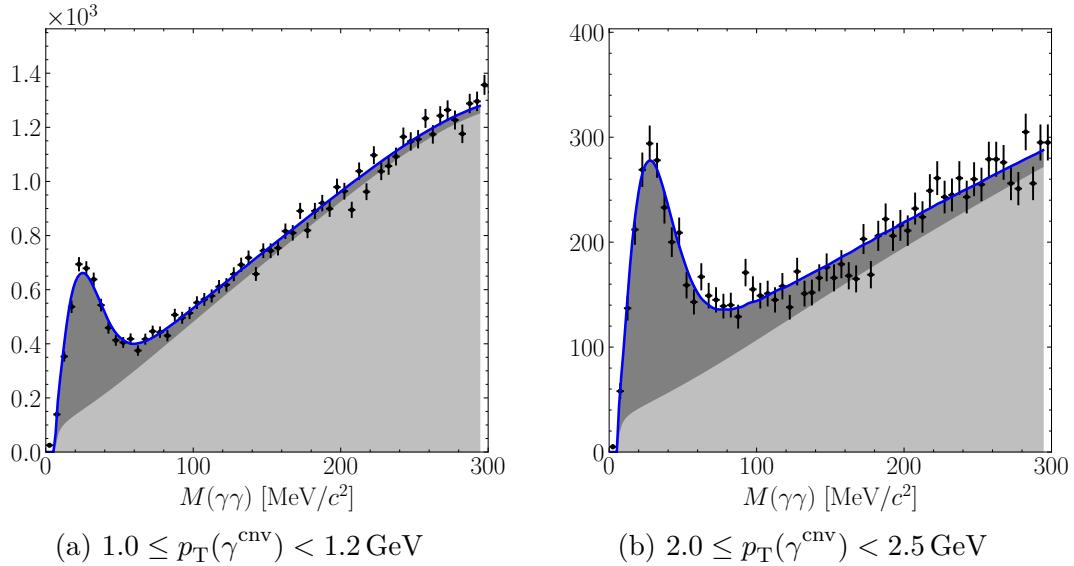


Figure 6-4: Example fits background to 8.16 TeV $p\text{Pb}$ simulation. The combinatoric background is shown in light gray, while the bremsstrahlung background is shown in dark gray. The blue line shows the full background fit.

beam configuration	rapidity
$p\text{Pb}$	$2.5 < y^* < 3.5$
$\text{Pb}p$	$-4.0 < y^* < -3.0$

Table 6.3: Center-of-mass rapidity regions used for measuring $R_{\text{pPb}}^{\pi^0}$.

$p\text{Pb}$ and $\text{Pb}p$ configurations. Yields are measured in the rapidity windows shown in Table 6.3. The pp π^0 yields are measured in both of these rapidity regions.

6.5.2 Efficiency Corrections

The measured π^0 yields are corrected for detector inefficiencies using efficiency maps constructed from simulation. Due to the relatively small size of the pp simulation sample, both pp data samples are corrected using efficiency maps from $p\text{Pb}$ simulation. The $p\text{Pb}$ and $\text{Pb}p$ samples are corrected using the corresponding simulation sets. Efficiency maps are shown in Figure 6-7.

The pp , $p\text{Pb}$ and $\text{Pb}p$ data samples will have different detector occupancies and, as a consequence, different π^0 reconstruction efficiencies. The π^0 reconstruction efficiency is the product of the ECAL and converted photon reconstruction efficiencies. The LHCb ECAL photon reconstruction efficiency ϵ^{cal} has been studied using $B^+ \rightarrow J/\psi K^{*+} (\rightarrow K^+ \pi^0)$ decays, indicating agreement between data and simulation to within about 4% [179]. However, this only holds for samples with similar detector occupancies. The efficiency is corrected for occupancy differences by calculating the reconstruction efficiency as a function of the number of hits in the SPD in the event (N_{SPD}), which provides an estimate of the total number of charged particles produced in LHCb acceptance. The average ECAL efficiency for each dataset is then calculated using the N_{SPD} distribution from data and the efficiency map from simulation.

The converted photon reconstruction efficiency is studied using the ratio

$$\frac{\epsilon^{\text{cav}}}{\epsilon^{\text{cal}}} = \frac{\pi^0 \rightarrow \gamma^{\text{cav}} \gamma^{\text{cav}}}{\pi^0 \rightarrow \gamma^{\text{cav}} \gamma^{\text{cal}}}, \quad (6.10)$$

where ϵ^{cav} and ϵ^{cal} are the converted and ECAL photon efficiencies, respectively. Data and simulation can be compared using the double ratio R^ϵ given by

$$R^\epsilon = \left(\frac{\epsilon^{\text{cav}}}{\epsilon^{\text{cal}}} \right)_{\text{sim}} \Big/ \left(\frac{\epsilon^{\text{cav}}}{\epsilon^{\text{cal}}} \right)_{\text{data}} \quad (6.11)$$

The measured π^0 yields are then corrected with factors given by

$$\frac{\epsilon_{\text{sim}}^{\pi^0}}{\epsilon_{\text{data}}^{\pi^0}} = R^\epsilon \left(\frac{\epsilon_{\text{sim}}^{\text{cal}}}{\epsilon_{\text{data}}^{\text{cal}}} \right)^2. \quad (6.12)$$

The ECAL and converted photon efficiency corrections are summarized in Table 6.4.

Yields are extracted using a similar fit model to that presented in Section 6.4.

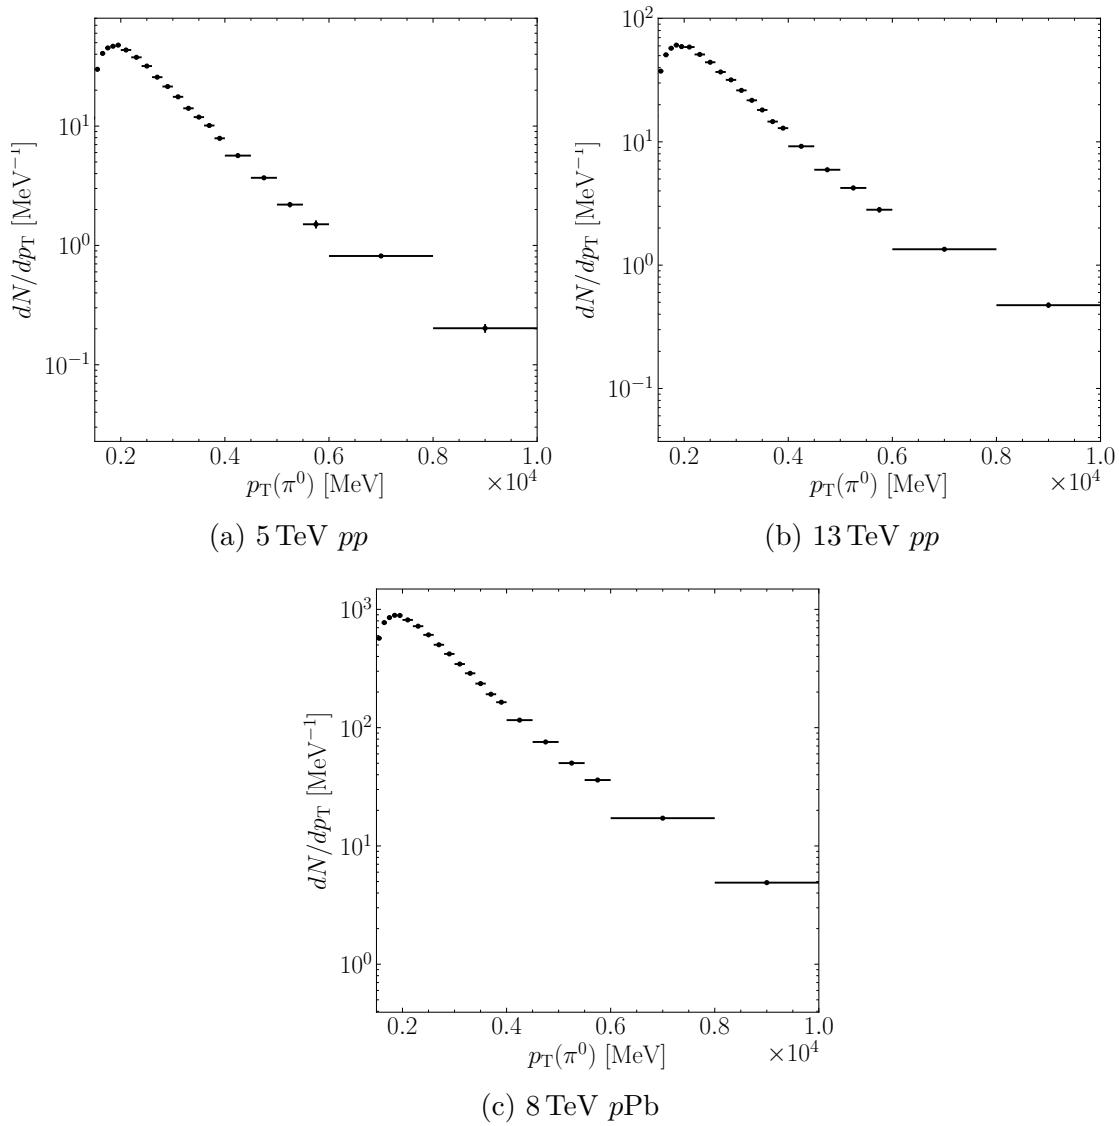


Figure 6-5: Uncorrected π^0 yields in the forward rapidity window.

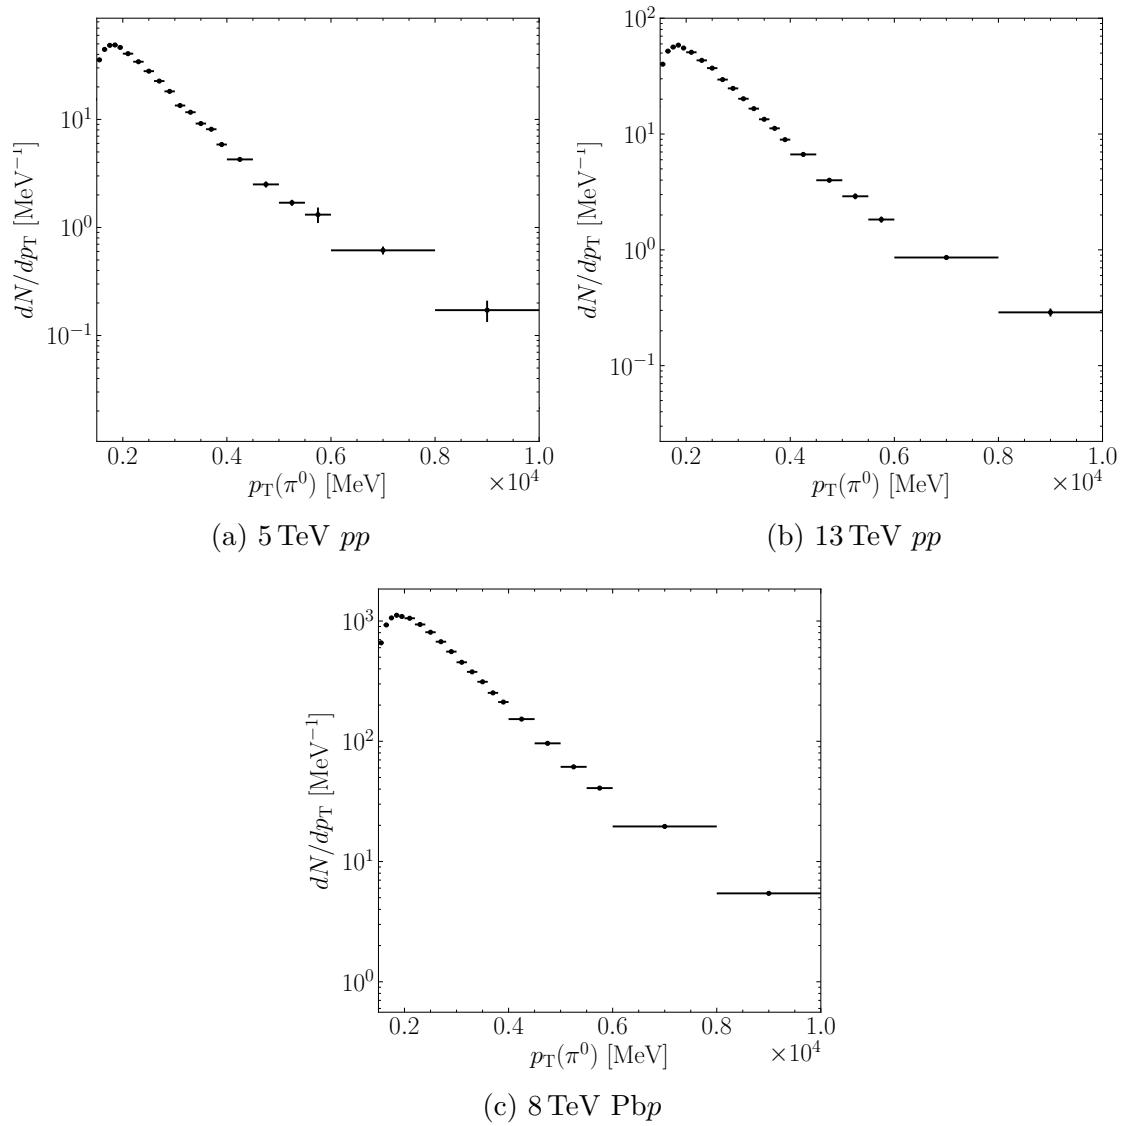
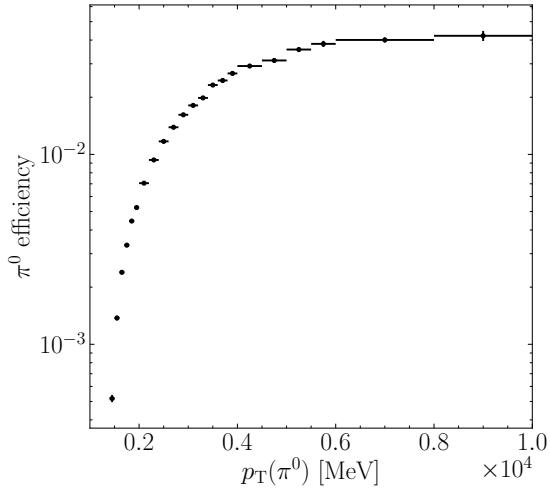
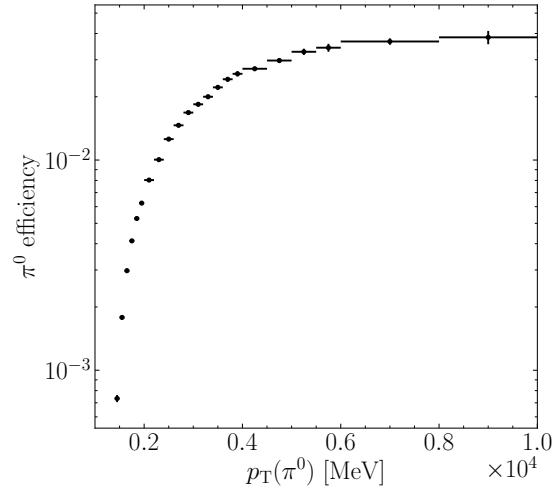


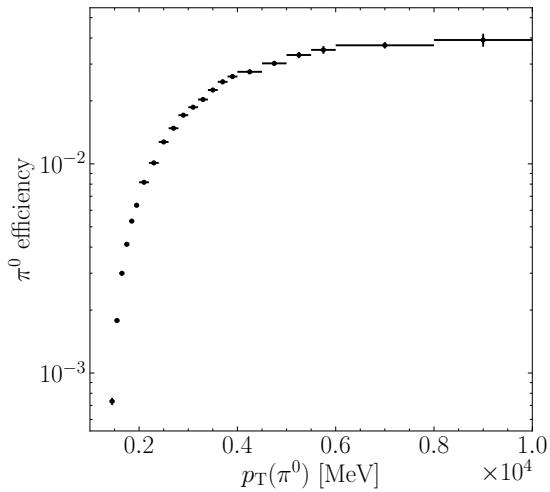
Figure 6-6: Uncorrected π^0 yields in the backward rapidity window.



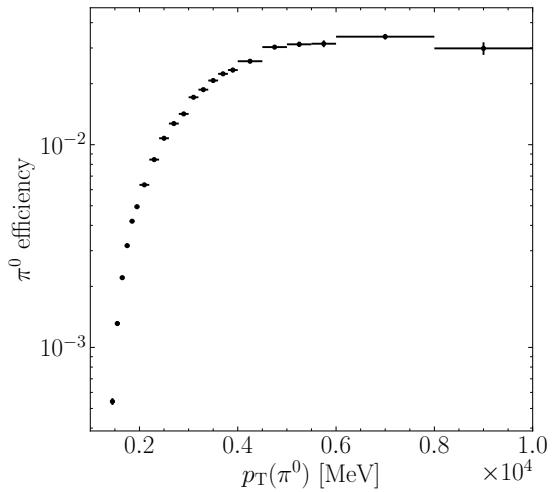
(a) pp , Forward



(b) pp , Backward



(c) pPb



(d) Pbp

Figure 6-7: Efficiency maps used to correct measured π^0 yields. The maps shown in a and b are used to correct both 13 TeV and 5.02 TeV pp measurements.

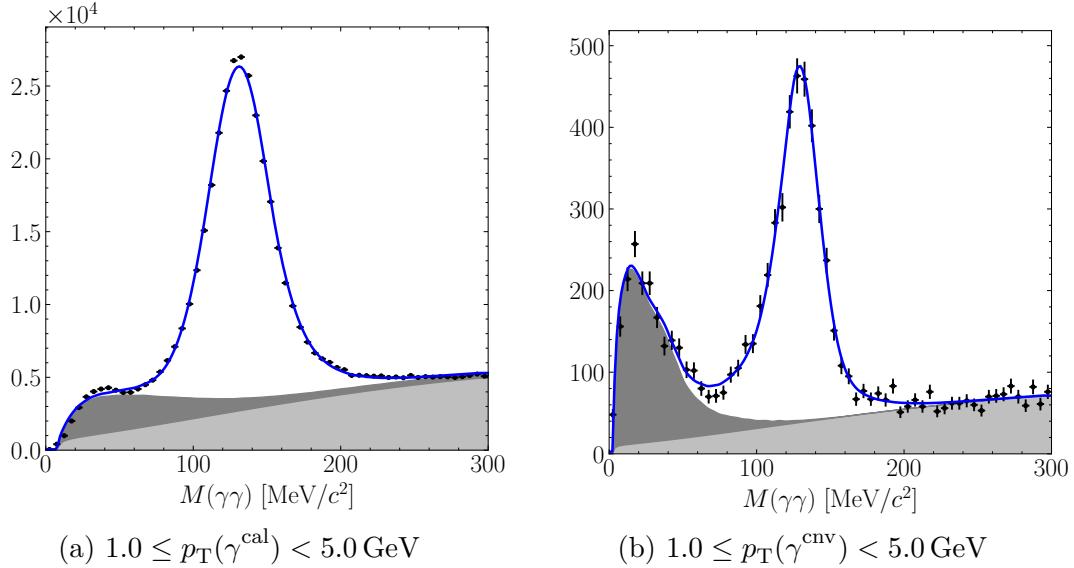


Figure 6-8: Fits to diphoton mass spectra for $\gamma^{\text{cnv}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cnv}}\gamma^{\text{cnv}}$ candidates (b) in 13 TeV pp data.

	$\epsilon_{\text{sim}}^{\text{cal}}/\epsilon_{\text{data}}^{\text{cal}}$	R^{ϵ}	$\epsilon_{\text{sim}}^{\pi^0}/\epsilon_{\text{data}}^{\pi^0}$
13 TeV pp	0.995 ± 0.015	1.076 ± 0.072	1.065 ± 0.075
5.02 TeV pp	0.977 ± 0.014	0.701 ± 0.054	0.669 ± 0.040
8.16 TeV $p\text{Pb}$	1.044 ± 0.016	0.926 ± 0.058	1.009 ± 0.057
8.16 TeV $\text{Pb}p$	1.065 ± 0.016	0.881 ± 0.053	0.999 ± 0.064

Table 6.4: Efficiency corrections to the measured π^0 yields.

The signal tail parameters are fixed using fits to $\pi^0 \rightarrow \gamma^{\text{cnv}}\gamma^{\text{cal}}$ and $\pi^0 \rightarrow \gamma^{\text{cnv}}\gamma^{\text{cnv}}$ decays in simulation. The bremsstrahlung peak in the $\gamma^{\text{cnv}}\gamma^{\text{cnv}}$ is modelled using pairs of converted photons that share a track. The resulting fits are shown in Figures 6-8, 6-9, 6-10, and 6-11. The measured correction factors are shown in Table 6.4. The corrections indicate that the π^0 reconstruction efficiency in 5.02 TeV pp data is about 50% larger than that in 8.16 TeV $p\text{Pb}$ simulation due to the smaller detector occupancy in 5.02 TeV pp collisions. On the other hand, simulation overestimates the converted photon efficiency in 13 TeV pp . The 13 TeV pp data uses an ECAL particle identification configuration developed for Run 1 8 TeV pp data taking. The ECAL PID configuration was updated before the 5.02 TeV and proton-lead runs, resulting in improved electron reconstruction performance.

6.5.3 Luminosity Calibration

The pp data samples consist of collisions between leading bunches. The selection of leading bunch crossings biases the dataset, requiring an independent luminosity calibration. The 5.02 TeV pp leading bunch sample has been measured with a precision of about 2.5%. No such calibration exists for the 13 TeV pp leading bunch data, so

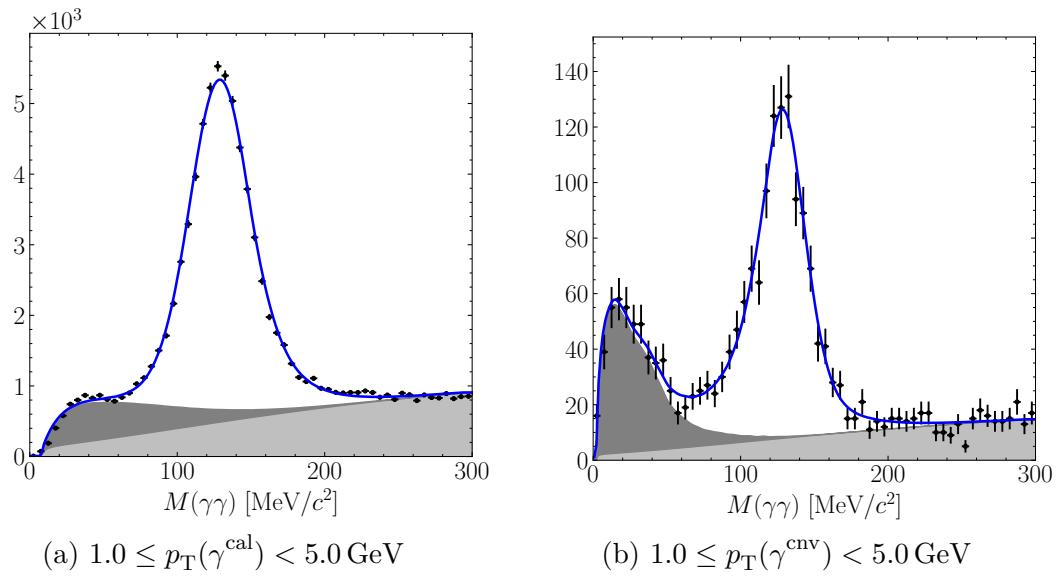


Figure 6-9: Fits to diphoton mass spectra for $\gamma^{\text{cnv}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cnv}}\gamma^{\text{cnv}}$ candidates (b) in 5.02 TeV pp data.

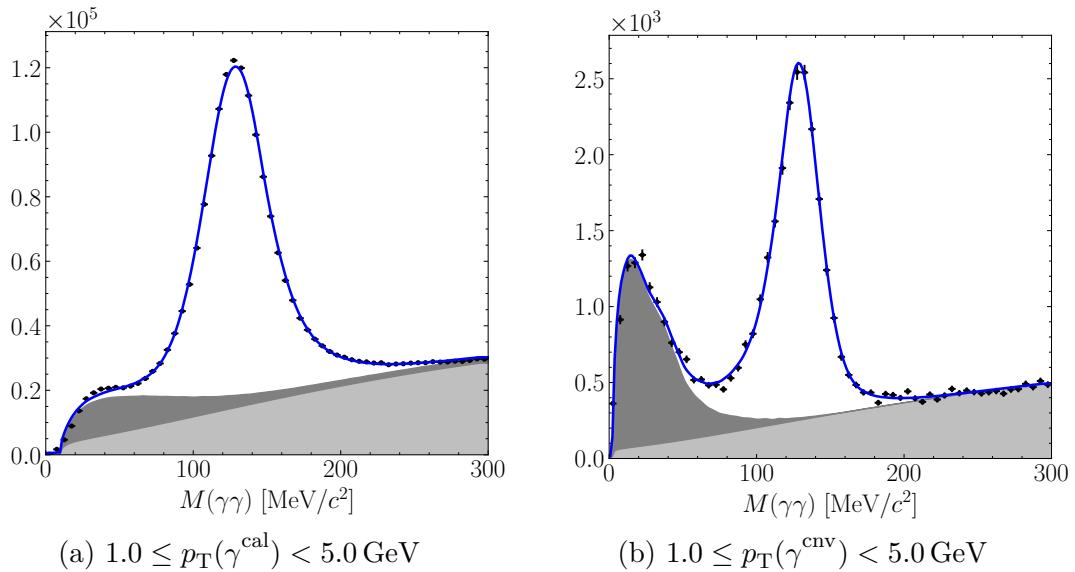


Figure 6-10: Fits to diphoton mass spectra for $\gamma^{\text{cnv}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cnv}}\gamma^{\text{cnv}}$ candidates (b) in 8.16 TeV $p\text{Pb}$ data.

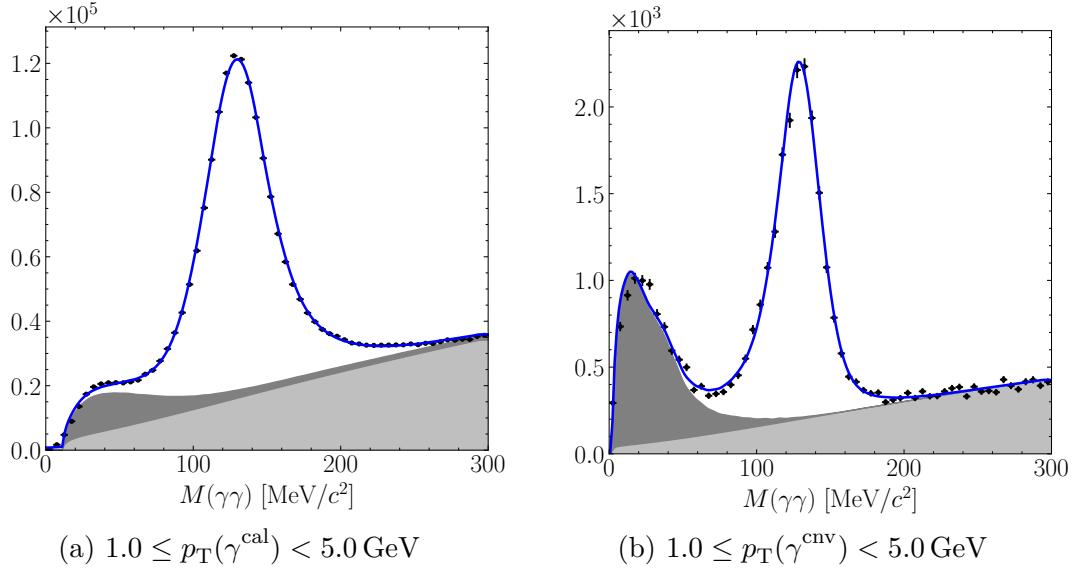


Figure 6-11: Fits to diphoton mass spectra for $\gamma^{\text{cav}}\gamma^{\text{cal}}$ candidates (a) and $\gamma^{\text{cav}}\gamma^{\text{cav}}$ candidates (b) in 8.16 TeV $\text{Pb}p$ data.

an independent luminosity calibration is performed.

The method used here closely follows the luminosity calibration presented in Ref. [180]. The integrated luminosity of a no bias sample is given by

$$\mathcal{L} = \frac{\mu N_{\text{BX}}}{\sigma_{\text{vis}}^{pp}}, \quad (6.13)$$

where μ is the pile-up or interaction rate, N_{BX} is the number of bunch crossings in the data sample, and σ_{vis}^{pp} is the visible pp inelastic cross section. The number of pp interactions per bunch crossing n is Poisson-distributed according to the probability mass function

$$P(n, \mu) = \frac{\mu^n e^{-\mu}}{n!}, \quad (6.14)$$

where μ is the interaction rate. The probability of a bunch crossing producing no interactions is then given by

$$P_0 \equiv P(0, \mu) = e^{-\mu}. \quad (6.15)$$

From Eq. 6.15, the interaction rate μ is given by

$$\mu = -\ln P_0. \quad (6.16)$$

P_0 is calculated by counting the number of empty events in the data sample. An empty event is defined as an event in which a given luminosity counter has a value of zero. Luminosity counters include the number of VELO tracks, number of PVs, and number of SPD hits in the event. Each luminosity counter requires a different value of σ_{vis}^{pp} . Here the number of PVs is used to calculate P_0 . The interaction rate is then

dataset	μ_{MB}	$\sigma_{\text{vis}}^{pp,p\text{Pb}}$ [mb]	\mathcal{L} [nb $^{-1}$]
5.02 TeV pp	-	-	3.82 ± 0.08
13 TeV pp	0.901 ± 0.005	57.85 ± 2.26	9.77 ± 0.38
8.16 TeV $p\text{Pb}$	1.025 ± 0.010	2220 ± 50	0.426 ± 0.012
8.16 TeV Pbp	1.025 ± 0.010	2125 ± 53	0.534 ± 0.014

Table 6.5: Luminosities for each sample used to calculate $R_{p\text{Pb}}^{\pi^0}$. The reported μ value for the $p\text{Pb}$ and Pbp samples is the visible minimum bias interaction rate μ_{MB} .

given by

$$\mu = -\ln \left(1 - \frac{N_{\text{vis}} - N_{\text{bkg}}}{N_{\text{BX}}} \right), \quad (6.17)$$

where N_{vis} is the number of visible events and N_{bkg} is the number of background events. N_{bkg} is measured using bunch crossings in which one beam is empty. The background contribution is found to be much less than 1% and much smaller than the uncertainty associated with the σ_{vis}^{pp} measurement.

The visible cross section σ_{vis}^{pp} is measured in precision luminosity studies [105]. LHCb has performed luminosity measurements in both 13 TeV and 5 TeV pp collisions using both the Van der Meer scan method and beam gas imaging using the SMOG system. These methods have been used to measure σ_{vis}^{pp} to a precision of about 4%.

Similar luminosity calibrations are performed using the $p\text{Pb}$ and Pbp data samples. Both proton-lead datasets consist of minimum bias events that require at least one VELO track. As a result, the zero counting method described in this section cannot be used directly. Instead, the interaction rate is calculated using a separate sample of luminosity triggers. The interaction rate is measured with relatively poor precision. However, because the interaction rate is small, this leads to sub-percent level luminosity uncertainties. Because no empty events are included in the proton-lead data samples, the integrated luminosity is calculated using only visible events. The number of interactions per visible event is given by

$$\mu_{\text{MB}} = \frac{\sum_{n=1}^{\infty} n P(n, \mu)}{\sum_{n=1}^{\infty} P(n, \mu)}. \quad (6.18)$$

The integrated luminosity of the studied proton-lead samples is then given by

$$\mathcal{L} = \frac{\mu_{\text{MB}} N_{\text{vis}}}{\sigma_{\text{vis}}^{p\text{Pb}}}. \quad (6.19)$$

In practice, μ_{MB} can be calculated to the needed precision using only $n \leq 3$. As for pp collisions, $\sigma_{\text{vis}}^{p\text{Pb}}$ was measured using the Van der Meer scan method. The measured integrated luminosities and associated uncertainties are shown in Table 6.5.

	efficiency	luminosity	fit model	interpolation
forward	8.3%	3.1%	0.1 – 0.4%	1.5 – 3.0%
backward	8.9%	3.1%	0.1 – 0.4%	1.5 – 3.0%

Table 6.6: Summary of the systematic uncertainties in the $p\text{Pb}$ (forward) and Pbp (backward) directions as explained in the text.

6.5.4 Systematic Uncertainties

The $R_{p\text{Pb}}^{\pi^0}$ systematic uncertainties are summarized in Table 6.6. The individual sources of uncertainty are explained below.

Efficiency Corrections

The dominant source of systematic uncertainty is the efficiency correction measurement. The uncertainties associated with the correction factors $\epsilon_{\text{sim}}^{\pi^0}/\epsilon_{\text{data}}^{\pi^0}$ are given in Table 6.4. These are taken as global systematic uncertainties. In addition, a systematic uncertainty of 4% is assigned to the ECAL efficiency. The combination of the efficiency correction uncertainty and ECAL uncertainty make up the total “efficiency” uncertainty given in Table 6.6.

Luminosity Calibration

The uncertainty in the luminosity calibration is dominated by the uncertainty in the measured visible cross sections, which is about 2.0%, 3.9%, 2.2%, and 2.2% in 5.02 TeV pp , 13 TeV pp , $p\text{Pb}$, and Pbp , respectively. In 13 TeV pp , an additional uncertainty of 0.5% is included to account for background contributions in the calculation of N_{vis} . In proton-lead, an additional uncertainty of 1% is included to account for the roughly 40% uncertainty in μ .

Fit Model

The systematic uncertainty associated with the π^0 fit model is determined by varying the fixed fit parameters by their uncertainties given by fits to simulation. This includes the tail parameters of the signal peak and bremsstrahlung background. This causes the π^0 yield to vary by about 3 – 4%. As a cross-check, the fit error is checked in fits to simulation, where the error is given by the difference between the measured and true yields. These provide consistent estimates of the systematic uncertainty of the π^0 yields. The yield uncertainty largely cancels in $R_{p\text{Pb}}^{\pi^0}$, resulting in uncertainties much smaller than 1%.

Reference Interpolation

No large 8 TeV pp minimum or no bias data sample is available, so the pp reference is determined by interpolating the π^0 cross section between 5.02 and 13 TeV. The interpolation using both a linear function and a power law. Interpolation results are

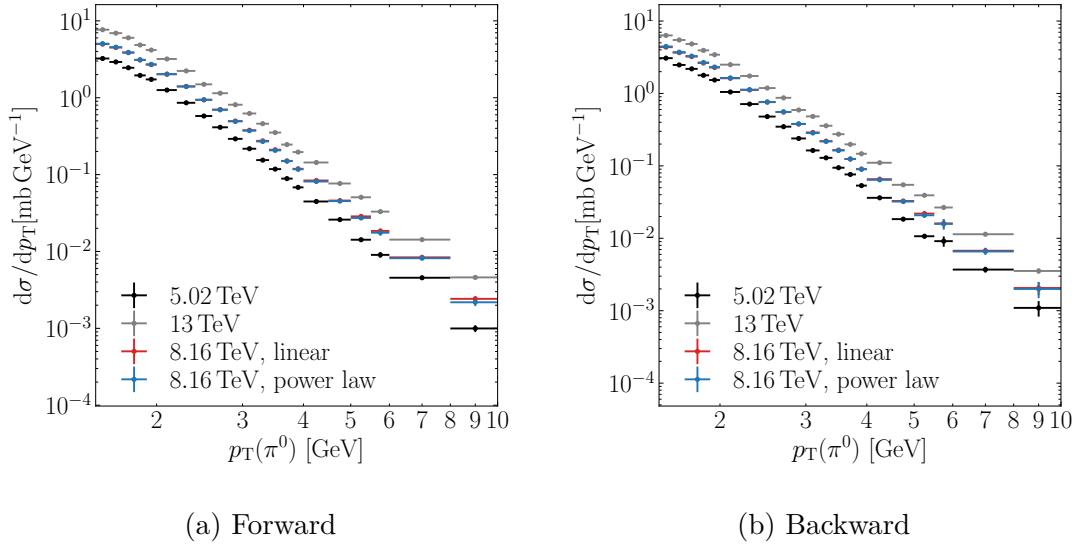


Figure 6-12: Measured 5.02 and 13 TeV $pp \pi^0$ production cross sections and interpolation results.

shown in Figure 6-12. The results differ by 1.5 – 3.0% in both the forward and backward directions. The nominal result is calculated using the power law interpolation and the difference between the two interpolations is taken as a systematic uncertainty.

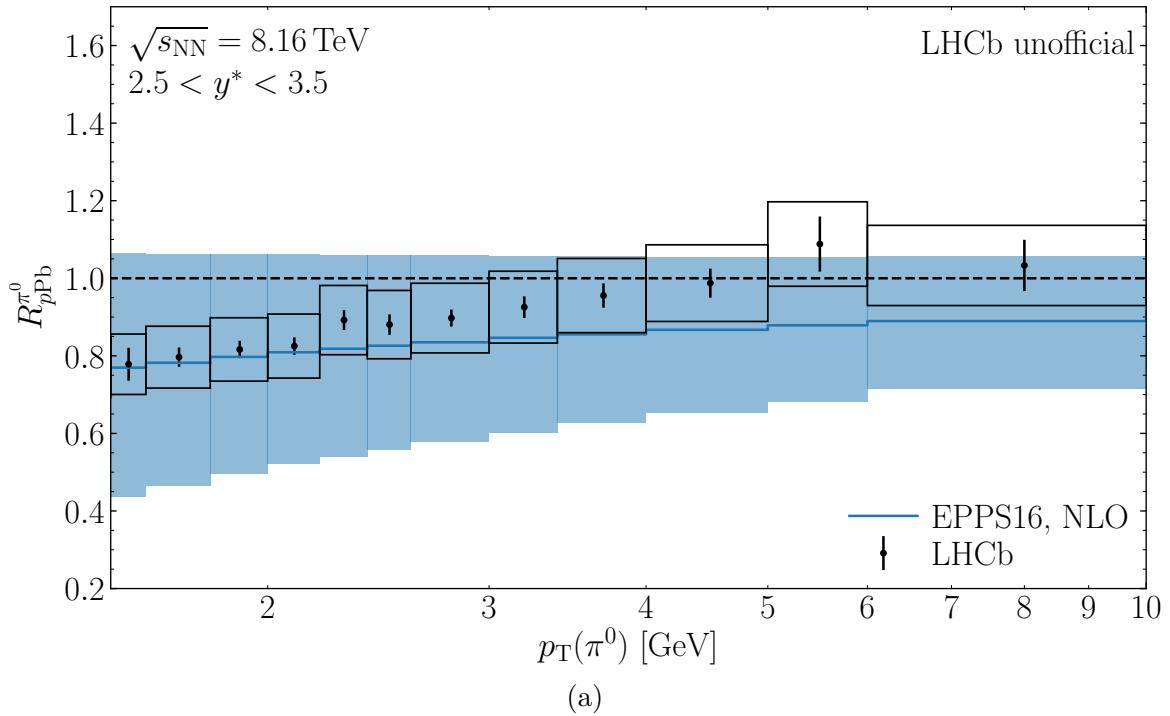
6.5.5 Results

The measured nuclear modification factors are shown in Figure 6-13. Also shown are pQCD predictions calculated at NLO [58]. The predictions are calculated using the CT14 proton PDF set, EPPS16 nPDF set [47], and DSS14 parton-to-pion fragmentation functions [181]. The forward measurement shows suppression relative to the pp reference, which is qualitatively consistent with nuclear shadowing at low x . The forward measurement agrees well with the EPPS16 prediction, although the nPDF uncertainties are large. The backward measurement demonstrates a large enhancement relative to the pp reference at $p_T > 2$ GeV. The enhancement is qualitatively consistent with anti-shadowing, but much larger than the EPPS16 prediction.

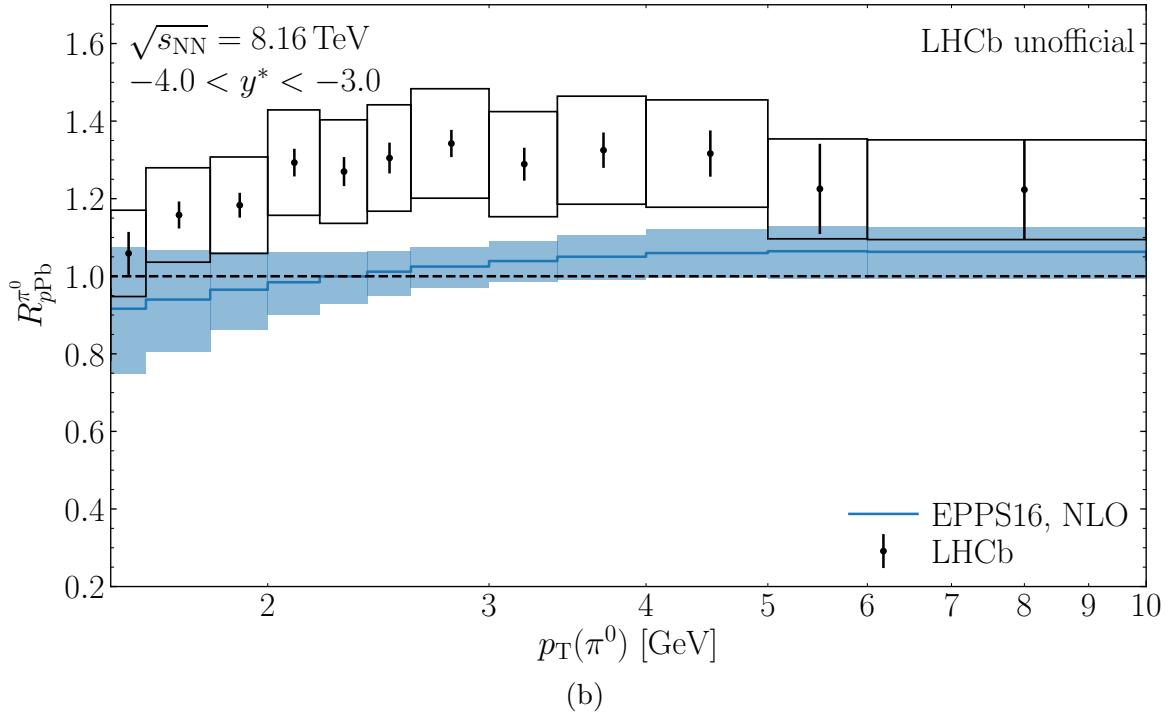
6.6 R^γ Measurement

6.6.1 Photon and π^0 Yields

The photon selections used in the R^γ measurement are identical to those used in the $R_{p\text{Pb}}^{\pi^0}$ measurement. Only DD photons in events with a single PV are considered. Inclusive photon yields are shown in Figure 6-14. Neutral pion yields are extracted in bins of $p_T(\gamma^{\text{cnv}})$. The resulting π^0 yields are shown in Figure 6-15.



(a)



(b)

Figure 6-13: Neutral pion nuclear modification factors measured in $p\text{Pb}$ (a) and $\text{Pb}p$ (b) collisions at $\sqrt{s_{\text{NN}}} = 8.16 \text{ TeV}$ with the LHCb detector. The statistical uncertainties are illustrated using error bars and the systematic uncertainties are shown as boxes. NLO pQCD calculations are shown in blue. The shaded regions show the nPDF uncertainties.

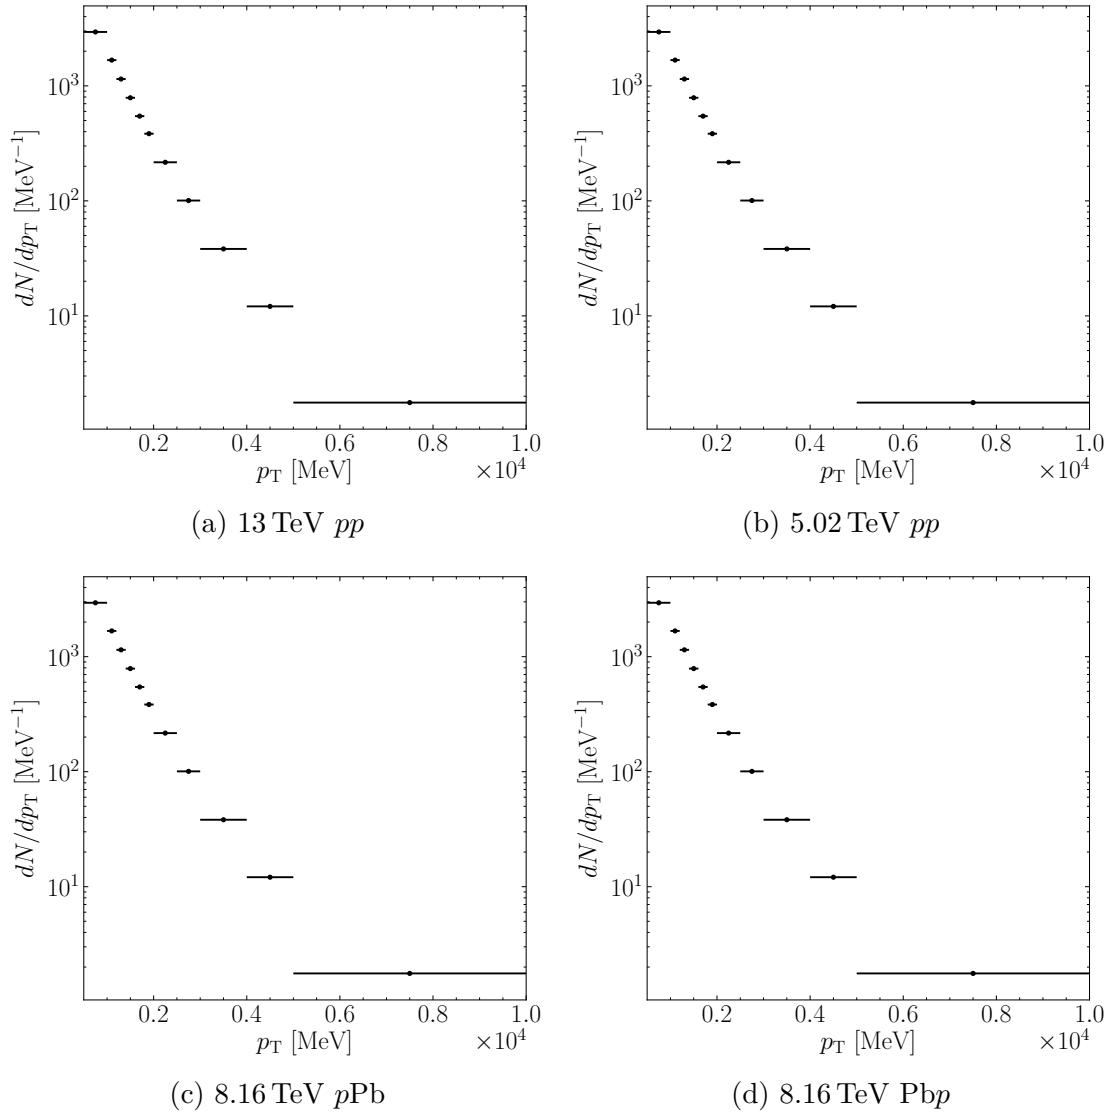


Figure 6-14: Uncorrected inclusive converted photon yields used to measure R^γ .

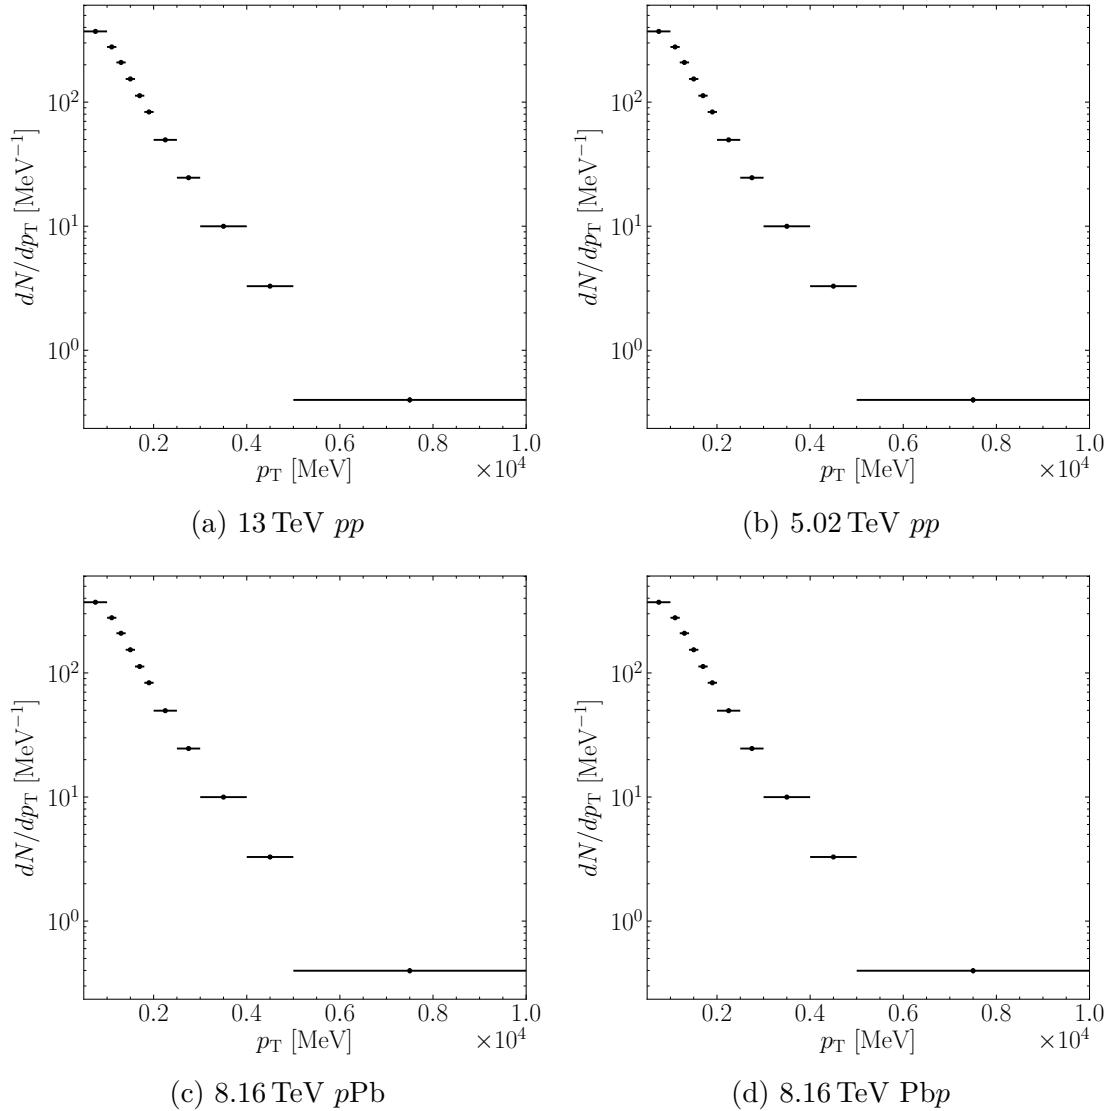


Figure 6-15: Uncorrected π^0 yields used to measure R^γ .

6.6.2 Efficiency and Acceptance Corrections

The converted photon efficiency in both the data and reference samples cancels in R^γ , but the ECAL photon efficiencies may differ between the data and reference samples. The π^0 yields are corrected to account for ECAL inefficiencies. The efficiency ϵ^{tot} is defined as

$$\epsilon^{\text{tot}} = \mathcal{P}(\pi^0|\gamma^{\text{cnv}}), \quad (6.20)$$

where $\mathcal{P}(\pi^0|\gamma^{\text{cnv}})$ is the probability of reconstructing a π^0 given that one of the π^0 decay photons is reconstructed as a converted photon. The π^0 photon that is not reconstructed as a converted photon is referred to as the spectator photon.

The total efficiency consists of two factors, ϵ^{acc} and ϵ^{rec} . The acceptance efficiency ϵ^{acc} is the probability that the spectator photon will pass the fiducial requirements of $p_T > 400$ MeV and $2.0 < \eta < 4.5$. The reconstruction efficiency ϵ^{rec} is the probability that a spectator photon passing the ECAL fiducial selection will be reconstructed as an ECAL cluster. Both factors are determined as a function of the converted photon transverse momentum.

For π^0 decays producing a converted photon with $2.5 < \eta < 4.0$ and $p_T > 500$ MeV, all spectator photons passing the momentum requirement of $p_T > 400$ MeV will satisfy the requirement $2.0 < \eta < 4.5$. As a result, ϵ^{acc} is the probability that the spectator photon will have $p_T > 400$ MeV. The momentum distribution of the spectator photons is determined by the π^0 p_T spectrum. The π^0 p_T spectra are measured in both data and simulation. The π^0 yields are extracted as a function of the π^0 p_T , and the yields are corrected for detector inefficiencies using the efficiency maps from Section 6.5. The measured π^0 spectra are then fit with a Hagedorn function [182] given by

$$\frac{dN}{dp_T} = A \times p_T \times \left(\exp(ap_T + bp_T^2) + \frac{p_T}{p_0} \right)^{-n}, \quad (6.21)$$

where A , a , b , p_0 and n are free parameters in the fit. The resulting π^0 p_T spectra and fit results are shown in Figure 6-16 for data and Figure 6-17 for simulation. The Hagedorn fits are used to generate π^0 candidates, which then decay to two photons. One of the simulated decay photons is chosen as the converted photon candidate, while the other is chosen as the spectator. Both photon momenta are smeared by the momentum resolution obtained from simulation. The smeared momenta are used to calculate the π^0 acceptance as a function of $p_T(\gamma^{\text{cnv}})$. The measured acceptances from data are shown in Figure 6-18, and those from simulation are shown in Figure 6-19.

The ECAL reconstruction efficiency ϵ^{rec} is determined from simulation. The ECAL efficiencies are similar in each simulation sample, so ϵ^{rec} is determined using 8.16 TeV $p\text{Pb}$ and $\text{Pb}p$ simulation. The resulting total ϵ^{rec} and ϵ^{tot} are shown in Figure 6-20 for data and in Figure 6-21 for simulation. The ECAL efficiency is further corrected for occupancy differences using the correction factors given in Table 6.4.

The acceptance determination is validated by comparing the measured and true acceptances in simulation. The measured π^0 acceptance is first calculated without momentum smearing and compared to the true acceptance in simulation. The ratios of measured to true acceptance in simulation are shown in Figure 6-22. The mea-

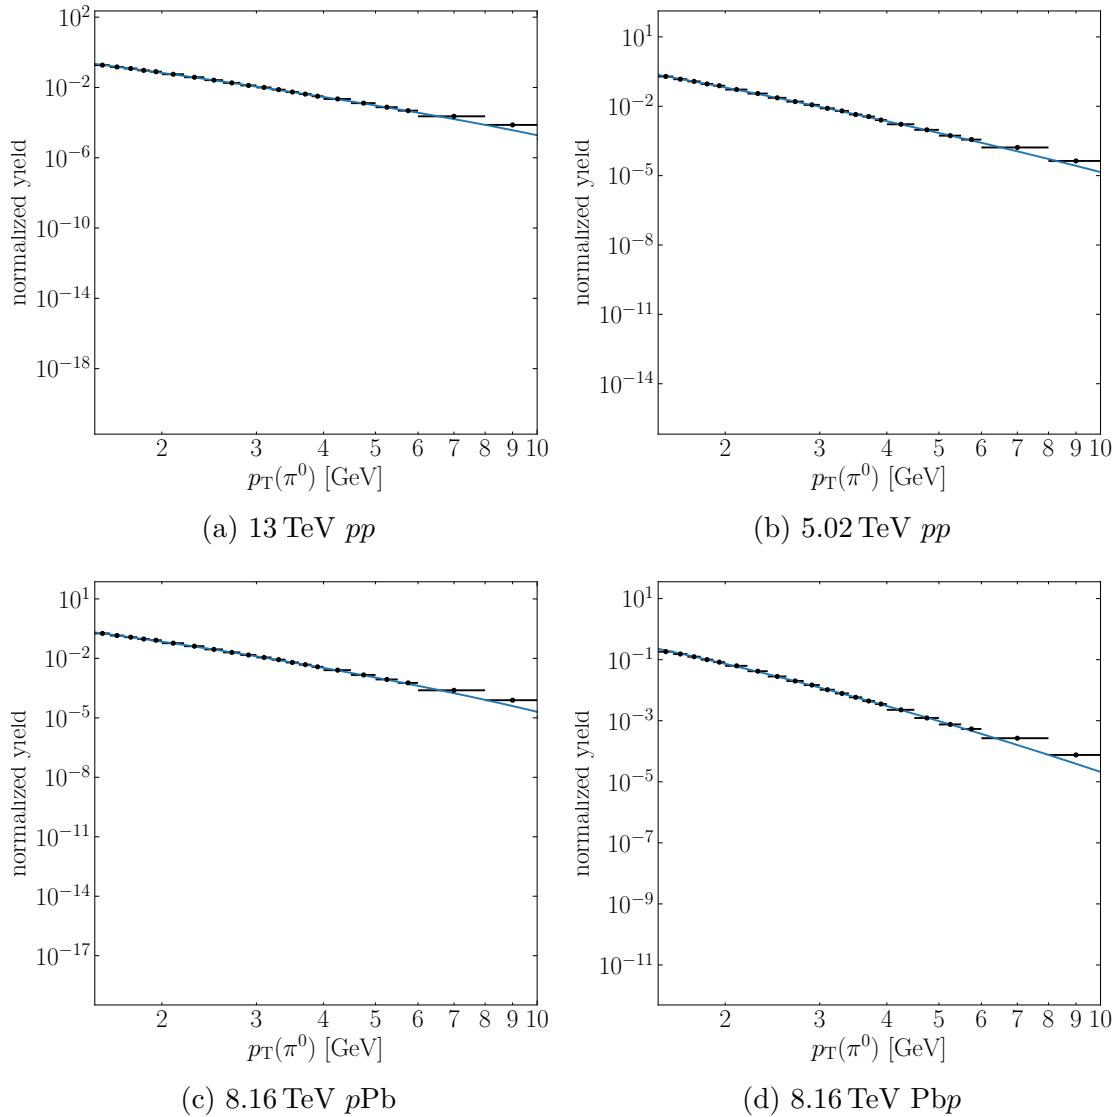


Figure 6-16: Efficiency corrected π^0 yields from data. Hagedorn fit results are shown as blue curves.

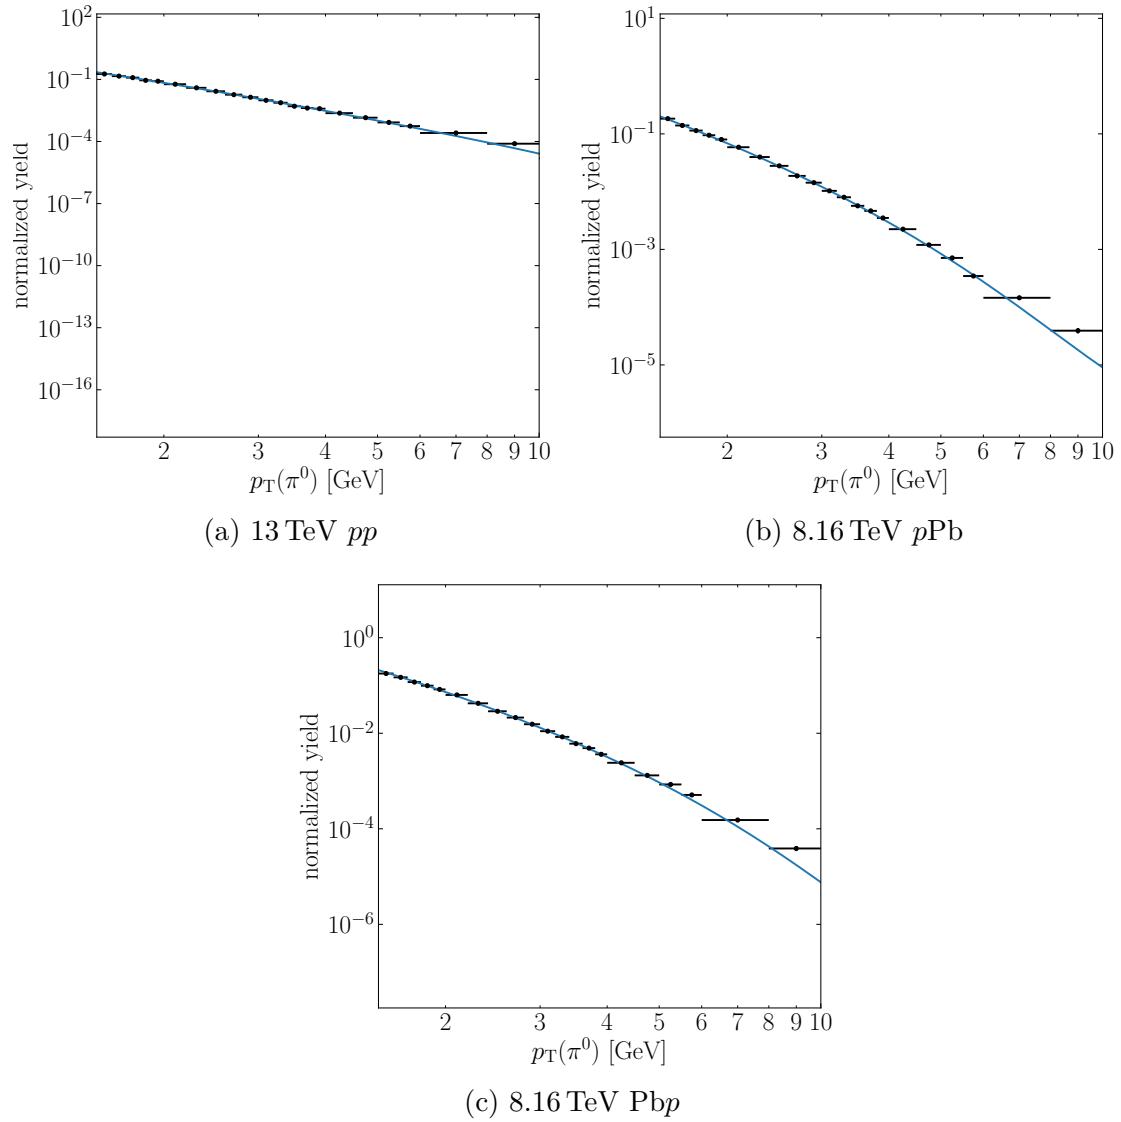


Figure 6-17: Efficiency corrected π^0 yields from simulation. Hagedorn fit results are shown as blue curves.

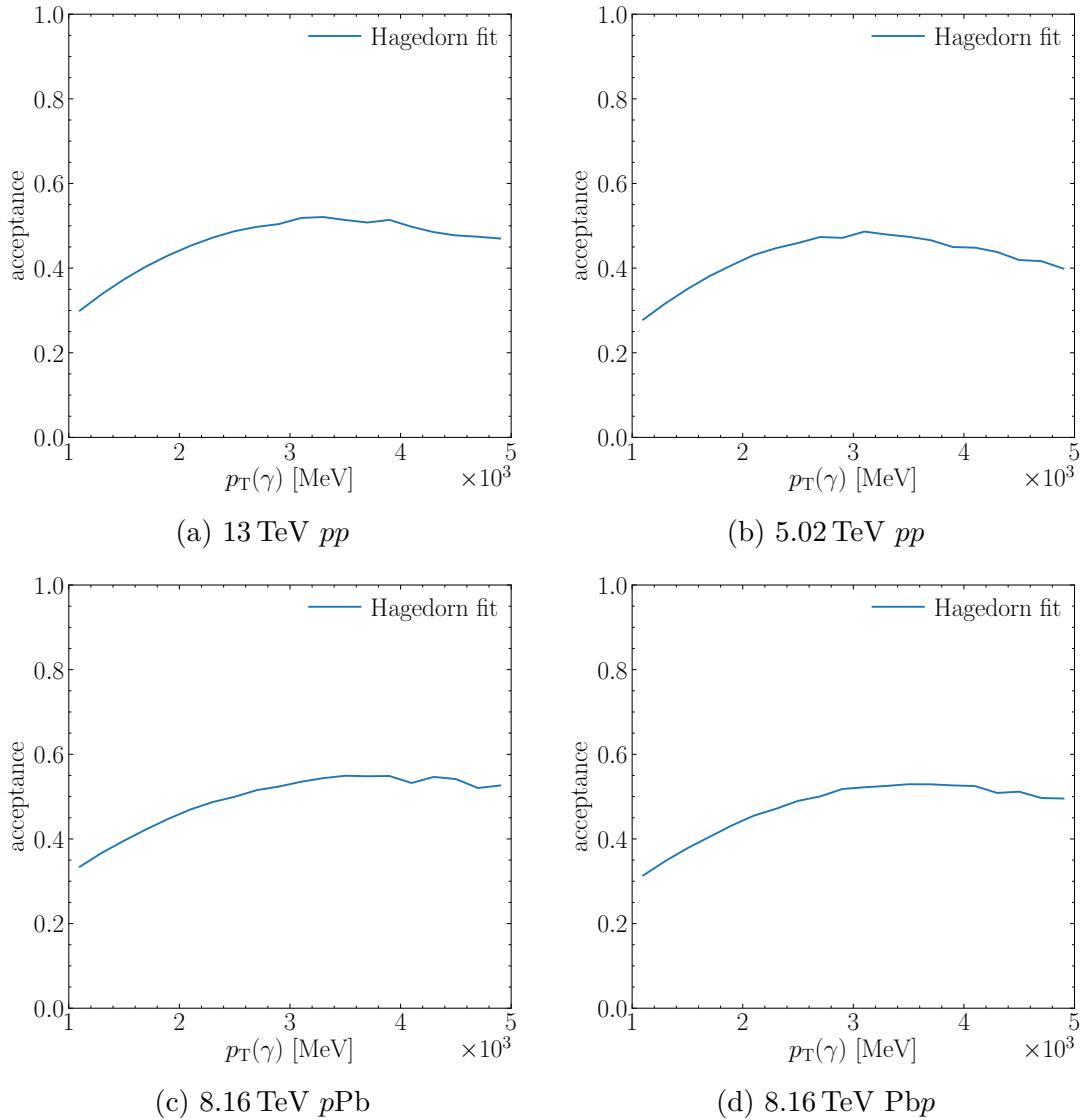


Figure 6-18: Calculated π^0 acceptances from data.

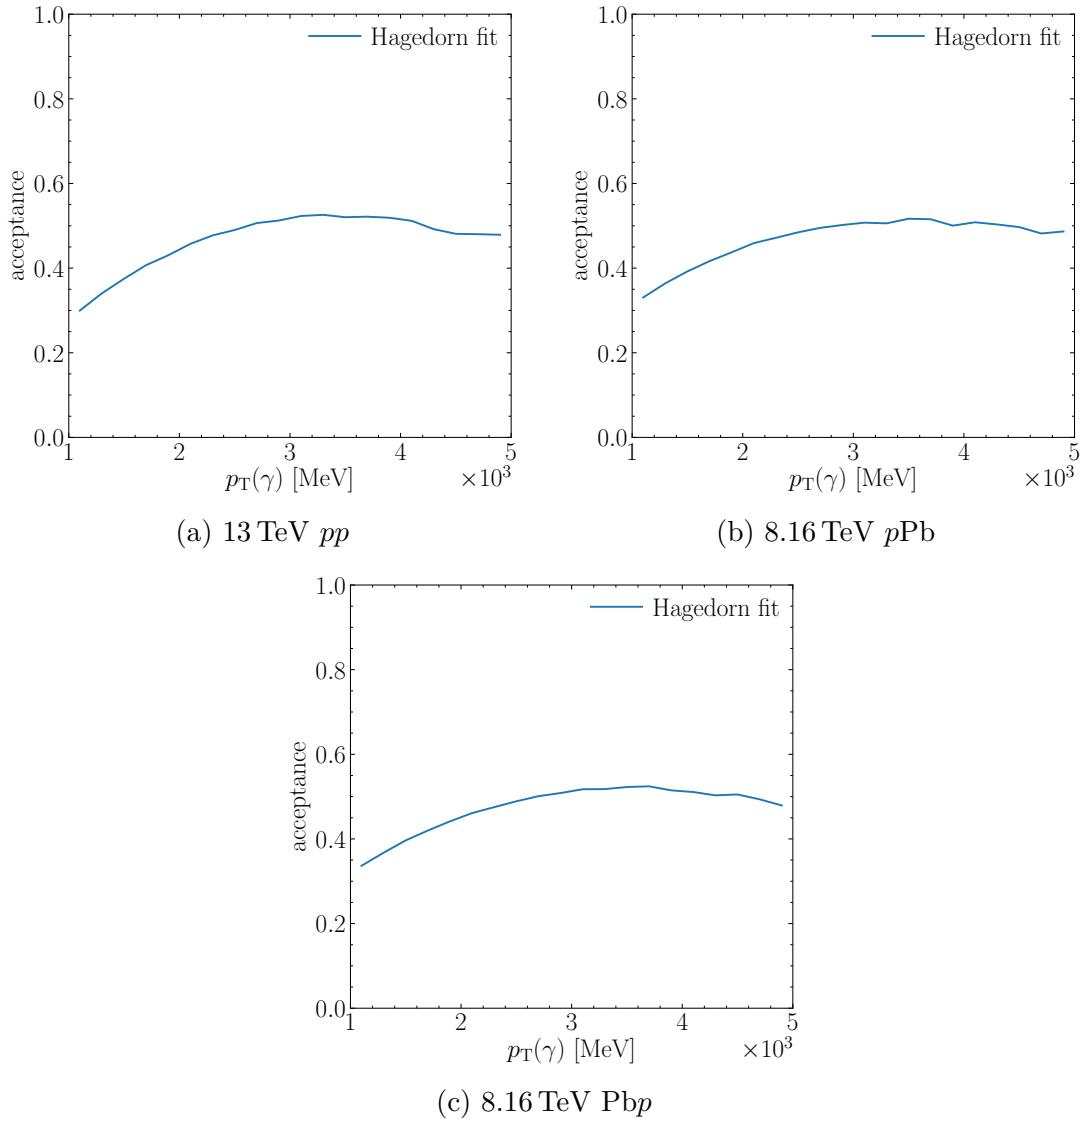


Figure 6-19: Calculated π^0 acceptances from simulation.

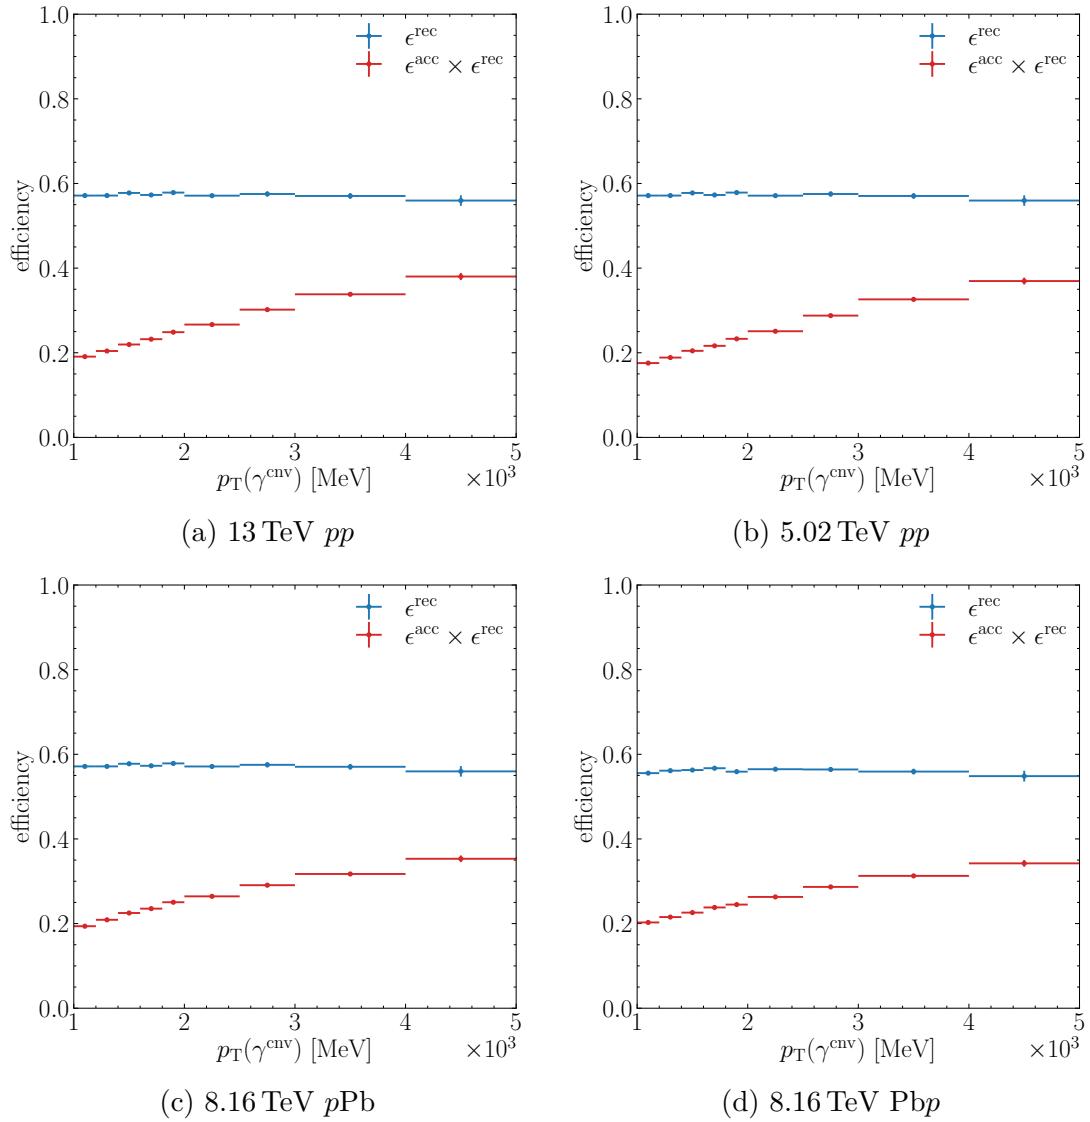


Figure 6-20: Reconstruction and total π^0 efficiencies determined for each data sample.

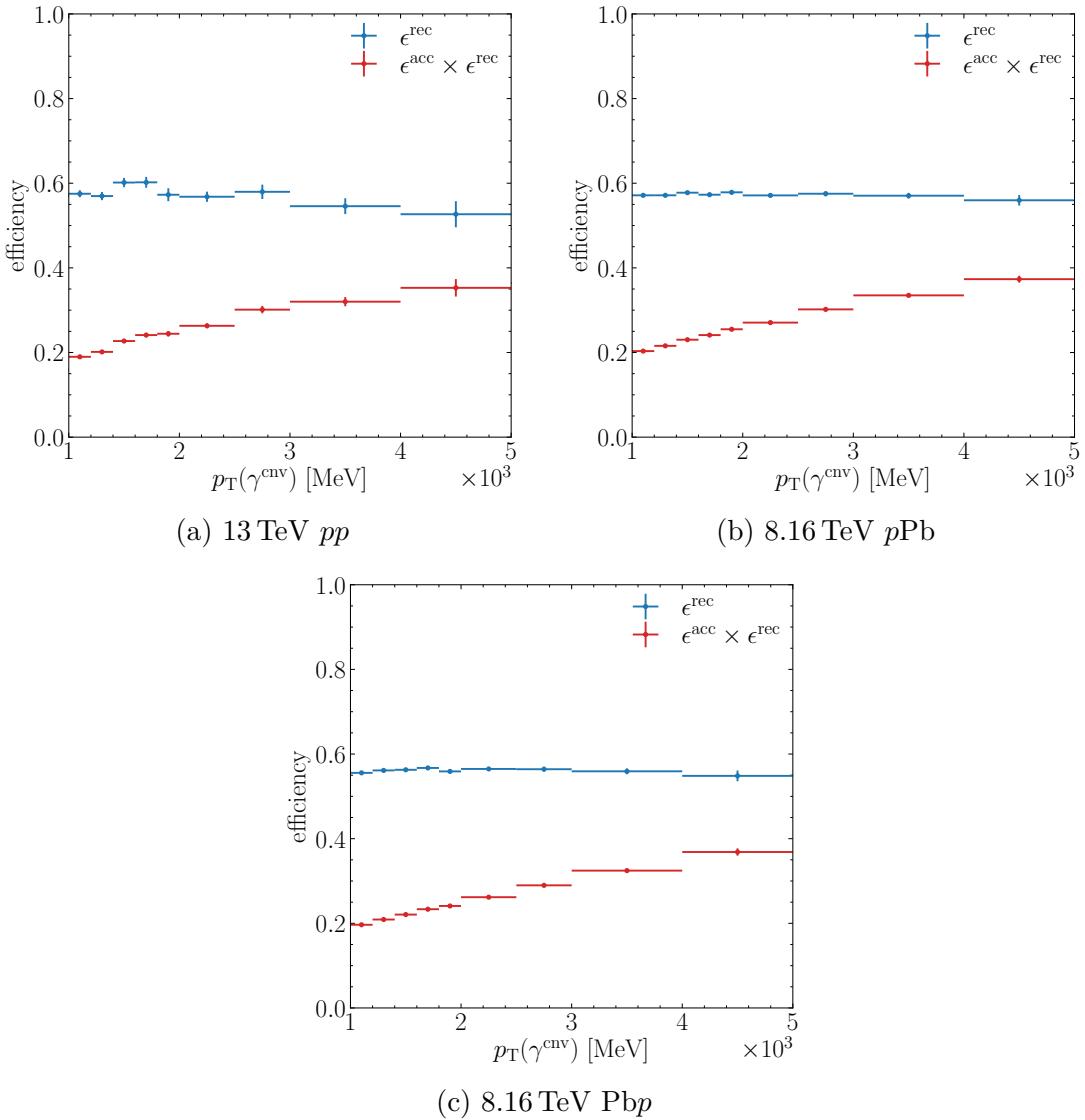


Figure 6-21: Reconstruction and total π^0 efficiencies determined for each simulation sample.

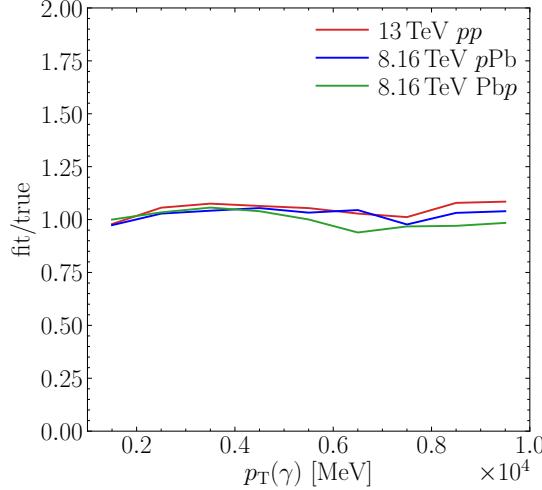


Figure 6-22: Ratios of measured to true acceptance in simulation.

sured and fit values agree to within about 5% and the ratio agrees between datasets to within about 1% at low p_T and 3% at high p_T . The observed agreement indicates that the measured π^0 spectrum accurately reproduces the π^0 decay photon p_T distribution. The total efficiency determination is then validated by comparing the corrected π^0 photon yields to the number of reconstructed converted photons from true π^0 decays in simulation. Comparisons of the corrected and true π^0 photon yields are shown in Figure 6-23. Ratios of the measured to true yields are shown in Figure 6-24. The measured yield overestimates the true yield at high and low p_T . Because good agreement between measured and true acceptances is seen in Figure 6-22 and ϵ^{rec} is almost constant in p_T , the yield discrepancy is likely due to the momentum dependence of the converted photon p_T resolution. However, the disagreement is similar between simulation samples and will largely cancel in R^γ .

6.6.3 π^0 Decay Photon Fraction

The fraction of decay photons originating from π^0 decays, f^{π^0} , may differ in the data and reference samples. About 85% of photons are expected to originate from π^0 decays, and about 10% originate from η decays. As a result, f^{π^0} can be approximated as

$$f^{\pi^0} \approx \frac{N_{\pi^0}^\gamma}{N_{\pi^0}^\gamma + N_\eta^\gamma} = \left(1 + \frac{N_\eta^\gamma}{N_{\pi^0}^\gamma}\right)^{-1}, \quad (6.22)$$

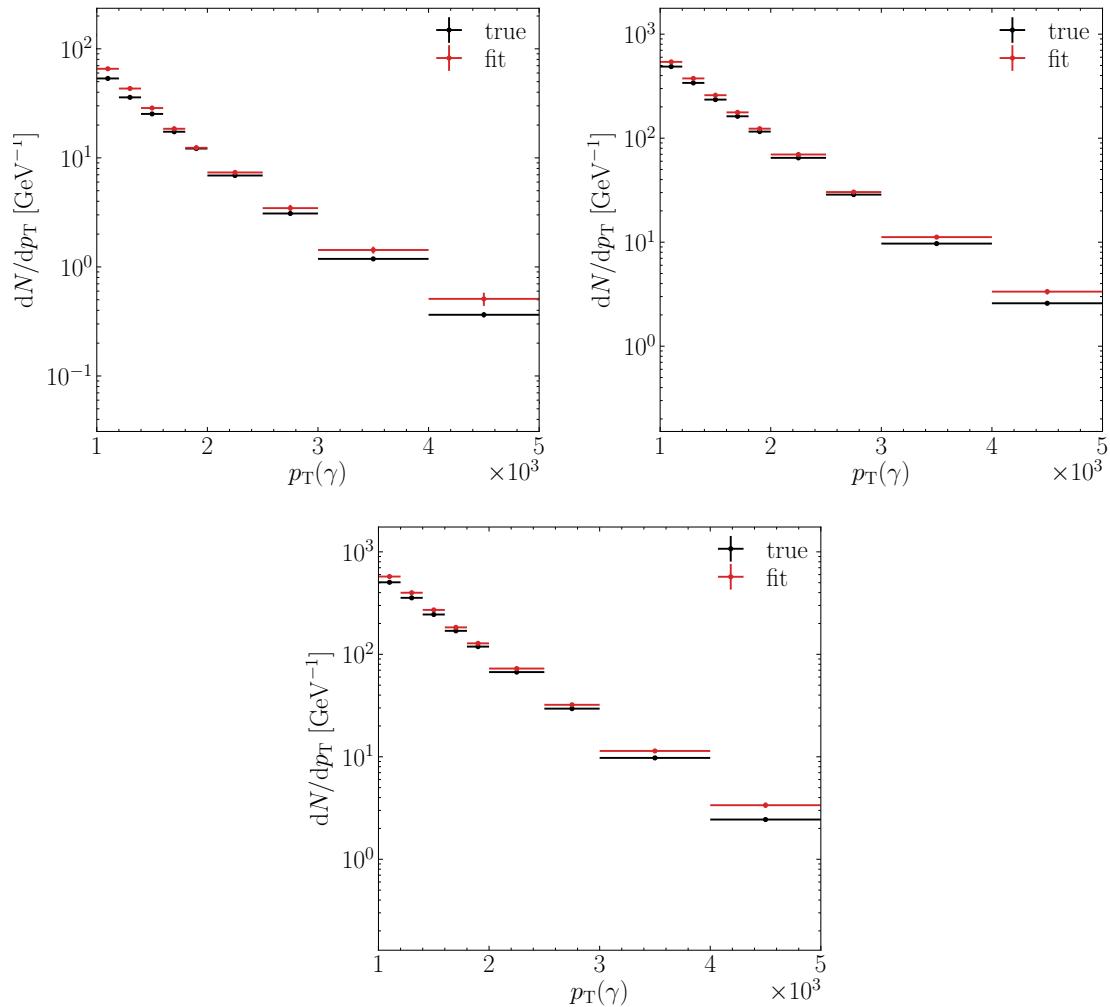


Figure 6-23: True and fully corrected measured π^0 yields in each simulation sample.

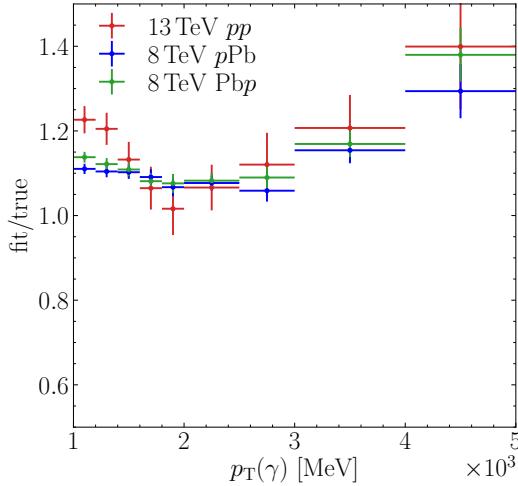


Figure 6-24: Ratio of fully corrected measured π^0 yields to the true yields in simulation.

where $N_{\pi^0}^\gamma$ and N_η^γ are the photon yields from π^0 and η decays, respectively. The factor $f_{\text{data}}^\pi/f_{\text{ref}}^{\pi^0}$ is then given by

$$\frac{f_{\text{data}}^{\pi^0}}{f_{\text{ref}}^{\pi^0}} \approx \frac{\left(1 + N_\eta^\gamma/N_{\pi^0}^\gamma\right)_{\text{ref}}}{\left(1 + N_\eta^\gamma/N_{\pi^0}^\gamma\right)_{\text{data}}}. \quad (6.23)$$

All other particle decays are expected to contribute to the total decay photon yield at a level of around 2%. As a result, we can expect any correction due to differences in these contributions to be a factor of five to ten smaller than the correction due to the η contribution. This will be less than 1% and much smaller than other sources of systematic uncertainty.

For each data sample, η and π^0 yields are extracted using fits to the diphoton mass spectra. Candidates are constructed from combinations of converted and ECAL photons. The large combinatorial background in the η mass region is suppressed by requiring $p_T(\gamma^{\text{cal}}) > 1 \text{ GeV}$. The π^0 yields in simulation are determined using fits to the diphoton mass spectra, while the η yields are determined using the simulation truth record. The π^0 and η acceptances differ. The π^0/η acceptance ratio is calculated using large samples of η and π^0 mesons generated with PYTHIA8 [79]. The acceptance ratio is shown in Fig. 6-25. The acceptance is then used to correct the measured η/π^0 ratio, resulting in a less than 1% correction to $f_{\text{data}}^{\pi^0}/f_{\text{ref}}^{\pi^0}$. The measured η/π^0 ratios are shown in Figure 6-26 for data and Figure 6-27 for simulation. The η/π^0 ratio is

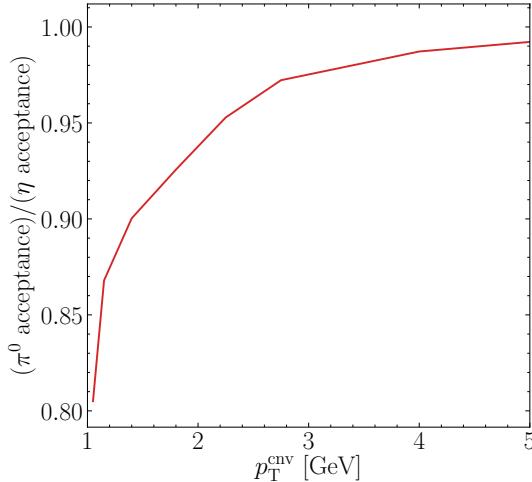


Figure 6-25: Ratio of π^0 to η acceptance from PYTHIA8.

fit to a function inspired by m_T scaling [182] given by

$$\frac{\eta}{\pi^0}(p_T) = A \exp \left(\frac{\sqrt{p_T^2 + m_{\pi^0}^2} - \sqrt{p_T^2 + m_\eta^2}}{T} \right), \quad (6.24)$$

where A and T are fit parameters, and m_{π^0} and m_η are the π^0 and η meson masses, respectively. The fit results are used to calculate f^{π^0} in the data and reference samples.

The resulting values of $f_{\text{data}}^{\pi^0}/f_{\text{ref}}^{\pi^0}$ are shown in Figure 6-28. The f^{π^0} fits vary between datasets by about 1%. The small variation of f^{π^0} is consistent with data from other experiments, which indicate that the η/π^0 ratio is invariant with beam species and center-of-mass energy [183]. The $p\text{Pb}$ and $\text{Pb}p$ simulation samples have smaller values of f^{π^0} than that determined from pp simulation, which is due to differences between EPOS and PYTHIA8.

6.6.4 Resolution Corrections

The inclusive photon and π^0 yields are corrected for detector resolution effects using Bayesian unfolding [184] as implemented in the PyUnfold package [185]. Unfolding is performed in $p_T(\gamma^{\text{cinv}})$. Detector response maps are constructed from simulation and are shown in Figure 6-29. The map constructed from $p\text{Pb}$ simulation is used as the nominal unfolding map and the others are used for systematic checks. The unfolded inclusive photon spectra are shown in Figure 6-30, and the unfolded π^0 photon spectra are shown in Figure 6-31. The effects of unfolding on the measured yields are significant, but the results largely cancel in R^γ . Unfolding ultimately leads to a 1 – 5% effect on the final result which increases as p_T increases.

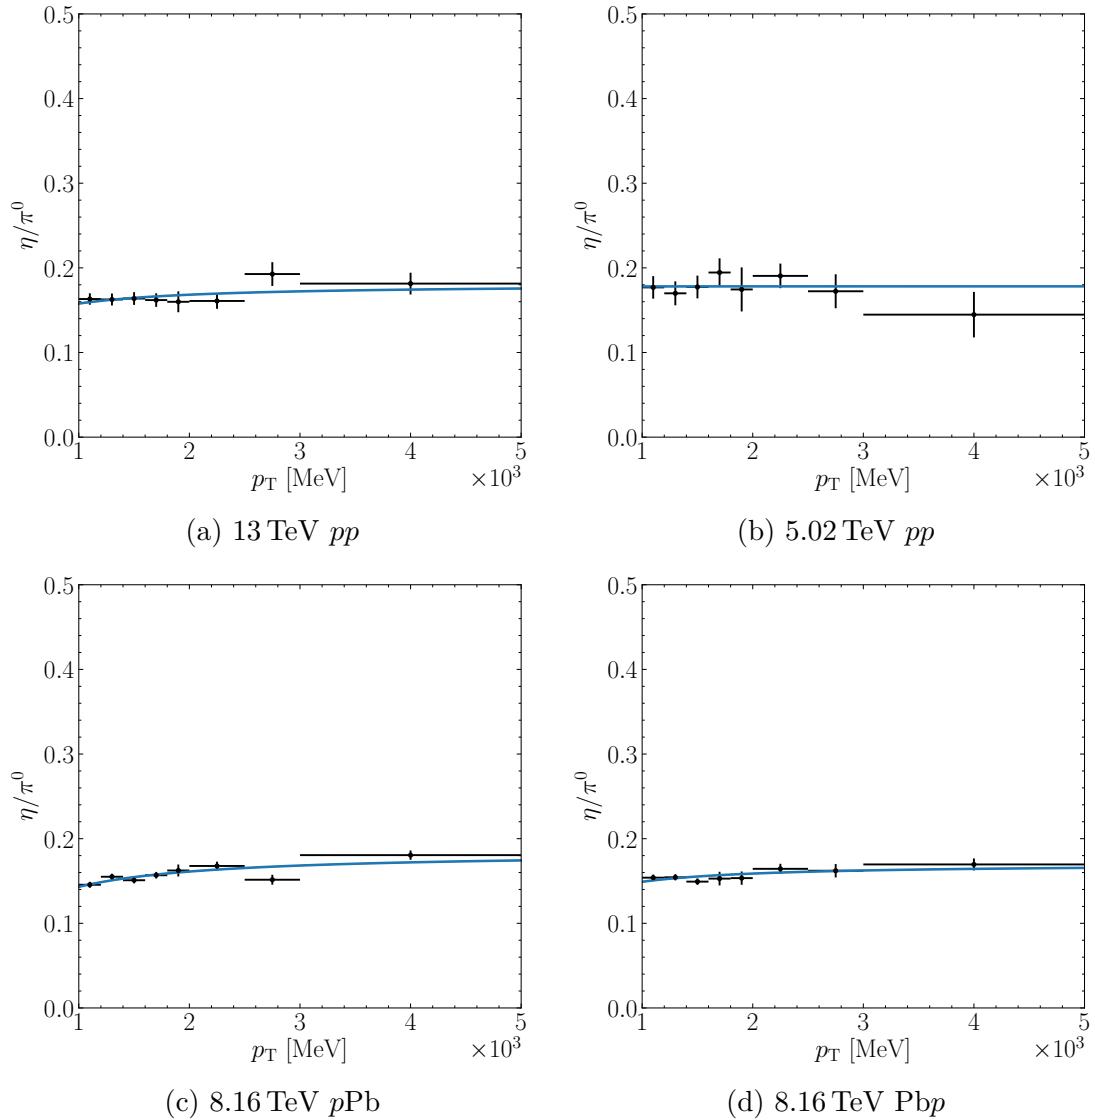


Figure 6-26: Measured η/π^0 ratios in data. Fit results are illustrated by the blue curves.

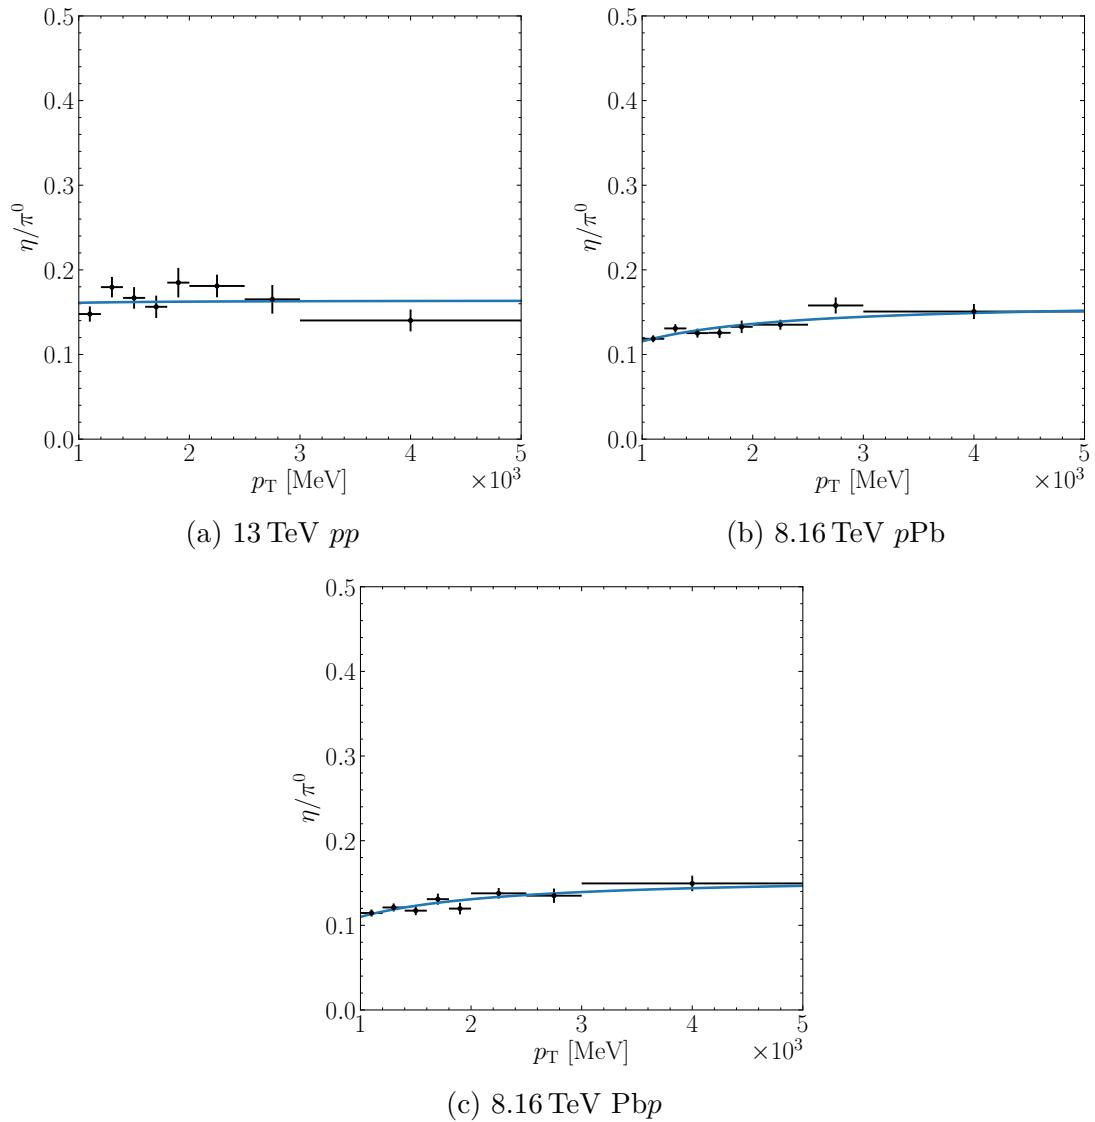


Figure 6-27: Measured η/π^0 ratios in simulation. Fit results are illustrated by the blue curves.

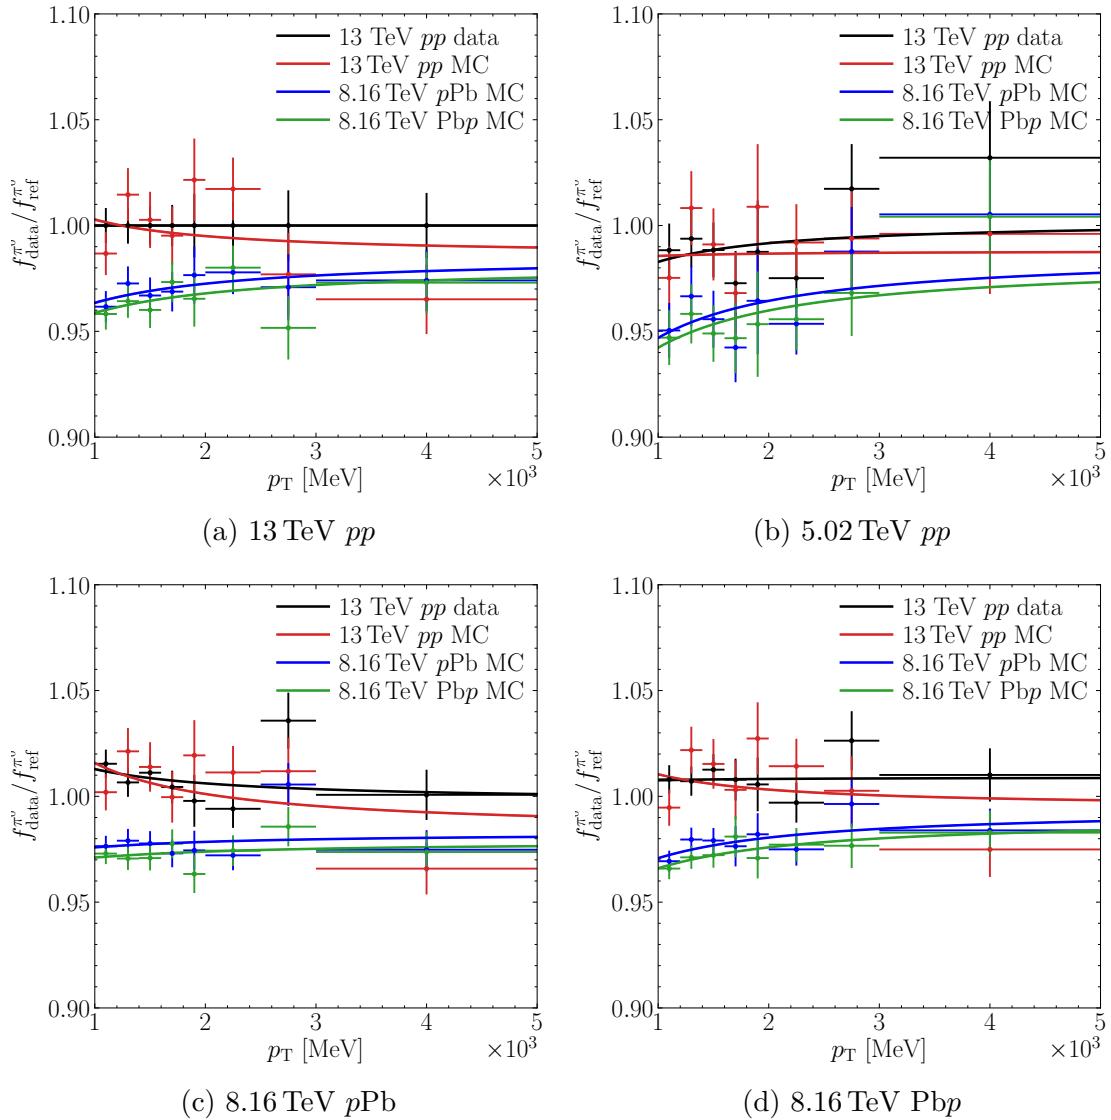


Figure 6-28: Measured $f_{\text{data}}^{\pi^0}/f_{\text{ref}}^{\pi^0}$ for each combination of data and reference sample. The data points show the measured values, while the curves show the fit results.

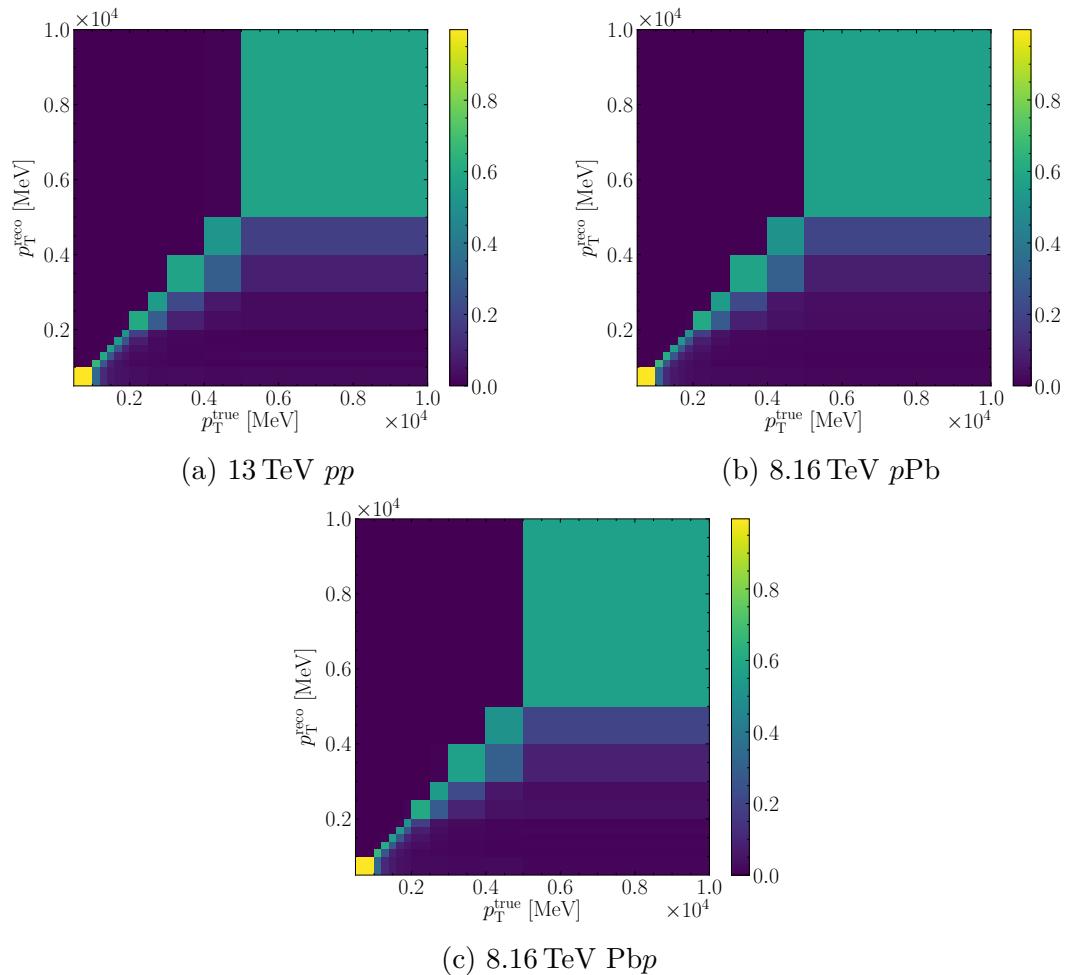


Figure 6-29: Unfolding maps obtained from each simulation sample.

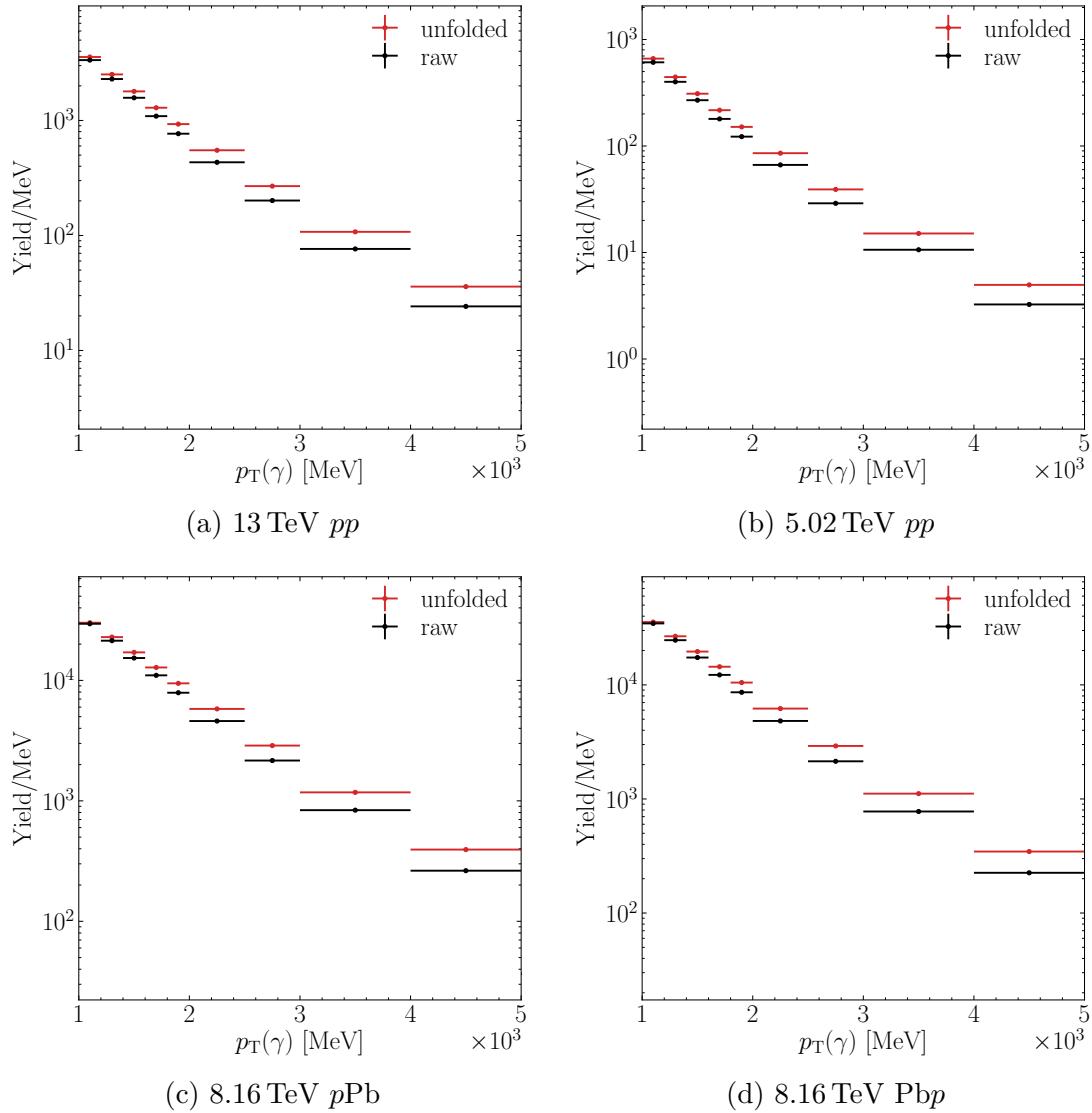


Figure 6-30: Raw and unfolded converted photon yields from data.

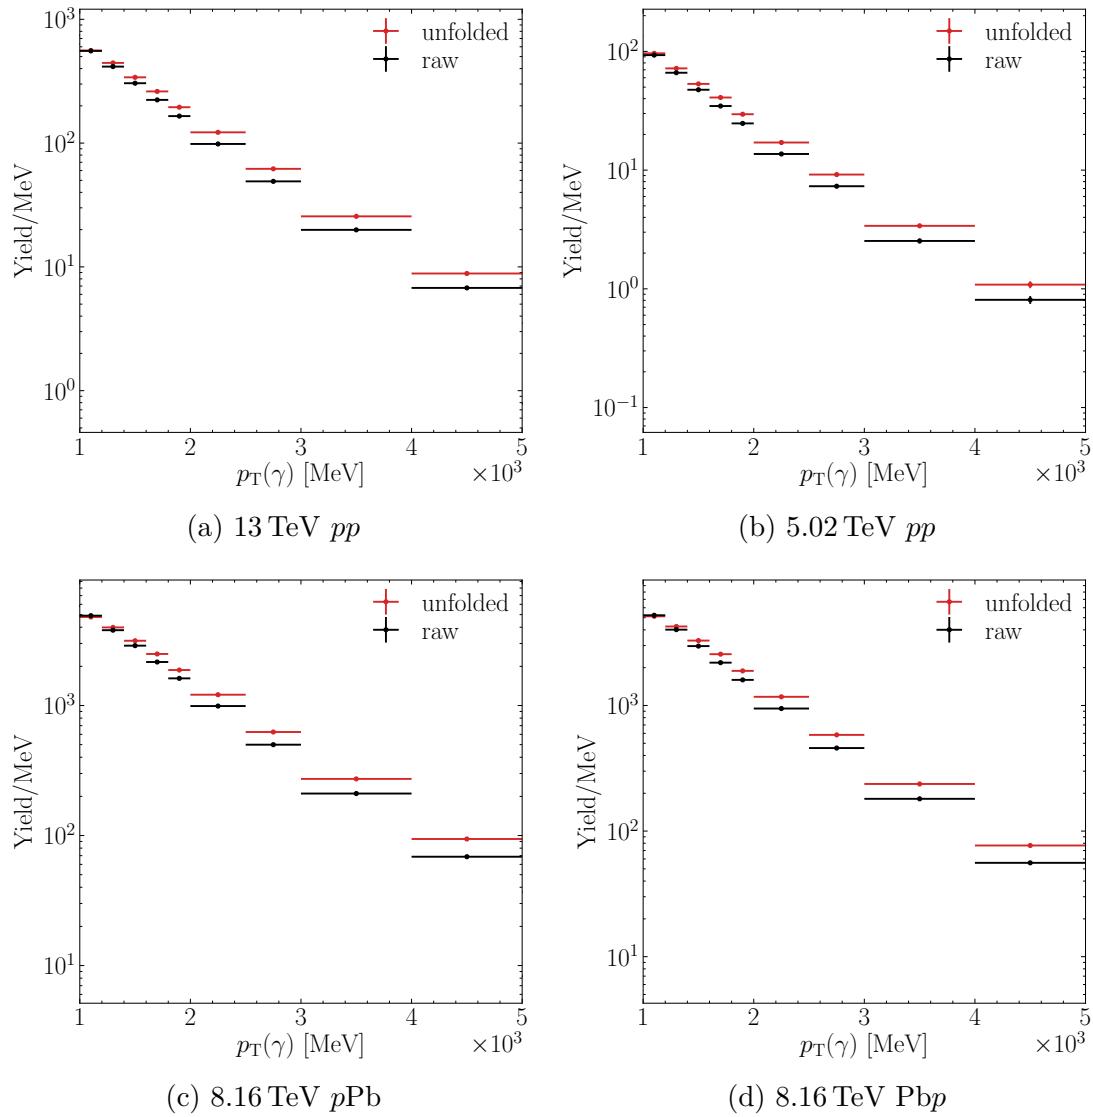


Figure 6-31: Raw and unfolded π^0 yields from data.

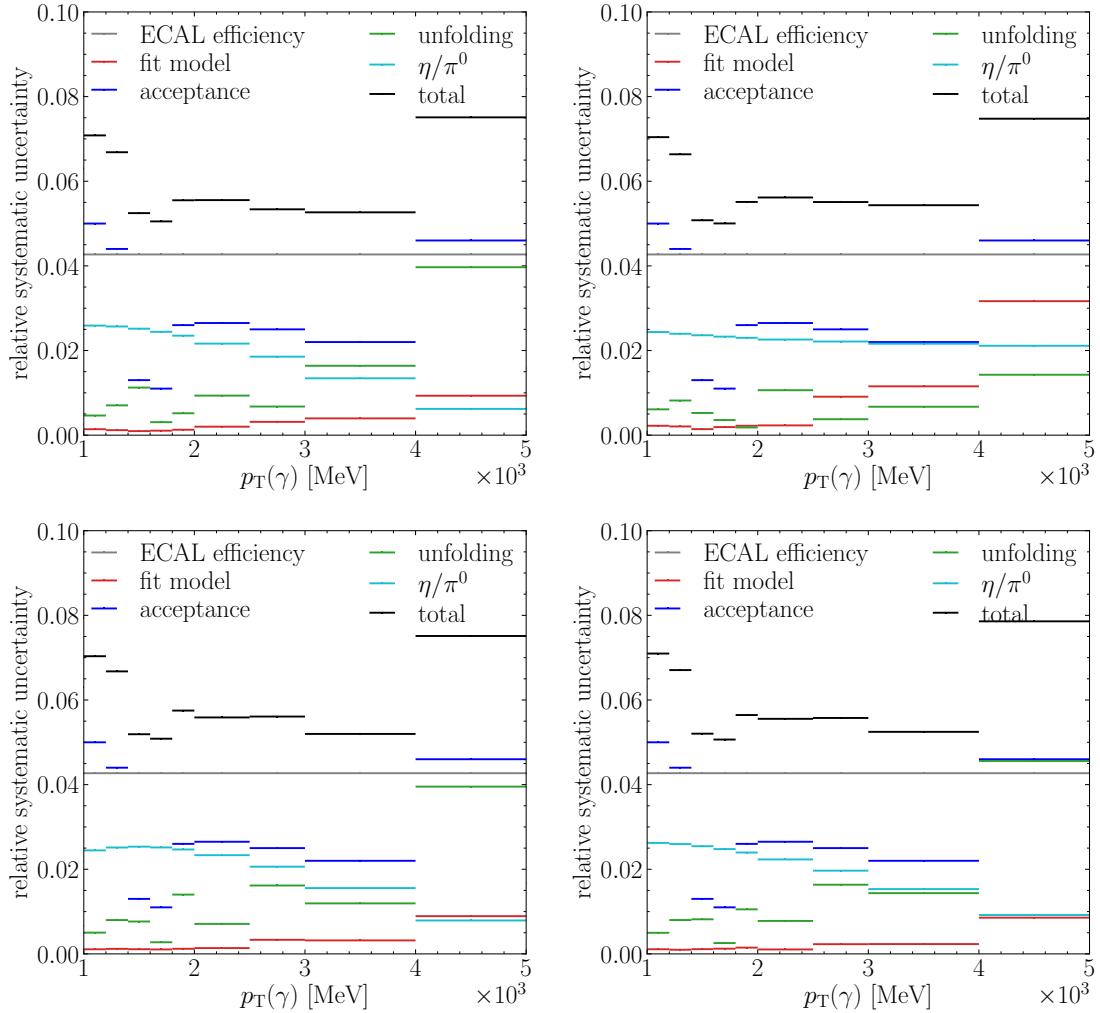


Figure 6-32: Summary of systematic uncertainties for each R^γ measurement.

6.6.5 Systematic Uncertainties

The systematic uncertainties for each R^γ measurement are summarized in Figure 6-32. The uncertainties shown are evaluated using the nominal reference sample, which is a combination of all available simulation samples. The individual sources of systematic uncertainty are described below.

Efficiency Corrections

The ECAL photon efficiency has been measured in 7 TeV and 8 TeV data and was found to agree with simulation to within 4% [179]. As a result, a systematic uncertainty of 4% is assigned to the ECAL photon efficiency. Additionally, the occupancy corrections are calculated with a precision of 1.5% for each data sample, which is included in the ECAL efficiency uncertainty. The resulting total ECAL efficiency uncertainty is 4.2%. The ECAL efficiency is the largest source of systematic uncer-

tainty over most of the p_T range in each measurement. The uncertainty due to the ϵ^{acc} and ϵ^{rec} correction method is determined using the validation study presented in Section 6.6.2. The standard deviation of the ratio measurements shown in Figure 6-24 is taken as a systematic uncertainty.

Fit Model

The systematic uncertainty associated with the π^0 fit model is determined similarly to the fit model uncertainty in Section 6.5. The fixed fit parameters are varied by their uncertainties. As in Section 6.5, the π^0 yields vary by around 4%. The variation largely cancels in the double ratio, resulting in systematic uncertainties of less than 1% over the entire p_T range.

Determination of f^{π^0}

The π^0 and η fit model uncertainties will largely cancel in the determination of $f_{\text{data}}^{\pi^0}/f_{\text{ref}}^{\pi^0}$. Furthermore, the resulting uncertainty on R^γ will be suppressed by a factor of $1/(1+N^{\pi^0}/N^\eta) \approx 0.1$. For this reason the fit model uncertainty is neglected here. An additional uncertainty arises from the choice of η/π^0 ratio fit model. As a systematic check, η/π^0 is fit using a linear function in addition to the nominal m_T -scaling fit. The difference between the nominal and linear fits is taken as a systematic uncertainty. The resulting uncertainty is largest at low p_T , where it is about 2.5%. The uncertainty decreases to about 1% as p_T increases.

Momentum Resolution

Differences in momentum resolution between data and simulation are determined using fits to the mass spectrum of π^0 's reconstructed using pairs of converted photons. The fitting procedure consists of three steps.

- The mass spectrum of π^0 candidates matched to true π^0 mesons in simulation is fit.
- The mass spectrum of reconstructed π^0 candidates in data is fit with a fixed signal shape determined by the fit to simulation.
- The data spectrum is fit again with background parameters fixed and signal parameters free.

The fit results are used to create an ensemble of unfolding maps. The converted photon momentum resolution $\Delta p_T/p_T$ in simulation is fit to a Crystal Ball. The parameters of the Crystal Ball are varied by the relative difference of the corresponding parameter between π^0 mass fits in data and simulation. Each variation is used to smear the generated photon momenta to produce an unfolding map. The procedure is repeated with each simulation sample. Unfolding is performed using each map variation and the standard deviation of the result is taken as the systematic uncertainty. The associated uncertainty is denoted “unfolding” in Figure 6-32.

6.6.6 Results

R^γ is measured using multiple reference samples. Measurements with each combination of data and reference sample are shown in Figure 6-33. The measurements using reference samples from simulation agree within statistical uncertainties in most bins and are combined for the nominal measurement. In addition, the measurements using the 13 TeV pp data sample as a reference agree with those using reference samples from simulation.

Nominal results are calculated using a reference sample consisting of a combination of all available simulated samples. The nominal results are shown in Figures figs. 6-34 to 6-37. Results are compared to perturbative QCD predictions calculated using JETPHOX [65–67]. The predictions are calculated using the CTEQ6 PDF set [68], the EPS09 nuclear PDF set [46], and the BFG II photon fragmentation functions [69]. The predictions include both prompt and fragmentation contributions to the total direct photon cross section. The decay photon cross section is calculated using EPOS [174]. The uncertainties shown include only the proton PDF uncertainties, which are expected to be much smaller than the uncertainties associated with the nuclear PDF, as well as the renormalization, factorization, and fragmentation scales. In addition, no uncertainty is included for the decay photon cross section, which is likely large.

No significant direct photon signal is observed in 5.02 or 13 TeV pp . In 13 TeV pp , R^γ agrees with unity at a level much smaller than the estimated systematic uncertainties at low p_T , where the direct photon contribution is expected to be negligible. In particular, the agreement is much better than the 4.1% ECAL efficiency uncertainty, which is estimated from previous LHCb studies. The observed agreement indicates that the R^γ measurement presented here could serve as the most precise measurement of the LHCb ECAL photon efficiency to date.

The $p\text{Pb}$ and $\text{Pb}p$ measurements are both consistent with pQCD predictions over most of the studied p_T range. Although both measurements are consistent with unity within systematic uncertainties at $p_T < 2$ GeV, the central values are consistently greater than one and are roughly constant for $p_T < 4$ GeV. Similar behavior would be expected in the presence of a thermal photon contribution to R^γ of about 5%. An improved estimate of the LHCb ECAL efficiency, as well as dedicated studies of the π^0 and η meson p_T distributions, would reduce the dominant systematic uncertainties. As a result, a future measurement has the potential to provide decisive evidence for direct photon production for p_T less than 2 GeV.

6.7 Conclusion

This chapter has presented measurements of $R_{p\text{Pb}}^{\pi^0}$ in proton-lead collisions at $\sqrt{s_{\text{NN}}} = 8.16$ TeV. These are the only measurements of $R_{p\text{Pb}}^{\pi^0}$ at far forward and backward rapidities at the LHC. As such, the $R_{p\text{Pb}}^{\pi^0}$ measurements presented here provide information about the partonic structure of the nucleus in a previously unexplored kinematic regime. The forward $p\text{Pb}$ measurement shows clear shadowing behavior and will

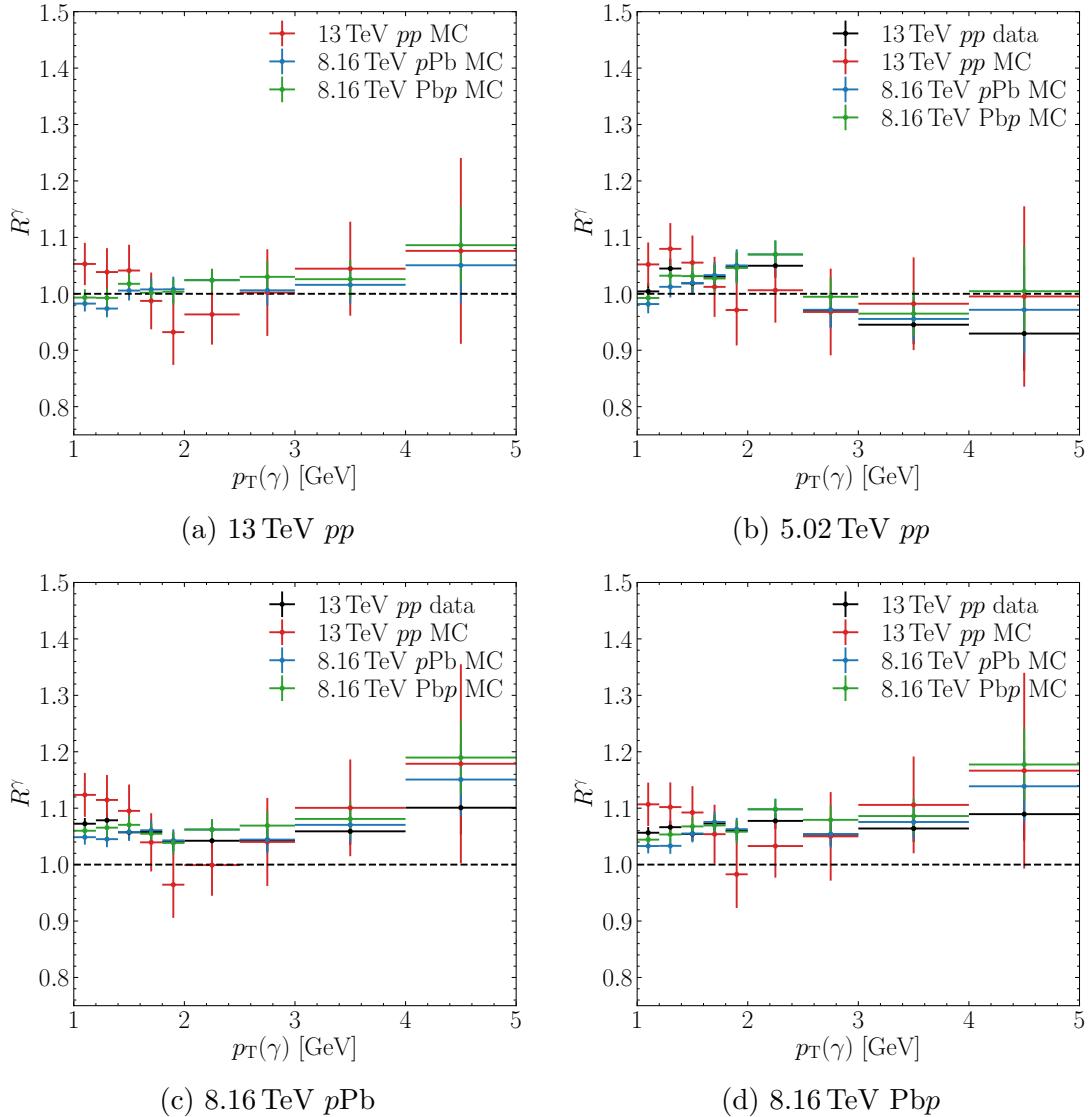
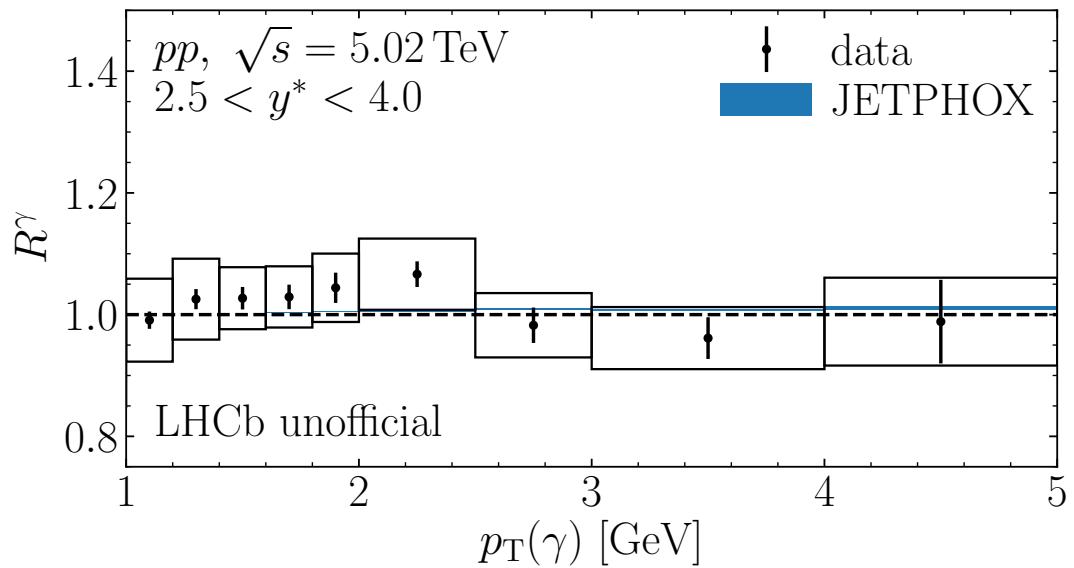
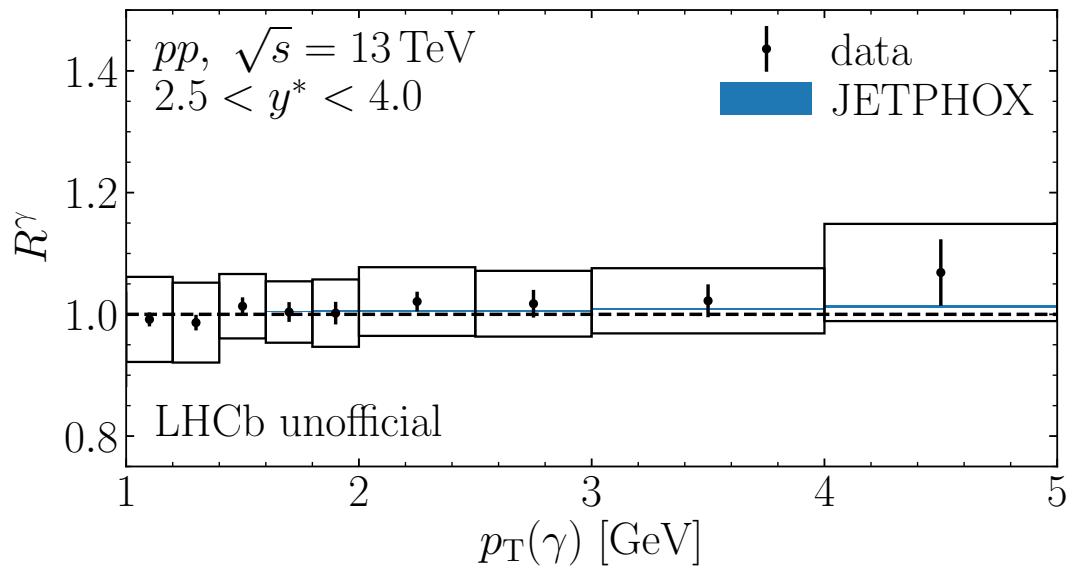


Figure 6-33: R^γ measurements performed with each reference sample. The 13 TeV pp data reference sample is omitted from a, as this measurement is trivially unity.



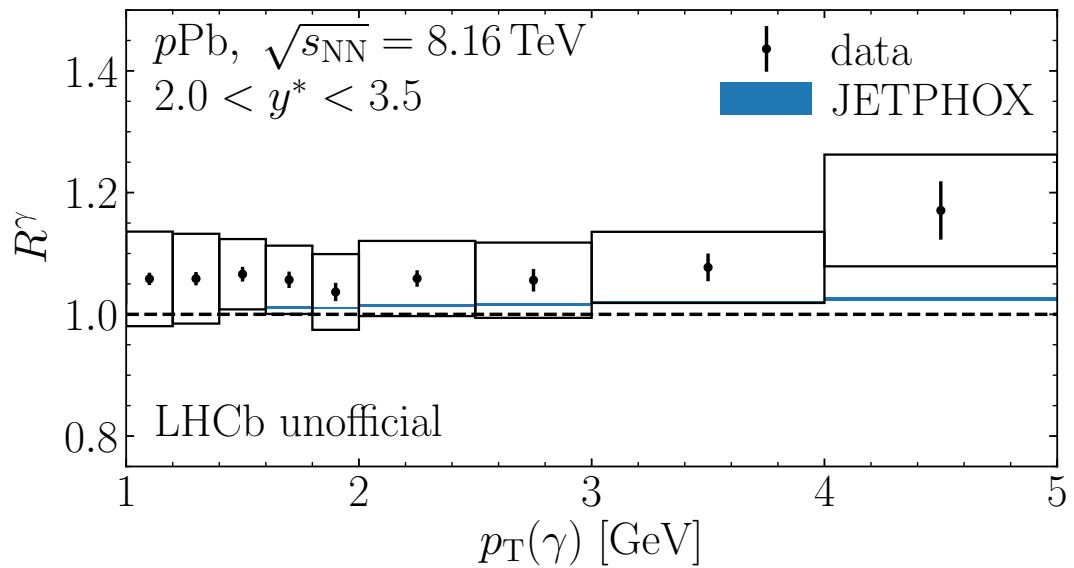


Figure 6-36: R_γ measured in $p\text{Pb}$ collisions at $\sqrt{s_{\text{NN}}} = 8.16 \text{ TeV}$.

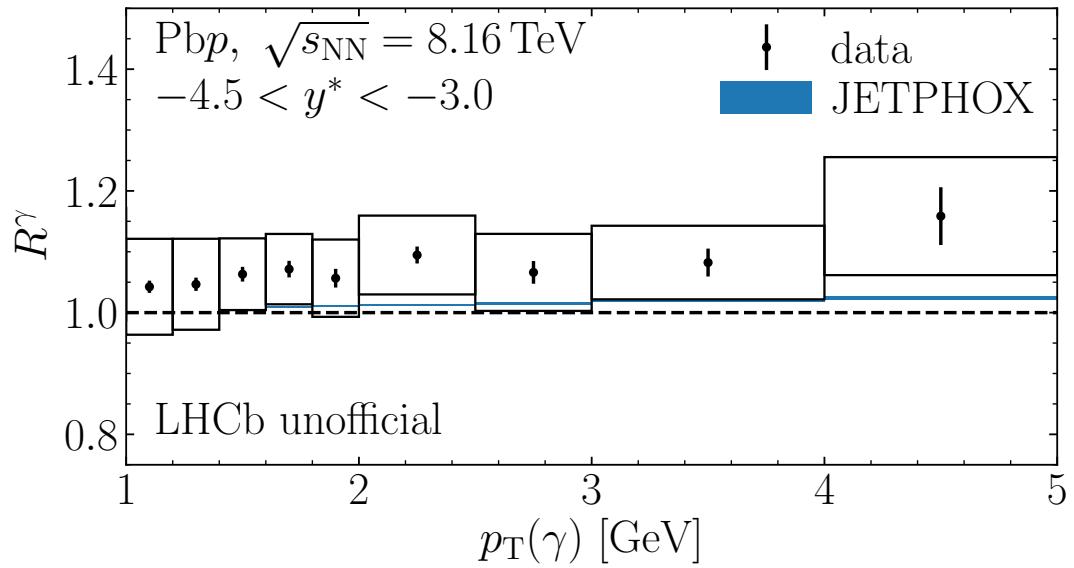


Figure 6-37: R_γ measured in Pbpb collisions at $\sqrt{s_{\text{NN}}} = 8.16 \text{ TeV}$.

constrain the gluon nPDF at low- x . The backward measurement of $R_{p\text{Pb}}^{\pi^0}$ shows a large enhancement in the anti-shadowing region. This enhancement is much larger than that predicted by pQCD calculations. The excess is similar to those observed in measurements of the charged particle nuclear modification factor at CMS [186] and charged hadron nuclear modification factor at PHENIX [187]. The CMS measurement at $p_{\text{T}} > 50 \text{ GeV}$ and the PHENIX measurement at $1 < p_{\text{T}} < 4 \text{ GeV}$ and backward rapidities probe a similar region in x as the measurement presented here. Taken together, these measurements present consistent experimental evidence for a large enhancement in the anti-shadowing region beyond that predicted by nPDFs. Measurements such as the nuclear modification of charged particles at backward rapidities with the LHCb detector would provide additional probes of the same physical processes with almost completely independent systematic uncertainties.

The R^{γ} measurements are also the first of their kind at far forward and backward rapidities at the LHC, providing the first constraints on direct photon production in the LHCb fiducial region. No significant direct photon signal is observed for $p_{\text{T}} < 4 \text{ GeV}$. The pp measurements demonstrate the potential of these measurements to improve the understanding of the LHCb photon reconstruction performance and can be used to improve the systematic uncertainties in future R^{γ} measurements in proton-lead collisions.

Chapter 7

Conclusion

Measurements of neutral pion and photon production with the LHCb detector were presented. These measurements are the first of their kind at far forward and backward rapidities at LHC energies. The π^0 nuclear modification factor demonstrates suppression in the forward region consistent with nuclear shadowing at low- x . This measurement provides constraints on the nuclear gluon density in unexplored kinematic territory. Enhancement is observed in the backward direction, consistent with anti-shadowing at high- x . This measurement is part of a consistent experimental picture pointing to enhancements of π^0 and charged particle production at both RHIC and the LHC beyond that predicted by nPDF calculations.

Additionally, measurements of the direct photon excess ratio R^γ were presented. No significant direct photon signal is observed in proton-proton collisions. The pp measurements at $\sqrt{s} = 13$ TeV and 5 TeV agree with pQCD calculations. The 13 TeV measurement in particular agrees with unity at a level much smaller than systematic uncertainties, indicating that measurements of R^γ in pp collisions could be used to measure the LHCb ECAL performance in the future. Additionally, hints of direct photon production are seen in proton-lead collisions. The R^γ measurements presented here lay the groundwork for better understanding the LHCb detector's photon detection performance, allowing for future improvements of systematic uncertainties. Future measurements have the potential to either reveal thermal photon production in proton-lead collisions or tightly constrain it. The measurements presented in this thesis are some of the only studies of all-photon final states using the LHCb detector. The methods developed for this work open the door to a wide variety of future measurements, such as studies of η meson and isolated photon production.

In addition to measurements of π^0 and photon production, the Allen project was also presented. Allen is the first full trigger stage at a high-energy particle physics experiment designed for GPUs. The resulting performance gains will help broaden the LHCb physics program in Run 3. The version of Allen presented in this thesis is a prototype designed to demonstrate the feasibility of running a full high level trigger on GPUs. Since Allen was selected as the nominal LHCb HLT1 in the summer of 2020, work has begun to take Allen from prototype to production software ready for real data taking in 2022.

Appendix A

Kalman Filter Parameterizations

The parameterized Kalman filter presented in Section 4.3.4 uses different parameterizations for each subdetector, as well as for the VELO-UT and UT-SciFi regions. Only the UT-SciFi extrapolation explicitly accounts for energy loss as the particle passes through the detector material. Otherwise, the extrapolations are tuned to simulations that include energy loss and consequently implicitly capture its effects. Consequently the value q/p used in the track state will always be that of the state closest to the beamline.

A.1 Propagation in the VELO, UT, and SciFi

In general effect of the LHCb magnetic field within each subdetector is approximated as a single kick. The position of the kick along the z axis, as well as the magnitudes of the changes in track slopes t_x and t_y , are provided by parameterizations. The weak magnetic field in the VELO region allows for further simplification. Here a straight line extrapolation is used for y and a small first order correction in q/p is applied to t_x at the midpoint of the extrapolation.

Extrapolations within the UT and SciFi are handled similarly. In both cases the predicted t_x takes the form

$$t_{x,k} = t_{x,k-1} + \left(P_0 \frac{q}{p} + P_1 \left(\frac{q}{p} \right)^3 + P_2 y_{k-1}^2 \frac{q}{p} \right) \Delta z, \quad (\text{A.1})$$

where the coefficients P_n take different values depending on the subdetector and layer. The predicted t_y is given by

$$t_{y,k} = t_{y,k-1} + P \left(\frac{q}{p} \right)^n y_{k-1}, \quad (\text{A.2})$$

where $n = 1$ in the UT and $n = 2$ in the SciFi. The kicks in the x and y directions

are applied at different positions in z . These z positions are given by parameters for each layer of each detector subsystem, $z_{k-1 \rightarrow k}^x$ and $z_{k-1 \rightarrow k}^y$.

A.2 Propagation in the VELO-UT Region

The extrapolation between the last layer of the VELO and the first layer of the UT is again performed using a single kick approximation. The extrapolated t_x is given by

$$t_{x,k} = \frac{a_{k-1}}{\sqrt{(1 - a_{k-1}^2)(1 + t_{y,k}^2)}}, \quad (\text{A.3})$$

where a_i is given by

$$a_i = \frac{t_{x,i}}{\sqrt{1 + t_{x,i}^2 + t_{y,i}^2}} - \frac{q}{p} (P_0 + P_1 z_0 + P_2 t_{y,i}^2). \quad (\text{A.4})$$

Here z_0 is the starting point of the extrapolation. Similarly to the UT and SciFi regions, the predicted t_y is given by

$$t_{y,k} = t_{y,k-1} + P \frac{q}{p} t_{x,k-1}. \quad (\text{A.5})$$

The z position of the kick in x is given by

$$z_{k-1 \rightarrow k}^x = P_0 + P_1 z_0 + P_2 z_0^2 + P_3 t_{y,k-1}^2, \quad (\text{A.6})$$

while the z position of the kick in y is given by a single parameter $z_{k-1 \rightarrow k}^y$.

A.3 Propagation in the UT-SciFi Region

The UT-SciFi region includes the dipole magnet and most of the track's deflection. The extrapolation in this region is handled using an expansion around a hypothetical reference tracks originating from $x = y = z = 0$. First an updated $(q/p)_{\text{cor}}$ is calculated to account for energy loss using a second order expansion in q/p . The reference tracks are then extrapolated using a fourth order expansion in $(q/p)_{\text{cor}}$. The coefficients of this expansion are then corrected using a first order expansion in the deviations δ_{t_x} and δ_{t_y} of the slopes of the true track from the reference track. It should be noted that the expansion around reference tracks originating from the origin means that this method is only valid in the VELO \rightarrow SciFi direction. The backwards extrapolation is performed using the inverse of the Jacobian used for the forward extrapolation.

A.4 Multiple Scattering Noise

In each extrapolation region the multiple scattering noise matrix Q_k is given by

$$Q_{k,ij} = C_{ij}\sigma_i\sigma_j$$

for $i, j \neq q/p$, where $C_{ii} = 1$. The other correlation coefficients C_{ij} are empirical parameters. The single variable uncertainties are given by

$$\sigma_{x,y} = P_{x,y} \left| \frac{q}{p} \Delta z \right|, \quad (\text{A.7})$$

$$\sigma_{t_x,t_y} = P_{t_x,t_y} \left| \frac{q}{p} \right|, \quad (\text{A.8})$$

where P_x , P_y , P_{t_x} , and P_{t_y} are separate parameters.

Appendix B

Diphoton Mass Fits

B.1 Fits for $R_{p\text{Pb}}^{\pi^0}$

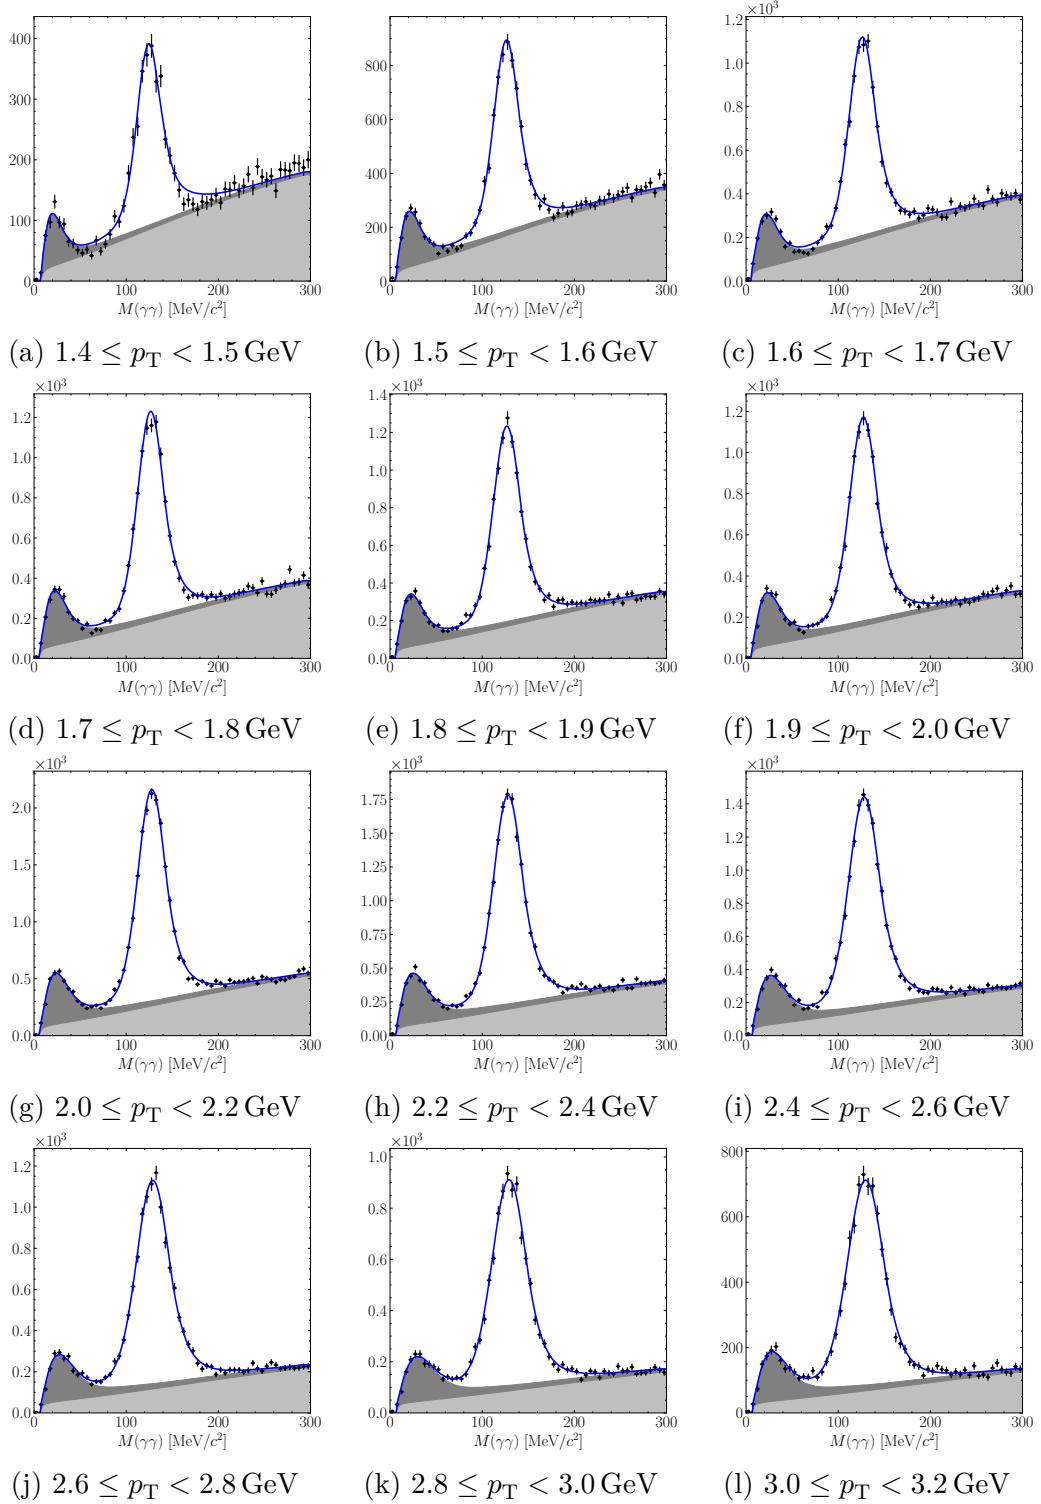


Figure B-1: π^0 mass fits in 13 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$.

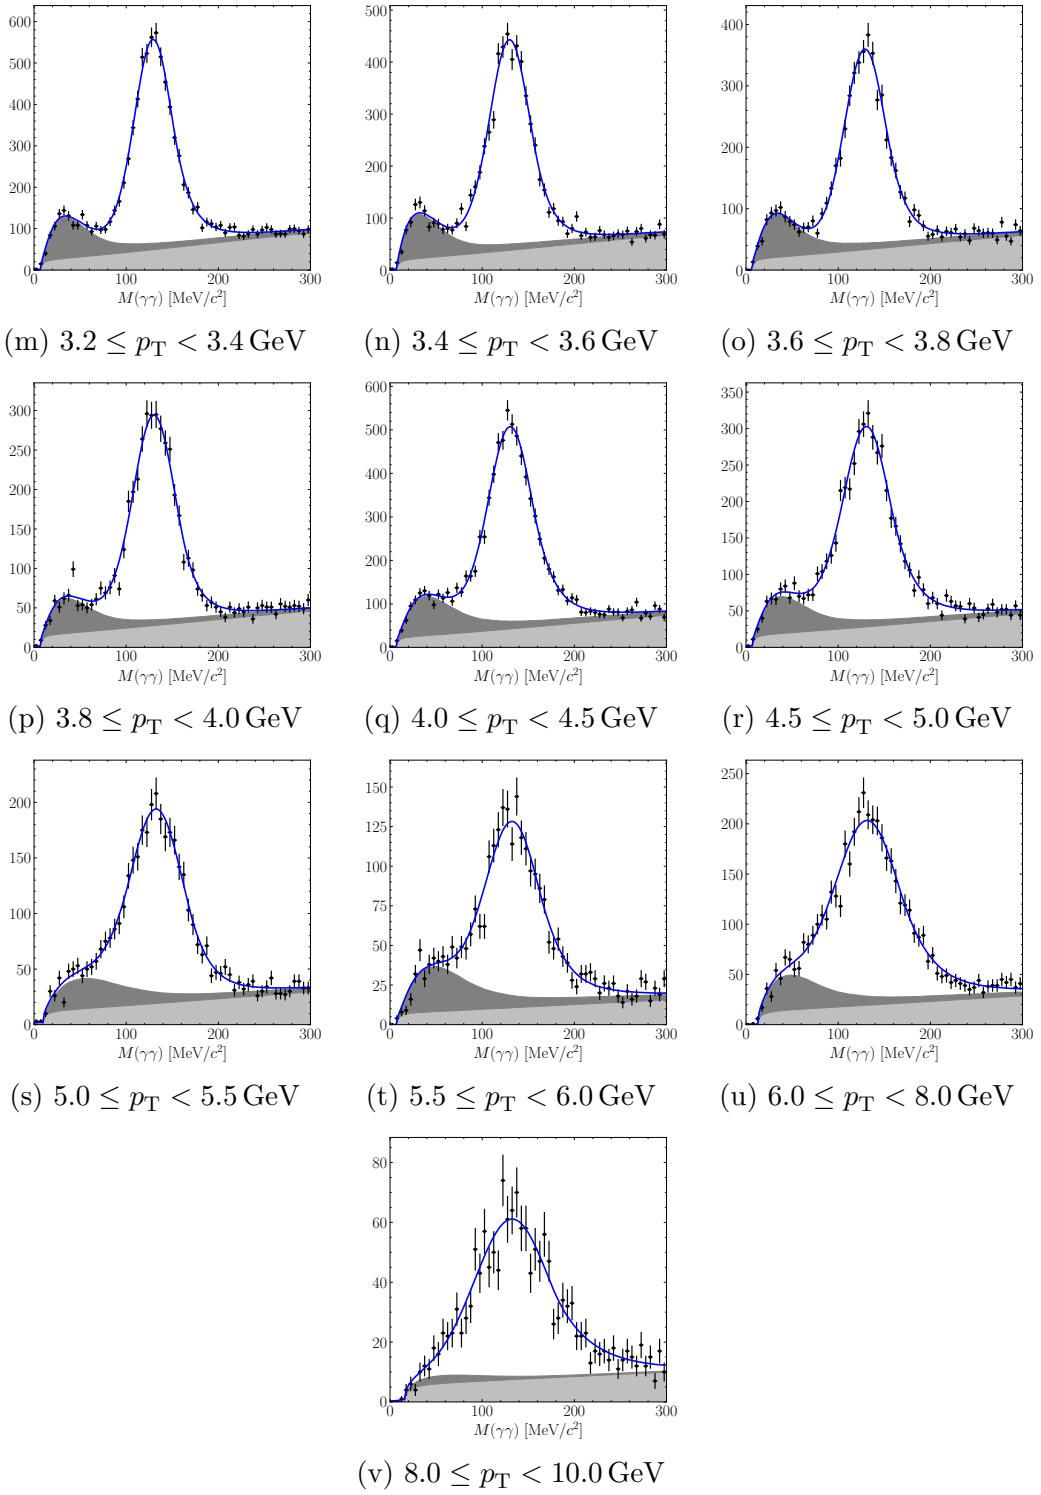


Figure B-1: π^0 mass fits in 13 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).

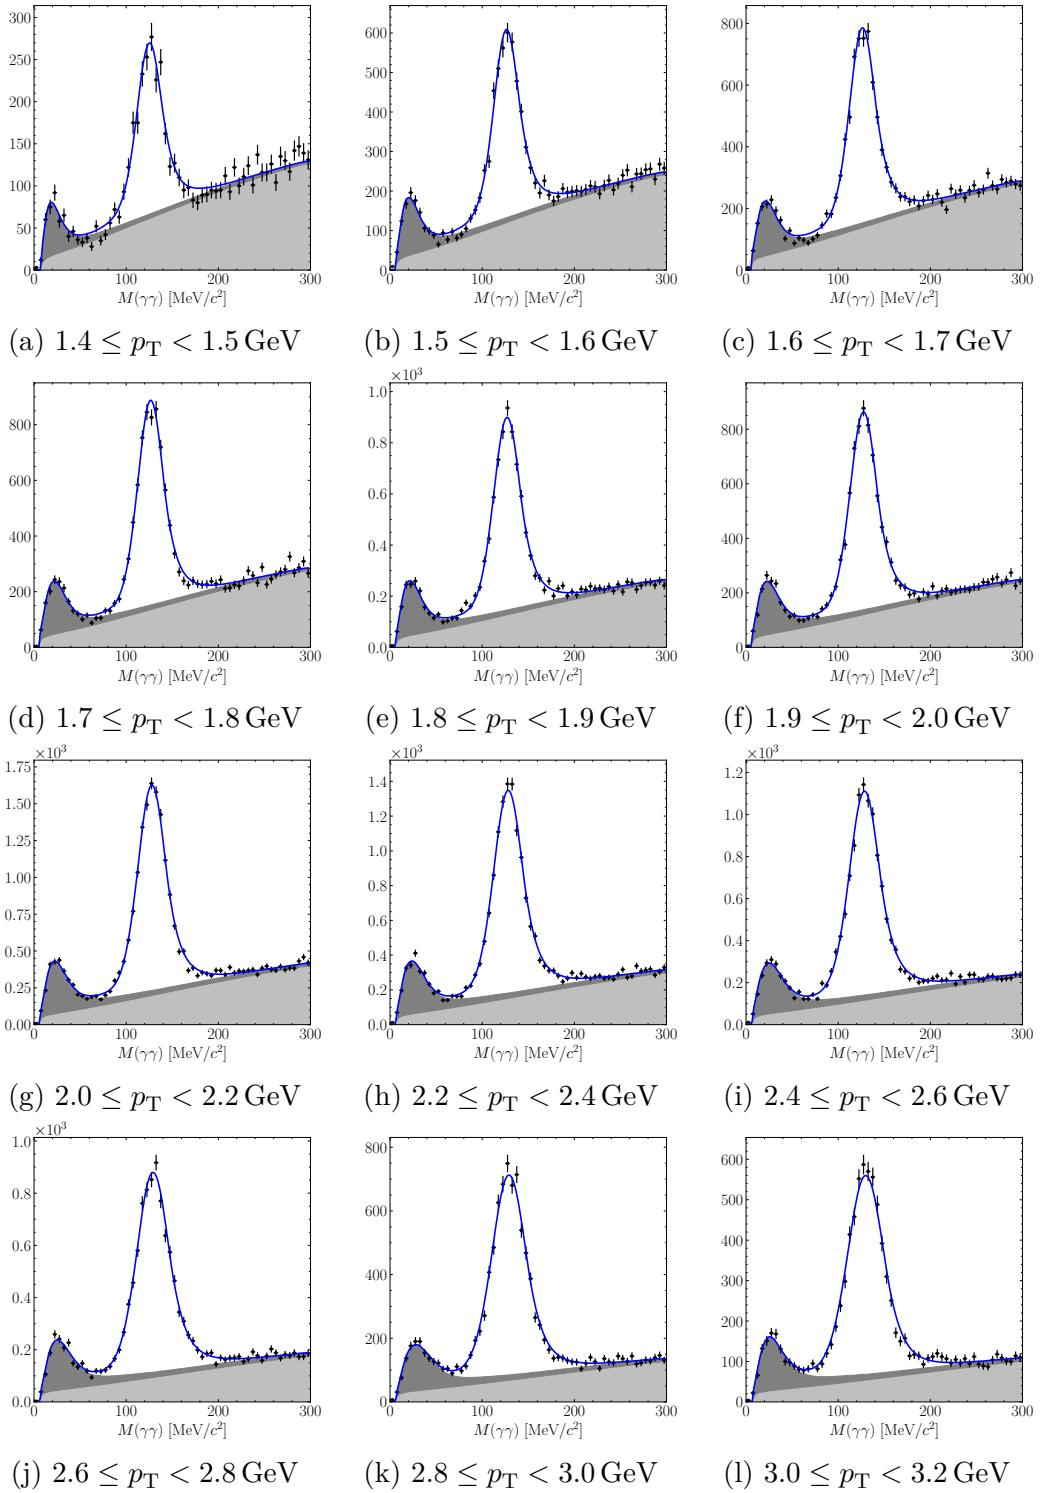


Figure B-2: π^0 mass fits in 13 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$.

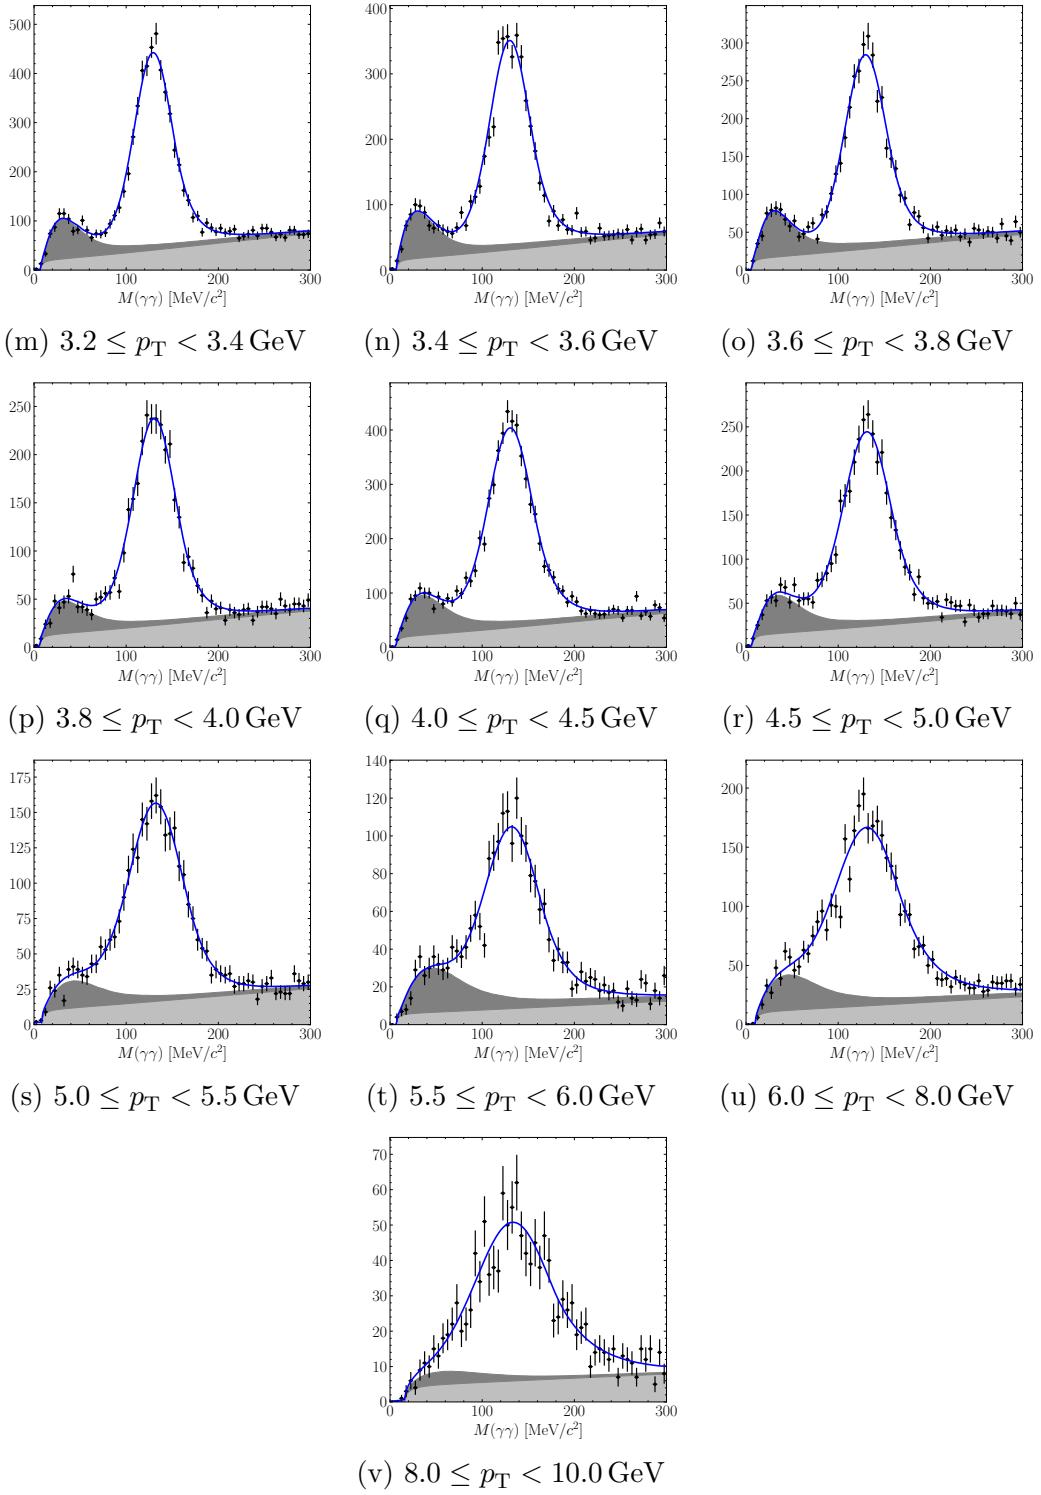


Figure B-2: π^0 mass fits in 13 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$ (cont.).

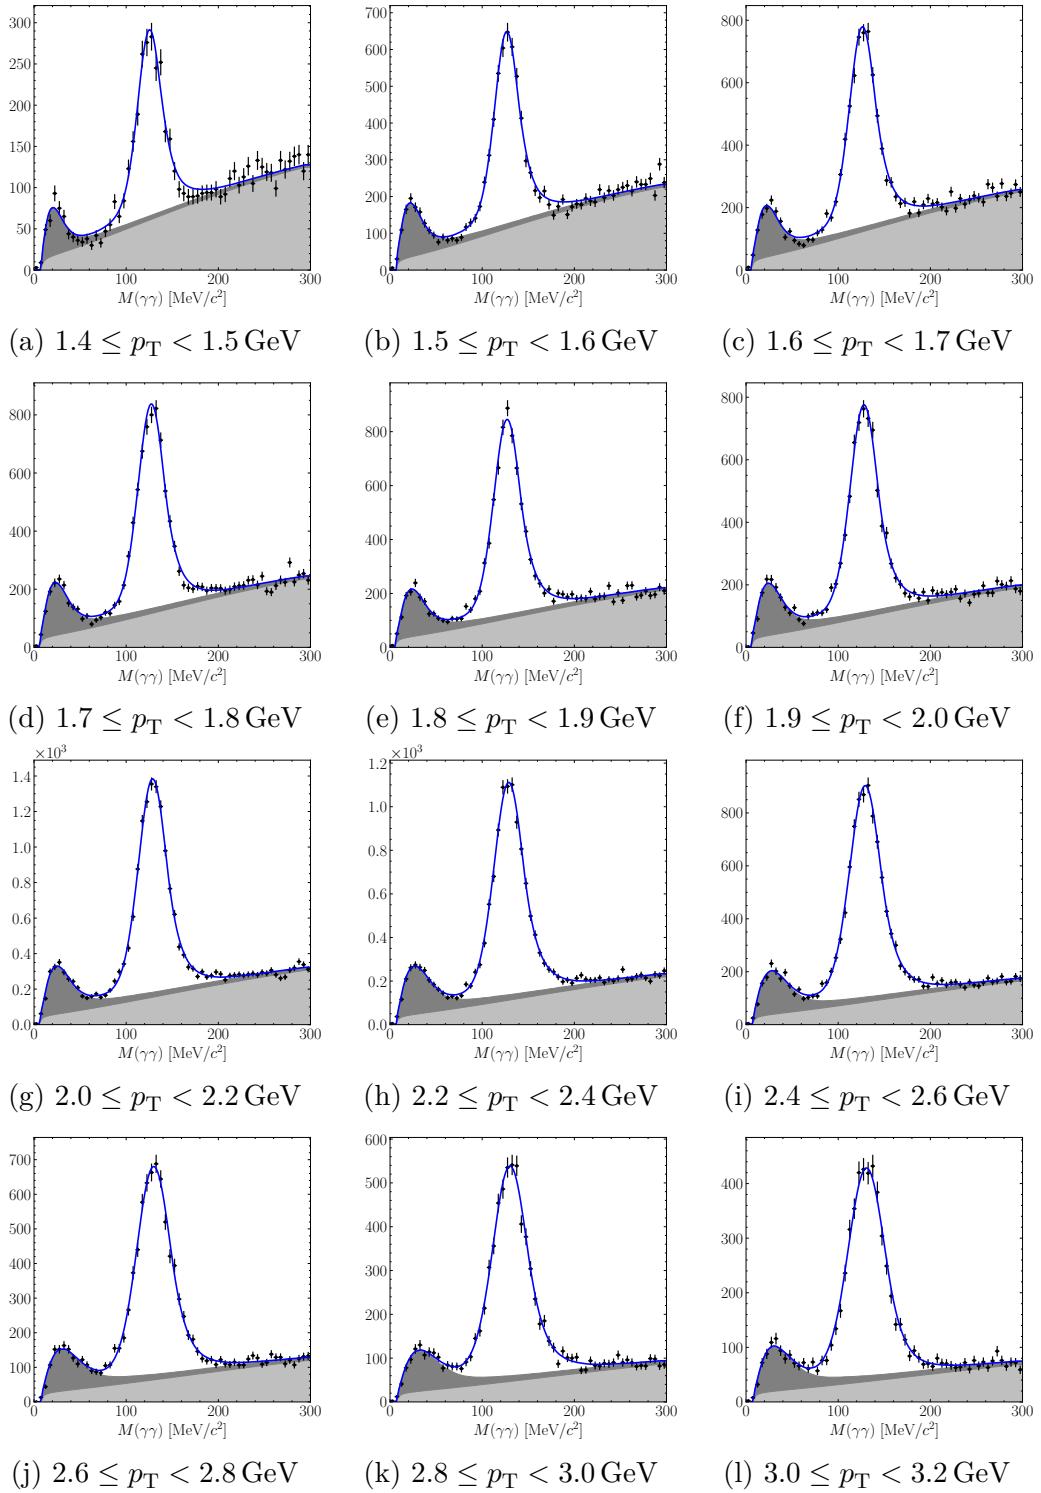


Figure B-3: π^0 mass fits in 13 TeV pp data as a function of $p_T(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$.

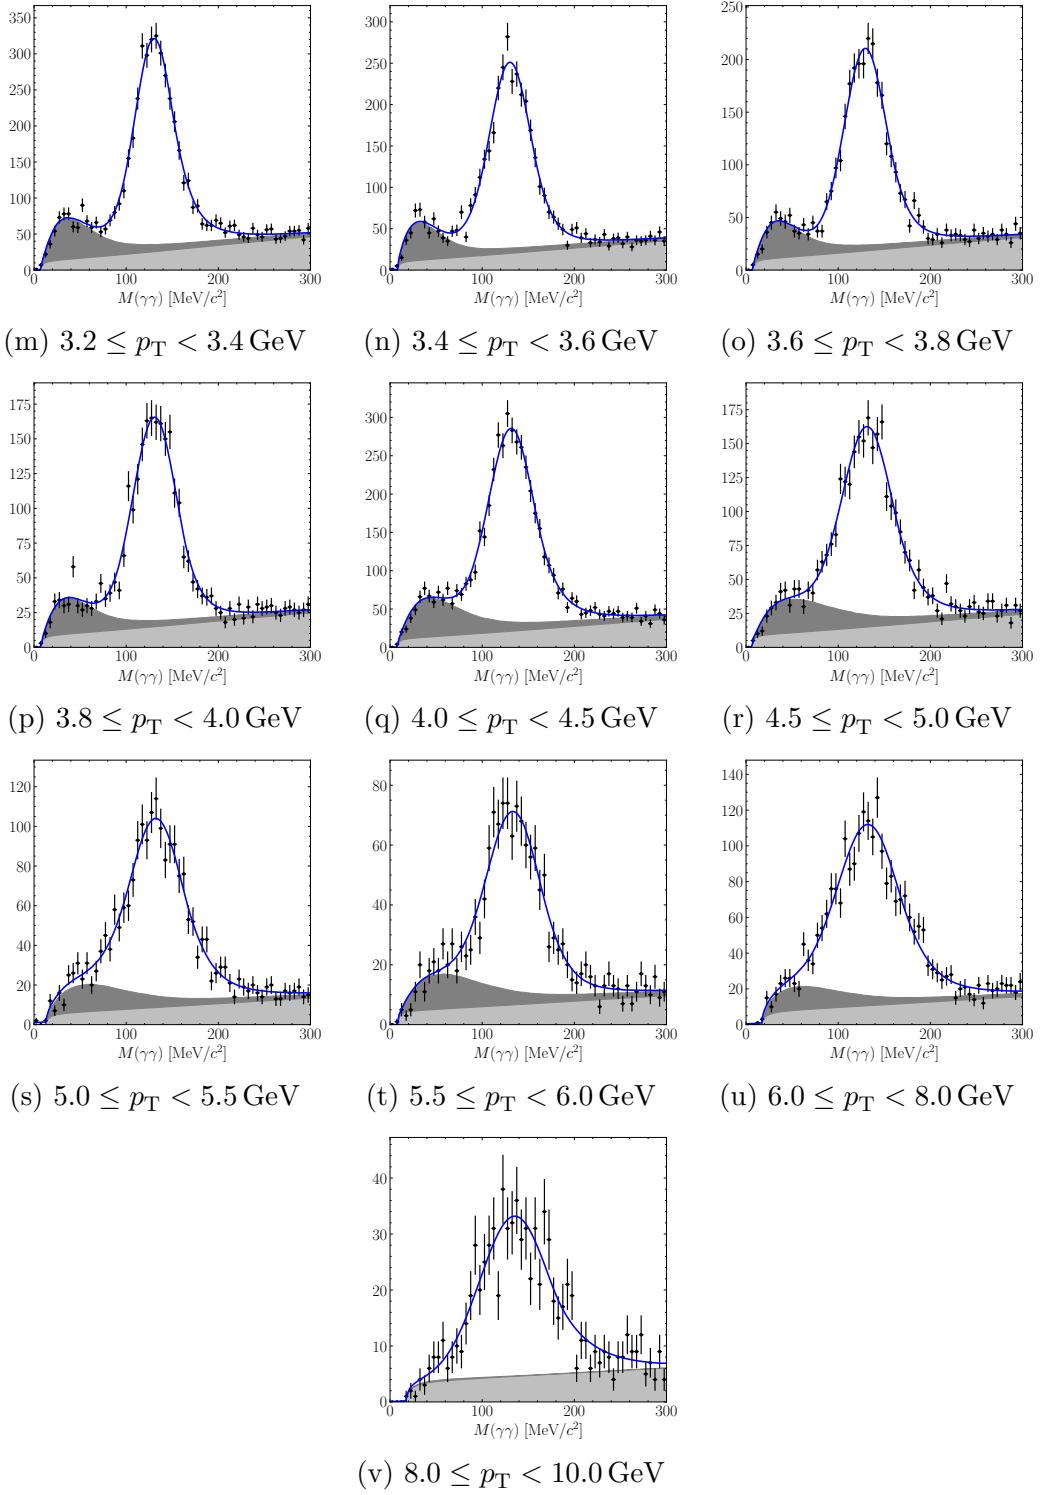


Figure B-3: π^0 mass fits in 13 TeV pp data as a function of $p_T(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$ (cont.).

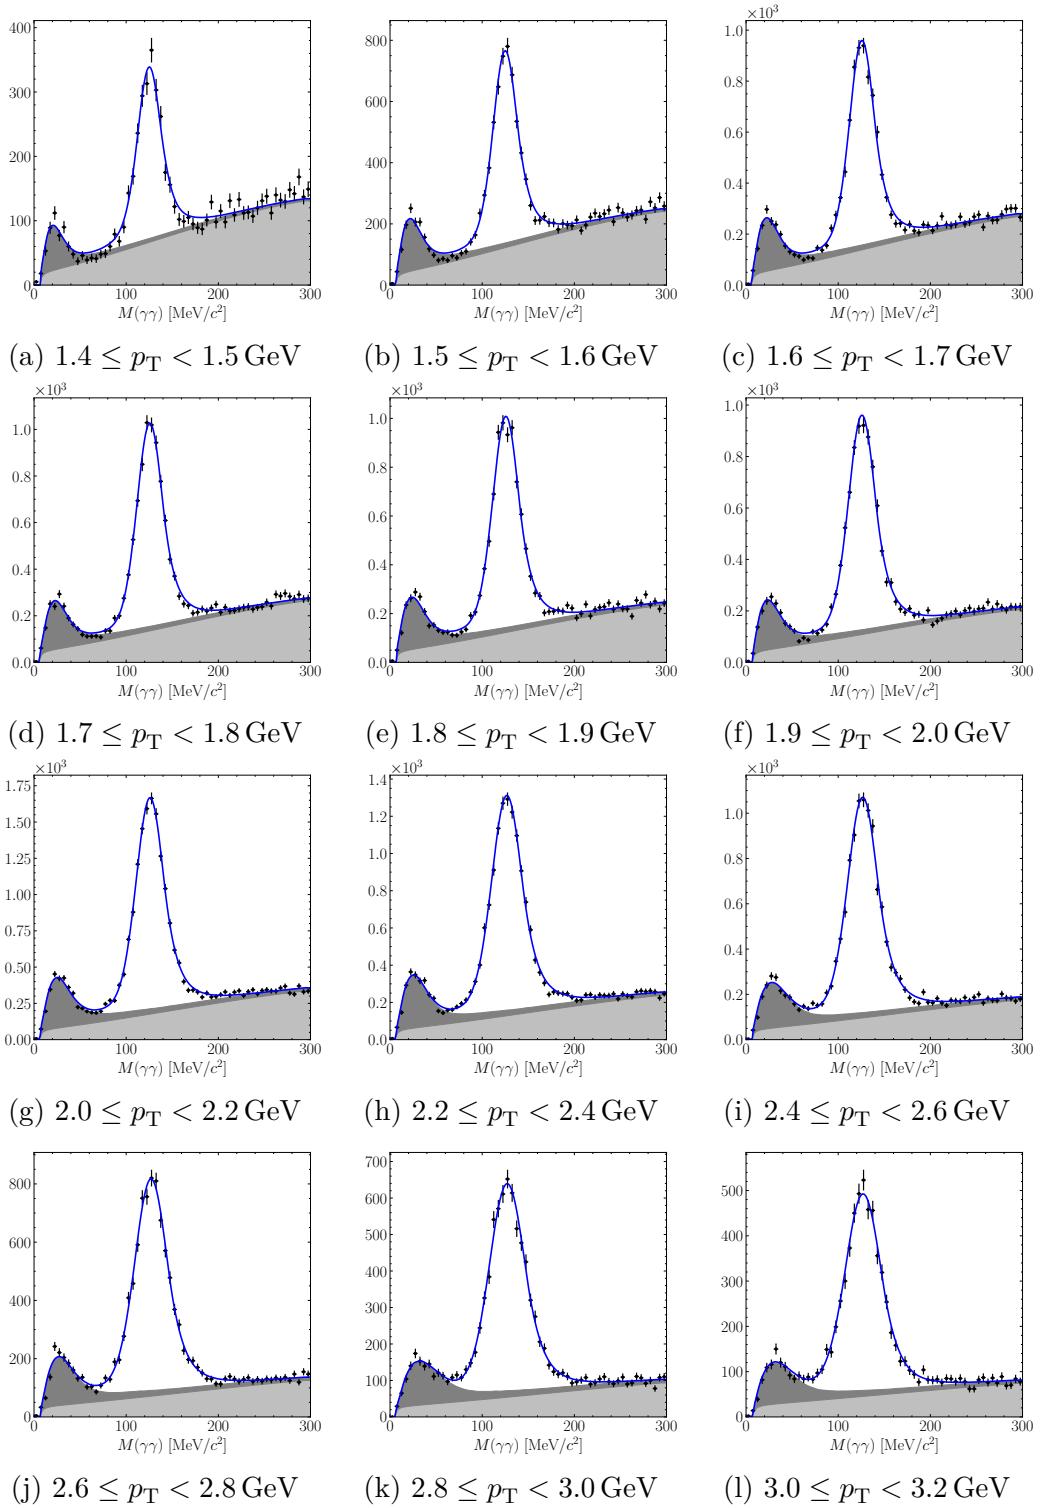


Figure B-4: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$.

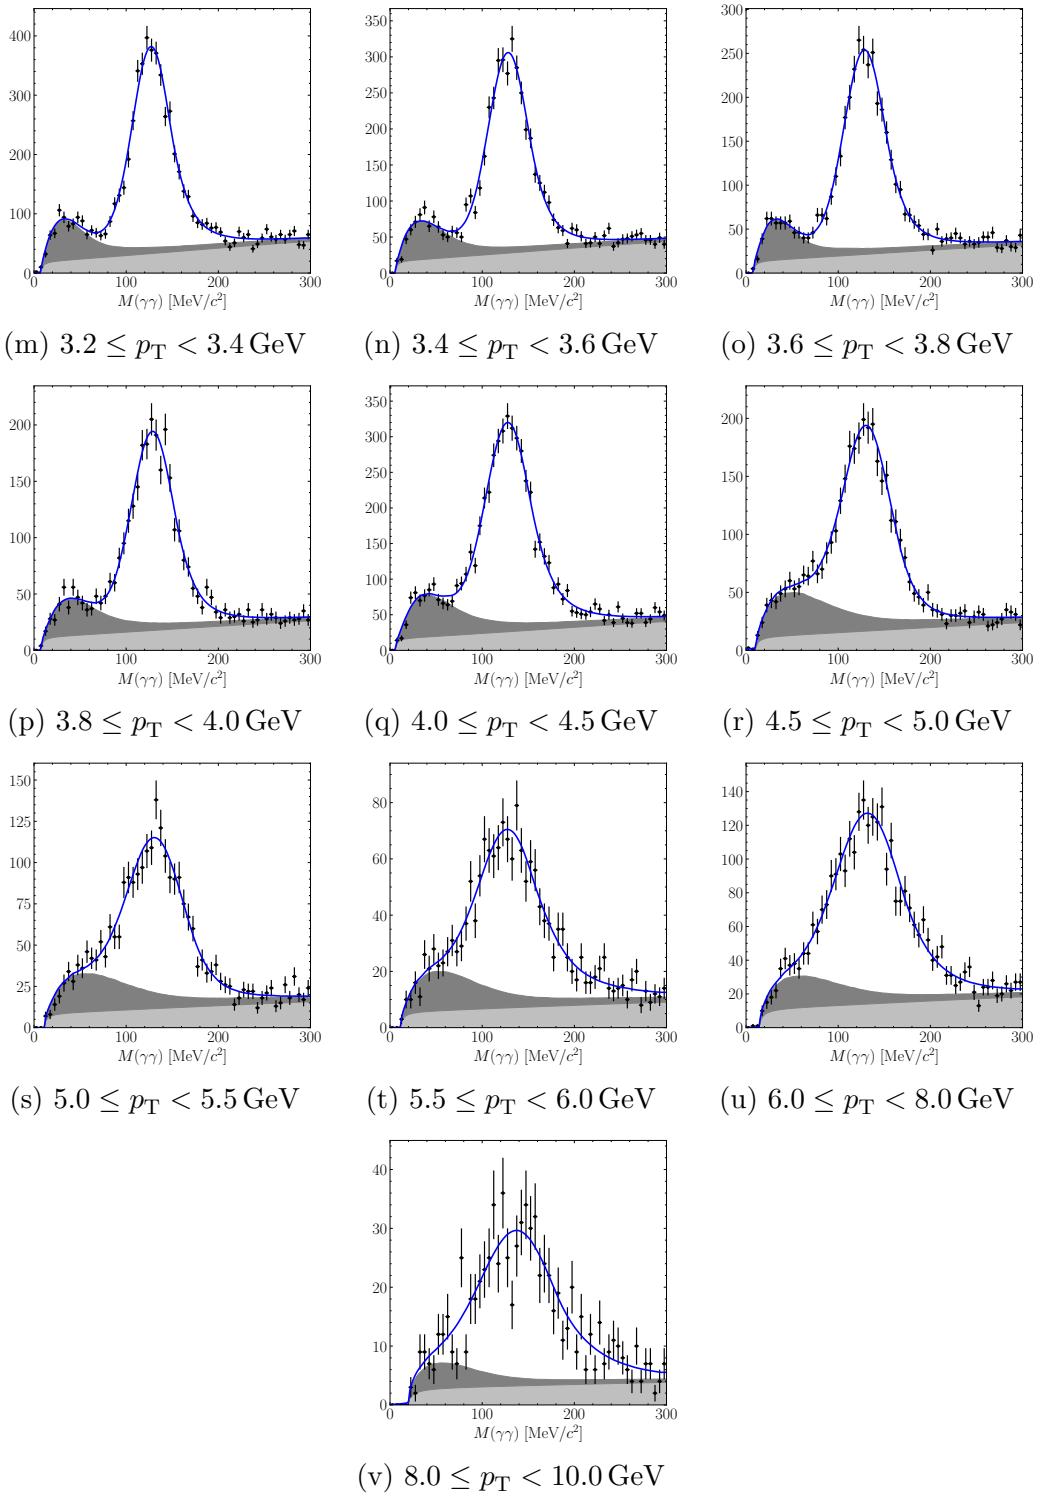


Figure B-4: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).

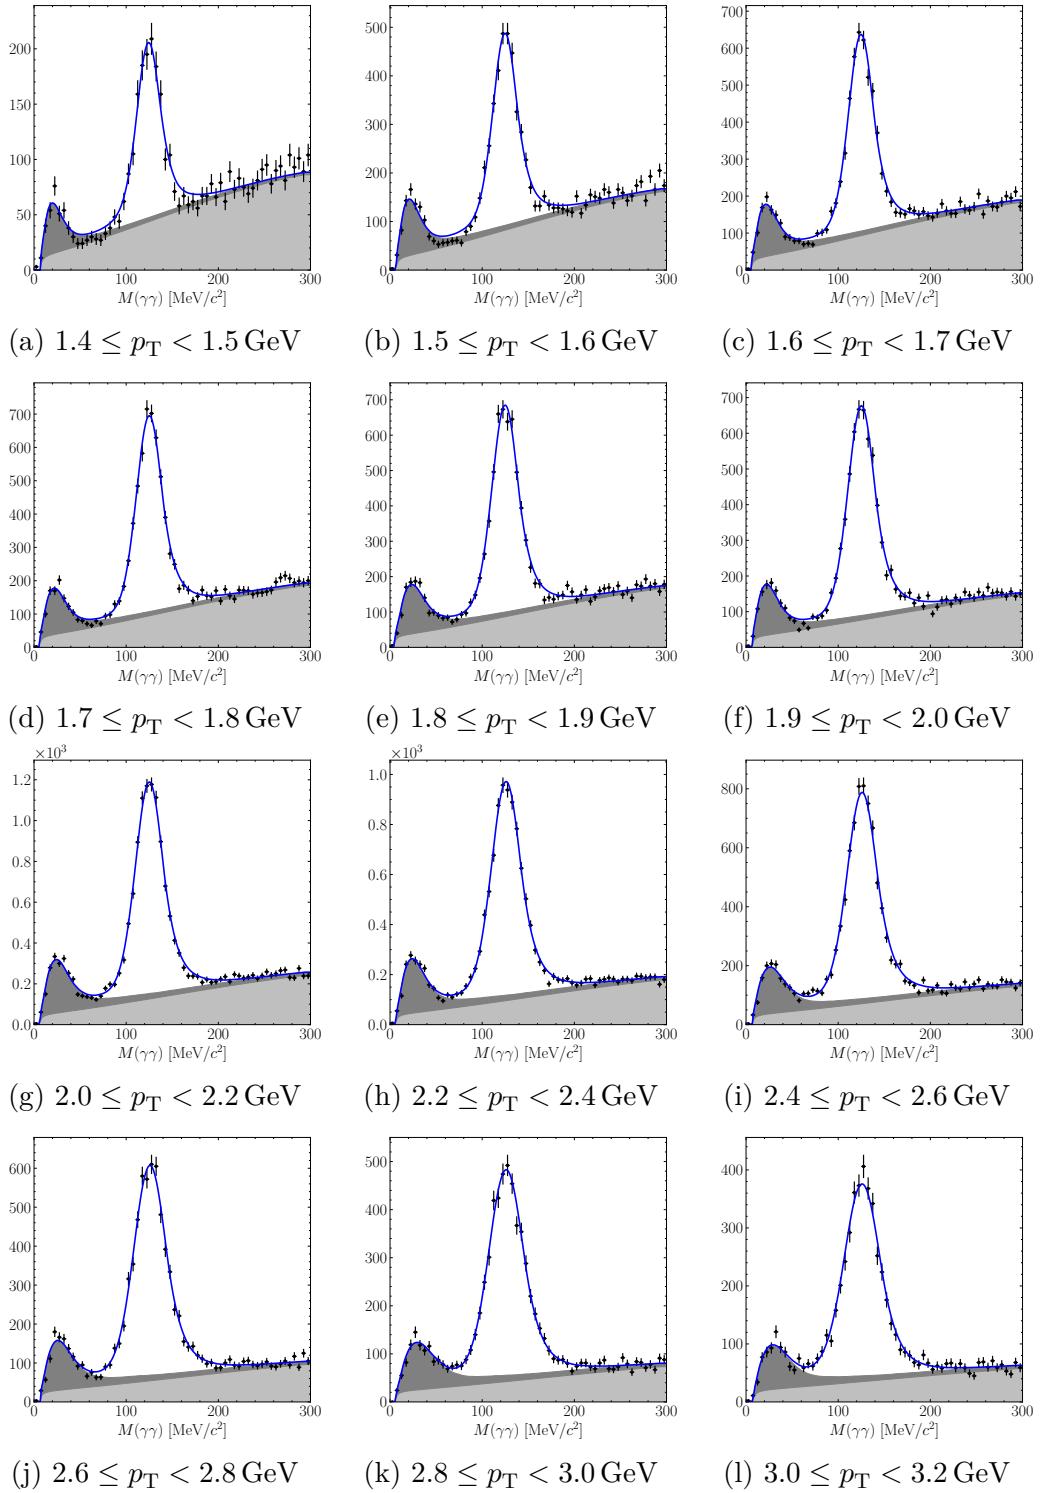


Figure B-5: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$.

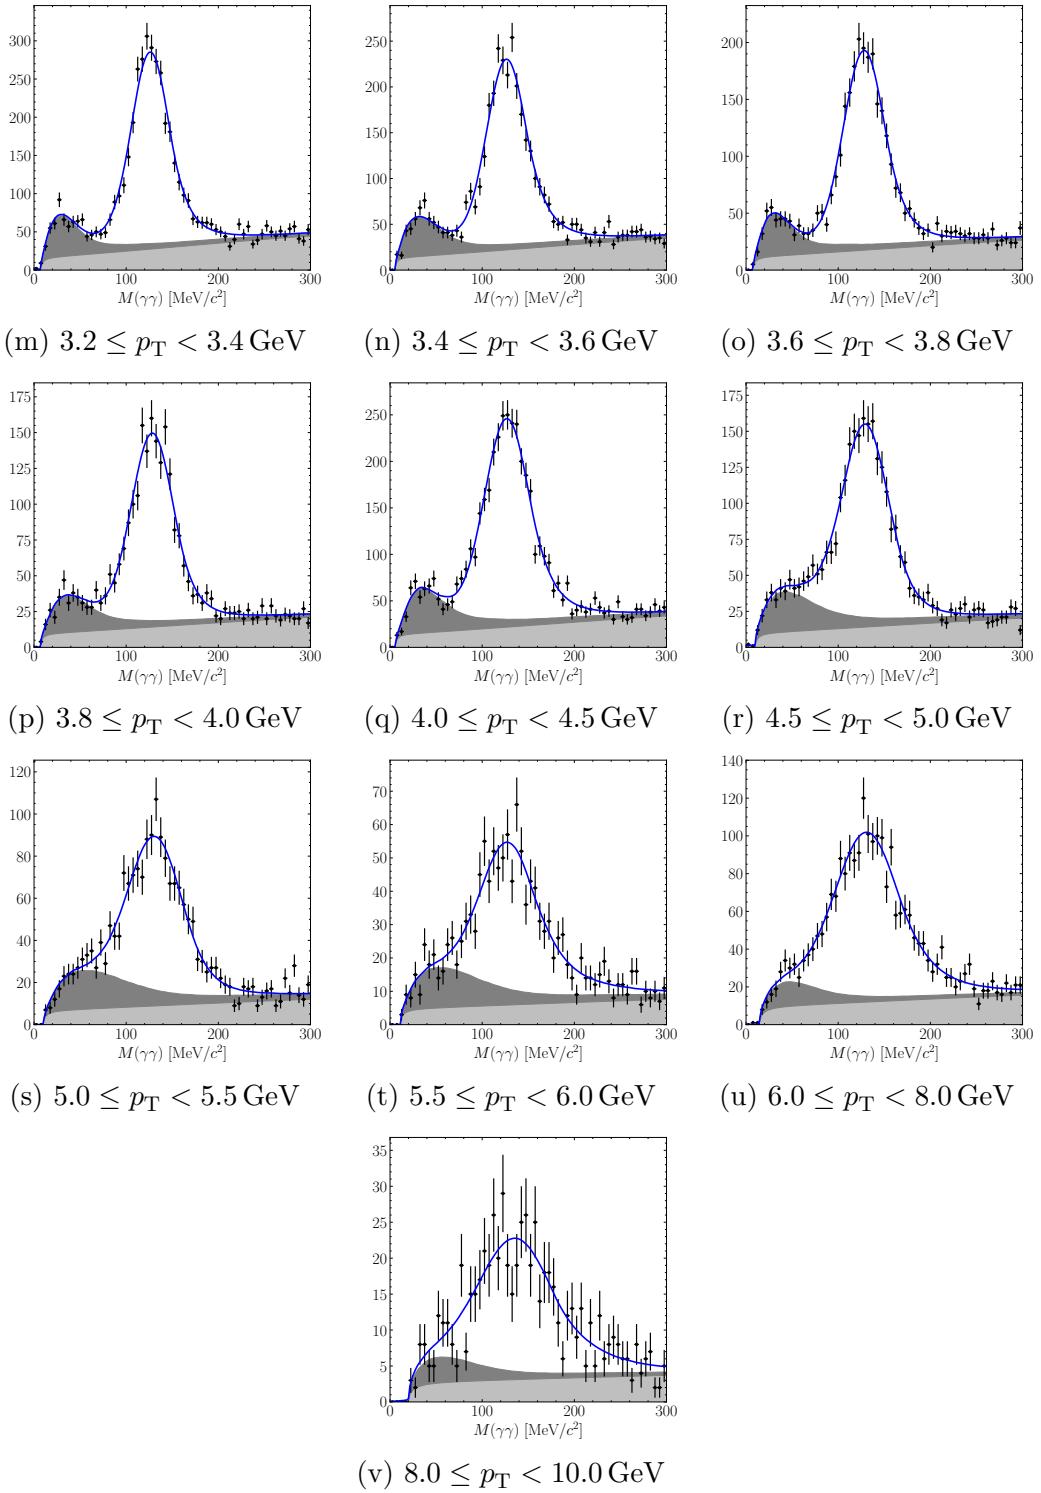


Figure B-5: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$ (cont.).

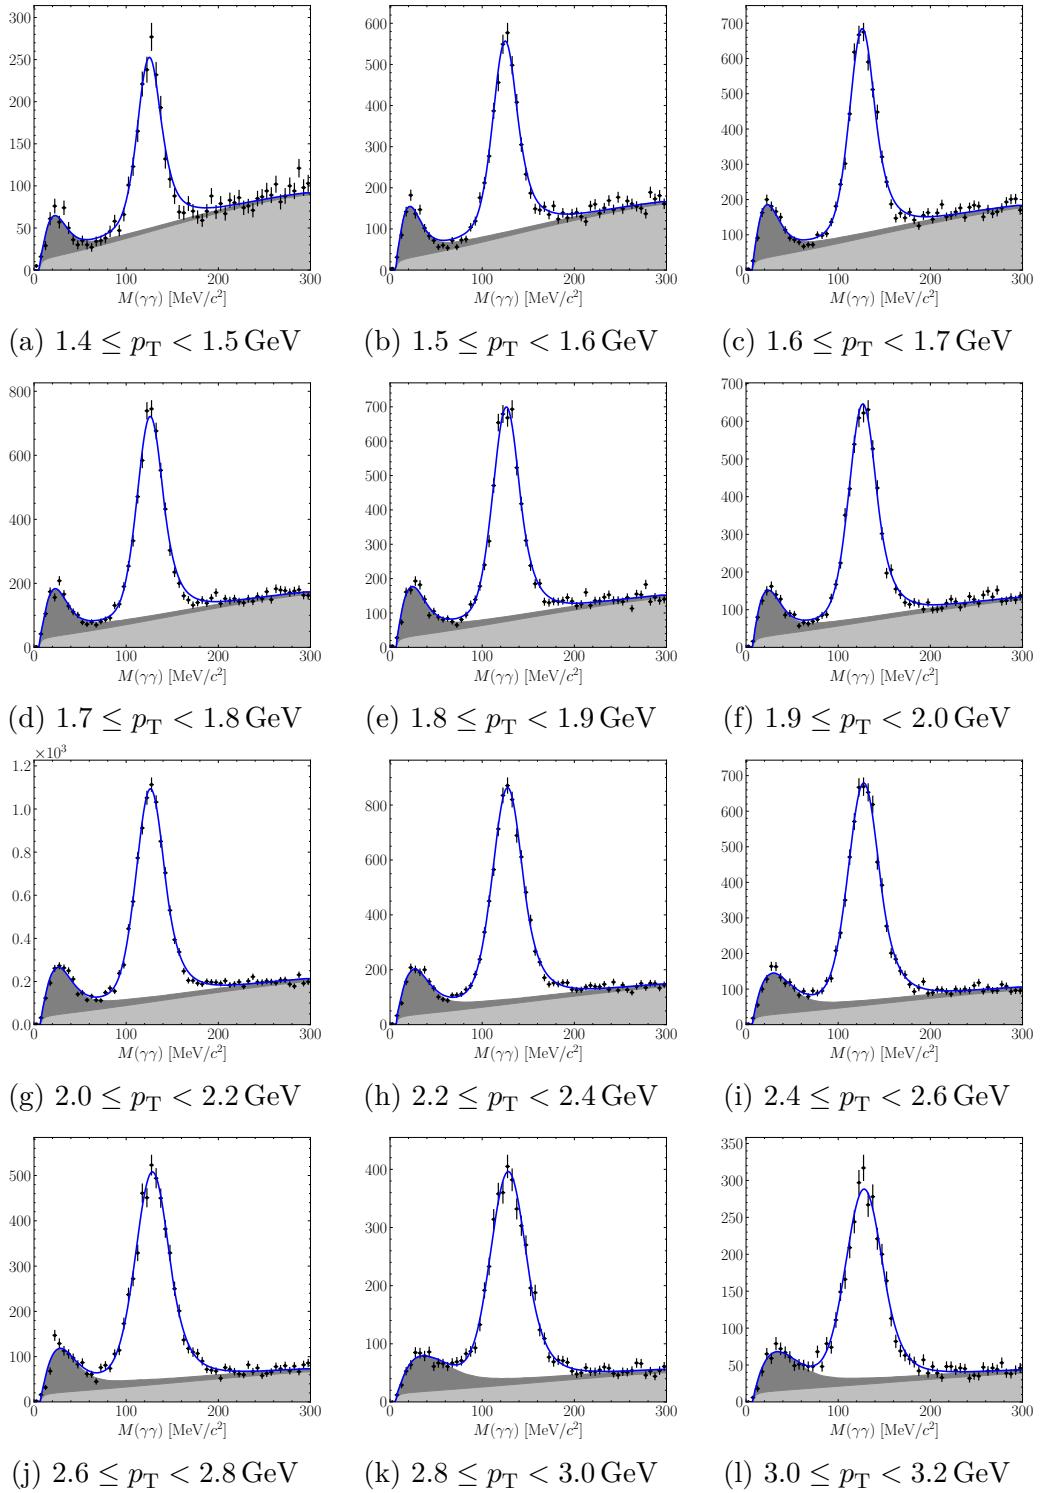


Figure B-6: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$.

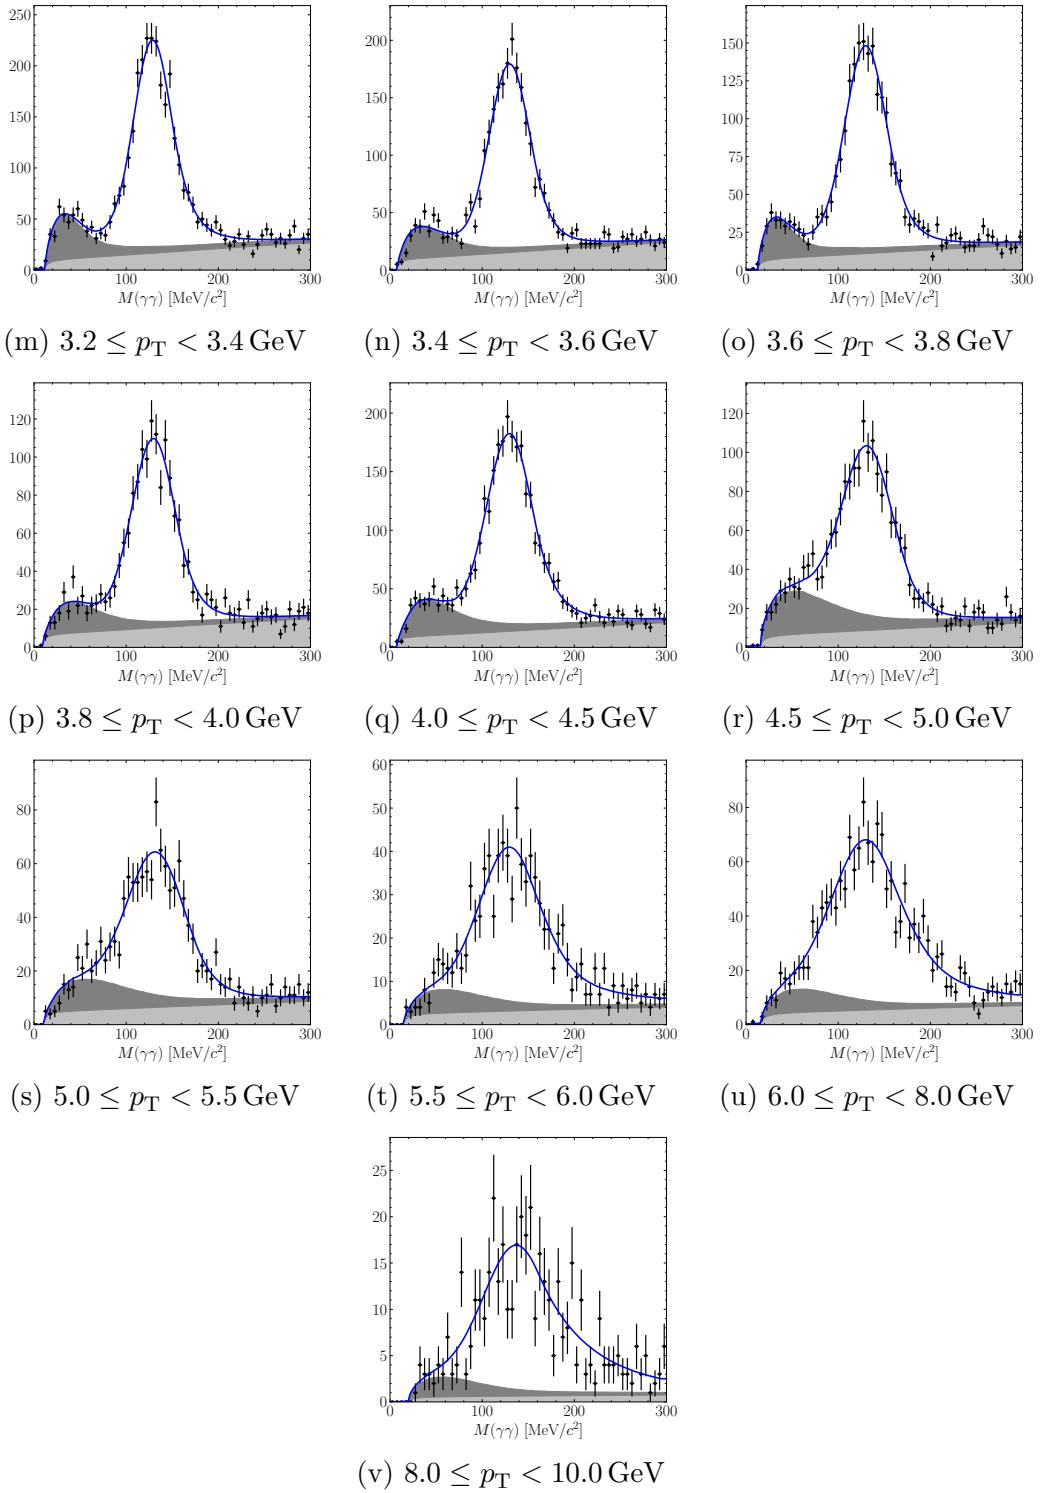


Figure B-6: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\pi^0)$ for $3.0 < \eta(\pi^0) < 4.0$ (cont.).

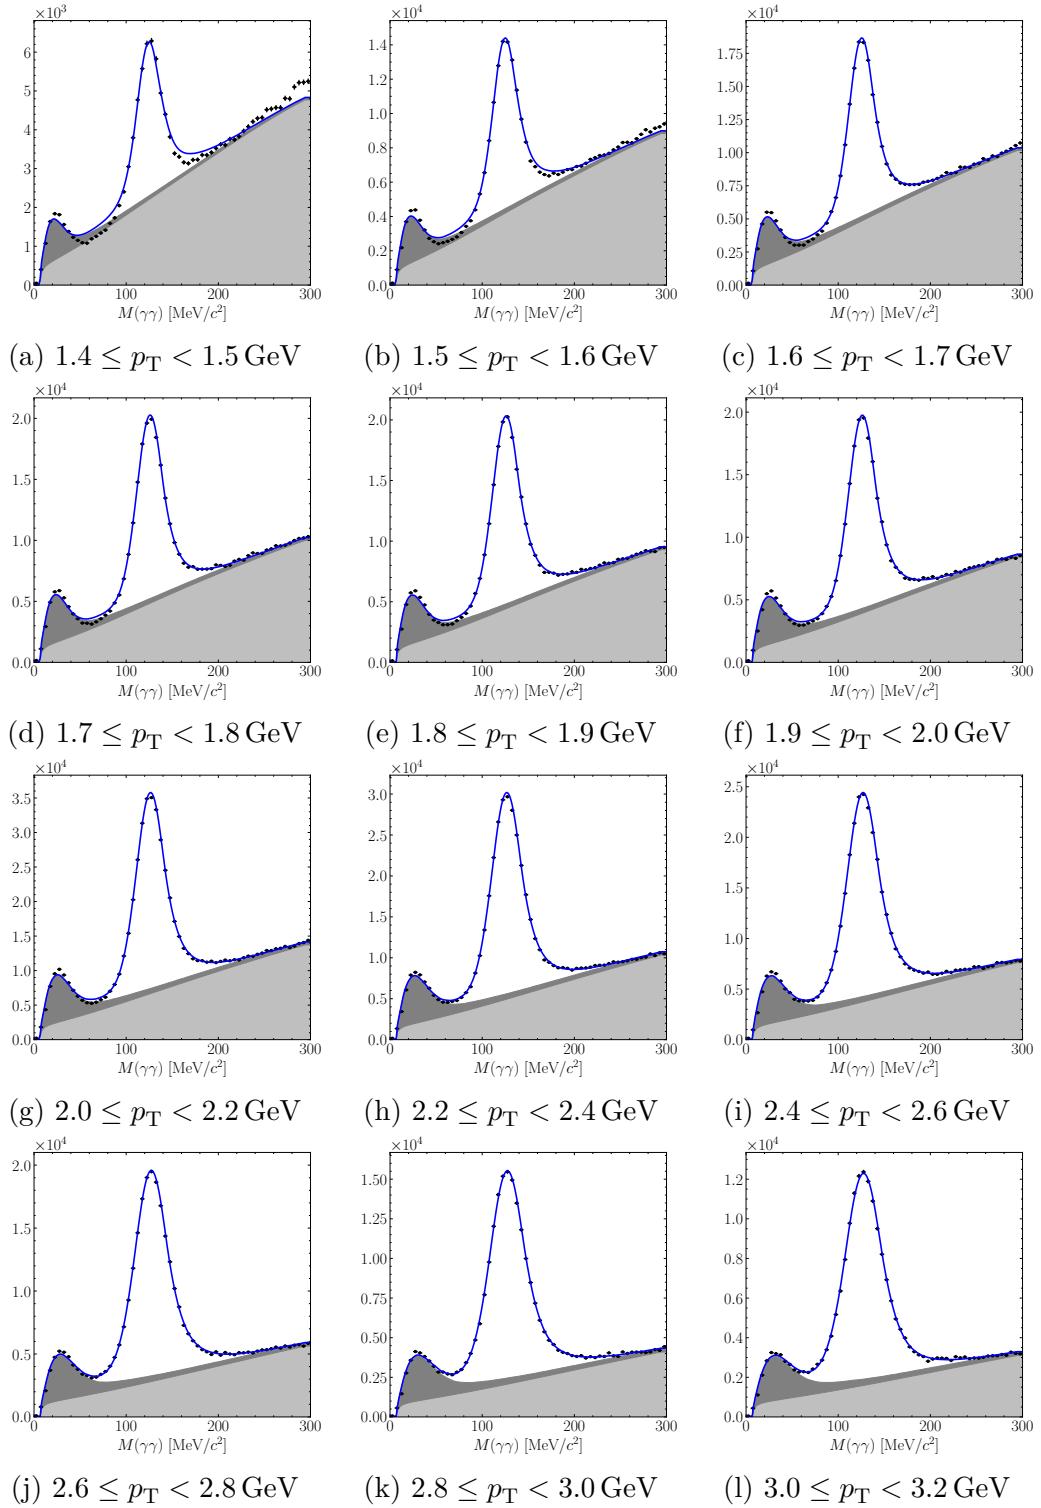


Figure B-7: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$.

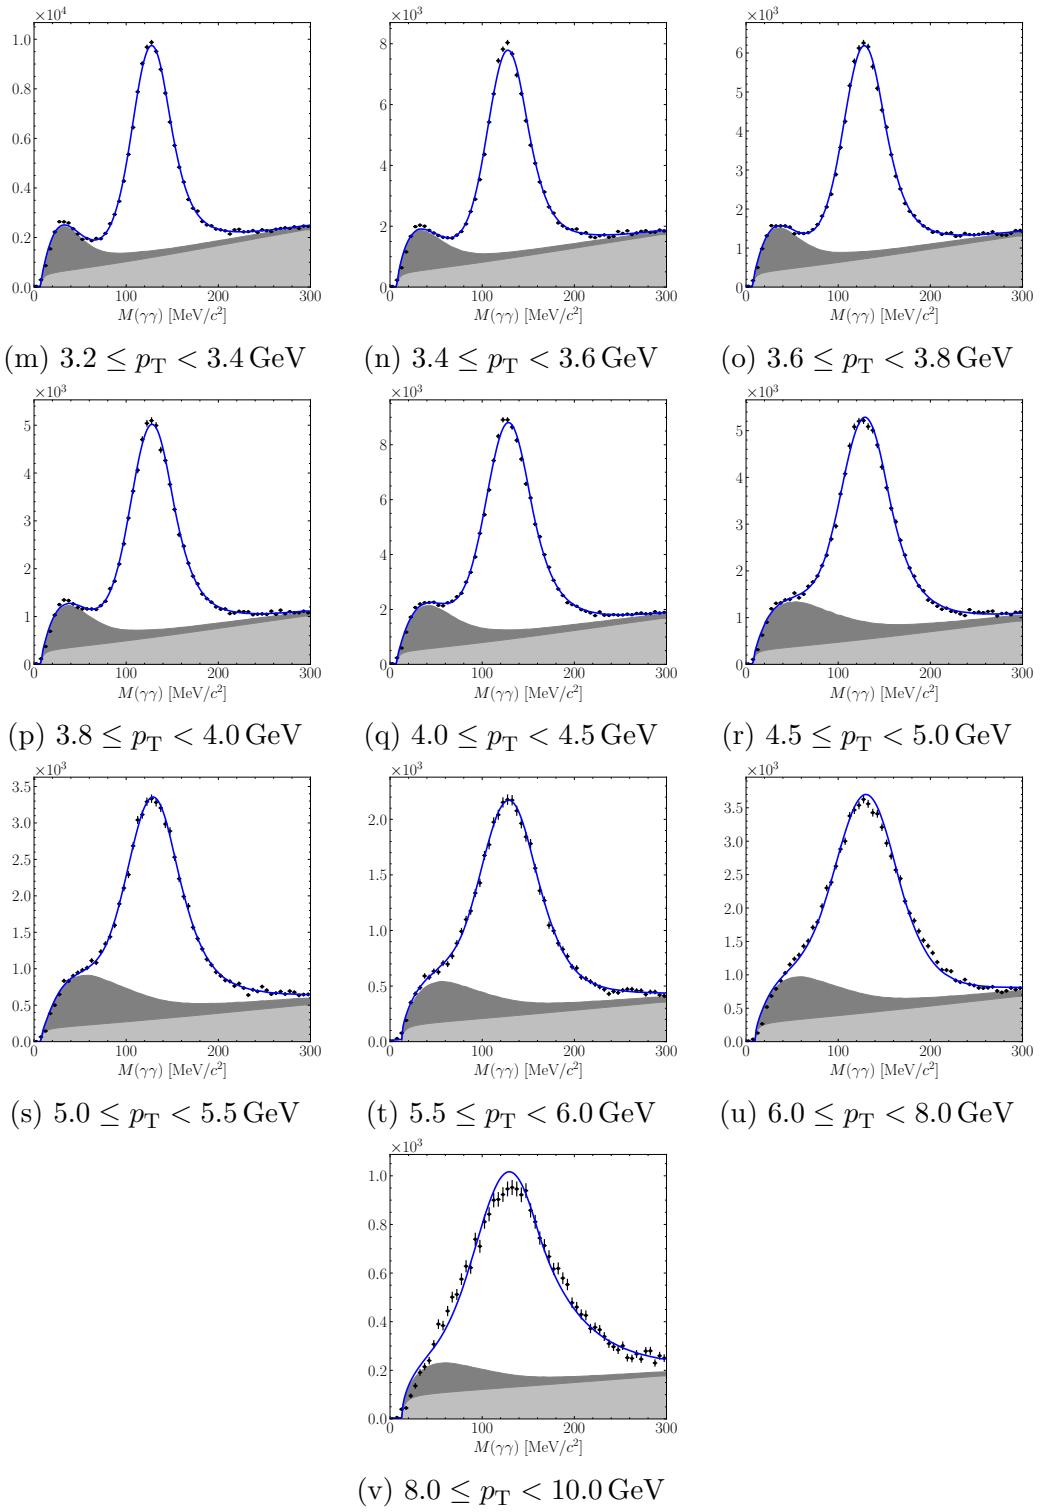


Figure B-7: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).

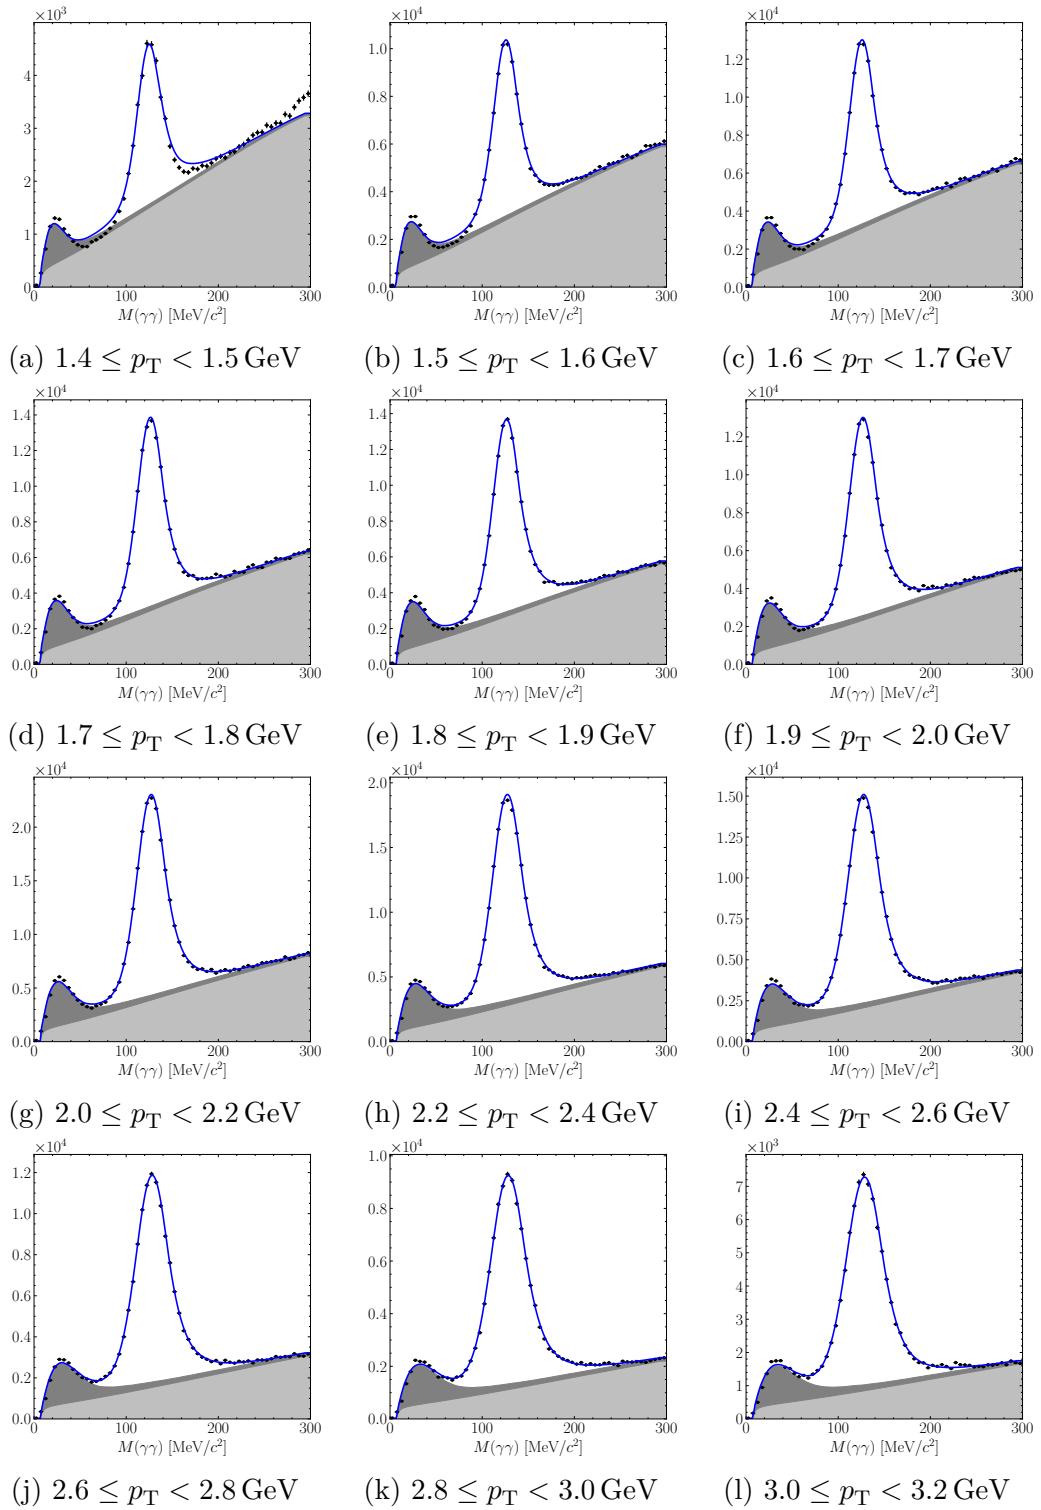


Figure B-8: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$.

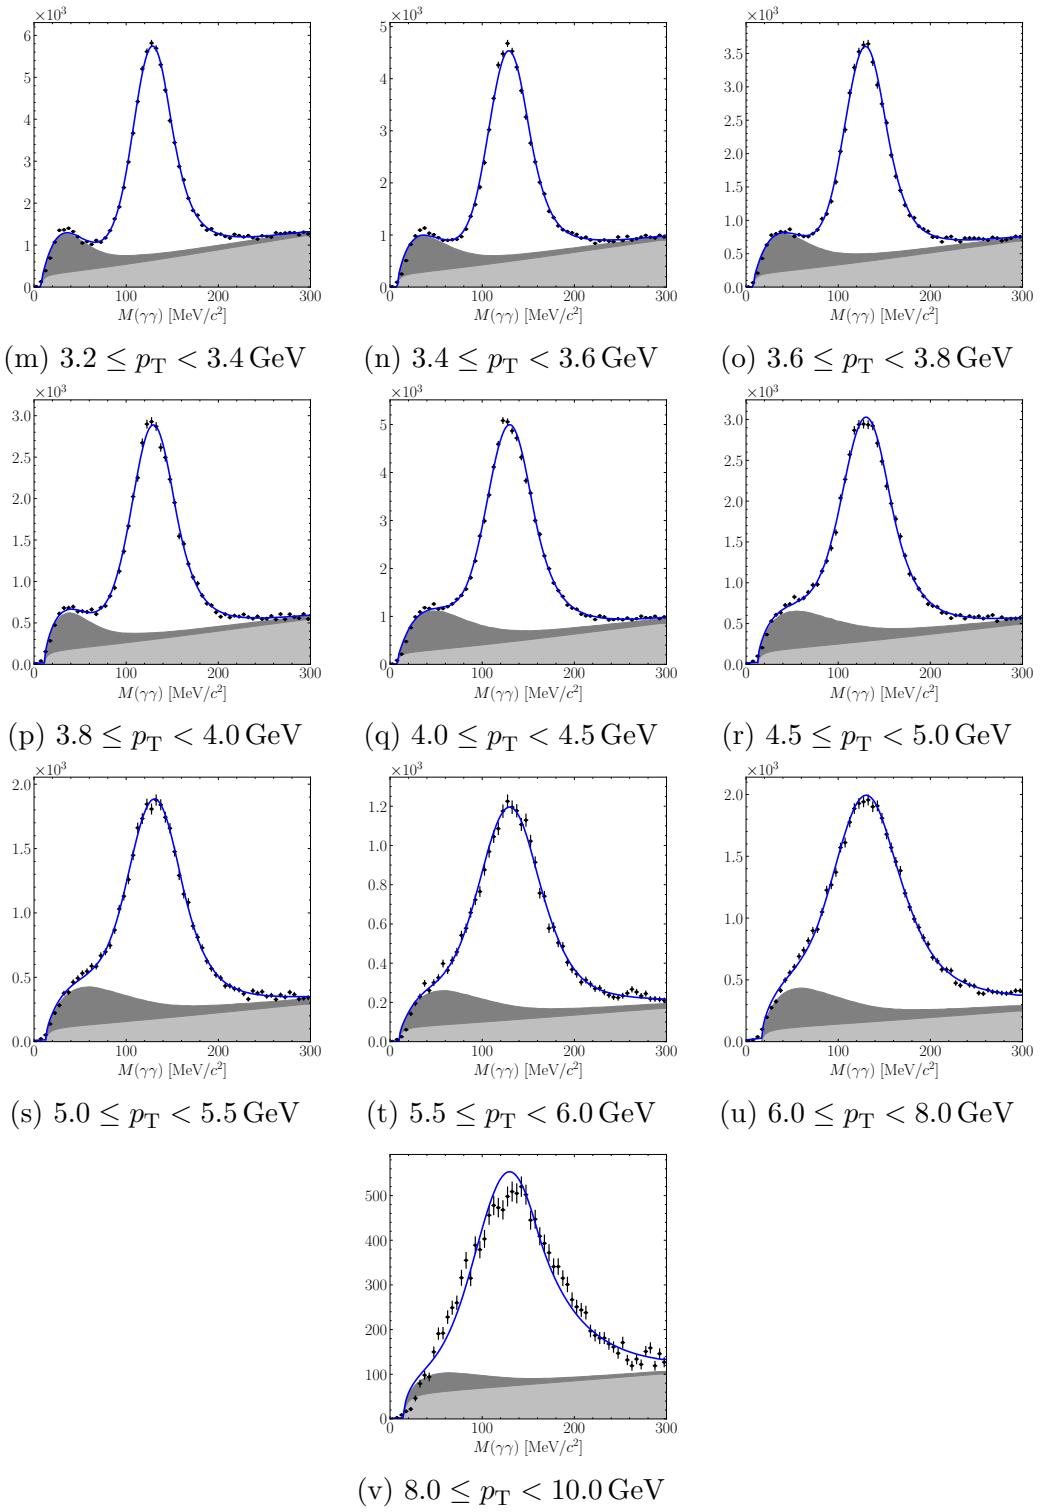


Figure B-8: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < y^*(\pi^0) < 3.5$ (cont.).

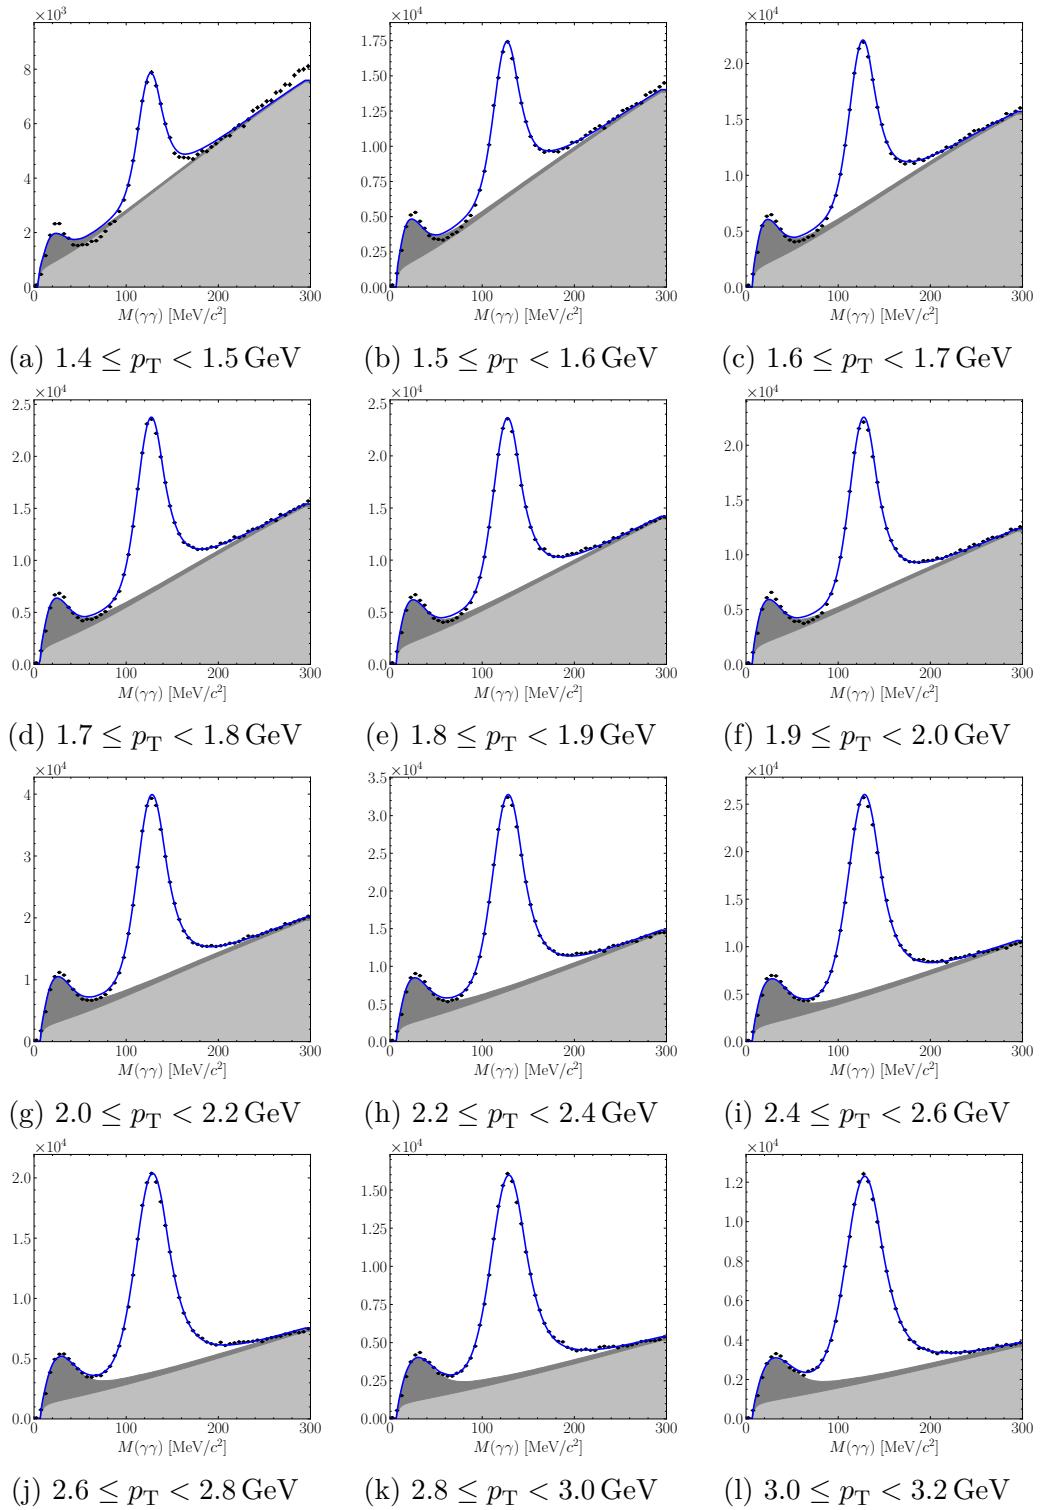


Figure B-9: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$.

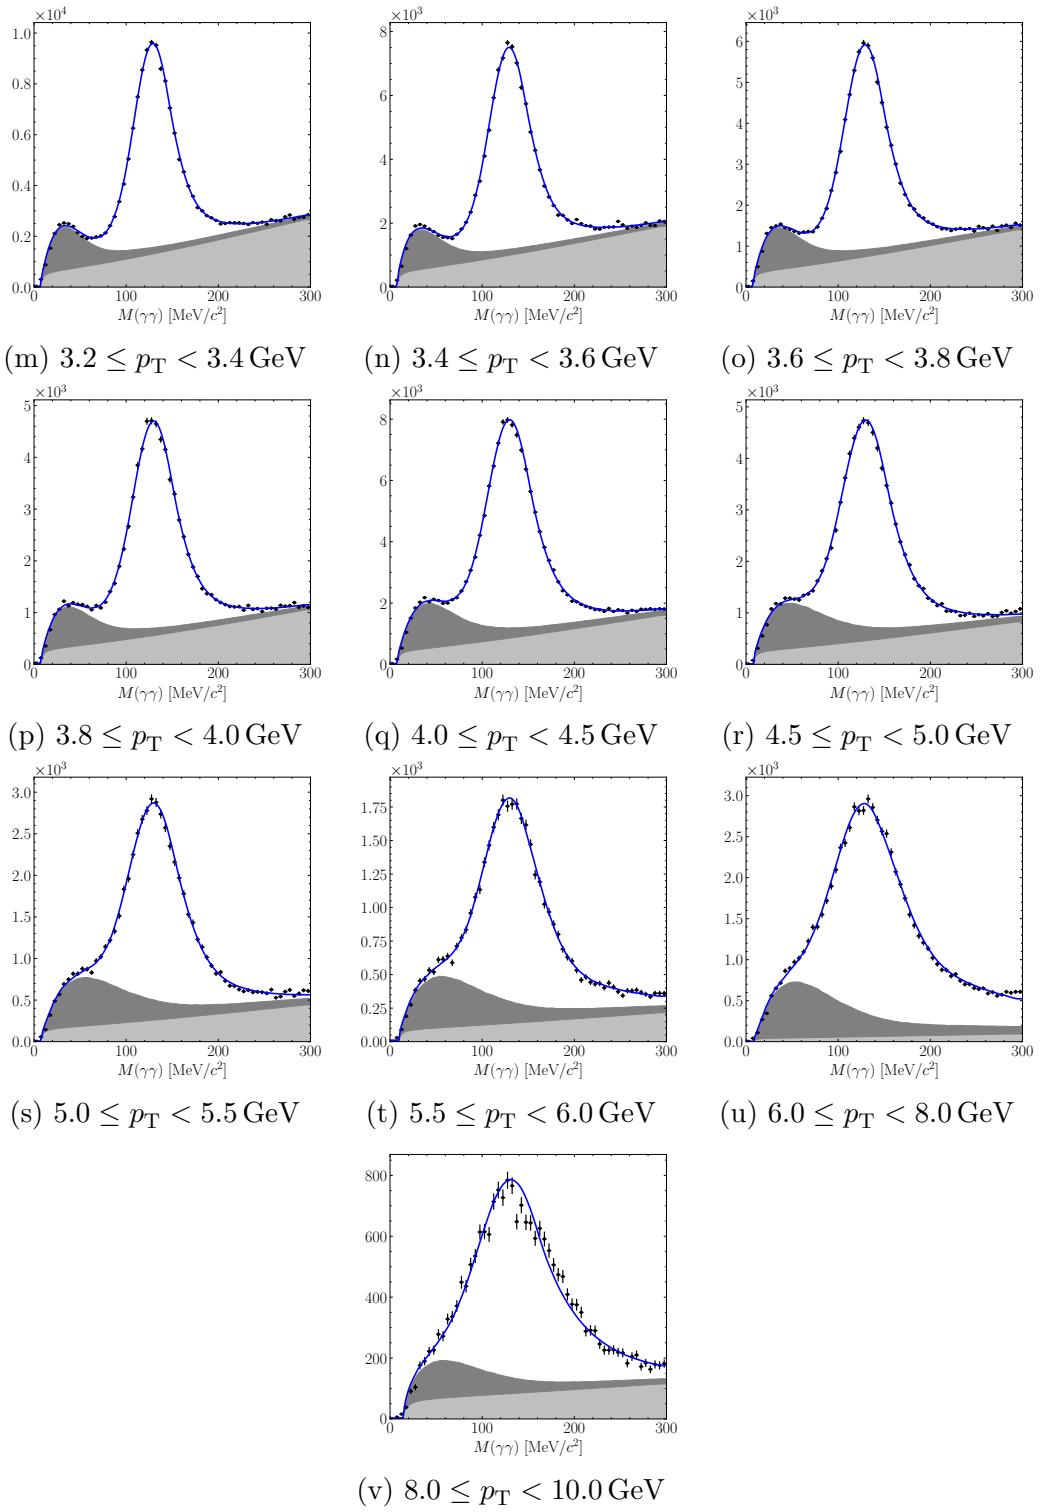


Figure B-9: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 4.0$ (cont.).

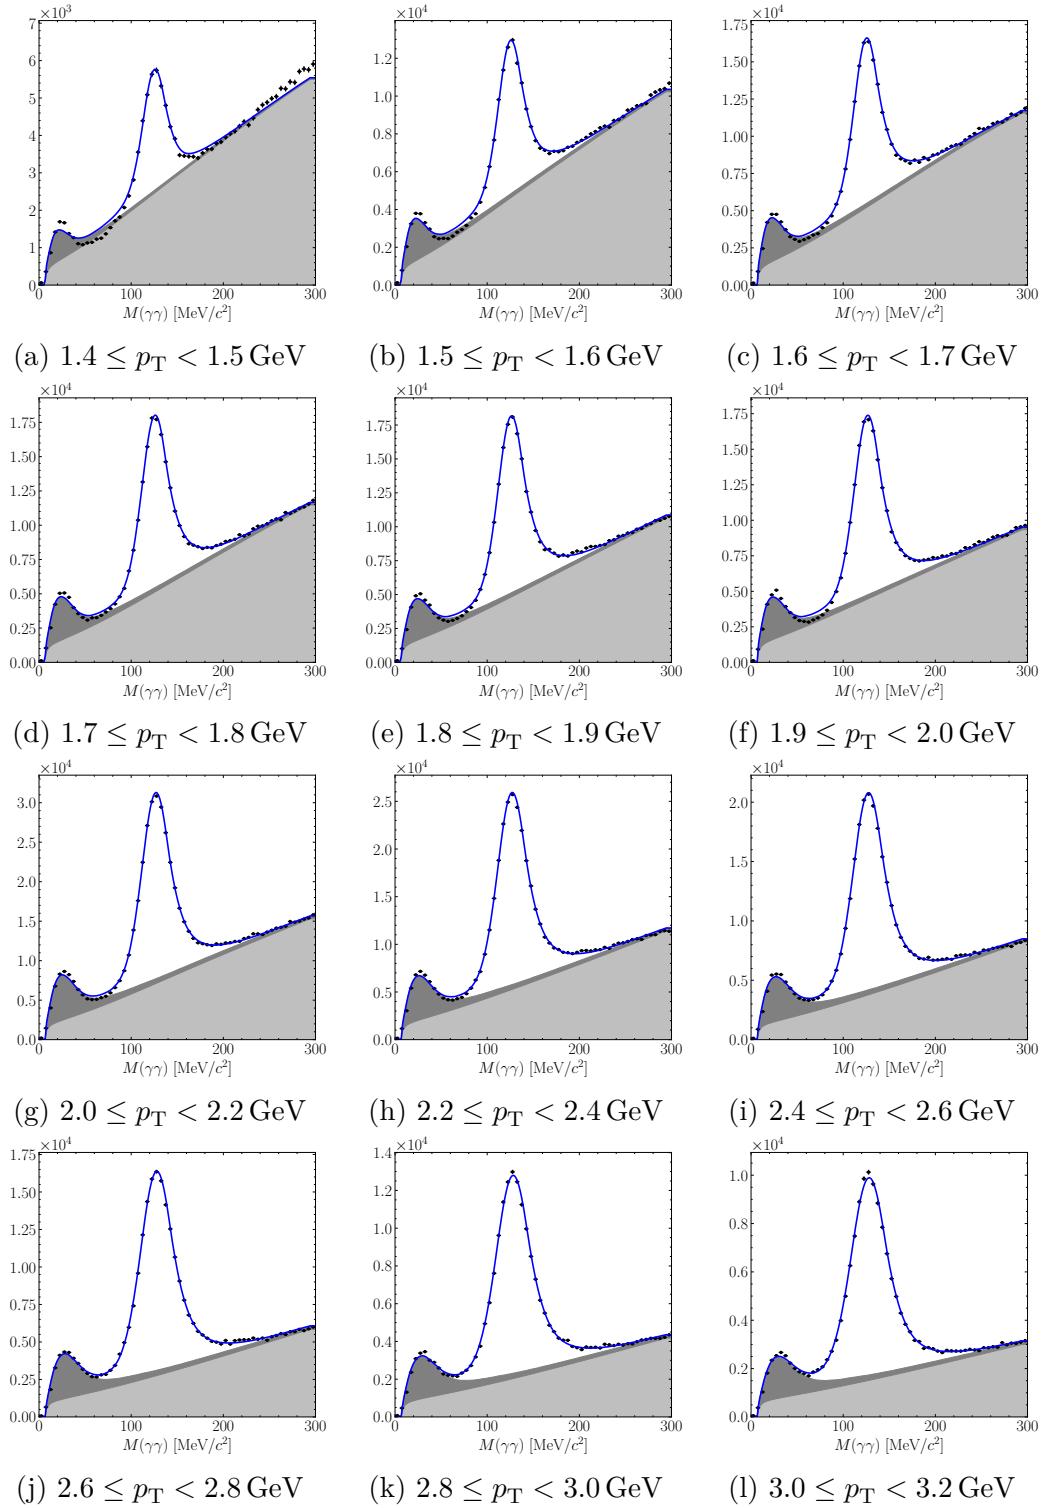


Figure B-10: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $2.5 < \eta(\pi^0) < 3.5$.

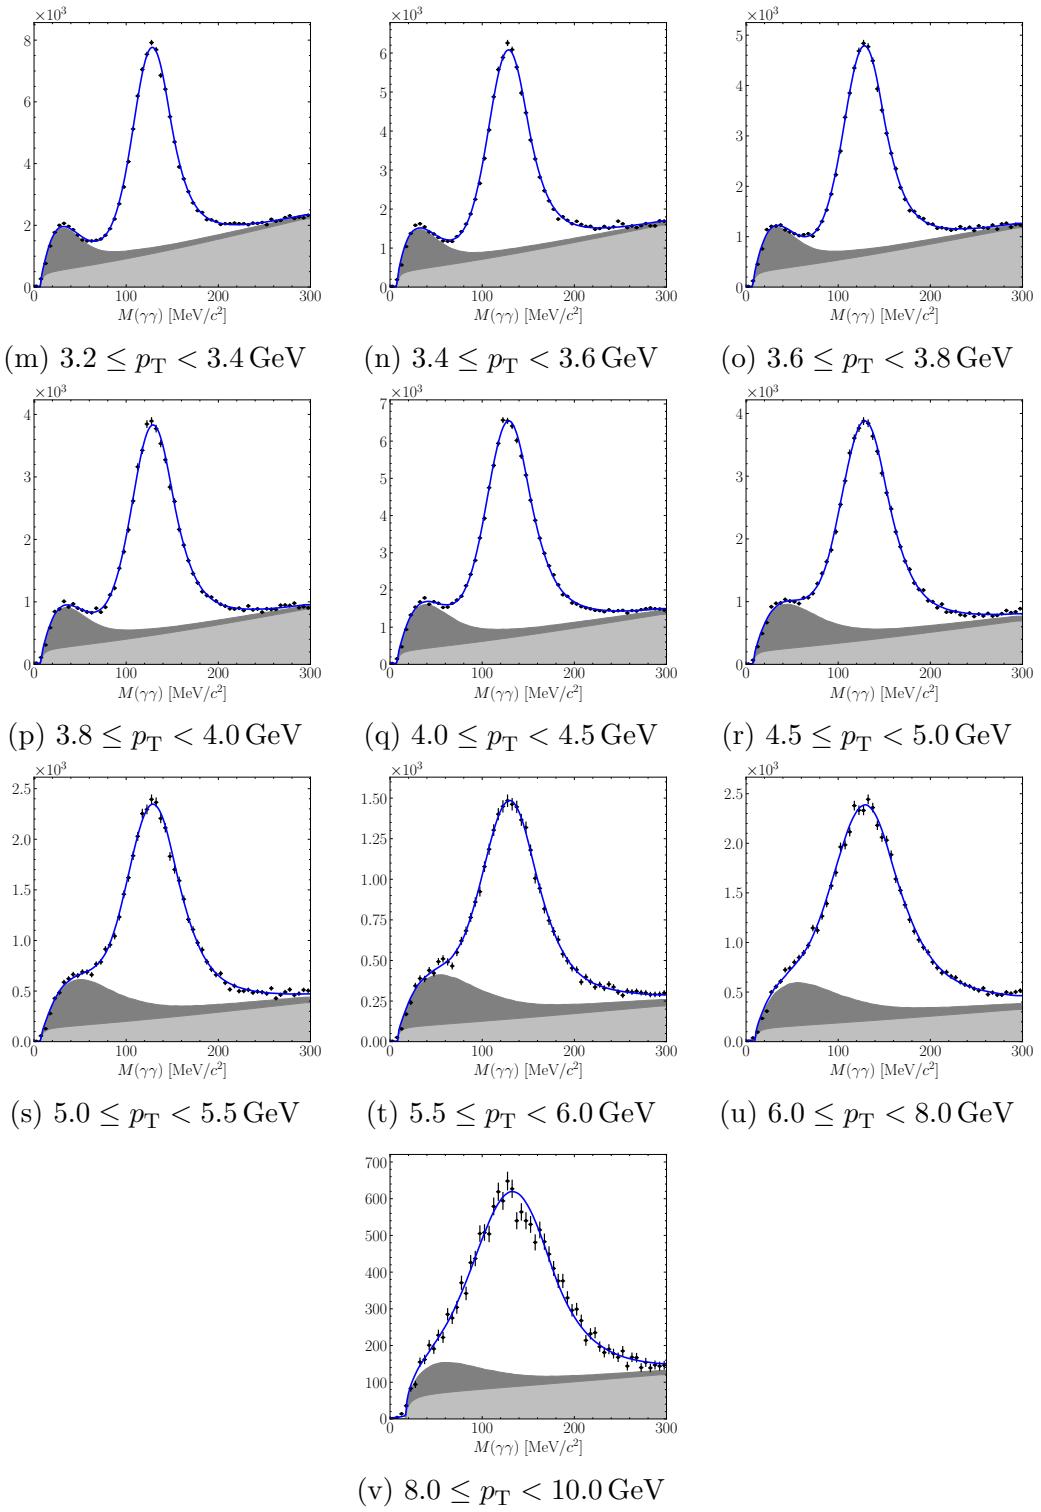


Figure B-10: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\pi^0)$ for $-4.0 < y^*(\pi^0) < -3.0$ (cont.).

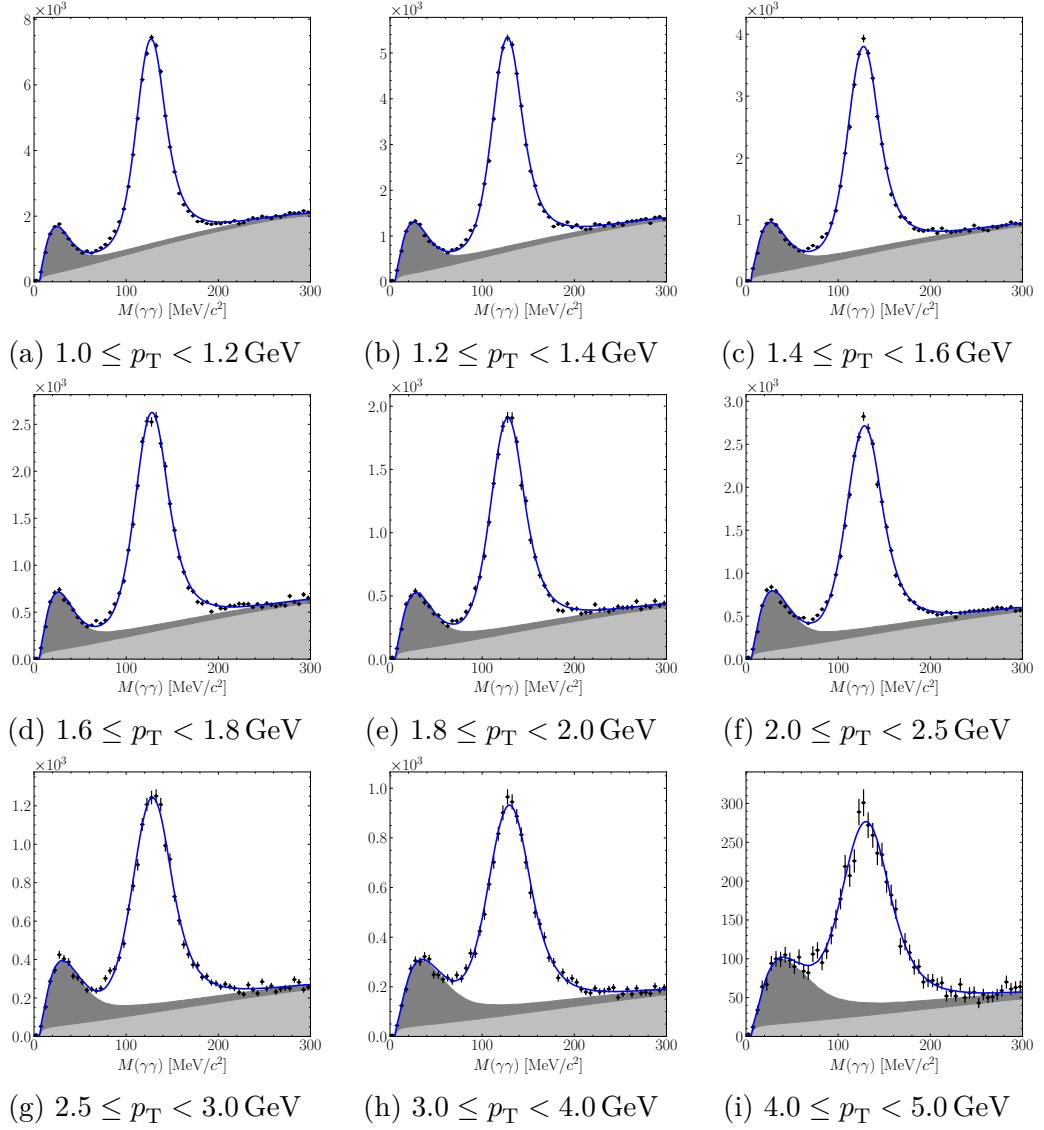


Figure B-11: π^0 mass fits in 13 TeV pp data as a function of $p_T(\gamma^{\text{cnv}})$.

B.2 Fits for R^γ

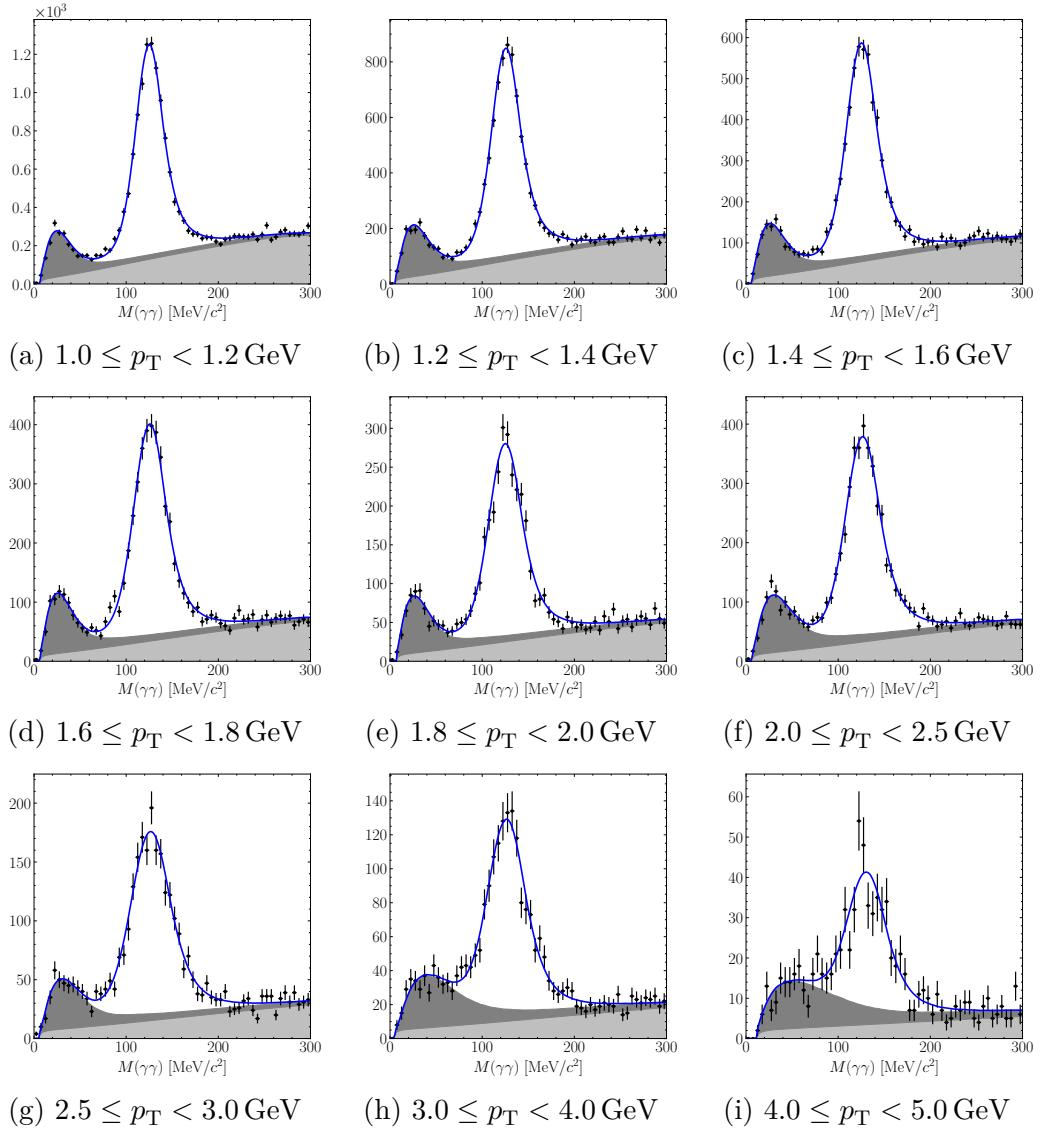


Figure B-12: π^0 mass fits in 5.02 TeV pp data as a function of $p_T(\gamma^{\text{cnv}})$.

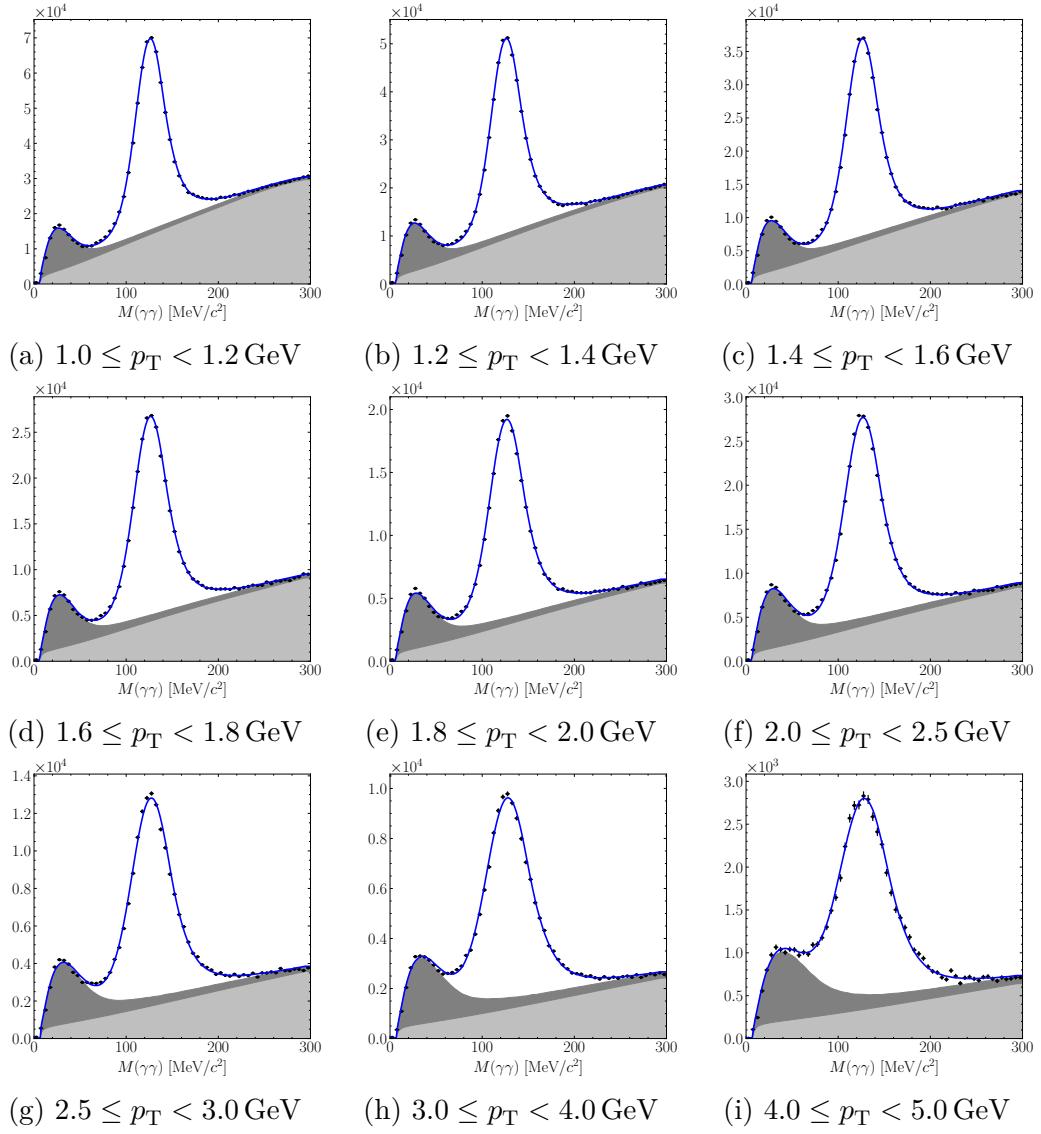


Figure B-13: π^0 mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\gamma^{\text{conv}})$.

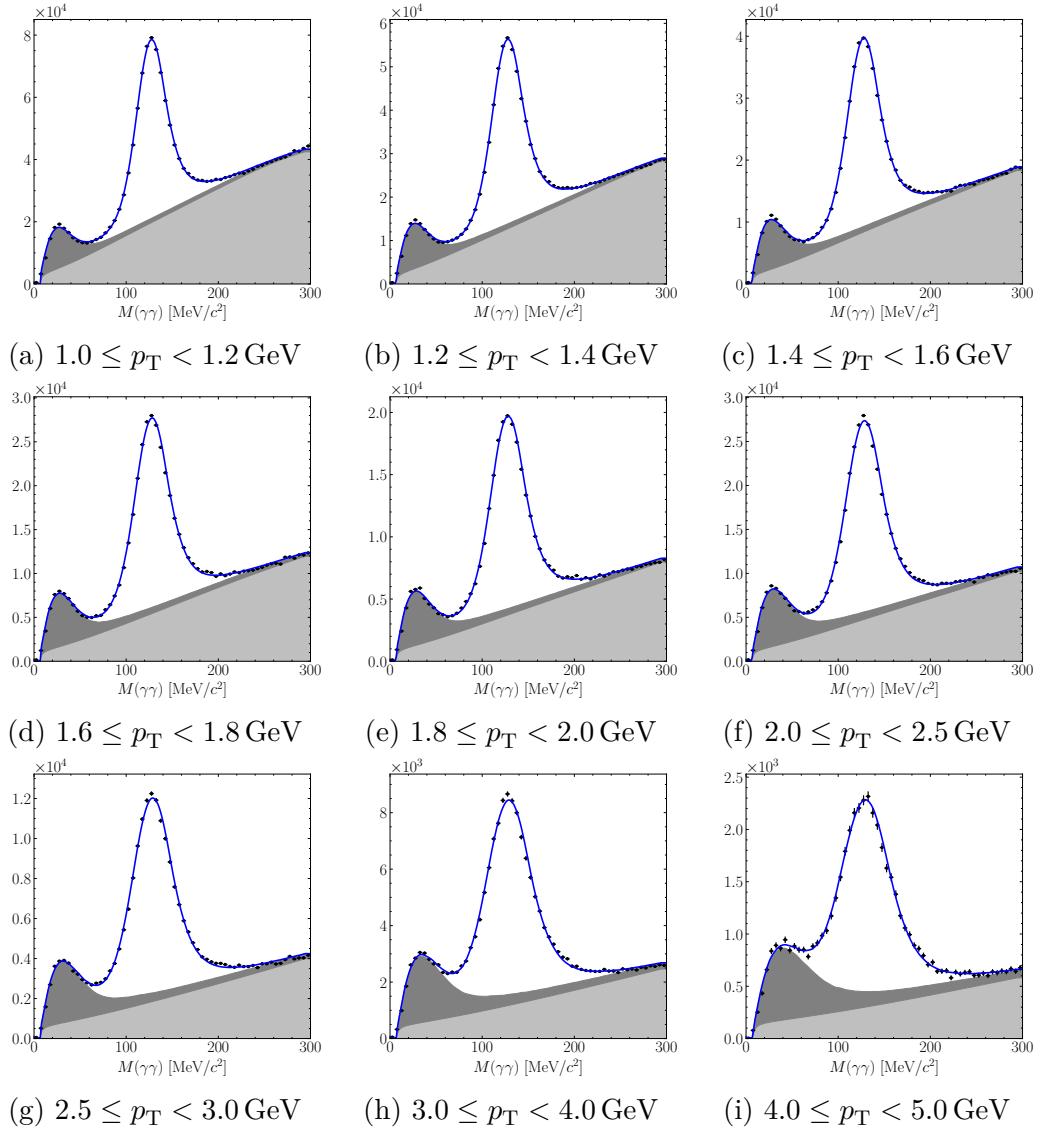


Figure B-14: π^0 mass fits in 8.16 TeV Pbp data as a function of $p_T(\gamma^{\text{cnv}})$.

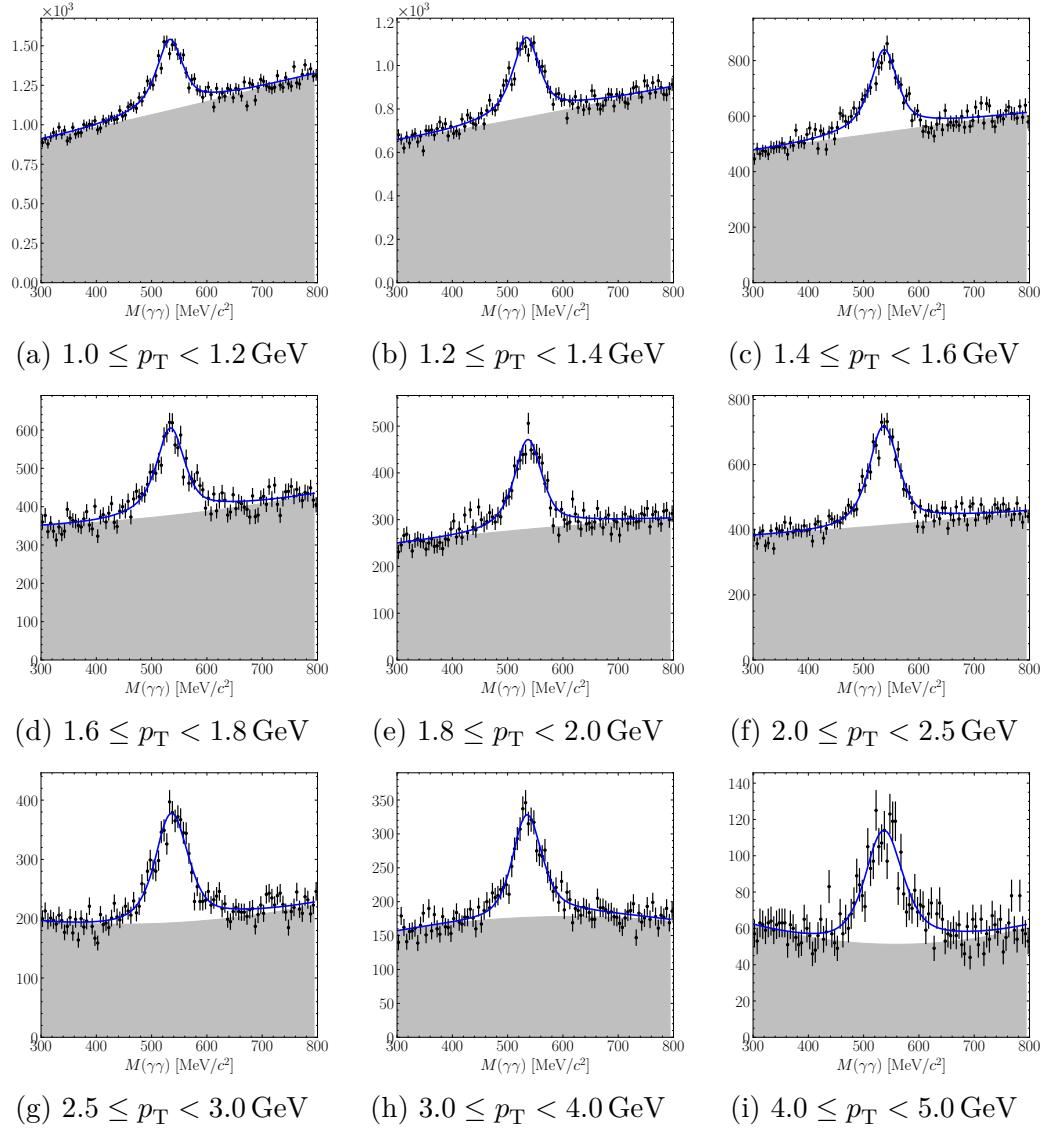


Figure B-15: η mass fits in 13 TeV pp data as a function of $p_T(\eta^{\text{cnv}})$.

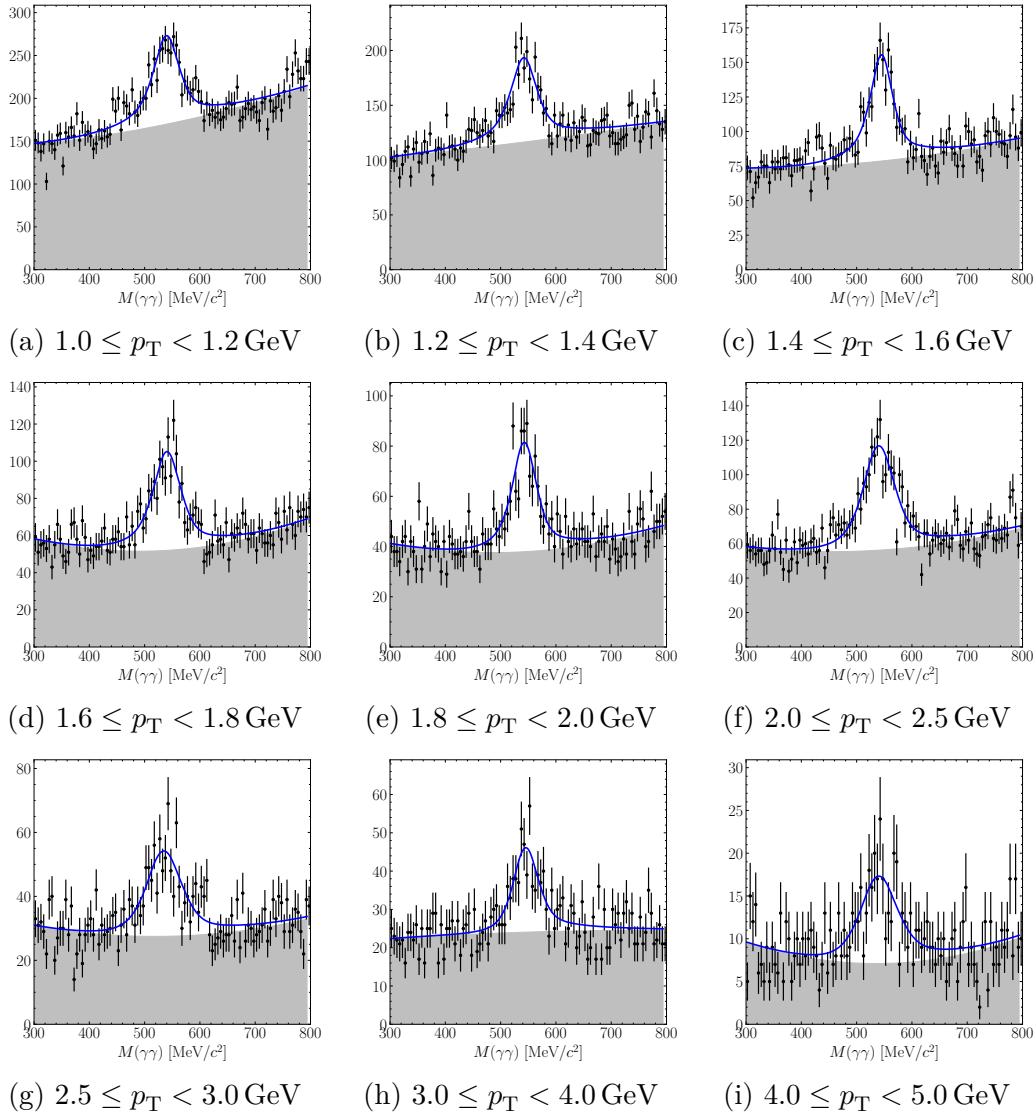


Figure B-16: η mass fits in 5.02 TeV pp data as a function of $p_T(\gamma^{\text{conv}})$.

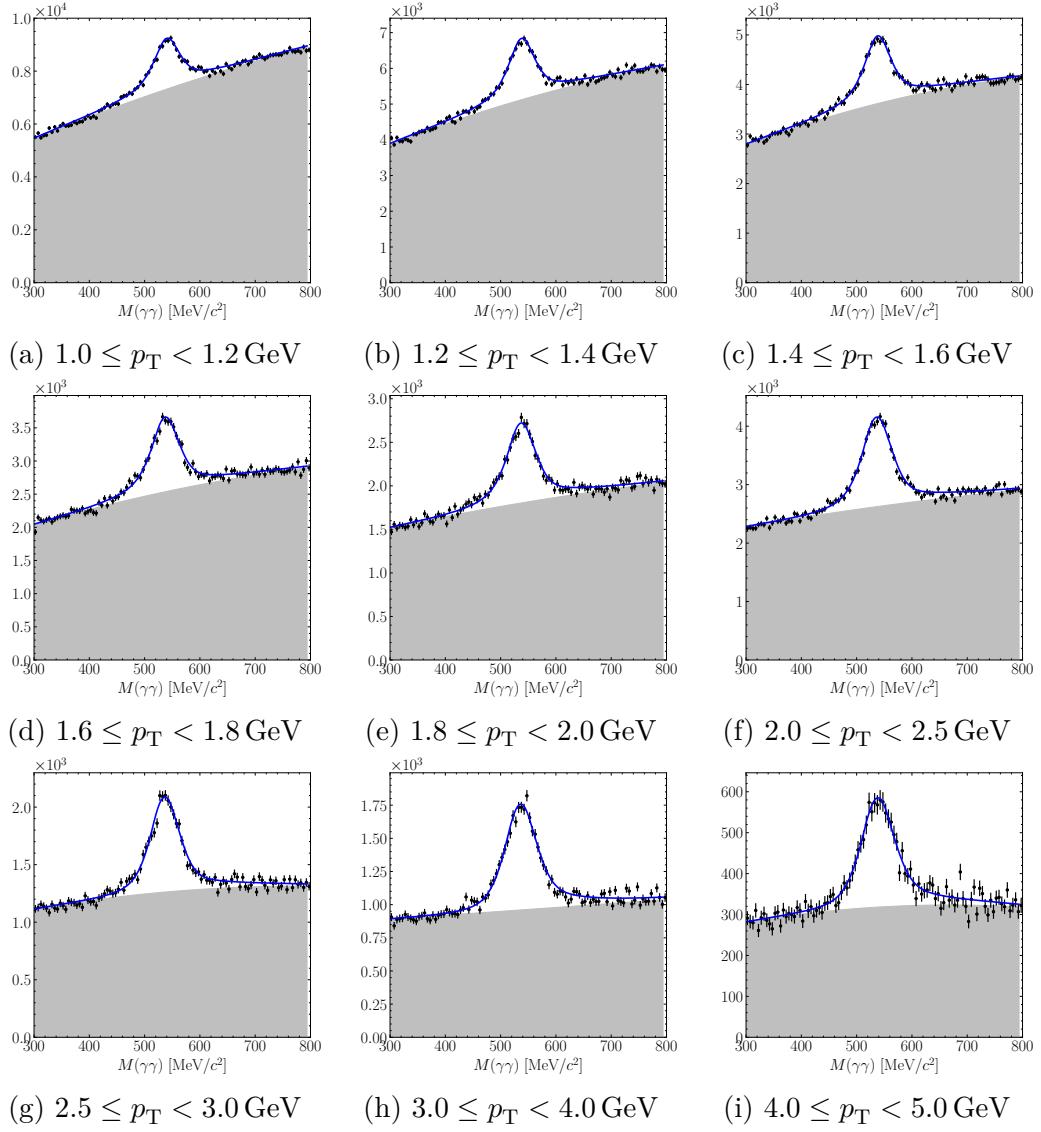


Figure B-17: η mass fits in 8.16 TeV $p\text{Pb}$ data as a function of $p_T(\gamma^{\text{cnv}})$.

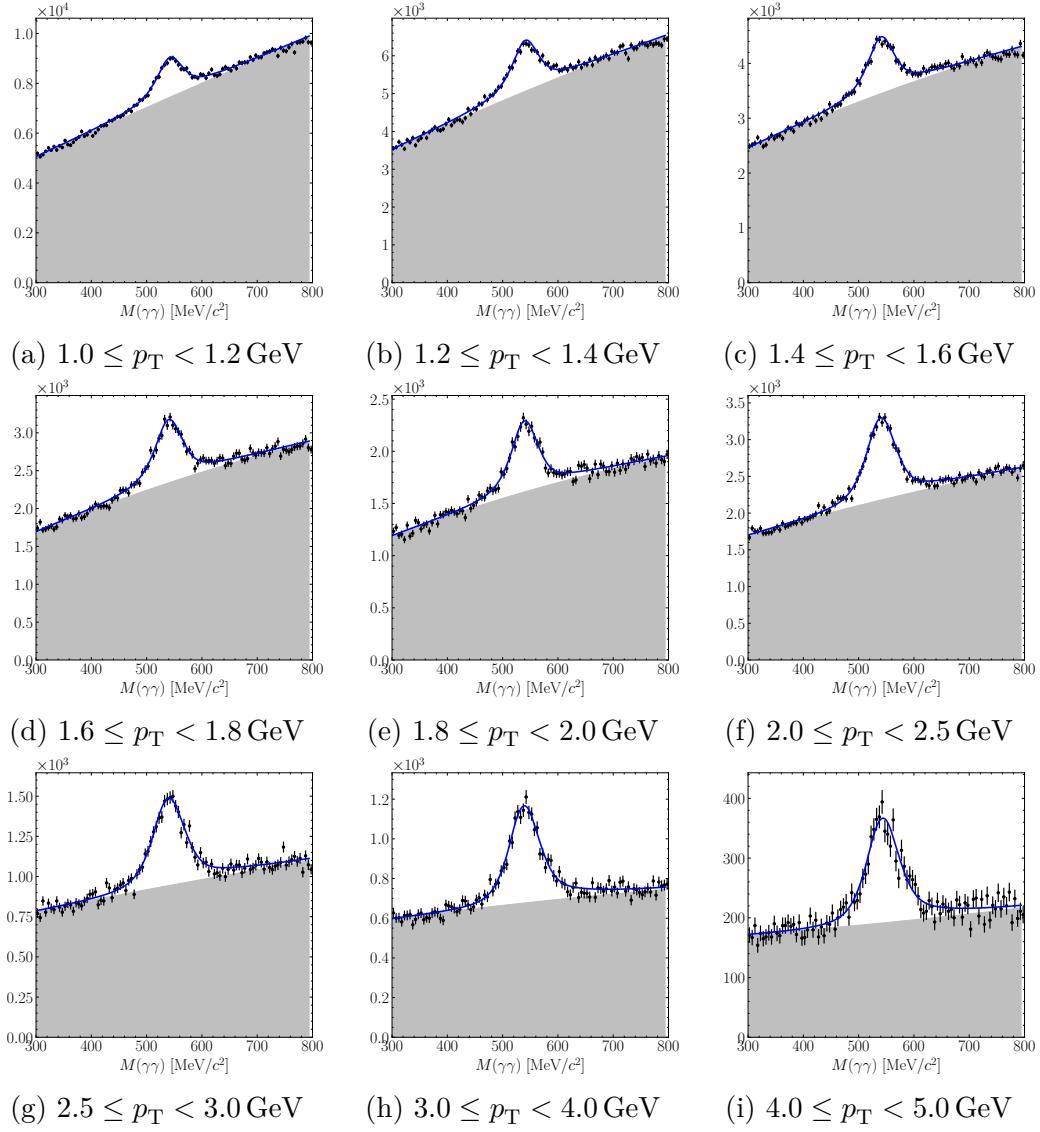


Figure B-18: η mass fits in 8.16 TeV $Pb\bar{p}$ data as a function of $p_T(\gamma^{\text{cnv}})$.

Bibliography

- [1] CMS, S. Chatrchyan *et al.*, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, Phys. Lett. B **716** (2012) 30, [arXiv:1207.7235](https://arxiv.org/abs/1207.7235).
- [2] ATLAS, G. Aad *et al.*, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B **716** (2012) 1, [arXiv:1207.7214](https://arxiv.org/abs/1207.7214).
- [3] J. C. Collins and M. J. Perry, *Superdense Matter: Neutrons Or Asymptotically Free Quarks?*, Phys. Rev. Lett. **34** (1975) 1353.
- [4] N. Cabibbo and G. Parisi, *Exponential Hadronic Spectrum and Quark Liberation*, Phys. Lett. B **59** (1975) 67.
- [5] E. V. Shuryak, *Theory of Hadronic Plasma*, Sov. Phys. JETP **47** (1978) 212.
- [6] P. Foka and M. A. Janik, *An overview of experimental results from ultra-relativistic heavy-ion collisions at the CERN LHC: Hard probes*, Rev. Phys. **1** (2016) 172, [arXiv:1702.07231](https://arxiv.org/abs/1702.07231).
- [7] P. Foka and M. A. Janik, *An overview of experimental results from ultra-relativistic heavy-ion collisions at the CERN LHC: Bulk properties and dynamical evolution*, Rev. Phys. **1** (2016) 154, [arXiv:1702.07233](https://arxiv.org/abs/1702.07233).
- [8] D. W. Higinbotham *et al.*, *Proton radius from electron scattering data*, Phys. Rev. C **93** (2016) 055207, [arXiv:1510.01293](https://arxiv.org/abs/1510.01293).
- [9] H1, ZEUS, H. Abramowicz *et al.*, *Combination of measurements of inclusive deep inelastic $e^\pm p$ scattering cross sections and QCD analysis of HERA data*, Eur. Phys. J. C **75** (2015) 580, [arXiv:1506.06042](https://arxiv.org/abs/1506.06042).
- [10] D. Griffiths, *Introduction to elementary particles*, 2008.
- [11] M. Thomson, *Modern particle physics*, Cambridge University Press, New York, 2013.
- [12] Super-Kamiokande, Y. Fukuda *et al.*, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81** (1998) 1562, [arXiv:hep-ex/9807003](https://arxiv.org/abs/hep-ex/9807003).

- [13] SNO, Q. R. Ahmad *et al.*, *Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by 8B solar neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **87** (2001) 071301, [arXiv:nucl-ex/0106015](https://arxiv.org/abs/nucl-ex/0106015).
- [14] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*, Addison-Wesley, Reading, USA, 1995.
- [15] P. Braun-Munzinger, V. Koch, T. Schäfer, and J. Stachel, *Properties of hot and dense matter from relativistic heavy ion collisions*, Phys. Rept. **621** (2016) 76, [arXiv:1510.00442](https://arxiv.org/abs/1510.00442).
- [16] J. Aebischer, J. Kumar, and D. M. Straub, *Wilson: a Python package for the running and matching of Wilson coefficients above and below the electroweak scale*, Eur. Phys. J. C **78** (2018) 1026, [arXiv:1804.05033](https://arxiv.org/abs/1804.05033).
- [17] STAR, J. Adams *et al.*, *Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR Collaboration's critical assessment of the evidence from RHIC collisions*, Nucl. Phys. A **757** (2005) 102, [arXiv:nucl-ex/0501009](https://arxiv.org/abs/nucl-ex/0501009).
- [18] PHENIX, K. Adcox *et al.*, *Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration*, Nucl. Phys. A **757** (2005) 184, [arXiv:nucl-ex/0410003](https://arxiv.org/abs/nucl-ex/0410003).
- [19] BRAHMS, I. Arsene *et al.*, *Quark gluon plasma and color glass condensate at RHIC? The Perspective from the BRAHMS experiment*, Nucl. Phys. A **757** (2005) 1, [arXiv:nucl-ex/0410020](https://arxiv.org/abs/nucl-ex/0410020).
- [20] PHOBOS, B. B. Back *et al.*, *The PHOBOS perspective on discoveries at RHIC*, Nucl. Phys. A **757** (2005) 28, [arXiv:nucl-ex/0410022](https://arxiv.org/abs/nucl-ex/0410022).
- [21] ALICE, K. Aamodt *et al.*, *Elliptic flow of charged particles in Pb-Pb collisions at 2.76 TeV*, Phys. Rev. Lett. **105** (2010) 252302, [arXiv:1011.3914](https://arxiv.org/abs/1011.3914).
- [22] CMS, S. Chatrchyan *et al.*, *Observation and studies of jet quenching in PbPb collisions at nucleon-nucleon center-of-mass energy = 2.76 TeV*, Phys. Rev. C **84** (2011) 024906, [arXiv:1102.1957](https://arxiv.org/abs/1102.1957).
- [23] CMS, S. Chatrchyan *et al.*, *Indications of suppression of excited Υ states in PbPb collisions at $\sqrt{S_{NN}} = 2.76$ TeV*, Phys. Rev. Lett. **107** (2011) 052302, [arXiv:1105.4894](https://arxiv.org/abs/1105.4894).
- [24] ATLAS, G. Aad *et al.*, *Measurement of the centrality dependence of J/ψ yields and observation of Z production in lead-lead collisions with the ATLAS detector at the LHC*, Phys. Lett. B **697** (2011) 294, [arXiv:1012.5419](https://arxiv.org/abs/1012.5419).
- [25] ATLAS, G. Aad *et al.*, *Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.77$ TeV with the ATLAS Detector at the LHC*, Phys. Rev. Lett. **105** (2010) 252303, [arXiv:1011.6182](https://arxiv.org/abs/1011.6182).

- [26] ATLAS, G. Aad *et al.*, *Measurement of the pseudorapidity and transverse momentum dependence of the elliptic flow of charged particles in lead-lead collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector*, Phys. Lett. B **707** (2012) 330, [arXiv:1108.6018](https://arxiv.org/abs/1108.6018).
- [27] PHENIX, C. Aidala *et al.*, *Creation of quark-gluon plasma droplets with three distinct geometries*, Nature Phys. **15** (2019) 214, [arXiv:1805.02973](https://arxiv.org/abs/1805.02973).
- [28] CMS, V. Khachatryan *et al.*, *Observation of Long-Range Near-Side Angular Correlations in Proton-Proton Collisions at the LHC*, JHEP **09** (2010) 091, [arXiv:1009.4122](https://arxiv.org/abs/1009.4122).
- [29] LHCb, R. Aaij *et al.*, *Measurements of long-range near-side angular correlations in $\sqrt{s_{NN}} = 5$ TeV proton-lead collisions in the forward region*, Phys. Lett. B **762** (2016) 473, [arXiv:1512.00439](https://arxiv.org/abs/1512.00439).
- [30] ATLAS, G. Aad *et al.*, *Observation of Associated Near-Side and Away-Side Long-Range Correlations in $\sqrt{s_{NN}} = 5.02$ TeV Proton-Lead Collisions with the ATLAS Detector*, Phys. Rev. Lett. **110** (2013) 182302, [arXiv:1212.5198](https://arxiv.org/abs/1212.5198).
- [31] ALICE, B. Abelev *et al.*, *Long-range angular correlations on the near and away side in p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, Phys. Lett. B **719** (2013) 29, [arXiv:1212.2001](https://arxiv.org/abs/1212.2001).
- [32] ATLAS, *Two-particle azimuthal correlations in photo-nuclear ultra-peripheral Pb+Pb collisions at 5.02~TeV with ATLAS*.
- [33] A. Badea *et al.*, *Measurements of two-particle correlations in e^+e^- collisions at 91 GeV with ALEPH archived data*, Phys. Rev. Lett. **123** (2019) 212002, [arXiv:1906.00489](https://arxiv.org/abs/1906.00489).
- [34] Belle, A. Abdesselam *et al.*, *Measurement of two-particle correlations in hadronic e^+e^- collisions at Belle*, 2020, [arXiv:2008.04187](https://arxiv.org/abs/2008.04187).
- [35] ZEUS, I. Abt *et al.*, *Two-particle azimuthal correlations as a probe of collective behaviour in deep inelastic ep scattering at HERA*, JHEP **04** (2020) 070, [arXiv:1912.07431](https://arxiv.org/abs/1912.07431).
- [36] R. Sahoo, *Possible Formation of QGP-droplets in Proton-Proton Collisions at the CERN Large Hadron Collider*, AAPPS Bull. **29** (2019) 16, [arXiv:1908.10566](https://arxiv.org/abs/1908.10566).
- [37] S. Schlichting and P. Tribedy, *Collectivity in Small Collision Systems: An Initial-State Perspective*, Adv. High Energy Phys. **2016** (2016) 8460349, [arXiv:1611.00329](https://arxiv.org/abs/1611.00329).
- [38] C. Bierlich, S. Chakraborty, G. Gustafson, and L. Lönnblad, *Setting the string shoving picture in a new frame*, [arXiv:2010.07595](https://arxiv.org/abs/2010.07595).

[39] E. D. Bloom *et al.*, *High-Energy Inelastic e p Scattering at 6-Degrees and 10-Degrees*, Phys. Rev. Lett. **23** (1969) 930.

[40] J. C. Collins, D. E. Soper, and G. Sterman, *Factorization of hard processes in qcd* , arXiv:hep-ph/0409313.

[41] T.-J. Hou *et al.*, *New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC*, arXiv:1912.10053.

[42] NNPDF, R. D. Ball *et al.*, *Parton distributions from high-precision collider data*, Eur. Phys. J. C **77** (2017) 663, arXiv:1706.00428.

[43] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and collider physics*, vol. 8, Cambridge University Press, 2011.

[44] A. Buckley *et al.*, *LHAPDF6: parton density access in the LHC precision era*, Eur. Phys. J. C **75** (2015) 132, arXiv:1412.7420.

[45] European Muon, J. J. Aubert *et al.*, *The ratio of the nucleon structure functions F_{2n} for iron and deuterium*, Phys. Lett. B **123** (1983) 275.

[46] K. J. Eskola, H. Paukkunen, and C. A. Salgado, *Eps09 - a new generation of nlo and lo nuclear parton distribution functions*, arXiv:0902.4154arXiv:0902.4154.

[47] K. J. Eskola, P. Paakkinen, H. Paukkunen, and C. A. Salgado, *Epps16: Nuclear parton distributions with lhc data*, arXiv:1612.05741arXiv:1612.05741.

[48] S. J. Brodsky and H. J. Lu, *Shadowing and Antishadowing of Nuclear Structure Functions*, Phys. Rev. Lett. **64** (1990) 1342.

[49] N. Armesto, *Nuclear shadowing*, J. Phys. G **32** (2006) R367, arXiv:hep-ph/0604108.

[50] O. Hen, E. Piasetzky, and L. B. Weinstein, *New data strengthen the connection between Short Range Correlations and the EMC effect*, Phys. Rev. C **85** (2012) 047301, arXiv:1202.3452.

[51] K. Kovarik *et al.*, *nCTEQ15 - Global analysis of nuclear parton distributions with uncertainties in the CTEQ framework*, Phys. Rev. D **93** (2016) 085037, arXiv:1509.00792.

[52] J. L. Albacete and C. Marquet, *Gluon saturation and initial conditions for relativistic heavy ion collisions*, Prog. Part. Nucl. Phys. **76** (2014) 1, arXiv:1401.4866.

[53] E. Iancu and R. Venugopalan, in *The Color glass condensate and high-energy scattering in QCD*, R. C. Hwa and X.-N. Wang, eds., pp. 249–3363, 2003. arXiv:hep-ph/0303204.

[54] D. d'Enterria and J. Rojo, *Quantitative constraints on the gluon distribution function in the proton from collider isolated-photon data*, Nucl. Phys. B **860** (2012) 311, [arXiv:1202.1762](https://arxiv.org/abs/1202.1762).

[55] J. M. Campbell, J. Rojo, E. Slade, and C. Williams, *Direct photon production and PDF fits reloaded*, Eur. Phys. J. C **78** (2018) 470, [arXiv:1802.03021](https://arxiv.org/abs/1802.03021).

[56] F. Arleo and T. Gousset, *Measuring gluon shadowing with prompt photons at RHIC and LHC*, Phys. Lett. B **660** (2008) 181, [arXiv:0707.2944](https://arxiv.org/abs/0707.2944).

[57] F. Arleo, K. J. Eskola, H. Paukkunen, and C. A. Salgado, *Inclusive prompt photon production in nuclear collisions at rhic and lhc*, [arXiv:1103.1471](https://arxiv.org/abs/1103.1471)[arXiv:1103.1471](https://arxiv.org/abs/1103.1471).

[58] I. Helenius, K. J. Eskola, and H. Paukkunen, *Probing the small- x nuclear gluon distributions with isolated photons at forward rapidities in $p+pb$ collisions at the lhc*, [arXiv:1406.1689](https://arxiv.org/abs/1406.1689)[arXiv:1406.1689](https://arxiv.org/abs/1406.1689).

[59] S. De, *Extent of sensitivity of single photon production to parton distribution functions*, [arXiv:1305.0624](https://arxiv.org/abs/1305.0624)[arXiv:1305.0624](https://arxiv.org/abs/1305.0624).

[60] J. Jalilian-Marian and A. H. Rezaeian, *Prompt photon production and photon-hadron correlations at rhic and the lhc from the color glass condensate*, [arXiv:1204.1319](https://arxiv.org/abs/1204.1319)[arXiv:1204.1319](https://arxiv.org/abs/1204.1319).

[61] A. H. Rezaeian, *Semi-inclusive photon-hadron production in pp and pa collisions at rhic and lhc*, [arXiv:1209.0478](https://arxiv.org/abs/1209.0478)[arXiv:1209.0478](https://arxiv.org/abs/1209.0478).

[62] G. S. dos Santos, G. G. da Silveira, and M. V. T. Machado, *Prompt photon production in high-energy pa collisions at forward rapidity*, [arXiv:2006.15743](https://arxiv.org/abs/2006.15743).

[63] ALICE, S. Acharya *et al.*, *Measurement of the inclusive isolated photon production cross section in pp collisions at $\sqrt{s} = 7$ TeV*, Eur. Phys. J. C **79** (2019) 896, [arXiv:1906.01371](https://arxiv.org/abs/1906.01371).

[64] ATLAS, G. Aad *et al.*, *Measurement of isolated-photon plus two-jet production in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **03** (2020) 179, [arXiv:1912.09866](https://arxiv.org/abs/1912.09866).

[65] S. Catani, M. Fontannaz, J. P. Guillet, and E. Pilon, *Cross-section of isolated prompt photons in hadron hadron collisions*, JHEP **05** (2002) 028, [arXiv:hep-ph/0204023](https://arxiv.org/abs/hep-ph/0204023).

[66] P. Aurenche *et al.*, *A New critical study of photon production in hadronic collisions*, Phys. Rev. D **73** (2006) 094007, [arXiv:hep-ph/0602133](https://arxiv.org/abs/hep-ph/0602133).

[67] Z. Belghobsi *et al.*, *Photon - Jet Correlations and Constraints on Fragmentation Functions*, Phys. Rev. D **79** (2009) 114024, [arXiv:0903.4834](https://arxiv.org/abs/0903.4834).

[68] J. Pumplin *et al.*, *New generation of parton distributions with uncertainties from global QCD analysis*, JHEP **07** (2002) 012, [arXiv:hep-ph/0201195](https://arxiv.org/abs/hep-ph/0201195).

[69] L. Bourhis, M. Fontannaz, and J. P. Guillet, *Quarks and gluon fragmentation functions into photons*, Eur. Phys. J. C **2** (1998) 529, [arXiv:hep-ph/9704447](https://arxiv.org/abs/hep-ph/9704447).

[70] P. V. Ruuskanen, *Electromagnetic probes of quark - gluon plasma in relativistic heavy ion collisions*, Nucl. Phys. A **544** (1992) 169.

[71] F. D. Steffen and M. H. Thoma, *Hard thermal photon production in relativistic heavy ion collisions*, [arXiv:hep-ph/0103044](https://arxiv.org/abs/hep-ph/0103044)[arXiv:hep-ph/0103044](https://arxiv.org/abs/hep-ph/0103044).

[72] S. Turbide, C. Gale, E. Frodermann, and U. Heinz, *Electromagnetic radiation from nuclear collisions at ultrarelativistic energies*, Phys. Rev. C **77** (2008) 024909.

[73] C. Shen, U. W. Heinz, J.-F. Paquet, and C. Gale, *Thermal photons as a quark-gluon plasma thermometer revisited*, [arXiv:1308.2440](https://arxiv.org/abs/1308.2440)[arXiv:1308.2440](https://arxiv.org/abs/1308.2440).

[74] C. Gale, *Photon production in hot and dense strongly interacting matter*, [arXiv:0904.2184](https://arxiv.org/abs/0904.2184)[arXiv:0904.2184](https://arxiv.org/abs/0904.2184).

[75] D. G. d'Enterria and D. Peressounko, *Probing the QCD equation of state with thermal photons in nucleus-nucleus collisions at RHIC*, Eur. Phys. J. C **46** (2006) 451, [arXiv:nucl-th/0503054](https://arxiv.org/abs/nucl-th/0503054).

[76] F.-M. Liu and K. Werner, *Direct photons at low transverse momentum – a qgp signal in pp collisions at lhc*, [arXiv:1102.1052](https://arxiv.org/abs/1102.1052)[arXiv:1102.1052](https://arxiv.org/abs/1102.1052).

[77] H. W. Lewis, J. R. Oppenheimer, and S. A. Wouthuysen, *The Multiple Production of Mesons*, Phys. Rev. **73** (1948) 127.

[78] P. Ilten, *CIMBA: fast Monte Carlo generation using cubic interpolation*, Comput. Phys. Commun. **258** (2021) 107622, [arXiv:1908.08353](https://arxiv.org/abs/1908.08353).

[79] T. Sjöstrand *et al.*, *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191** (2015) 159, [arXiv:1410.3012](https://arxiv.org/abs/1410.3012).

[80] WA98, M. M. Aggarwal *et al.*, *Observation of direct photons in central 158-A-GeV Pb-208 + Pb-208 collisions*, Phys. Rev. Lett. **85** (2000) 3595, [arXiv:nucl-ex/0006008](https://arxiv.org/abs/nucl-ex/0006008).

[81] PHENIX, A. Adare *et al.*, *Enhanced production of direct photons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV and implications for the initial temperature*, Phys. Rev. Lett. **104** (2010) 132301, [arXiv:0804.4168](https://arxiv.org/abs/0804.4168).

[82] PHENIX, A. Adare *et al.*, *Centrality dependence of low-momentum direct-photon production in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **91** (2015) 064904, [arXiv:1405.3940](https://arxiv.org/abs/1405.3940).

[83] STAR, L. Adamczyk *et al.*, *Direct virtual photon production in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Lett. B **770** (2017) 451, [arXiv:1607.01447](https://arxiv.org/abs/1607.01447).

[84] PHENIX, A. Adare *et al.*, *Beam Energy and Centrality Dependence of Direct-Photon Emission from Ultrarelativistic Heavy-Ion Collisions*, Phys. Rev. Lett. **123** (2019) 022301, [arXiv:1805.04084](https://arxiv.org/abs/1805.04084).

[85] PHENIX, A. Adare *et al.*, *Low-momentum direct photon measurement in Cu+Cu collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **98** (2018) 054902, [arXiv:1805.04066](https://arxiv.org/abs/1805.04066).

[86] ALICE, J. Adam *et al.*, *Direct photon production in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Lett. B **754** (2016) 235, [arXiv:1509.07324](https://arxiv.org/abs/1509.07324).

[87] PHENIX, A. Adare *et al.*, *Direct photon production in d+Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **87** (2013) 054907, [arXiv:1208.1234](https://arxiv.org/abs/1208.1234).

[88] V. Canoa Roman, *Direct Photons at the PHENIX Experiment: From Large to Small Systems*, MDPI Proc. **10** (2019) 32.

[89] ALICE, S. Acharya *et al.*, *Direct photon production at low transverse momentum in proton-proton collisions at $\sqrt{s} = 2.76$ and 8 TeV*, Phys. Rev. C **99** (2019) 024912, [arXiv:1803.09857](https://arxiv.org/abs/1803.09857).

[90] ALICE, D. Blau, *Direct photon production in pp, p-pb and pb-pb collisions: results from alice*, EPJ Web Conf. **222** (2019) 02001.

[91] PHENIX, A. Adare *et al.*, *Azimuthally anisotropic emission of low-momentum direct photons in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **94** (2016) 064901, [arXiv:1509.07758](https://arxiv.org/abs/1509.07758).

[92] ALICE, S. Acharya *et al.*, *Direct photon elliptic flow in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Lett. B **789** (2019) 308, [arXiv:1805.04403](https://arxiv.org/abs/1805.04403).

[93] C. Gale, *Direct photon production in relativistic heavy-ion collisions - a theory update*, PoS **High-pT2017** (2019) 023, [arXiv:1802.00128](https://arxiv.org/abs/1802.00128).

[94] *LHC Machine*, JINST **3** (2008) S08001.

[95] ATLAS, G. Aad *et al.*, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.

[96] CMS, S. Chatrchyan *et al.*, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S08004.

[97] ALICE, K. Aamodt *et al.*, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002.

[98] LHCb, J. Alves, A. Augusto *et al.*, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005.

- [99] *LHC Design Report Vol.1: The LHC Main Ring*, .
- [100] R. Alemany-Fernandez *et al.*, *Operation and Configuration of the LHC in Run 1*, .
- [101] J. Wenninger, *Operation and Configuration of the LHC in Run 2*, .
- [102] F. Follin and D. Jacquet, *Implementation and experience with luminosity levelling with offset beam*, in *ICFA Mini-Workshop on Beam-Beam Effects in Hadron Colliders*, 183–187, 2014, [arXiv:1410.3667](https://arxiv.org/abs/1410.3667).
- [103] LHCb, R. Aaij *et al.*, *LHCb Detector Performance*, *Int. J. Mod. Phys. A* **30** (2015) 1530022, [arXiv:1412.6352](https://arxiv.org/abs/1412.6352).
- [104] LHCb, R. Aaij *et al.*, *Design and performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC*, *JINST* **14** (2019) P04013, [arXiv:1812.10790](https://arxiv.org/abs/1812.10790).
- [105] LHCb, R. Aaij *et al.*, *Precision luminosity measurements at LHCb*, *JINST* **9** (2014) P12005, [arXiv:1410.0149](https://arxiv.org/abs/1410.0149).
- [106] BaBar, Belle, A. J. Bevan *et al.*, *The Physics of the B Factories*, *Eur. Phys. J. C* **74** (2014) 3026, [arXiv:1406.6311](https://arxiv.org/abs/1406.6311).
- [107] C. T. Hill, *B physics in hadron colliders*, in *Theoretical Advanced Study Institute (TASI 93) in Elementary Particle Physics: The Building Blocks of Creation - From Microfermions to Megaparsecs*, 0127–138, 1993, [arXiv:hep-ph/9309309](https://arxiv.org/abs/hep-ph/9309309).
- [108] P. Nason, S. Dawson, and R. K. Ellis, *The Total Cross-Section for the Production of Heavy Quarks in Hadronic Collisions*, *Nucl. Phys. B* **303** (1988) 607.
- [109] ATLAS, M. Aaboud *et al.*, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673, [arXiv:1704.07983](https://arxiv.org/abs/1704.07983).
- [110] CMS, S. Chatrchyan *et al.*, *Description and performance of track and primary-vertex reconstruction with the CMS tracker*, *JINST* **9** (2014) P10009, [arXiv:1405.6569](https://arxiv.org/abs/1405.6569).
- [111] R. D. Ball *et al.*, *Parton distributions with LHC data*, *Nucl. Phys. B* **867** (2013) 244, [arXiv:1207.1303](https://arxiv.org/abs/1207.1303).
- [112] LHCb, R. Aaij *et al.*, *Measurement of B_s^0 and D_s^- meson lifetimes*, *Phys. Rev. Lett.* **119** (2017) 101801, [arXiv:1705.03475](https://arxiv.org/abs/1705.03475).
- [113] R. Aaij *et al.*, *Performance of the LHCb Vertex Locator*, *JINST* **9** (2014) P09007, [arXiv:1405.7808](https://arxiv.org/abs/1405.7808).

- [114] J. Gassner, M. Needham, and O. Steinkamp, *Layout and Expected Performance of the LHCb TT Station*, .
- [115] J. Gassner, F. Lehner, and F. Steiner, *The mechanical design of the LHCb silicon trigger tracker*, .
- [116] LHCb, *LHCb: Inner tracker technical design report*, .
- [117] *LHCb: Outer tracker technical design report*, .
- [118] LHCb, *LHCb calorimeters: Technical design report*, .
- [119] A. Arefev *et al.*, *Beam test results of the LHCb electromagnetic calorimeter*, .
- [120] LHCb, *LHCb muon system technical design report*, .
- [121] T. Head, *The LHCb trigger system*, JINST **9** (2014) C09015.
- [122] LHCb, *LHCb: RICH technical design report*, .
- [123] T. Likhomanenko *et al.*, *LHCb Topological Trigger Reoptimization*, J. Phys. Conf. Ser. **664** (2015) 082025, [arXiv:1510.00572](https://arxiv.org/abs/1510.00572).
- [124] R. Aaij *et al.*, *Tesla : an application for real-time data analysis in High Energy Physics*, Comput. Phys. Commun. **208** (2016) 35, [arXiv:1604.05596](https://arxiv.org/abs/1604.05596).
- [125] V. V. Gligorov and M. Williams, *Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree*, JINST **8** (2013) P02013, [arXiv:1210.6861](https://arxiv.org/abs/1210.6861).
- [126] R. E. Kalman, *A new approach to linear filtering and prediction problems*, Journal of Basic Engineering **82** (1960) 35.
- [127] R. Aaij *et al.*, *Allen: A high level trigger on gpus for lhcb*, [arXiv:1912.09161](https://arxiv.org/abs/1912.09161)[arXiv:1912.09161](https://arxiv.org/abs/1912.09161).
- [128] LHCb, *LHCb Upgrade GPU High Level Trigger Technical Design Report*, CERN-LHCC-2020-006. LHCB-TDR-021, CERN, Geneva, 2020.
- [129] C. Fitzpatrick and V. V. Gligorov, *Anatomy of an upgrade event in the upgrade era, and implications for the lhcb trigger*, LHCb-PUB-2014-027. CERN-LHCb-PUB-2014-027, CERN, Geneva, 2014.
- [130] J. Alwall *et al.*, *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, JHEP **07** (2014) 079, [arXiv:1405.0301](https://arxiv.org/abs/1405.0301).
- [131] R. Frederix *et al.*, *The automation of next-to-leading order electroweak calculations*, [arXiv:1804.10017](https://arxiv.org/abs/1804.10017)[arXiv:1804.10017](https://arxiv.org/abs/1804.10017).
- [132] R. D. Ball *et al.*, *Parton distributions from high-precision collider data*, The European Physical Journal C **77** (2017) .

- [133] LHCb, *Lhcb trigger and online upgrade technical design report* , 2014.
- [134] LHCb, I. Bediaga, *Lhcb velo upgrade technical design report* , 2013.
- [135] LHCb, *Lhcb tracker upgrade technical design report* , 2014.
- [136] LHCb, R. Aaij *et al.*, *Search for Dark Photons Produced in 13 TeV pp Collisions*, Phys. Rev. Lett. **120** (2018) 061801, [arXiv:1710.02867](https://arxiv.org/abs/1710.02867).
- [137] LHCb, *Lhcb pid upgrade technical design report* , 2013.
- [138] D. Rohr *et al.*, *Alice hlt tpc tracking of pb-pb events on gpus*, [arXiv:1712.09407](https://arxiv.org/abs/1712.09407)[arXiv:1712.09407](https://arxiv.org/abs/1712.09407).
- [139] D. Funke *et al.*, *Parallel track reconstruction in CMS using the cellular automaton approach*, Journal of Physics: Conference Series **513** (2014) 052010.
- [140] D. vom Bruch, *Online data reduction using track and vertex reconstruction on GPUs for the mu3e experiment*, EPJ Web of Conferences **150** (2017) 00013.
- [141] V. Singhal, S. Chattpadhyay, and V. Friese, *Investigation of heterogeneous computing platforms for real-time data analysis in the cbm experiment*, [arXiv:1810.11966](https://arxiv.org/abs/1810.11966)[arXiv:1810.11966](https://arxiv.org/abs/1810.11966).
- [142] E. Lindholm, J. Nickolls, S. Oberman, and J. Montrym, *Nvidia tesla: A unified graphics and computing architecture*, IEEE Micro **28** (2008) 39.
- [143] J. Sanders and E. Kandrot, *CUDA by Example: An Introduction to General-Purpose GPU Programming*, Addison-Wesley Professional, 1st ed., 2010.
- [144] *Cuda toolkit documentation v11.1.1*, .
- [145] D. H. C. Perez, N. Neufeld, and A. R. Nunez, *A fast local algorithm for track reconstruction on parallel architectures*, in *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, IEEE, 2019.
- [146] P. F. Declaro *et al.*, *A parallel-computing algorithm for high-energy physics particle tracking and decoding using GPU architectures*, IEEE Access **7** (2019) 91612.
- [147] F. Archilli *et al.*, *Performance of the Muon Identification at LHCb*, JINST **8** (2013) P10020, [arXiv:1306.0249](https://arxiv.org/abs/1306.0249).
- [148] C. Fitzpatrick *et al.*, *Upgrade trigger: Bandwidth strategy proposal*, LHCb-PUB-2017-006. CERN-LHCb-PUB-2017-006, CERN, Geneva, 2017.
- [149] LHCb, R. Aaij *et al.*, *Test of lepton universality with $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays*, JHEP **08** (2017) 055, [arXiv:1705.05802](https://arxiv.org/abs/1705.05802).
- [150] T. Boettcher, P. Ilten, and M. Williams, *Direct probe of the intrinsic charm content of the proton*, Phys. Rev. D **93** (2016) 074008, [arXiv:1512.06666](https://arxiv.org/abs/1512.06666).

- [151] NNPDF, R. D. Ball *et al.*, *A Determination of the Charm Content of the Proton*, Eur. Phys. J. C **76** (2016) 647, [arXiv:1605.06515](https://arxiv.org/abs/1605.06515).
- [152] European Muon, J. J. Aubert *et al.*, *Production of charmed particles in 250-GeV μ^+ - iron interactions*, Nucl. Phys. B **213** (1983) 31.
- [153] S. J. Brodsky, P. Hoyer, C. Peterson, and N. Sakai, *The Intrinsic Charm of the Proton*, Phys. Lett. B **93** (1980) 451.
- [154] MILC, W. Freeman and D. Toussaint, *Intrinsic strangeness and charm of the nucleon using improved staggered fermions*, Phys. Rev. D **88** (2013) 054503, [arXiv:1204.3866](https://arxiv.org/abs/1204.3866).
- [155] IceCube, M. G. Aartsen *et al.*, *Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector*, Science **342** (2013) 1242856, [arXiv:1311.5238](https://arxiv.org/abs/1311.5238).
- [156] R. Gauld *et al.*, *The prompt atmospheric neutrino flux in the light of LHCb*, JHEP **02** (2016) 130, [arXiv:1511.06346](https://arxiv.org/abs/1511.06346).
- [157] P. Jimenez-Delgado, T. J. Hobbs, J. T. Londergan, and W. Melnitchouk, *New limits on intrinsic charm in the nucleon from global analysis of parton distributions*, Phys. Rev. Lett. **114** (2015) 082002, [arXiv:1408.1708](https://arxiv.org/abs/1408.1708).
- [158] S. J. Brodsky and S. Gardner, *Comment on “New Limits on Intrinsic Charm in the Nucleon from Global Analysis of Parton Distributions”*, Phys. Rev. Lett. **116** (2016) 019101, [arXiv:1504.00969](https://arxiv.org/abs/1504.00969).
- [159] P. Jimenez-Delgado, T. J. Hobbs, J. T. Londergan, and W. Melnitchouk, *Reply to Comment on “New limits on intrinsic charm in the nucleon from global analysis of parton distributions”*, Phys. Rev. Lett. **116** (2016) 019102, [arXiv:1504.06304](https://arxiv.org/abs/1504.06304).
- [160] S. Dulat *et al.*, *New parton distribution functions from a global analysis of quantum chromodynamics*, Phys. Rev. D **93** (2016) 033006, [arXiv:1506.07443](https://arxiv.org/abs/1506.07443).
- [161] NNPDF, R. D. Ball *et al.*, *Parton distributions for the LHC Run II*, JHEP **04** (2015) 040, [arXiv:1410.8849](https://arxiv.org/abs/1410.8849).
- [162] S. Alioli, P. Nason, C. Oleari, and E. Re, *Vector boson plus one jet production in POWHEG*, JHEP **01** (2011) 095, [arXiv:1009.5594](https://arxiv.org/abs/1009.5594).
- [163] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, JHEP **11** (2004) 040, [arXiv:hep-ph/0409146](https://arxiv.org/abs/hep-ph/0409146).
- [164] D. J. Lange, *The EvtGen particle decay simulation package*, Nucl. Instrum. Meth. A **462** (2001) 152.
- [165] P. Golonka and Z. Was, *PHOTOS Monte Carlo: A Precision tool for QED corrections in Z and W decays*, Eur. Phys. J. C **45** (2006) 97, [arXiv:hep-ph/0506026](https://arxiv.org/abs/hep-ph/0506026).

[166] R. Frederix and S. Frixione, *Merging meets matching in MC@NLO*, JHEP **12** (2012) 061, [arXiv:1209.6215](https://arxiv.org/abs/1209.6215).

[167] M. Cacciari, G. P. Salam, and G. Soyez, *FastJet User Manual*, Eur. Phys. J. C **72** (2012) 1896, [arXiv:1111.6097](https://arxiv.org/abs/1111.6097).

[168] M. Cacciari, G. P. Salam, and G. Soyez, *The anti- k_t jet clustering algorithm*, JHEP **04** (2008) 063, [arXiv:0802.1189](https://arxiv.org/abs/0802.1189).

[169] LHCb, R. Aaij *et al.*, *Identification of beauty and charm quark jets at LHCb*, JINST **10** (2015) P06013, [arXiv:1504.07670](https://arxiv.org/abs/1504.07670).

[170] NNPDF, R. D. Ball *et al.*, *Reweighting NNPDFs: the W lepton asymmetry*, Nucl. Phys. B **849** (2011) 112, [arXiv:1012.0836](https://arxiv.org/abs/1012.0836), [Erratum: Nucl.Phys.B 854, 926–927 (2012), Erratum: Nucl.Phys.B 855, 927–928 (2012)].

[171] R. D. Ball *et al.*, *Reweighting and Unweighting of Parton Distributions and the LHC W lepton asymmetry data*, Nucl. Phys. B **855** (2012) 608, [arXiv:1108.1758](https://arxiv.org/abs/1108.1758).

[172] C. Loizides, J. Kamin, and D. d'Enterria, *Improved Monte Carlo Glauber predictions at present and future nuclear colliders*, Phys. Rev. C **97** (2018) 054910, [arXiv:1710.07098](https://arxiv.org/abs/1710.07098), [Erratum: Phys.Rev.C 99, 019901 (2019)].

[173] G. Corti *et al.*, *Software for the LHCb experiment*, IEEE Trans. Nucl. Sci. **53** (2006) 1323.

[174] S. Porteboeuf, T. Pierog, and K. Werner, *Producing Hard Processes Regarding the Complete Event: The EPOS Event Generator*, in *45th Rencontres de Moriond on QCD and High Energy Interactions*, 135–140, Gioi Publishers, 2010, [arXiv:1006.2967](https://arxiv.org/abs/1006.2967).

[175] GEANT4, S. Agostinelli *et al.*, *GEANT4—a simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) 250.

[176] LHCb, R. Aaij *et al.*, *Measurement of the relative rate of prompt χ_{c0} , χ_{c1} and χ_{c2} production at $\sqrt{s} = 7\text{TeV}$* , JHEP **10** (2013) 115, [arXiv:1307.4285](https://arxiv.org/abs/1307.4285).

[177] M. De Cian, S. Farry, P. Seyfert, and S. Stahl, *Fast neural-net based fake track rejection in the LHCb reconstruction*, LHCb-PUB-2017-011. CERN-LHCb-PUB-2017-011, CERN, Geneva, 2017.

[178] T. Skwarnicki, *A study of the radiative CASCADE transitions between the Upsilon-Prime and Upsilon resonances*, PhD thesis, Cracow, INP, 1986.

[179] E. Govorkova, *Study of π^0/γ efficiency using B meson decays in the LHCb experiment*, Phys. Atom. Nucl. **79** (2016) 1474, [arXiv:1505.02960](https://arxiv.org/abs/1505.02960).

[180] LHCb, R. Aaij *et al.*, *Measurement of the inelastic pp cross-section at a centre-of-mass energy of 13 TeV*, JHEP **06** (2018) 100, [arXiv:1803.10974](https://arxiv.org/abs/1803.10974).

- [181] D. de Florian *et al.*, *Parton-to-Pion Fragmentation Reloaded*, Phys. Rev. D **91** (2015) 014035, [arXiv:1410.6027](https://arxiv.org/abs/1410.6027).
- [182] L. Altenkämper, F. Bock, C. Loizides, and N. Schmidt, *Applicability of transverse mass scaling in hadronic collisions at energies available at the CERN Large Hadron Collider*, Phys. Rev. C **96** (2017) 064907, [arXiv:1710.01933](https://arxiv.org/abs/1710.01933).
- [183] ALICE, S. Acharya *et al.*, *Neutral pion and η meson production in p -Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, Eur. Phys. J. C **78** (2018) 624, [arXiv:1801.07051](https://arxiv.org/abs/1801.07051).
- [184] G. D'Agostini, *A multidimensional unfolding method based on Bayes' Theorem*, DESY-94-099, DESY, Hamburg, 1994.
- [185] J. Bourbeau and Z. Hampel-Arias, *Pyunfold: A python package for iterative unfolding*, The Journal of Open Source Software **3** (2018) 741.
- [186] CMS, V. Khachatryan *et al.*, *Charged-particle nuclear modification factors in $PbPb$ and pPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV*, JHEP **04** (2017) 039, [arXiv:1611.01664](https://arxiv.org/abs/1611.01664).
- [187] PHENIX, C. Aidala *et al.*, *Nuclear-modification factor of charged hadrons at forward and backward rapidity in $p+Al$ and $p+Au$ collisions at $\sqrt{s_{NN}} = 200$ GeV*, Phys. Rev. C **101** (2020) 034910, [arXiv:1906.09928](https://arxiv.org/abs/1906.09928).