

Abstract

Novel Uses of Machine Learning for Differential Jet Quenching Measurements at the LHC

Hannah J. Bossi
2023

At sufficiently high temperatures and pressures, QCD matter becomes a hot and dense deconfined medium known as the [Quark-Gluon Plasma \(QGP\)](#). Collisions of relativistic heavy ions are used to recreate the [QGP](#), providing a rich laboratory for exploring the mysteries of the strong interaction. The intrinsic and dynamic properties of the [QGP](#) are probed with jets, narrow cones of particles resulting from the scattering of quarks and gluons with a high momentum transfer. In heavy-ion collisions, jets interact with the [QGP](#) as they traverse it, leading to jet energy loss and modification of the jet's internal structure. The [ALICE](#) detector at the [LHC](#) is optimized for measurements in the heavy-ion collision environment and allows for the reconstruction of jets at relatively low transverse momentum. In this thesis, the most differential measurement of jet energy loss ever made by the [ALICE](#) collaboration, as well as the first measurement of jets in heavy-ion collisions using machine learning techniques, will be described. The cone-size dependence of jet quenching is investigated, revealing hints that wider jets lose more energy; an intriguing observation consistent with numerous jet quenching models.



Novel Uses of Machine Learning for Differential Jet Quenching
Measurements at the LHC

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Hannah J. Bossi

Dissertation Director: John W. Harris

May 2023

Copyright © 2023 by Hannah J. Bossi
All rights reserved.



Novel Uses of Machine Learning for Differential Jet Quenching Measurements at the LHC

Hannah J. Bossi

Acknowledgments

During the many trials of graduate school, Caitie and I used to remind one another: “We love physics. Of all the things, we love physics the most.” I do indeed love physics, which should become clear in the remainder of this thesis. But, when it comes to what I love most, of all the things, it’s the wonderful people I get to experience life with. It is in this section that I will attempt to express this love.

Firstly, I’d like to thank my advisor Prof. John Harris and my co-advisor Prof. Helen Caines. John, thank you for being an unwavering advocate and continually supporting my goals. No matter what ambitious travel itinerary, research plan, or timeline I gave you, you always said okay and that you knew I could get it done. Thank you for trusting me to mentor Mary, which has allowed me to grow as an independent researcher. Helen, my fellow Leo, I have had an absolute blast working with you. Thank you for providing fantastic advice and guidance with your truly impressive problem-solving abilities that make the impossible possible. I admire so many things about you both as scientists, the most important of which is how you value the people that you do science with. You both have changed the field, making it a much more inclusive and productive space, and I am honored to be a part of your legacy.

I am incredibly grateful to my RHIG academic family (ordered by office location); Caitie, Audrey, Youqi, Fernando, Tong, Dan, Dave, Laura, Mary, Raymond, Michael, Nikolai, Helen, Raghav, Isaac, Ágnes, Mike, Andrew, Sierra, Ananya, Paula, and John (as well as Mesut and Rüediger who were based at CERN). Your intelligence and kindness are evident in every meeting, hallway coffee, and lunchtime conversation. Today, I feel just as heard, respected, and valued as I did on my first day in the group. Wherever I go, I will strive to recreate the culture we have here. Within the group, there are a few individuals I would like to thank. Mary, I am so proud of the scientist you have become and I am so lucky to work with such a bright, kind, and persistent student. Raymond, from starting my first train to writing the paper and everywhere in between, thank you for always being someone I can go to with any question. Ágnes, thank you for inspiring me to think big and enriching my life with the full-spirited nature by which you live yours. Mike, I am so thankful for our many-year friendship and the common interests we share. Thank you for serving as my “bodyguard” to keep me safe and providing a home away from home whenever I was at CERN. Paula, thank you for making the many “world tours” completed during this thesis possible and for always brightening my morning by saying hello.

I owe my first scientific experiences to Ruth, Apryl, Paige, and the rest of the “380 trouble squad”. Thank you for teaching me that persistence (and good documentation) make a scientist’s life much easier, and for exposing me to the joy that a lifelong dedication to science can bring.

I have wholeheartedly enjoyed my time working alongside many ALICE colleagues both as a graduate student and as an undergraduate. I owe this largely to the whole 587 crew; Markus, Florian, Friederike, Nico, Constantin, Leticia, and the many others that came and went. Thank you for making CERN feel like a second home and for always welcoming me back with open arms. Markus, as my summer student advisor you were responsible for my beginnings in ALICE. Thank you for your patience and support as I

found my way and for your invaluable contributions to my service work project. Florian, though we do not share a thesis advisor, I view you as my academic brother. I am so happy that you have been a constant throughout my time on ALICE. Friederike, thank you for being my friend and mentor and for always pushing me to do my best. Meeting you changed the course of my life and I would not be writing this today if it were not for you. I am also thankful to Carolina, Luca, and Zuzana for their friendship and their service to the ALICE collaboration. Being a junior representative was a stressful part of this job, but as a result, we formed an incredibly close friendship that I treasure.

I am thankful for my dear friends from Colby (in the order that I met them); Sam, Chloe, Abby, Haley, Sophie, Michael, and Lilly (and of course Brittany too). Thank you for loving me as I am and being there with me from the very day I found out I was accepted to Yale. As we have grown into adults, I am so happy that we have maintained the lifelong bonds we made at Colby. Thank you for celebrating with me when things go well, comforting me when things don't, and providing constant laughter and joy.

I would also like to thank my friends from graduate school. I would like to thank the members of my cohort, specifically Ridge, James, Giacomo, and Trevor. You have all helped me a tremendous amount, particularly during the first few years of classes. I am also thankful to Samantha, who is perhaps the most emotionally intelligent person I have ever met. Thank you for inspiring me to think beyond this thesis and for helping me define my own vision of what it means to be a scientist.

I am also incredibly grateful for Aniket. Thank you for helping me take both literal and figurative leaps of faith that allow me to grow into a better person. Thank you for the countless weekends you spent working with me while writing this thesis so that my inner extrovert would not get lonely. Thank you for every hand-written note, cup of coffee, and pint of ice cream. This thesis would not have been possible without your love and support.

So much of who I am today is due to my loving and supportive family. The single hardest part of getting a Ph.D. in physics has been that it meant spending time away from you all. Mom and Dad, thank you for providing me with the confidence of knowing that whatever I chose to do in this life, you would support me. Thank you for instilling values of kindness, positivity, and persistence. Brittany, you are (and will always be) my favorite person in this universe. This thesis is perhaps the greatest gift I will ever give another person, and I am so happy to dedicate it to you. Jacob, thank you for always protecting me and for inspiring me with the excitement with which you approach your many adventures. Eleanor and Michael, thank you for the love and happiness that you add to our family. I am so happy to have you two as siblings. I am also grateful for the support and friendship of my extended family; my grandparents, my aunts and uncles, my cousins (and their children). In particular, I would like to thank my grandmother, who passed away during the completion of this thesis, for providing me with truly unconditional love. I wish that you could be here to see the conclusion of this journey. I would also like to thank my Uncle Bucky and Ellen for always making me laugh and for the constant praise of saying that I am "the smartest girl they know".

I will conclude these acknowledgments by thanking the two women who have most greatly shaped who I am as a physicist. Caitie, you have been my greatest ally and champion while also being my best friend. It means so much to me to have you by my side for

big life-changing days like this one, but what I will miss most are the everyday moments; the morning walks, the coffee chats, and the spontaneous adventures. You are a blessing in my life that has made my time at Yale better in every way and I am honored to have shared this journey with you. Laura, I am forever thankful for the many roles you play in my life (mentor, role-model, teammate, colleague, and dear friend). I can see so much of you in the physicist I have become and I am so proud of those parts of me. Thank you for always valuing me as a person first and for being someone I can trust completely. You have contributed so significantly to my education that this thesis is also yours to celebrate. My greatest hope for the future is that one day I can do for someone else what you have done for me.

For Brittany.

The greatest thing I will ever be, I have always been.

Your twin.

Contents

1	Introduction	1
2	Background	3
2.1	The Standard Model of Particle Physics	3
2.2	Quantum Chromodynamics	4
2.2.1	Jet Production Cross Section	6
2.2.2	Non-perturbative QCD	8
2.2.3	The Quark-Gluon Plasma	9
2.3	Heavy-Ion Collisions	10
2.3.1	Evolution of a heavy-ion collision	10
2.3.2	Geometry of a heavy-ion collision	13
2.3.3	Experimental Facilities	15
2.3.4	Experimental Probes of the QGP	18
2.4	Jets in Vacuum	21
2.4.1	Theoretical Definition of a Jet	21
2.4.2	Jet Finding Algorithms	24
2.4.3	Jet as a Proxy for the Parton	26
2.4.4	Event Generators in pp	28
2.4.5	Jet Measurements in Vacuum	30
2.5	Jets in Heavy-ion Collisions	33
2.5.1	Jets as a Probe of the QGP	33
2.5.2	Expectations of energy loss	34
2.5.3	Experimental Observables	36
2.5.4	Modeling Jet Quenching	39
2.6	Machine Learning	42
2.6.1	What is Machine Learning?	42
2.6.2	Decision Trees and Neural Networks	45
2.6.3	ML for jets	48

3	The ALICE Detector	50
3.1	Overview	50
3.2	ALICE Tracking	52
3.2.1	Inner Tracking System	52
3.2.2	Time Projection Chamber	52
3.2.3	Track Reconstruction	53
3.2.4	Particle Identification	54
3.3	ALICE EMCal	54
3.4	ALICE in Run 3	56
4	Online Calibration of the Electromagnetic Calorimeter	58
4.1	Bad-Channel Calibration in Run 2	59
4.2	Bad-Channel Calibration in Run 3	60
4.2.1	Differences between Run 2 and Run 3 Approach	60
4.2.2	Online Bad-Channel Calibration	60
4.2.3	Validation of the Run 3 Calibration	61
5	R-dependence of the R_{AA}	63
5.1	Motivation and Previous Measurements	63
5.2	Data Selection	65
5.2.1	Track Selection	65
5.2.2	Embedding	66
5.3	Jet Reconstruction	67
5.4	Background Subtraction	67
5.5	ML-Based Background Subtraction	69
5.5.1	Introduction	69
5.5.2	Training Data Set	70
5.5.3	Regression Target	71
5.5.4	Input Parameters	72
5.5.5	ML Methods	73
5.5.6	Model Investigations	74
5.5.7	Performance	75
5.6	Jet Performance	76
5.7	Unfolding	77
5.7.1	Introduction	77
5.7.2	Forming the Response Matrix	80
5.7.3	Checks of the Unfolding Procedure	82
5.8	Quantification of the Fragmentation Bias	84
5.8.1	Using JEWEL Fragmentation for Training	87
5.9	Systematic Uncertainties	89

5.9.1	Sources of Systematic Uncertainty for Charged Jets	89
5.9.2	Combining Systematic Uncertainties for the R_{AA} Ratios	90
5.10	Results	92
5.10.1	Inclusive Jet Transverse Momentum Spectra	92
5.10.2	Nuclear Modification Factors	94
5.10.3	Jet Cross-Section Ratios	95
5.10.4	Nuclear Modification Factor Double Ratios	97
5.11	Comparisons to Theory	99
6	Full Jet R_{AA}	102
6.1	Data Selection and Embedding	102
6.1.1	Cluster Selection	102
6.2	Jet Reconstruction	103
6.3	Background Subtraction	104
6.4	ML-Based Background Subtraction	104
6.4.1	Training Data Set	104
6.4.2	Input Parameters	104
6.4.3	Model Investigations	105
6.4.4	Performance	107
6.5	Jet Performance	108
6.6	Sources of Systematic Uncertainty for Full Jets	108
6.7	Results	109
6.7.1	Inclusive Jet Transverse Momentum Spectra	109
6.7.2	Nuclear Modification Factors	110
6.8	Fragmentation Bias	111
6.9	Theoretical Comparisons	113
7	Interpreting the R-dependence of the R_{AA}	115
7.1	Simulation Setup	115
7.2	Steepness of the Jet Spectra in pp	116
7.3	Flavor-dependent Quenching	117
8	Conclusions and Future Work	119
8.1	Conclusions	119
8.2	Future Work	120
8.2.1	Direct Extensions of this Result	120
8.2.2	Overcoming the Survivor Bias	121
8.2.3	New Observables	121

A	Additional Checks of the R-dependence Result	122
A.1	Checking Correlation of Input Features	122
A.2	Influence of a Flat Training Spectrum	123
A.3	Embedded MC and Thermal Toy Comparisons	124
A.4	Investigations on the ATLAS Discrepancy	125
B	Analysis Code	129
B.1	Charged-Particle Jet Analysis Code	129
B.2	Full Jet Analysis Code	130
C	Supplemental Figures for the Charged Jet Analysis	131
C.1	Performance	131
C.2	Jet Performance Plots	132
C.3	Unfolding Plots	134
C.3.1	Trivial Closure Test	134
C.3.2	Split MC Test	136
C.3.3	Stability with the Number of Iterations	137
C.3.4	Refolded to Raw Distributions	138
C.3.5	Pearson Correlation Coefficients	139
C.4	Systematics	140
	Bibliography	145

List of Figures

2.1	Summary of the standard model of particle physics, originally appearing in [1].	4
2.2	The three possible QCD interaction vertices, where spring-like lines represent gluons and straight lines represent quarks.	5
2.3	Summary of experimental measurements of the QCD coupling constant, α_s as a function of the momentum transferred Q , originally taken from [2]. M_z is the mass of the Z boson.	7
2.4	The factorization of the jet production cross section as written in Equation 2.5.	8
2.5	PDFs calculated at NNLO for $Q = 3.2$ GeV (left) and $Q = 100$ GeV (right), originally appearing in [3].	8
2.6	Space-time grid used for calculations in Lattice QCD.	9
2.7	Left: Phase diagram of QCD matter as a function of temperature and baryon chemical potential [4]. Right: Lattice QCD calculations of thermodynamic properties compared to hadron resonance gas calculations [5].	10
2.8	The evolution of a heavy-ion collision, originally appearing in [6].	11
2.9	Sample initial energy density profiles for the MC Glauber model (left) and the IP Glasma model (right), originally appearing in [7].	12
2.10	(Top): Shear and bulk viscosity. (Bottom): Radial and anisotropic flow. .	13
2.11	(Left): Geometry both before and after a heavy-ion collision. The impact parameter, b , as well as spectators (gray) and participants (in color) are shown. (Right): Side view of colliding nuclei for central and peripheral collisions with impact parameter b	14
2.12	V0M signal amplitude with a Glauber+NBD fit [8], originally appearing in [6].	15

2.13	The accelerator complex of the Large Hadron Collider (not to scale), adapted from [9] last updated in 2022. The path taken by various species as they enter the different complexes is outlined on the sides of the image, where in the case of heavy ions units of energy are given per atomic mass unit (u) for each beam.	17
2.14	Visual representation of the Fourier decomposition of the flow harmonics. Here Φ_n reflects the n^{th} order harmonic. Originally appearing in [10].	19
2.15	The hard scattering and subsequent fragmentation and hadronization of partons, appearing in [11].	22
2.16	Parton splitting process for $i \rightarrow jk$ with the momentum fraction z , azimuthal angle ϕ , and opening angle θ	22
2.17	Comparison of spatial distribution (ϕ , y) of jets found in a parton-level event using the k_T (left), C/A (middle), and anti- k_T (right) sequential recombination algorithms, adapted from [12].	25
2.18	Illustrations of two different hadronization models. Left: the string hadronization model and Right: the cluster hadronization model. Figure originally appearing in [13].	29
2.19	Left: ALICE pp jet cross-section ratios of $R = 0.2$ jets to other radii, originally appearing in [14]. These results are additionally compared to NLO+NLL+NP predictions [15–17]. Right: CMS pp cross sections of jets of different R as compared to $R = 0.4$ (here abbreviated AK4) jets, originally appearing in [18]. These results are additionally compared to NLO predictions [19].	31
2.20	The primary Lund jet plane.	32
2.21	Left: ALICE fully-corrected primary Lund plane density for $R = 0.4$ jets with $20 < p_{T,\text{jet}} < 120$ GeV/c in pp collisions at $\sqrt{s} = 13$ TeV, originally appearing in [20]. Right: ATLAS fully-corrected Lund plane for $R = 0.4$ jets with $p_T > 675$ GeV/c in pp collisions at $\sqrt{s} = 13$ TeV, originally appearing in [21].	33
2.22	(Left): Back-to-back jets in pp collisions and heavy-ion(AA) collisions. Jets in pp exist in vacuum whereas jets in AA traverse the QGP medium and lose energy. (Right): Feynman diagrams for the energy loss ΔE of a quark with energy E due to collisional (left diagram) and radiative (right diagram) processes, originally appearing in [22].	34
2.23	Momentum broadening (left) and the deflection of the jet centroid (right). The image for the background is from the MADAI collaboration.	35
2.24	Decoherence (left) vs. coherence (right).	36
2.25	The jet p_T spectrum in pp and heavy-ion (A+A) collisions and the origin of jet suppression.	37

2.26	The weakly-coupled medium response via recoils (left) and a strongly-coupled medium response via a hydrodynamic wake (right).	40
2.27	A categorization scheme of various different available jet quenching models.	41
2.28	The general workflow for a given ML application.	43
2.29	True positive rate as a function of the false positive rate, also called the ROC curve.	45
2.30	General structure of a feed-forward neural network, where the arrows represent the direction of the flow of information.	47
3.1	Schematic view of the ALICE detector during Run 2 of the LHC [23]. . .	51
3.2	Left: Schematic of the ALICE TPC, originally appearing in [24]. Right: dE/dx (see Equation 3.1) distribution as measured in the TPC as a function of momentum, originally appearing in [25]. The black lines correspond to the expected mean energy loss according to Equation 3.1.	53
3.3	Schematic of ALICE PID and reconstruction where the p_T ranges shown correspond to those in published measurements, adapted from [6].	55
3.4	The ALICE EMCal (top) and DCal (bottom), originally appearing in [26]. The PHOS calorimeter is shown in tan near the DCal.	56
4.1	Bad-channel map as calculated from the same input data using the Run 2 procedures (left) vs. the Run 3 procedure (right). Cells filled green indicate those labeled as good, those filled gray indicate those labeled as bad and cells filled red indicate those marked as bad.	62
4.2	Left: A comparison between the Run 2 and the Run 3 calibrations, cells marked in green are cells marked as good in both calibrations, cells marked in red are cells marked as bad in both calibrations, and cells marked in gray are flagged differently in each calibration. Right: A comparison of the energy distribution before the masking of bad channels (blue), after masking in the Run 2 calibration (red) and after masking in the Run 3 calibration (green).	62
5.1	Previous measurements of the jet R_{AA} for $R = 0.4$ 0–10% jets from RHIC [27] and the LHC [14, 28, 29].	64
5.2	Azimuthal distribution for selected hybrid tracks before the fiducial acceptance cut.	66
5.3	Residual p_T -distributions of embedded jet probes of known transverse momentum into Pb–Pb collision data. Left: Comparison of the distributions for the area-based and ML-based background estimators. Note the lines connecting the points do not represent a fit and are only present to guide the eye. Right: Radius dependence of the width of the distributions, where the error bars come from the uncertainty in calculating the width.	69

5.4	Transverse momentum distribution for charged tracks from the toy model.	70
5.5	(Left): Comparison between the track p_T distribution within jets in the toy model and the embedded MC for $R = 0.4$ charged jets. (Right): Distribution of true p_T 's for the training sample for $R = 0.6$ jets as seen by the ML.	71
5.6	The Pearson correlation of input parameters for central (0-10%) $R = 0.6$ jets. Note that the names of variables here correspond to how they are written in the analysis code. See Appendix B for more details.	73
5.7	Visualization of a decision tree for the random forest estimator trained on $R = 0.6$ jets in central (0-10%) collisions. For visualization purposes, only the first few layers of the decision tree are included.	75
5.8	The δp_T distributions for $R = 0.4$ (left) and $R = 0.6$ (right) jets in central (0-10%) collisions using the area-based (dashed lines) and ML-based (solid markers) corrections.	76
5.9	Left: JES for $R = 0.6$ jets in 0–10% central collisions. Right: JER for $R = 0.6$ jets in 0–10% central collisions	77
5.10	Overview of how comparisons between data and theory jets are made. . .	78
5.11	Jet reconstruction efficiency for $R = 0.6$ jets.	81
5.12	The hybrid distributions for the various p_T hard bins for $R = 0.6$ jets in central (0-10%) collisions. These distributions are plotted without the $p_{T,\text{part}} > 10 \text{ GeV}/c$ cut (left) and with the $p_{T,\text{part}} > 10 \text{ GeV}/c$ cut (right). .	81
5.13	Left: (Top) folded distributions for each iteration of the Bayesian unfolding procedure for the Trivial Test. (Bottom) Ratio of the unfolded to the true distribution as a function of the number of iterations for $R = 0.6$ charged jets in the 0–10% most central Pb–Pb collisions. Right: (Top) Unfolded distributions for each iteration of the Bayesian unfolding procedure for the split MC test. (Bottom) Ratio of the unfolded to the true distribution for $R = 0.4$ 0–10% full jets where the reconstructed distribution that we unfold is made up of the jets that make up the response. . . .	82
5.14	Left: (Top) Refolded distribution as a function of the iteration of the Bayesian unfolding procedure. (Bottom) The ratio of the refolded to raw distribution for $R = 0.6$ charged jets in the 0–10% most central Pb–Pb collisions. Right: (Top) Unfolded distribution as a function of the iteration of the Bayesian unfolding procedure. (Bottom) Ratio of the unfolded result at a given iteration to the result obtained with 8 iterations for $R = 0.6$ charged jets in the 0–10% most central collisions.	83
5.15	The refolded spectrum of the ML-based $R = 0.6$ result in 0-10% collisions as compared to the reconstructed area-based spectrum.	84

5.16	Left: The R_{AA}^{toy} distribution for $R = 0.4$ jets using the various modifications. Right: The fragmentation functions at low jet p_T ($40 < p_{T,\text{true}} < 100$ GeV/ c , lower right panel) and high jet p_T ($100 < p_{T,\text{true}} < 200$ GeV/ c , upper right panel) for 0–10% central Pb–Pb collisions. The ratio of the fragmentation functions measured in Pb–Pb and pp collisions are shown for jets with $R = 0.4$ and $p_T > 100$ GeV/ c [30] (ATLAS), and for jets with $R = 0.3$ and $p_T > 30$ GeV/ c recoiling from a photon with $E_T > 60$ GeV/ c [31] (CMS).	87
5.17	δp_T distributions for $R = 0.4$ jets both without recoils (left) and with recoils (right) for the area-based (AB) and ML-based (ML) background correction methods. In the right panel vacuum JEWEL, similar to PYTHIA, is also shown.	88
5.18	Nuclear modification factor for $R = 0.4$ charged-particle jets in 0–10% Pb–Pb collisions obtained with the ML-based background correction trained on PYTHIA, compared to results obtained with the ML-based background correction trained on JEWEL both with and without recoils.	89
5.19	The charged-particle jet p_T spectra in pp collisions for $R = 0.2, 0.3, 0.4, 0.5, 0.6$. The vertical bars denote statistical uncertainties and the vertical extent of the boxes denotes systematic uncertainties.	93
5.20	The p_T -differential inclusive charged-particle jet yield distributions as a function of p_T for different values of R in three centrality classes: Left: 0–10% collisions, Middle: 30–50%, Right: 60–80% collisions. The peripheral spectra were measured using the AB method for the background correction. All other reported spectra were corrected with the ML-based background estimator.	93
5.21	Nuclear modification factors of inclusive charged-particle jets as a function of p_T for $R = 0.2, R = 0.4$, and $R = 0.6$, shown for 0–10%, 30–50% and 60–80% central Pb–Pb collisions for the ML-based method compared to results obtained with the area-based method where applicable.	94
5.22	Nuclear modification factors for jets with $R = 0.6$ in 0-10% (left) and 30-50% (right) central Pb–Pb collisions outlining the impact of the various fragmentation models on the final result. Note that the systematic uncertainties are drawn both with and without the fragmentation uncertainties in the empty and filled boxes, respectively.	96
5.23	Jet cross-section ratios for $\sigma(R = 0.2)/\sigma(R = 0.4)$ (left) and $\sigma(R = 0.2)/\sigma(R = 0.6)$ (right).	96
5.24	Double ratio of jet nuclear modification factors using $R_{AA}^{R=0.2}$ as the denominator and using $R = 0.4$ (left) and $R = 0.6$ (right) as the numerator for central (top row) and semi-central (bottom row) collisions.	98

5.25	Nuclear modification factors for $R = 0.2$, $R = 0.4$, and $R = 0.6$, shown for 0–10% and 30–50% central Pb–Pb collisions compared to theoretical calculations incorporating jet quenching (see text for details).	100
5.26	Double ratio of jet nuclear modification factors using $R_{AA}^{R=0.2}$ as denominator and using $R = 0.4$ (left) and $R = 0.6$ (right) as the numerator compared to model predictions for central (top row) and semi-central (bottom row) collisions. Note that a comparison to JEWEL with recoils was omitted from the right-hand plot as its prediction is out of scale.	101
6.1	Comparison of the δp_T distributions for $R = 0.4$ jets in central collisions (0–10%) trained using the tracks and clusters selected separately and selecting from a full constituent list for the neural network estimator. The AB method is shown for the purposes of comparison.	105
6.2	(Left): Comparison of the δp_T distributions for $R = 0.4$ jets in central collisions (0–10%) trained using the matching regression target and the true p_T fraction regression target (see Section 5.5.3) for the ML estimator. (Right): Comparison of the δp_T distributions for $R = 0.4$ jets in central collisions (0–10%) trained using the simplified input parameter list and the full parameter list (see Section 5.5.4) for the ML estimator. In both figures, the AB method is shown for the purposes of comparison.	106
6.3	(Left): δp_T distribution for full jets with $R = 0.4$ in central (0–10%) collisions. Comparison of various ML-based and AB methods. (Right): Comparison of the width of the δp_T distributions for charged jets with several resolution parameters. Comparison of ML-based and area-based correction. The results in both panels are shown for $p_{T,\text{det}} > 40$ GeV/c to avoid edge effects to get a symmetric distribution even in the case of strong fluctuations.	107
6.4	JES and JER (as defined in Equations 5.5 and 5.6 respectively) for $R = 0.4$ jets in central (0–10%) collisions.	108
6.5	The relative uncertainties of various systematic uncertainties.	109
6.6	The full jet p_T spectra in pp collisions for $R = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6$, originally appearing in [14].	110
6.7	Left: The p_T -differential inclusive full jet yield for $R = 0.4$ jets in the 0–10% most central Pb–Pb collisions for results obtained with the ML-based and area-based correction methods. Right: Nuclear modification factor of inclusive full jets as a function of p_T for $R = 0.4$ for 0–10% central Pb–Pb collisions for the ML-based method compared to results obtained with the area-based method where applicable.	111

6.8	Ratio of the modified to unmodified fragmentation functions for the three different modifications presented in Section 6.8. These modifications are shown for both the eight leading particles (left panel) and all particles (right panel).	112
6.9	The nuclear modification factor with ML-based and area-based corrections are shown for $R = 0.4$ full jets in 0-10% collisions. Left: Comparison of the fragmentation bias curves for a toy model with three different modifications. Right: Comparison to various models. The models that are shown here also do not have this leading track bias requirement.	114
7.1	Left: The jet cross section ratios in PYTHIA simulations as compared to ALICE pp data. Right: PYTHIA jet p_T spectra generated for a variety of R for charged-particle jets and full jets.	117
7.2	Quark fraction in PYTHIA as a function of R	118
7.3	Predictions for R_{AA} ratios for $R = 0.4/R = 0.2$ (left) and for $R = 0.6/R = 0.2$ (right) from the flavor-independent quenching model and the flavor-dependent quenching model as compared to the ALICE data. . . .	118
A.1	Visual representation of the correlation of parameters for central (0–10%) $R = 0.6$ jets.	123
A.2	(Left): Training distribution for the completely flat training test. (Right): Performance comparison of the flat and nominal training samples.	124
A.3	(Top): The unfolded result achieved for $R = 0.6$ jets in central collisions using the nominal training sample and a flat training sample. (Bottom): Ratio of the two unfolded distributions in the top panel.	124
A.4	(Top) Comparison of the fully-unfolded results for $R = 0.6$ jets in central collisions using both the thermal background (with the true p_T fraction as the target) and the embedded MC background (with the geometrically-matched detector-level jet p_T as the target) (Bottom): The ratio between these two quantities.	125
A.5	Comparison of different measurements of the R -dependence of jet suppression by the ALICE (this thesis), ATLAS, and CMS collaborations. Figure courtesy of Prof. Dennis Perepelitsa.	126
A.6	Quark fraction in both ALICE (left) and ATLAS (right) kinematics.	127
A.7	The jet cross section ratios for $\sigma(R = 0.2)/\sigma(R = 0.4)$ (left) and $\sigma(R = 0.2)/\sigma(R = 0.6)$ for the nominal ($p_{T, \text{const}} > 150 \text{ MeV}/c$) and ATLAS ($p_T > 700 \text{ MeV}/c$). In the right panel, the ATLAS results are shown correcting back to 700 MeV/c and 0 MeV/c in the unfolding procedure.	128
C.1	δp_T distributions for $R = 0.4$ central (0-10%) jets in various bins of $p_{T, \text{det}}$	131
C.2	δp_T distributions for $R = 0.6$ central (0-10%) jets in various bins of $p_{T, \text{det}}$	131

C.3	The JES (as defined in Equations 5.5) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	132
C.4	The JER (as defined in Equations 5.6) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	132
C.5	The JES (as defined in Equations 5.5) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions. . . .	133
C.6	The JER (as defined in Equations 5.6) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions. . . .	133
C.7	The JES (as defined in Equations 5.5) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions. . . .	133
C.8	The JER (as defined in Equations 5.6) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions. . . .	134
C.9	Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	134
C.10	Trivial closure tests for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.	135
C.11	Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.	135
C.12	Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	136
C.13	Trivial closure tests for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.	136
C.14	Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.	136
C.15	Unfolding stability for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	137
C.16	Unfolding stability for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.	137
C.17	Unfolding stability for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.	137
C.18	Refolded to raw ratio for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	138
C.19	Refolded to raw ratio for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.	138
C.20	Refolded to raw ratio for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.	138
C.21	Pearson coefficients for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.	139

C.22	Pearson coefficients for $R = 0.6$ jets corrected by the ML-based method in central (0-10%) collisions with iterations 6, 7, 9, 10. Note that results with 8 iterations can be found in Figure C.21.	139
C.23	Pearson coefficients for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.	139
C.24	Pearson coefficients for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in peripheral(60-80%) collisions.	140
C.25	Relative systematic uncertainties for jet spectrum, $R = 0.2$ (left: 0-10%, right: 30-50%). The total uncertainty, which is the quadratic sum of the individual uncertainties, is also shown.	140
C.26	Relative systematic uncertainties for jet spectrum, $R = 0.4$ (left: 0-10%, right: 30-50%). The total uncertainty, which is the quadratic sum of the individual uncertainties, is also shown.	140
C.27	Relative systematic uncertainties for jet spectrum, $R = 0.6$ (left: 0-10%, right: 30-50%). The total uncertainty, which is the quadratic sum of the individual uncertainties, is also shown.	141

List of Tables

3.1	Description of the various sub-detectors of the ALICE detector. Numbers in parenthesis (ex: (1)) correspond with the numbering scheme in Figure 3.1.	51
5.1	Features utilized in the random forest model trained on $R = 0.4$ jets in central (0–10%) collisions as well as their feature importances, gathered using <code>scikit-learn</code>	74
5.2	Relative systematic uncertainties (%) for jet spectra for 0–10% central Pb–Pb collisions and all resolution parameters. The maximum uncertainties for low p_T ($p_{T,jet} < 50$ GeV/ c) and high p_T ($p_{T,jet} > 50$ GeV/ c) are shown. The direction of asymmetric uncertainties is indicated with a + or – sign. The combined uncertainty is the quadrature sum of individual uncertainties.	91
5.3	Relative systematic uncertainties (%) for jet spectra for 30-50% central Pb–Pb collisions and all resolution parameters. The maximum uncertainties for low p_T ($p_{T,jet} < 50$ GeV/ c) and high p_T ($p_{T,jet} > 50$ GeV/ c) are shown. The direction of asymmetric uncertainties is indicated with a + or – sign. The combined uncertainty is the quadrature sum of individual uncertainties.	91
5.4	Relative systematic uncertainties (%) for jet spectra for 60-80% central Pb–Pb collisions and all used resolution parameters. The maximum uncertainties for low p_T ($p_{T,jet} < 50$ GeV/ c) and high p_T ($p_{T,jet} > 50$ GeV/ c) are shown. In this centrality interval the spectra are measured with the area-based method, the uncertainties related to the fragmentation functions adopted in the machine learning algorithm are not included for this case. The direction of asymmetric uncertainties is indicated with a + or – sign. The combined uncertainty is the quadrature sum of individual uncertainties.	92

6.1	Percentage of jets which contain a leading, sub-leading, and sub-sub-leading track or cluster compared in embedded MC and data from $R = 0.4$ jets in 0-10% central collisions.	106
A.1	Values for different experimental parameters for the ALICE, ATLAS, and CMS measurements shown in Figure A.5. Note that for CMS the constituent $p_{T,\min}$ was not studied for this comparison, and is therefore listed as unknown.	127

Chapter 1

Introduction

Arguably, the most important part of being a physicist is not what you know, but what you can figure out. Over the course of human history, both physicists and non-physicists alike have managed to figure out a great deal about the universe. The composition of the universe has been studied at both its smallest scales of $\sim 10^{-18}$ m and its largest scales of 13.7 billion light years. These milestone scientific achievements cannot be attributed to a single “lone-genius”, but represent the combined effort of many.¹

Today, physicists around the world are collaborating to discover new aspects of the universe, the pieces of which it is composed, and how those pieces came to be. One of the largest collaborative scientific efforts to ever exist occurs at the [Conseil Européen pour la Recherche Nucléaire \(CERN\)](#) where thousands of physicists analyze collisions from the [Large Hadron Collider \(LHC\)](#) to extend human knowledge. There are four main experiments at the LHC, each with complimentary physics capabilities and goals. “[A Large Ion Collider Experiment](#)” ([ALICE](#)) is one of these four experiments, optimized for the study of the collisions of heavy lead ions. These collisions recreate the conditions of the early universe, where matter existed as a soup called the [Quark-Gluon Plasma \(QGP\)](#). Scientists at the LHC study the [QGP](#) to gain insight into how the fundamental building blocks of the universe combine to form everyday matter.

Sometimes in high-energy collisions, the fundamental building blocks of matter collide, creating a cone-shaped spray of particles called a jet. When jets are formed in the collisions of heavy ions, they interact with the [QGP](#). Theory predicts that these interactions should result in an overall energy loss of the jet and a modification of its substructure, which has been confirmed by experimental measurements. However, the precise interplay of the various physical mechanisms that contribute to these observations is unknown. In this thesis, a measurement of the dependence of the overall energy loss (called jet suppression) on the cone size of the jet, which aims to determine the relative impact of these mechanisms, will be discussed. To do this, a novel [machine learning \(ML\)](#) method

¹See, “The Lone Genius Paradigm and Our Infatuation With Intellectual Heroes” by Ágnes Mócsy.

was employed, marking the first application of **ML** to an experimental jet measurement in heavy-ion collisions and enabling the most differential measurement of jet suppression ever measured by the **ALICE** collaboration. This measurement indicates that wider jets lose more energy, shedding light on how the substructure of a jet could impact the way it loses energy. Some simulations of jets in heavy-ion collisions can also reproduce this trend, highlighting the importance of how a jet's interaction with the **QGP** is modeled.

This thesis is broken into eight different chapters, of which this general introduction is the first. Background information relevant to the content of this thesis will be discussed in Chapter 2. The **ALICE** detector will be discussed in Chapter 3. The online calibration of the **ALICE** Electromagnetic Calorimeter will be discussed in Chapter 4. A measurement of the dependence of jet suppression on the cone size of the jet for charged-particle jets using a novel **ML** technique will be discussed in Chapter 5. A measurement of jet suppression for full jets using this novel **ML** technique will be discussed in Chapter 6. Simulations performed to understand the results of this thesis will be discussed in Chapter 7. Finally, a discussion of these results and an outlook for future measurements will be discussed in Chapter 8. Similar to the results presented in this thesis, the cover also employs machine learning and was generated at <https://dream.ai/create> using the title of this thesis and an **ALICE** event display as input.

Chapter 2

Background

2.1 The Standard Model of Particle Physics

All visible matter in the universe consists of fundamental particles governed by four known fundamental forces: the gravitational force, the electromagnetic force and the weak force,¹ and the strong force. Each of these forces acts over different length scales with varying strengths. The strong force, aptly named, is the strongest of the four forces, acting over very small distances of 10^{-15} m or 1 fm. The electromagnetic force is about 100 times weaker² than the strong force, but acts over an infinite range. The weak force is 1 million times weaker than the strong force and acts over small distances of 10^{-3} fm. The gravitational force is significantly weaker than the other fundamental forces (about 10^{43} times weaker) and also acts over an infinite range.

The Standard Model of particle physics is the fundamental theory that describes three of these four fundamental forces (the electromagnetic, the weak, and the strong force) and all known particles as quantized fields and their excitations. It is the most precise physical theory to date, agreeing with experimental measurements across 10 orders of magnitude [32]. This section will outline the basics of the Standard Model. The Standard Model is organized according to the categories shown in Figure 2.1. Fundamental particles are sorted into two main groups: fermions that have half-integer spin and obey Fermi-Dirac statistics and bosons that have integer spin and obey Bose-Einstein statistics. Fermions are then subdivided into two categories: quarks (and anti-quarks) that have a color charge and participate in the strong interaction and leptons that carry no color charge and participate in electroweak interactions. There are six types (also called flavors) of quarks and six flavors of leptons sorted into pairs of particles exhibiting similar physical behavior called

¹The electromagnetic and the weak forces are low-energy manifestations of the unified “electroweak” interaction.

²The notion of the “strength” of a force is ambiguously defined, so different numbers for these relative strengths may be quoted elsewhere.




Figure 2.1: Summary of the standard model of particle physics, originally appearing in [1].

generations, where lighter and more stable particles are in the first generation and heavier particles are in the second and third generations. Bosons are then also sorted into two categories: gauge (or vector) bosons that have spin 1 and are force carriers, and scalar bosons that have spin 0.

2.2 Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the quantum field theory of the strong force, which mediates the interaction between partons (quarks and gluons). Gluons mediate strong interactions while also carrying color charge, a striking difference from the quantum field theory of electromagnetism, **Quantum Electrodynamics (QED)**, which is mediated by the photon that carries no electric charge. As a result of this feature of **QCD**, in nature quarks and gluons are not observed in isolation and instead are confined inside color-neutral mesons ($\bar{q}q$ bound state) and baryons (qqq bound state) that are collectively referred to as hadrons. **QCD** is a part of the $SU(3)$ symmetry group, meaning that the theory is invariant under local $SU(3)$ transformations. The Lagrangian for **QCD** is written in Equation 2.1. Here $\psi_{q,a}$ are the spinors for a quark field with flavor q , color a , and mass m_q . In




Figure 2.2: The three possible **QCD** interaction vertices, where spring-like lines represent gluons and straight lines represent quarks.

this theory, the number of colors N_c is three³ and the number of possible flavors is six.⁴ The Dirac γ matrices are represented by γ^μ [33]. The quantity g_s is related to the strong coupling constant as $\alpha_s = \frac{g_s^2}{4\pi}$, which is the **QCD** analog of the fine-structure constant in **QED**.

$$\mathcal{L}_{\text{QCD}} = \sum_q \bar{\psi}_{q,a} (i\gamma^\mu \partial_\mu \delta_{ab} - g_s \gamma^\mu t_{ab}^C A_\mu^C - m_q \delta_{ab}) \psi_{q,b} - \frac{1}{4} F_{\mu\nu}^A F_A^{\mu\nu}. \quad (2.1)$$

The four terms in the Equation 2.1 correspond to the free quark kinetic term, the quark-gluon interaction term, the mass term for a quark of flavor q with mass m_q , and the gluon kinetic term, respectively. The strength of the quark-gluon interaction in the second term of Equation 2.1 is specified by g_s . The kinetic term of the gauge gluon field (the fourth term in Equation 2.1), contains the gluon field tensor defined in Equation 2.2, $F_A^{\mu\nu}$. The gluon field tensor is analogous to the electromagnetic tensor in **QED**, but with an additional term to account for the gluon carrying a color charge that allows for the gluon self-interaction vertices shown in Figure 2.2.

$$F_{\mu\nu}^A = \partial_\mu A_\nu^A - \partial_\nu A_\mu^A - g_s f^{ABC} A_\mu^B A_\nu^C \quad (2.2)$$

The **QCD** Lagrangian leads to three possible interaction vertices in **QCD**, a 3-gluon and 4-gluon vertex as well as a quark-anti-quark gluon vertex, which are shown (respectively from left to right) in Figure 2.2.

When **QCD** theory is renormalized, this renormalization depends on a corresponding renormalization scale. When this scale is taken to be close to the momentum transfer in a given process (Q), $\alpha_s(Q)$ as defined in Equation 2.3 represents the strength of the strong interaction in that given process. Here, Λ_{QCD} is a characteristic momentum scale of QCD

³Red, green, and blue.

⁴Up, down, charm, strange, top, and bottom.

that takes an approximate value of 200 MeV [33].⁵

$$\alpha_s(Q) = \frac{2\pi}{(11 - \frac{2}{3}n_f)\log(\frac{Q}{\Lambda_{\text{QCD}}})} \quad (2.3)$$

Figure 2.3 shows α_s as a function of Q . A distinguishing feature of QCD, called *asymptotic freedom*, is that the coupling becomes asymptotically weaker for high energies (small distance scales, high Q^2). The coupling increases for low energies (large distance scales, small Q^2). This can also be seen from the effective QCD potential as a function of the distance r between quarks, shown in Equation 2.4. At large distances (low energies), this potential grows linearly with r such that it eventually becomes favorable to form a new $q\bar{q}$ pair from the vacuum, leading to the confinement of partons inside hadrons.

$$V(r) = -\frac{4}{3} \frac{\alpha_s}{r} + kr \quad (2.4)$$

At large Q , where $\alpha_s \ll 1$, QCD predictions can be performed using a perturbative expansion in powers of α_s . Often, only the first, second, or third orders in α_s , referred to as **Leading Order (LO)**, **Next to Leading Order (NLO)**, and **Next to Next to Leading Order (NNLO)** respectively, are computed. When the coupling is $\mathcal{O}(1)$ the perturbative approach breaks down and non-perturbative techniques (discussed in Section 2.2.2) must be applied. The perturbative terms can be approximately⁶ factorized from the non-perturbative ones as they happen at different energy scales, called the QCD Factorization Theorem.

2.2.1 Jet Production Cross Section

A particularly relevant QCD calculation to this thesis is that of the production cross section for inclusive jets (jets will be described in Section 2.4.1) in **proton-proton (pp)** collisions, which can be factorized in the limit of sufficiently small jet cone sizes as in Equation 2.5 from Ref. [34] and Figure 2.4. Each of the individual terms of this equation will be described below, where \otimes represents integrals over the parton momentum fractions.

$$\frac{d\sigma^{\text{pp} \rightarrow \text{jet} + \text{X}}}{dp_T d\eta} = \sum_{ab} f_a(x_a, Q^2) \otimes f_b(x_b, Q^2) \otimes \left(\sum_c \hat{\sigma}_{ab \rightarrow c}(z, \mu) \otimes J_c \right) \quad (2.5)$$

The terms $f_a(x_a, Q^2)$ and $f_b(x_b, Q^2)$ represent the probability of finding a parton inside of a nucleon with a momentum fraction $x_i = p_{\text{parton},i}/p_{\text{nucleon}}$ given a momentum transfer Q^2 , described by the **Parton Distribution Function (PDF)**. PDFs are non-perturbative, but can be evolved to another scale using perturbative evolution equations, referred to as the **Dokshitzer Gribov Lipatov Altarelli Parisi (DGLAP)** evolution equations named for

⁵Different values for Λ_{QCD} , such as $\Lambda_{\text{QCD}} \approx 1$ GeV, may be quoted depending on the source [2].

⁶See <https://royalsocietypublishing.org/doi/10.1098/rsta.2021.0058> for more detail as to why this is approximate.




Figure 2.3: Summary of experimental measurements of the QCD coupling constant, α_s as a function of the momentum transferred Q , originally taken from [2]. M_Z is the mass of the Z boson.

its authors [35–37]. PDFs can be experimentally determined via deep inelastic scattering processes where electrons are scattered off of protons. Figure 2.5 shows the PDFs calculated as a function of x for different partons at $Q = 3.2$ GeV and $Q = 100$ GeV. At very low values of x , the gluon PDF is dominant whereas at high values of x the valence quark PDFs (particularly u and d quarks) dominate, indicating that the soft contributions to a nucleon tend to come from gluons and the hard contributions tend to come from quarks. Eventually, there can be no more gluons and the gluon PDF is expected to saturate, which can be described by a Color Glass Condensate (CGC) [38]. Understanding saturation effects, and their role in the collisions of high-energy hadrons and nuclei, is one of the main physics goals of the upcoming Electron-Ion Collider (EIC).⁷

The term $\hat{\sigma}_{ab \rightarrow c}(z, \mu)$ in Equation 2.5 represents the production of a parton c from the hard scattering of two partons a and b at a hard-scattering scale $\mu \sim p_T$ with $z = p_{T,c}/p_T$. This can be calculated perturbatively at different orders in α_s . The term J_c in Equation 2.5, referred to as the jet function, is the term that describes the formation of the jet from parton c . The DGLAP evolution equations can be used to evolve jet functions from the scale of the jet mass to the scale of the hard scattering.⁸

⁷<https://www.bnl.gov/eic/goals.php>

⁸For the analogous case of the formation of a single hadron from a parton, so-called fragmentation functions are used that non-perturbative in nature and need to be extracted from data.




Figure 2.4: The factorization of the jet production cross section as written in Equation 2.5.




Figure 2.5: PDFs calculated at NNLO for $Q = 3.2$ GeV (left) and $Q = 100$ GeV (right), originally appearing in [3].

2.2.2 Non-perturbative QCD

Non-perturbative techniques become a necessity when α_s is on the $\mathcal{O}(1)$ at very low Q . There are a few existing methods of performing non-perturbative calculations, the most prominent is Lattice QCD. Lattice QCD is a lattice gauge theory formulated on a grid (or lattice) in space and time as shown in Figure 2.6. Quarks are at the lattice sites and gluons are the links. The size of the spacing depends on the energy scale. When the spacing (a) is zero and the size of the lattice (L) is taken infinitely large, the continuum is recovered. Lattice QCD calculations are often applied to the QCD phase diagram and performed for $\mu_B = 0$. For lattice calculations where $\mu_B \neq 0$ other techniques are used to approximate a solution.

Another technique for calculations in non-perturbative QCD (npQCD) is Anti-de Sitter/conformal field theory (AdS/CFT), which is based on a supposed correspondence be-




Figure 2.6: Space-time grid used for calculations in Lattice QCD.

tween [AdS](#) theories of quantum gravity and [CFTs](#). This correspondence relies on a weak-strong duality where calculations that are difficult in strongly-coupled gauge theories are easy in a semi-classical gravitational picture. Refer to [39] for more information about [AdS/CFT](#) correspondence.

One additional way to perform [npQCD](#) calculations is [Soft Collinear Effective Theory](#) ([SCET](#)). [SCET](#) is an effective field theory for soft or collinear particles that can be used for [npQCD](#) calculations of interactions between particles of different energies, namely high-energy quarks with soft gluons. Refer to [40] for more information about [SCET](#).

2.2.3 The Quark-Gluon Plasma

At low temperatures, such as those that exist in nature, [QCD](#) matter exists as confined colorless hadrons. Correspondingly, at sufficiently high temperatures [QCD](#) matter becomes a hot and dense deconfined medium known as the [Quark-Gluon Plasma](#) ([QGP](#)). If the temperature is high enough, [QCD](#) predicts that partons will be weakly coupled with one another as a consequence of asymptotic freedom. This suggests two states of [QCD](#) matter that can be described by a phase diagram and a corresponding phase transition, as shown as a function of temperature and baryon chemical potential (μ_B) in the left panel of Figure 2.7. A [QGP](#) with very low μ_B and a very high temperature is thought to have been formed in the primordial universe lasting until roughly 10^{-6} seconds after the Big Bang. Deconfined quark-matter at very low temperatures and very high μ_B is thought to be formed in the core of neutron stars [41].

The transition from a hadron resonance gas to a [QGP](#) occurs in the non-perturbative regime, where lattice calculations are often used to probe this transition region. These

calculations utilize an equation of state that describes the equilibrium properties of **QCD** matter by relating thermodynamic properties such as temperature (T), pressure (p), and energy density (ϵ). The normalized pressure ($3p/T^4$), normalized energy density (ϵ/T^4) and entropy density ($3s/4T^3$) can be plotted as a function of temperature, as shown in the right panel of Figure 2.7. As a result of the transition from hadronic to partonic states, additional degrees of freedom are created, as seen by the rapid rise in energy density and entropy in Figure 2.7. Lattice calculations also indicate that at $\mu_B = 0$ and a temperature of $T_c \approx 150$ MeV [5] this transition exists as a smooth crossover [42] while at higher values of μ_B and at lower temperatures it exists as a first-order phase transition [43]. An open question is where in the phase diagram the **QCD** critical point, which marks the end point of a first-order phase transition line, exists.

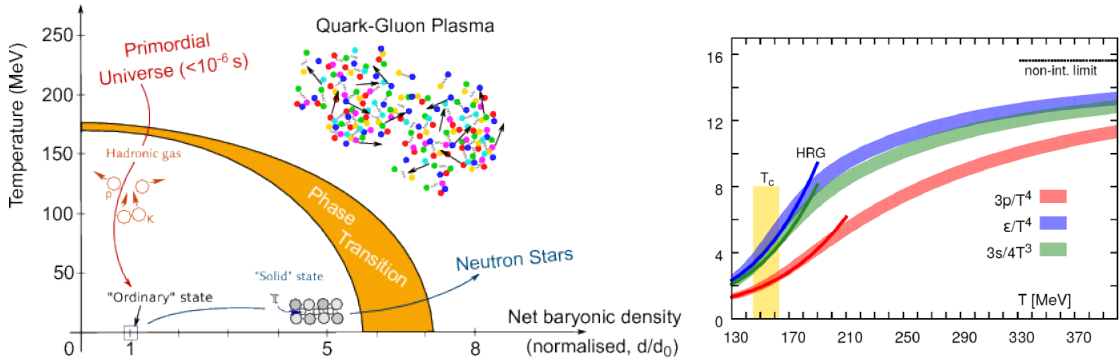


Figure 2.7: Left: Phase diagram of **QCD** matter as a function of temperature and baryon chemical potential [4]. Right: Lattice **QCD** calculations of thermodynamic properties compared to hadron resonance gas calculations [5].

The hot and dense **QGP** medium can be experimentally recreated through the collisions of relativistic heavy ions that generate the highest artificial temperatures on Earth of more than 10^{12} K.⁹ Measurements of collective behavior in heavy-ion collisions indicate a strongly-coupled “perfect” liquid¹⁰ instead of the weak coupling associated with asymptotic freedom. This surprising result indicates that the **QGP** provides a rich laboratory for studying the mysteries of the strong interaction.

2.3 Heavy-Ion Collisions

2.3.1 Evolution of a heavy-ion collision

The evolution of a heavy-ion collision can be broken into factorizable stages as shown in Figure 2.8. In this section, each stage of a heavy-ion collision will be briefly described in

⁹<https://www.guinnessworldrecords.com/world-records/highest-man-made-temperature>

¹⁰The **QGP** has the lowest shear viscosity to entropy density (η/s) ratio ever observed, $\eta/s \sim 1/4\pi$.

chronological order.

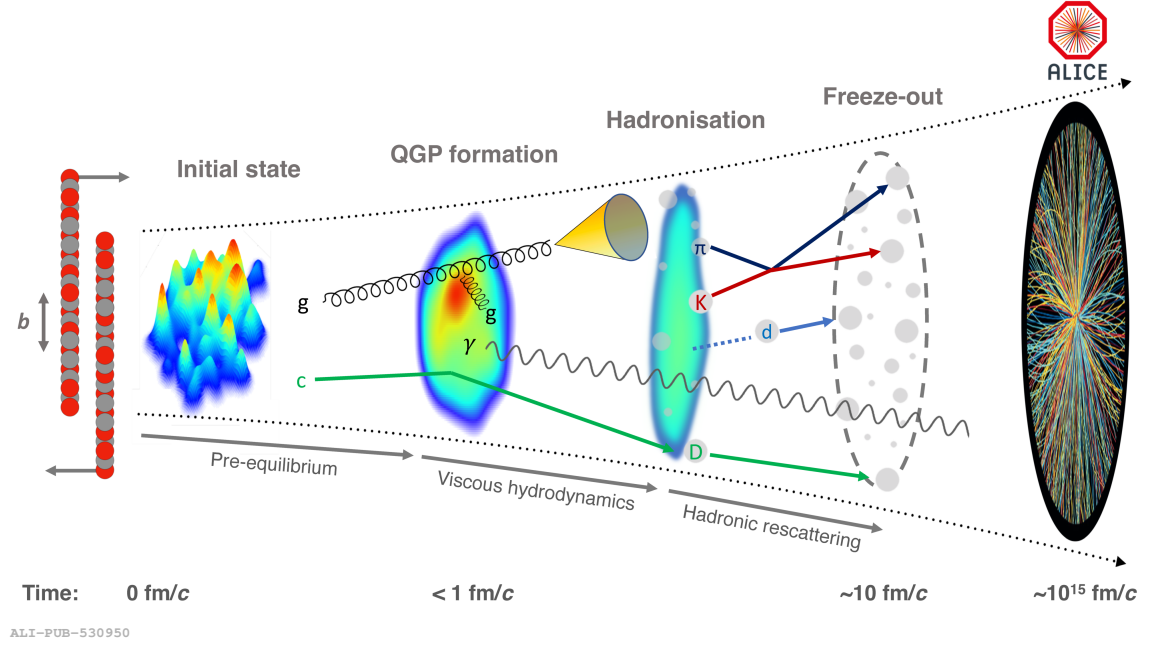


Figure 2.8: The evolution of a heavy-ion collision, originally appearing in [6].

Initial State

Each incoming nucleus begins as a Lorentz contracted disk moving at a speed close to the speed of light. These two discs then collide with a given collision geometry (see Section 2.3.2) where small- Q^2 interactions form a weakly-coupled pre-equilibrium parton gas, thus entering the stage where the energy density of the collision at its fluctuations are created. As the nuclei collide there can also be large- Q^2 scattering processes between partons, which will be discussed at length in Section 2.4.

The initial state can be modeled in a variety of ways. One way to model the initial state is as a superposition of independent nucleons, using a [Monte Carlo \(MC\)](#) Glauber Model [44] where the density of nucleons is typically sampled from a Woods-Saxon distribution.¹¹ Despite ignoring any binding of the individual nucleons, the Glauber model convoluted with a particle production model is very successful at predicting distributions sensitive to the initial state (for example the charged-particle multiplicity, see Figure 2.12). For this reason it is used in determinations of the event centrality in experiment, as described in Section 2.3.2. The second way of dynamically modeling the initial state is by

¹¹The Woods-Saxon distribution is the potential distribution of nucleons in the atomic nucleus in the shell model of nucleus given as $V(r) = -\frac{V_0}{1+\exp\{\frac{r-R}{a}\}}$, where $V_0 \approx 50$ MeV is the depth of the potential well, $a \approx 0.5$ fm is the thickness of the surface of the nucleus and $R = r_0 A^{1/3}$ is the nuclear radius where $r_0 = 1.25$ fm and A is the atomic mass.

using CGC effective theory in an Impact Parameter Glasma (abbreviated IP Glasma) approach to describe the initial state as a coherent wall of gluons [7]. CGC models include fluctuations both in the distribution of color charges in the nucleus as well as quantum fluctuations of the nucleon distribution in the nuclear wave functions. Sample distributions for the initial energy densities for the MC Glauber model as well as the IP Glasma model are shown in Figure 2.9. As seen in Figure 2.9, the IP Glasma model has finer structure due to the smaller length scales of the fluctuations [7].




Figure 2.9: Sample initial energy density profiles for the MC Glauber model (left) and the IP Glasma model (right), originally appearing in [7].

QGP Phase

Due to the high temperatures and energy densities in the initial state, the system thermalizes and then forms the QGP. The precise mechanism for this thermalization is unknown. During this phase, the temperature of the medium is above the critical temperature of $T_c \approx 150$ MeV [5] predicted by lattice QCD calculations. During its evolution, the medium will expand and cool until hadronization occurs.

The QGP acts as a strongly-coupled fluid and its evolution can be described by relativistic viscous hydrodynamics with hydrodynamical models such as MUSIC [45], iEBE-VISHNU [46], and CLVisc [47] or transport models such as AMPT [48]. In this picture the expansion is highly driven by the non-uniform spatial energy distribution in the initial state described above, resulting in pressure gradients in the QGP phase which then results in *anisotropic flow*. This is influenced by the shear viscosity of the QGP, or the resistance to fluid deformation. Greater pressure at the center of the QGP medium than the outskirts results in *radial flow*. This is influenced by the bulk viscosity of the QGP, or the resistance to expansion, which controls the rate of expansion. Both the shear and the bulk viscosity as well as anisotropic and radial flow are shown in Figure 2.10.

Hadronization and Freezeout

Hadronization is a non-perturbative process that happens when the energy density of the system becomes sufficiently small and the temperature is sufficiently cool such that partons




Figure 2.10: (Top): Shear and bulk viscosity. (Bottom): Radial and anisotropic flow.

combine into hadrons. Modeling of the hadronization process in the [QGP](#) medium differs from that for pp collisions as the medium particles must also hadronize, where the modeling of hadronization in pp collisions is described in Section 2.4.4. Hadronization in heavy-ion collisions can be modeled using statistical hadronization models [49, 50] where partons form hadrons on a statistical basis or by so-called quark recombination/coalescence models where quarks and anti-quarks can recombine into hadrons. The temperature continues to cool over the process of hadronization. The temperature at which the hadron species becomes fixed is called the chemical freezeout temperature. The kinetic freezeout temperature refers to the temperature at which the exchange of energy and momentum stops.

2.3.2 Geometry of a heavy-ion collision

When two nuclei collide, only some of the nucleons participate in the collision. The number of participating nucleons, referred to as N_{part} , is defined as the number of nucleons that undergo one (or more) binary collision(s) with nucleons of the other nucleus. The total number of binary collisions is referred to as N_{coll} . The number of nucleons that do not participate in the collision (or “spectate”) is referred to as N_{spec} . Therefore, by definition, $N_{\text{part}} = 2A - N_{\text{spec}}$, where A is the atomic mass of the colliding nucleus.¹² The number of participating nucleons is driven in part by the impact parameter, b , of the collision, which is defined as the distance between the centers of the two colliding nuclei in the transverse plane. For example, when there is larger overlap between the two nuclei, N_{part} will be

¹²Where $A = 208$ in the case of Pb–Pb collisions.

larger. The geometry of a heavy-ion event both before and after the collision can be found in Figure 2.11.




Figure 2.11: (Left): Geometry both before and after a heavy-ion collision. The impact parameter, b , as well as spectators (gray) and participants (in color) are shown. (Right): Side view of colliding nuclei for central and peripheral collisions with impact parameter b .

As effects due to the QGP medium may vary as a function of the impact parameter, it is common to perform analyses on collections of events with a similar collision geometry, determined by the so-called “centrality”. Due to the experimentally inaccessible nature of b , experimental analogs are used instead. One such analog is the average charged-particle multiplicity ($\langle N_{\text{ch}} \rangle$) that increases monotonically with b . In ALICE, the centrality is determined by the measured signal in the V0 detector¹³ that is correlated with the charged-particle multiplicity. This distribution, shown for ALICE data in Figure 2.12 is then fitted with the Glauber model [44] coupled to a particle production model in order to determine the corresponding N_{spec} and N_{part} . Smaller percentile values of centrality (for example 0–10% or central collisions) refer to collisions with more overlap and therefore smaller values of b and larger percentile values of centrality (for example 60–80% or peripheral collisions) refer to collisions with less overlap and therefore larger values of b as shown in the right panel of Figure 2.11. Another experimental observable that can be related to the number of spectators is the energy carried by particles close to the beam direction and deposited in so-called Zero-Degree Calorimeters (ZDC, see Chapter 3 for more details).

As the phenomena studied in heavy-ion collisions have been observed in increasingly smaller systems (at ALICE, pp or p-Pb collisions) [51], this approach to determining centrality has been applied to studies of these smaller systems as well. While the term “centrality” may still be used, this is perhaps better understood as event activity, since the correlation between multiplicity and impact parameter is weaker in these systems.

¹³See Chapter 3 for more details.




Figure 2.12: V0M signal amplitude with a Glauber+NBD fit [8], originally appearing in [6].

Coordinate System of a Heavy-Ion Collision

The coordinate system used to characterize the high-energy collisions is such that the x-axis points radially inward toward the center of the collider, the y-axis points upward and the z-axis is along the beam line. The momentum carried by each particle is then written in terms of the four-vector $p^\mu = (E, p_x, p_y, p_z)$ where in this thesis the transverse momentum $p_T = \sqrt{p_x^2 + p_y^2}$ is most commonly used. For the position of the particles in the cylindrical detector, the common choice of variables is the azimuthal angle ϕ and the pseudorapidity η . Here η is defined in Equation 2.6 where θ is the polar angle relative to the beam (z) axis. In the massless limit, the pseudorapidity is equivalent to the rapidity y (not to be confused with the y-axis), which is also defined in Equation 2.6.

$$\eta = -\ln\left(\tan \frac{\theta}{2}\right), y = \frac{1}{2} \ln\left(\frac{E + p_z}{E - p_z}\right) \quad (2.6)$$

The center-of-mass energy for a pp collision is typically written as \sqrt{s} , where $s = (p_1 + p_2)^2$ is one of the Mandelstam variables where p_1 and p_2 are the four momenta of the incoming protons. In heavy-ion collisions this is written as the center-of-mass-energy per nucleon-nucleon collision, $\sqrt{s_{NN}}$.

2.3.3 Experimental Facilities

The [Relativistic Heavy-Ion Collider \(RHIC\)](#) at [Brookhaven National Laboratory \(BNL\)](#) and the [Large Hadron Collider \(LHC\)](#) at [CERN](#) are two existing facilities that collide heavy ions. In this section, each of these facilities will be briefly described, with a focus on the [LHC](#).

Relativistic Heavy-Ion Collider

The **Relativistic Heavy-Ion Collider (RHIC)** based at **BNL** is a dedicated heavy-ion collider that has been in operation since 2000. **RHIC** consists of two concentric rings, each 3834 m in circumference. As of the time of the writing of this thesis there are two main experiments at **RHIC**; **STAR** and **sPHENIX**.¹⁴ **RHIC** produces **QCD** matter at high T and a variety of μ_B s, depending on the collision energy. **RHIC** has collided several different species at many different energies, something particularly useful for an experimental search for the **QCD** critical point. The results from **RHIC** most relevant to this thesis are collisions of gold (Au) ions at $\sqrt{s_{NN}} = 200$ GeV.

In 2020, **BNL** was chosen as the site of the planned **EIC** [52], meaning that the **RHIC** facility will be transformed into **electron ring at RHIC (eRHIC)** upon the completion of its scheduled runtime.

Large Hadron Collider

The **Large Hadron Collider (LHC)** based at **CERN** is the largest particle accelerator in the world.¹⁵ It contains two concentric rings with a 27 km circumference that are located near the France-Switzerland border in Geneva, Switzerland. In order to avoid the Jura mountain range in this region, the **LHC** sits at a slant with a depth ranging from ~ 50 to 175 meters. For most of the year, the **LHC** records data from the collisions of protons, which are collided with energies up to $\sqrt{s} = 13.6$ TeV. Around 1 month per year the **LHC** also records heavy-ion collision data using Xe and Pb ions accelerated up to $\sqrt{s_{NN}} = 5.36$ TeV.¹⁶ The **LHC** produces **QCD** matter at high T and near-zero μ_B .

There are four main experiments located at different interaction points (**IPs**) at the **LHC**. The “**A Toroidal LHC Apparatus**” (**ATLAS**) experiment is located at interaction point 1 (**IP1**) in Meyrin, Switzerland and is a general-purpose particle detector. The “**A Large Ion Collider Experiment**” (**ALICE**) sits at **IP2** in St. Genis, France and is optimized for the study of heavy-ion collisions. The “**Compact Muon Solenoid**” (**CMS**) experiment sits at **IP5** in Cessy, France and is also a general-purpose experiment. The “**LHC Beauty**” (**LHCb**) experiment sits at **IP8** in Ferney-Voltaire, France and is optimized for studies of the bottom quark. Injection of the clockwise beam, so-called Beam 1 occurs at **IP2** and the injection of the anti-clockwise beam, so-called Beam 2 occurs at **IP8**. The beam dumping system for both beams is located at **IP6**. The main beam collimator installations are located at **IP3** and **IP7**. The radio frequency system is located at **IP4**.

The **LHC** accelerator complex is shown in Figure 2.13. In order to collide at near the speed of light, the beams must first be accelerated. For the collisions of Pb ions ($Z =$

¹⁴**sPHENIX** is scheduled to take data for the first time in the spring of 2023.

¹⁵It is also the largest machine ever built.

¹⁶Historically this occurs around November. At the time of writing this thesis, the next heavy-ion run will occur in October 2023.

82, $A = 208$), the source enters the injection chain at **Linear Accelerator (LINAC) 3** [53] and is then accelerated to 4.2 MeV/u. Several **LINAC 3** pulses are then accumulated into bunches that are accelerated to 72 MeV/u in the **Low-Energy Ion Ring (LEIR)** [54]. Then the beams are accelerated both in the **Proton Synchrotron (PS)** [55] and the **Super Proton Synchrotron (SPS)** [56] where they reach energies of 5.9 GeV/u and 177 GeV/u, respectively, before reaching the **LHC**. In the **LHC**, each beam is accelerated to its target energy of 2.51 TeV/u, resulting in a collision energy per nucleon of $\sqrt{s_{NN}}$ of 5.02 TeV.¹⁷ For protons, this process begins at **LINAC 4** [57] where H^- is accelerated to 160 MeV. Then the beam enters the **Proton Synchrotron Booster (PSB)** [58] where it is stripped of its electron, leaving only a beam of protons. Then the beam is accelerated both in the **PS** and the **SPS** where it reaches energies of 26 GeV and 450 GeV, respectively, before reaching the **LHC** and being accelerated to its target energy.

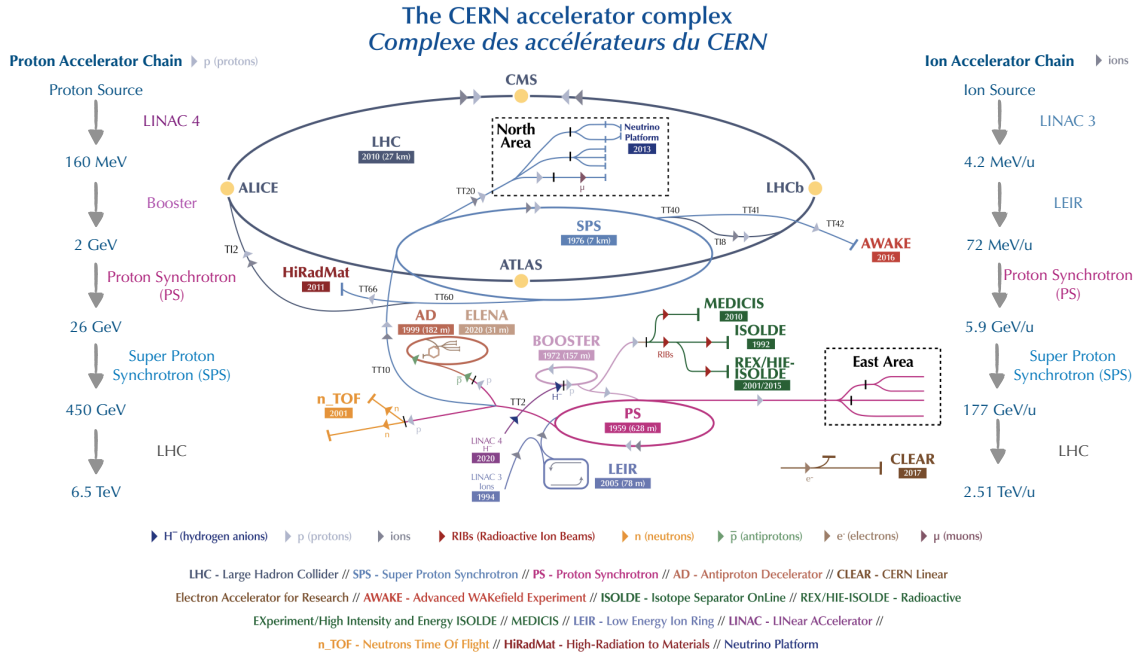


Figure 2.13: The accelerator complex of the Large Hadron Collider (not to scale), adapted from [9] last updated in 2022. The path taken by various species as they enter the different complexes is outlined on the sides of the image, where in the case of heavy ions units of energy are given per atomic mass unit (u) for each beam.

Collisions at the **LHC** occur based on a previously-specified filling scheme that specifies the number of bunches to be collided at each experiment, the spacing between them, and other relevant information. For protons, one bunch contains around 100 billion protons that are typically spaced 25 ns apart. For the case of heavy ions, one bunch contains

¹⁷Therefore the total energy in the collision is $\sqrt{s} = 1.04$ PeV.

around 10 million heavy ions with a typical spacing of 150 ns between them.¹⁸ As a result of this setup, multiple collisions can occur in the crossing of two bunches (aptly referred to as a bunch crossing). This leads to an effect called pileup where one single event recorded in the detector can contain information from multiple collisions. Pileup can either occur for two collisions in the same bunch crossing (called in-bunch pileup) or two collisions in different bunch crossings (called out-of-bunch pileup).

2.3.4 Experimental Probes of the QGP

There are two main aspects of the QGP of interest. The first is the intrinsic properties of the QGP, such as the equation of state, the shear viscosity, and the bulk viscosity. The second is the dynamic processes of the QGP or, in other words, how various processes such as the evolution of partons or hadronization occur in the QGP. The experimental probes used to study both the properties and dynamics of the QGP can be divided into three categories, which are discussed below.

Soft Probes

The first category is composed of soft probes,¹⁹ concerning the hadronization products of the QGP medium. Soft probes generally relate to the collective properties of the QGP. One exception includes measurements of the enhancement of (multi-)strange hadrons [59], which are expected to be created throughout the evolution of the QGP.

The primary way to measure this is via measurements of radial flow and anisotropic flow, which are sensitive to the bulk and shear viscosity of the QGP, respectively. Measuring the mean p_T of identified particles is a clean probe of radial flow, where higher mass particles are expected to exhibit a higher boost in the mean p_T sensitive to the degree of radial flow. Measurements of the anisotropic flow are separated into the different harmonics in order to characterize the shape of the anisotropies as shown in Figure 2.14, which relates to different spatial anisotropies in the initial state. The Fourier decomposition is written out in Equation 2.7 [60]. In general, the CGC model of the initial state better predicts the momentum anisotropies seen in data than the MC Glauber model of the initial state.

$$\frac{dN}{d\Delta\phi} \propto (1 + 2 \sum_{n=1}^{\infty} v_n \cos(n(\phi - \Psi_n))) \quad (2.7)$$

In Equation 2.7, $\frac{dN}{d\Delta\phi}$ represents the azimuthal distribution, Ψ_n is the symmetry plane angle, ϕ represents the angular component of the transverse momentum vector, and n represents the order of the harmonic.²⁰ The momentum anisotropies, measured by v_n , arise from

¹⁸The exact values depend on the particular filling scheme used.

¹⁹The term soft here indicates lower momentum.

²⁰When $n = 2$, this is referred to as *elliptic flow*.

spatial anisotropies in the initial state, given by the eccentricities ε_n , related by the coefficient κ_n as $v_n \approx \kappa_n \varepsilon_n$ [6]. As in the case of radial flow, the mass dependence of the anisotropic flow is sensitive to the flow velocity. Correlations between the different n^{th} order harmonics are also of interest and provide important input for constraining model calculations.




Figure 2.14: Visual representation of the Fourier decomposition of the flow harmonics. Here Φ_n reflects the n^{th} order harmonic. Originally appearing in [10].

There are a number of other soft probes of the collective behavior of the QGP. One common example is via measurements of two-particle correlations, where short-range correlations are expected to be caused by the products of hard-scattered partons and long-range correlations are expected to be caused by collective effects. Event-by-event fluctuations can also probe collective effects via the n th order cumulant distributions for identified particles, such as protons. These distributions have also been used in experimental searches for the QCD critical point. Measurements of the global polarization can also be used to probe the vorticity of the QGP medium, generated by offsets in the colliding nuclei. However, a generated magnetic field in the QGP, possibly the largest magnetic field in the universe, may cause similar polarization effects [61].

Non-Interacting Probes

The second category of observables is *non-interacting* (also called electroweak) observables that concern particles that weakly interact with the QGP medium since their mean-free-path is larger than the size of the QGP. Though there are other types of non-interacting probes, this section will focus on the measurements of real photons and virtual photons measured via lepton pairs (called dileptons), which can be measured by the ALICE experiment.

There are many different types of both real and virtual photons that can be measured in heavy-ion collisions, each of which probes different effects. Photons produced via the decay of other particles are called decay photons, and all other photons are referred to as direct photons. At low transverse momenta, the direct photon spectrum is dominated by photons radiated during QGP evolution, or so-called “thermal photons”. Experimental measurements of thermal photons provide powerful constraints on the temperature and

the space-time evolution of the QGP as well as useful input for the QCD equation of state.²¹ Dileptons originating from various sources can be separated via the mass of the dilepton pair, where at low invariant mass the dominant contribution is from thermal radiation from the hot hadron gas. At larger masses there are contributions both from semi-leptonic decays of open heavy-flavor hadrons and thermal radiation from the QGP. Current ALICE measurements of dileptons are not yet sensitive to signals of thermally produced dileptons, but such measurements may be possible in Run 3 and Run 4 of the LHC [6].

Hard Probes

The third category, which will be discussed at length in this thesis, concerns *hard*²² observables that are the result of high momentum transfer scatterings of partons (as shown in the absence of the QGP in Figure 2.15) whose production is calculable in pQCD. As these partons are formed early in the collisions before QGP formation they experience its full evolution and serve as a colored probe of the colored medium.²³ The partons are expected to lose energy via interactions with the QGP and this energy loss can be useful in the determination of the various intrinsic and dynamic QGP properties. The most relevant QGP property related to the study of hard probes is the transport coefficient, $\hat{q} = d\langle k_{\perp}^2 \rangle / dL$, which quantifies the momentum transfer between a hard parton and the soft QGP medium per unit length.

Different partons may interact differently with the QGP and therefore it is useful to isolate and study the interactions of different types of partons, for example heavy quarks, with the QGP. Heavy quarks are expected to interact differently with the QGP than light quarks and therefore measurements of heavy quarks can be used to constrain intrinsic properties of the QGP, such as the spatial diffusion coefficient $D_s = 4T^2/\hat{q}$, and various dynamical properties of the QGP. Heavy-flavor measurements can be broken up into two categories; open heavy-flavor measurements and hidden heavy-flavor measurements. Open heavy-flavor measurements refer to the direct measurement of a heavy-flavor meson (mesons containing c and b quarks, such as D and B mesons, respectively), serving as a proxy for a hard-scattered heavy quark (c and b).²⁴ Hidden heavy-flavor, also called as quarkonia, refer to bound $q\bar{q}$ states (so $c\bar{c}$ and $b\bar{b}$ states). Common examples include the charmonium states of J/ψ and $\psi(2s)$, and the bottomonium states of $\Upsilon(1s)$, $\Upsilon(2s)$, and $\Upsilon(3s)$. The production of quarkonia is expected to melt in a sequential fashion²⁵ as the

²¹<https://cerncourier.com/a/measurement-of-photons-stimulates-quest-for-qgp-temperature/>

²²The term hard here indicates higher momentum.

²³Whereas non-interacting probes as mentioned above can be thought of as non-colored probes of the colored medium.

²⁴Measurements of jets initiated by a heavy quark [62] can also be used for this purpose, though this comes with additional experimental complications.

²⁵The dissociation or “melting” of a quarkonium state is expected to occur when the Debye screening radius is similar to the binding radius of the state [63].

temperature increases, where states with a smaller binding energy would be more suppressed. This effect can be experimentally measured as a sequential suppression of the quarkonium yield in heavy-ion collisions as compared to pp collisions and is a key probe of the temperature evolution of the QGP [64, 65]. However, the quarkonium yield could also increase in the QGP due to the recombination of un-correlated quarks, which can complicate measurements.

An additional hard probe of the QGP is the measurement of hard charged hadrons. These roughly correspond to the hadronization products originating from the hard-scattered parton, and therefore are also sensitive to parton energy loss effects in the QGP. Another proxy for the dynamics of the hard-scattered parton is via jets, which will be discussed at length in the next section and the remainder of this thesis.

2.4 Jets in Vacuum

In this section jets in vacuum, i.e. in the absence of the QGP medium, will be discussed. Jets in pp collisions are assumed to be vacuum-like in this context. ²⁶

2.4.1 Theoretical Definition of a Jet

In high-energy collisions, partons inside the colliding nucleons can scatter off one another with a high momentum transfer (Q^2) in a process referred to as a hard scattering. A majority of hard scatterings in the vacuum are $2 \rightarrow 2$, resulting in high- p_T partons traveling 180 degrees apart in the transverse plane with approximately equal p_T . The computation of the production cross-section of jets in QCD is discussed in Section 2.2.1. The outgoing partons then fragment via a parton shower and hadronize into a spray of particles called a jet as shown in Figure 2.15. In this section, the physics of the parton shower and hadronization will be briefly described.

Parton shower

Analogous to the emission of photons from an accelerated electric charge, accelerated color charges can also emit gluons. However, these gluons themselves carry a color charge (unlike photons) and can themselves emit further gluons. In addition, they can also produce quark-antiquark pairs. The combination of these two processes is referred to as the parton shower. This manifests itself as higher-order corrections to the original hard process described in Section 2.2.1. Due to the difficulty of the explicit calculation, the dominant contributions at each order, mostly soft gluon emission and collinear splittings, are instead approximated.

²⁶Some recent measurements [51] indicate that small QGP-like droplets could be formed in smaller collisions systems (such as high-multiplicity pp collisions), but such results are not the focus of this discussion.




Figure 2.15: The hard scattering and subsequent fragmentation and hadronization of partons, appearing in [11].

For collinear splittings the probability for a split, where a parton splits into two partons, $P_{i \rightarrow jk}$, as shown in Figure 2.16, is dictated by the DGLAP [35, 36] (also sometimes called the Altarelli-Parisi (AP) [37]) splitting functions as written in Equation 2.8. The splitting function is written as a function of z , the fraction of momentum of the original parton carried by the split parton (where in addition, $C_F = 4/3$, $C_A = 3$ and $T_R = 1/2$).




Figure 2.16: Parton splitting process for $i \rightarrow jk$ with the momentum fraction z , azimuthal angle ϕ , and opening angle θ .

$$P_{g \rightarrow qq}(z) = T_R(z^2 + (1-z)^2), \quad P_{q \rightarrow qq}(z) = C_F \left(\frac{1 + (1-z)^2}{z} \right)$$

$$P_{g \rightarrow gg}(z) = 2C_A \left(\frac{z}{(1-z)_+} + \frac{1-z}{z} + z(1-z) \right) + \delta(1-z) \frac{11C_A - 4n_f T_R}{6} \quad (2.8)$$

$$P_{q \rightarrow gq} = C_F \left(\frac{1+z^2}{1-z} \right)_+$$

Equation 2.8 additionally utilizes the plus (+) prescription as written in Equation 2.9.

$$\int_0^1 dz [g(z)]_+ f(z) = \int_0^1 dz g(z) f(z) - \int_0^1 dz g(z) f(1) \quad (2.9)$$

Here, effects from virtual quantum loops are included in the probability of *not splitting* during an evolution between two scales Q_1^2 to Q_2^2 . This is given by the Sudakov form factor as written in Equation 2.10.

$$\Delta_i(Q_1^2, Q_2^2) = \exp\left\{-\int_{Q_2^2}^{Q_1^2} \frac{dQ^2}{Q^2} \frac{\alpha_s}{2\pi} \int_{Q_0^2}^{1-Q_0^2} dz \int_0^{2\pi} d\phi P_{ij}(z, \phi)\right\} \quad (2.10)$$

In this case, the z limits of integration specify the limit where the splitting is hard enough such that it can be resolved and therefore included in the shower. The generation of a parton shower from an initial scale, Q^2 , is performed by solving $\Delta_1(Q^2, Q_1^2) = \mathcal{R}$, where \mathcal{R} is a random number from 0 to 1. This process is then repeated until the splitting becomes un-resolvable.

Hadronization

As mentioned in Section 2.2.1, the transition from partons to hadrons is **non-perturbative (NP)** in nature and therefore difficult to calculate. To evolve parton-level predictions to the hadron-level and compare to experimental measurements, correction factors that estimate the size of **NP** effects are typically applied. These correction factors are determined by taking the ratio of predictions from event generators including and not-including **NP** effects such as hadronization and multi-parton interactions. For more information on how each of these **NP** effects is modeled in these event generators, see Section 2.4.4.

As evidenced by the above description, jets are theoretically complex objects representative of the original parton kinematics and sensitive to various physical scales from the hard-scattering scale set by the amplitude (or matrix elements) for the process to the hadronization scale. Jets are useful for three main avenues of study, which are listed below.

1. **Studying processes of a specific origin:** Jets are useful objects to tag processes of a specific type or origin. Some examples include decays of the Higgs Boson, beyond the Standard Model processes, dynamics of quarks or gluons, and decays of heavy quarks.
2. **Tests of fundamental QCD:** The production of the high p_T parton that forms the jet is calculable in **pQCD**. As jets in vacuum serve as a proxy for these partons, jets are ideal probes of **pQCD**. Additionally, jets can be used to probe **npQCD** via measurements that isolate the **NP** process of hadronization.
3. **Probing the QGP in heavy-ion collisions:** Jets in vacuum additionally serve as a reference for jets in heavy-ion collisions, where any difference between the two systems serves as an in-medium effect. This will be discussed in Section 2.5.

2.4.2 Jet Finding Algorithms

The definition of a jet is the result of a jet-finding (also called a jet clustering) algorithm. A jet-finding algorithm attempts to recluster all fragmentation particles of a given parton - providing experimental access to the full energy and momentum of the original parton.

Jet-finding algorithms should be fast, sensitive to the underlying physics and have a property called infrared and collinear safety (abbreviated **IRC** safety) [66]. For a jet clustering algorithm to be “collinear safe” the algorithm must be unaffected by a collinear splitting of a particle where $(\Delta\phi)^2 + (\Delta\eta)^2 = 0$. Infrared safety means that the jet clustering algorithm is insensitive to the emission of a soft particle (for example the emission of a low-energy gluon), which can occur as a result of both perturbative and non-perturbative effects. Theoretical calculations of jets rely on collinear splittings to be included in the same jet such that infrared divergences from the real emission cancel with divergences resulting from virtual corrections, as dictated by the KLN theorem [67]. Experimental jets should not be sensitive to such emissions.

The two main classes of jet-finding algorithms are jet cone algorithms and sequential recombination algorithms. Sequential recombination algorithms, which are more commonly used, will be discussed in this section. For more information on jet cone algorithms and an example of an **IRC** safe jet cone algorithm, see [68]. Sequential recombination algorithms are a class of algorithms that first identify a pair of particles that are closest in some distance metric, recombining them to form pseudo-jets, and then iteratively repeating this procedure until a stopping condition is reached. The distance between particles is generally calculated as

$$d_{ij} = \min(k_{T,i}^{2p}, k_{T,j}^{2p}) \frac{\Delta_{ij}^2}{R^2} \quad (2.11)$$

where $\Delta_{ij}^2 = (\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2$, p is the parameter in this process that specifies the type of sequential recombination algorithm, k_T is the transverse momentum of the particle, and R is the resolution parameter that roughly corresponds to the radius of the jet cone. For comparison, the distance between each particle and the beam is calculated as

$$d_{iB} = k_{T,i}^{2p}. \quad (2.12)$$

When considering a given pair of particles, the distances specified by Equations 2.11 and 2.12 will be compared. When Equation 2.11 is the smallest of the two, these particles will then be grouped into a pseudo-jet. When Equation 2.12 is smaller, the i th pseudo-jet is removed and named a jet. This procedure is then repeated until all particles in the event are clustered into jets, or until an alternate stopping condition specified by the user is reached. In order to combine particles into pseudo-jets, there are two different recombination schemes that can be used. The first is the E -scheme recombination in which the four-vectors of individual particles are combined in order to form a (pseudo)jet

four-vector. The second is the p_T -scheme recombination that imposes an additional scaling that makes the jet's energy equivalent to its three-momentum.

The sequential recombination algorithm that is most commonly used for the reconstruction of signal jets at the LHC is the anti- k_T algorithm [12], which corresponds to $p = -1$ in the above procedure. As a consequence of the choice of $p = -1$, the anti- k_T algorithm clusters high p_T particles first, a feature that creates “soft-resilient” jets that are relatively insensitive to the presence of an underlying event (UE) or pileup. For this reason the anti- k_T algorithm is favorable in noisy environments such as the LHC. The k_T algorithm [69] corresponds to $p = 1$ in the above procedure. The k_T algorithm favors the clustering of lower p_T particles. As a result of this feature, the k_T algorithm tends to create “soft adaptable” jets that are more influenced by the presence of an UE or pileup. The k_T algorithm is useful for studies of the UE due to its tendency to naturally organize a uniform soft background into clusters of variable area. The third and final type of sequential recombination algorithm is the Cambridge/Aachen (C/A) algorithm [70], which corresponds to $p = 0$ in the above procedure. As a consequence, the C/A algorithm does not take p_T into account and is purely geometric. Though this tends to produce jets with irregular shapes, the C/A algorithm is useful to enforce angular ordering that mimics the QCD branching sequence where each emission is expected to be at a smaller angle than the previous emission. A comparison of the jets formed by the k_T , C/A, and anti- k_T algorithms is shown in Figure 2.17.




Figure 2.17: Comparison of spatial distribution (ϕ , y) of jets found in a parton-level event using the k_T (left), C/A (middle), and anti- k_T (right) sequential recombination algorithms, adapted from [12].

There are two main types of jets that can be reconstructed. The first is charged-particle jets,²⁷ which are composed of only charged constituents. Charged-particle jets are commonly measured, particularly for measurements of jet substructure, in ALICE in order to take advantage of its precise tracking. Though charged jets have the clear disadvantage that they are missing neutral information, the probability distributions for charged particles compared to all particles are found to be very similar [71]. In addition, analytical calculations are typically performed only when both charged and neutral information is

²⁷Herein referred to as charged-particle jets or charged jets.

included, but track functions [72, 73] offer a promising avenue for the analytical computation of track-based observables. Jets that contain both charged and neutral constituents, commonly called full jets, can also be formed. Measurements of full jets have the added benefit that they are closer to the definition of a jet and its originating parton. However, their precision, particularly for substructure observables, is limited by the detector resolution.

The above procedure remains identical in both pp and heavy-ion collisions. However, there are a few additional considerations for the jet population in heavy-ion collisions due to the presence of the QGP medium. In heavy-ion collisions the radiated partons (as well as the original parton) can interact with the medium, making even the theoretical definition of a jet ambiguous. Additionally, in heavy-ion collisions there is a large fluctuating background of particles that do not originate from the hard scattering thus complicating jet measurements, which will be discussed in Section 5.4. These fluctuations can be uncorrelated, emerging from Poissonian fluctuations, or correlated, emerging from flow-like effects. In order to suppress these background contributions, additional restrictions are often imposed on the experimental jet definition in heavy-ion collisions.²⁸ These restrictions bias the jet sample towards harder more collimated jets, potentially selecting quark-initiated jets²⁹ at a larger rate than gluon-initiated jets. These considerations add additional challenges for measuring jets in heavy-ion collisions and should be kept in mind in experimental measurements and the corresponding theoretical comparisons.

2.4.3 Jet as a Proxy for the Parton

As mentioned in Section 2.4.1, a jet serves as a direct proxy for a hard-scattered parton. However, jets are not a perfect representation of the dynamics of the initial parton. This is due to the fact that different effects can move energy into or outside of the jet cone, causing the energy of the jet to differ from the energy carried by the hard-scattered parton. Below, each of these differences and a general formulation of the corresponding p_T loss (or gain) as a function of the resolution parameter of the jet will be discussed.

One such effect that can move energy from the hard-scattered parton out of the jet cone is the parton shower. This process is dictated by the QCD splitting function and is calculable in pQCD. Using the LO result with a small R approximation, the amount of lost p_T for quark jets is given in Equation 2.13 and that for gluon jets is given in Equation 2.14 [74–76]. In these equations, $C_F = \frac{4}{3}$ and $C_A = 3$ are the Casimir color factors for quarks and gluons, respectively. Additionally, $T_R = \frac{1}{2}$ is a normalization convention, and n_f is the number of active³⁰ quark flavors. The dependence on R of the p_T lost due to the

²⁸One common example is to require the jet to have a constituent above a certain p_T .

²⁹Quark(gluon)-initiated jets will herein be referred to as quark(gluon) jets.

³⁰Active in this context refers to the number of flavors with a mass below the energy scale of the process.

parton shower can be simplified to $\delta p_T \propto \ln(\frac{1}{R})$.

$$(\delta p_T)_q = -C_F \frac{\alpha_s p_T}{\pi} \ln\left(\frac{1}{R}\right) (2 \ln(2) - \frac{3}{8}) + \dots \quad (2.13)$$

$$(\delta p_T)_g = -\frac{\alpha_s p_T}{\pi} \ln\left(\frac{1}{R}\right) [C_A (2 \ln(2) - \frac{43}{96}) + T_R n_f \frac{7}{48}] + \dots \quad (2.14)$$

Hadronization, the process describing the transition of partons into hadrons, can also cause energy to move outside of the jet cone. Some theoretical models [77] parameterize this effect by taking $\alpha_s(\mu) = \mu_1 \delta(\mu - \mu_1)$ where μ_1 corresponds to the Landau pole, which gives the approximate p_T loss as written in Equation 2.15. Here $C = C_A$ for gluon jets and $C = C_F$ for quark jets and $\mathcal{A}(\mu_1)$ is a scale relevant for processes of hadronization. This dependence can be simplified to $\delta p_T \propto \frac{1}{R}$.

$$\delta p_T \approx -\frac{2C\mathcal{A}(\mu_1)}{\pi R} \quad (2.15)$$

The **Underlying Event (UE)**, which refers to any process that is not the result of the hard scattering, can cause energy not associated with the original parton to move inside the jet cone.³¹ The **UE** is assumed to be distributed roughly uniformly and therefore the dependence can be written as in Equation 2.16 [76, 18] where Λ_{UE} is the energy density per unity rapidity. Here, this dependence can be simplified to $\delta p_T \propto R^2$.

$$(\delta p_T)_{\text{UE}} \approx \frac{1}{2} \Lambda_{\text{UE}} R^2 \quad (2.16)$$

The overall difference between the p_T of the experimental jet and that of the parent parton is then a convolution of these three effects. Jets with $R = 0.4$ are thought to be most closely associated with the dynamics of the parent parton at the **LHC** as they optimize p_T lost due to the parton shower and hadronization moving energy outside of the jet cone with effects of the **UE** pushing energy into the jet cone [18]. As a result, jets with $R = 0.4$ are the most commonly measured in experiment.

Measuring the dependence of the jet p_T spectrum on R in vacuum is a good way to isolate and test both the perturbative regime of the parton shower and the non-perturbative regimes of hadronization and the **UE** [18]. Varying the p_T of the jet is in addition a good way of varying the perturbative and **NP** scales, where at lower p_T **NP** effects become dominant and at higher p_T jets become more perturbative. This is due to the fact that the p_T is a rough proxy for the Q^2 of the interaction, which dictates the running of the strong coupling constant $\alpha_s(Q)$ (see Figure 2.3).

³¹The **UE** will be described in more detail in Section 2.4.4.

2.4.4 Event Generators in pp

Monte-Carlo models (MCs) in HEP utilize random sampling techniques in order to perform simulations of high-energy collisions. Models that simulate jets in vacuum typically include simulation of four different effects; the parton shower, hadronization, underlying event, and hadronic decays. In this section, the ways that available MCs model each of these effects and the underlying physics mechanisms at play will be discussed.

Parton Shower

The process described above in Section 2.4.1 is for a Q^2 - or virtuality-ordered shower, where the hardest interactions occur first. However, there can be other types of showers that may be employed in MCs. One such shower is a coherent shower, which produces an angular-ordered shower with less soft gluon emission as compared to a virtuality-ordered shower. This manifests itself in a different number and energy distribution of hadrons after hadronization. Another approach is a dipole shower where gluon emission is generated by the dipole radiation pattern of the parton pair. In this picture, each quark/anti-quark is uniquely connected to a color partner and gluons are connected to multiple color partners. Each pair of color partners forms a dipole that splits into two upon the emission of a gluon.

Underlying event

The term UE in the context of jet measurements³² refers to anything that does not originate from the hard scattering of interest. One source of the UE is the fragmentation of non-colliding partons. Another source of UE is collisions between partons in the incoming hadrons that do not directly participate in the hard scattering. The extent of these multi-parton interactions (MPIs) and therefore the impact of the UE depends on the impact parameter structure of the hadron-hadron collision. The overlap of the partonic structures of the proton is directly related to the number of MPI, and therefore is important to model accurately. Both in- and out-of-bunch pileup³³ can also contribute to the UE and may also be taken into account in simulations.

Hadronization Models

All partons in the final state, whether originating from the hard scattering of interest or the UE, must hadronize. There are two common models for hadronization; the string hadronization model [78] and the cluster hadronization model [79]. The string model is based on the observation from Lattice QCD that at large distances the potential energy of color sources increases linearly with their separation, creating a string (or tube). For

³²In the context of other measurements, such as those of soft probes, what is referred to here as the UE would instead be the signal of interest.

³³See Section 2.3.3 for a definition of pileup.

example, when quarks and anti-quarks are separated a narrow flux tube (also called a string) forms between them with a constant strong force along the tube. As the distance between the two quarks becomes larger and larger, it becomes energetically favorable for an additional $q\bar{q}$ pair to arise from the vacuum as opposed to the flux tube getting larger. This process will then repeat itself until all available energy has been converted into $q\bar{q}$ pairs, which can be identified with mesons. Baryons are formed in a similar manner, but with a quark-diquark pair. Another illustration of the string hadronization model can be found in the left panel of Figure 2.18.

The cluster hadronization model is based on the property of QCD where at evolution scales much less than the hard process scale, the partons in a shower are clustered into colorless groups. These groupings will only depend on q and Λ_{QCD} , and will be independent of the hard subprocesses. These clusters can be identified at the hadronization scale as proto-hadrons that decay into the observed final-state hadrons. An illustration of the cluster hadronization model can be found in the right panel of Figure 2.18.




Figure 2.18: Illustrations of two different hadronization models. Left: the string hadronization model and Right: the cluster hadronization model. Figure originally appearing in [13].

Hadronic Decays

Following the hadronization step, the decay of any unstable hadrons must also be modeled. The decay products of excited hadronic states make up a large fraction of the observed final state, thus all excited states and all of the potential decay modes must be modeled with great accuracy. This process can be difficult as some properties of these hadronic decays are not yet established experimentally.

Example event generators

- **HERWIG (Fortran) [80] or HERWIG++ (C++) [81]:** Angular-ordered parton shower using the cluster model of hadronization. Built-in modeling of MPIs.

- **PYTHIA 6 (Fortran) [82] or PYTHIA 8 (C++) [83]:** The most widely used and well established event generator. Both versions implement a so-called “simple” shower with p_T -ordering where PYTHIA8 also includes other optional variations on this parton shower. Both versions use the Lund string model for hadronization and a highly developed multiple interaction model for the underlying event.
- **SHERPA [84]:** Utilizes a dipole shower and a cluster model for hadronization. Uses the PYTHIA formulation for multiple parton interaction modeling.

2.4.5 Jet Measurements in Vacuum

In this section, measurements of jets in vacuum (in this context, meaning pp collisions), will be discussed. Note that this is not meant to be an exhaustive list,³⁴ but instead meant to highlight a selection of vacuum jet measurements relevant to the work in this thesis. For clarity, these will be divided into two general categories; radial scans and the Lund plane.

Radial scans

As mentioned in Section 2.4.3, measuring the R -dependence of the jet p_T spectrum in vacuum can be a good way of isolating and exploring the perturbative and the non-perturbative scales. This type of measurement also provides the baseline for R -dependent measurements in heavy-ion collisions, which will be discussed in Chapter 5, so an understanding of these baselines is critical to interpreting these results. Two existing results [14, 18] of the radial scan of the p_T -dependent jet yields, shown as the ratio of jet cross sections for a nominal radius and other jet radii, can be found in Figure 2.19.

The ALICE results [14], shown in the left panel of Figure 2.19 indicate a large change in the jet cross-section ratio at low p_T that grows with R demonstrating the influence of the UE, which goes as R^2 as given by Equation 2.16. Additionally, the agreement with analytical calculations at NLO with a resummation of jet R and threshold logarithms to Next to Leading Log (NLL) accuracy with NP corrections, abbreviated NLO+NLL+NP [15–17] is worst at low p_T for small R , likely due to the stronger influence of hadronization in this regime. The CMS results [18], shown in the right panel of Figure 2.19 exhibit similar features, with overall flatter ratios than those of ALICE. This can be understood by the fact that jets at higher p_T are more collimated, causing less energy to move out of the cone as a result of hadronization and the parton shower. In addition, jets at higher p_T carry much greater energy than the background, causing the decreased relative influence of the UE. In this case the largest disagreement with NLO predictions [19] occurs at large R and low p_T , likely due to the stronger influence of the UE in this regime. Overall, both measurements

³⁴For a somewhat exhaustive list of jet measurements in vacuum at the LHC, see the following Twiki page: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/LHCJetSubstructureMeasurements>

indicate the jet p_T spectrum in vacuum has an R -dependence, which will be relevant when interpreting the results in Chapter 5.

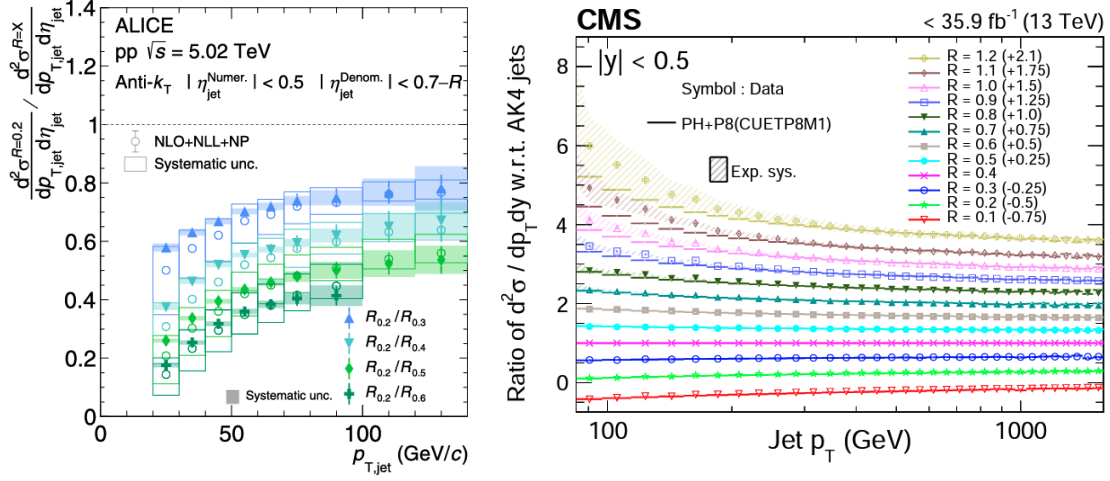


Figure 2.19: Left: **ALICE** pp jet cross-section ratios of $R = 0.2$ jets to other radii, originally appearing in [14]. These results are additionally compared to NLO+NLL+NP predictions [15–17]. Right: **CMS** pp cross sections of jets of different R as compared to $R = 0.4$ (here abbreviated AK4) jets, originally appearing in [18]. These results are additionally compared to NLO predictions [19].

Lund plane

As mentioned in Section 2.4.1, the vacuum can induce gluon radiation and quark pair production via a parton shower that can be described by the QCD splitting function. These partonic splittings, as they fragment and hadronize, give the jet a hard substructure, which manifests itself as subjets within the jet (see Figure 2.15). These splittings can be fully characterized by the opening angle and momentum fraction (or degree of symmetry) of the split. These two intrinsic axes can then be plotted in a way that covers the full phase space of jet splittings, which is introduced in Figure 2.20. This is known as the Lund plane (sometimes also referred to as the Lund jet plane) [85]. The primary Lund plane refers to the Lund plane filled with the emissions from the hardest prong. Additionally, at LO the density of emissions in the primary Lund plane is directly proportional to $\alpha_s(Q)$, meaning that the running of the strong coupling constant shapes the plane [86].

In the Lund jet plane (Figure 2.20) the angular axis is given as $\ln(1/\Delta R)$ where ΔR is the opening angle of the split and the axis that represents symmetry of the split is written in terms of the relative transverse momentum of the split $\ln(k_T)$ where $k_T = p_{T,2} \sin(\Delta R)$ and $p_{T,2}$ is the p_T of the sub-leading subjet in the splitting. This axis is sometimes also written in terms of the shared momentum fraction of the split, z which is defined as $z = \frac{p_{T,2}}{p_{T,1} + p_{T,2}}$. Given the axes definitions as shown in Figure 2.20, lines of constant z are

diagonal lines. The Lund plane can also be used to differentiate regions of wide-angle radiation from regions of collinear radiation, which can occur on both soft and hard scales. This can also be used to separate the perturbative (large $\ln(k_T)$) from the non-perturbative regions (small $\ln(k_T)$), where $k_T \sim \Lambda_{\text{QCD}}$ marks the perturbative to non-perturbative transition region.




Figure 2.20: The primary Lund jet plane.

The Lund jet plane is not only a theoretical representation of the spatial and temporal evolution of the jet, but it can also be measured in experiment via C/A reclustering (as defined in Section 2.4.2), which exploits the angular-ordered nature of QCD evolution [86]. In some cases, a grooming procedure is also applied to the jet in order to mitigate the impact of NP effects such as hadronization and the underlying event. This focuses the measurement on the hard or perturbative aspects of the jet, allowing for precise comparisons to pQCD. One common grooming method is Soft Drop (SD) grooming [87], where hard splittings are selected by the cut as shown in Equation 2.17. Here z_{cut} serves as a soft threshold and β serves as the angular exponent, where $\beta = 0$ is typically used.

$$z > z_{\text{cut}} \left(\frac{\Delta R}{R} \right)^\beta \quad (2.17)$$

Other grooming methods such as Dynamical Grooming [88] are also useful and employed in experimental vacuum jet measurements [89]. However, when there is interest in the NP region, or the transition between the perturbative and the NP regions, typically no grooming is applied.

There have been two measurements of the Lund plane in experiment, both shown in Figure 2.21. The left panel of Figure 2.21 shows the primary Lund plane as measured by ALICE [20] at low charged jet p_T s from 20-120 GeV/ c and the right panel shows the

Lund plane as measured by [ATLAS](#) [21] for higher full jet p_T 's with $p_T > 675$ GeV/ c . Note that the axes definitions differ in these two measurements, which results in a different qualitative shape of the distributions. Both measurements are complementary to each other, exploring different features of the perturbative to non-perturbative transition. Projections of the Lund plane are additionally very useful to isolate regions of phase space and compare them to models. Differential measurements identical to these projections, but for jets with a grooming procedure applied, have been made at [RHIC](#) [90, 91] and the [LHC](#) [89, 92]. In addition, projections of the Lund plane for jets initiated by a heavy quark have also been useful in exposing the long-sought-after Dead Cone effect [93] where emissions are suppressed in a cone depending on the m/E of the emitter.

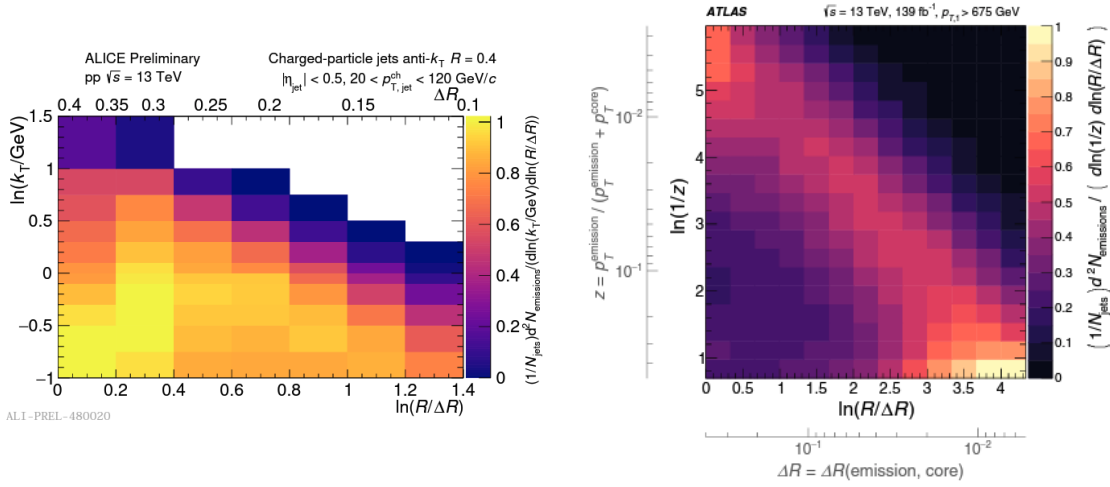


Figure 2.21: Left: ALICE fully-corrected primary Lund plane density for $R = 0.4$ jets with $20 < p_{T,\text{jet}} < 120$ GeV/ c in [pp](#) collisions at $\sqrt{s} = 13$ TeV, originally appearing in [20]. Right: ATLAS fully-corrected Lund plane for $R = 0.4$ jets with $p_T > 675$ GeV/ c in [pp](#) collisions at $\sqrt{s} = 13$ TeV, originally appearing in [21].

Overall, the Lund plane is an incredibly useful tool to study the substructure of jets. A clear and thorough understanding of the substructure in vacuum is essential to the results presented in this thesis, as the modification of jets in the [QGP](#) may differ for jets with a different substructure. One new set of observables that offers similar information as the Lund plane are Energy-Energy Correlators ([EECs](#)), which will be discussed in Chapter 8.

2.5 Jets in Heavy-ion Collisions

2.5.1 Jets as a Probe of the QGP

In heavy-ion collisions, the high p_T partons will interact with the colored medium resulting in jet energy loss and substructure modification, a phenomenon collectively referred to as

jet quenching [94]. The formation of partons via the hard-scattering occurs early in the collision before QGP formation, meaning that the partons (or jets) experience the full evolution of the medium. This feature of jets also makes them a powerful probe of the medium and its microscopic properties, such as the jet transport coefficient defined in Section 2.3.4. Another such property is the resolution length, which refers to the distance over which the QGP can “resolve” the color structure of the jet. Additionally, studies of jet quenching can be used to determine the extent to which the colored medium resolves the colored structure (or substructure) of the jet. For jet quenching measurements, jets in vacuum (pp) generally serve as a reference for jets in heavy-ion collisions, where any difference between the two systems serves as an in-medium effect. A cartoon diagram of quenched jets in heavy-ion collisions compared to un-quenched jets in pp collisions can be found in the left panel of Figure 2.22.

In the weakly-coupled limit [22], parton energy loss can take the form of elastic collisions with medium partons (collisional energy loss) or inelastic medium-induced gluon radiation (radiative energy loss). The Feynman diagrams for these processes are included in Figure 2.22. In the strongly-coupled limit, jets may lose energy via a drag force when traveling through the continuous liquid based on the conjectured AdS/CFT correspondence [95, 96].³⁵




Figure 2.22: (Left): Back-to-back jets in pp collisions and heavy-ion(AA) collisions. Jets in pp exist in vacuum whereas jets in AA traverse the QGP medium and lose energy. (Right): Feynman diagrams for the energy loss ΔE of a quark with energy E due to collisional (left diagram) and radiative (right diagram) processes, originally appearing in [22].

2.5.2 Expectations of energy loss

The expectations of energy loss discussed in this section are (1) suppression of jet yields (2) modification of jet (sub)structure and (3) a deflection of the jet centroid. This format will be used to refer to these expectations in the following sections.

³⁵See Section 2.2.2 for more details on the AdS/CFT correspondence.

Jet quenching has three main categories of expected signatures. One such expectation is that jets will exhibit an overall energy loss, resulting in a decrease of jet yield at a fixed p_T in heavy-ion collisions as compared to pp collisions, a phenomena referred to as jet suppression [97]. Jet suppression can be impacted by other factors such as the fraction of quark and gluon jets (see Figure 7.2) and the shape of the spectrum (see Figure 2.19), making many differential measurements of jet suppression useful for disentangling energy loss effects. Such effects will be explored using experimental data in Chapter 5 and in simulations in Chapter 7.

The second expected signature is the modification of the complex internal structure of the jet. These modifications can come from a variety of sources. Soft radiative and collisional interactions with the medium shift energy from higher momentum final-state particles to lower-momentum particles and broadens the jet, a phenomena referred to as momentum broadening [98], shown in the left panel of Figure 2.23. The medium can also induce partonic splittings, further modifying the substructure of the jet. Additionally, the QGP medium is expected to add a wake of soft particles to the jet cone in response to the presence of the jet itself (see Figure 2.26 and Ref. [95]).

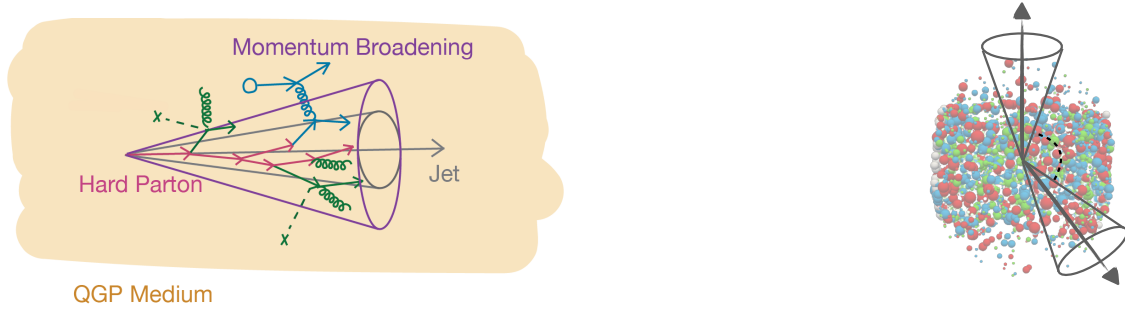


Figure 2.23: Momentum broadening (left) and the deflection of the jet centroid (right). The image for the background in the right panel is from the MADAI collaboration [99].

Depending on the resolution length, L_{res} , of the QGP, the colored medium can either resolve or not resolve the colored substructure of the jet introduced by partonic splittings, referred to as decoherence or coherence, respectively. These two effects are shown in Figure 2.24. When the resolution length of the medium is smaller than the distance between two subjects, the QGP resolves them as two effective energy loss sources, resulting in greater overall energy loss. When the resolution length is larger than the distance between two subjects, the QGP resolves them as one effective energy loss source, resulting in less overall energy loss.

The third expected signature is a deflection of the jet centroid (experimentally the jet axis) due to scatterings that can either occur as a single hard “Moliere” (point-like) scattering [100] or multiple soft scatterings with QGP quasi-particles [101]. Jet deflection is shown in the right panel of Figure 2.23.




Figure 2.24: Decoherence (left) vs. coherence (right).

One final expectation is that different jets with different partonic structures in vacuum, flavors, kinematics, or path lengths through the medium³⁶ etc. will all lose energy differently. Even the same jets (i.e. identical jets in vacuum) can lose energy differently due to fluctuations in energy loss. These differences can complicate the interpretation of inclusive jet measurements, but these differences can also be exploited in order to isolate and better understand various energy loss mechanisms. For example, looking at the dependence of jet-quenching effects on the path length can help disentangle the dominant form of energy loss. Collisional energy loss is expected to scale as the path length through the medium, L [94], radiative energy loss is expected to scale as L^2 [22], and energy loss in the strong coupling limit is expected to scale as L^3 [95] assuming a static medium. As jets are rare probes, making further selections to exploit these differences often requires new techniques, cleverly designed observables, and a wealth of available data. Recent years have shown drastic improvements in measuring these different jets, ushering in a new era of jet quenching measurements.

2.5.3 Experimental Observables

The experimental observables discussed in this section are sorted according to the corresponding expectation of jet quenching they probe (as listed in Section 2.5.2). This is not meant to be an exhaustive list, but rather to provide context and background for the results presented in this thesis. For a more complete review than what is provided here, see [102].

As listed in Section 2.5.2, jet suppression is the first experimental expectation of jet quenching. This is shown in Figure 2.25, where jets in heavy-ion (A+A) collisions lose energy (ΔE) as compared to the pp spectrum, resulting in a suppression of jet yields at a fixed p_T . The degree of suppression is determined not only by the energy loss ΔE , but also the steepness of the pp jet p_T spectrum.³⁷

Jet suppression effects can be quantified through the nuclear modification factor (R_{AA}), defined in Equation 2.18 as the ratio of the per-event jet yield in Pb–Pb and the cross section in pp collisions scaled by the average nuclear thickness ($\langle T_{AA} \rangle = \frac{\langle N_{coll} \rangle}{\sigma_{inelastic}^{NN}}$) computed as the ratio of the number of binary collisions between nucleons and the inelastic

³⁶For an image of this effect, see Figure 2.22.

³⁷The influence of this steepness will be discussed in detail in Chapter 5 and Chapter 8.




Figure 2.25: The jet p_T spectrum in pp and heavy-ion (A+A) collisions and the origin of jet suppression.

nucleon-nucleon cross section in a Glauber model [103] in order to account for the differing collision geometry.

$$R_{AA} = \frac{\frac{1}{\langle T_{AA} \rangle} \frac{1}{N_{\text{events, PbPb}}} \frac{d^2 N_{\text{jet}}^{AA}}{dp_T d\eta}}{\frac{d^2 \sigma_{pp}}{dp_T d\eta}} \quad (2.18)$$

The R_{AA} is a key signature for QGP formation that probes the combined effect of all medium modifications to an object [22]. Here, the jet R_{AA} will be discussed, but measurements of R_{AA} for other objects, for example charged hadrons, can also be useful. Differential measurements of jet suppression, for instance, as a function of substructure [104] or path length [105], can be especially useful in isolating different energy loss mechanisms. Another differential measurement of the R_{AA} depending on the resolution parameter will be discussed in Chapter 5.

The second experimental expectation is the modification of the internal structure of the jet. Jet substructure measurements can be categorized into two distinct types: measurements of jet structure and measurements of jet substructure. Jet structure variables are those that probe the distribution of radiation within a jet at the hadron level and jet substructure variables focus on the hard substructure of the jet at the parton level.

Some examples of jet structure variables include fragmentation functions, shown in Equation 2.19 that are defined using the longitudinal momentum fraction of a constituent $z = \frac{p_{T,i}}{p_{T,\text{jet}}} \cos(\Delta R)$,³⁸ and reflect how constituents are distributed with respect to the longitudinal momentum.

$$D(z) \equiv \frac{1}{N_{\text{jet}}} \frac{dn_{ch}}{dz} \quad (2.19)$$

The fragmentation functions of jets has been measured many times: for inclusive jets, photon-tagged jets, and D^0 mesons within jets. [30, 31, 106], showing a significant modification in heavy-ion collisions of this observable. The radial fragmentation function measures a similar quantity differentially in angle by looking at the number density in

³⁸In some cases, the $\cos(\Delta R)$ term is omitted in this definition.

transverse rings outward from the jet axis [107]. Jet shapes (also referred to as the jet profile) measures the radial distribution of momentum carried by the constituents of the jet as defined in Equation 2.20. Jet shapes have been measured many times: for inclusive jets [108, 109], photon-tagged jets [110], and to measure the distribution of the radial distribution of D^0 mesons within jets [111].

$$\rho(r) = \frac{1}{\delta r} \frac{1}{N_{\text{jet}}} \sum_{\text{jets}} \frac{\sum_{\text{tracks} \in (r_a, r_b)} p_{T, \text{track}}}{p_{T, \text{jet}}} \quad (2.20)$$

In addition to the per-constituent-particle observables mentioned above, there also exists jet structure observables that are measured per jet. The generalized angularities [112], defined in Equation 2.21, reflect a whole phase space of jet structure observables, each with a different influence of the momentum and angular components depending on the values of κ, α . Common jet structure observables, such as the jet mass ($\kappa = 1, \alpha = 2$), and jet girth ($\kappa = 1, \alpha = 1$) all represent different parts of this phase space. Measurements of the generalized angularities will be key to determining the extent that these observables are different and can help to resolve why some jet structure variables, such as the girth, show a strong modification [113], but others, such as the mass, do not [114].³⁹

$$\lambda_{\alpha}^{\kappa} = \left(\frac{p_{T, i}}{p_{T, \text{jet}}} \right)^{\kappa} \left(\frac{\Delta R_{i, \text{jet}}}{R} \right)^{\alpha} \quad (2.21)$$

Jet substructure observables probe the partonic splittings via subjects (or “prongs”) within the jet. As mentioned in Section 2.4.5, a grooming procedure is typically applied in order to reduce sensitivity to the non-perturbative effects and focus on the hard substructure of the jet. Some common groomed jet splitting observables are the groomed jet radius ($\theta_g \equiv \frac{R_g}{R} \equiv \frac{\sqrt{\Delta y_{1,2}^2 + \Delta_{1,2} \phi^2}}{R}$, where 1 and 2 refer to the leading and sub-leading subjects of the first splitting that satisfies the grooming condition, respectively), the groomed momentum splitting fraction ($z_g \equiv \frac{p_{T,2}}{p_{T,1} + p_{T,2}}$, also known as the splitting function), and the number of splittings in a jet that pass SD (n_{SD}). These observables have been measured in various different experimental contexts [92, 116, 117], where clear modifications are seen in the case of the groomed jet radius, but no significant modification of the splitting function or n_{SD} is seen. Comparisons with models for the suppression of wide angle splittings as seen in R_g are consistent with models including decoherence, but could also be consistent with coherence effects with a high quark fraction. Another jet substructure observable is the N-subjettiness, which reflects the rate of jets with N subjects (or prongs), as defined within C/A declustering, in the inclusive jet sample. When measured [118], the N -subjettiness showed no strong modification of the rate of 2-pronged subjects in heavy-ion collisions as

³⁹Measurements of the full phase space of generalized angularities have not been completed in heavy-ion collisions at the time of writing this thesis, though a measurement does exist in pp [115].

compared to vacuum. The fragmentation of subjects has also been measured [119], providing a test of the universality of jet fragmentation, and demonstrating no strong modification of these distributions in the presence of the QGP.

The third experimental expectation of jet quenching, as listed in Section 2.5.2 is the deflection of the jet centroid. The central observable to measure this is the jet acoplanarity (Δ_{recoil}), which is defined as the difference between the trigger-normalized recoiling jet distributions in signal and reference-trigger track p_T intervals. Measurements from both RHIC [120]⁴⁰ and the LHC⁴¹ indicate a medium-induced broadening of this distribution for large R and low p_T . Although recent work suggests that such observations may not be due to in-medium scatterings, but instead, are due to the wake [121].⁴² Experimentally, in-medium scatterings could also manifest themselves as an enhancement of large groomed k_T ⁴³ splittings in heavy-ions, though this observable is also sensitive to the substructure modification of the groomed jet observables mentioned previously. When measured experimentally, no such enhancement of large $k_{T,g}$ splittings is found, but a hint of a modification at low $k_{T,g}$ is observed [122].⁴⁴ More work is therefore needed to determine definitively the nature and extent of jet deflection in the QGP.

The goal of experimental jet observables is to determine for every jet in Pb–Pb collisions, what would the equivalent jet look like in pp, and therefore how was it modified. As a result, many of the observables discussed above study modification by taking the ratio of distributions in heavy-ion and pp collisions at a fixed p_T . However, the fraction of quark or gluon jets, which may be differently modified due to their differing color factors, may change at a fixed p_T . This alone could create behavior that mimics a modification [123], making the interpretation of such variables difficult. This effect is called the *survivor bias* and presents an ongoing challenge for heavy-ion jet physics [124–126]. Some ways to overcome this are to use γ - or Z -tagged jets [127], where the electroweak probe serves as a proxy for the unmodified p_T . Additionally, making differential measurements is a good way to probe modification over many regions in phase space, where such effects may have a different influence.

2.5.4 Modeling Jet Quenching

In order to model jet-quenching effects, two main categories of effects can be included. The first is the impact of the QGP on the jet, which can be formulated via weakly- or strongly-coupled mechanisms. The weakly-coupled limit consists of radiative and collisional energy loss mechanisms as described earlier and shown in Figure 2.22. Various

⁴⁰Note this work is preliminary and not yet peer-reviewed.

⁴¹<https://rb.gy/yxpiu6> Note this work is preliminary and not yet peer-reviewed.

⁴²Note this work is preliminary and not yet peer-reviewed.

⁴³Where $k_{T,g}$ is the k_T of the splitting, as defined in Section 2.4 after a grooming procedure has been applied.

⁴⁴Note this work is preliminary and not yet peer-reviewed.

pQCD formalisms implement these mechanisms either as multiple soft scatterings such as in the Baier-Dokshitzer-Mueller-Peigne-Schiff-Zakharov (BDMPS-Z) [128, 129], Arnold-Moore-Yaffe (AMY) [130], and Amesto-Salgado-Wiedemann (ASW) [131] formalisms, or as few hard scatterings as in the Higher Twist (HT) [132] and Gyulassy-Levai-Vitev (GLV) [133] formalisms. The formalism used will determine the extent of the quenching, and each will result in slightly different predictions for energy loss. The strong coupling limit utilizes AdS/CFT where the main mechanism for energy loss is a drag force. Different models use different combinations of the above implementations.

The influence of the jet on the medium can also impact jet quenching as the response of the medium can result in additional particles being added to the jet cone. The medium response can also be implemented in a weakly- or strongly-coupled way. In the case of the weakly-coupled limit, the medium response is implemented using a kinetic-theory-based approach where medium partons pick up energy scattered from the jet and recoil [22]. Correspondingly after the recoil there remains a hole in place of the parton. This process is illustrated in the left panel of Figure 2.26. In the strongly-coupled limit, the medium response is implemented via hydrodynamic theory where the evolution is a bulk medium with a diffusive wake. In addition, there can also be a negative wake formed behind a hard parton, which is analogous to holes in the weakly-coupled case. An illustration of the hydrodynamic medium response via a wake can be found in the right panel of Figure 2.26.




Figure 2.26: The weakly-coupled medium response via recoils (left) and a strongly-coupled medium response via a hydrodynamic wake (right).

These different modeling choices can result in many different models, each of which with varying predictions for jet-quenching effects. Models can therefore be categorized by the effects implemented, as shown schematically in Figure 2.27. Note that this is only one of the ways to categorize jet quenching models. Below, some of the available jet quenching models will be described and categorized in this scheme.

The Hybrid Model [95] is the only strongly-coupled approach with an AdS/CFT-inspired drag force and a hydrodynamic wake for the medium response. Another common model, JEWEL [134, 135], consists of a weakly-coupled MC implementation of BDMPS-based medium-induced gluon radiation in a medium modeled with a Bjorken expansion. JEWEL additionally includes the option for whether or not to implement the medium re-

sponse via recoils [136]. The **Linear Boltzman Transport (LBT)** model [137, 138] uses linear Boltzmann equations⁴⁵ in order to describe the weakly-coupled energy loss of partons in the **QGP**. The effect of the medium response is modeled via recoiling thermal partons. **LIDO** [139], also uses linear Boltzmann equations for describing jet energy loss, but instead a hydrodynamic medium response is included. Mehtar-Tani et al. [140] is a first-principle analytical calculation for energy loss including soft energy flow and a hydrodynamic medium evolution. **MARTINI** [141] embeds partons into a hydrodynamic medium with a modified parton shower. Qiu et al. [34], herein referred to as the Factorization Model, is based on a factorization approach inspired by phenomenological considerations. **JETSCAPE** [142], is a modular framework that can include a medium-modified shower as well as the medium response in both the weakly- and strongly-coupled approaches. Therefore, JETSCAPE is not included in Figure 2.27 as it can in principle cover any part of the phase space, depending on the choice of the user.

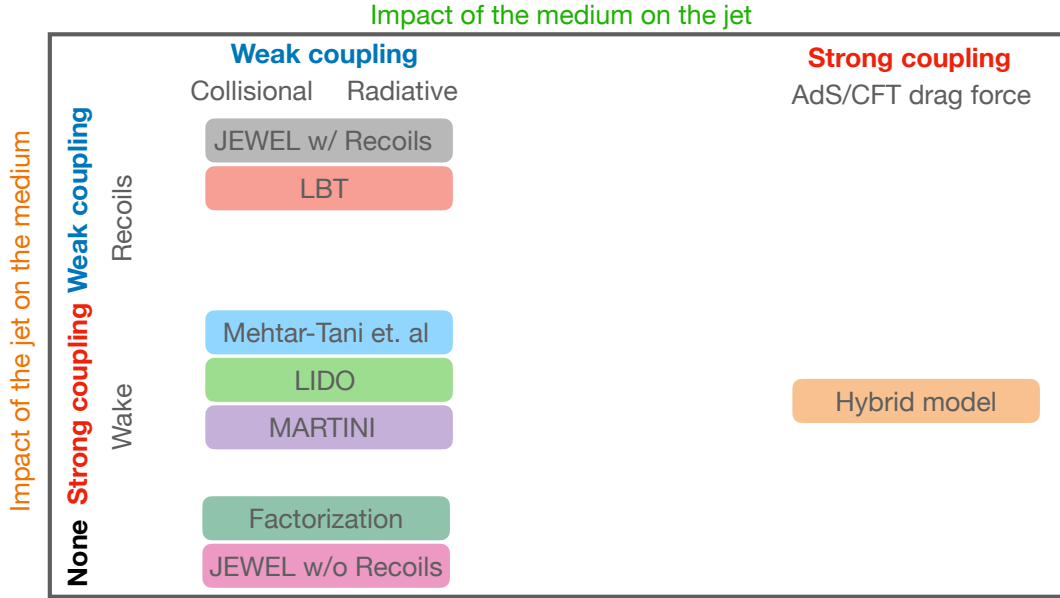


Figure 2.27: A categorization scheme of various different available jet quenching models.

As of the time of the writing of this thesis, there is no clear choice for the “best” model of jet quenching. Theoretical predictions and experimental measurements evolve together. Therefore, in order to have a better theoretical understanding there must also be new experimental jet measurements. The R -dependence of the R_{AA} , as discussed in Chapter 5, is an incredibly discriminative observable and is therefore critical in this effort.

⁴⁵The Boltzmann equations describe the behavior of a thermodynamic system not in equilibrium.

2.6 Machine Learning

Results in this thesis employ the modern analysis technique of [machine learning \(ML\)](#). This section serves as a brief overview of [ML](#) and the concepts related to the analysis. A special focus will be given to applying [ML](#) techniques for jet physics. For a more extensive overview on [ML](#) in high-energy physics than what can be provided here see [\[143\]](#).

2.6.1 What is Machine Learning?

This section will provide a general introduction to [ML](#). Some algorithms will be mentioned in this Section, but not described, for illustrative purposes. For an in-depth explanation of these algorithms, refer to Section [2.6.2](#) or the corresponding citation.

Machine learning, most simply defined, is any algorithm that imitates human learning, meaning that it improves in accuracy over time. This learning can be roughly grouped into three different types - supervised learning, unsupervised learning, and reinforcement learning. Supervised learning is when the algorithm will learn from a labeled set of true values. In this case, the learning is driven by the algorithm's ability to perform the task. Unsupervised learning is when the algorithm finds structure in the data without knowing the desired outcome. In this case, the learning is driven by the data. Reinforcement learning is when the algorithm will learn in a reward-based system to determine a series of actions. In this case, the learning is driven by the reward.

[ML](#) problems can also be divided into three classes: classification problems, regression problems, and generation (or synthesis) problems. Classification problems seek to group objects into predefined classes, whereas regression problems seek to assign a value to a given sample. Generation (or synthesis) problems seek to generate a sample that mimics desired features.

The general procedure for all machine learning applications can be found in Figure [2.28](#). The first step of this procedure is to collect good and reliable input data. Any learning algorithm will learn solely based on the data that is provided to it, so any [ML](#) application can only perform as well as the input data provided to it. For example, if non-realistic or unintended features are present in the data, then these features will be learned by the [ML](#). Commonly, this is summarized with the saying “garbage in, garbage out”. For the current analysis, the data sample will be discussed in Section [5.2](#).

The next step of the procedure is to design a minimal set of discriminative features. The learning algorithm seeks to create a relation between the features (or input variables/parameters) and the target (or truth definition). Therefore, it is very important to select unbiased input parameters that are relevant to the problem at hand.

In general, there is no standard procedure for choosing input variables for the [ML](#). One useful method to do this is to first create a simplistic model with good interpretability

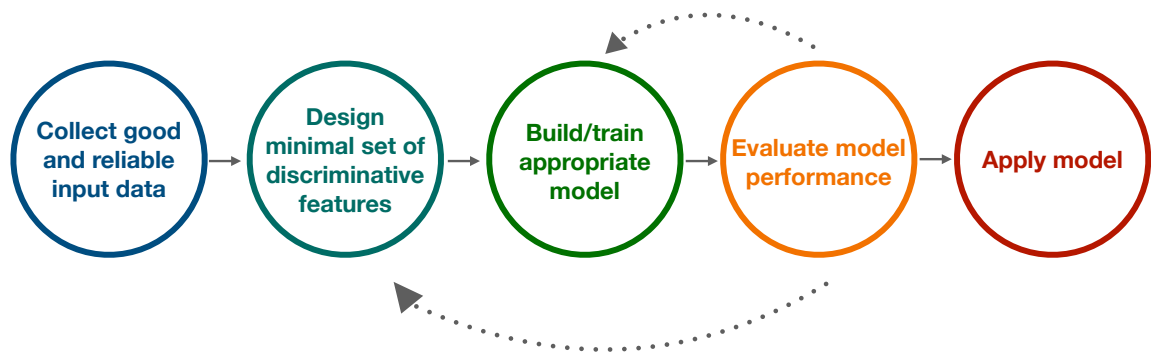


Figure 2.28: The general workflow for a given ML application.

tools and then utilize these tools to select the optimal set of parameters. This procedure is motivated by the fact that, to first order, different models faced with the same problem will learn similar features. The first step of this procedure is to create a model with all possible input features. Then for each potential variable, one can ask two useful questions:

1. How correlated is the variable with other variables in the model?
2. How important is the variable to the overall performance of the model?

The correlation between variables can be visualized via so-called heat maps (see Figure 5.6) or simply via 2D scatter plots of the direct correlations (see Figure A.1). Highly correlated variables should be removed as this can mask correlations between input parameters and the target and can result in a decrease in performance. The importance of different variables to the model can be quantified via these different interpretability tools. For example, random forest models [144] implemented in `scikit-learn`[145] come with built-in interpretability tools such as the feature importances (Gini-index[146], see Table 5.1) or a visualization of the decision tree. Features that are important to the model will have a high feature importance and will be used to make decisions often in the decision tree. Input variables that do not have a high importance should be removed as they increase the bias on the training sample while not providing useful information. After removal of unimportant or correlated parameters, a new model should be created and tested in an iterative procedure until the optimal parameter list is achieved.

The next step in the generalized procedure of an ML application is to build and train an appropriate model. There are many different learning algorithms available each with many different possible architectures. A subset of these algorithms is described in greater detail in Section 2.6.2. No specific ML algorithm will consistently outperform all other algorithms in all different cases as dictated by the No Free Lunch Theorem [147]. In other words, different algorithms are better suited for different problems. For example, tools like neural networks (NNs) [148] are great for regression tasks, tools like random forests or boosted decision trees [149] are great for performing classification tasks, and tools like

Generative Adversarial Networks (GANs) [150] are useful in generation problems. In order to choose the appropriate model, one should not only consider the problem, but also the available input. For example, tools such as convolutional neural networks (CNNs) [151] are useful tools for processing images and tools such as recurrent neural networks (RNNs) are useful in processing time-dependent input. The best approach is to first choose a number of potential models, then compare the results from each. Note that the ML-based result should not depend strongly on the learning algorithm used (within uncertainties), but it is normal for different algorithms to demonstrate different performance.

In any ML application, there are also many hyperparameters, i.e. parameters that change the model, but cannot be learned from data. Choosing the correct values for these hyperparameters generally consists of defining all possible “reasonable” values and then sampling them while evaluating the performance of the model for each set of hyperparameters. The optimal set of hyperparameters is the set that gives the best performance. Specific methods and tools for hyperparameter optimization will not be discussed here; see Ref. [152] for a complete review.

Once an appropriate model has been chosen, one must evaluate its performance. The type of performance evaluation will depend on the problem, but in general the performance of any supervised ML algorithm should reflect its accuracy (how well the model reproduces the target) and its robustness (how well the model performs on unseen data). One common performance evaluation technique for a classifier distinguishing between two classes is the area under the Receiver Operating Characteristic (ROC) curve [153]. The ROC curve is shown in Figure 2.29. This curve is given as the true positive rate, or the rate of correct classification, as a function of the false positive rate, or the rate of incorrect classification. For other physics applications, the evaluation of model performance is usually treated the same as the non-ML method. For more information, see Section 5.6.

The final step of the general ML procedure is to apply the model, which will greatly vary depending on the application. The application of ML to unseen data should not occur until input parameter optimization, model design, and model evaluation have been performed. Often this process requires many iterations and careful thought before proceeding. Another important detail of model application is to quantify the systematic uncertainty associated with the ML-based method, which is also application specific.

Overall, ML is a useful tool for solving complex problems in the modern world. Its impact has been seen in applications ranging from personal assistant software, to video surveillance and more. Of interest to this thesis is the application of ML for physics applications. The application of ML for High-Energy Physics (HEP) and for jet physics in particular, will be discussed in Section 2.6.3.




Figure 2.29: True positive rate as a function of the false positive rate, also called the ROC curve.

2.6.2 Decision Trees and Neural Networks

As mentioned in Section 2.6.1, there are many different *ML* algorithms available for learning, each with its advantages and disadvantages. As neural networks and decision trees are employed in this thesis, a more in-depth explanation of how these methods work will be provided here.⁴⁶

Decision Trees

One easy analogy in order to understand how decision trees work is to first think of the game popular since the 19th century, *Twenty Questions*.⁴⁷ This game involves two parties, the questioner and the answerer. The answerer first thinks of an object and the questioner then asks questions and the answerer must answer “yes”, “no”, “sometimes”, or “unknown” (where the possible answers can vary based on the variant). In order to win the game, the questioner must guess the object in twenty questions, otherwise the answerer wins. It is advantageous for the questioner to ask questions that cover a broad scope to significantly narrow the possible number of objects, a task that is difficult to develop on instinct alone. Decision trees, in essence, are the algorithmic equivalent of an optimal questioner that asks the best possible question at each point to “identify” the class or value

⁴⁶For a more detailed explanation, see <https://cs.colby.edu/courses/S17/cs251/CS251-S17-Lectures.pdf>

⁴⁷<http://www.20q.net/>

of an object.

Decision trees, as the name suggests, are set up as tree-like structures made of decision nodes, each of which subdivides the data into two or more parts. The root node refers to the first node in the tree from which all other nodes branch out and the term leaf nodes refer to nodes at the edges of trees that assign a predicted class to an object (for classification problems) or a value (for regression problems). These nodes can either be enumerated types, where there is a branch for each type, or numeric types that represent a threshold test with a single constant value, which subdivides the data into two parts.

In training, the goal is to create the set of decisions (or questions) that can be made in order to maximize the information gain with each question. The information content of a node is measured via entropy, as defined in Equation 2.22 where $E(x)$ is the entropy and $p(x)$ is the probability of having the specified feature in data.

$$E(x) = - \sum p(x) \log_2(p(x)) \quad (2.22)$$

The metric used in order to evaluate potential decisions (or questions) is the **information gain (IG)**. This is mathematically specified as the entropy of the parent node (calculated with Equation 2.22) minus the weighted entropies of the child nodes as defined in Equation 2.23. Here the weight (w_x) corresponds to the relative size (in number of samples) of the child node relative to the parent.

$$IG = E(\text{parent}) - \sum w_x E(\text{child}) \quad (2.23)$$

The general procedure of training a decision tree is to evaluate all possible splits (or decisions) and then choose the optimal split of the data set at each step. This type of algorithm is referred to as a greedy algorithm. Greedy algorithms make optimal local decisions whether or not these decisions result in optimal global decisions. For this reason, algorithms that combine many different individual decision trees either in series (boosted decision trees) or in parallel (random forest decision trees) are typically preferred to reduce potential biases and improve performance.

Neural Networks

As mentioned in Section 2.6.1, **ML** is defined as any algorithm that imitates human learning. Neural networks seek to do this via a structure that mimics how information flows in the human brain. In a feed-forward **NN** the flow of information happens between nodes from the input layer to the output layer, as visualized in Figure 2.30. The input layer contains the input parameters that the network will use to learn and the output layer provides the result of the network.

As shown in Figure 2.30, each node is fully connected to every node in the next layer. These connections also have a weight, which reflects the node's importance in determining




Figure 2.30: General structure of a feed-forward neural network, where the arrows represent the direction of the flow of information.

the output value. Each node also contains a bias that can be thought of as a preconceived notion of how much that node should factor into the overall decision. The output of a given node (O_j) is generally written as the sum of weighted inputs ($w_{ij}O_i$) plus a bias term (B_j), which is given in Equation 2.24.

$$O_j = \left(\sum_{i=0}^{N-1} w_{ij}O_i \right) + B_j \quad (2.24)$$

In order to introduce non-linearity into the network that allows it to learn more complex relations, the results of this process are additionally passed through an activation function, which is a function that produces a fixed range of outputs from an arbitrary range of inputs. There are many potential choices of the activation function for NNs, where in this analysis a ReLU activation function⁴⁸ is used.

In training, the NN seeks to determine the optimal set of weights and biases that yields the best performance. In the initial case, the set of weights and biases are random and the information is propagated forward in the NN, where the output of a given node is calculated via the above-mentioned procedure. This output is then evaluated via a loss function, which evaluates the performance of the NN by how far off the result from the NN is from the desired values. Then, in a process called backpropagation the weights and biases will be updated moving from the output to the input in order to improve the model. The backpropagation process is controlled via parameters such as the learning rate, which

⁴⁸<https://machinelearningmastery.com/rectified-linear-activation-function-for-deep-learning-neural-networks/>

specifies how much each step should impact the weights/biases and the momentum, which specifies how much previous outcomes should impact the new weights/biases. Note that each of these parameters that cannot be directly learned from the data (hyperparameters) should be separately tuned in a process referred to as a hyperparameter tuning procedure, which will not be covered here. See Ref. [152] for a full review.

There are many different types of NNs that can be used in a variety of contexts. Networks with one hidden layer, as visualized in Figure 2.30, are referred to as shallow NNs. Networks with more than one hidden layer are referred to as deep neural networks (DNNs). The deeper a DNN becomes, the more abstract the features become and the more complex relationships between variables that can be learned. For this reason, DNNs are commonly used both in academia and in industry. RNNs allow information to travel in both directions by introducing loops into the node structure. This is suitable for sequential input data such as data with a time dependence. CNNs are used for images as inputs where, with the help of a filter, they can be used to determine whether or not a specific feature is present. NNs can also be combined into more complex structures such as GANs where two NNs will compete with one another in a game in order to generate samples with a set of desired features. Additionally, two NNs can be combined to form an autoencoder [154], particularly useful in denoising data.

2.6.3 ML for jets

Typically in HEP analyses, a series of boolean decisions are employed to select data after which a statistical analysis is performed. However, when data sets and physics objects increase in complexity, the optimal set of decisions is difficult to derive from expert knowledge alone. This difficulty prompted the application of algorithms that select multiple variables simultaneously, inspiring countless ML HEP analyses [155]. The large and complex data sets available in HEP make it an ideal environment for ML applications. The LHC data volume represents a medium-sized ML application as compared with some industry applications. ML has proven to be useful in quality control, detector simulation, triggering, event classification, object identification, background reduction and more.

ML is also an ideal tool for jet physics as it can exploit the multi-dimensional complex nature of jets in order to learn. Jets can additionally be represented in many ways for ML, allowing a variety of learning algorithms to be used. For example, jets can be represented either as a collection of objects via constituent parameters such as the declustering history, ordering of constituent momenta, or other related variables. Jets can also be represented as a single object via properties of the jet itself, such as the jet transverse momentum, jet mass, radial moment or other variables. In order to use advanced ML tools developed for image processing, there has also been some success in representing the jet via a so-called jet image formed by taking the constituents of the jet and mapping its energy into pixels in (η, ϕ) space [156, 125, 157]. The many different possible representations of jets as well

as their complexity allow them to be an ideal playground for [ML](#) applications.

There have been many successful applications of [ML](#) for jet physics. Most of these applications can be categorized as simulation-based inference, where the [ML](#) learns in a supervised manner from a [MC](#) simulation. The most common use of [ML](#) for jet measurements in pp collisions is to use ML techniques such as [DNNs](#) or boosted decision trees to tag different jet topologies, such as jets originating from a heavy-quark [[158–161](#)]. Another common application is to use [ML](#) techniques in order to distinguish between quark and gluon jets. Such applications have proven to be successful in simulation [[162–165](#)] (for a complete list, see Ref [[155](#)]), but applications to experimental data are still quite rare. Recently, [ML](#) has also been utilized to correct for detector effects in high-energy collisions using the `OmniFold` technique [[166](#)]. This technique has been applied to archived H1 e^+e^- data to measure lepton-jet correlations [[167](#)] and by the [LHCb](#) collaboration to measure distributions of charged hadrons in Z -tagged jets [[168](#)].

Though much progress has been made towards [ML](#) applications for jet physics in both pp collisions and e^+e^- collisions, applications of [ML](#) to jets in heavy-ion collisions are still quite rare. This is due largely to the complexity of these collisions and lack of simulations that accurately describe jet-quenching effects observed in data. Additionally, jets in heavy-ion collisions have a large uncorrelated background, which adds an additional difficulty for [ML](#). Despite these difficulties, there have been an increasing number of approaches with demonstrated success in simulations [[125](#), [169](#)]. The results mentioned in this thesis represent the first application of [ML](#) used to make a jet measurement with heavy-ion data. The difficulties mentioned above are overcome by utilizing simple and shallow ML techniques to avoid a large dependence on simulation with increased interpretability. For more information on this technique, refer to Section [5.4](#).

In the future, [ML](#) applications in [HEP](#) and in jet physics will only continue to increase. These efforts, as indicated by recent feasibility studies [[170](#)], will likely include not only simulation-based inference but also data-based learning. In addition, due to the large volumes of data foreseen in the next decades, [ML](#) will become an essential tool in order to reduce the data volume.

Chapter 3

The ALICE Detector

3.1 Overview

The conditions mentioned here reflect the detector conditions in place during Run 2 of the LHC. For upgrades related to the ongoing Run 3 of the LHC, refer to Section 3.4.

“A Large Ion Collider Experiment” (ALICE) is an experiment designed and optimized for the study of ultra-relativistic heavy-ion collisions at the LHC [171]. Located at Point 2 in St. Genis, France, the ALICE detector sits 56 m below the surface, weighs $\sim 10,000$ tons¹ and is 26 m long, 16 m high, and 16 m wide. To accomplish its physics goals, the ALICE detector must be able to measure particle p_T over a few orders of magnitude ($0.15 < p_T < 100$ GeV/c) and accurately perform Particle Identification (PID), all in very high number (or multiplicity) of particles $\sim 10^3$ per unit rapidity.

The ALICE detector is composed of numerous sub-detectors, each of which contributes unique information regarding the collision that is then combined for the purposes of physics analysis. A schematic of these sub-detectors is shown in Figure 3.1. The main sub-detectors are located in the central barrel in the mid-rapidity region ($|\eta| \lesssim 0.9$) immersed in a magnetic field of 0.5 T. Particularly relevant to this thesis are the tracking and calorimetry systems used for the reconstruction of charged jets and full jets.² These will be discussed in more detail in Sections 3.2 and 3.3. In Table 3.1, a brief description of sub-detectors is provided, numbered according to the scheme in Figure 3.1.

¹For reference, the Eiffel Tower weighs about 7,300 tons.

²See Section 2.4 for more information on jets.

THE ALICE DETECTOR




Figure 3.1: Schematic view of the ALICE detector during Run 2 of the LHC [23].

Table 3.1: Description of the various sub-detectors of the ALICE detector. Numbers in parenthesis (ex: (1)) correspond with the numbering scheme in Figure 3.1.

Name	Detector Description	Acceptance	Physics Purpose
ITS(1)	6-layer silicon tracker	$ \eta < 0.9$	tracking
FMD(2)	silicon strip detector	$-3.4 < \eta < -1.7$ $1.7 < \eta < 5$	charged particle multiplicity
T0 [172](2)	Cherenkov detector	$-3.3 < \eta < -3$ $4.6 < \eta < 4.9$	timing & position of the interaction point
V0 [173](2)	scintillator detector	$-3.7 < \eta < -1.7$, $2.8 < \eta < 5.1$	centrality and multiplicity, defines the min bias trigger
TPC(3)	time projection chamber	$ \eta < 0.9$	tracking/PID
TRD [174](4)	transition radiation detector	$ \eta < 0.84$	e^- PID at intermediate p_T
TOF [175](5)	Multigap Resistive Plate Chamber strip	$ \eta < 0.9$	time of flight
HMPID(6)	Ring Imaging Cherenkov Detector	$ \eta < 0.6$ $\Delta\phi = 57.6^\circ$	PID for high p_T charged particles
EMCal(7)	sampling calorimeter	$ \eta < 0.7$ $80^\circ < \phi < 187^\circ$	photons, electrons, and neutral mesons
DCal(8)	sampling calorimeter	$ \eta < 0.7$ $253^\circ < \phi < 320^\circ$	dijet measurements
PHOS(9)	sampling calorimeter	$ \eta < 0.12$, $\Delta\phi = 70^\circ$	photons, electrons, and neutral mesons
CPV(9)	proportional chambers w/ cathode readout	$ \eta < 0.12$ $\Delta\phi = 60^\circ$	veto of charged particles
L3(10)	solenoid magnet	full acceptance	generates 0.5 T magnetic field
MS(11-15)	cathode pad/strip chambers	$-4 < \eta < -2.5$	measurement of high p_T muons
PMD(16)	preshower/gas counter	$2.3 < \eta < 3.9$	photon multiplicity
AD(17)	plastic scintillators	$-6.9 < \eta < -4.9$ $4.7 < \eta < 6.3$	improve measurement of diffractive pp scatterings
ZDC(18)	calorimeter	$4.8 < \eta < 5.7$	detecting spectators, rejecting pileup
ACORDE(19)	plastic scintillator	2 188 x 20 cm ² paddles	high-energy cosmic rays

3.2 ALICE Tracking

The **ALICE** tracking system seeks to identify the location, momentum, and identity of charged particles by measuring the curvature of a charged track in the presence of a magnetic field. First, the detector components of the tracking system will be described, then the particle identification and track reconstruction algorithm will also be described.

3.2.1 Inner Tracking System

The **ITS** is the first detector located radially outward from the interaction point. The central motivation of the **ITS** is to determine the primary interaction vertex and to separate this vertex from the secondary decay vertex of particles with a short lifetime, such as D and B mesons. The **ITS** consists of six silicon layers, chosen for the speed and good spatial resolution, covering the full pseudorapidity range ($|\eta| < 0.9$). These six layers are composed of two Silicon Pixel Detectors (**SPDs**), two Silicon Drift Detectors (**SDDs**), and two Silicon Strip Detectors (**SSDs**). The first two **SPD** layers, which extend to $|\eta| < 2$, utilize a logical readout for the precise determination of the primary vertex. The next two **SDD** layers have good spatial resolution ($35\ \mu\text{m}$ along the azimuthal direction and $25\ \mu\text{m}$ along the beam direction) and can be used for a measurement of dE/dx (as defined in Equation 3.1). The outer two **SDD** layers provide a spatial resolution of $20\ \mu\text{m}$ along the azimuthal direction and $830\ \mu\text{m}$ along the beam direction and can still be utilized to measure the dE/dx .

3.2.2 Time Projection Chamber

TPCs are widely utilized in modern physics experiments for the detection of charged particles. Generally, **TPCs** at colliders consist of a cylindrical detection volume located symmetrically around the beam axis filled with a gas mixture (usually 90% noble gas and 10% quencher gas). Typically the **TPC** volume contains a high-voltage electrode that divides the volume in half. This is accompanied by multi-proportional wire chambers that sit on either side as end-plates with a readout system, creating an electric field. As a charged particle traverses the **TPC** it will ionize the gas in the drift volume along its trajectory. These electrons then drift towards the anode and are read out, allowing for the measurement of the charged particle's trajectory. The positive ion also remains in the drift volume, but drifts at a much slower rate than the ionized electrons, causing an accumulation of so-called space charge. This space charge can distort the electric field, an effect that must be corrected for in experiments.

The **ALICE TPC** [24], follows this general scheme, occupying a cylindrical shape (2π in azimuth and $|\eta| < 0.9$) with a volume of $\sim 90\ \text{m}^3$ that is divided into two drift regions by an electrode located at the center of the chamber, as shown in the left panel of Figure

3.2. The gas utilized in the ALICE TPC for Run 2 is a mixture of Ar (88%) and CO₂ (12%).³

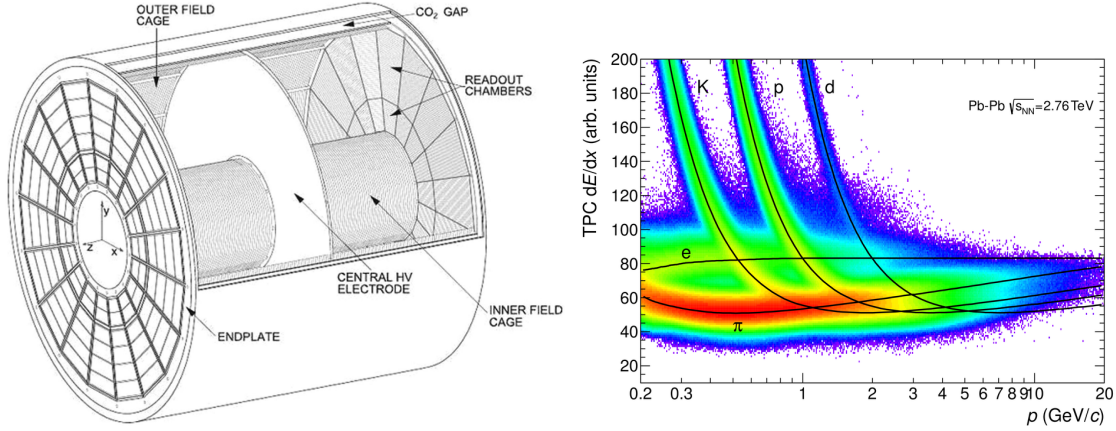


Figure 3.2: Left: Schematic of the ALICE TPC, originally appearing in [24]. Right: dE/dx (see Equation 3.1) distribution as measured in the TPC as a function of momentum, originally appearing in [25]. The black lines correspond to the expected mean energy loss according to Equation 3.1.

3.2.3 Track Reconstruction

The ALICE track reconstruction procedure is done using a Kalman Filter approach [176]. The track reconstruction process begins with a clusterization procedure performed separately in each detector to form “clusters” with a characteristic position, signal amplitude, signal time, and error. So-called tracklets in the SPD are then formed by connecting clusters in the two layers of the SPD. The single space point where a maximum number of tracklets converge in the SPD is then used as a preliminary interaction vertex. Cluster information is then combined to form a track in a procedure with three iterations. The first iteration starts at the TPC and works inwards, matching TPC tracks to those in the ITS in order to construct a vertex. Then in the second iteration tracks are propagated from this vertex outwards through the ITS, TPC and to outer detectors such as the TRD and TOF. The third iteration moves from the outward TRD and TOF and moves inward, further constraining the primary vertex. This results in a tracking efficiency that is approximately 67% for track $p_T = 0.15$ GeV/c, increasing to approximately 84% at track $p_T = 1$ GeV/c and remaining above 75% at higher track p_T [14]. The tracking efficiency in 0–10% Pb–Pb collisions is estimated using simulations as an approximately 2% reduction in the tracking efficiency as compared to pp, independent of the track p_T . The momentum resolution in pp collisions is about 1% at a track p_T of 1 GeV/c and about 4% at 50 GeV/c. In heavy-ion

³Note these values were for Run 2 of the LHC - different running periods may have different gas mixtures.

collisions, the momentum resolution at high track p_T is approximately 10–15% worse than in pp collisions.

3.2.4 Particle Identification

The charge of a particle can be determined via the direction of curvature of the charged track in the presence of the magnetic field, but this information alone is not sufficient to identify the particle. One useful technique to identify particles is via the dE/dx distributions, which are derived by measuring the energy loss of particles. This energy loss will be according to the Bethe-Bloch formula [177], written in Equation 3.1 where K is a constant, z is the charge number of the incident particle, Z is the atomic number of the absorber, I is the mean excitation energy, $\delta(\beta\gamma)$ is the density effect correction to ionization energy loss, and W_{\max} is the maximum energy transfer for a single collision.

$$\left\langle -\frac{dE}{dx} \right\rangle = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (3.1)$$

Note that Equation 3.1 only depends on the velocity (β), not the mass or momentum. In order to identify particles, this dE/dx distribution is written as a function of momentum as shown in the black lines in the right panel of Figure 3.2.

In ALICE, the PID is performed using the ITS, TPC, TOF and the HMPID. As mentioned in Section 3.2.1, the outer four layers of the ITS can provide a measure of the dE/dx that is useful for the measure of low p_T tracks ($p_{T,\text{track}} < 0.7 \text{ GeV}/c$). The TPC also provides dE/dx that is useful over broad ranges of momenta, as shown in the right panel of Figure 3.2. The TOF detector is dedicated to PID, measuring the arrival time of particles that is useful for kaon/proton separation for track p_T 's of up to approximately 4 GeV/ c . The timing resolution of the TOF detector for the 2015 Pb–Pb data-taking period, as analyzed in this thesis, was 56 ps. The HMPID is also used for kaon/proton separation up to high track p_T 's of approximately 5 GeV/ c . The combination of information from each of these detectors allows for good particle PID across broad p_T ranges as shown in Figure 3.3 for pions, kaons, and protons in the light green bands and for tritium, deuterium and Helium-3 in the dark blue bands.

3.3 ALICE EMCal

Electromagnetic calorimeters are designed to measure the energy of photons, neutral mesons, and electrons as they interact with the electrically-charged particles in matter. Photons and electrons will produce electromagnetic showers when they encounter a high-density material, depositing all of their energy in its volume. These electromagnetic showers can occur via two processes; pair production in which a photon can convert into an electron-positron pair and Bremsstrahlung radiation where these electrons and positrons can emit




Figure 3.3: Schematic of ALICE [PID](#) and reconstruction where the p_T ranges shown correspond to those in published measurements, adapted from [6].

further photons. The combination of these two processes continues in a shower until the energy decreases such that pair production is not possible and Bremsstrahlung dominates until all available energy is radiated. Charged hadrons can also deposit some amount of energy in the calorimeter by hadronically scattering with the calorimeter material. There are two main types of calorimeters, sampling and non-sampling. Sampling calorimeters are typically layered with a passive or “absorbing” high density material (like Pb) and an active medium, such as plastic scintillators, serves as the detector. Sampling calorimeters only sample part of the shower but can be more cost-effective. Non-sampling calorimeters utilize a single homogeneous medium (such as lead-tungstate) as both the absorber and the detector.

The [ALICE EMCal](#) [26] is a sampling calorimeter made up of layered Pb-scintillator with a partial acceptance of $\Delta\phi = 107^\circ$ and $|\eta| < 0.7$. The [EMCal](#) is subdivided into 12 supermodules, 10 full-sized modules and two one-third-sized modules. The full-sized supermodules are each composed of $12 \times 24 = 288$ modules arranged in 24 strip modules of 12×1 modules each. Correspondingly, the one-third size modules contain $4 \times 24 = 96$ modules. In total (not including [DCal](#)), the [EMCal](#) has 12,288 towers (also referred to as cells), each spanning $\Delta\eta \times \Delta\phi = 0.014 \times 0.014$. The [Dijet Calorimeter \(DCal\)](#) was added for the purposes of measuring dijet and hadron-jet correlations.⁴ This additional arm adds 5,376 towers for a total of 17,664 towers between the two calorimeters. The [DCal](#) is divided into 3 main regions with the following acceptances: (1): $-0.7 < \eta < -0.22$ and $260^\circ < \phi < 320^\circ$, (2): $0.22 < \eta < 0.7$ and $260^\circ < \phi < 320^\circ$, (3): $-0.7 < \eta < 0.7$ and $320^\circ < \phi < 327^\circ$. A schematic diagram of the [EMCal](#) and the [DCal](#) are shown in Figure 3.4. The [PHOS](#) calorimeter sits inside a hole in the [DCal](#), shown in tan in Figure

⁴Neither the [DCal](#) nor the [PHOS](#) detectors were utilized in the full jet measurement as discussed in Chapter 6 as the additional coverage was not needed.

3.4, covering an acceptance of $|\eta| < 0.12$, $220^\circ < \phi < 320^\circ$.

When a particle interacts with the cell material it will then shower as described above spreading its energy into neighboring cells. In order to reconstruct the energy deposited from a single particle, information from adjacent cells is combined to form a so-called cluster using a specific clusterization algorithm. Note that in this procedure it is possible that multiple particles could contribute energy to a single cluster, unlike tracking where one track unambiguously originates from one particle. The raw signal amplitude for each cell is digitized into ADC counts, where one count corresponds to 16 MeV on average. This average correspondence is determined via a test beam and the exact per-cell correspondence is handled in the energy calibration procedure (see Chapter 4). The arrival time of the signal is also recorded, which has an average value of 600 ns due to the length of the cables. The readout time of approximately $1\mu s$ is much greater than the space between bunch crossings (25 ns) and therefore multiple collisions will occur in the same readout window. Therefore, there must be a correction for this delay in arrival time, which is handled via the time calibration procedure (see Chapter 4).




Figure 3.4: The ALICE EMCal (top) and DCal (bottom), originally appearing in [26]. The PHOS calorimeter is shown in tan near the DCal.

3.4 ALICE in Run 3

Run 3 of the LHC began in 2022, resulting in pp collisions at a record center-of-mass energy of $\sqrt{s} = 13.6$ TeV.⁵ During the long shutdown following Run 2 of the LHC, ALICE

⁵<https://home.cern/news/news/physics/lhc-run-3-physics-record-energy-starts-tomorrow>

underwent major upgrades that may allow a 50-fold increase in the number of recorded collisions. These upgrades consist of an upgrade to the **TPC**, **ITS**, as well as a new **Muon Forward Tracker (MFT)** and **Fast Interaction Trigger (FIT)**.⁶ To accommodate the higher volumes of data, a new **Online-Offline (O^2)** software framework⁷ is used for reconstruction.

The most significant of these upgrades is the upgrades to the **TPC**, which result in a new continuous readout system capable of handling 50 kHz interaction rates for Pb–Pb runs. In order to perform this continuous operation, the **TPC** multi-proportional wire chambers were replaced with a **Gas Electron Multiplier (GEM)** system that has a small anode-cathode distance with a large electric field, producing large gas amplification over a small distance and short time. Additionally, in order to handle the increased outgoing data rate (up to 3 TBs/s) new **GPU**-based data reduction and specialized readout hardware was installed. Another necessary upgrade for the implementation for continuous readout is an upgrade to the **ITS** for faster readout along with more precise resolution with less material.

These large increases in the outgoing data rate poses a significant challenge for both the online and offline computing systems for **ALICE**. O^2 was developed with the aim of maximal reduction of the data as early as possible during the data flow. Towards this aim, the O^2 data flow begins its reconstruction synchronously with the data-taking. This is referred to as the synchronous reconstruction, where a full **TPC** reconstruction using online cluster finding as well as a partial reconstruction of other detectors will take place, resulting in an overall reduction of the data volume as raw detector information is no longer needed. An additional reconstruction is performed asynchronously⁸ with data-taking where the final calibrations will be applied and then data of the desired quality will be stored for further analysis. For an in-depth description of the calibration of the **EMCal** in Run 3, refer to Chapter 4. Additionally, the data structure will be upgraded from object-oriented containers to lighter-weight inter-connected tables. This new paradigm will allow for bulk operations to be performed easily on the table rows and columns instead of time-consuming loops over the containers, resulting in a greater efficiency designed to handle the greater volume.

⁶These upgrades are so significant that **ALICE** is now referred to as **ALICE 2**, signifying a completely new detector.

⁷<https://github.com/AliceO2Group/AliceO2/>

⁸This is referred to as the asynchronous reconstruction.

Chapter 4

Online Calibration of the Electromagnetic Calorimeter

This chapter describes work completed over the course of this thesis in order to perform the bad-channel calibration in Run 3 of the [LHC](#). The work was done as a part of (and in collaboration with) the [EMCal](#) Collaboration.

As described in Chapter 3, the [Electromagnetic Calorimeter \(EMCal\)](#) is used in order to perform measurements of photons, neutral mesons, dileptons and the neutral component for full jet measurements. However, before the data can be used for analysis, it must be properly calibrated. There are several different calibrations that are needed for the [EMCal](#) [26] to ensure that the measured signal best reflects the true physical conditions, which are listed and briefly described below.

- **Energy Calibration:** The energy calibration seeks to determine a coefficient for each cell that, when multiplied by the cell response gives the absolute cell energy. The correction factors for each cell are obtained in an iterative procedure aiming to select the correction factors which best constrain the invariant mass of the π^0 .
- **Time Calibration:** The [EMCal](#) is a triggered detector and the time is measured relative to the trigger time, where ideally the signal time in the [EMCal](#) is identical to the trigger time. However, due to effects such as cable length, response times of the electronics, and any time shifts due to clock phase differences, the arrival time may vary cell by cell. This correction is done by subtracting the average cell time over a period of data taking, which should represent the trigger peak, from the measured cell time.
- **Bad-Channel Calibration:** The bad-channel calibration seeks to remove cells (also called channels) that give an improper response to a hit, are noisy, or for some

reason have a discontinuous energy spectrum. This will be discussed in detail in this chapter.

- **Temperature Calibration:** The gain of each [avalanche photo diode \(APD\)](#) has a linear dependence on the temperature. The temperature calibration seeks to account for this effect, as the data were taken in varying conditions.

Data-taking in [ALICE](#) looks very different in Run 3 than in previous runs of the [LHC](#), as described in Chapter 3. In order to handle the volume of data taken at a higher rate, the implementation of the calibrations must also change as compared to the process in Run 2. This chapter will first describe the bad-channel masking procedure in Run 2 (see Section 4.1), and then the bad-channel masking procedure in Run 3 (see Section 4.2) will also be described.

4.1 Bad-Channel Calibration in Run 2

The principle task in the bad-channel calibration procedure is to identify cells as either good, bad, or dead. Cells that are labeled as bad or dead are referred to as *masked* and are not considered for further analysis. Cells that are labeled as dead refer to cells with zero recorded hits. This typically can occur when front-end electronics cards are switched off for replacement. Cells are identified as good or bad using the mean energy per cell and the number of hits per cell. To account for the fact that cells behind the [TRD](#) support structures have a lower-than-average number of hits scale factors derived from detector simulations are applied in an iterative procedure for the mean number of hits per row/column. After this scaling is applied the mean energy per cell and the mean number of hits per cell are calculated for the given run, where cells that are outside of $\mu \pm 5\sigma$ are labeled as bad. As the time calibration was completed after the bad-channel calibration, cells that fired outside the normal time window with high frequency were also labeled as bad.

Once this procedure was completed for a given period containing many runs, these runs were then divided into runblocks each with a different list of bad channels. This was done in order to ensure enough statistics to reliably extract the mean energy for the calibration. In order to determine these runblocks, only cells that were labeled as bad in over 20% of runs were considered and then sets of runs with similar patterns of bad channels were grouped together into a runblock. After the determination of runblocks, the procedure of calculating the mean energy per cell and the mean number of hits is repeated, this time for each runblock.

Finally, once all other calibrations described above are completed, the bad-channel masking procedure as described above is repeated one final time in order to check for additional bad channels. This procedure is sometimes referred to as the Analysis-Level

QA. The masking of cells is then completed on a run-by-run basis offline for each runblock as determined by the above procedure.

4.2 Bad-Channel Calibration in Run 3

4.2.1 Differences between Run 2 and Run 3 Approach

Due to the increased data volume and data-taking rate, the procedure utilized in Run 2 is no longer possible in Run 3 due to general changes to the data-taking procedure, resulting in a few fundamental changes in the bad channel masking procedure. Firstly, the data-taking in Run 3 will be done in time-frames instead of runs due to the newly-adopted continuous readout of the TPC. Therefore, the calibration must also be performed in time-frames instead of runs or runblocks. Secondly, in Run 3 there will be much higher data volumes than existed previously. Due to this fact, the determination of bad channels must be performed online during the synchronous reconstruction phase. The bad-channel masking will then take place soon after in the asynchronous reconstruction phase. Therefore, it is critical that the calibration in Run 3 be fast and memory efficient while also being accurate as once a channel is masked, the information cannot be recovered without an additional asynchronous reconstruction pass.

There are many different possible approaches to perform the bad-channel calibration in Run 3 including approaches utilizing ML or other new techniques. However, the first approach to this calibration is done in a similar spirit to the calibration in Run 2. This was done so that results can be more easily compared and this new algorithm could be implemented in time for the beginning of Run 3 data taking. Possible methods using ML are left as future work.

4.2.2 Online Bad-Channel Calibration

The algorithm utilized in Run 3 for the identification of bad channels online during the synchronous reconstruction pass will be conceptually outlined in this section.¹ The principle data structure utilized in this procedure is Boost Histograms² due to their lightweight nature. This allows for information to be passed smoothly during the data-taking period.

Firstly, after the accumulation of a reasonable amount of statistics a boost histogram was filled with the cell ID (an identifier associated with each cell in the EMCal) and the cell energy. The sufficient amount of statistics can change for each data-taking configuration,

¹For the code and more detailed information than can be provided in this context, please see: <https://github.com/AliceO2Group/AliceO2/tree/dev/Detectors/EMCAL/calibration>

²See documentation for Boost Histogram: [here](#)

but was initially tested and ran with 10 million events ³ as the required number of events for sufficient statistics. Additionally, this boost histogram was filled with scaling factors to account for the reduced number of hits in cells behind TRD support structures.

Following this step the distribution was divided into energy slices, where a cell will be marked as bad if it is marked as bad in any of the chosen energy slices. This procedure was initially tested and ran with the following energy slices: 100-300 MeV, 300-500 MeV, 500 MeV-1 GeV, and 1-4 GeV. Then for each energy slice the average energy and the average number of hits was calculated. In this case, the average values were calculated as truncated mean (and sigma) [178] in order to mitigate the impact of bad channels that have not yet been removed have on the calculation. Then, a good cell window is calculated for the mean energy per cell and the mean number of hits per cell with the window calculated as $\mu \pm n_\sigma \sigma$ where n_σ is a configurable parameter, but chosen as 5 for the purposes of initial testing and application.

Then each cell is marked as good or bad where, as before, cells are additionally marked as dead if they contain zero hits. For the remaining cells, the average cell energy and number of hits is then calculated and compared to the good cell window. A cell is labeled as bad if its average energy or number of hits falls outside of the outside of the good cell window for any of the considered energy ranges. Cells that are not flagged as bad or dead are then considered as good. If timing information is available, cells can additionally be marked as bad if the cell's average time is outside of a window given by $\mu \pm 5\sigma$ again calculated using the truncated mean. The set of masked cells will then be applied in the asynchronous reconstruction step removing masked cell signals from the data set. A sample status of various cells following the bad-channel calibration in the Run 2 and Run 3 schemes is shown in Figure 4.1. The energy distribution before and after the masking of bad channels is shown in the right panel Figure 4.2, where any prominent bumps/outliers in the energy distribution are removed by the calibration of bad channels.

4.2.3 Validation of the Run 3 Calibration

Due to the fact that the calibration in Run 3 is performed entirely online and only calibrated information is stored following the asynchronous reconstruction, it is very important that this procedure is validated and functioning as intended. This is done using data from Run 2 for which the bad channels are known. This data is then converted into the format utilized for Run 3 and then the above-mentioned procedure is applied. The two bad-channel maps for each of these different procedures can be found in Figure 4.1. By eye these distributions are very comparable, though it is difficult to compare on an ensemble basis by eye.

In addition, a visual representation of the direct comparison between these distributions is made as shown in Figure 4.2. This figure demonstrates the bad-channel calibrations in Run

³This number was determined based on Run 2 studies, where this number of events is reached much faster in Run 3 due to the larger readout rate.

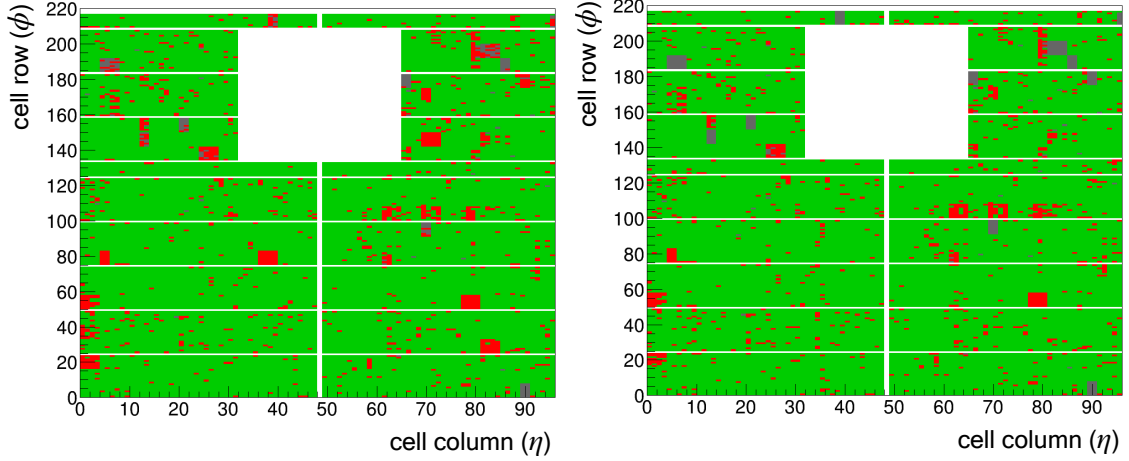


Figure 4.1: Bad-channel map as calculated from the same input data using the Run 2 procedures (left) vs. the Run 3 procedure (right). Cells filled green indicate those labeled as good, those filled gray indicate those labeled as bad and cells filled red indicate those marked as bad.

2 and Run 3 yield comparable results. In addition, the energy distribution after masking in the Run 2 and Run 3 schemes is compared in the right panel of Figure 4.2. Here, it can be seen that the energy distributions are roughly similar, indicating consistent results between the two schemes. As a result of these tests, the new calibration was validated and employed online for Run 3 pp data-taking in 2023.

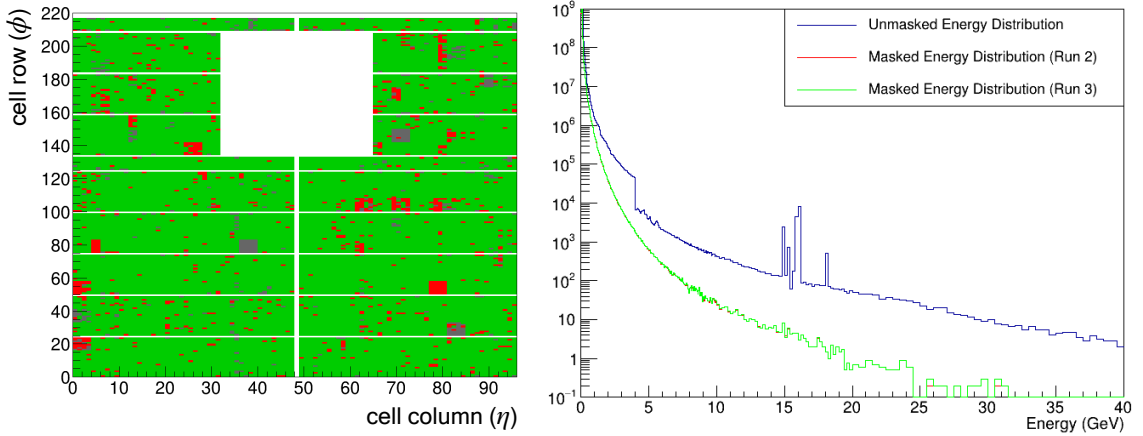


Figure 4.2: Left: A comparison between the Run 2 and the Run 3 calibrations, cells marked in green are cells marked as good in both calibrations, cells marked in red are cells marked as bad in both calibrations, and cells marked in gray are flagged differently in each calibration. Right: A comparison of the energy distribution before the masking of bad channels (blue), after masking in the Run 2 calibration (red) and after masking in the Run 3 calibration (green).

Chapter 5

R -dependence of the R_{AA}

The work reported in this section is based on work performed as a part of the ALICE collaboration with charged-particle jets [122, 179]. Measurements in many parts of phase space will be presented in this chapter due to the differential nature of this work, but an emphasis will be placed on results that yield new physics, such as the measurements at large R and the corresponding R_{AA} ratios.

5.1 Motivation and Previous Measurements

As mentioned in Chapter 2, the R_{AA} is a key experimental signature of jet quenching. There have been many measurements of the jet R_{AA} , both at RHIC [27] and the LHC [14, 28, 29, 180–182], a selection of which are shown in Figure 5.1. All measurements shown in Figure 5.1 report an R_{AA} for $R = 0.4$ jets significantly below unity consistent with model predictions, indicative of jet quenching. Though these measurements span two orders of magnitude, there still remains a small kinematic gap between measurements at RHIC and the LHC that this thesis will bridge.

Of recent interest is the dependence of jet suppression effects on the jet radius, R , which is shown to help disentangle the many competing aspects of energy loss [183].¹ This is a direct consequence of the idea that the nature and extent of energy loss is expected to vary across different R and p_T scales. For example, energy that moves to larger angles (as shown in Ref. [109] and the left panel of Figure 2.23) is expected to be recovered as R increases, causing jets at larger R to exhibit less suppression. The medium-induced wake that adds soft particles to the jet cone could cause a similar trend in jet suppression, as the contribution of the wake grows with R . However, large R jets may have more effective energy loss sources, as indicated by their more complex substructure, which would result in larger R jets exhibiting more suppression. Similar effects could also be caused by larger R jets having a higher gluon-to-quark ratio at a fixed p_T due to the expectation that

¹These energy loss sources are discussed in detail in Section 2.5.




Figure 5.1: Previous measurements of the jet R_{AA} for $R = 0.4$ 0–10% jets from RHIC [27] and the LHC [14, 28, 29].

gluon jets will lose more energy due to their color factor. Additionally, the jet transverse momentum spectrum in vacuum also exhibits an R -dependence [14, 18], which may also propagate to the R_{AA} and introduce an R -dependence.

Each of the effects mentioned above may exist to varying extents in different kinematic ranges, rapidity intervals, centrality intervals etc as the jet population changes with each of these factors. Therefore, differential measurements of the R_{AA} with R are useful for determining which effects may be dominant in different regions of phase space. The R -dependence of the R_{AA} has been measured recently both at RHIC [27] and the LHC [28, 184].² The ATLAS collaboration was the first to measure this dependence via the ratio of the inclusive jet p_T spectra in central and peripheral collisions, the R_{cp} , which indicated that jets of intermediate p_T for the LHC at large- R ($R = 0.5$) jets exhibit less suppression than small- R jets [184]. The CMS collaboration recently measured the R -dependence of very-large- R jets ($R = 1.0$) relative to small R jets at very high p_T and observed no R -dependence, indicating a cancellation of effects. The STAR collaboration has also measured the R -dependence for jets with $R = 0.4$ compared to jets of $R = 0.2$. This measurement shows a significant R -dependence at very low p_T with larger R jets being less suppressed when compared to a pp reference with the no track p_T requirement on the leading track of the jet.³ However, the jet population at RHIC differs from that at the LHC, where jets at RHIC contain more quark-initiated jets, which prevents a direct

²For a comparison of these results to the results shown in this thesis, refer to Appendix A.

³Note that such conclusions would change if the pp reference is applied consistently in the Au+Au and pp jet p_T spectra. This can be seen from taking the ratios between R of Figure 9.8 in <https://drupal.star.bnl.gov/STAR/files/phd.thesis-rusnak.pdf>.

comparison between measurements at the two facilities.

There are a number of theoretical predictions for the R -dependence of the R_{AA} . Some models predict a rise in the R_{AA} with increasing R [183, 185, 186]. Alternatively, other models predict that the R_{AA} is roughly constant with varying R [187], reflecting a cancellation of competing effects. Finally, some models show a decrease in the R_{AA} with increasing R [34, 140, 183]. Therefore, the R -dependence of the R_{AA} is a discriminative observable that, when compared to model predictions containing different physical processes, can better indicate the relative abundance of different energy loss mechanisms than the R_{AA} alone.

The theoretical and experimental considerations mentioned above motivate an experimental measurement of the R -dependence of the R_{AA} at lower jet p_T at the LHC. ALICE is optimally situated to make measurements in this regime due to its precise tracking down to lower p_T in high-multiplicity environments. However, new techniques need to be developed in order to make measurements over large ranges in R and p_T to accommodate the difficulty in correcting for the large fluctuating underlying event in heavy-ion collisions. The following sections outline the procedure for performing this measurement and its results.

5.2 Data Selection

Pb–Pb collision data from an ALICE data-taking period in 2015⁴ were utilized in this analysis. Minimum bias⁵ events were selected where the absolute value of the z_{vtx} is less than 10 cm and the vertex has at least one contributor. For collisions within the 0–10%, 30–50%, and 60–80% centrality ranges, about 5, 10, and 10 million events, respectively, were accepted for further analysis following these selections.

5.2.1 Track Selection

Charged-particle jets (also called track-based jets) cluster charged tracks into a jet via a jet-finding algorithm. Since jet reconstruction can be sensitive to holes or lowered efficiencies in the acceptance, tracks with a uniform efficiency across the full acceptance should be used. Here, so-called hybrid tracks, consisting of global tracks that have at least one hit in the SPD combined with tracks that include the primary-vertex, are used. As seen in Figure 5.2, these hybrid tracks are uniformly distributed in azimuth (ϕ).

⁴Internally referred to as LHC15o.

⁵As the name suggests, the term minimum bias refers to events selected with a minimal preference to physics at play.




Figure 5.2: Azimuthal distribution for selected hybrid tracks before the fiducial acceptance cut.

5.2.2 Embedding

One additional sample utilized in this analysis is the sample that combines a known signal from simulation with real Pb–Pb minimum bias events in a process referred to as embedding. In order to ensure a high-statistics sample a $p_{T,\text{hard}}$ -binned MC production (PYTHIA 8, Monash 2013 Tune [83]) with a full GEANT3 [188] detector simulation was used.⁶ Each of the 20 $p_{T,\text{hard}}$ bins is populated with approximately 1,000,000 events with the corresponding edges of [5, 7, 9, 12, 16, 21, 28, 36, 45, 57, 70, 85, 99, 115, 132, 150, 169, 190, 212, 235, 1000] GeV/ c . The MC production is anchored to the runlist from the Pb–Pb data set as described in Section 5.2.

Once the signal events are generated and the detector simulation has been run they are then combined at the detector level with the Pb–Pb minimum bias data background to form a hybrid event. To account for the worsened tracking efficiency in Pb–Pb as compared to pp collisions an additional 2% of tracks in PYTHIA are randomly rejected when forming the hybrid event.⁷ As tracks contain an unambiguous mapping to a single particle, tracks in the hybrid event are simply the combination of tracks in the signal and background samples.

In this analysis the embedded sample is utilized to evaluate the impact of the background in a context where the true signal is known. This occurs both when constructing the response matrix (discussed in Section 5.7) and also when creating the sample for training the ML estimator (discussed in Section 5.5).

⁶Internally referred to as LHC16j5, see: <https://alice.its.cern.ch/jira/browse/ALIROOT-6905>

⁷Note that this does not indicate a tracking efficiency of 98% as the 2% rejection is in addition to the existing tracking efficiency in pp.

5.3 Jet Reconstruction

Charged-particle jets are reconstructed using the anti- k_T algorithm with E -scheme recombination [12] in the FastJet package [189].⁸ Charged-particle jets utilize charged tracks with $p_{T,\text{track}} > 0.15 \text{ GeV}/c$ as constituents. Jets are required to be contained within $|\eta_{\text{jet}}| < 0.9 - R$, the fiducial acceptance of the TPC. A jet area cut of $A_{\text{jet}} > 0.56\pi R^2$ is applied to suppress contamination by non-physical jets [190]. Jets containing a track $p_T > 100 \text{ GeV}/c$ are additionally removed due to reduced tracking performance in this region. For the background jets used in the calculation of ρ (see Section 5.4) in this analysis the k_T algorithm with the E -scheme recombination scheme in the FastJet package [189] is utilized. For the ρ calculation, a smaller jet radius of $R = 0.2$ is used regardless of the radius of the desired signal jet to be less sensitive to fluctuations, though this is shown to have a small effect.

In the embedded sample, three types of signal jets are reconstructed; particle(-level), detector(-level), and hybrid(-level) jets. All three types form jets using the same algorithmic definition and the same fiducial and constituent cuts as the signal-jets described above. Particle-level jets are reconstructed using charged particles in the PYTHIA [83] event as constituents. Detector-level jets are reconstructed using charged tracks in the PYTHIA event propagated through a full detector simulation using GEANT3 [188]. Hybrid-level jets use the charged tracks in the hybrid event as constituents. Jets used for the calculation of ρ in the case of the embedded sample have the same background-jet definition as described above while using charged tracks in the hybrid event as constituents.

5.4 Background Subtraction

One of the largest obstacles to reconstructing the jet p_T in heavy-ion collisions is the large fluctuating background due to the UE. For example, the fluctuations in the charged particle momentum density per event in central (0–10%) Pb–Pb collisions at the LHC are $\approx 18 \text{ GeV}/c$ [191]. Upward fluctuations of the UE can themselves be reconstructed as jets, commonly called fakes. Fakes contaminate the signal at low jet p_T . On average this background is proportional to the jet area, indicating that the influence of this background will be largest for jet measurements at larger R and lower jet p_T .

Measurements of the R -dependence of the R_{AA} rely on the ability to measure over a broad range in R and p_T . In order to perform this measurement, the background must be accounted for and removed. Such a correction procedure must not only correct for the fluctuating background that smears the jet p_T , but it must also remove the signal-contaminating contributions of fakes. One common treatment of this background (herein referred to as the area-based or AB method) is to subtract off the event-averaged momentum density (ex-

⁸See Section 2.4.2.

cluding the two leading jets) multiplied by the jet area from the uncorrected jet p_T [191]. Mathematically, this is given in Equation 5.1 where $p_{T,\text{corr}}$ and $p_{T,\text{raw}}$ are the corrected and raw jet transverse momenta, ρ is the event-averaged median background density, and A_{jet} is the area of the jet. Note that before the calculation of ρ the two jets with the highest p_T in the event are removed to reduce the contamination from the signal in the background calculation.

$$p_{T,\text{corr}} = p_{T,\text{raw}} - \rho A_{\text{jet}} \quad (5.1)$$

Typically, a leading track cut is applied to remove the fake jet contribution. This is done at the cost of an introduced bias on the fragmentation of jets selected. While the AB method effectively corrects for the average background in its pedestal-subtraction, it does not account for region-to-region fluctuations. These residual fluctuations are commonly handled in an unfolding procedure, which will be discussed in Section 5.7.

The performance of the background subtraction procedure is typically evaluated using jets reconstructed in the hybrid events resulting from the embedding procedure described in Section 5.2.2. The background procedure is then applied to reconstructed jets in the hybrid event. The closer the corrected jet p_T is to the p_T of the true jet from simulation, the better the performance of the background subtraction procedure. This performance is quantified via the width of the δp_T distributions as defined in Equation 5.2. Here, $p_{T,\text{rec}}$ is given as the reconstructed jet p_T in the hybrid event after the background subtraction procedure is applied and $p_{T,\text{true}}$ is the true jet p_T from simulation.

$$\delta p_T = p_{T,\text{rec}} - p_{T,\text{true}} \quad (5.2)$$

The performance of the background subtraction procedure additionally dictates the kinematic reach of the measurement as the input p_T range to the unfolding procedure (see Section 5.7) is dictated by $\pm 5\sigma$ of the δp_T distribution. Therefore, the better the performance of the background subtraction, the larger the p_T range given as input to the unfolding procedure - resulting in a larger kinematic reach of the measurement.

Previously, ALICE has made a measurement of the R_{AA} using the AB subtraction for $R = 0.4$ jets down to 60 GeV/ c in central (0–10%) collisions at $\sqrt{s_{NN}} = 5.02$ TeV [14]. However, the performance of the AB technique limits the extension of this result to larger R and lower in p_T . For an R -dependent measurement of the R_{AA} covering a wide range in R and p_T , a new ML-based background correction is employed, which is discussed in Section 5.5. The AB method will be used here for the purpose of comparison with the new ML-based approach, and also as the background subtraction technique used for the R_{AA} in peripheral (60–80%) collisions. For the δp_T distributions for the AB curves compared to the ML-based approach, see Figure 5.3.⁹ For the δp_T distributions as a function of p_T , refer to Appendix C.

⁹This figure will be discussed in more depth in the following section.




Figure 5.3: Residual p_T -distributions of embedded jet probes of known transverse momentum into Pb–Pb collision data. Left: Comparison of the distributions for the area-based and ML-based background estimators. Note the lines connecting the points do not represent a fit and are only present to guide the eye. Right: Radius dependence of the width of the distributions, where the error bars come from the uncertainty in calculating the width.

5.5 ML-Based Background Subtraction

5.5.1 Introduction

In both of the charged and full jet analyses, a novel ML-based background subtraction was employed [192]. This regression task¹⁰ seeks to create a mapping between the measured and corrected jet p_T using information about the individual jet’s properties and the properties of its constituents. Machine learning is a natural solution for this problem as the ML can utilize the complexity of the jet and the arrangement of its constituents in order to derive a relation that is difficult to form with expert knowledge alone. Not only does this mapping correct the jet p_T for background effects, but it also automatically corrects for fake jets by correcting their p_T to zero. Such an approach yields a considerable improvement in performance, which will be discussed in detail in Section 5.6. In this section, the details of this ML-based approach to background subtraction will be discussed. Differences between the full and charged jet approach will be discussed in Chapter 6.

¹⁰Refer to Section 2.6 for more information.

5.5.2 Training Data Set

Creating a realistic event for the ML algorithm to train on is a critical step in this procedure. This ML-based estimator represents a supervised learning problem, where the algorithm learns the desired relations from a labeled set of true values. Therefore, the training data must not only have a well-defined truth value, but it must also have a realistic background similar to the conditions seen in data. To do this one can embed a known signal into a realistic Pb–Pb background using the process described in Section 5.2.2. For the signal a PYTHIA sample as described in Section 5.2 is used. For a description of the regression target used, refer to Section 5.5.3.

For the charged jet results a thermal toy model was utilized. The thermal background is created by randomly distributing charged particles according to a flat particle multiplicity distribution ranging from 0 to 3000 tracks and with a realistic (quasi-thermal) momentum distribution, as given in Figure 5.4. A flat multiplicity distribution contains both sparse and dense events with equal weights and the maximum track count of 3000 corresponds to the most central events. The momentum distribution is drawn from a modified power-law function that goes as p_T^{-6} (Tsallis fit to data) and defined such that it coincides with the track momentum distribution at low p_T , but falls much faster for higher particle momenta roughly above 4 GeV/c. Using a toy model for the background has the advantage that it is easy to modify the background to test the model dependence on certain features, at the potential consequence of the model not including subtleties present in an actual heavy-ion background.




Figure 5.4: Transverse momentum distribution for charged tracks from the toy model.

Another potential choice for the background in the training set, which was employed in the full jet analysis, is to use Pb–Pb minimum bias data. This method has the advantage that it is a realistic background by construction. However, one potential downside of this approach is that the Pb–Pb minimum bias data will also include the production of jets, meaning that the background will contain signal-like objects. Jets are rare-enough

probes that this is thought to not have a significant effect. A comparison between the track p_T distribution within jets in the toy and the embedded MC is shown in the left panel of Figure 5.5. These distributions are roughly comparable, but show some deviations at lower p_T . Despite these differences, results achieved using these two different choices for the background of the training data set are comparable (see Figure A.4).

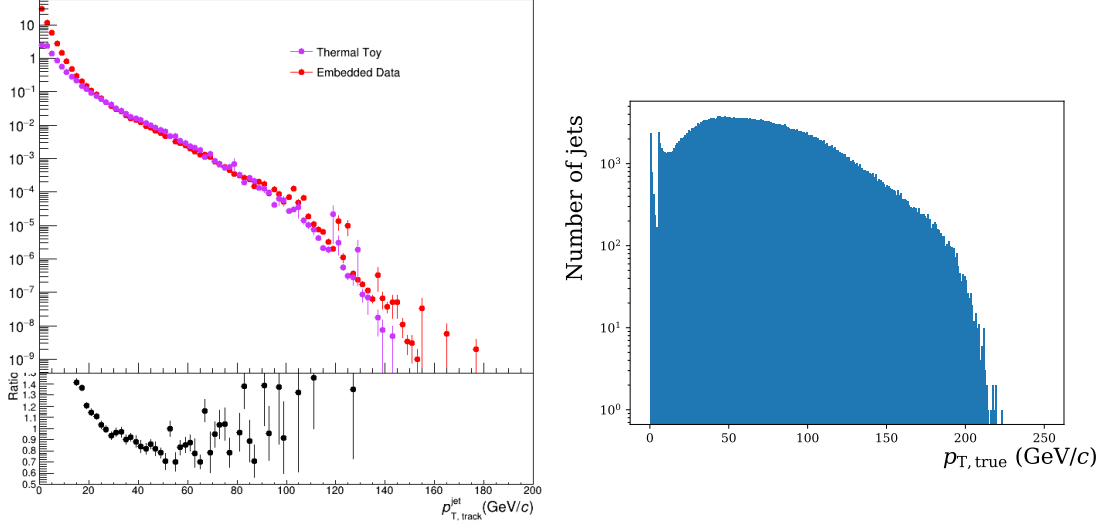


Figure 5.5: (Left): Comparison between the track p_T distribution within jets in the toy model and the embedded MC for $R = 0.4$ charged jets. (Right): Distribution of true p_T 's for the training sample for $R = 0.6$ jets as seen by the ML.

Note that the PYTHIA events were generated in $p_{T,\text{Hard}}$ bins for the simulated jet samples used in training. These $p_{T,\text{Hard}}$ bins are then scaled down to create realistic spectra. The ML algorithm is trained on the $p_{T,\text{Hard}}$ -binned sample before any reweighting correction is applied. As a result, a higher emphasis is placed on higher p_T contributions than the natural abundances. This is done intentionally to ensure that the estimator does not over-learn features of low p_T jets due to their higher abundance. An additional precaution to avoid over-fitting is that 95% of jets below 5 GeV/c are discarded for training. For the true p_T distribution of jets used for training, i.e. unscaled by any reweighting correction, refer to the right panel of Figure 5.5.

5.5.3 Regression Target

For the training, the supervised learning techniques that are applied need a truth value assigned to each sample (jet). There are two possible choices for this truth value, both of which yield similar results. The first possible truth definition involves matching the hybrid-level jet to a PYTHIA jet at the detector level using a geometric matching procedure, described in Section 5.7. The principal advantage of this method is that it has a

clear physical meaning. However, this choice of regression target adds an additional parameter of the geometric matching radius into the model, which can cause it to be less robust. The second possible definition is referred to as the true p_T fraction. This is defined in Equation 5.3, where $p_{T,\text{jet}}^{\text{raw}}$ represents the reconstructed full jet transverse momentum before any background subtraction is applied, $\sum_i p_{T,\text{const } i}^{\text{PYTHIA}}$ represents the sum of the p_T of the PYTHIA constituents and $\sum_i p_{T,\text{const } i}$ is the sum of the p_T of the jet's constituent.

$$p_{T,\text{jet}}^{\text{true}} = p_{T,\text{jet}}^{\text{raw}} \times \frac{\sum_i p_{T,\text{const } i}^{\text{PYTHIA}}}{\sum_i p_{T,\text{const } i}}. \quad (5.3)$$

5.5.4 Input Parameters

When using ML techniques, as described in Section 2.6.1, it is essential to minimize the number of parameters (also called features) from which the model will learn. This is done to make the model robust as possible and therefore less sensitive to subtleties in the training data set. The possible input parameters fall into four different categories: event properties (such as the centrality or the background density), jet properties (such as the AB corrected jet p_T , uncorrected jet p_T , or various jet shape ¹¹ parameters), or constituent properties (such as the p_T of the track or the number of tracks in a jet).

The minimal parameter list was chosen by iteratively removing unimportant or highly correlated features while ensuring that the performance was not significantly reduced. The degree of correlation between input features can be quantified by the Pearson correlations (Figure 5.6) or by visualizing them directly (Figure A.1). For example, from Figure 5.6 it can be seen that the number of tracks in a jet is highly correlated with the p_T of tracks with the ~ 10 th highest p_T and above. Therefore, softer tracks can be removed safely from the input parameter list, as the same information is captured in the number of tracks. The feature importances of the random forest model described in Section 5.5.5 were used in order to evaluate the importance of different parameters. This is mathematically defined as the mean decreased impurity or the number of times a feature is used to split a node weighted by the number of samples it splits. For example, the corrected and uncorrected jet p_T are highly correlated and the corrected jet p_T has a higher feature importance. Therefore, the uncorrected jet p_T was removed as an input parameter. The feature importances from the first iteration can be found in Table 5.1. Based on these considerations, the following input features were selected: the jet p_T , the first radial moment of constituent momenta (jet angularity), the number of constituents within the jet, and the p_T of the eight leading (highest p_T) particles within the jet. ¹²

¹¹For a definition of jet shapes, see Section 2.5.3.

¹²If certain input parameters are not available, e.g. if a jet does not have enough constituents, zero is used as the input value. This procedure is called zero-padding.




Figure 5.6: The Pearson correlation of input parameters for central (0-10%) $R = 0.6$ jets. Note that the names of variables here correspond to how they are written in the analysis code. See Appendix B for more details.

5.5.5 ML Methods

In this analysis, the ML algorithm used was a shallow, multi-layered NN with three layers with [100,100,50] nodes implemented in scikit-learn [145]. The ReLU activation function was used with an ADAM optimizer for the stochastic gradient descent algorithm. All unlisted parameters take the default values in scikit-learn. Scikit-learn version 0.20.3 was utilized in this analysis. Though using deep learning techniques would be possible, shallow methods were chosen for transparency and simplicity.

Table 5.1: Features utilized in the random forest model trained on $R = 0.4$ jets in central (0–10%) collisions as well as their feature importances, gathered using `scikit-learn`.

Feature	Importance	Feature	Importance
Jet p_T (area-based corr)	0.2611	$p_{T,\text{constituent}}^6$	0.0390
Number of tracks	0.0026	$p_{T,\text{constituent}}^7$	0.0502
Number of clusters	0.0017	$p_{T,\text{constituent}}^8$	0.0096
Mean of track p_T 's	0.1808	$p_{T,\text{constituent}}^9$	0.0072
Jet Angularity	0.0308	$p_{T,\text{constituent}}^{10}$	0.0053
$p_{T,\text{constituent}}^1$	0.0356	$p_{T,\text{constituent}}^{11}$	0.0051
$p_{T,\text{constituent}}^2$	0.0480	$p_{T,\text{constituent}}^{12}$	0.0047
$p_{T,\text{constituent}}^3$	0.1312	$p_{T,\text{constituent}}^{13}$	0.0032
$p_{T,\text{constituent}}^4$	0.0870	$p_{T,\text{constituent}}^{14}$	0.0039
$p_{T,\text{constituent}}^5$	0.0897	$p_{T,\text{constituent}}^{15}$	0.0031

5.5.6 Model Investigations

In this section, some investigations of the ML-based correction are investigated. Note, that this is not an exhaustive list, and the remainder of such checks are listed in [Appendix A](#).

Looking into the Decision Tree

Many ML algorithms allow the user to access information about how the algorithm is learning the features in the training set. One example of this is the feature importances mentioned in [Section 5.5.4](#). Another example is looking into the actual decision tree itself. This provides an unambiguous look into how the algorithm is making its decisions. [Figure 5.7](#) provides a visual representation into one of the decision trees included in the random forest estimator. Though the final estimator applied to data will be an average of all of the individual decision trees, this is a useful indicator of the types of decisions that these trees make. Notably, constituents have a high importance for low p_T jets whereas at higher jet p_T the AB corrected jet p_T has a higher importance. This confirms the initial assumption that the ML is utilizing constituent information in regions where it is normally difficult to distinguish signal from the background, i.e. at low p_T .

Impact of Flow on the ML Estimator

One type of fluctuation in the heavy-ion background is correlated fluctuations due to flow-like effects.¹³ These flow-like effects generate a momentum-space anisotropy of particles, oriented with respect to the event plane. The magnitude of this correlated background grows with the radius of the jet due to the increased space for fluctuations. This magnitude is also expected to depend on the angle between the jet axis and the event plane, given by [Equation 5.4](#) where ψ_{EP} represents the event-plane orientation. Jets classified as in-

¹³The other type are purely-statistical Poissonian fluctuations. For more on flow, refer to [Section 2.3.4](#).

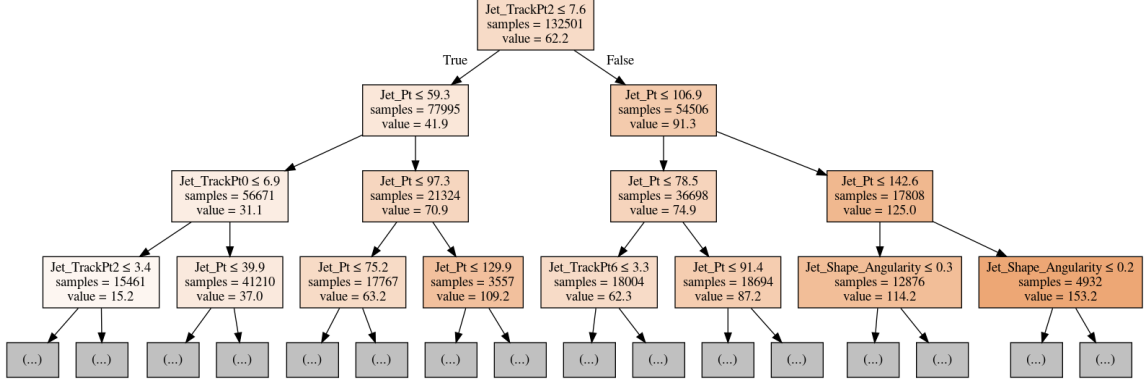


Figure 5.7: Visualization of a decision tree for the random forest estimator trained on $R = 0.6$ jets in central (0-10%) collisions. For visualization purposes, only the first few layers of the decision tree are included.

plane $\Delta\phi < 30^\circ$ would have a higher degree of correlated fluctuations as compared to jets out-of-plane ($\Delta\phi > 60^\circ$), where

$$\Delta\phi = \phi_{\text{jet}} - \psi_{\text{EP}}. \quad (5.4)$$

The standard [AB](#) method does not remove flow-like effects, which causes a dependence of the width of the δp_T distributions on $\Delta\phi$. This can be seen in the dashed lines in Figure 5.8, where the mean of the δp_T distributions is larger for the in-plane case. Such an observation is consistent with the background subtraction not fully-correcting for these correlated fluctuations. This test for the ML-based method is also shown in Figure 5.8, where the performance of the ML-based method exhibits very minimal event-plane dependence. This indicates that although the estimator was trained on a thermal background that does not include correlated fluctuations, it can properly correct for these effects. Such an observation also indicates that the two possible backgrounds used for training yield similar results.

5.5.7 Performance

The performance of the ML-based background subtraction method is measured with the same metric as any background subtraction method, the δp_T distributions as defined in Equation 5.2. The δp_T distributions for the ML-based background subtraction compared to that of the [AB](#) subtraction for $R = 0.4$ jets in central collisions is shown in the left panel of Figure 5.3. The width of these distributions, which roughly corresponds to the amount of remaining residual fluctuations after background subtraction, is significantly reduced in the ML-based method as compared to the [AB](#) method. This indicates an improved performance. The right panel of Figure 5.3 shows the standard deviation of the δp_T distributions as a function of R for the ML-based and [AB](#) methods for central and semi-central colli-

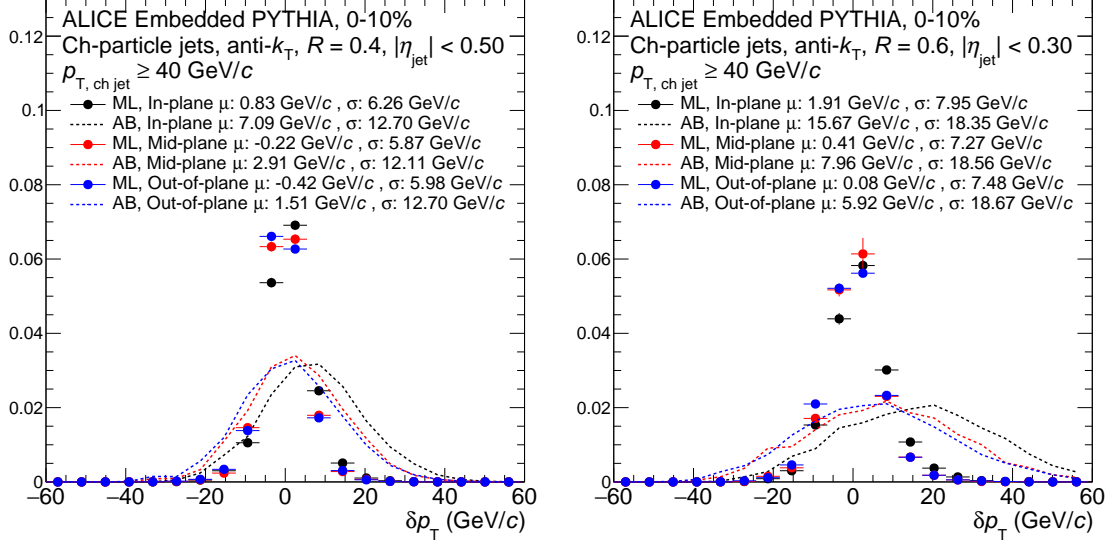


Figure 5.8: The δp_T distributions for $R = 0.4$ (left) and $R = 0.6$ (right) jets in central (0-10%) collisions using the area-based (dashed lines) and ML-based (solid markers) corrections.

sions. The ML-based method outperforms the AB method in all explored regions of phase space and this improvement is largest for large R and low p_T . These metrics indicate that the ML-based background subtraction yields a significant performance improvement over traditional techniques, which will enable measurements to lower in p_T and larger in R than previously possible. This is due to the 5σ cut discussed in Section 5.4, where for the AB curve shown in the left panel of Figure 5.3 $5\sigma \sim 60$ GeV/c and for the ML curve $5\sigma \sim 30$ GeV/c. As a result, in the ML case, jets with much lower p_T will be given to the unfolding procedure, which should lead to a larger kinematic range in the resulting measurement.

5.6 Jet Performance

Thorough evaluation of the jet performance is a critical step in any jet measurement and evaluates how well that jet measurement can be done. In this case, comparing the jet performance of the AB and ML-based corrections is an essential metric in evaluating how an ML-based background correction will improve jet measurements. The following includes the jet performance of the AB method in comparison to the ML-based corrections. Included in the main text are plots for central (0-10%) charged jets. The jet performance plots for semi-central and peripheral charged jets with $R = 0.2, 0.4, 0.6$ are included in Appendix C. Jet performance is quantified in terms of the Jet Energy Scale (JES) and the Jet Energy Resolution (JER). These quantities are mathematically defined in Equations 5.5

and 5.6.

$$\Delta_{\text{JES}} = \left\langle \frac{p_{\text{T,rec}} - p_{\text{T,true}}}{p_{\text{T,true}}} \right\rangle \quad (5.5)$$

$$\text{JER} = \frac{\sigma(p_{\text{T,rec}})}{p_{\text{T,true}}} \quad (5.6)$$

In the above equations, $p_{\text{T,true}}$ refers to the “true” transverse momentum or the transverse momentum at the particle level. The **JES** and **JER** are shown for $R = 0.6$ jets in central (0–10%) collisions. The influence of background fluctuations is largest at lower p_{T} , as evidenced by the larger **JER** in this region. Note that the bumpiness in this distribution is due to the fitting procedure to calculate the **JES** and the **JER**.




Figure 5.9: Left: JES for $R = 0.6$ jets in 0–10% central collisions. Right: JER for $R = 0.6$ jets in 0–10% central collisions

5.7 Unfolding

5.7.1 Introduction

The detectors used for **HEP** measurements have a finite resolution, which smears the quantities of interest. Unfolding, or deconvolution, is a technique to recover the “true information” that would have been measured with a perfect detector, which is referred to as a particle-level jet. This “true information” is then used to make comparisons between jets measured in experiment and jets calculated in theory, as shown in Figure 5.10. Additionally, this “true information” is also what is used to make comparisons between different experimental measurements. Theoretical calculations are typically performed at the parton

level, where **NP** corrections for effects like hadronization and the **UE** must be applied to compare to the particle-level jet measured in experiment.




Figure 5.10: Overview of how comparisons between data and theory jets are made.

Unfolding is used following background subtraction to remove residual fluctuations, uncorrelated effects of the underlying event, and detector effects. The detector effects for which unfolding corrects include tracking inefficiency, track p_T resolution, missing long-lived neutral particles, gaps in acceptance, etc. The better the correction for the background is, the more mathematically stable the unfolding procedure will be.

Though unfolding is an effective tool, there are a few ambiguities in the procedure. Firstly, unfolding assumes the shape of the spectrum and the fragmentation pattern. In unfolding, it is assumed that observed jets have the same shape and fragmentation pattern as simulation, which is not true, especially in heavy-ion collisions. The second ambiguity is that there is not always a one-to-one correspondence between a particle-level (or truth) jet and the jet as measured in the detector. A number of checks of the unfolding procedure are completed to ensure that these ambiguities do not impact the final result. These checks are described in detail in Section 5.7.3.

One possible alternative procedure is to smear particle-level theory jets using a detector response to create theory predictions at the reconstructed level that can be compared directly to the reconstructed distribution. The disadvantage of this method is that it is more difficult for theorists to make comparable predictions as the detector response must be provided in advance of publication. Comparing experimental measurements is also more difficult in this scheme.

The unfolding procedure works by creating a mapping from true jets to reconstructed jets that is achieved by a **MC** embedding procedure as described in Sections 5.2.2 and 5.7.2. This mapping is represented in matrix form and is commonly referred to as the response matrix, which is mathematically defined in Equation 5.7.

$$Rp_{T,MC}^{\text{true}} = p_{T,MC}^{\text{rec}} \quad (5.7)$$

Here, mathematically the response matrix (R) represents the probability as given in Equation 5.8.

$$R_{ij} = \text{Prob}(\text{observed in bin } i | \text{true in bin } j). \quad (5.8)$$

The overall goal of the unfolding procedure is to create a mapping from the reconstructed jet to the true jet, which can then be applied to reconstructed data. To do this, the mapping is inverted by taking the inverse of the response matrix, as shown in Equation 5.9.

$$p_{T,\text{data}}^{\text{true}} = R^{-1} p_{T,\text{data}}^{\text{rec}} \quad (5.9)$$

At its core, unfolding is a matrix inversion. Hence, if there are strong non-diagonal components to the response matrix, then the matrix inversion, and therefore the unfolding, is made difficult. This is one reason why it is necessary to perform the background subtraction prior to the unfolding procedure. Taking the inverse also often results in an instability with respect to the statistical fluctuations. Within statistical uncertainties the smeared data can be explained by the actual physical solution, but also by a large family of unphysical solutions. In a physical solution, a degree of smoothness can be expected, which is imposed in the unfolding procedure through a process called regularization. There are many different unfolding methods, each of which handles regularization differently. The two most common methods of unfolding are Bayesian unfolding [193] and [Single Value Decomposition \(SVD\)](#) unfolding [194]. The work presented in this thesis utilizes Bayesian unfolding. OmniFold [166], an ML-based unfolding method has also become popular in recent years, but will not be discussed in detail here.

Bayesian unfolding uses Bayes Theorem to estimate the true distribution given the prior, the response matrix and the measured distribution. Regularization is introduced into this procedure via its iterative nature, where the true distribution for one iteration will be taken as the “prior” for the next iteration. [SVD](#) unfolding is based on the single value decomposition method that decomposes a given matrix A with dimensions $n \times m$ into three matrices as written in Equation 5.10. Here U is an $n \times n$ matrix with n orthogonal columns, S is an $m \times m$ diagonal matrix where the diagonal elements are singular values that are ≥ 0 , and V is an $m \times m$ matrix with m orthogonal columns.

$$A = USV^t \quad (5.10)$$

The [SVD](#) method is easily adaptable for unfolding problems as it is easy to compute the inverse by taking the inverse of the different decomposed matrices. This process is easier to perform on the decomposed matrices, exploiting for example the fact that the inverse of an orthogonal matrix is its transpose. This is written mathematically in Equation 5.11.

$$A^{-1} = VS^{-1}U^t \quad (5.11)$$

Regularization is imposed in this procedure by a regularization parameter, k , which prevents high-frequency singular values in the inverse by replacing each singular value, s_i with $s_i^2 \rightarrow \frac{s_i^2}{s_i^2 + s_k^2}$. When the value of k is too small the result can be dominated by information from the prior and when the value of k is too large the result can be dominated by statistical fluctuations. All unfolding methods should yield consistent results within their associated uncertainties, which is described in Section 5.9. Each method must also pass the various checks of the unfolding procedure in order to ensure that the result is not dependent on the parameters associated with that method, or the simulation utilized for the unfolding procedure. Details of these checks are included in Section 5.8.

5.7.2 Forming the Response Matrix

The response matrix, as defined in Equation 5.7 reflects a mapping between jets from the embedded event at the true (particle) and observed (hybrid) levels. The response matrix is formed by matching particle-level and hybrid-level jets in the following two-step procedure. First, the hybrid-level jet is geometrically matched to a detector-level jet, where only matches with a maximum distance of $0.75 \times R$ are accepted. The hybrid and detector-level jets are additionally required to share 50% of the detector-level jet p_T . Then the detector-level jet is geometrically matched with a maximum distance of $0.75 \times R$ to a particle-level jet. These matched jets form a correspondence between the particle- and hybrid-level jets, which is then used to fill a response matrix. The jet reconstruction efficiency, defined as the ratio of the number of accepted detector-level jets geometrically matched to a particle-level jet and the number of particle-level jets in a given p_T^{true} interval written in Equation 5.12, accounts also for the efficiency of matching jets.

$$\varepsilon_{\text{rec}}(p_{T, \text{ch jet}}^{\text{true}}) = N_{\text{matched}}(p_{T, \text{ch jet}}^{\text{true}}) / N_{\text{generated}}(p_{T, \text{ch jet}}^{\text{true}}) \quad (5.12)$$

The jet reconstruction efficiency is high in all regions of phase space, as shown in Figure 5.11, and is used to correct the jet p_T spectrum after unfolding. In addition, a correction for the kinematic efficiency, given by the ratio of the truth jets in the measured kinematic range over the truth jets in the unfolded kinematic range, is also applied. The final result is reported in regions where this kinematic efficiency is large ($> 70\%$).

When forming the response matrix from the embedded event, additional care must be taken to account for the fact that the tails of the hybrid distribution are dominated by the data distribution. This arises from the fact that low $p_{T, \text{hard}}$ bins, which have a high weight when the $p_{T, \text{hard}}$ -binned spectrum is combined, are dominated by low p_T MC jets and therefore the data distribution dominates a high p_T . To remove this contribution, a $p_T > 10 \text{ GeV}/c$ cut is applied at the particle level. Such a cut has a negligible impact on the final result due to its kinematic range. The hybrid distribution for the various $p_{T, \text{hard}}$ bins as well as the combined distribution is shown in Figure 5.12. The particle-level cut




Figure 5.11: Jet reconstruction efficiency for $R = 0.6$ jets.

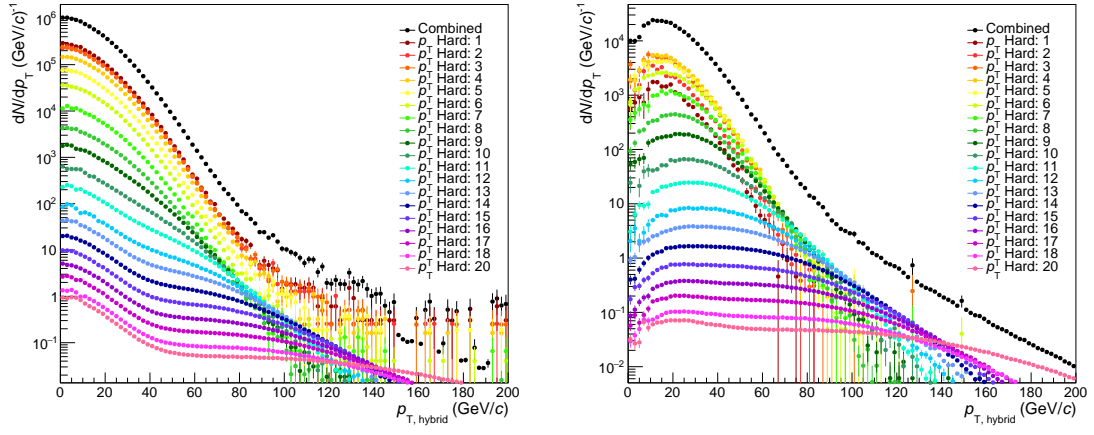


Figure 5.12: The hybrid distributions for the various p_T hard bins for $R = 0.6$ jets in central (0-10%) collisions. These distributions are plotted without the $p_{T,\text{part}} > 10$ GeV/c cut (left) and with the $p_{T,\text{part}} > 10$ GeV/c cut (right).

effectively removes high p_T outliers, making the combined distribution smooth.

5.7.3 Checks of the Unfolding Procedure

For checks of the unfolding, plots from a single jet definition are shown in order to demonstrate the check. For the result of unfolding checks for other parts of phase space, please refer to Appendix C.

Trivial Test

The trivial closure test is the most basic test of the mechanics of the unfolding itself. This test works by unfolding the exact same distribution that is used to fill the response. This means that, almost trivially, the unfolded result should agree exactly with the true distribution that fills the response as shown in the left panel of Figure 5.13.




Figure 5.13: Left: (Top) folded distributions for each iteration of the Bayesian unfolding procedure for the Trivial Test. (Bottom) Ratio of the unfolded to the true distribution as a function of the number of iterations for $R = 0.6$ charged jets in the 0–10% most central Pb–Pb collisions. Right: (Top) Unfolded distributions for each iteration of the Bayesian unfolding procedure for the split MC test. (Bottom) Ratio of the unfolded to the true distribution for $R = 0.4$ 0–10% full jets where the reconstructed distribution that we unfold is made up of the jets that make up the response.

Split MC Test

The Split MC test is similar to the trivial closure test in that the embedded MC is being unfolded, but it tests the statistical stability of the numeric unfolding procedure. To complete this test, the embedded MC is split and a fraction fills the response and the remaining fraction fills a “pseudo-data” distribution that is later unfolded. The percentage used to split the embedded MC is typically set by the available statistics in the MC and data samples.

The pseudo-data is then unfolded and checked for consistency with the true distribution. The results from the split **MC** closure test for $R = 0.6$ charged jets in the 0–10% most central collisions, which indicate good closure, can be found in the right panel of Figure 5.13.

Refolded to Raw Test

The refolded to raw ratio test ensures that the unfolding procedure has cyclic closure. Cyclic closure implies that the reverse of the unfolding procedure, called refolding, can be applied on the unfolded distribution and returns the data distribution originally used as input as written in Equation 5.7. The ratio of the refolded to raw distribution for $R = 0.6$ charged jets in the 0–10% most central collisions, which indicate good closure, can be found in the left panel of Figure 5.14.

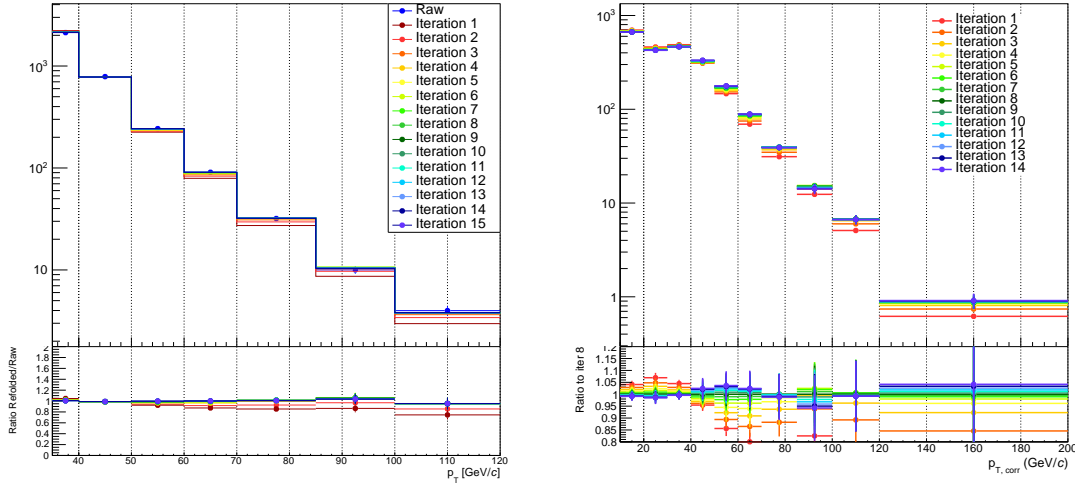


Figure 5.14: Left: (Top) Refolded distribution as a function of the iteration of the Bayesian unfolding procedure. (Bottom) The ratio of the refolded to raw distribution for $R = 0.6$ charged jets in the 0–10% most central Pb–Pb collisions. Right: (Top) Unfolded distribution as a function of the iteration of the Bayesian unfolding procedure. (Bottom) Ratio of the unfolded result at a given iteration to the result obtained with 8 iterations for $R = 0.6$ charged jets in the 0–10% most central collisions.

Stability with the Number of Iterations

A key feature of an unfolding procedure that is operating normally is that the resulting distribution is independent of the regularization used. For Bayesian unfolding this is the number of iterations used. Generally, the metric for stable unfolding is that the deviations between iterations remain 10% or less, however, any sort of pattern or other sort of

structure can also be a cause for concern. The results from the test of the stability with the number of iterations, which indicate that the result is stable, can be found in the right panel of Figure 5.14.

Refolding Check of the ML Estimator

The AB estimator serves as a natural cross-check of the ML estimator in kinematic regions where both background subtraction methods can be used (see Figure 5.21). However, for $R = 0.6$ jets, this is not possible as the AB method cannot be used in this region. To cross-check the ML-based result for $R = 0.6$ jets in central (0–10%) collisions, the ML-based spectrum can instead be refolded with the AB response. The refolded result is then compared with the AB corrected jet spectrum at the reconstructed level, which is shown in Figure 5.15. As seen in Figure 5.15, these spectra have a strong agreement, indicating that the ML-based result for $R = 0.6$ jets in central 0–10% collisions is consistent with the result that would have been achieved with the AB method.




Figure 5.15: The refolded spectrum of the ML-based $R = 0.6$ result in 0-10% collisions as compared to the reconstructed area-based spectrum.

5.8 Quantification of the Fragmentation Bias

The ML estimator is trained on jets generated with PYTHIA 8 simulations for pp collisions, where the fragmentation¹⁴ is known to differ quantitatively from the fragmentation in Pb–Pb data to which the estimator is applied [30, 31, 106, 195]. The inclusion of

¹⁴For a definition of the fragmentation function, refer to Equation 2.19.

constituent information in training the ML estimator introduces an explicit fragmentation dependence. It is common for background subtraction methods to introduce a fragmentation bias by selecting jets with a harder core that are more easily distinguished from the background. For example, the AB correction is often combined with a requirement on the p_T of the leading charged constituent in order to suppress the background contribution, which itself biases the fragmentation of the jet sample. The impact of training the estimator on jets with a different fragmentation than those to which the estimator is applied must be quantified. Overall, there are two points in the analysis procedure that are sensitive to the simulated fragmentation: the response matrix and the training.

The dependence on this fragmentation is quantified by varying the fragmentation in physics-motivated ways consistently in the response and training. These phenomenological modifications are performed by radiating additional hadrons from existing jet constituents in simulation, thereby changing the final-state hadron distribution.¹⁵ The modifications are governed by tunable parameters specifying p_{loss} , f_{loss} , and ΔR , where for each constituent a particle is radiated with probability p_{loss} , carrying a specified fraction f_{loss} of the constituents original p_T at an angle randomly sampled from a uniform distribution between 0 and ΔR . When applicable, the three-momentum of the radiated particle is then subtracted from the original constituent three-momentum, and the radiated particle is added to the list of jet constituents if it falls within the jet cone. Three different shower modifications were studied using this framework, with each variation modeling a different aspect of in-medium jet modification discussed in Section 2.5:

1. **Fractional Collinear:** the radiated particle carries a fraction of the original constituent's energy (specified by f_{loss}) and is emitted predominantly within the jet cone, by using $\Delta R = 0.1, 0.2$, and 0.4 for $R = 0.2, 0.4$, and 0.6 , respectively. This mimics the effect of collisional and radiative energy loss that remains inside the jet cone.
2. **Fractional Large Angle:** the radiated particle carries a fraction of the original constituent's energy (specified by f_{loss}) and frequently radiates outside the jet cone by using $\Delta R = 0.4, 0.6$, and 0.8 for $R = 0.2, 0.4$, and 0.6 , respectively. This mimics the effect of collisional and radiative energy loss that moves outside the jet cone.
3. **Medium Response:** the emission occurs as described for the Fractional Collinear case, but the original jet constituent p_T is unmodified, emulating the addition of particles from the medium into the jet.

The values of the tunable parameters used in each of these modifications were guided by observations of modification in data. Specifically, the p_{loss} values were determined by evaluating the excess particle yield for jets in Pb–Pb collisions compared to those in pp collisions using the jet radial profiles for $R = 0.4$ inclusive jets above $100 \text{ GeV}/c$ [109]. Each

¹⁵This can be done easily in the framework of embedding, as the combination of signal and background consists of simply adding together the list of tracks.

of the modifications uses the p_{loss} value extracted from data, which combines contributions from all of the above-mentioned effects. This approach overestimates the contribution of each individual effect, which helps to account for any differences when extrapolating these values to a new region in phase space. The values of f_{loss} were set to 25% and 10%, which provide a charged hadron R_{AA} of comparable magnitude to the measured values in the 0–10% and 30–50% centrality ranges, respectively [196].

Another variation of the fragmentation model is obtained by only utilizing either quark or gluon jets. Quark jets tend to be narrower and have fewer constituents, each of them carrying a significant fraction of the jet’s momentum (harder fragmentation), while gluon jets tend to be wider and have more constituents carrying smaller fractions of the jet’s energy, as seen in Figure 5.16. The inclusive jet population contains a mixture of quark and gluon jets, meaning that utilizing only quark or gluon jets provides significant variation in the fragmentation. The modified and the unmodified jet fragmentation distributions are compared by their ratio, $R_{\text{AA}}^{\text{toy}}$ which is given in Equation 5.13.

$$R_{\text{AA}}^{\text{toy}} = \frac{Y_{\text{modified}}}{Y_{\text{unmodified}}} \quad (5.13)$$

The $R_{\text{AA}}^{\text{toy}}$ distributions are shown for $R = 0.4$ jets as a function of p_{T} in the left panel of Figure 5.16. The medium response adds energy to the jet cone, which therefore results in $R_{\text{AA}}^{\text{toy}} > 1$. For the fractional collinear model, the jet does not lose energy as the energy “lost” is retained within the cone, which results in $R_{\text{AA}}^{\text{toy}} = 1$. In the case of the fractional large angle model, the jet will lose energy, resulting in $R_{\text{AA}}^{\text{toy}} < 1$.

To quantify the modifications introduced by the various fragmentation scenarios, the ratio between modified and unmodified jet fragmentation functions as a function of z is shown in the right panels of Figure 5.16. Both panels include comparisons to the measured ratio of fragmentation functions in Pb–Pb and pp collisions for $R = 0.4$ inclusive jets from ATLAS [30] and $R = 0.3$ photon-tagged jets from CMS [31]. The kinematic region of the ATLAS measurement is the only region where the fragmentation function of inclusive jets has been measured. The CMS measurement is a quark-dominated sample and does not fully describe the phase space measured in this analysis, but it is still useful for the purpose of comparing the magnitude of the induced variations in the toy models. The toy model variations qualitatively cover the modification of fragmentation function in previously-measured regions of phase space. Unfortunately, due to the lack of fragmentation measurements in the phase space covered by this measurement, a direct comparison is not possible.

For each variation, both the training and the response matrix were varied to quantify the effect of a different fragmentation model through the full analysis chain. In principle, any unfolded heavy-ion measurement could have a fragmentation bias that is inherent to the unfolding procedure that assumes a fragmentation model. Unique to the ML method is the fragmentation dependence of the training, but the results varied minimally when only




Figure 5.16: Left: The R_{AA}^{toy} distribution for $R = 0.4$ jets using the various modifications. Right: The fragmentation functions at low jet p_T ($40 < p_{T, \text{true}} < 100$ GeV/c, lower right panel) and high jet p_T ($100 < p_{T, \text{true}} < 200$ GeV/c, upper right panel) for 0–10% central Pb–Pb collisions. The ratio of the fragmentation functions measured in Pb–Pb and pp collisions are shown for jets with $R = 0.4$ and $p_T > 100$ GeV/c [30] (ATLAS), and for jets with $R = 0.3$ and $p_T > 30$ GeV/c recoiling from a photon with $E_T > 60$ GeV/c [31] (CMS).

the training sample fragmentation was varied (see Figure 5.18), indicating that the main effect originates from the fragmentation in the response matrix.

5.8.1 Using JEWEL Fragmentation for Training

Note that the investigations using JEWEL discussed here are not included in the final result, like the investigations discussed above, but instead are included as an additional check.

As an additional check of the fragmentation, simulations from JEWEL [134, 135]¹⁶ both with and without recoils, which includes jet quenching effects, were used. Though JEWEL is not a perfect model for jets in heavy-ion collisions, it can adequately demonstrate the effect that a varied fragmentation has on the training, performance, and final results achieved with the ML-based estimator. Note that these studies were only completed with $R = 0.4$ jets as the treatment of recoils in the case of larger R jets is non-trivial.

The first metric used to compare the ML estimator trained with JEWEL (both with and

¹⁶For more information about JEWEL, see Section 2.5.4.

without recoils) to the ML estimator trained with PYTHIA is the performance of these estimators. This is quantified via the δp_T distributions as defined in Equation 5.2. The δp_T distributions for $R = 0.4$ charged-particle jets in 0–10% collisions are shown in Figure 5.17 for JEWEL with (right panel) and without (left panel) recoils. The ML-based method trained on JEWEL exhibits an improvement in performance, evidenced by a decrease in the width, over the AB method for JEWEL with and without recoils. In the right panel of Figure 5.17, comparisons between the performance of the model trained on JEWEL in vacuum, similar to PYTHIA, are also shown. Note that the performance of the model is similar between vacuum JEWEL (PYTHIA) and JEWEL without recoils, indicating a similar performance regardless of the training data used. JEWEL with recoils exhibits a slightly worse performance than vacuum JEWEL (PYTHIA) and JEWEL without recoils. However, all ML methods that are compared show a performance improvement over the AB method.

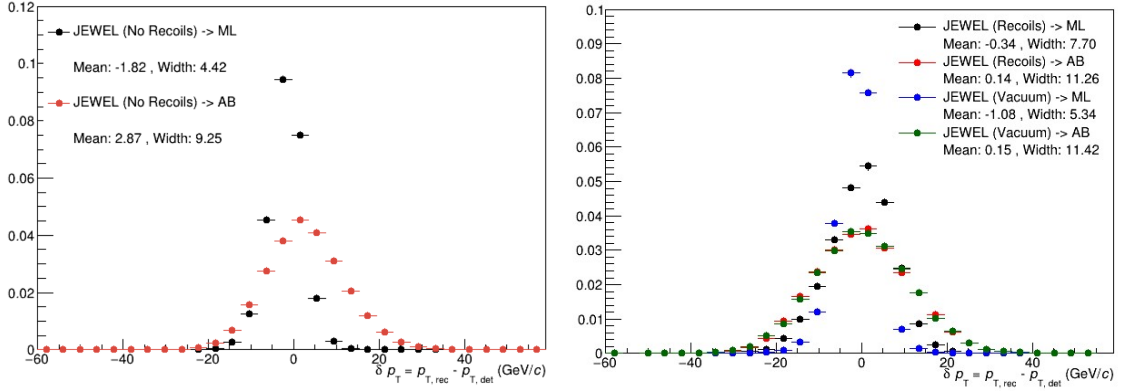


Figure 5.17: δp_T distributions for $R = 0.4$ jets both without recoils (left) and with recoils (right) for the area-based (AB) and ML-based (ML) background correction methods. In the right panel vacuum JEWEL, similar to PYTHIA, is also shown.

Another test was to compare the nuclear modification factors achieved with models trained on JEWEL both with and without recoils to that trained on PYTHIA (nominal). This comparison is shown in Figure 5.18. See Section 5.10 for more details on the nominal nuclear modification factor to which this test is compared. Results trained using JEWEL both with and without recoils are consistent with the results trained using PYTHIA, indicating that the ML estimator is robust to the simulation used in training. Note that this test demonstrates the sensitivity of the ML-based result to the fragmentation used in training only. However, the analysis is sensitive to the PYTHIA fragmentation not only in the training of the ML estimator, but also in the unfolding procedure. For comparisons of a full simulation of JEWEL to the unfolded result, refer to Section 5.11.




Figure 5.18: Nuclear modification factor for $R = 0.4$ charged-particle jets in 0–10% Pb–Pb collisions obtained with the ML-based background correction trained on PYTHIA, compared to results obtained with the ML-based background correction trained on JEWEL both with and without recoils.

5.9 Systematic Uncertainties

In this section, the sources of systematic uncertainty of the presented measurements will be discussed. Many systematic uncertainties are common between the full and charged jet analyses, where sources unique to the full jet analysis will be discussed in Section 6.6. How the systematic uncertainties were propagated during the calculation of the R_{AA} and jet cross-section ratios is discussed in Section 5.9.2.

5.9.1 Sources of Systematic Uncertainty for Charged Jets

The systematic uncertainties of the inclusive jet spectrum arise from three main sources: the tracking efficiency, the unfolding procedure, and the fragmentation dependence in the ML method. The combined systematic uncertainty on the jet spectrum is a quadratic sum of the individual uncertainties. The systematic uncertainties are summarized for 0–10%, 30–50%, and 60–80% central collisions in Tables 5.2, 5.3, 5.4 respectively.

Tracking efficiency uncertainty: tracking inefficiency in an analysis can result in a reduction in jet p_T , which corresponds to a reduction in measured yield in a given p_T interval. In pp collisions, this efficiency is estimated using pp events generated with the PYTHIA8 (Monash 2013 tune) [83] generator combined with the GEANT3-based detector simulation and response model of ALICE [197]. The tracking efficiency in pp collisions is approximately 67% for track $p_T = 0.15$ GeV/c, rises to approximately 84% at track $p_T = 1$

GeV/ c and remains above 75% for all higher track p_{T} s [14]. The tracking efficiency in 0–10% Pb–Pb collisions is estimated by comparing central to peripheral HIJING+GEANT 3 [198] events, resulting in an approximately 2% reduction in the tracking efficiency as compared to pp, independent of the track p_{T} . For the systematic uncertainty associated with the tracking efficiency, the nominal Pb–Pb tracking efficiency for jets in the response matrix was reduced by an extra 4% as compared to pp [199, 200].

Regularization parameter: for the Bayesian unfolding procedure utilized in this analysis, a regularization procedure, as described in Section 5.7, is applied, where the number of iterations is used as the regularization procedure. For the associated systematic uncertainty, the number of iterations used in the unfolding procedure is varied by ± 1 .

Prior: for the Bayesian unfolding procedure, as described in Section 5.7 a prior distribution is needed, which is nominally chosen to be the PYTHIA8 jet p_{T} spectrum in this analysis. The systematic uncertainty associated with the sensitivity of the unfolded result to the prior was evaluated by reweighting the PYTHIA8 spectrum by the parameterized ratio of the embedded MC to Pb–Pb data, accounting for any shape differences between the two. The difference between the result unfolded with the reweighted response and with the nominal response was taken as a systematic uncertainty.

Measured p_{T} -range: the minimum p_{T} of the jets that enter the unfolding procedure is determined by five times the residual fluctuations σ , which suppresses the fake jet yield. The systematic uncertainty associated with this low- p_{T} cut-off of the data that serves as input to the unfolding procedure was evaluated by varying this cutoff by ± 5 GeV/ c .

Fragmentation: the background estimator is trained using the jet spectrum from simulated pp collisions using PYTHIA8. Variations in the results obtained using different fragmentation models were used to estimate the systematic uncertainties, considering the following alternatives: q/g fragmentation, Medium Response, Fractional Collinear, and Fractional Large Angle. The variations are added in quadrature and the corresponding uncertainties are considered to be asymmetric.¹⁷ For more detail, see Section 5.8.

The uncertainties of the R_{AA} are calculated using the systematic uncertainties of the Pb–Pb jet p_{T} spectra and the uncertainties of the pp reference, propagated to the ratio. The systematics for the pp reference are evaluated using the same procedure as previous measurements [201], which will not be discussed in this thesis. Included in the R_{AA} uncertainty is an additional normalization uncertainty associated with the calculation of the T_{AA} [103].

5.9.2 Combining Systematic Uncertainties for the R_{AA} Ratios

For definitions of the variables discussed here, refer to Section 5.10.

¹⁷This is because different fragmentation models will always push the result in one direction.

Table 5.2: Relative systematic uncertainties (%) for jet spectra for 0–10% central Pb–Pb collisions and all resolution parameters. The maximum uncertainties for low p_T ($p_{T,\text{jet}} < 50$ GeV/c) and high p_T ($p_{T,\text{jet}} > 50$ GeV/c) are shown. The direction of asymmetric uncertainties is indicated with a + or – sign. The combined uncertainty is the quadrature sum of individual uncertainties.

Resolution parameter	0.2		0.4		0.6	
p_T	Low p_T	High p_T	Low p_T	High p_T	Low p_T	High p_T
Tracking eff.	21	18	24	12	12	34
Regularization param.	(< 2)	(< 2)	(< 2)	2	2	(< 2)
Unfolding prior	6	16	8	(< 2)	4	(< 2)
Measured p_T range	46	4	8	(< 2)	6	8
Fractional Collinear	+30	+12	+12	+16	+8	+20
Fractional Large Angle	+10	+10	+6	+10	+8	+14
Fractional Medium Response	+28	+14	+20	+14	+22	+14
Quarks/Gluon	-8	-6	-12	-12	-14	-12
Combined	58	32	32	26	28	44

Table 5.3: Relative systematic uncertainties (%) for jet spectra for 30-50% central Pb–Pb collisions and all resolution parameters. The maximum uncertainties for low p_T ($p_{T,\text{jet}} < 50$ GeV/c) and high p_T ($p_{T,\text{jet}} > 50$ GeV/c) are shown. The direction of asymmetric uncertainties is indicated with a + or – sign. The combined uncertainty is the quadrature sum of individual uncertainties.

Resolution parameter	0.2		0.4		0.6	
p_T	Low p_T	High p_T	Low p_T	High p_T	Low p_T	High p_T
Tracking eff.	10	12	12	16	12	14
Regularization param.	(< 2)	(< 2)	(< 2)	(< 2)	(< 2)	(< 2)
Unfolding prior	(< 2)	4	6	2	(< 2)	6
Measured p_T range	28	(< 2)	10	(< 2)	20	(< 2)
Fractional Collinear	+12	+6	+22	+14	+26	+24
Fractional Large Angle	+8	+8	+8	+10	+26	+24
Fractional Medium Response	+8	+8	+14	+10	+22	+20
Quarks/Gluon	-8	-6	-8	-6	-12	-6
Combined	34	22	32	26	42	40

The systematic uncertainties on the R_{AA} double ratios and the cross-section ratios for different R values are evaluated by separately treating correlated and uncorrelated uncertainties. All unfolding uncertainties for the spectra in Pb–Pb collisions are treated as uncorrelated and added in quadrature. The tracking efficiency uncertainty and the uncertainties due to the fragmentation dependence are treated as correlated by evaluating the double ratio for variation and calculating the difference from the nominal double ratio. The deviations from the nominal value for each of these cases are then added in quadrature to

Table 5.4: Relative systematic uncertainties (%) for jet spectra for 60-80% central Pb–Pb collisions and all used resolution parameters. The maximum uncertainties for low p_T ($p_{T,\text{jet}} < 50$ GeV/c) and high p_T ($p_{T,\text{jet}} > 50$ GeV/c) are shown. In this centrality interval the spectra are measured with the area-based method, the uncertainties related to the fragmentation functions adopted in the machine learning algorithm are not included for this case. The direction of asymmetric uncertainties is indicated with a + or – sign. The combined uncertainty is the quadrature sum of individual uncertainties.

Resolution parameter	0.2		0.4		0.6	
p_T	Low p_T	High p_T	Low p_T	High p_T	Low p_T	High p_T
Tracking eff.	4	10	2	10	10	14
Regularization param.	(< 2)	(< 2)	(< 2)	(< 2)	(< 2)	(< 2)
Unfolding prior	6	(< 2)	10	6	(< 2)	8
Measured p_T range	32	4	6	4	36	14
Quarks/Gluon	-2	-4	-4	-4	(< -2)	-4
Combined	46	14	14	10	50	16

obtain the final correlated uncertainty on the R_{AA} double ratios and the jet cross-section ratios.

5.10 Results

In this section, various measurements of charged jets are presented [122]. These include the inclusive jet transverse momentum spectra, the nuclear modification factors, jet cross-section ratios, and the ratios of the nuclear modification factors at different R . The results from pp collisions utilized as the reference in the nuclear modification factors, though not a part of this thesis work, are additionally shown for completeness.

5.10.1 Inclusive Jet Transverse Momentum Spectra

The inclusive jet transverse momentum spectra are fundamental components of any jet analysis, particularly for analyses involving jet suppression. The p_T -differential charged-particle jet cross-section in pp collisions is shown in Figure 5.19 for $R = 0.2$ to $R = 0.6$. Jets with larger R show higher yields due to more energy being captured in a larger radius at fixed jet p_T , with this effect being largest at low p_T as discussed in Section 2.4.5. The vacuum jet p_T spectra exhibit an R -dependent steepness that should be considered when interpreting the R -dependence of the R_{AA} .¹⁸

The inclusive charged-particle jet transverse momentum spectra in Pb–Pb collisions are presented as an event-normalized yield divided by the $\langle T_{AA} \rangle$ [103] of the given centrality class, as written in Equation 5.14. Note that the jet p_T spectra are reported differen-

¹⁸This is also discussed at length in Chapter 2 and Chapter 8.




Figure 5.19: The charged-particle jet p_T spectra in pp collisions for $R = 0.2, 0.3, 0.4, 0.5, 0.6$. The vertical bars denote statistical uncertainties and the vertical extent of the boxes denotes systematic uncertainties.

tially in η , meaning that they are divided by the fiducial η window. The inclusive charged jet p_T spectra in the 0–10%, 30–50%, 60–80% central Pb–Pb collisions are shown in the left, middle, and right panels of Figure 5.20 for $R = 0.2, 0.4$, and 0.6 .

$$\frac{1}{\langle T_{AA} \rangle} \frac{d^2 N_{\text{ch jet}}}{dp_{T, \text{ch jet}} d\eta_{\text{jet}}} [\text{mb (GeV/c)}^{-1}]. \quad (5.14)$$




Figure 5.20: The p_T -differential inclusive charged-particle jet yield distributions as a function of p_T for different values of R in three centrality classes: Left: 0–10% collisions, Middle: 30–50%, Right: 60–80% collisions. The peripheral spectra were measured using the [AB](#) method for the background correction. All other reported spectra were corrected with the ML-based background estimator.

5.10.2 Nuclear Modification Factors

As mentioned in Section 2.5.3, the nuclear modification factor, as defined in Equation 2.18, is the primary observable for experimental observations of jet suppression. An essential test of the new ML-based correction is to compare the results achieved with this method to those compared with the standard AB method. The ML-based method additionally enables measurements at lower jet p_T and for large R ($R = 0.6$) in the 0–10% most central Pb–Pb collisions.

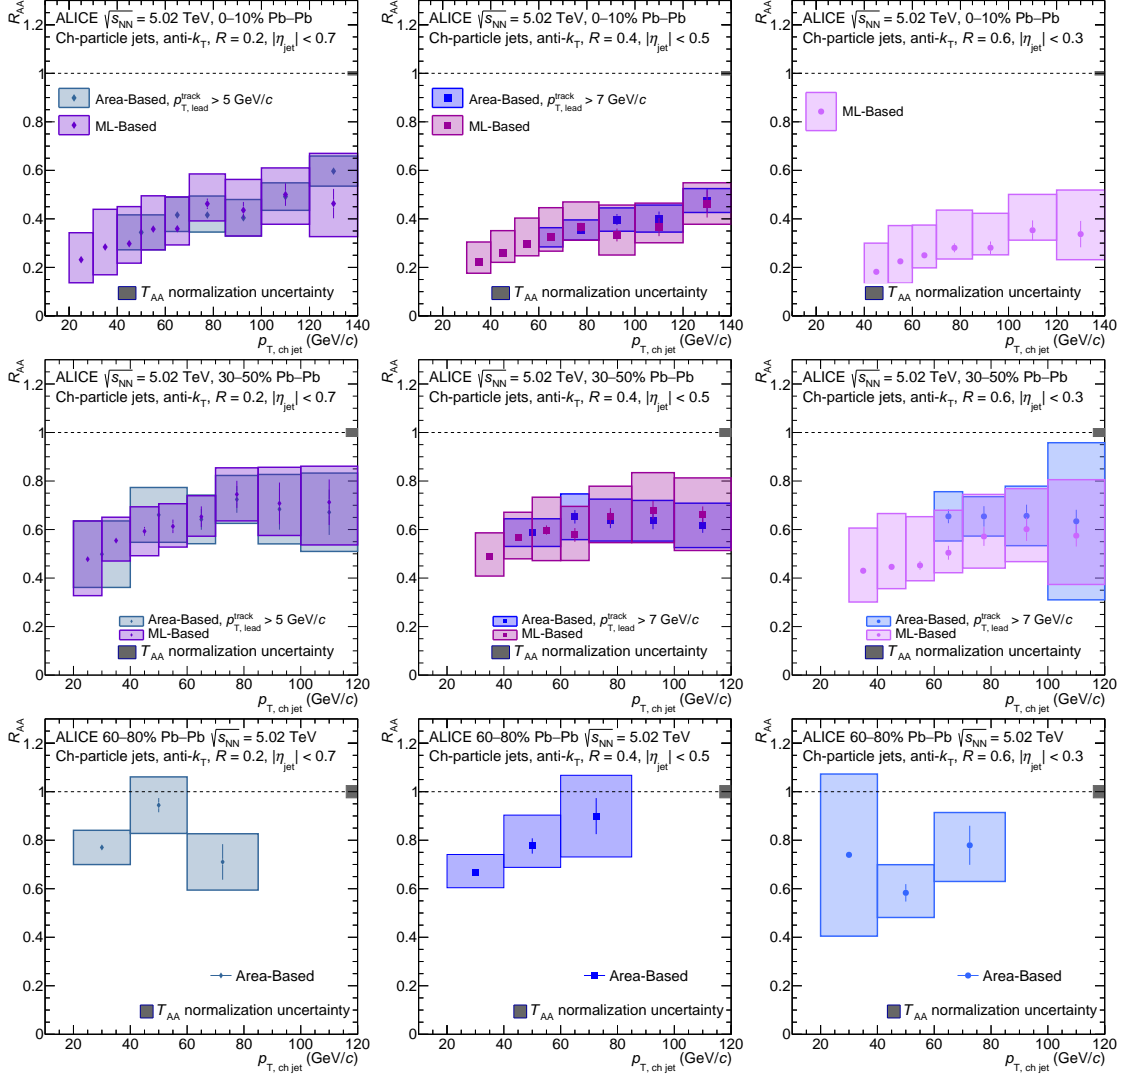


Figure 5.21: Nuclear modification factors of inclusive charged-particle jets as a function of p_T for $R = 0.2$, $R = 0.4$, and $R = 0.6$, shown for 0–10%, 30–50% and 60–80% central Pb–Pb collisions for the ML-based method compared to results obtained with the area-based method where applicable.

Figure 5.21 shows the nuclear modification factors using both the new ML-based es-

timator and the established [AB](#) estimator for the 0–10%, 30–50%, and 60–80% centrality classes for charged jets with $R = 0.2, 0.4$, and 0.6 . Due to its extensive coverage in phase space, these results represent the most differential measurement of jet suppression effects ever made in [ALICE](#). Measurements in the 0–10% collisions reveal an R_{AA} significantly below unity, consistent with the expectations of jet quenching and previous measurements. The degree of suppression, and therefore energy loss is less for measurements in the 30–50% and the 60–80% as expected due to the smaller size of the formed [QGP](#), and therefore the traversed path length, being less in these collision systems.¹⁹ Additionally, the R_{AA} measured using the [AB](#) and [ML](#)-based background subtraction methods is consistent within uncertainties in all kinematic regions of overlap. Notably, the [ML](#)-based method also allows for the extension of the measurement to lower in p_T than the [AB](#) method, particularly in more central collisions or for larger R . Note that in some cases the magnitude of the systematic uncertainty is larger for the [ML](#)-based method due to the additional fragmentation uncertainty not included in the [AB](#) method. For the case of peripheral collisions, this measurement was done using the [AB](#) method alone, as the [ML](#)-based method would not have considerably improved the result. Theoretical comparisons with these results are shown in Section 5.11.

Of additional interest is the relative contribution of the systematic uncertainty introduced by the fragmentation variations as discussed in Sections 5.8 and 5.9. This is shown in Figure 5.22 for $R = 0.6$ jets in both central and semi-central collisions. The relative contribution from the fragmentation systematic is largest for large R jets at lower jet p_T and is larger for central than semi-central collisions. The contribution of the fractional collinear model is the largest contribution to these fragmentation systematics both in central and semi-central collisions. The quark-only fragmentation variations are the only ones that result in a lower R_{AA} as it is the only fragmentation variation where the overall fragmentation is hardened.

5.10.3 Jet Cross-Section Ratios

The jet cross-section ratio is defined as the ratio of the per-event jet yields measured in the same collision system for different resolution parameters:

$$\sigma(R = R_1)/\sigma(R = R_2) = \frac{dN_{R_1}}{dp_{T, \text{ch jet}}} / \frac{dN_{R_2}}{dp_{T, \text{ch jet}}}. \quad (5.15)$$

In the ratio, the tracking and fragmentation uncertainties are highly correlated, while all other uncertainties are considered as uncorrelated. Measuring the inclusive jet cross-section as a function of R provides information on how the jet population changes with R and has been measured both at [RHIC](#) [27] and the [LHC](#) [14, 182, 190, 202, 203].

¹⁹Attributes such as the density and temperature of the [QGP](#) may also play a role.

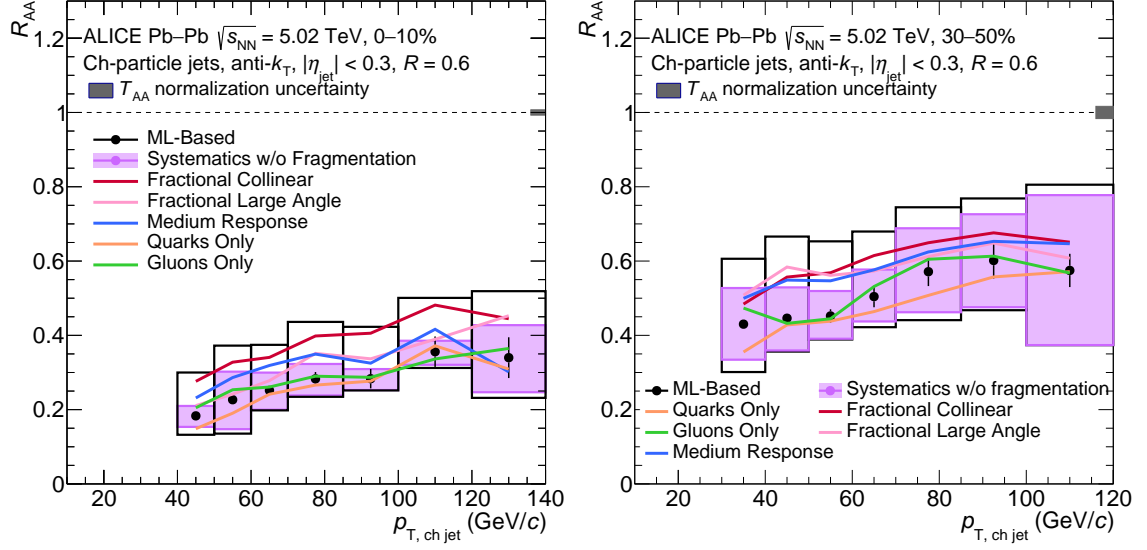


Figure 5.22: Nuclear modification factors for jets with $R = 0.6$ in 0-10% (left) and 30-50% (right) central Pb-Pb collisions outlining the impact of the various fragmentation models on the final result. Note that the systematic uncertainties are drawn both with and without the fragmentation uncertainties in the empty and filled boxes, respectively.




Figure 5.23: Jet cross-section ratios for $\sigma(R = 0.2)/\sigma(R = 0.4)$ (left) and $\sigma(R = 0.2)/\sigma(R = 0.6)$ (right).

The jet cross-section ratios are shown in Figure 5.23 for **pp** collisions and for Pb-Pb collisions in the 0-10% and 30-50% centrality intervals. The left panel of Figure 5.23 presents the ratios for $R = 0.2$ and $R = 0.4$ and the right panel for $R = 0.2$ and $R = 0.6$. The jet cross-section ratio in the right panel of Figure 5.23 for Pb-Pb collisions is significantly larger than that of **pp** collisions, taking into account the uncertainties. This

suggests a narrowing of the intra-jet energy distribution in Pb–Pb collisions. No strong centrality dependence is observed in the jet cross-section ratios. There is a small p_T -dependence in the jet cross-section ratios in [pp](#) collisions, which becomes stronger at lower jet p_T , indicating that the jet population changes with both p_T and R even in vacuum.²⁰ However, no such dependence on the jet p_T is observed in Pb–Pb collisions. This indicates that there is an R -dependence in the jet cross-section in [pp](#) collisions, which may impact the R -dependence of the R_{AA} as discussed below.

5.10.4 Nuclear Modification Factor Double Ratios

The double ratio of the nuclear modification factor, which compares R_{AA}^R for different R to $R_{AA}^{R=0.2}$, is used to quantify the variation of the nuclear modification factor with respect to the jet resolution parameter. This ratio is defined in Equation 5.16.

$$R_{AA}^{R/0.2} = \frac{R_{AA}^R}{R_{AA}^{0.2}} = \frac{\sigma_{AA}(R)}{\sigma_{AA}(0.2)} \bigg/ \frac{\sigma_{pp}(R)}{\sigma_{pp}(0.2)} \quad (5.16)$$

The double ratio defined here is not only a double ratio of nuclear modification factors, but also of jet cross-section ratios as defined in Equation 5.15. The R_{AA} double ratio is a key observable to quantify the R -dependence of energy loss: when this ratio is less than unity, jets with larger R are more suppressed; when it is consistent with unity, there is no R -dependence (or a cancellation of effects); and when it is greater than unity, larger R jets are less suppressed.

²⁰See Section 2.4.5 for more details.

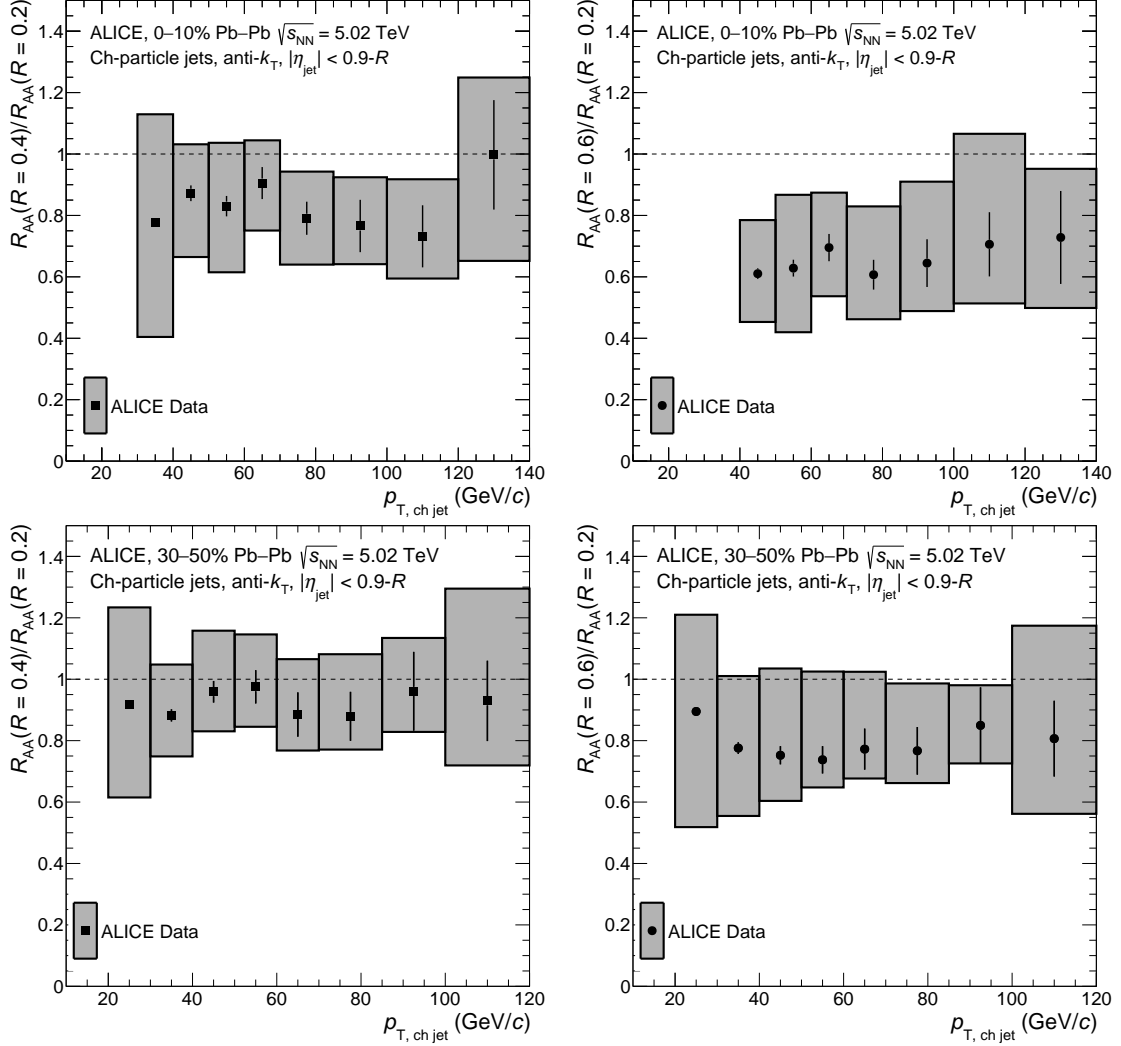


Figure 5.24: Double ratio of jet nuclear modification factors using $R_{AA}^{R=0.2}$ as the denominator and using $R = 0.4$ (left) and $R = 0.6$ (right) as the numerator for central (top row) and semi-central (bottom row) collisions.

The measured R_{AA} double ratios are shown in Figure 5.24 for 0–10% and 30–50% central Pb–Pb collisions. An R -dependence is observed for 0–10% central collisions where wider jets lose more energy, while no strong R -dependence is seen for 30–50% central Pb–Pb collisions. Many competing effects can be considered when interpreting this ratio. As observed in Ref. [109] and shown in Figure 2.23, the jet energy is lost predominantly to soft particles at large angles relative to the jet axis. Thus, as the jet R is increased, the energy lost outside of a smaller- R jet cone should be recovered in a larger- R jet cone. Additionally, effects from the wake could particles back inside the jet cone that would cause an increase in the R_{AA} with increasing R . Recent jet substructure measurements, such as those in Ref. [104, 117], show that the yield of jets with a more complex substructure, for

example wider jets, is suppressed at a fixed jet p_T compared to narrower jets. This may arise due to the greater number of effective energy loss sources in jets with a more complex substructure. The larger- R jets at a given p_T could be a population of jets with a greater fraction of gluon-initiated jets as compared to smaller- R jets, which may be more strongly quenched in the medium due to their larger color factor. Such an effect would cause the R_{AA} to decrease with increasing R . Finally, in vacuum the jet spectrum is steeper for large R jets as seen in Figure 5.23 and Ref. [18]. Thus, even for the same energy loss, some small decrease in the R_{AA} with increasing R would be observed.

The results shown in Figure 5.24 indicate a larger influence of wider jets with a more complex substructure experiencing more suppression and/or differences in jet populations of the spectra in pp collisions. This explanation may be compatible with the CMS measurement [18], which did not observe an R -dependence, due to differences in steepness of the jet spectra in pp collisions being less significant at higher jet p_T than lower jet p_T .²¹ Additionally, quenching effects are expected to be larger at lower jet p_T [183]. The results presented in Figure 5.26 are in contrast with the R -dependence of the ATLAS R_{cp} [184], which is discussed at length in Appendix A. Work to further interpret these results is presented in Chapter 7. Theoretical comparisons with this result are shown in Section 5.11.

5.11 Comparisons to Theory

In this section, all theoretical comparisons that have been provided are included. As a result, not all model comparisons are available for all parts of phase space or for all measurements presented in Section 5.10.

In Figure 5.25, the nuclear modification factors for $R = 0.2, 0.4$, and 0.6 jets are compared to theoretical models incorporating jet quenching, which are discussed in more detail in Section 2.5.4. Results are compared with: the Hybrid Model [95], which implements a energy loss with an AdS/CFT-inspired dependence on L and E , as well as a response of the medium to the lost energy; the Linear Boltzman Transport (LBT) model [137, 138] and LIDO [139], which use linear Boltzmann equations to describe the transport of partons in the QGP; JETSCAPE [142], which includes a medium-modified parton shower at high parton virtuality via MATTER [204], switching to the LBT model at low virtuality [137] (JETSCAPEv3.5 AA22 tune); Mehtar-Tani et al. [140], which is a first-principles analytical calculation of the single-inclusive jet spectrum using quenching factors; MARTINI [141], which embeds partons into a hydrodynamic medium with a modified parton shower; and JEWEL [134, 135], which consists of a MC implementation of BDMPS-based medium-induced gluon radiation in a medium modeled with a Bjorken expansion, including calculations both with and without recoils [136].

²¹For a comparison between the results in this thesis and the CMS measurement, see Figure A.5.

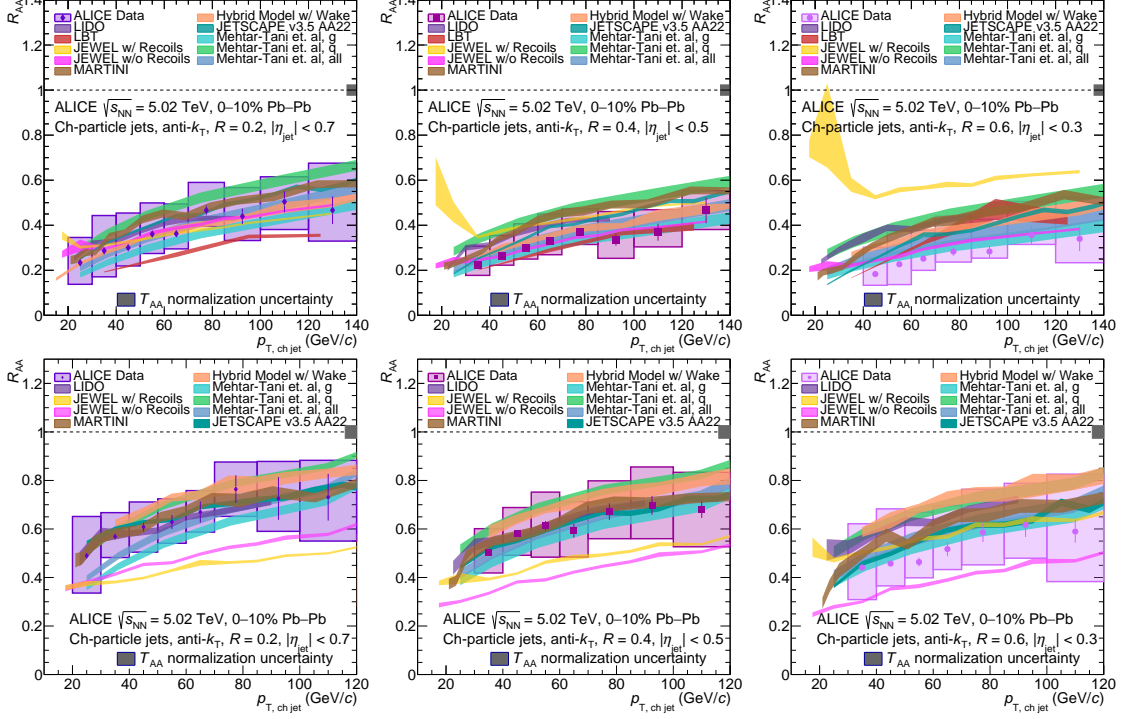


Figure 5.25: Nuclear modification factors for $R = 0.2$, $R = 0.4$, and $R = 0.6$, shown for 0–10% and 30–50% central Pb–Pb collisions compared to theoretical calculations incorporating jet quenching (see text for details).

These calculations generally describe the data in central collisions across all presented values of R with the exception of JEWEL with recoils, which overestimates the result at low jet p_T . This overestimation grows with R , consistent with an over-estimation of medium-response effects. The calculations span a larger range for the semi-central collisions but still mostly describe the data within uncertainties. These results are consistent with the idea that the individual R_{AA} 's are not a particularly discriminative observable.

The double ratio is also compared in Figure 5.26 to calculations that exhibit different dependencies with the jet R depending on the relative contributions of various energy loss mechanisms, making this a potentially discriminative observable. In addition to the models previously compared to the jet R_{AA} , further comparisons are made with: Qiu et al. [34], which is based on a factorization approach inspired by phenomenological considerations. JEWEL with recoils shows an increasing R_{AA} with increasing R due to the medium response, which is in contrast with the data, especially for large R . LBT also shows an increasing trend with increasing R , as well as a jet p_T dependence at large R , which is not supported by the data. As these two models both include recoiling medium partons, this indicates that the relative influence of recoiling medium partons may be overestimated. The Hybrid Model and LIDO only slightly change with resolution parameter. JETSCAPE, JEWEL without recoils, all variations of Mehtar-Tani et al., and the factorization model

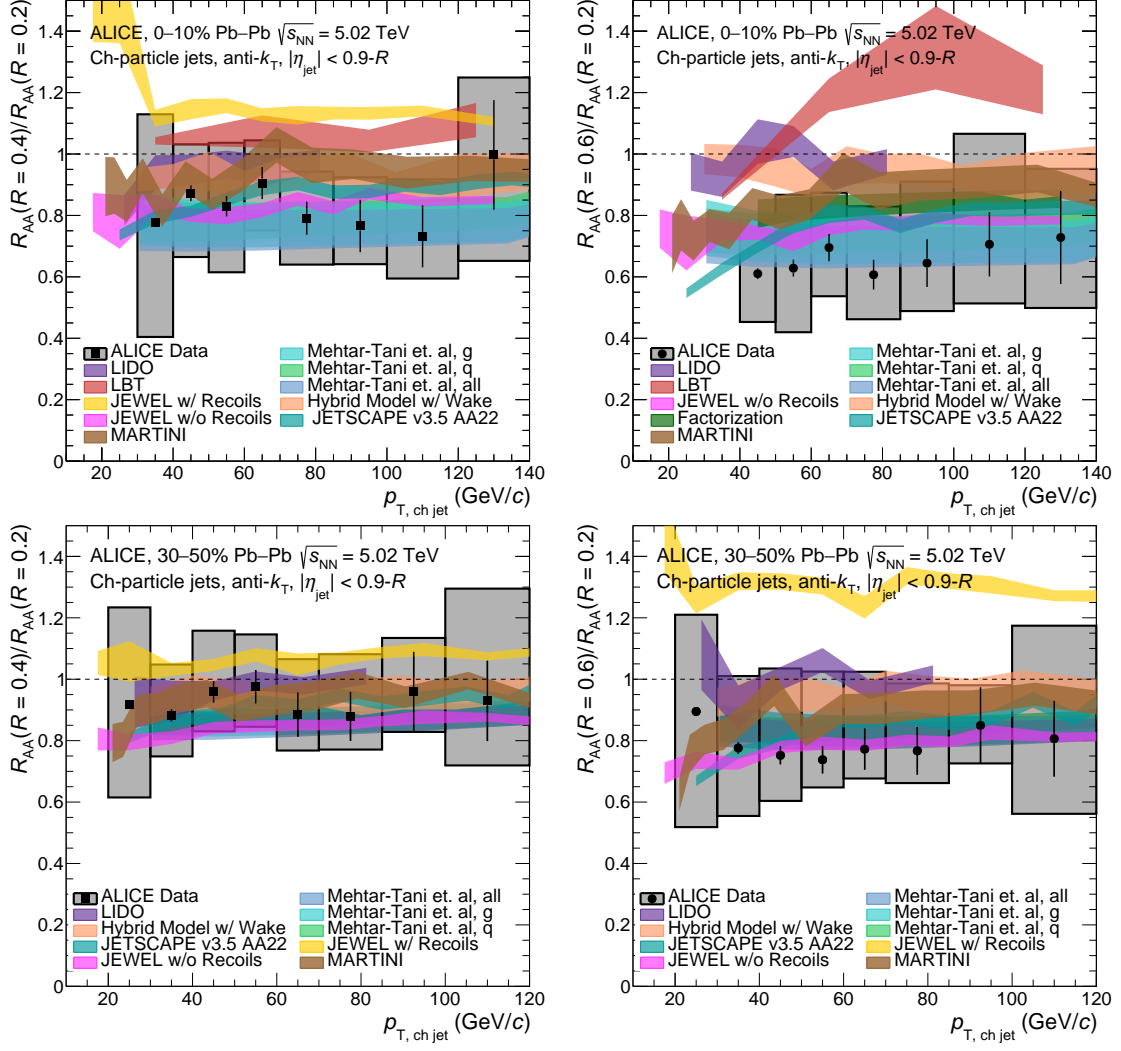


Figure 5.26: Double ratio of jet nuclear modification factors using $R_{AA}^{R=0.2}$ as denominator and using $R = 0.4$ (left) and $R = 0.6$ (right) as the numerator compared to model predictions for central (top row) and semi-central (bottom row) collisions. Note that a comparison to JEWEL with recoils was omitted from the right-hand plot as its prediction is out of scale.

show the R_{AA} decreasing with increasing R , describing the trend in the data.

Chapter 6

Full Jet R_{AA}

The work reported in this section is based on work performed as a part of the ALICE collaboration with full jets [205]. As much of the information related to this analysis overlaps with that of the charged jet analysis, the emphasis here is made on the differences between the two procedures. Details not mentioned in this chapter can be assumed to be identical between the two analyses. This result was completed for full jets with $R = 0.4$ in central collisions and was approved as [ALICE Preliminary](#). More differential extensions of this result are left for future work.

6.1 Data Selection and Embedding

The data selection for the full jet analysis is the same as described in Section 5.2, the main difference being that full jets contain both charged and neutral information. Neutral particles are reconstructed via electromagnetic clusters in the [EMCal](#), the selections for which are discussed below. For the embedded sample the information from electromagnetic clusters in the signal and background samples is combined as the direct sum of energy deposited in calorimeter towers (cells). Following this combination, the clustering algorithm is run again to form the clusters of the hybrid event. The track selection is identical between the full and charged jet analyses and will not be discussed here.

6.1.1 Cluster Selection

Relevant to the full jet results mentioned in this chapter is the identification and selection of electromagnetic clusters for the measurement of neutral particles. A single electromagnetic shower deposits energy over a few neighboring cells of the calorimeter. In order to remove cells that are not suitable for further analysis, a bad-channel calibration procedure is applied based on the average occupancy and energy deposition in a given cell. ¹ Fol-

¹Details of this bad-channel calibration procedure is discussed in Chapter 4.

lowing this calibration and those listed in Chapter 4, cells are clustered together using a clustering algorithm ² with $E_{\text{cell}} = 100 \text{ MeV}$ and $E_{\text{seed}} = 300 \text{ MeV}$, standard values for jet analyses. ³ There was also a timing requirement on the leading cell of the cluster so that $t_{\text{clus}} \in [-50 \times 10^{-9}, 100 \times 10^{-9}]$ seconds.

Charged particles also deposit energy in the calorimeter. These particles can either deposit part or all of their energy in the calorimeter. Tracks are propagated to the average shower depth within the EMCal and “matched” to a cluster. Tracks are considered to be matched if they are within a specified (η, ϕ) range of the cluster. This range can either be p_T -independent ($\Delta\eta < 0.015$, $\Delta\phi < 0.03$), or include a dependence on the p_T of the matched track. In the nominal case a p_T -dependent track matching is used, and the p_T independent track matching is included as a systematic uncertainty in the full jet analysis. Clusters are allowed to have multiple track matches, meaning that a single cluster can contain energy from multiple particles. In order to avoid double counting, clusters that contain matched tracks are corrected as written in Equation 6.1, where f nominally takes the value of 1 and specifies the fraction of energy to be subtracted.

$$E_{\text{corr}} = E_{\text{reco}} - f \sum_i p_i \quad (6.1)$$

Consider the case where the particle deposits only a fraction of its energy. In this case E_{corr} would be negative and as a result, the cluster is rejected. An improvement to such a correction would be to include a p_T dependence in f , but such studies are not included here. Other over-corrections will be handled via the unfolding, such as the case of multiple matched particles, which is a small effect due to the granularity of the calorimeter.

An additional correction is applied for the non-linearity of the detector response. This correction also includes an energy calibration of clusters such that $E_{\text{rec}}/E_{\text{in}} = 1$. In the case of embedding this correction is complicated by the fact that different non-linearity corrections are normally applied to MC and data, but the embedded cells contain contributions from both MC and data. To account for this difference, a correction is applied at the cell level where either the data cells are scaled to the level of MC or vice versa. The variation of this choice is included as a systematic.

6.2 Jet Reconstruction

Jets utilized as the signal discussed in this chapter are reconstructed using the anti- k_T method, with $R = 0.4$ with the p_T -scheme recombination scheme. Full jets utilize both charged tracks (with $p_{T,\text{track}} > 0.15 \text{ GeV}/c$) and neutral clusters (with $E_{\text{cluster}} > 300$

²Internally the algorithm used is referred to as the V2 clustering algorithm - the V3 algorithm is not yet compatible with the embedding framework.

³These selections are different for neutral meson analyses, see [here](#) for the neutral meson defaults.

MeV) as constituents. Jets both in data and simulation are required to be contained within $|\eta_{\text{jet}}| < 0.7 - R$, the fiducial acceptance of the calorimeter. Charged jets reconstructed with the k_T algorithm and the p_T recombination scheme are used for the calculation of ρ . A scaling factor to calculate the background for full jets will be later applied, as described in Section 6.3. For the ρ calculation, a jet radius of $R = 0.2$ is used regardless of the signal size in order to be less sensitive to fluctuations, though this is shown to have a small effect. Before the calculation of ρ the two jets with the highest p_T in the event are removed in order to reduce the contamination from the signal in the background calculation.

6.3 Background Subtraction

The AB background subtraction procedure for full jets is also given by Equation 5.1. However, the calculation of ρ must be adapted such that the event-averaged momentum density calculated using tracks additionally includes contributions from neutral particles. To do this, the value of ρ is scaled to account for the ratio of possible constituents in full jets as compared to charged jets. This scaling factor is centrality-dependent as written in Equation 6.2, where A_{calo} represents the geometric $\eta - \phi$ acceptance of the electromagnetic calorimeter. In this analysis the scaling factors from Ref. [14] were utilized.

$$s(C) = \frac{(\sum p_{T,\text{track}}^{\text{calo}} + E_{T,\text{cluster}}^{\text{calo}})/A_{\text{calo}}}{\sum p_{T,\text{Track}}^{\text{TPC}}/A_{\text{TPC}}} \quad (6.2)$$

6.4 ML-Based Background Subtraction

6.4.1 Training Data Set

In the full jet case, the fact that both true and background particles can contribute energy to the same cluster makes the creation of a thermal background for the purposes of training the ML more difficult. However, these effects are naturally included when using reconstructed data as the background. For the full jet analysis the training data set was created by embedding a PYTHIA signal into a Pb–Pb minimum bias data background. The regression target used was the matched PYTHIA detector-level jet p_T . The estimator trained using these two truth definitions yields similar performance, as shown in Figure 6.2, indicating that the results are not sensitive to this choice.

6.4.2 Input Parameters

The input parameters of the ML algorithm are largely the same as the charged jet case, but with the addition of neutral constituents in the input parameters. As charged jets contain $\sim 2/3$ the energy of full jets, 12 constituents (instead of 8) are used for full jets. The

constituent with the highest p_T is taken regardless of its identity as a charged or neutral constituent. The performance of the model treating charged and neutral constituents separately is shown in Figure 6.1. The input parameters chosen for the training are: the [AB](#) corrected jet p_T , mean of constituent track p_T , the jet angularity, p_T of the 12 leading constituents, and the number of tracks within the jet.

6.4.3 Model Investigations

The model investigations presented here are those that are relevant for the study of full jets. Additional model investigations are discussed in Section 5.5.6 and Appendix A.

Study of Separating Tracks and Clusters or Using a Constituent List

One essential investigation of the full jet [ML](#) estimator is how clusters and tracks should be included in training. For example, is it advantageous to use the six highest track and six highest cluster p_T s or simply to use the 12 highest constituent p_T s, regardless of whether or not the constituent is charged or neutral. The desired mechanism for choosing constituents is that which yields the best performance. The performance of these two choices is demonstrated in Figure 6.1, indicating that the constituents chosen from a single constituent list demonstrate an improved performance.




Figure 6.1: Comparison of the δp_T distributions for $R = 0.4$ jets in central collisions (0–10%) trained using the tracks and clusters selected separately and selecting from a full constituent list for the neural network estimator. The [AB](#) method is shown for the purposes of comparison.

This improvement arises from the fact that choosing the highest p_T constituents selects a different set of constituents in both embedded [MC](#) and data, as shown in Table 6.1. Training on a separate constituent list of clusters introduces a slight systematic bias due to the differences between the p_T -ordering of constituents in embedded [MC](#) being different from that in data. Training with a single constituent list allows for the p_T ordering to be different in embedded [MC](#) and data, avoiding such a bias. Due to this bias and the difference in performance, a single constituent list was used for training in this analysis.

Table 6.1: Percentage of jets which contain a leading, sub-leading, and sub-sub-leading track or cluster compared in embedded MC and data from $R = 0.4$ jets in 0-10% central collisions.

% of Jets	Embedded MC	Data
% with a leading Cluster	13.7%	33.8%
% with a leading Track	86.3%	66.2%
% with a sub-leading Cluster	19.6%	33.5%
% with a sub-leading Track	80.4%	66.4%
% with a sub-sub-leading Cluster	22.8%	32.8%
% with a sub-sub-leading Track	77.2%	67.2%

Comparison of the Two Truth Definitions

As mentioned in Section 5.5.3 there are two separate choices for the regression target; the true p_T fraction and the matched PYTHIA detector-level jet p_T . To ensure a robust estimator, it is important that these two definitions yield similar results. The left panel of Figure 6.2 indicates that both definitions result in similar performance. Though the true p_T fraction case exhibits slightly improved performance, the full jet analysis utilizes the matching regression target due to the difficulties of accurately computing the true p_T fraction for embedded clusters.

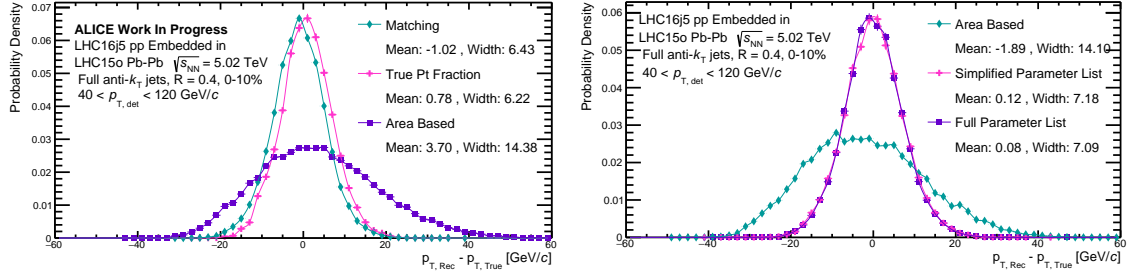


Figure 6.2: (Left): Comparison of the δp_T distributions for $R = 0.4$ jets in central collisions (0-10%) trained using the matching regression target and the true p_T fraction regression target (see Section 5.5.3) for the ML estimator. (Right): Comparison of the δp_T distributions for $R = 0.4$ jets in central collisions (0-10%) trained using the simplified input parameter list and the full parameter list (see Section 5.5.4) for the ML estimator. In both figures, the AB method is shown for the purposes of comparison.

Impact of Using a Simplified Input Parameter List

In creating a model for ML-based background estimators, simple models can help avoid a dependence on data subtleties. Therefore, as described in Section 5.5.4, a simplified input parameter list was created with the hopes of creating a simple model without significantly decreasing performance. The right panel of Figure 6.2 demonstrates that the choice of the

simplified input parameter list exhibits comparable performance to the performance using the full parameter list. Therefore, using the simplified input parameter list is preferable.

6.4.4 Performance

The performance of the **ML** estimator in the full jet case is evaluated via the δp_T distributions, as defined in Equation 5.2. The left panel of Figure 6.3 shows these distributions for the **AB** background correction as compared to the **NN**, random forest, and linear regression approaches. Overall, the ML-based approaches have a significantly reduced width in δp_T as compared with the **AB** estimator, indicating an improved performance. Additionally, all ML-based methods exhibit a similar performance, indicating that the **ML** methods are all picking up on similar features.

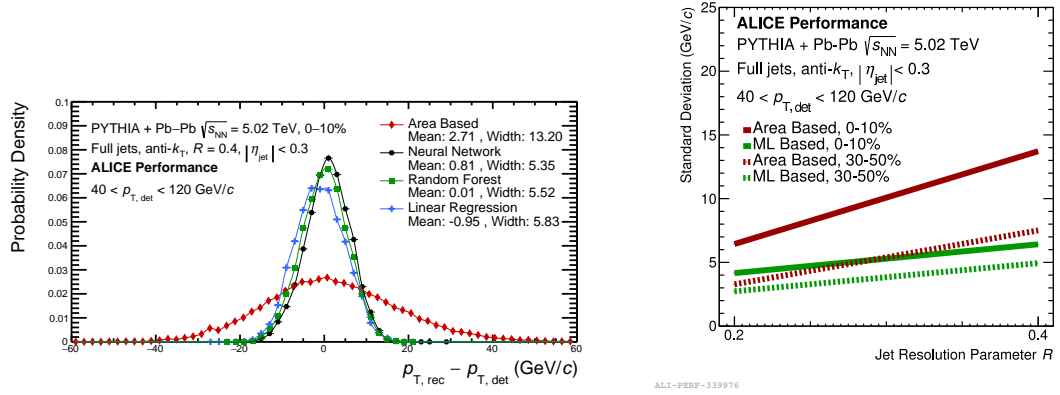


Figure 6.3: (Left): δp_T distribution for full jets with $R = 0.4$ in central (0–10%) collisions. Comparison of various ML-based and **AB** methods. (Right): Comparison of the width of the δp_T distributions for charged jets with several resolution parameters. Comparison of ML-based and area-based correction. The results in both panels are shown for $p_{T,det} > 40$ GeV/c to avoid edge effects to get a symmetric distribution even in the case of strong fluctuations.

The δp_T distributions were also evaluated for the case of full jets for various resolution parameters and centrality windows, which is summarized via the widths of these distributions in the right panel of Figure 6.3. In all explored regions of phase space, the ML-based method shows a reduced width as compared to the **AB** correction indicating improved performance. This improvement increases with R and is greatest for more central collisions, where the influence of the background is higher.

6.5 Jet Performance

The jet performance in the case of full jets is also evaluated via the **JES** and the **JER** (as defined in Equations 5.5 and 5.6, respectively). The **JES** and the **JER** are shown in Figure 6.4 for the **AB** correction as compared to the various ML-based corrections. The ML-based corrections all yield similar performance, which indicates that the ML estimators are all picking up on similar features. Of additional interest is the improvement in the **JER** as a result of the ML-based correction methods. The improvement in the **JER** shown in Figure 6.4 confirms the conclusions from the δp_T distributions of improved performance, which should lead to an improved jet measurement. Note that the value of the **JES** shown here is more negative than in the case of charged jets due to the additional missing hadronic energy from neutral hadrons, which is corrected for in the unfolding procedure.

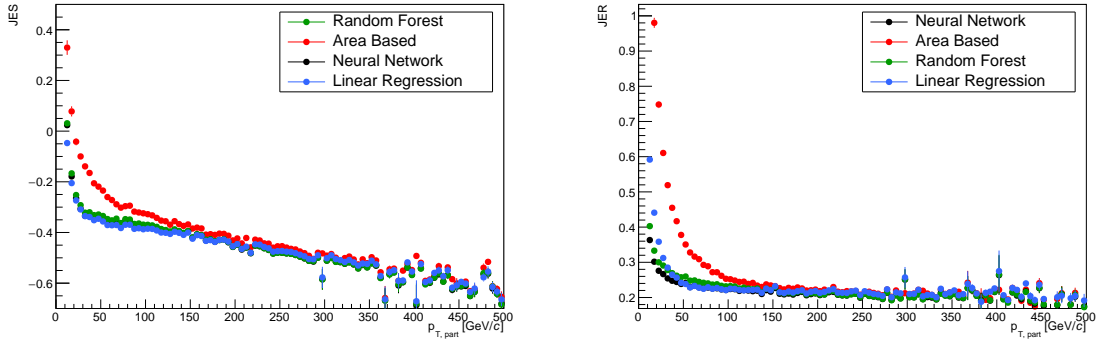


Figure 6.4: JES and JER (as defined in Equations 5.5 and 5.6 respectively) for $R = 0.4$ jets in central (0–10%) collisions.

6.6 Sources of Systematic Uncertainty for Full Jets

All systematic uncertainties mentioned in Section 5.9.1 are also applied in the full jet case with the exception of the toy model fragmentation variations (though the quark/gluon systematic is still applied). The quantification of the fragmentation bias on this result is discussed in Section 6.8.

Track Matching Procedure: as mentioned in Section 6.1.1, charged tracks are propagated to the average shower depth within the **EMCal** and matched to the cluster if they are in the specified (η, ϕ) range. The uncertainty associated with this matching procedure was evaluated by varying the p_T -dependent track matching procedure to be a p_T -independent track matching procedure, where tracks are required to be within $\Delta\eta < 0.015$ and $\Delta\phi < 0.03$ of a cluster in order to match.

Non-linearity: the response of the **EMCal**, used to measure neutral particles, is non-linear. For the case of embedded **MC**, as used in the response matrix, the correction to the combined calorimeter clusters can either be made to the data cells or the embedded **MC** cells. The associated systematic uncertainty for the non-linearity correction is evaluated by varying whether or not to correct the energy of the **MC** or the data cells in the embedded data used to construct the response matrix.

EMCal Response: there is an additional systematic uncertainty included for the description of **EMCal** response in **MC**. For the evaluation of this systematic uncertainty it is assumed that the **MC** describes the pp-like response well and uses the uncertainty calculated for the pp analysis at 5.02 TeV, as in Ref. [14].




Figure 6.5: The relative uncertainties of various systematic uncertainties.

The relative uncertainty resulting from each of the individual contributions is shown in Figure 6.5. The relative contribution of the non-linearity uncertainty is the largest in the full jet case. The systematic uncertainty on the R_{AA} is separated out into correlated and shape uncertainties for full jets. The correlated uncertainties refer to uncertainties that are positively correlated among all jet p_T bins. This includes the tracking efficiency, the non-linearity correction, the track matching procedure, and the uncertainty associated with the **EMCal** response. Shape uncertainties encompass those that impact the final shape of the jet p_T spectrum, which includes all other uncertainties.

6.7 Results

6.7.1 Inclusive Jet Transverse Momentum Spectra

The vacuum jet p_T spectrum used as a reference in this analysis is taken from previously published results [14], as shown in Figure 6.6. The full jet vacuum p_T spectra additionally vary in steepness depending on the R . This change in steepness is stronger for charged

jets in pp collisions than that of full jets (see Figure 7.1), which also may contribute to different R -dependencies in jet suppression.




Figure 6.6: The full jet p_T spectra in pp collisions for $R = 0.1, 0.2, 0.3, 0.4, 0.5, 0.6$, originally appearing in [14].

Full jets must also be corrected for the fiducial rapidity acceptance as $\Delta\eta_{\text{EMCal}} - 2R$. The spectra must also be corrected for the partial ϕ acceptance of the [EMCal](#), which is done by multiplying by the factor given in Equation 6.3.

$$\frac{2\pi}{\Delta\phi_{\text{EMCal}} - 2R} \quad (6.3)$$

The inclusive full jet transverse momentum spectra in Pb–Pb collisions are presented in a manner analogous to that of charged jets as written in Equation 6.4.

$$\frac{1}{\langle T_{\text{AA}} \rangle} \frac{d^2 N_{\text{jet}}}{dp_{T, \text{jet}} d\eta_{\text{jet}}} [\text{mb (GeV/c)}^{-1}] \quad (6.4)$$

The inclusive full jet p_T spectrum for $R = 0.4$ jets in the 0–10% most central Pb–Pb collisions is shown in the left panel of Figure 6.7 for results obtained with the ML-based and the [AB](#) methods. Here, the ML-based and [AB](#) methods are consistent with one another over the reported kinematic range, serving as a good crosscheck of the [ML](#) method.

6.7.2 Nuclear Modification Factors

The nuclear modification factor is defined in Equation 2.18. The nuclear modification factor for $R = 0.4$ full jets in the 0–10% most central Pb–Pb collisions is shown in the right panel of Figure 6.7. The R_{AA} exhibits a suppression (i.e. is significantly below

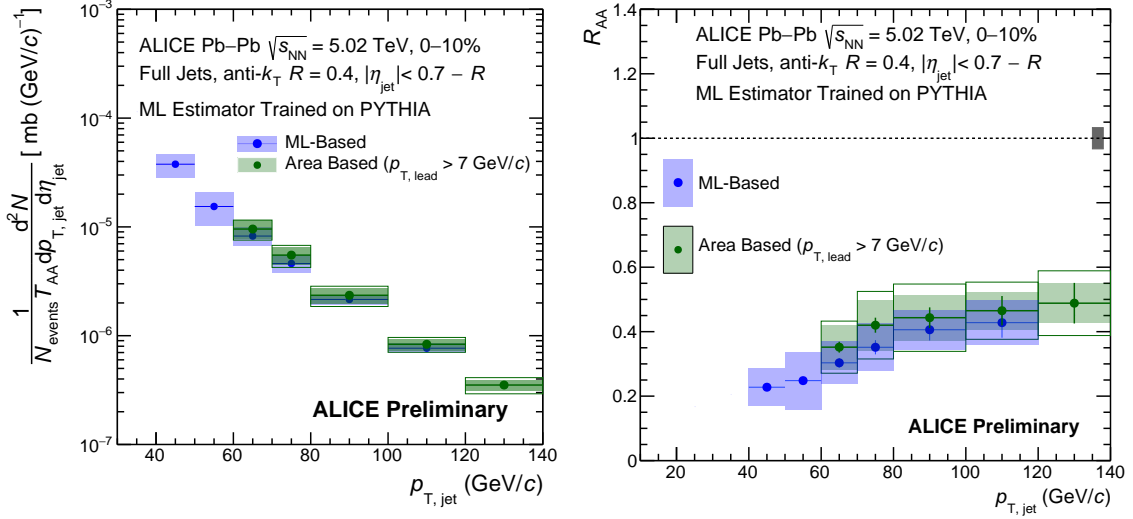


Figure 6.7: Left: The p_T -differential inclusive full jet yield for $R = 0.4$ jets in the 0–10% most central Pb–Pb collisions for results obtained with the ML-based and area-based correction methods. Right: Nuclear modification factor of inclusive full jets as a function of p_T for $R = 0.4$ for 0–10% central Pb–Pb collisions for the ML-based method compared to results obtained with the area-based method where applicable.

unity), consistent with the expectation of jet quenching. Additionally, results for the AB and the ML-based methods are consistent in the kinematic regions of overlap. The ML-based method allows for the extensions of results to lower in p_T with reduced systematic uncertainties, allowing for potential overlap with the lower p_T jets measured at RHIC.

6.8 Fragmentation Bias

Note that the fragmentation studies discussed here occurred before the more sophisticated fragmentation procedures discussed in Section 5.8. Extending the studies discussed in Section 5.8 to full jets is left for future work.

To study the fragmentation bias for the case of full jets, three different modifications to the constituents of the jet and therefore the fragmentation function were studied. These include, (1) Fractional Collinear with $f_{\text{loss}} = 0.1$, (2) Fractional Large Angle with $f_{\text{loss}} = 0.1$, and (3) the BDMPS motivated modification to the fragmentation function. As the BDMPS-motivated modification is not discussed previously in this thesis, it will be the only model described in detail in this section. The studies here are conducted with charged jets due to the relative ease of modifying the embedded distribution and then the effects are extrapolated to full jets, as described below.

For this modification, a BDMPS model for gluon emission [206] was utilized. To do

this a 2D **BDMPS** energy distribution, as defined in Equation 6.5, was sampled in order to determine the angle θ_g and energy ω of the emission. For the distribution in Equation 6.5, $\hat{q} = 2 \text{ GeV}^2/\text{fm}$ ⁴ and $L = 7 \text{ fm}$ with $p_{\text{loss}} = 1.0$. Note that this case is not entirely analogous to the original intention of the model as the “emissions” are applied to tracks, but is suitable for the investigative purposes presented here.

$$P(\theta_g, \omega) = \alpha \omega \theta_g^3 \sqrt{\frac{2\omega}{\hat{q}}} L e^{\frac{-\theta_g^2 \omega^2}{\sqrt{2\omega \hat{q}}}} \quad (6.5)$$

In Figure 6.8 the ratio of modified to unmodified fragmentation functions for both the leading eight particles (left panel) and inclusive particles (right panel) is shown for the different modifications. Comparisons with the leading 8 particles are shown as this reflects the input to the **ML** in the charged jet case. Note that the **BDMPS** model demonstrates the largest modification at low values of z , whereas the fractional models exhibit a larger modification at higher values of z .

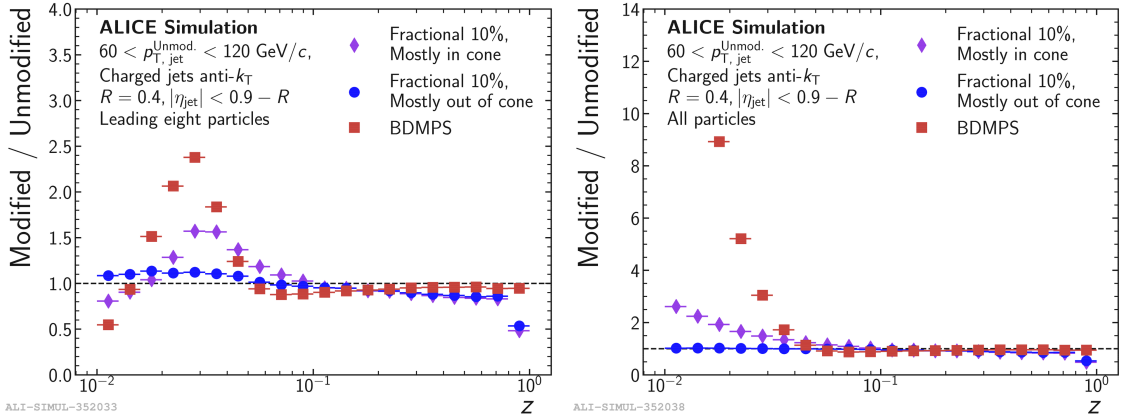


Figure 6.8: Ratio of the modified to unmodified fragmentation functions for the three different modifications presented in Section 6.8. These modifications are shown for both the eight leading particles (left panel) and all particles (right panel).

To determine how such modifications would impact the final result with the **ML** estimator, the training is performed using the modified PYTHIA fragmentation instead of the nominal unmodified PYTHIA and then this result is propagated to the R_{AA} . As these studies were performed with charged jets, a scaling factor was calculated by shifting the R_{AA} in p_T (to account for the difference in energy between charged and full jets) and then taking the ratio between this shifted R_{AA} and the un-shifted R_{AA} . This ratio then was applied to the full jet result to approximate the potential bias. The final bias lines are shown in the left panel of Figure 6.9 in comparison to the nominal result. However, the response matrix in this case was filled using the thermal toy instead of the embedded **MC**. Later investigations revealed that this difference alone causes a bulk of the observed difference

⁴This value corresponds to a plasma with $T = 400 \text{ MeV}$.

between the modified and the nominal fragmentation patterns in the final result. Regardless, in the left panel of Figure 6.9 one sees that the [ML](#) is relatively robust to the explored modifications, with all curves falling within the systematic errors of the [ML](#) measurement, or the systematic errors of the [AB](#) measurement.

6.9 Theoretical Comparisons

The full jet results discussed in this section are compared to theoretical predictions in a two-paneled plot in the left panel Figure 6.9 with the fragmentation bias comparisons in the right panel. Below, the central discriminative features of these models will be described, where more information can be found in Section 2.5.4 and the corresponding references.

These results are compared to the hybrid model [183], which implements medium response via the wake, with [AdS/CFT](#) describing the non-perturbative regime. Additionally shown are comparisons to JEWEL with and without recoils for jets with $|\eta| < 0.9$. Comparisons with the [Linear Boltzman Transport \(LBT\)](#) [138] are also shown for jets in rapidity $|y| < 2.1$. Results are also compared with [SCETg](#) ([SCET](#) with Glauber gluons), including collisional energy loss [207]. In general, the predictions in Figure 6.9 show good agreement with data. Hints of tension at lower jet p_T values are observed with [LBT](#) and JEWEL with recoils.

Future extensions of this work include measurements larger in R and other centrality classes. For full jets, extending in R much past $R = 0.5$ is limited by the η acceptance of the [EMCal](#) ($|\eta| < 0.7 - R$). However, full jets may be able to reach higher in p_T due to the [EMCal](#) jet triggers, which select events with jets at a higher rate. As full jets have a different steepness in vacuum, such studies would allow the influence of this effect to be studied in greater detail.

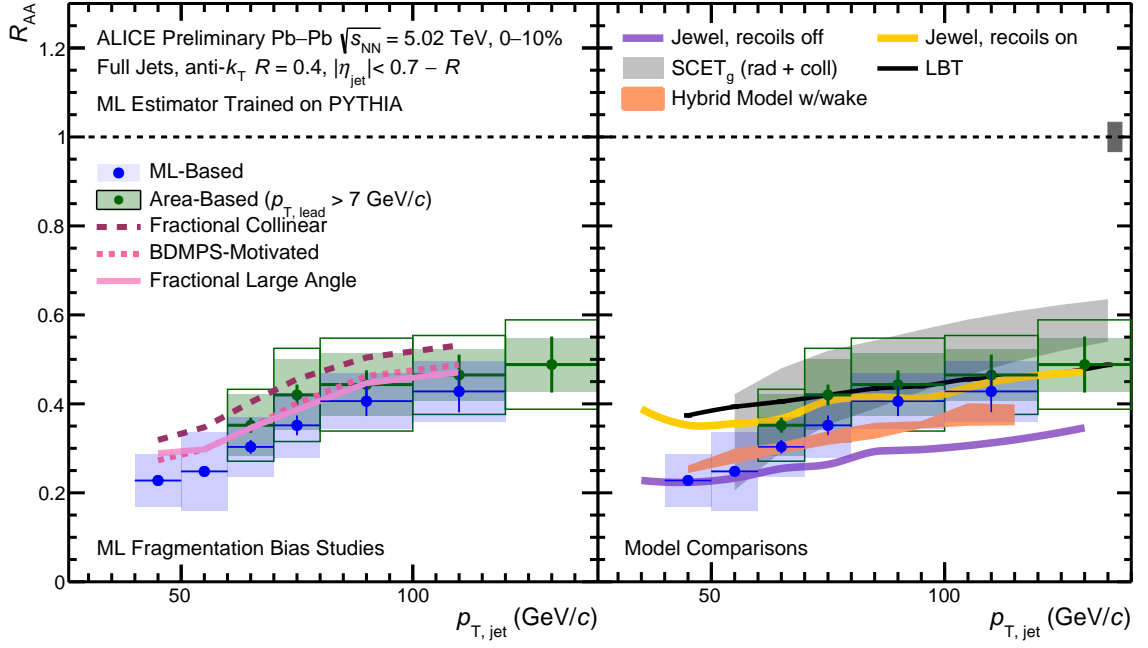


Figure 6.9: The nuclear modification factor with ML-based and area-based corrections are shown for $R = 0.4$ full jets in 0-10% collisions. Left: Comparison of the fragmentation bias curves for a toy model with three different modifications. Right: Comparison to various models. The models that are shown here also do not have this leading track bias requirement.

Chapter 7

Interpreting the R -dependence of the R_{AA}

The studies reported in this section are based on work performed at Yale in collaboration with Mary Zhang (Colby College) and Dr. Laura Havener.

The results shown in Chapter 5 reveal a hint of an R -dependence in measurements of jet suppression for charged-particle jets. Of interest is the relative presence of competing effects that can fix the value of the R_{AA} and its progression with R . Effects that would cause the R_{AA} to decrease with increasing R include the steepness of the jet p_T spectrum in pp, differing quark and gluon fractions, and larger R jets having more effective energy loss sources. These effects are balanced by those that would cause the R_{AA} to increase with increasing R such as the recovery of wide-angle radiation at larger R and the increased influence of the response of the QGP medium. The results presented in this thesis show a decrease in the R_{AA} with increasing R , which indicates a larger influence of effects that can cause such a decrease. However, effects such as the medium response and recovery of wide-angle radiation can still occur, but would not be dominant. This chapter presents work to disentangle the various mechanisms mentioned above through PYTHIA simulations.

7.1 Simulation Setup

In order to investigate the R -dependence of the R_{AA} , PYTHIA simulations for pp collisions at $\sqrt{s} = 5.02$ TeV covering the fiducial acceptance of ALICE ($|\eta_{\text{jet}}| < 0.9$) were used. The PYTHIA events were generated in \hat{p}_T bins of [5, 11, 21, 36, 52, 70, 84, 100, 121, 136, 156, 181, 200, 240, 5020] GeV/ c . As a crosscheck, the R -dependence of the pp jet cross section was compared to ALICE data from pp collisions as shown in the left panel of Figure 7.1. PYTHIA describes the R -dependent trend seen in ALICE pp data,

and therefore serves as a reasonable baseline for these studies. There is some small disagreement for larger R values that is likely due to the modeling of the UE, which has a large effect in this kinematic region.

In order to build in the effect of quenching on this sample, a similar approach is followed to Ref. [123] where quenching is applied in a p_T -dependent manner such that $p_{T,\text{quenched}} = p_{T,\text{unquenched}} - S$. Here the energy loss shift S is given by Equation 7.1, where s' is the shift in p_T when $p_{T,\text{unquenched}} = p_{T_0}$, p_{T_0} is a constant used to parameterize the transverse momentum and α is a constant that influences the shape of the spectrum and is set equal to 0.55 as in Ref. [123]. The values for s' and p_{T_0} are described in the following section. The flavor-dependent case will be discussed in Section 7.3.

$$S = s' \left(\frac{p_{T,\text{unquenched}}}{p_{T_0}} \right)^\alpha \quad (7.1)$$

Note that this approach is a simplistic way to implement quenching that does not reflect the complexities present in data. This is intentional and allows for the isolation of individual effects such as the influence of the steepness in pp (see Section 7.2), flavor (see Section 7.3), detector acceptances, center-of-mass energy, jet definition, etc. The isolation of such effects is relevant for the interpretation of the R -dependence of the R_{AA} results.

7.2 Steepness of the Jet Spectra in pp

Due to the fact that the R_{AA} is the ratio of spectra at a fixed p_T , given the same absolute energy loss, jets that have a steeper initial spectrum will appear more suppressed. It can be seen in both panels of Figure 7.1 that the spectra for larger- R jets are steeper than that of smaller- R jets. In the right panel of Figure 7.1 it can be seen that the jet p_T spectra for charged jets are steeper than for full jets. Therefore, larger R jets may appear to be more suppressed given the same absolute energy loss. The extent to which this influences the R -dependent trends seen in data is investigated in this chapter.

In order to perform this test, jets were quenched in a manner dictated by Equation 7.1 with $s' = 4.25$ and $p_{T_0} = 8.5$ GeV/c. These values were chosen in order to provide the best fits of the ALICE data. Due to the fact that the same parameters were used for each R , the same absolute degree of quenching is applied. Therefore, any R -dependence in suppression effects is entirely due to the R -dependence of the jet p_T spectrum in pp collisions.

The R_{AA} ratios for this flavor-independent quenching model are shown in the gray points in Figure 7.3 as compared to ALICE data. As expected, this model does result in the same qualitative trend as seen in ALICE data where larger- R jets are more suppressed. There is additionally quantitative agreement with the results for $R_{AA}(R = 0.4)/R_{AA}(R = 0.2)$. However, for $R_{AA}(R = 0.6)/R_{AA}(R = 0.2)$ this model slightly underestimates the




Figure 7.1: Left: The jet cross section ratios in PYTHIA simulations as compared to ALICE pp data. Right: PYTHIA jet p_T spectra generated for a variety of R for charged-particle jets and full jets.

trend. This indicates that although the R -dependence in vacuum does have an effect, some R -dependence in the quenching mechanisms is needed to reproduce the ALICE results.

7.3 Flavor-dependent Quenching

One potential mechanism that could result in larger- R jets being more suppressed is differences in quenching effects between quark and gluon jets. Quark and gluon jets are expected to exhibit different amounts of energy loss due to their differing color charges, where gluon jets would experience more energy loss. The fraction of quark jets is expected to be less in larger- R jets, as shown for PYTHIA jets in Figure 7.2. Here, jets are identified as quark or gluon jets based on the distance between the initiating parton and the jet. Larger- R jets may be more suppressed simply due to the fact that there are fewer quark jets and more gluon jets at larger R . This section investigates the influence of these effects.

In order to apply a flavor-dependent quenching, two different S values were applied for quark and gluon jets as written in Equations 7.2 and 7.3. Here, the only difference between the two quenching types is the $\frac{9}{4}$ factor accounting for the ratios of the color factors.

$$S_{\text{quark}} = s' \left(\frac{p_{T, \text{unquenched}}}{p_{T_0}} \right)^\alpha \quad (7.2)$$

$$S_{\text{gluon}} = \left(\frac{9}{4} \right) s' \left(\frac{p_{T, \text{unquenched}}}{p_{T_0}} \right)^\alpha \quad (7.3)$$

The R_{AA} ratios for this flavor-dependent quenching model are shown in the blue points in Figure 7.3 as compared to ALICE data. Note that in this scheme, effects due to the




Figure 7.2: Quark fraction in PYTHIA as a function of R .

steepness of the spectrum in vacuum will still be present, so any effect due solely to flavor-dependence is taken as the difference between the gray and the blue points in Figure 7.3. For smaller R the combination of these two effects shows both a qualitative agreement with the ALICE data, but some level of disagreement emerges for the larger R results.




Figure 7.3: Predictions for R_{AA} ratios for $R = 0.4/R = 0.2$ (left) and for $R = 0.6/R = 0.2$ (right) from the flavor-independent quenching model and the flavor-dependent quenching model as compared to the ALICE data.

Overall, these simulations indicate that some R -dependence in the quenching is needed in order to produce the observed R -dependence of suppression effects. Introducing a flavor-dependent quenching shows a better agreement with ALICE data. However, in order to provide more definitive conclusions additional quenching effects, such as momentum broadening and the response of the QGP medium, should also be implemented.

Chapter 8

Conclusions and Future Work

8.1 Conclusions

This thesis studies [QCD](#) matter pushed to its extremes to form a [Quark-Gluon Plasma \(QGP\)](#), which has not existed in nature since the very beginnings of the universe, by measuring the collisions of relativistic heavy ions with the [ALICE](#) detector. As the lifetime of this hot and dense medium is quite short (10^{-15} fm/c), internally-generated probes are useful in determining the intrinsic and dynamic properties of the [QGP](#) medium. One such probe is narrow cones of particles initiated from the hard-scattering of colliding partons, and subsequent fragmentation and hadronization. When jets traverse the [QGP](#), they lose energy and their substructure becomes modified, referred to as jet quenching. Different jets, with different structures, flavors, kinematics, path lengths through the medium, etc. will all lose energy differently. Due to the increasing amounts of new techniques and available data at both [RHIC](#) and the [LHC](#), this attribute of jets that once presented a challenge is now an opportunity. Differential measurements of jets can now be used to determine how quenching effects depend on the attributes of the jet population as well as differentiating the relative strength of different energy loss mechanisms.

In this thesis, the dependence of jet energy loss on the jet cone size (or radius) was measured at the [ALICE](#) experiment at the [LHC](#). To perform this measurement, a novel [ML](#) algorithm was developed to better correct the jet p_T in comparison to the standard [ALICE](#) approach. New methods were also employed to study the fragmentation bias of this approach, which could also be applied in future contexts. In this measurement, a hint of an R -dependence was observed, where charged jets with a larger radius lose more energy in the [QGP](#). Though generally consistent with predictions from jet quenching models, these results also reveal a tension with models where the response of the [QGP](#) medium plays a large role, especially at lower p_T . These results may arise from wider jets having a more complex substructure and differences in the vacuum jet population. For wider jets to exhibit a greater degree of suppression, these effects must outweigh the recovery of

energy lost to the medium and the medium response. Simulations assuming a flavor- and p_T -dependent energy loss can reproduce an R -dependence in energy loss similar to the observed magnitude, indicating the potential size of these effects in isolation. The degree of energy loss for full jets was also measured at the [ALICE](#) experiment at the [LHC](#) down to very low jet p_T . Similar to charged jets this measurement also indicates a tension at lower jet p_T with models where the response of the [QGP](#) medium plays a large role. Performing a study of the R -dependence of full jets is left for future work.

In summary, the work presented in this thesis represents a valuable contribution to studies of jet quenching in heavy-ion collisions discussed in Chapter 2 and to the [ALICE](#) experiment discussed in Chapter 3. In Chapter 4 a new online calibration algorithm for the [EMCal](#) was discussed, which is currently being employed to ensure reliable [EMCal](#) data in Run 3 of the [LHC](#). In Chapter 5 the most differential measurement of jet suppression ever made by [ALICE](#) and the first-ever experimental measurement of jets in heavy-ion collisions using [ML](#) was presented. In Chapter 6 a complementary measurement using full jets was discussed, measuring down to the lowest full jet p_T ever measured at the [LHC](#). Finally, Chapter 8 presents studies to understand the role of the jet population in the observed jet-quenching effects presented in this thesis. Ultimately, such studies have shed light on the role that the jet radius plays in jet energy loss, and have provided an intriguing basis for future investigations.

8.2 Future Work

In this section future measurements will be discussed that could extend the effort set forth in this thesis to discriminate between various energy loss effects. These will be divided up into three categories: (1) direct extensions of this result, (2) measurements to overcome the survivor bias, and (3) new observables. Though this will be a short and biased perspective, it is intended to provide an overview of potential future avenues of research.

8.2.1 Direct Extensions of this Result

One obvious extension of the results presented in Chapter 6 would be to measure the radial R -dependence of jet suppression for full jets. This would provide an additional probe of the influence of the steepness of the jet spectrum in pp collisions on the final results, as the R -dependence for full jets in pp collisions is slightly stronger. An additional potential extension is to measure the R -dependence of jet suppression to larger- R at [RHIC](#). This would provide an additional probe of the influence of the fraction of quark and gluon jets, where the fraction of quark jets is much larger at [RHIC](#).

Further extensions of this result include more differential measurements of jet suppression that depend not on R , but on other variables such as substructure. Such efforts have already begun, but performing these measurements more differentially in lower p_T regions

would be of interest. The ML-based method pioneered in this thesis could be extended to enable such differential measurements. This will also be aided by the large amount of data that will be collected in Run 3 of the LHC.

8.2.2 Overcoming the Survivor Bias

A focus of jet quenching measurements in the coming years will be to overcome the survivor bias as discussed in Section 2.5. One example of such effects is measurements of heavy-flavor-, photon-, or Z -tagged jets and their substructure, which selects a sample dominated by quark jets and allows for a selection on the initial p_T of the jet. Another possibility would be to use ML to identify the initial jet p_T , allowing jets to be sorted by the degree of quenching they experience [208]. However, the model dependence of such an approach would need to be studied in-depth prior to its application to experimental data. The fragmentation bias studies outlined in this thesis offer one potential avenue for studying the model dependence of ML applications to experimental measurements in heavy-ion collisions.

8.2.3 New Observables

Due in part to the new data collected in Run 3, many new observables will be key to jet-quenching studies in the coming years. For example, measurements of heavy-flavor jets will greatly benefit from the new upgrades of the ALICE ITS. These measurements are critical to determining the mass dependence of the jet quenching effects. In addition, measurements of different regions of the Lund plane in heavy-ion collisions will likely be possible, which could shed light on how the QGP impacts the Dead-Cone effect.

One new observable likely to play a key role in precision measurements of jets and their quenching are Energy-Energy Correlators (EECs) [209]. This new class of observables directly connects experimental measurements to theoretical calculations and offers a natural separation between the perturbative and non-perturbative scales. In addition, such tools can also be used to analyze the energy flow of jets initiated by a charm or beauty quark [210]. In the coming years, such tools will likely also be used to study the modification of the energy flow of jets in heavy-ion collisions and shed insight into QGP properties, such as the resolution length [211].

Appendix A

Additional Checks of the R -dependence Result

In this appendix, additional checks of the R -dependent result discussed in Chapter 5 will be presented. Other checks are included in the main body of this thesis in Section 5.5.6.

A.1 Checking Correlation of Input Features

One common metric in order to select input features for training is to determine the correlation between input features. An additional check of the ML estimator is to look at the 2D correlations between variables, which is shown in Figure A.1, to ensure variables are not strongly correlated with one another. In this figure, all compared variables have only minimal correlations with other input features.




Figure A.1: Visual representation of the correlation of parameters for central (0–10%) $R = 0.6$ jets.

A.2 Influence of a Flat Training Spectrum

The training distribution of jets for the nominal ML estimator is shown in the right panel of Figure 5.5. This training spectrum is mostly flat in order to allow the ML to learn features of both low and high p_T jets in the same relative abundances. In the event that a realistic spectrum was used, the estimator would focus on the features of lower p_T jets, which are in greater abundance due to the steeply falling spectrum, and would not perform well on higher p_T jets. To confirm that such an effect is not occurring in the mostly-flat spectrum used in training, results were compared with those using a completely flat training sample. The new training distribution is shown in the left panel of Figure A.2. A comparison of the performance with the flat and nominal training samples is shown in the right panel of Figure A.2. This comparison indicates that the ML estimator has a similar performance for a completely flat training sample and a mostly-flat training sample, indicating that any deviations from a flat spectrum do not strongly impact the ML algorithm. The unfolded results for $R = 0.6$ jets in central collisions achieved with these two different training samples are shown in Figure A.3. Here, good agreement is seen between the results achieved with a completely flat training sample and the nominal training sample, indicating that

any deviations from a flat spectrum in the training sample do not impact the final result achieved with the ML algorithm.




Figure A.2: (Left): Training distribution for the completely flat training test. (Right): Performance comparison of the flat and nominal training samples.




Figure A.3: (Top): The unfolded result achieved for $R = 0.6$ jets in central collisions using the nominal training sample and a flat training sample. (Bottom): Ratio of the two unfolded distributions in the top panel.

A.3 Embedded MC and Thermal Toy Comparisons

One essential check to perform in this analysis is to show that the results do not strongly depend on the chosen truth definition or the background used in training. For the charged

jet results presented in Chapter 5 of this thesis, the true p_T fraction defined in Equation 5.3 was used as the truth definition and the thermal toy background as discussed in Section 5.2 were utilized. For the case of full jets, the true p_T fraction definition is ambiguous as energy from the PYTHIA and Pb–Pb event can contribute energy to a single calorimeter cluster. Therefore, in this case the truth definition of the geometrically matched detector-level jet p_T in the embedded MC event is used. An important test of the robustness of this method is that the result does not strongly depend on the choice of regression target or background. A comparison of the unfolded result for $R = 0.6$ jets in central collisions can be found in Figure A.4. These two results show good agreement, indicating that the results achieved with the ML-based estimator do not strongly depend on this choice.




Figure A.4: (Top) Comparison of the fully-unfolded results for $R = 0.6$ jets in central collisions using both the thermal background (with the true p_T fraction as the target) and the embedded MC background (with the geometrically-matched detector-level jet p_T as the target) (Bottom): The ratio between these two quantities.

A.4 Investigations on the ATLAS Discrepancy

As mentioned in Section 5.1, the R -dependence of jet suppression effects has been measured at the LHC by the CMS [28] and ATLAS [184] experiments and at RHIC by the STAR [27] experiment. The results presented in this thesis exhibit both quantitative and qualitative differences from the ATLAS results, as shown in Figure A.5. The purpose of this section is to discuss the differences between these two measurements. In order to begin to determine the discrepancy between the results presented in this thesis and the ATLAS results, a natural place to look is at the various experimental parameters which may

differ between the two results. A summary of these differences is included in Table A.1. Each of these differences will be discussed below.




Figure A.5: Comparison of different measurements of the R -dependence of jet suppression by the ALICE (this thesis), ATLAS, and CMS collaborations. Figure courtesy of Prof. Dennis Perepelitsa.

The first difference listed is the type of jet where, for the reasons specified in Section 2.4.2, ALICE measures charged jets while both ATLAS and CMS measure full jets. In Figure 7.1, the steepness of the jet p_T spectra for a variety of R for both charged and full jets is shown. Here it can be seen that, especially at low p_T , charged jets exhibit a greater steepness than that of full jets, which could impact the resulting R -dependence. Though this effect could have an impact on the R -dependence, it is not large enough to produce the qualitative difference between the ALICE and ATLAS data seen here. In addition to differences in pp , the charged and the neutral components of jets may lose energy differently, and could result in a different R -dependence. The optimal way to test this is to perform an R -dependent measurement for the full jet case, which is listed as future work. The difference in observable (R_{AA} vs. R_{CP}) was also tested by making the R_{CP} using ALICE data, but such tests were inconclusive due to the large statistical error bars on the peripheral result.

The next difference listed is the rapidity, where both ATLAS and CMS have significantly larger detector acceptances than the ALICE results. The rapidity is also a factor that could change the steepness of the spectrum, driven by a change in the quark fraction, and therefore the R -dependence. The quark fractions in PYTHIA for both ALICE and ATLAS kinematics are shown in Figure A.6. Note that this comparison also includes many of the differences listed in Table A.1, including the center-of-mass energy, the R , as well as the constituent cut.¹ In comparing the quark fractions, the ALICE kinematics yield a higher quark fraction in PYTHIA for all R . A higher quark fraction in ALICE would not

¹Note, the constituent cuts used here for ATLAS kinematics are not the exact same as the actual constituent cut of 700 MeV, but this difference will be discussed later on in this section.

Table A.1: Values for different experimental parameters for the ALICE, ATLAS, and CMS measurements shown in Figure A.5. Note that for CMS the constituent $p_{T,\min}$ was not studied for this comparison, and is therefore listed as unknown.

Parameter	ALICE	ATLAS	CMS
Jet Type	Ch-particle jets	full jets	full jets
Rapidity	$ \eta < 0.9$	$ \eta < 2.1$	$ \eta < 2.0$
$\sqrt{s_{NN}}$ (TeV)	5.02	2.76	5.02
Observable	R_{AA}	R_{cp}	R_{AA}
R	$R = 0.6$	$R = 0.5$	$R = 0.6$
Constituent $p_{T,\min}$	$p_{T,\text{const}} > 150 \text{ MeV}$	$p_{T,\text{const}} > 700 \text{ MeV}$	Unknown

explain the greater degree of suppression observed for large R , as quark jets are thought to lose less energy in the QGP. Therefore, these kinematic differences do not explain the difference in the ALICE compared to the ATLAS results.




Figure A.6: Quark fraction in both ALICE (left) and ATLAS (right) kinematics.

As discussed above, one difference between the ALICE results presented in this thesis and the ATLAS results is the constituent cut used. ALICE can measure to lower constituent p_T 's of 150 MeV/c whereas ATLAS cannot measure constituents below 700 MeV/c due to detector constraints. In ATLAS this is not an explicit constraint imposed in jet finding (as in the case in ALICE), but rather an implicit constraint. In order to test the effect of such a cutoff, the ALICE analysis was rerun with a higher 700 MeV cutoff. The jet cross section ratios as a result of this test are shown in Figure A.7 both for smaller radii of $R = 0.4$ and larger radii of $R = 0.6$.² Results are compatible between the two minimum constituent p_T cuts, indicating that the ALICE results are robust to the choice of minimum p_T and also that this difference does not account for the difference in these two results. In addition, for the larger R results, the effect of correcting back to 0 MeV/c or 700 MeV/c in the

²The pp analysis was not repeated with the 700 MeV/c cut, which is why the comparison is made using the jet cross section ratios as opposed to the R_{AA} .

unfolding procedure is examined, and is found to have a minimal effect on the jet cross section ratio.




Figure A.7: The jet cross section ratios for $\sigma(R = 0.2)/\sigma(R = 0.4)$ (left) and $\sigma(R = 0.2)/\sigma(R = 0.6)$ for the nominal ($p_{T,\text{const}} > 150 \text{ MeV}/c$) and ATLAS ($p_T > 700 \text{ MeV}/c$). In the right panel, the ATLAS results are shown correcting back to 700 MeV/c and 0 MeV/c in the unfolding procedure.

Overall, none of the investigated differences in measurement parameters can account for the quantitative and qualitative differences shown above. One remaining un-investigated difference is a difference in the background subtraction procedures. Due to the compatibility between the ML-based and area-based results, such a difference must be present between the ATLAS procedure and both ALICE procedures. Unfortunately, as the ATLAS background subtraction is performed at the level of calorimeter towers, a concrete investigation of these differences in experiment would be quite difficult. Therefore, more measurements of jet suppression at the two experiments are needed in order to reconcile these differences.

Appendix B

Analysis Code

In this appendix, details of the internal analysis code used in this thesis are described. Due to its internal nature, some of this information is not publicly accessible. Regardless, such information is included in order to provide guidance for future work to extend or reproduce these results.

B.1 Charged-Particle Jet Analysis Code

The charged-particle jet analysis described in Chapter 6 was processed with the ALICE Lego train system in the `Jets_EMCPbPb` train. The relevant internal trains numbers used for the results in this analysis are included below. The analysis code used to extract jets in trees is located in `AliPhysics` in the analysis task referred to as `AliAnalysisTaskJetExtractor`. The machine learning code used in this analysis as well as the post-processing code are located in a private repository, available for reference upon request. Before we perform the machine learning we store the data files in `pandas` ¹ dataframes, utilizing the HDF5 format ² for storage of these dataframes. Note that the process of extracting and concatenating all of the data from each $p_{T,\text{hard}}$ bin is very computationally expensive and time consuming. For that reason, we utilize a machine with 32 GB of RAM for this process. In order to interface the `pandas` software with the trees that were originally written in `ROOT`, we utilize `root pandas` ³. The code to perform this extraction, the training of the ML, and its application can be found at <https://github.com/raymondEhlers/jet-background-ml>. Code for post-processing and unfolding is available upon request.

$R = 0.2, 0.4, 0.6$, ML Trains

¹<https://pandas.pydata.org/>

²<https://www.hdfgroup.org/solutions/hdf5/>

³<https://pypi.org/project/root-pandas/>

- Embedded MC Training Data set: 6954-6973
- Embedded MC Tracking Systematic: 7400-7419
- LHC15o Data: 6953

$R = 0.2, 0.4, 0.6$, AB Trains

- Embedded MC Training Data set: 7053-7072
- Embedded MC Tracking Systematic: 7460-7479
- LHC15o Data: 7052

The reconstruction efficiency used in this analysis was run in the `Jets_EMCPp_MC` train number 2523 for the area-based analysis and 2566 for the ML-based analysis.

B.2 Full Jet Analysis Code

The full jet analysis described in Chapter 6 was processed with the ALICE Lego train system in the `Jets_EMCPbPb` train, in a manner identical to the charged-particle jet analysis. The relevant internal trains numbers used for the results in this analysis are included below.

- Embedded MC Training Dataset: 5630-5649
- Embedded MC Tracking Systematic: 5762-5779, 5821
- Embedded MC Matching Systematic: 5822 - 5841
- Embedded MC Non Linearity Systematic: 5905 - 5924
- LHC15o Data: 5629

The reconstruction efficiency used in this analysis was run in the `Jets_EMCPp_MC` train number 2136.

Appendix C

Supplemental Figures for the Charged Jet Analysis

Figures shown in this appendix seek to supplement the information provided in Chapter 5. Namely, this appendix includes both jet performance and unfolding plots for all parts of phase space not included in the main body of the text. As a result, this section will include minimal textual explanations.

C.1 Performance




Figure C.1: δp_T distributions for R = 0.4 central (0-10%) jets in various bins of $p_{T,det}$.




Figure C.2: δp_T distributions for R = 0.6 central (0-10%) jets in various bins of $p_{T,det}$.

C.2 Jet Performance Plots




Figure C.3: The JES (as defined in Equations 5.5) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.




Figure C.4: The JER (as defined in Equations 5.6) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.




Figure C.5: The JES (as defined in Equations 5.5) for $R = 0.2$, 0.4 , and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.




Figure C.6: The JER (as defined in Equations 5.6) for $R = 0.2$, 0.4 , and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.




Figure C.7: The JES (as defined in Equations 5.5) for $R = 0.2$, 0.4 , and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.




Figure C.8: The JER (as defined in Equations 5.6) for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.

C.3 Unfolding Plots

C.3.1 Trivial Closure Test




Figure C.9: Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.




Figure C.10: Trivial closure tests for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central ($30\text{-}50\%$) collisions.




Figure C.11: Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central ($60\text{-}80\%$) collisions.

C.3.2 Split MC Test




Figure C.12: Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.




Figure C.13: Trivial closure tests for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.




Figure C.14: Trivial closure test for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.

C.3.3 Stability with the Number of Iterations




Figure C.15: Unfolding stability for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.




Figure C.16: Unfolding stability for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.




Figure C.17: Unfolding stability for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (60-80%) collisions.

C.3.4 Refolded to Raw Distributions




Figure C.18: Refolded to raw ratio for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central ($0\text{-}10\%$) collisions.




Figure C.19: Refolded to raw ratio for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central ($30\text{-}50\%$) collisions.




Figure C.20: Refolded to raw ratio for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central ($60\text{-}80\%$) collisions.

C.3.5 Pearson Correlation Coefficients

In this section the Pearson correlation coefficients, which captures the degree of correlations between bins, are shown for the different bin numbers (as shown in Section 5.7). These correlation coefficients were calculated using the Bayesian iteration used in the nominal result. The desired effect being that bins which are far away from one another do not appear to be correlated.




Figure C.21: Pearson coefficients for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in central (0-10%) collisions.

The Pearson correlation coefficient for different iterations for $R = 0.6$ in 0-10% collisions are shown in Figure C.22. Note that the first 3 bins are not considered in the final measurement.




Figure C.22: Pearson coefficients for $R = 0.6$ jets corrected by the ML-based method in central (0-10%) collisions with iterations 6, 7, 9, 10. Note that results with 8 iterations can be found in Figure C.21.




Figure C.23: Pearson coefficients for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in semi-central (30-50%) collisions.




Figure C.24: Pearson coefficients for $R = 0.2, 0.4$, and 0.6 jets corrected by the ML-based method in peripheral(60-80%) collisions.

C.4 Systematics

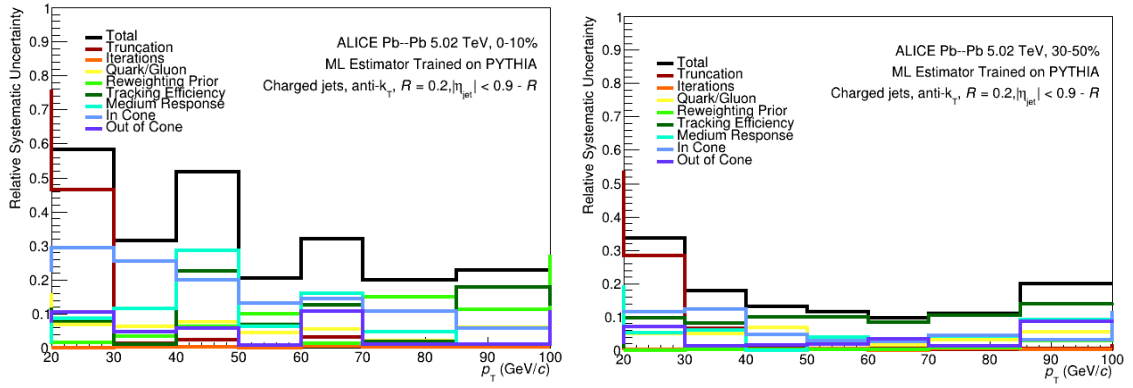


Figure C.25: Relative systematic uncertainties for jet spectrum, $R = 0.2$ (left: 0-10%, right: 30-50%). The total uncertainty, which is the quadratic sum of the individual uncertainties, is also shown.

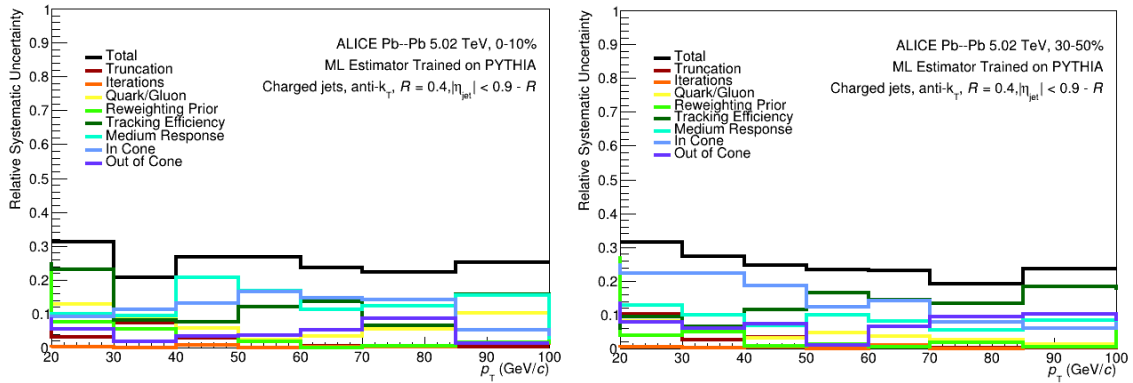


Figure C.26: Relative systematic uncertainties for jet spectrum, $R = 0.4$ (left: 0-10%, right: 30-50%). The total uncertainty, which is the quadratic sum of the individual uncertainties, is also shown.

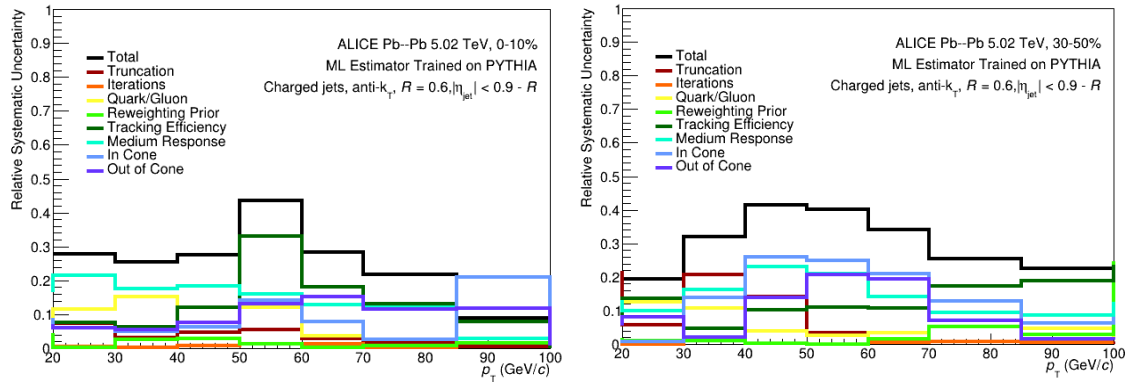


Figure C.27: Relative systematic uncertainties for jet spectrum, $R = 0.6$ (left: 0-10%, right: 30-50%). The total uncertainty, which is the quadratic sum of the individual uncertainties, is also shown.

Acronyms

O^2 Online-Offline. [57](#)

AB area-based. [xv](#), [xvi](#), [67](#), [68](#), [72](#), [74–76](#), [84](#), [85](#), [88](#), [93–95](#), [104–108](#), [110](#), [111](#), [113](#)

ACORDE ALICE Cosmic Ray Detector. [51](#)

ADC Analog-Digital Converter. [56](#)

AdS Anti-de Sitter. [8](#), [9](#), [34](#), [40](#), [99](#), [113](#)

ALICE “A Large Ion Collider Experiment”. [xii](#), [xiii](#), [1](#), [2](#), [14](#), [16](#), [19](#), [20](#), [25](#), [30–32](#), [50](#), [52–57](#), [59](#), [65](#), [68](#), [89](#), [95](#), [102](#), [115–121](#)

AMY Arnold-Moore-Yaffe. [40](#)

AP Altarelli-Parisi. [22](#)

APD avalanche photo diode. [59](#)

ASW Amesto-Salgado-Wiedemann. [40](#)

ATLAS “A Toroidal LHC Apparatus”. [16](#), [33](#), [64](#), [86](#), [99](#)

BDMPS Baier-Dokshitzer-Mueller-Peigne-Schiff. [40](#), [99](#), [111](#), [112](#)

BDMPS-Z Baier-Dokshitzer-Mueller-Peigne-Schiff-Zakharov. [40](#)

BNL Brookhaven National Laboratory. [15](#), [16](#)

CERN Conseil Européen pour la Recherche Nucléaire. [1](#), [15](#), [16](#)

CFT Conformal Field Theory. [8](#), [9](#), [34](#), [40](#), [99](#), [113](#)

CGC Color Glass Condensate. [7](#), [12](#), [18](#)

CMS “Compact Muon Solenoid”. [xii](#), [16](#), [30](#), [31](#), [64](#), [86](#), [99](#)

CNN Convolutional Neural Network. [44](#), [48](#)

CPV Charged Particle Veto. [51](#)

DCal Dijet Calorimeter. [xiii](#), [51](#), [55](#), [56](#)

DGLAP Dokshitzer Gribov Lipatov Altarelli Parisi. [6](#), [7](#), [22](#)

DNN Deep Neural Network. [48](#), [49](#)

EEC Energy-Energy Correlator. [33](#), [121](#)

EIC Electron-Ion Collider. [7](#), [16](#)

EMCal Electromagnetic Calorimeter. [xiii](#), [51](#), [55–58](#), [60](#), [102](#), [108–110](#), [113](#), [120](#)

eRHIC election ring at RHIC. [16](#)

FIT Fast Interaction Trigger. [57](#)

FMD Forward Multiplicity Detector. [51](#)

GAN Generative Adversarial Network. [44](#), [48](#)

GEM Gas Electron Multiplier. [57](#)

GLV Gyulassy-Levai-Vitev. [40](#)

GPU Graphics Processing Unit. [57](#)

HEP High-Energy Physics. [28](#), [44](#), [48](#), [49](#), [77](#)

HMPID High Momentum PID. [51](#), [54](#)

HT Higher Twist. [40](#)

IG information gain. [46](#)

IP Interaction Point. [16](#)

IRC Infrared and Collinear. [24](#)

ITS Inner Tracking System. [51–54](#), [57](#), [121](#)

JER Jet Energy Resolution. [76](#), [77](#), [108](#)

JES Jet Energy Scale. [76](#), [77](#), [108](#)

LBT Linear Boltzman Transport. [41](#), [99](#), [100](#), [113](#)

LEIR Low-Energy Ion Ring. [17](#)

LHC Large Hadron Collider. [1](#), [15–17](#), [20](#), [25](#), [27](#), [33](#), [39](#), [48](#), [50](#), [53](#), [56](#), [58](#), [59](#), [63–65](#), [67](#), [95](#), [119–121](#)

LHCb “LHC Beauty”. [16](#), [49](#)

LINAC Linear Accelerator. [17](#)

LO Leading Order. [6](#), [26](#), [31](#)

MC Monte Carlo. [11](#), [12](#), [18](#), [28](#), [40](#), [49](#), [66](#), [71](#), [78](#), [80](#), [82](#), [83](#), [90](#), [99](#), [103](#), [105](#), [109](#), [112](#)

MFT Muon Forward Tracker. [57](#)

ML machine learning. [xvi](#), [1](#), [2](#), [42–46](#), [48](#), [49](#), [60](#), [66](#), [69–74](#), [76](#), [84–86](#), [88](#), [89](#), [94](#), [95](#), [104–108](#), [110](#), [112](#), [113](#), [119–121](#)

MPI multi-parton interaction. [28](#)

MS Muon Spectrometer. [51](#)

NLL Next to Leading Log. [30](#)

NLO Next to Leading Order. [xii](#), [6](#), [30](#), [31](#)

NN Neural Network. [43](#), [46–48](#), [73](#), [107](#)

NNLO Next to Next to Leading Order. [xi](#), [6](#), [8](#)

NP non-perturbative. [23](#), [27](#), [30](#), [32](#), [78](#)

npQCD non-perturbative QCD. [8](#), [9](#), [23](#)

PDF Parton Distribution Function. [xi](#), [6–8](#)

PHOS Photon Spectrometer. [xiii](#), [51](#), [55](#), [56](#)

PID Particle Identification. [xiii](#), [50](#), [51](#), [54](#), [55](#)

PMD Photon Multiplicity Detector. [51](#)

pp proton-proton. [xii](#), [xv](#), [6](#), [21](#), [26](#), [30](#), [33–36](#), [39](#), [92](#), [93](#), [96](#), [97](#), [99](#), [126](#)

pQCD Perturbative QCD. [20](#), [23](#), [26](#), [32](#), [40](#)

PS Proton Synchrotron. [17](#)

PSB Proton Synchrotron Booster. [17](#)

QA Quality Assurance. [60](#)

QCD Quantum Chromodynamics. [xi](#), [4–10](#), [12](#), [16](#), [19–21](#), [23](#), [25](#), [26](#), [28](#), [29](#), [31](#), [119](#)

QED Quantum Electrodynamics. [4](#), [5](#)

QGP Quark-Gluon Plasma. [1](#), [2](#), [9](#), [10](#), [12–14](#), [18–21](#), [23](#), [26](#), [33–35](#), [37](#), [39](#), [41](#), [95](#), [99](#), [115](#), [118–121](#)

RHIC Relativistic Heavy-Ion Collider. [15](#), [16](#), [33](#), [39](#), [63](#), [64](#), [95](#), [111](#), [119](#), [120](#)

RNN Recurrent Neural Network. [44](#), [48](#)

ROC Receiver Operating Characteristic. [44](#)

SCET Soft Collinear Effective Theory. [9](#), [113](#)

SD Soft Drop. [32](#), [38](#)

SDD Silicon Drift Detector. [52](#)

SPD Silicon Pixel Detector. [52](#), [53](#), [65](#)

sPHENIX super Pioneering High Energy Nuclear Interaction eXperiment. [16](#)

SPS Super Proton Synchrotron. [17](#)

SSD Silicon Strip Detector. [52](#)

STAR Solenoid Tracker at RHIC. [16](#), [64](#)

SVD Single Value Decomposition. [79](#)

TOF Time of Flight. [51](#), [53](#), [54](#)

TPC Time Projection Chamber. [51–54](#), [57](#), [60](#), [67](#)

TRD Transition Radiation Detector. [51](#), [53](#), [59](#), [61](#)

UE Underlying Event. [25](#), [27](#), [28](#), [30](#), [67](#), [78](#), [116](#)

ZDC Zero Degree Calorimeter. [14](#), [51](#)

Bibliography

- [1] Wikimedia Commons. Standard model of elementary particles. https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg, 2019. [xi](#), [4](#)
- [2] R. L. Workman et al. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2022(8):083C01, 08 2022. 083C01. [xi](#), [6](#), [7](#)
- [3] Richard D. Ball et al. The path to proton structure at 1% accuracy. *Eur. Phys. J. C*, 82(5):428, 2022. [xi](#), [8](#)
- [4] Antonin Maire. Phase diagram of QCD matter : Quark-Gluon Plasma. <https://cds.cern.ch/record/2025215>, 2015. General Photo. [xi](#), [10](#)
- [5] A. Bazavov et al. Equation of state in (2+1)-flavor QCD. *Phys. Rev. D*, 90:094503, 2014. [xi](#), [10](#), [12](#)
- [6] The ALICE experiment – A journey through QCD. 11 2022. [xi](#), [xiii](#), [11](#), [15](#), [19](#), [20](#), [55](#)
- [7] Bjoern Schenke, Prithwish Tribedy, and Raju Venugopalan. Fluctuating Glasma initial conditions and flow in heavy ion collisions. *Phys. Rev. Lett.*, 108:252301, 2012. [xi](#), [12](#)
- [8] Centrality determination in heavy ion collisions. <https://cds.cern.ch/record/2636623>, 2018. [xi](#), [15](#)
- [9] Ewa Lopienska. The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022. 2022. General Photo. [xii](#), [17](#)
- [10] Zhi Qiu. *Event-by-Event Hydrodynamic Simulations for Relativistic Heavy-ion Collisions*. PhD thesis, Ohio State U., Columbus, 2013. [xii](#), [19](#)
- [11] André H. Hoang. The Top Mass: Interpretation and Theoretical Uncertainties. In *7th International Workshop on Top Quark Physics*, 12 2014. [xii](#), [22](#)
- [12] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti- k_t jet clustering algorithm. *JHEP*, 04:063, 2008. [xii](#), [25](#), [67](#)

- [13] B. Webber. Parton shower Monte Carlo event generators. *Scholarpedia*, 6(12):10662, 2011. revision #128236. [xii](#), [29](#)
- [14] Shreyasi Acharya et al. Measurements of inclusive jet spectra in pp and central Pb–Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Rev. C*, 101(3):034911, 2020. [xii](#), [xiii](#), [xvi](#), [30](#), [31](#), [53](#), [63](#), [64](#), [68](#), [90](#), [95](#), [104](#), [109](#), [110](#)
- [15] Zhong-Bo Kang, Felix Ringer, and Ivan Vitev. The semi-inclusive jet function in SCET and small radius resummation for inclusive jet production. *JHEP*, 10:125, 2016. [xii](#), [30](#), [31](#)
- [16] Xiaohui Liu, Sven-Olaf Moch, and Felix Ringer. Threshold and jet radius joint resummation for single-inclusive jet production. *Phys. Rev. Lett.*, 119:212001, Nov 2017.
- [17] Xiaohui Liu, Sven-Olaf Moch, and Felix Ringer. Phenomenology of single-inclusive jet production with jet radius and threshold resummation. *Phys. Rev. D*, 97:056026, Mar 2018. [xii](#), [30](#), [31](#)
- [18] Albert M Sirunyan et al. Dependence of inclusive jet production on the anti- k_T distance parameter in pp collisions at $\sqrt{s} = 13$ TeV. *JHEP*, 12:082, 2020. [xii](#), [27](#), [30](#), [31](#), [64](#), [99](#)
- [19] Stefano Frixione, Paolo Nason, and Carlo Oleari. Matching NLO QCD computations with Parton Shower simulations: the POWHEG method. *JHEP*, 11:070, 2007. [xii](#), [30](#), [31](#)
- [20] Physics Preliminary Summary: Measurement of the primary Lund plane density in pp collisions at $\sqrt{s} = 13$ TeV with ALICE. <https://cds.cern.ch/record/2759456>, 2021. [xii](#), [32](#), [33](#)
- [21] Georges Aad et al. Measurement of the Lund Jet Plane Using Charged Particles in 13 TeV Proton-Proton Collisions with the ATLAS Detector. *Phys. Rev. Lett.*, 124(22):222002, 2020. [xii](#), [33](#)
- [22] David d’Enterria. Jet quenching. *Landolt-Bornstein*, 23:471, 2010. [xii](#), [34](#), [36](#), [37](#), [40](#)
- [23] Arturo Tauro. ALICE Schematics. General Photo, 2017. [xiii](#), [51](#)
- [24] J. Alme et. al. The ALICE TPC, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 622(1):316–367, oct 2010. [xiii](#), [52](#), [53](#)
- [25] Betty Bezverkhny Abelev et al. Performance of the ALICE Experiment at the CERN LHC. *Int. J. Mod. Phys. A*, 29:1430044, 2014. [xiii](#), [53](#)

- [26] Performance of the ALICE Electromagnetic Calorimeter. 9 2022. [xiii](#), [55](#), [56](#), [58](#)
- [27] J. Adam et al. Measurement of inclusive charged-particle jet production in Au+Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV. *Phys. Rev. C*, 102(5):054913, 2020. [xiii](#), [63](#), [64](#), [95](#), [125](#)
- [28] Albert M Sirunyan et al. First measurement of large area jet transverse momentum spectra in heavy-ion collisions. *JHEP*, 05:284, 2021. [xiii](#), [63](#), [64](#), [125](#)
- [29] Morad Aaboud et al. Measurement of the nuclear modification factor for inclusive jets in Pb+Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV with the ATLAS detector. *Phys. Lett. B*, 790:108–128, 2019. [xiii](#), [63](#), [64](#)
- [30] Morad Aaboud et al. Measurement of jet fragmentation in Pb+Pb and *pp* collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV with the ATLAS detector. *Phys. Rev. C*, 98(2):024908, 2018. [xv](#), [37](#), [84](#), [86](#), [87](#)
- [31] Albert M Sirunyan et al. Observation of medium-induced modifications of jet fragmentation in Pb–Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV Using isolated photon-tagged jets. *Phys. Rev. Lett.*, 121(24):242301, 2018. [xv](#), [37](#), [84](#), [86](#), [87](#)
- [32] Tatsumi Aoyama, Masashi Hayakawa, Toichiro Kinoshita, and Makiko Nio. Tenth-order electron anomalous magnetic moment: Contribution of diagrams without closed lepton loops. *Phys. Rev. D*, 91:033006, Feb 2015. [3](#)
- [33] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Addison-Wesley, Reading, USA, 1995. [5](#), [6](#)
- [34] Jian-Wei Qiu, Felix Ringer, Nobuo Sato, and Pia Zurita. Factorization of jet cross sections in heavy-ion collisions. *Phys. Rev. Lett.*, 122(25):252301, 2019. [6](#), [41](#), [65](#), [100](#)
- [35] Yuri L. Dokshitzer. Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics. *Sov. Phys. JETP*, 46:641–653, 1977. [7](#), [22](#)
- [36] V. N. Gribov and L. N. Lipatov. Deep inelastic $e p$ scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972. [22](#)
- [37] Guido Altarelli and G. Parisi. Asymptotic Freedom in Parton Language. *Nucl. Phys. B*, 126:298–318, 1977. [7](#), [22](#)
- [38] Edmond Iancu and Raju Venugopalan. *The Color glass condensate and high-energy scattering in QCD*, pages 249–3363. 3 2003. [7](#)
- [39] Jorge Casalderrey-Solana, Hong Liu, David Mateos, Krishna Rajagopal, and Urs Achim Wiedemann. *Gauge/String Duality, Hot QCD and Heavy Ion Collisions*. Cambridge University Press, 2014. [9](#)

- [40] Christian W. Bauer, Sean Fleming, Dan Pirjol, and Iain W. Stewart. An Effective field theory for collinear and soft gluons: Heavy to light decays. *Phys. Rev. D*, 63:114020, 2001. [9](#)
- [41] Eemeli Annala, Tyler Gorda, Alekski Kurkela, Joonas Nättilä, and Alekski Vuorinen. Evidence for quark-matter cores in massive neutron stars. *Nature Phys.*, 16(9):907–910, 2020. [9](#)
- [42] Heng-Tong Ding, Frithjof Karsch, and Swagato Mukherjee. Thermodynamics of strong-interaction matter from Lattice QCD. *Int. J. Mod. Phys. E*, 24(10):1530007, 2015. [10](#)
- [43] Shinji Ejiri. Canonical partition function and finite density phase transition in lattice qcd. *Phys. Rev. D*, 78:074507, Oct 2008. [10](#)
- [44] Michael L. Miller, Klaus Reygers, Stephen J. Sanders, and Peter Steinberg. Glauber modeling in high energy nuclear collisions. *Ann. Rev. Nucl. Part. Sci.*, 57:205–243, 2007. [11](#), [14](#)
- [45] Björn Schenke, Sangyong Jeon, and Charles Gale. Elliptic and triangular flow in event-by-event $D = 3 + 1$ viscous hydrodynamics. *Phys. Rev. Lett.*, 106:042301, Jan 2011. [12](#)
- [46] Chun Shen, Zhi Qiu, Huichao Song, Jonah Bernhard, Steffen Bass, and Ulrich Heinz. The iEBE-VISHNU code package for relativistic heavy-ion collisions. *Comput. Phys. Commun.*, 199:61–85, 2016. [12](#)
- [47] Long-Gang Pang, Hannah Petersen, and Xin-Nian Wang. Pseudorapidity distribution and decorrelation of anisotropic flow within the open-computing-language implementation clvisc hydrodynamics. *Phys. Rev. C*, 97:064918, Jun 2018. [12](#)
- [48] Zi-Wei Lin, Che Ming Ko, Bao-An Li, Bin Zhang, and Subrata Pal. Multiphase transport model for relativistic heavy ion collisions. *Phys. Rev. C*, 72:064901, Dec 2005. [12](#)
- [49] J. Cleymans and K. Redlich. Unified description of freeze-out parameters in relativistic heavy ion collisions. *Phys. Rev. Lett.*, 81:5284–5286, Dec 1998. [13](#)
- [50] Anton Andronic, Peter Braun-Munzinger, Krzysztof Redlich, and Johanna Stachel. Decoding the phase structure of QCD via particle production at high energy. *Nature*, 561(7723):321–330, 2018. [13](#)
- [51] James L. Nagle and William A. Zajc. Small System Collectivity in Relativistic Hadronic and Nuclear Collisions. *Ann. Rev. Nucl. Part. Sci.*, 68:211–235, 2018. [14](#), [21](#)
- [52] A. Accardi et al. Electron Ion Collider: The Next QCD Frontier: Understanding the glue that binds us all. *Eur. Phys. J. A*, 52(9):268, 2016. [16](#)

- [53] Linear accelerator 3. <https://cds.cern.ch/record/1997426>, 2012. [17](#)
- [54] The Low Energy Ion Ring. <https://cds.cern.ch/record/1997352>, 2012. [17](#)
- [55] Jean-Paul Burnet et al. Fifty years of the CERN Proton Synchrotron : Volume 2. 8 2013. [17](#)
- [56] The Super Proton Synchrotron. <https://cds.cern.ch/record/1997188>, 2012. [17](#)
- [57] J. Vollaie et al. *Linac4 design report*, volume 6/2020 of *CERN Yellow Reports: Monographs*. CERN, Geneva, 9 2020. [17](#)
- [58] The Proton Synchrotron Booster. <https://cds.cern.ch/record/1997372>, 2012. [17](#)
- [59] Betty Bezverkhny Abelev et al. Multi-strange baryon production at mid-rapidity in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 728:216–227, 2014. [Erratum: *Phys.Lett.B* 734, 409–410 (2014)]. [18](#)
- [60] S. Voloshin and Y. Zhang. Flow study in relativistic nuclear collisions by Fourier expansion of Azimuthal particle distributions. *Z. Phys. C*, 70:665–672, 1996. [18](#)
- [61] L. Adamczyk et al. Global Λ hyperon polarization in nuclear collisions: evidence for the most vortical fluid. *Nature*, 548:62–65, 2017. [19](#)
- [62] Serguei Chatrchyan et al. Evidence of b-Jet Quenching in PbPb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Rev. Lett.*, 113(13):132301, 2014. [Erratum: *Phys.Rev.Lett.* 115, 029903 (2015)]. [20](#)
- [63] F. Karsch, M. T. Mehr, and H. Satz. Color Screening and Deconfinement for Bound States of Heavy Quarks. *Z. Phys. C*, 37:617, 1988. [20](#)
- [64] A. Andronic et al. Heavy-flavour and quarkonium production in the LHC era: from proton-proton to heavy-ion collisions. *Eur. Phys. J. C*, 76(3):107, 2016. [21](#)
- [65] Albert M Sirunyan et al. Measurement of nuclear modification factors of $\Upsilon(1S)$, $\Upsilon(2S)$, and $\Upsilon(3S)$ mesons in PbPb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Lett. B*, 790:270–293, 2019. [21](#)
- [66] Gavin P. Salam. Towards Jetography. *Eur. Phys. J. C*, 67:637–686, 2010. [24](#)
- [67] T. Kinoshita. Mass singularities of Feynman amplitudes. *J. Math. Phys.*, 3:650–677, 1962. [24](#)
- [68] Gavin P Salam and Grégory Soyez. A practical seedless infrared-safe cone jet algorithm. *Journal of High Energy Physics*, 2007(05):086–086, may 2007. [24](#)
- [69] Stephen D. Ellis and Davison E. Soper. Successive combination jet algorithm for hadron collisions. *Phys. Rev. D*, 48:3160–3166, Oct 1993. [25](#)

- [70] Yu.L Dokshitzer, G.D Leder, S Moretti, and B.R Webber. Better jet clustering algorithms. *Journal of High Energy Physics*, 1997(08):001–001, aug 1997. [25](#)
- [71] Georges Aad et al. Measurement of soft-drop jet observables in pp collisions with the ATLAS detector at $\sqrt{s}=13$ TeV. *Phys. Rev. D*, 101(5):052007, 2020. [25](#)
- [72] Hsi-Ming Chang, Massimiliano Procura, Jesse Thaler, and Wouter J. Waalewijn. Calculating Track-Based Observables for the LHC. *Phys. Rev. Lett.*, 111:102002, 2013. [26](#)
- [73] Yibei Li, Ian Moulton, Solange Schrijnder van Velzen, Wouter J. Waalewijn, and Hua Xing Zhu. Extending Precision Perturbative QCD with Track Functions. *Phys. Rev. Lett.*, 128(18):182001, 2022. [26](#)
- [74] F. Aversa, P. Chiappetta, Mario Greco, and J. P. Guillet. QCD Corrections to Parton-Parton Scattering Processes. *Nucl. Phys. B*, 327:105, 1989. [26](#)
- [75] Daniel de Florian and Werner Vogelsang. Resummed cross-section for jet production at hadron colliders. *Phys. Rev. D*, 76:074031, 2007.
- [76] Mrinal Dasgupta, Lorenzo Magnea, and Gavin P. Salam. Non-perturbative QCD effects in jets at hadron colliders. *JHEP*, 02:055, 2008. [26](#), [27](#)
- [77] M. Beneke. Renormalons. *Phys. Rept.*, 317:1–142, 1999. [27](#)
- [78] Bo Andersson, G. Gustafson, G. Ingelman, and T. Sjostrand. Parton Fragmentation and String Dynamics. *Phys. Rept.*, 97:31–145, 1983. [28](#)
- [79] Bryan R Webber. A QCD model for jet fragmentation including soft gluon interference. *Nucl. Phys. B*, 238:492–528, 1984. [28](#)
- [80] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour, and B. R. Webber. HERWIG 6.5 release note. 10 2002. [29](#)
- [81] M. Bahr et al. Herwig++ Physics and Manual. *Eur. Phys. J. C*, 58:639–707, 2008. [29](#)
- [82] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. PYTHIA 6.4 Physics and Manual. *JHEP*, 05:026, 2006. [30](#)
- [83] Torbjörn Sjöstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159–177, 2015. [30](#), [66](#), [67](#), [89](#)
- [84] T. Gleisberg, Stefan. Hoeche, F. Krauss, M. Schonherr, S. Schumann, F. Siegert, and J. Winter. Event generation with SHERPA 1.1. *JHEP*, 02:007, 2009. [30](#)

- [85] Bo Andersson, Gosta Gustafson, Leif Lonnblad, and Ulf Pettersson. Coherence Effects in Deep Inelastic Scattering. *Z. Phys. C*, 43:625, 1989. [31](#)
- [86] Frédéric A. Dreyer, Gavin P. Salam, and Grégory Soyez. The Lund Jet Plane. *JHEP*, 12:064, 2018. [31](#), [32](#)
- [87] Andrew J. Larkoski, Simone Marzani, Gregory Soyez, and Jesse Thaler. Soft Drop. *JHEP*, 05:146, 2014. [32](#)
- [88] Yacine Mehtar-Tani, Alba Soto-Ontoso, and Konrad Tywoniuk. Dynamical grooming of QCD jets. *Phys. Rev. D*, 101(3):034004, 2020. [32](#)
- [89] Measurements of the groomed jet radius and momentum splitting fraction with the soft drop and dynamical grooming algorithms in pp collisions at $\sqrt{s} = 5.02$ TeV. 4 2022. [32](#), [33](#)
- [90] Raghav Kunnawalkam Elayavalli. Measurement of splittings along a jet shower in $\sqrt{s} = 200$ GeV pp collisions at STAR. *SciPost Phys. Proc.*, 8:059, 2022. [33](#)
- [91] Raghav Kunnawalkam Elayavalli. Exploring and exploiting various regimes within the jet shower. *Rev. Mex. Fis. Suppl.*, 3(4):040903, 2022. [33](#)
- [92] Albert M Sirunyan et al. Measurement of the Splitting Function in pp and Pb-Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Rev. Lett.*, 120(14):142302, 2018. [33](#), [38](#)
- [93] S. Acharya et al. Direct observation of the dead-cone effect in quantum chromodynamics. *Nature*, 605(7910):440–446, 2022. [Erratum: Nature 607, E22 (2022)]. [33](#)
- [94] J. D. Bjorken. Energy Loss of Energetic Partons in Quark - Gluon Plasma: Possible Extinction of High p(t) Jets in Hadron - Hadron Collisions. 8 1982. [34](#), [36](#)
- [95] Jorge Casalderrey-Solana, Doga Can Gulhan, José Guilherme Milhano, Daniel Pablos, and Krishna Rajagopal. A hybrid strong/weak coupling approach to jet quenching. *JHEP*, 10:019, 2014. [Erratum: JHEP09,175(2015)]. [34](#), [35](#), [36](#), [40](#), [99](#)
- [96] Hong Liu, Krishna Rajagopal, and Urs Achim Wiedemann. Wilson loops in heavy ion collisions and their calculation in ads/cft. *Journal of High Energy Physics*, 2007(03):066–066, Mar 2007. [34](#)
- [97] Miklos Gyulassy and Michael Plumer. Jet Quenching in Dense Matter. *Phys. Lett. B*, 243:432–438, 1990. [35](#)
- [98] Francesco D’Eramo, Mindaugas Lekaveckas, Hong Liu, and Krishna Rajagopal. Momentum Broadening in Weakly Coupled Quark-Gluon Plasma (with a view to finding the quasiparticles within liquid quark-gluon plasma). *JHEP*, 05:031, 2013. [35](#)

- [99] Madhura N. Phadke, Lifford Pinto, Oluwafemi Alabi, Jonathan Harter, Russell M. Taylor II, Xunlei Wu, Hannah Petersen, Steffen A. Bass, and Christopher G. Healey. Exploring ensemble visualization. In Pak Chung Wong, David L. Kao, Ming C. Hao, Chaomei Chen, Robert Kosara, Mark A. Livingston, Jinah Park, and Ian Roberts, editors, *Visualization and Data Analysis 2012*, volume 8294, page 82940B. International Society for Optics and Photonics, SPIE, 2012. [35](#)
- [100] Francesco D’Eramo, Krishna Rajagopal, and Yi Yin. Molière scattering in quark-gluon plasma: finding point-like scatterers in a liquid. *JHEP*, 01:172, 2019. [35](#)
- [101] João Barata, Yacine Mehtar-Tani, Alba Soto-Ontoso, and Konrad Tywoniuk. Re-visiting transverse momentum broadening in dense QCD media. *Phys. Rev. D*, 104(5):054047, 2021. [35](#)
- [102] Leticia Cunqueiro and Anne M. Sickles. Studying the QGP with Jets at the LHC and RHIC. *Prog. Part. Nucl. Phys.*, 124:103940, 2022. [36](#)
- [103] Constantin Loizides, Jason Kamin, and David d’Enterria. Improved Monte Carlo Glauber predictions at present and future nuclear colliders. *Phys. Rev.*, C97(5):054910, 2018. [erratum: *Phys. Rev.C*99,no.1,019901(2019)]. [37](#), [90](#), [92](#)
- [104] Measurement of substructure-dependent jet suppression in Pb+Pb collisions at 5.02 TeV with the ATLAS detector. 11 2022. [37](#), [98](#)
- [105] Measurements of the suppression and correlations of dijets in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. 5 2022. [37](#)
- [106] Serguei Chatrchyan et al. Measurement of Jet Fragmentation in PbPb and pp Collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Rev. C*, 90(2):024908, 2014. [37](#), [84](#)
- [107] Georges Aad et al. Measurement of angular and momentum distributions of charged particles within and around jets in Pb+Pb and *pp* collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector. *Phys. Rev. C*, 100(6):064901, 2019. [Erratum: *Phys.Rev.C* 101, 059903 (2020)]. [38](#)
- [108] Serguei Chatrchyan et al. Modification of Jet Shapes in PbPb Collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 730:243–263, 2014. [38](#)
- [109] Albert M Sirunyan et al. Jet properties in PbPb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *JHEP*, 05:006, 2018. [38](#), [63](#), [85](#), [98](#)
- [110] Albert M Sirunyan et al. Jet Shapes of Isolated Photon-Tagged Jets in Pb-Pb and pp Collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Rev. Lett.*, 122(15):152001, 2019. [38](#)
- [111] Albert M Sirunyan et al. Studies of charm quark diffusion inside jets using PbPb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV. *Phys. Rev. Lett.*, 125(10):102001, 2020. [38](#)

- [112] Andrew J. Larkoski, Jesse Thaler, and Wouter J. Waalewijn. Gaining (Mutual) Information about Quark/Gluon Discrimination. *JHEP*, 11:129, 2014. [38](#)
- [113] Shreyasi Acharya et al. Medium modification of the shape of small-radius jets in central Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. *JHEP*, 10:139, 2018. [38](#)
- [114] S. Acharya et al. First measurement of jet mass in Pb-Pb and p-Pb collisions at the LHC. *Phys. Lett. B*, 776:249–264, 2018. [38](#)
- [115] Shreyasi Acharya et al. Measurements of the groomed and ungroomed jet angularities in pp collisions at $\sqrt{s} = 5.02$ TeV. *JHEP*, 05:061, 2022. [38](#)
- [116] Shreyasi Acharya et al. Exploration of jet substructure using iterative declustering in pp and Pb-Pb collisions at LHC energies. *Phys. Lett. B*, 802:135227, 2020. [38](#)
- [117] Shreyasi Acharya et al. Measurement of the groomed jet radius and momentum splitting fraction in pp and Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. *Phys. Rev. Lett.*, 128(10):102001, 2022. [38](#), [98](#)
- [118] Shreyasi Acharya et al. First measurements of N-subjettiness in central Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. *JHEP*, 10:003, 2021. [38](#)
- [119] Measurement of inclusive and leading subjet fragmentation in pp and Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. 4 2022. [39](#)
- [120] Derek Anderson. Measurement of medium-induced modification of jet yield and acoplanarity using semi-inclusive γ_{dir} +jet and π^0 +jet distributions in $p+p$ and central Au+Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV by STAR. 12 2022. [39](#)
- [121] Z. Hulcher, D. Pablos, and K. Rajagopal. Sensitivity of jet observables to the presence of quasi-particles in QGP. In *29th International Conference on Ultra-relativistic Nucleus-Nucleus Collisions*, 8 2022. [39](#)
- [122] Hannah Bossi. R -dependence of inclusive jet suppression and groomed jet splittings in heavy-ion collisions with ALICE. In *29th International Conference on Ultra-relativistic Nucleus-Nucleus Collisions*, 8 2022. [39](#), [63](#), [92](#)
- [123] Martin Spousta and Brian Cole. Interpreting single jet measurements in Pb-Pb collisions at the LHC. *Eur. Phys. J.*, C76(2):50, 2016. [39](#), [116](#)
- [124] Jasmine Brewer, José Guilherme Milhano, and Jesse Thaler. Sorting out quenched jets. *Phys. Rev. Lett.*, 122(22):222301, 2019. [39](#)
- [125] Yi-Lun Du, Daniel Pablos, and Konrad Tywoniuk. Deep learning jet modifications in heavy-ion collisions. *JHEP*, 21:206, 2020. [48](#), [49](#)
- [126] Megan Connors, Christine Nattrass, Rosi Reed, and Sevil Salur. Jet measurements in heavy ion physics. *Rev. Mod. Phys.*, 90:025005, 2018. [39](#)

- [127] Jasmine Brewer, Quinn Brodsky, and Krishna Rajagopal. Disentangling jet modification in jet simulations and in Z+jet data. *JHEP*, 02:175, 2022. [39](#)
- [128] R. Baier, Yuri L. Dokshitzer, Alfred H. Mueller, S. Peigne, and D. Schiff. Radiative energy loss of high-energy quarks and gluons in a finite volume quark - gluon plasma. *Nucl. Phys. B*, 483:291–320, 1997. [40](#)
- [129] B. G. Zakharov. Fully quantum treatment of the Landau-Pomeranchuk-Migdal effect in QED and QCD. *JETP Lett.*, 63:952–957, 1996. [40](#)
- [130] Peter Brockway Arnold, Guy D. Moore, and Laurence G. Yaffe. Photon emission from ultrarelativistic plasmas. *JHEP*, 11:057, 2001. [40](#)
- [131] Urs Achim Wiedemann. Gluon radiation off hard quarks in a nuclear environment: Opacity expansion. *Nucl. Phys. B*, 588:303–344, 2000. [40](#)
- [132] Xin-Nian Wang and Xiao-feng Guo. Multiple parton scattering in nuclei: Parton energy loss. *Nucl. Phys. A*, 696:788–832, 2001. [40](#)
- [133] Miklos Gyulassy, Peter Levai, and Ivan Vitev. Jet quenching in thin quark gluon plasmas. 1. Formalism. *Nucl. Phys. B*, 571:197–233, 2000. [40](#)
- [134] Korinna C. Zapp, Frank Krauss, and Urs A. Wiedemann. A perturbative framework for jet quenching. *JHEP*, 1303:080, 2013. [40](#), [87](#), [99](#)
- [135] Korinna C. Zapp. JEWEL 2.0.0: directions for use. *Eur.Phys.J.*, C74(2):2762, 2014. [40](#), [87](#), [99](#)
- [136] Raghav Kunnawalkam Elayavalli and Korinna Christine Zapp. Medium response in JEWEL and its impact on jet shape observables in heavy ion collisions. *JHEP*, 07:141, 2017. [41](#), [99](#)
- [137] Yayun He, Tan Luo, Xin-Nian Wang, and Yan Zhu. Linear Boltzmann Transport for Jet Propagation in the Quark-Gluon Plasma: Elastic Processes and Medium Recoil. *Phys. Rev. C*, 91:054908, 2015. [Erratum: *Phys.Rev.C* 97, 019902 (2018)]. [41](#), [99](#)
- [138] Yayun He, Shanshan Cao, Wei Chen, Tan Luo, Long-Gang Pang, and Xin-Nian Wang. Interplaying mechanisms behind single inclusive jet suppression in heavy-ion collisions. *Phys. Rev.*, C99(5):054911, 2019. [41](#), [99](#), [113](#)
- [139] Weiyao Ke and Xin-Nian Wang. QGP modification to single inclusive jets in a calibrated transport model. *JHEP*, 05:041, 2021. [41](#), [99](#)
- [140] Yacine Mehtar-Tani, Daniel Pablos, and Konrad Tywoniuk. Cone-Size Dependence of Jet Suppression in Heavy-Ion Collisions. *Phys. Rev. Lett.*, 127(25):252301, 2021. [41](#), [65](#), [99](#)
- [141] Bjoern Schenke, Charles Gale, and Sangyong Jeon. MARTINI: An Event generator for relativistic heavy-ion collisions. *Phys. Rev. C*, 80:054913, 2009. [41](#), [99](#)

- [142] J. H. Putschke et al. The JETSCAPE framework. 3 2019. 41, 99
- [143] Kim Albertsson et al. Machine Learning in High Energy Physics Community White Paper. *J. Phys. Conf. Ser.*, 1085(2):022008, 2018. 42
- [144] Leo Breiman. Random forests. *Machine Learning*, 45(1):5–32, 2001. 43
- [145] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011. 43, 73
- [146] Gilles Louppe. *Understanding Random Forests: From Theory to Practice*. PhD thesis, University of Liege, Belgium, 10 2014. arXiv:1407.7502. 43
- [147] Shai Shalev-Shwartz and Shai Ben-David. *Understanding Machine Learning - From Theory to Algorithms*. Cambridge University Press, 2014. 43
- [148] Jürgen Schmidhuber. Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117, jan 2015. 43
- [149] Yann Coadou. Boosted decision trees. 3 2022. 43
- [150] Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks, 2014. 44
- [151] Laith Alzubaidi, Jinglan Zhang, Amjad J. Humaidi, Ayad Al-dujaili, Ye Duan, Omran Al-Shamma, Jesus Santamaría, Mohammed Abdulraheem Fadhel, Muthana Al-Amidie, and Laith Farhan. Review of deep learning: concepts, cnn architectures, challenges, applications, future directions. *Journal of Big Data*, 8, 2021. 44
- [152] Tong Yu and Hong Zhu. Hyper-parameter optimization: A review of algorithms and applications, 2020. 44, 48
- [153] Andrew P. Bradley. The use of the area under the roc curve in the evaluation of machine learning algorithms. *Pattern Recognition*, 30(7):1145–1159, 1997. 44
- [154] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016. <http://www.deeplearningbook.org>. 48
- [155] Matthew Feickert and Benjamin Nachman. A Living Review of Machine Learning for Particle Physics. 2 2021. 48, 49
- [156] Luke de Oliveira, Michael Kagan, Lester Mackey, Benjamin Nachman, and Ariel Schwartzman. Jet-images – deep learning edition. *JHEP*, 07:069, 2016. 48
- [157] Michael Kagan. Image-Based Jet Analysis. 12 2020. 48

- [158] A. M. Sirunyan et al. Identification of heavy-flavour jets with the CMS detector in pp collisions at 13 TeV. *JINST*, 13(05):P05011, 2018. [49](#)
- [159] Georges Aad et al. Identification of Jets Containing b -Hadrons with Recurrent Neural Networks at the ATLAS Experiment. Technical report, CERN, Geneva, 2017. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2017-003>.
- [160] Georges Aad et al. Search for new resonances in mass distributions of jet pairs using 139 fb^{-1} of pp collisions at $\sqrt{s} = 13 \text{ TeV}$ with the ATLAS detector. *JHEP*, 03:145, 2020.
- [161] Rüdiger Haake. Machine and deep learning techniques in heavy-ion collisions with ALICE. *PoS EPS-HEP2017 (2017)*, 9 2017. [49](#)
- [162] Leif Lönnblad, Carsten Peterson, and Thorsteinn Rognvaldsson. Finding gluon jets with a neural trigger. *Phys. Rev. Lett.*, 65:1321–1324, Sep 1990. [49](#)
- [163] Jon Pumplin. How to tell quark jets from gluon jets. *Phys. Rev. D*, 44:2025–2032, Oct 1991.
- [164] Georges Aad et al. Quark versus Gluon Jet Tagging Using Jet Images with the ATLAS Detector. Technical report, CERN, Geneva, 2017. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2017-017>.
- [165] Jason Sang Hun Lee, Inkyu Park, Ian James Watson, and Seungjin Yang. Quark-Gluon Jet Discrimination Using Convolutional Neural Networks. *J. Korean Phys. Soc.*, 74(3):219–223, 2019. [49](#)
- [166] Anders Andreassen, Patrick T. Komiske, Eric M. Metodiev, Benjamin Nachman, and Jesse Thaler. OmniFold: A Method to Simultaneously Unfold All Observables. *Phys. Rev. Lett.*, 124(18):182001, 2020. [49](#), [79](#)
- [167] V. Andreev et al. Measurement of Lepton-Jet Correlation in Deep-Inelastic Scattering with the H1 Detector Using Machine Learning for Unfolding. *Phys. Rev. Lett.*, 128(13):132002, 2022. [49](#)
- [168] Multidifferential study of identified charged hadron distributions in Z -tagged jets in proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$. 8 2022. [49](#)
- [169] Yang-Ting Chien and Raghav Kunnawalkam Elayavalli. Probing heavy ion collisions using quark and gluon jet substructure. 3 2018. [49](#)

- [170] Yue Shi Lai, Duff Neill, Mateusz Płoskoń, and Felix Ringer. Explainable machine learning of the underlying physics of high-energy particle collisions. *Phys. Lett. B*, 829:137055, 2022. [49](#)
- [171] K. Aamodt et al. The ALICE experiment at the CERN LHC. *JINST*, 3:S08002, 2008. [50](#)
- [172] P. Cortese et. al. *ALICE forward detectors: FMD, TO and VO: Technical Design Report*. Technical design report. ALICE. CERN, Geneva, 2004. Submitted on 10 Sep 2004. [51](#)
- [173] E. Abbas et al. Performance of the ALICE VZERO system. *JINST*, 8:P10016, 2013. [51](#)
- [174] Shreyasi Acharya et al. The ALICE Transition Radiation Detector: construction, operation, and performance. *Nucl. Instrum. Meth. A*, 881:88–127, 2018. [51](#)
- [175] Francesca Carnesecchi. Performance of the ALICE Time-Of-Flight detector at the LHC. *JINST*, 14(06):C06023, 2019. [51](#)
- [176] Yu Belikov, K Safarik, and B Batyunya. Kalman Filtering Application for Track Recognition and Reconstruction in ALICE Tracking System. Technical report, CERN, Geneva, 1997. [53](#)
- [177] M. et. al Tanabashi. Review of particle physics. *Phys. Rev. D*, 98:030001, Aug 2018. [54](#)
- [178] M Jobs and H R Shaylor. Data analysis techniques in high energy physics. *Reports on Progress in Physics*, 35(3):1077, sep 1972. [61](#)
- [179] Measurement of the radius dependence of charged-particle jet suppression in Pb-Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV. 3 2023. [63](#)
- [180] Jaroslav Adam et al. Measurement of jet suppression in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 746:1–14, 2015. [63](#)
- [181] Georges Aad et al. Measurements of the nuclear modification factor for jets in Pb + Pb collisions at $\sqrt{s_{NN}}=2.76$ TeV with the ATLAS Detector. *Phys. Rev. Lett.*, 114(7):072302, 2015.
- [182] Vardan Khachatryan et al. Measurement of inclusive jet cross sections in pp and Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Rev. C*, 96(1):015202, 2017. [63](#), [95](#)
- [183] Daniel Pablos. Jet Suppression From a Small to Intermediate to Large Radius. *Phys. Rev. Lett.*, 124(5):052301, 2020. [63](#), [65](#), [99](#), [113](#)
- [184] Georges Aad et al. Measurement of the jet radius and transverse momentum dependence of inclusive jet suppression in lead-lead collisions at $\sqrt{s_{NN}}= 2.76$ TeV with the ATLAS detector. *Phys. Lett. B*, 719:220–241, 2013. [64](#), [99](#), [125](#)

- [185] Yang-Ting Chien and Ivan Vitev. Towards the understanding of jet shapes and cross sections in heavy ion collisions using soft-collinear effective theory. *JHEP*, 05:023, 2016. [65](#)
- [186] Yasuki Tachibana, Ning-Bo Chang, and Guang-You Qin. Full jet in quark-gluon plasma with hydrodynamic medium response. *Phys. Rev. C*, 95(4):044909, 2017. [65](#)
- [187] Chanwook Park et al. Constraints on jet quenching from a multi-stage energy-loss approach. *PoS, HardProbes2020*:150, 2021. [65](#)
- [188] R Brun, F Bruyant, M Maire, A C McPherson, and P Zanmarini. *GEANT 3: user's guide Geant 3.10, Geant 3.11; rev. version*. CERN, Geneva, 1987. [66](#), [67](#)
- [189] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. FastJet User Manual. *Eur.Phys.J.*, C72:1896, 2012. [67](#)
- [190] B. Abelev et al. Measurement of charged jet suppression in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 03:013, 2014. [67](#), [95](#)
- [191] Betty Abelev et al. Measurement of Event Background Fluctuations for Charged Particle Jet Reconstruction in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 03:053, 2012. [67](#), [68](#)
- [192] Rüdiger Haake and Constantin Loizides. Machine Learning based jet momentum reconstruction in heavy-ion collisions. *Phys. Rev. C*, 99(6):064904, 2019. [69](#)
- [193] Giulio D'Agostini. A multidimensional unfolding method based on Bayes' Theorem. Technical report, DESY, Hamburg, 1994. [79](#)
- [194] Andreas Hocker and Vakhtang Kartvelishvili. SVD approach to data unfolding. *Nucl. Instrum. Meth. A*, 372:469–481, 1996. [79](#)
- [195] Shreyasi Acharya et al. Measurement of jet radial profiles in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett. B*, 796:204–219, 2019. [84](#)
- [196] S. Acharya et al. Transverse momentum spectra and nuclear modification factors of charged particles in pp, p–Pb and Pb–Pb collisions at the LHC. *JHEP*, 11:013, 2018. [86](#)
- [197] René Brun, F Bruyant, Federico Carminati, Simone Giani, M Maire, A McPherson, G Patrick, and L Urban. *GEANT: Detector Description and Simulation Tool; Oct 1994*. CERN Program Library. CERN, Geneva, 1993. Long Writeup W5013. [89](#)
- [198] Xin-Nian Wang and Miklos Gyulassy. HIJING: A Monte Carlo model for multiple jet production in p p, p A and A A collisions. *Phys. Rev. D*, 44:3501–3516, 1991. [90](#)

- [199] Betty Bezverkhny Abelev et al. Performance of the ALICE experiment at the CERN LHC. *Int.J.Mod.Phys.*, A29:1430044, 2014. [90](#)
- [200] B. Abelev et al. Measurement of charged jet suppression in Pb–Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 1403:013, 2014. [90](#)
- [201] Shreyasi Acharya et al. Measurement of charged jet cross section in pp collisions at $\sqrt{s} = 5.02$ TeV. *Phys. Rev. D*, 100(9):092004, 2019. [90](#)
- [202] B. Abelev et al. Measurement of the inclusive differential jet cross section in pp collisions at $\sqrt{s} = 2.76$ TeV. *Phys.Lett.*, B722:262–272, 2013. [95](#)
- [203] Vardan Khachatryan et al. Measurement of the double-differential inclusive jet cross section in proton–proton collisions at $\sqrt{s} = 13$ TeV. *Eur. Phys. J. C*, 76(8):451, 2016. [95](#)
- [204] Shanshan Cao and Abhijit Majumder. Nuclear modification of leading hadrons and jets within a virtuality ordered parton shower. *Phys. Rev. C*, 101(2):024903, 2020. [99](#)
- [205] Hannah Bossi. Inclusive Jet Measurements in Pb-Pb Collisions at 5.02 TeV with ALICE using Machine Learning Techniques. *PoS, HardProbes2020*:135, 2021. [102](#)
- [206] R. Baier, Yuri L. Dokshitzer, Alfred H. Mueller, and D. Schiff. Quenching of hadron spectra in media. *JHEP*, 09:033, 2001. [111](#)
- [207] Hai Tao Li and Ivan Vitev. Inclusive heavy flavor jet production with semi-inclusive jet functions: from proton to heavy-ion collisions. *JHEP*, 07:148, 2019. [113](#)
- [208] Yi-Lun Du, Daniel Pablos, and Konrad Tywoniuk. Classification of quark and gluon jets in hot QCD medium with deep learning. *PoS, PANIC2021*:224, 2022. [121](#)
- [209] Patrick T. Komiske, Ian Moul, Jesse Thaler, and Hua Xing Zhu. Analyzing N-Point Energy Correlators inside Jets with CMS Open Data. *Phys. Rev. Lett.*, 130(5):051901, 2023. [121](#)
- [210] Evan Craft, Kyle Lee, Bianka Meçaj, and Ian Moul. Beautiful and Charming Energy Correlators. 10 2022. [121](#)
- [211] Carlota Andres, Fabio Dominguez, Raghav Kunnawalkam Elayavalli, Jack Holguin, Cyrille Marquet, and Ian Moul. Resolving the Scales of the Quark-Gluon Plasma with Energy Correlators. 9 2022. [121](#)