

Theory and Phenomenology of Heavy-Ion Collisions

by

Jasmine Therese Brewer

B.S., University of Colorado Boulder (2015)

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Author
Department of Physics
August 14, 2020

Certified by.....
Krishna Rajagopal
William A.M. Burden Professor of Physics
Thesis Supervisor

Accepted by
Nergis Mavalvala
Associate Department Head

Theory and Phenomenology of Heavy-Ion Collisions

by

Jasmine Therese Brewer

Submitted to the Department of Physics
on August 14, 2020, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

The quark–gluon plasma produced in high energy collisions between large nuclei provides a unique window into the rich many-body structure of the theory of the strong nuclear force, Quantum Chromodynamics (QCD). The theory and phenomenology of quark–gluon plasma is diverse and comprises varied conceptual frameworks for understanding the structure of this exotic material and how it forms and evolves in heavy-ion collisions. In this thesis, we cover aspects from three principal approaches: collectivity and the onset of hydrodynamic behavior in the quark–gluon plasma, energy loss and modification of high-energy particles through their interaction with the quark–gluon plasma, and a conjectured critical point in the phase diagram of QCD that may be passed in heavy-ion collisions. We first make a novel conjecture that the very early onset of hydrodynamic-like (attractor) behavior occurs because the system evolution is dominated by a reduced set of “pre-hydrodynamic” modes. In a simple model, these modes are the instantaneous ground state modes of an effective Hamiltonian that evolve into hydrodynamic modes in the hydrodynamic limit. Next we demonstrate an analysis method to compare jets in proton–proton and heavy-ion collisions that were more similar when they were produced. This method enables a direct measurement of the average energy loss of jets as a function of their energy, in addition to clarifying the interpretation of jet modification measurements. We further illustrate an analysis method to disentangle the modification of quark- and gluon-initiated jets in heavy-ion collisions. We provide a proof-of-concept that the quark and gluon fractions and the modification of quark and gluon jets can be separately estimated in heavy-ion collision experiments. Finally, we describe additional signatures of a critical point in the QCD phase diagram accessible in the Beam Energy Scan experiments by using a rapidity scan to probe more local structure of the phase diagram at each beam energy.

Thesis Supervisor: Krishna Rajagopal
Title: William A.M. Burden Professor of Physics

Acknowledgments

The mark of five years well spent is that I can hardly recognize the person I was five short years ago. The thanks for this transformation is due to many wonderful people and experiences, at MIT and beyond.

It has been an enormous honor and fortune to learn from and work with my thesis advisor, Krishna Rajagopal, over the past five years. I have benefited beyond measure from your keen physics insight; you always ask the right questions to set a project on the right path, and you can sniff out a mistake from a hundred miles away. I am still astonished by the number of opportunities you have given me to travel, to go to schools, to present my work, and to establish an extensive network in the heavy-ion community. You have been an unwavering supporter of me and are always quick to introduce me to people in the community and generously advertise me and my work. You have consistently been a thoughtful and deliberate mentor, and have given me the unique opportunity to try out my independence and grow as a physicist without ever having to feel lost or alone.

I am also extremely grateful to have had the opportunity to work with and learn from Jesse Thaler over the last several years. Your enthusiasm, breadth and depth of knowledge, and creative problem solving are inspirational. Your thrill in doing physics is contagious and it is an absolute pleasure to discuss physics with you. Your careful feedback on science, writing, and presentations has also been invaluable and has pushed me to uphold a higher standard in my work.

A list of my critical mentors is certainly not complete without my undergraduate advisor, Paul Romatschke. I will never know where I would be without you, but I can say with certainty that I would not be receiving a Ph.D. in heavy-ion physics from MIT. You were my first role model for what a great physicist should be: brilliant, creative, and enthusiastic while also being down-to-earth and generous with your time. You inspire me to this day and I am very grateful for the enormous amount of time and energy you invested in me. You told me at some point that as theorists we are lucky because we do not have any specific equipment required to do our jobs

that prevents us from thinking about something new every day; we are only limited by our willingness to go outside of our comfort zone and try new things. This lesson played no small part in the breadth of this thesis.

I have been fortunate to work with many other wonderful mentors and collaborators during my PhD. I am grateful to Yi Yin for his sharp physics insight and creativity that has led to several stimulating collaborations; I almost always see a new way of looking at a physics problem from discussing with you. I have benefited enormously from many thought-provoking discussions with Guilherme Milhano and Yen-Jie Lee. I am also thankful to Swagato Mukherjee, Andrey Sadofyev, Wilke van der Schee, Andrew Turner, and Li Yan for their crucial contributions to the work presented in this thesis, as well as to more recent collaborators Quinn Brodsky, Bruno Scheihing Hitschfeld, Weiyao Ke, Andrew Lin, Aleksas Mazeliauskas, Sohyun Park, and Urs Wiedemann for numerous enlightening discussions. Many thanks are also in order to an enumerable list of friends and colleagues in the heavy-ion community. You have all made travelling for “work” a joy that I always look forward to and I am happy to be part of such a warm and fun community of people.

Thanks to my wonderful CTP friends and cohort. I never fail to be astounded by how brilliant, caring, giving, and welcoming you have all been. I owe a great deal to the extensive help of my co-sufferers and teachers in the adventure that is Part III, especially Andrew Turner, Hongwan Liu, Gherardo Vita, Patrick Fitzpatrick, and Lina Necib. Every time I see a Quantum Field Theory problem I know which one of you I would ask about it. I am also thankful to Sarah Geller, Anthony Grebe, Greg Ridgway, and Yu-Chien Huang for making the CTP fun, and to Scott Morley, Charles Suggs, and Joyce Berggren for their friendly faces and for keeping us in line.

I am also grateful to the amazing friends who have kept life fun and have encouraged and supported me. Thanks to Simon Grosse-Holz and Isobel Kolbe for always being ready for any crazy adventure, to Luiz Gustavo Pimenta Martins, Alice Oliveira, and Anastasia Patterson for being warm and fun and always up to celebrate, to Samantha Strasser and Alex Leder for many culinary explorations, and Thomas and Amanda Lewis for being incredibly welcoming and supportive. Thanks to my

oldest friend, Devin Hansen, for growing up with me and for having a defining influence on my outlook and personality, and to my godmother Cathy Williams for a lifetime of love and friendship. I am especially indebted to the kind, genuine people who have kept me upright during the tough times and laughing during the good times over the last five years. Thanks to Sarah Geller for being a fiercely loyal friend and an incredible and inspiring person. Thanks to William Lewis for being a steadfast friend through many difficult times and for enjoying many fun adventures with me. Thanks to Andrew Turner for being one of my closest friends and an accomplice in many shenanigans; being your officemate made every day brighter and I will sincerely miss many great physics and non-physics conversations. I finally want to thank Lina Necib, who in addition to being a wonderful friend has been an exceptional mentor and role model. You are a constant source of strength and inspiration for me and I am extremely lucky that you adopted me as your protégé.

I would also like to thank the people who stood by me during the darkest and most challenging part of my PhD, particularly Aashka Dave, Zelda Mariet, Nuria Jane Chimeno, DiOnetta Jones Crayton, Anna Frebel, Cathy Modica, and Julie Shah. You taught me that by standing with someone and amplifying their voice you can lend them power when they can no longer find it in themselves. Through this thank you I commit to pushing myself to be for others what you have been for me. I also owe a great deal to the dear friends who stood unwaveringly by my side and offered countless hours of support in all possible ways, especially William Lewis, Lina Necib, and Andrew Turner.

Finally, to my parents. You taught me that the possibility of failure is never an excuse not to try something; that you never know if you will be bad at something until you try, and even if you are bad at it, if you like it you should do it anyway. You also taught me the value of hard work, perseverance, and responsibility. You let me be the sculptor of my own life from a very young age and I think genuinely never had your own mould of what you wanted me to be besides what I wanted to be. Thank you for your complete and unconditional love and support in everything I choose to do. I dedicate this thesis to you.

Contents

1	Introduction	17
2	Overview	37
2.1	This Thesis	37
2.2	Additional work	40
3	Adiabatic Hydrodynamization	43
3.1	Identification of pre-hydrodynamic mode(s)	46
3.2	Emergent dominance of pre-hydrodynamic modes at $\tau \ll \tau_C$	48
3.3	Implications of the dominance of pre-hydrodynamic modes	49
3.4	RTA as an example of adiabatic hydrodynamization	50
3.5	Adiabaticity in the rapidly-expanding QGP	54
3.6	Outlook	56
3.7	Ongoing work	56
4	A Quantile Procedure for Sorting Out Quenched Jets	59
5	Separating Quark and Gluon Jets	71
6	Rapidity Dependence of Cumulants	83
7	The Road Ahead	95

List of Figures

1-1	A schematic picture of the phase diagram of Quantum Chromodynamics (QCD). At low energies and densities, quarks and gluons are confined in hadrons like protons and neutrons, which themselves are the building blocks of the atomic nuclei that we see around us. At high energies or densities, quarks and gluons become deconfined and form a state of matter called quark-gluon plasma. These extremely high temperatures can be accessed in collisions of large nuclei at the Large Hadron Collider (LHC) and the Relativistic Heavy-Ion Collider (RHIC) as illustrated by the orange and yellow arrows. The conjectured (but not yet confirmed) critical point in the phase diagram being searched for at RHIC is illustrated by the white dot. Figure reproduced from Ref. [1].	19
1-3	Example displays of events containing jets in (a) proton–proton and (b) heavy-ion collisions as measured by the CMS experiment at the LHC. Jets are shown as large blue and red towers. While back-to-back jets in proton–proton collisions are relatively balanced in energy due to momentum conservation, a much higher asymmetry is observed in heavy-ion collisions due to energy loss of jets in the quark–gluon plasma.	22

- 1-5 Demonstration of attractor behavior for three microscopic theories: rBRSSS (causal version of second-order hydrodynamics, see Ref. [2]), Boltzmann (kinetic theory), and AdS/CFT ($\mathcal{N} = 4$ SYM). Shown is a quantity related to the ratio of the longitudinal pressure to the energy density as a function of the proper time τ rescaled by the temperature T (the combination τT is the inverse gradient strength). Numerical solutions to the full microscopic dynamics for many initial conditions (grey) converge to the attractor solution (black) while the system is still far from the hydrodynamic limit (red dashed). Solutions to first- and second-order hydrodynamics agree quantitatively with the attractor while the system is still far from equilibrium. Figure reproduced from Ref. [3]. 26
- 1-6 Net-proton number (a proxy for baryon number density) as a function of the rapidity y for several different energies in the range accessible at the RHIC beam energy scan. At low energies the baryon density is at midrapidity, and with increasing energy it moves toward successively higher rapidity. At LHC energies, which are substantially higher energy than any shown on this plot, the baryon density is at much higher rapidities and is essentially zero at midrapidity. Curves are hydrodynamic simulations from Ref. [4] and the figure is reproduced from Ref. [4]. 31
- 1-7 (a) A density plot showing the regions of positive (blue) and negative (red) kurtosis as a function of the Ising phase diagram variables (t, H). The dashed green line shows a possible freezeout curve in this phase diagram. The kurtosis along this line is shown in (b). Qualitatively, the non-monotonicity and the sign change of the kurtosis as a function of t shown are anticipated signatures of the criticality as a function of μ_B . Reproduced from Ref. [5] 32

3-1 We demonstrate the dominance of pre-hydrodynamic modes, and hence adiabatic hydrodynamization, for a Bjorken-expanding plasma in the relaxation time approximation. The red solid curve shows $\mathcal{E}_0(\lambda)$, which is the contribution to the percentage rate of change of the energy density coming from the pre-hydrodynamic mode (c.f. Eq. (3.2)). Black curves are $g(\lambda)$ obtained from solving Eq. (3.3) with constant τ_C for different initial conditions. After τ_{Redu} , they collapse onto the RTA attractor as obtained in Refs. [6, 3, 7]. The left and right shaded regions indicate $\tau \leq \tau_{\text{Redu}}$ and $\tau \geq \tau_{\text{Hydro}}$, respectively. 51

3-2 We show the fractional difference between \mathcal{E}_0 and g , which measures the relative importance of contributions from the ground state (pre-hydrodynamic) and excited modes. Red and dashed blue curves show this difference for representative initial conditions with constant and conformal τ_C , respectively. The fact that this quantity is small indicates the dominance of pre-hydrodynamic modes (i.e. adiabaticity) during the interval $\tau_{\text{Redu}} \leq \tau < \tau_{\text{Hydro}}$ 52

4-2 Distributions of (a) R_{AA} as a function of p_T^{jet} and (b) Q_{AA} as a function of p_T^{quant} , for the Z+jet (dashed) and di-jet (solid) samples in JEWEL. Although R_{AA} and Q_{AA} are derived from the same underlying jet p_T spectra, they provide different and complementary information. For example, the p_T dependence of R_{AA} is very different for Z+jet and di-jet events in JEWEL, while the average fractional p_T loss $1 - Q_{AA}$ is similar. Note that R_{AA} requires binning of the data, while Q_{AA} , which is based on the cumulative cross-section, can be plotted unbinned. 64

4-3 Mean of the jet p_T distribution compared to a baseline initial p_T (top), along with the corresponding standard deviation (bottom). Shown are (a) Z+jet events where the baseline is the physically observable p_T of the recoiling Z boson and (b) di-jet events where the baseline is the unphysical and unobservable p_T^{MC} of the initial hard scattering obtained from JEWEL. The reconstructed jet p_T for proton-proton and heavy-ion jets are shown in dashed black and blue, respectively. The p_T^{quant} of the heavy-ion sample, shown in red, more closely matches the initial jet p_T than the reconstructed heavy-ion p_T does. 66

4-4 Distribution of m/p_T for proton-proton (dashed black) and heavy-ion (blue) jets in di-jet events with reconstructed $p_T \in [100, 200]$ GeV. Heavy-ion jets with $p_T^{\text{quant}} \in [100, 200]$ GeV, corresponding to $p_T^{\text{AA}} \in [80, 173]$ GeV, are in red. The heavy-ion result is normalized to match the proton-proton baseline but the quantile result has the correct normalization by construction. Partially compensating for p_T migration via the quantile procedure shifts m/p_T towards being less modified. . . 68

5-1 Extracting quark-like and gluon-like jet topics from (top row) proton–proton collisions and (bottom row) heavy-ion collisions, as generated by PYTHIA and JEWEL, respectively. (a) Normalized distributions of jet constituent multiplicity for the γ + jet and dijet samples in proton–proton collisions. (b) The two underlying topics extracted from these distributions using the DEMIX method (colorful bands), compared to the self-normalized MC-level definition of quark- and gluon-initiated jets (black) for those jets within $\Delta R = 0.4$ of an initiating parton. (c) Fractions of the gluon-like topic in the γ + jet and dijet samples, compared to the MC-level quark- and gluon-initiated jet fractions. The corresponding results for heavy-ion collisions are shown in (d)–(f). The heavy-ion γ +jet statistics are chosen to match those available in Run 4 at the LHC. Possible reasons for the higher gluon-like topic fractions compared to the MC label fractions are provided in the text. 78

5-2 Gluon-like topic fractions for (a) proton–proton collisions and (b) heavy-ion collisions, as a function of jet p_T , for the γ + jet (hatched) and dijet (solid) samples. Colorful bands are the extracted gluon-like topic fractions from the full γ + jet and dijet distributions, as already shown in figs. 5-1c and 5-1f. Black bands show for comparison the gluon-like topic fractions extracted from (unphysically) restricted γ +jet and dijet samples including only jets whose axis is within $\Delta R = 0.4$ of a parton in the hard scattering matrix element. 79

5-3 Same as the bottom row of fig. 5-1, but using a factor of about 2.8 higher γ + jet statistics than expected after LHC Run 4. With more events, the agreement of the distribution of the quark-like topic (purple) relative to the MC-level definition (black) is somewhat improved compared to fig. 5-1e, though the gluon-like topic fraction remains high compared to the MC label fraction. 80

- 5-4 Constituent multiplicity distributions for (a) the quark-like topic and (b) the gluon-like topic as a function of jet p_T for p_T bins [100, 120] GeV, [120, 140] GeV, and [140, 160] GeV. Each violin plot has results for both (left side) proton–proton and (right side) heavy-ion collisions, and the change between the two sides illustrates the modification of the constituent multiplicity distribution for the corresponding topic. Horizontal lines indicate the median (solid) and 16% and 84% quantiles (dashed) of the multiplicity distributions. 81
- 6-1 In the left column we see that we have assumed the existence of a critical point (red dot) at $(\mu_B^c, T) = (260, 160)$ MeV whose critical region, bounded by the contour where $\xi = 1$ fm, is colored red and blue. The colors denote the sign of $\omega_{4,\sigma}$, with $\omega_{4,\sigma} > 0$ in blue and $\omega_{4,\sigma} < 0$ in red. Different rows correspond to different assumptions for where on the phase diagram a heavy ion collision freezes out, *cf.* collisions with varying beam energy. The black circles show where freezeout occurs at mid-rapidity, from top to bottom with $\mu_{B,0} = 200, 230, 240, 240$ MeV. The black dashed curves show how the freezeout conditions change with increasing spacetime rapidity, with the circle, square, and triangle indicating freezeout at $y_s = 0, 0.6,$ and $1.2,$ respectively. . . 92
- 6-2 We illustrate the behavior of $\omega_{4,\sigma}$ when the freezeout conditions are as in the first row of Figure 6-1 with $\mu_{B,0} = 200$ MeV and $\alpha = 50$ MeV, but where here we let the location of the critical point range from $\mu_B^c = 190$ MeV to $\mu_B^c = 300$ MeV. For each value of μ_B^c , namely for each horizontal slice across the figure, color indicates the value of $\omega_{4,\sigma}$ as a function of $|y_c|$, with $\Delta y = 0.4$ fixed as in Fig. 6-1. The slice indicated by the green dashed line corresponds to the top-right panel of Fig. 6-1. For $\mu_B^c < \mu_{B,0}$ ($\mu_B^c > \mu_{B,0}$), $\omega_{4,\sigma}$ decreases (increases) with increasing $|y_c|$ 93

Chapter 1

Introduction

Quantum Chromodynamics (QCD) is the theory of the strong nuclear force and governs the properties of almost all matter in our universe. The fundamental particles in QCD are quarks and gluons that carry and interact via the QCD charge, “color”. One of the most unique features of QCD is a coupling strength that is large for interactions with low momentum transfer and small for those with large momentum transfer [8, 9]. Strong interactions at low momentum transfer leads to the famous property of confinement in QCD: at low energies colored objects (quarks and gluons) are always confined within color-neutral states like protons, neutrons, and pions. Therefore all of the matter that we see in the universe is composite states of two, three, or more quarks and gluons. Sophisticated numerical simulations of QCD at low energies, lattice QCD, have been tremendously successful at describing the properties of these bound states from the first-principles dynamics of quarks and gluons [10]. On the other hand, the weak coupling of QCD for processes with high momentum transfer has enabled immense progress on understanding the interactions of quarks, gluons, and other fundamental particles in high energy proton-proton collisions at the Large Hadron Collider (LHC). Though the particles detected in these experiments are still hadrons, weakly-coupled QCD has been remarkably successful at describing these processes in terms of the interactions between quarks and gluons.

However, as encompassed concisely in the profound observation of Phil Anderson, “more is different” [11]. The aftermath of proton-proton collisions, for example at the

LHC, is relatively dilute with few total interactions between quarks and gluons. In high-energy collisions between larger systems like heavy nuclei, the situation is much different. The density of interactions is high enough that it makes sense not to talk about individual quarks and gluons but to talk about a *material* whose properties are the manifestation of a macroscopic number of those interactions. Unlike in more dilute systems, this material can be characterized by macroscopic properties like temperature and pressure, as for more everyday materials like water. At extremely high temperatures or (baryon number) densities, the hadronic matter that we experience in everyday life (like protons and neutrons) undergoes a “deconfinement transition”, where the quarks and gluons usually confined inside of these composite states become unbound, though their coupling is still strong. The resulting material formed from deconfined quarks and gluons is called quark-gluon plasma. This transition is illustrated in Fig. 1-1. Heavy-ion collisions are thus a unique opportunity to study the properties of a macroscopic droplet of a material whose underlying constituents are quarks and gluons.

Quark-gluon plasma has been created in high-energy collisions between gold nuclei at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Lab [12, 13, 14, 15, 16], and in collisions of lead nuclei at the Large Hadron Collider (LHC) at CERN [17, 18, 19]. This discovery ushered in an era of trying to understand the properties of this exotic material experimentally. This area is particularly exciting because first-principles QCD calculations in this regime are infeasible, and so in many ways everything about the properties of this matter is uncharted theoretical (and experimental) territory. In addition, the experimental and phenomenological challenges are significant: the quark-gluon plasma only exists for about 10^{-23} seconds before it turns back into the pions, protons, and other hadrons that are observed at the detector. This makes the data analysis and interpretation challenging because all of the dynamics of the collision must be inferred from particles measured long after the plasma is gone.

There are two principal strategies that can be used to study the properties of any fluid. The first is to study how it flows around obstacles or expands, which is





Figure 1-1: A schematic picture of the phase diagram of Quantum Chromodynamics (QCD). At low energies and densities, quarks and gluons are confined in hadrons like protons and neutrons, which themselves are the building blocks of the atomic nuclei that we see around us. At high energies or densities, quarks and gluons become deconfined and form a state of matter called quark-gluon plasma. These extremely high temperatures can be accessed in collisions of large nuclei at the Large Hadron Collider (LHC) and the Relativistic Heavy-Ion Collider (RHIC) as illustrated by the orange and yellow arrows. The conjectured (but not yet confirmed) critical point in the phase diagram being searched for at RHIC is illustrated by the white dot. Figure reproduced from Ref. [1].


sensitive to macroscopic properties of the fluid like shear viscosity. The second is to shoot high-energy particles through a vat of the fluid and to study how they lose energy and are deflected. This is a famous way to study the microscopic degrees of freedom of a material, and the wake of these high-energy particles in the fluid also probes its far-from-equilibrium response. These two conceptual approaches are both cornerstone methods for studying the properties of the quark-gluon plasma produced in heavy-ion collisions.

The flow and expansion of the droplet of quark-gluon plasma produced in a heavy-ion collision cannot be observed directly since real-time measurements are impossible due to its extremely short lifetime (around 10^{-23} seconds). However, correlations between particles produced from the quark-gluon plasma are sensitive to the dynamics that generated those correlations. In particular, it is anticipated that collisions between large nuclei at non-zero impact parameter have an elliptic overlap region (see Fig. 1-2a) that produces a dominant elliptic geometry of the quark-gluon plasma created in the collision. The correlations that result from this initial spatial anisotropy are sensitive to the underlying interactions. Without any interactions between particles (“free-streaming”), the expansion is uniform regardless of the initial spatial anisotropy. A characteristic feature of hydrodynamics, on the other hand, is that anisotropic pressure gradients resulting from spatial anisotropy drive faster expansion along the direction of higher pressure gradients. This translates spatial anisotropy into momentum anisotropy, with an efficiency dictated by dissipation in the fluid (e.g., shear viscosity). Specifically, dissipation tends to isotropize the momentum distribution by transferring spatial anisotropy into random thermal motion, while fluids with very low dissipation exhibit “collective flow” where spatial and momentum anisotropy are easily interchanged. Measurements of a large “ellipticity” coefficient v_2 of the momentum anisotropy of particles produced from the quark gluon plasma suggest a hydrodynamic description of a fluid with low dissipation, and put constraints on the ratio of shear viscosity to entropy density η/s for the produced fluid (see Fig. 1-2b for an example and Ref. [20] for a comprehensive review).

The short lifetime of the quark-gluon plasma also makes it impossible to use an



(a) Typical geometry of a Pb–Pb collision at the LHC from Monte Carlo Glauber calculations. Nucleons that participate in the collision are indicated by solid lines and spectator nucleons are represented by dashed lines. The non-zero impact parameter of the collision produces a dominant elliptical shape of participant nucleons. Figure reproduced from Ref. [21].



(b) Elliptic flow coefficient v_2 as a function of transverse momentum p_T of particles produced in the collision as measured by the STAR experiment Ref. [22] (black) compared to early hydrodynamic simulations with different values of the ratio of shear viscosity to entropy density, η/s , which characterizes dissipation in the fluid. Figure reproduced from Ref. [23]. The extraction of η/s became more precise with later 3+1-dimensional simulations, e.g., Ref. [24].

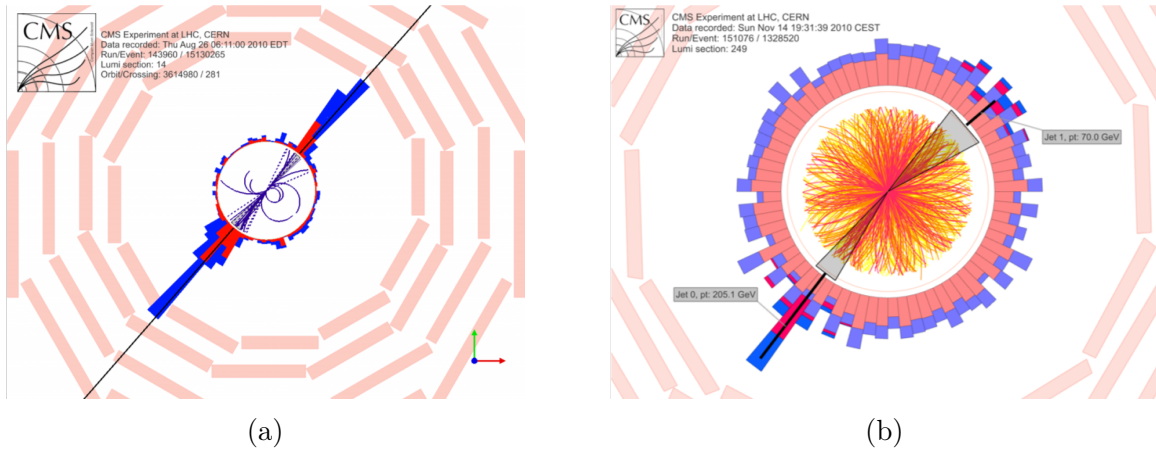


Figure 1-3: Example displays of events containing jets in (a) proton–proton and (b) heavy-ion collisions as measured by the CMS experiment at the LHC. Jets are shown as large blue and red towers. While back-to-back jets in proton–proton collisions are relatively balanced in energy due to momentum conservation, a much higher asymmetry is observed in heavy-ion collisions due to energy loss of jets in the quark–gluon plasma.

external high-energy probe (for example, an electron beam) to study its structure in the way that it is done for other materials. However, there are natural high-energy probes that are produced in the collision that can be used. In both proton–proton and heavy-ion collisions, there can be very high energy transfer interactions between constituent quarks or gluons in the participating nucleons. The result of these interactions between constituent quarks and gluons are collimated sprays of very high-energy particles, called jets, that are produced copiously in sufficiently-energetic proton–proton collisions. The production of jets in nucleus–nucleus collisions is very similar, but unlike in proton–proton collisions those sprays of particles produced in heavy-ion collisions must also punch through the quark–gluon plasma. Fig. 1-3 shows displays of events containing jets in proton–proton and heavy-ion collisions. Modifications of the yields and structure of jets in heavy-ion collisions compared to proton–proton collisions provides an important handle on the interaction of the quark–gluon plasma with high-energy particles that pass through it.


Over the years since its discovery, several important conclusions have been drawn about quark–gluon plasma:

1. **Quark–gluon plasma has an anomalously low ratio of shear viscosity**

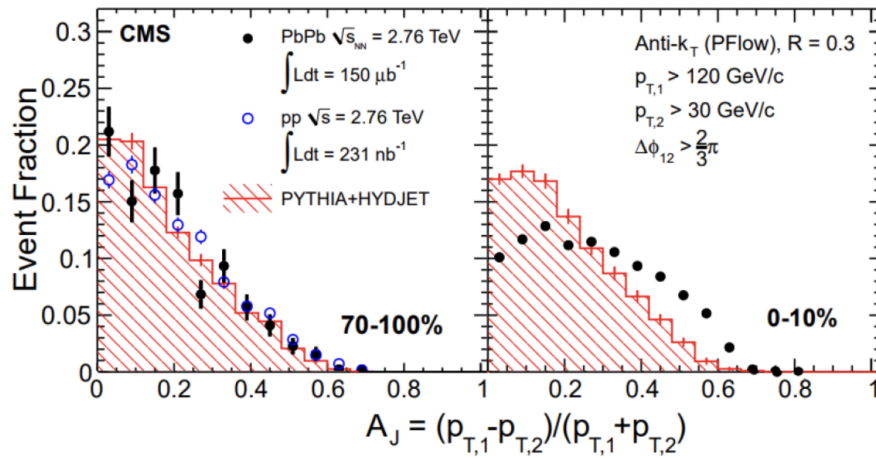
to entropy density. Fits of measured flow harmonics to sophisticated hydrodynamic simulations suggest values of $0 < \eta/s \lesssim 0.2$ in natural units (e.g., [24], [20]). This is the lowest value for this ratio observed in any fluid and suggests that quark–gluon plasma is a more perfect liquid even than, e.g., superfluid helium [25]. The shear viscosity η is actually large for quark–gluon plasma, but the dimensionless ratio η/s makes it possible to compare fluids across many orders of magnitude in the entropy density s . The uncertainty relation in quantum mechanics implies that there should be an absolute lower bound on η/s . Calculations in string theory have conjectured the bound $\eta/s \geq 1/4\pi \sim 0.08$ for any physical fluid [26], which appears to be nearly saturated by quark–gluon plasma. The extremely low value of η/s for quark–gluon plasma suggests that its constituents are very strongly-coupled among themselves; see Ref. [25] for a review.

2. **Quark–gluon plasma is semi-opaque to high-energy particles.** This is evidenced most cleanly by the observation of substantial suppression of the yield of high-energy jets in heavy-ion collisions compared to proton–proton collisions [27, 28, 29]. Fig. 1-4a shows measurements of the ratio of the yield of jets in heavy-ion collisions compared to proton–proton collisions at the LHC as a function of their transverse momentum p_T , normalized by the expected number of binary proton–proton collisions. The fact that this ratio is below one indicates that jets in heavy-ion collisions lose energy, so that at each p_T there are fewer jets in heavy-ion collisions compared to the proton–proton baseline. Another important signature is the observation of enhanced momentum asymmetry of back-to-back jet pairs in heavy-ion collisions compared to proton–proton collisions [18, 19, 30], indicating a mechanism for different energy loss of back-to-back jets in the quark–gluon plasma. Fig. 1-4b shows such a measurement from Ref. [31].

As with any exciting discovery, however, the discovery of quark-gluon plasma posed more questions than it solved. There remain many outstanding questions and



(a) Nuclear modification factor R_{AA} for jets in heavy-ion collisions as measured by the three LHC experiments. Figure made by Raghav Kunnawalkam Elayavalli and reproduced from Ref. [32].



(b) Enhanced asymmetry of the transverse momentum p_T of back-to-back dijet pairs in central heavy-ion collisions (right) compared to peripheral collisions (left) as measured in Ref. [31]. Figure reproduced from Ref. [33]

puzzles that motivate current and future work. In the following we will briefly discuss the context for several such questions in turn.

1. **What can be inferred about quark–gluon plasma from the success of hydrodynamic modelling?** Relativistic hydrodynamic modelling has been widely successful in describing varied facets of the phenomenology of particle production and correlations around the scale of the plasma temperature. However, hydrodynamics is formulated as a long-wavelength effective theory based on a gradient expansion around local equilibrium. The anticipated regime of applicability of hydrodynamics is therefore for systems that are close to local equilibrium. For the phenomenology of heavy-ion collisions, a hydrodynamic description of quark–gluon plasma is typically employed starting from very early times ($\sim 0.6 - 1$ fm/c) after the collision when presumably gradients are very large and the system is still very far from local equilibrium (see Ref. [20] for a review). The success of such hydrodynamic descriptions necessitates understanding the origin of this “unreasonable effectiveness” of hydrodynamics in heavy-ion collisions. In particular, does the success of hydrodynamics imply that the system is close to local equilibrium, or do the phenomenological consequences of hydrodynamics persist even when the system is not close to local equilibrium?

Crucial to this discussion is what is meant by hydrodynamics. Imagine a system that is “struck” with some external perturbation at $t = 0$. In general such a perturbation will excite many modes in the system. At late times, however, the response of the system will be governed just by the modes that are associated with conservation laws, like the conservation of energy and momentum. These are the so-called “hydrodynamic modes”. The transient modes excited at early times are “non-hydrodynamic modes” in the sense that they depend on the microscopic structure of the system and are not associated with any conservation laws. In practice, however, relativistic formulations of Navier-Stokes hydrodynamics are acausal, with the group velocity of the shear mode $v_g \sim k$

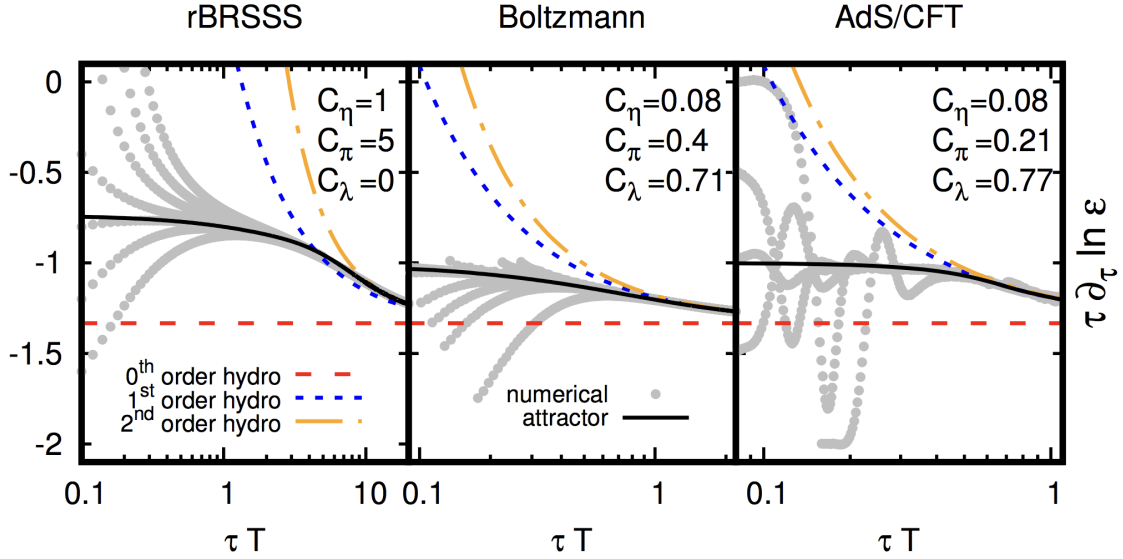


Figure 1-5: Demonstration of attractor behavior for three microscopic theories: rBRSSS (causal version of second-order hydrodynamics, see Ref. [2]), Boltzmann (kinetic theory), and AdS/CFT ($\mathcal{N} = 4$ SYM). Shown is a quantity related to the ratio of the longitudinal pressure to the energy density as a function of the proper time τ rescaled by the temperature T (the combination τT is the inverse gradient strength). Numerical solutions to the full microscopic dynamics for many initial conditions (grey) converge to the attractor solution (black) while the system is still far from the hydrodynamic limit (red dashed). Solutions to first- and second-order hydrodynamics agree quantitatively with the attractor while the system is still far from equilibrium. Figure reproduced from Ref. [3].

becoming arbitrarily large when $k \rightarrow \infty$ [20]. Since the regime of validity of hydrodynamics is small k , this is not a priori a problem, but it renders the theory unstable in practice [20]. All known stable and causal relativistic formulations of hydrodynamics contain a non-hydrodynamic mode or modes to regulate this behavior [20, 34].

In a groundbreaking paper (Ref. [35]), Heller and Spaliński showed that a system undergoing boost-invariant longitudinal (Bjorken) expansion in Müller-Israel-Stewart theory possesses an attractor solution. This attractor can be obtained by resumming the hydrodynamic gradient series to all orders, thus providing a far-from-equilibrium generalization of hydrodynamics. They observed that different initial conditions are described by a universal curve (the attractor) starting at early times when the system is still far from equilibrium. They

showed that in this case, solutions of first- and second-order hydrodynamics agree with the attractor while the system is still relatively far from equilibrium. Ref. [3] demonstrated similar behavior for conformal systems undergoing Bjorken expansion in several other microscopic theories, as seen in Fig. 1-5 (these results are particularly striking since the microscopic theories under consideration include both a weakly-coupled kinetic theory calculation and an infinitely strongly-coupled calculation). These and similar observations have suggested the possibility that hydrodynamics may provide an accurate phenomenological description outside of its naïve regime of applicability. However, the theoretical reason for this observation, and whether it extends to general systems (e.g. with less symmetry), remains an open question.

In Chapter 3 we present a novel perspective on the physical intuition of the attractor based on the adiabatic theorem in quantum mechanics. We show that, at early times, very rapid longitudinal expansion drives the system into the instantaneous ground state of an effective Hamiltonian describing its evolution. Once in this instantaneous ground state, the system evolves “adiabatically” in the sense that its evolution remains described by the (evolving) instantaneous ground state. In a simple model, we argue that the far-from-equilibrium evolution is described by a reduced set of “pre-hydrodynamic” modes, corresponding to the instantaneous ground state, that evolve into hydrodynamic modes in the hydrodynamic limit.

2. **What is the microscopic structure of quark–gluon plasma?** As discussed previously, measurements of flow coefficients combined with hydrodynamic modelling suggest that quark–gluon plasma is the most strongly-interacting fluid ever observed, which implies a lack of well-defined quasiparticle degrees of freedom. Since QCD is strongly coupled around the scale of the plasma temperature this may not be shocking. However, a defining property of QCD is that it is asymptotically free at high energies, meaning that interactions between quarks and gluons with large enough energy transfer should be weakly coupled.

This suggests that, if probed with high energy, the constituents of quark–gluon plasma should be weakly-coupled quarks and gluons. This is one motivation for using high-energy processes like jets to study the structure of quark–gluon plasma as a function of energy transfer (see, e.g., Ref. [33] for a review).

The suppression of jets in heavy-ion collisions compared to proton–proton collisions has been observed (e.g., Fig. 1-4a) implying that jets lose energy in the quark–gluon plasma. Many other measurements of the properties of jets in proton–proton and heavy-ion collisions suggest that jets are also modified in other ways (see Ref. [32] for a review). However, it has thus far proven difficult to make quantitative statements about quark–gluon plasma using jet measurements.

One reason for this is that the direct interpretation of jet modification measurements is obfuscated by multiple effects. Ideally, one would like to probe a material by first measuring the properties of the probe (e.g., its energy and direction), then shooting it through the material and measuring how its properties are modified. In heavy-ion collisions, however, there is commonly little or no direct information available about the properties of an individual probe *before* it was modified by the quark–gluon plasma. The main tool used to study jet modification is thus to compare the distributions of jet properties in heavy-ion collisions to a baseline of jet properties in proton–proton collisions. Energy loss itself makes this comparison complicated: most production properties of jets depend strongly on their energy, and so jets in proton–proton and heavy-ion collisions that have the same energy (in the heavy-ion case, after energy loss) had different properties when they were produced. This means that the “modification” that is measured includes both genuine modification of jets by the plasma and effects from the different samples of jets probed in the two cases.

Measurements being difficult to interpret directly is not necessarily a problem if there are reliable models that can incorporate any biases present in the measurement also in their predictions. Unfortunately, modelling the interaction

of jets with quark–gluon plasma is complicated and there is no first-principles solution. Jets are complicated and varied objects already in proton–proton collisions, and models of jet modification must provide a parametrization for how these objects develop in the presence of a fluctuating and highly-dynamical expanding droplet of quark–gluon plasma. Because QCD is strongly-coupled near the plasma temperature and weakly-coupled at high energies, both strong- and weak-coupling calculations fall short of describing the dynamics correctly at all relevant scales for the interaction of jets with the medium. Energy lost by a jet is deposited as a highly non-equilibrium perturbation to the plasma evolution. Experimentally it is only possible in very specific situations to distinguish particles that came from the high-energy fragmentation from particles from the plasma that are correlated with the jet, so the response of the medium generally impacts jet modification measurements.

Unfortunately, current measurements have not thus far been able to distinguish models with very different underlying models of the physics of the jet-medium interaction [32]. Part of the issue may be that the biases in jet modification observables discussed above result in measurements being most sensitive to jets that were least modified by the plasma, and therefore carry relatively little information about the jet-medium interaction.

There is substantial evidence that these effects can impact even qualitative interpretations of jet modification measurements. As an example, it has been observed in Ref. [36] that the core of heavy-ion jets are on average narrower than that of proton–proton jets with the same transverse momentum. This observation called into question whether this average narrowing implies that individual jets narrow from their interaction with quark–gluon plasma. Both in earlier work [37] and in work done in this Ph.D. [38, 39], it was shown that a narrowing of the average jet shape also happens in a model where every individual jet widens through its interaction with the plasma. Because of the bias, the average jet shape is dominated by those (typically narrow) jets that

lost little energy, and therefore can narrow due to sample migration even though no individual jet narrows.

Taken together, the presence of biases and the lack of availability of first-principles models make it particularly crucial to develop data analysis methods that enable jet modification measurements to be interpreted in a more model-independent way. In Chapter 4 we discuss a data-driven technique to alleviate some of these interpretation issues. We develop a new method to reduce the impact of energy loss on jet modification measurements by comparing proton–proton and heavy-ion jets in the same quantile of their production spectrum, instead of at the same energy. This corrects for the bias due to the average energy loss of jets of a particular energy, and enhances the sensitivity of measurements to jets that were more modified by the quark–gluon plasma.

A central question about the quark–gluon plasma produced in heavy-ion collisions is on what scale does it resolve the QCD color charge. Jets are an important tool to answer this question because they can be initiated by either a quark or gluon in the initial high-energy scattering, and subsequently will carry different total color charge. If the quark–gluon plasma sees a jet as a single extended object, then the ratio of energy loss of quark and gluon jets is expected to be equal to the ratio of their color charges. Direct measurements of different energy loss of quark and gluon jets have thus far proven elusive, however, because all jet measurements include (generally unknown) contributions from both types of jets. In Chapter 5 we illustrate a data-driven method to estimate the fraction of quark- and gluon-initiated jets and their separate modification in heavy-ion collisions. By separating quark and gluon contributions to jet observables via the method we show, it may be possible experimentally to determine whether they are modified differently by quark–gluon plasma.

3. **What is the phase diagram of Quantum Chromodynamics?** Fig. 1-1 shows a schematic of the phase diagram of QCD as a function of the baryon number density μ_B and temperature T . At zero baryon density it is possible

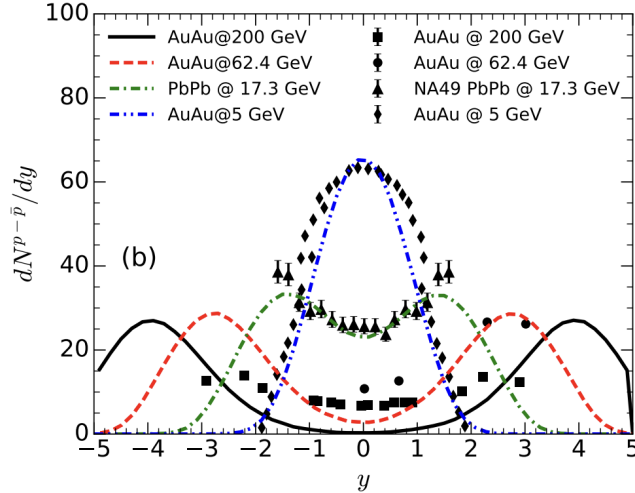


Figure 1-6: Net-proton number (a proxy for baryon number density) as a function of the rapidity y for several different energies in the range accessible at the RHIC beam energy scan. At low energies the baryon density is at midrapidity, and with increasing energy it moves toward successively higher rapidity. At LHC energies, which are substantially higher energy than any shown on this plot, the baryon density is at much higher rapidities and is essentially zero at midrapidity. Curves are hydrodynamic simulations from Ref. [4] and the figure is reproduced from Ref. [4].

to do first-principles calculations of QCD thermodynamics using lattice QCD simulations, which find that the transition from hadrons to quark–gluon plasma is a crossover [40]. At high densities and zero temperature, many models agree that the transition is first-order [41]. Other features shown on the phase diagram are conjecture, in particular the critical point shown in yellow that ends the first-order line. The reason for this is that at finite temperature and density lattice QCD calculations are currently infeasible due to the sign problem and so there are no first-principles QCD calculations. Simple model calculations generally have a critical point ending the first-order line, but disagree drastically on its location in the (T, μ_B) plane (see Ref. [41] and references therein). This motivates exploring the phase diagram experimentally using heavy-ion collision experiments.

A heavy-ion collision can be thought of as probing a trajectory through the phase diagram, since a droplet of quark–gluon plasma is created in a heavy-ion collision and undergoes a transition to hadrons as it expands and cools in the




Figure 1-7: (a) A density plot showing the regions of positive (blue) and negative (red) kurtosis as a function of the Ising phase diagram variables (t, H). The dashed green line shows a possible freezeout curve in this phase diagram. The kurtosis along this line is shown in (b). Qualitatively, the non-monotonicity and the sign change of the kurtosis as a function of t shown are anticipated signatures of the criticality as a function of μ_B . Reproduced from Ref. [5]

aftermath of the collision [42]. As illustrated in Fig. 1-1, the LHC experiments probe the low baryon density region of the phase diagram. At the high beam energies at the LHC, the baryons involved in the collision pass through each other and end up at very large rapidity, outside of the region accessible in measurements. In lower energy experiments like (particularly low energies) at RHIC, there are more baryons at midrapidity where measurements take place [43]. In fact, the baryon density at midrapidity in a heavy-ion collision can be tuned by changing the beam energy [44], as can be seen quantitatively in Fig. 1-6. This is the conceptual basis of the Beam Energy Scan experiment at RHIC, which aims to explore the phase diagram at higher μ_B by lowering the beam energy, with the goal of discovering (or ruling out) the location of the conjectured critical point; see Ref. [45].

A characteristic feature of a critical point in any system is the divergence of the correlation length. This divergence also causes a divergence in many fluctuation measures, for example the event-by-event fluctuations of particle multiplicities, near a critical point. These fluctuation measures are a crucial signature of the

QCD critical point that are accessible experimentally (e.g., Ref. [44]). Though theoretically the correlation length diverges near a critical point, in practice this divergence is cut off by both the finite system size and by critical slowing down, e.g. Ref. [46], and so neither the correlation nor the fluctuations grow without bound, potentially making the signatures more difficult to detect experimentally. Ref. [47] showed that non-Gaussian moments of these fluctuations scale with higher powers of the correlation length compared to Gaussian moments, and therefore are more sensitive to critical behavior. However, increasingly non-Gaussian moments also become increasingly difficult to measure experimentally; the kurtosis of the event-by-event fluctuations of the net proton multiplicity is a famous observable used to search for criticality at the Beam Energy Scan [48]. Non-monotonic dependence of the kurtosis on the beam energy (namely, baryon density) is a characteristic signature of critical behavior anticipated at the Beam Energy Scan. It was also shown in Ref. [5] that the kurtosis changes sign near a QCD critical point. A schematic picture of this is shown in Fig. 1-7. The non-monotonicity and the sign change of the kurtosis shown as a function of the Ising variable t is also anticipated in the QCD phase diagram as a function of μ_B .

In Chapter 6 we argue that, since the baryon density also has dependence on rapidity at fixed beam energy, a rapidity scan may be a valuable additional probe of critical behavior to the beam energy scan. We show that the rapidity dependence of fluctuations may have qualitatively similar features to those of the beam energy dependence, including the non-monotonicity and the sign change. If a rapidity scan is performed at each beam energy, it may also be possible to distinguish the location of a critical point lying between two beam energies.

In this thesis, we present recent progress on understanding quark–gluon plasma produced in heavy-ion collisions from three unique perspectives: its fluid-like behavior, quenching of high-energy particles, and its phase diagram. These three vantage points are foundational to how quark–gluon plasma is studied experimentally, phe-

nomenologically, and theoretically.

In Chapter 3, we develop a novel conjecture that the apparent phenomenological applicability of hydrodynamics at early times may be due to the emergent dominance of a reduced set of “pre-hydrodynamic” modes that evolve into hydrodynamic modes at late times. If confirmed in more realistic situations, this idea could be pivotal. It would suggest that far-from-equilibrium behavior in the quark–gluon plasma could be described by the substantially more simple physics of a small number of slow pre-hydrodynamic modes, rather than being sensitive to the full physics of the underlying microscopic description. In far-from-equilibrium situations, it may be more useful to formulate an expansion in the contributions from pre-hydrodynamic and higher modes than in terms of a gradient expansion. It would also imply that the phenomenological utility of hydrodynamic modelling does not necessarily imply that hydrodynamic modes are dominant.

In Chapter 4, we propose a method to compare proton–proton and heavy-ion jets that were produced with more similar transverse momentum p_T by comparing jets in the same quantile of the jet production spectrum, rather than at the same p_T . This makes it possible to correct for the average energy loss of jets in heavy-ion collisions when making jet modification measurements, and enables a measurement of the average energy loss of jets as a function of p_T . This method may be crucial in the interpretation of jet modification since it enhances the sensitivity of measurements to those jets that lost more energy to the plasma. In the future, we hope it may also be possible to extend this type of method to approximate the initial jet p_T more accurately than by the average energy loss of jets with that p_T , for example by using other information about the jet.

In Chapter 5 we propose a method to use measurable jet distributions, which are unknown mixtures of both quark- and gluon-initiated jets, to estimate the fraction of quark and gluon jets and their separate modification in heavy-ion collisions. This technique may open a door toward estimating the fraction of quark- and gluon-initiated jets and their separate modification in heavy-ion collisions. This is crucial to a quantitative understanding of differences in quark and gluon jet energy loss in

the quark–gluon plasma. Since quark and gluon jet fractions are different in heavy-ion collisions than in proton–proton collisions, separating the fraction modification from the distribution modification may also be important to the interpretation of jet modification observables.

In Chapter 6, we suggest that the Beam Energy Scan at RHIC could be supplemented by an additional rapidity scan which probes the more local structure of the phase diagram at each beam energy. This may provide signatures of critical behavior in the rapidity-dependence of cumulants, in addition to those anticipated in the beam energy-dependence. Particularly if the critical point lies between two of the beam energies possible at the Beam Energy Scan, this could be key to the observation of critical behavior in heavy-ion collision experiments.

This thesis only includes a subset of the work completed in these directions during this Ph.D., so we further refer the interested reader to Chapter 2 for abstracts of additional research.

Chapter 2

Overview

In the course of my Ph.D. I have completed six papers with various different collaborations, and one additional paper which is in the final stages of preparation but has yet to appear. Though this thesis will focus on a subset of this work, all are listed below with abstracts.

2.1 This Thesis

Adiabatic hydrodynamization in rapidly-expanding quark-gluon plasma [49]

In collaboration with Li Yan and Yi Yin

We propose a new scenario characterizing the transition of the quark–gluon plasma (QGP) produced in heavy-ion collisions from a highly non-equilibrium state at early times toward a fluid described by hydrodynamics at late times. In this scenario, the bulk evolution is governed by a set of slow modes, after an emergent time scale τ_{Redu} when the number of modes that govern the bulk evolution of the system is reduced. These slow modes are “pre-hydrodynamic” in the sense that they are initially distinct from, but evolve continuously into, hydrodynamic modes in hydrodynamic limit. This picture is analogous to the evolution of a quantum mechanical system that is governed by the instantaneous ground states under adiabatic evolution, and

will be referred to as “adiabatic hydrodynamization”. We shall illustrate adiabatic hydrodynamization using a kinetic description of weakly-coupled Bjorken expanding plasma. We first show the emergence of τ_{Redu} due to the longitudinal expansion. We explicitly identify the pre-hydrodynamic modes for a class of collision integrals and find that they represent the angular distribution (in momentum space) of those gluons that carry most of the energy. We use the relaxation time approximation for the collision integral to show quantitatively that the full kinetic theory evolution is indeed dominated by pre-hydrodynamic modes. We elaborate on the criterion for the dominance of pre-hydrodynamic modes and argue that the rapidly-expanding QGP could meet this criterion. Based on this discussion, we speculate that adiabatic hydrodynamization may describe the pre-equilibrium behavior of the QGP produced in heavy-ion collisions.

Sorting out quenched jets [50]

In collaboration with Guilherme Milhano and Jesse Thaler

We introduce a new “quantile” analysis strategy to study the modification of jets as they traverse through a droplet of quark-gluon plasma. To date, most jet modification studies have been based on comparing the jet properties measured in heavy-ion collisions to a proton-proton baseline at the same reconstructed jet transverse momentum (p_T). It is well known, however, that the quenching of jets from their interaction with the medium leads to a migration of jets from higher to lower p_T , making it challenging to directly infer the degree and mechanism of jet energy loss. Our proposed quantile matching procedure is inspired by (but not reliant on) the approximate monotonicity of energy loss in the jet p_T . In this strategy, jets in heavy-ion collisions ordered by p_T are viewed as modified versions of the same number of highest-energy jets in proton-proton collisions, and the fractional energy loss as a function of jet p_T is a natural observable (Q_{AA}). Furthermore, despite non-monotonic fluctuations in the energy loss, we use an event generator to validate the strong correlation between the p_T of the parton that initiates a heavy-ion jet and the p_T of the vacuum jet which corresponds to it via the quantile procedure (p_T^{quant}). We demonstrate that this strat-

egy both provides a complementary way to study jet modification and mitigates the effect of p_T migration in heavy-ion collisions.

Data-driven quark and gluon jet modification in heavy-ion collisions

In collaboration with Jesse Thaler and Andrew Patrick Turner; to appear

Whether quark- and gluon-initiated jets are modified differently by the quark–gluon plasma produced in heavy-ion collisions is a long-standing question that has thus far eluded a definitive experimental answer. A crucial complication for quark–gluon discrimination in both proton–proton and heavy-ion collisions is that all measurements necessarily average over the (unknown) quark–gluon composition of a jet sample. In the heavy-ion context, the simultaneous modification of both the fractions and substructure of quark and gluon jets by the quark–gluon plasma further obscures the interpretation. Here, we demonstrate a fully data-driven method for separating quark and gluon contributions to jet observables using a statistical technique called topic modeling. Assuming that jet distributions are a mixture of underlying “quark-like” and “gluon-like” distributions, we show how to extract quark and gluon jet fractions and constituent multiplicity distributions as a function of the jet transverse momentum. This proof-of-concept study is based on proton–proton and heavy-ion collision events from the Monte Carlo event generator JEWEL with statistics accessible in Run 4 of the Large Hadron Collider. These results suggest the potential for an experimental determination of quark and gluon jet modifications.

Searching for the QCD critical point via the rapidity dependence of cumulants [51]

In collaboration with Swagato Mukherjee, Krishna Rajagopal, and Yi Yin

The search for a possible critical point in the QCD phase diagram is ongoing in heavy ion collision experiments at RHIC which scan the phase diagram by scanning the beam energy; a coming upgrade will increase the luminosity and extend the rapid-

ity acceptance of the STAR detector. In fireballs produced in RHIC collisions, the baryon density depends on rapidity. By employing Ising universality together with a phenomenologically motivated freezeout prescription, we show that the resulting rapidity dependence of cumulant observables sensitive to critical fluctuations is distinctive. The dependence of the kurtosis (of the event-by-event distribution of the number of protons) on rapidity near mid-rapidity will change qualitatively if a critical point is passed in the scan. Hence, measuring the rapidity dependence of cumulant observables can enhance the prospect of discovering a critical point, in particular if it lies between two energies in the beam energy scan.

2.2 Additional work

Jet shape modifications in holographic dijet systems [52]

In collaboration with Andrey Sadofyev and Wilke van der Schee

We present a coherent model that combines jet production from perturbative QCD with strongly-coupled jet-medium interactions described in holography. We use this model to study the modification of an ensemble of jets upon propagation through a quark-gluon plasma resembling central heavy ion collisions. Here the modification of the dijet asymmetry depends strongly on the subleading jet width, which can therefore be an important observable for studying jet-medium interactions. We furthermore show that the modification of the shape of the leading jet is relatively insensitive to the dijet asymmetry, whereas the subleading jet shape modification is much larger for more imbalanced dijets.

Evolution of the Mean Jet Shape and Dijet Asymmetry Distribution of an Ensemble of Holographic Jets in Strongly Coupled Plasma [38]

(Including proceedings, Holographic Jet Shapes and their Evolution in Strongly Coupled Plasma [39])

In collaboration with Krishna Rajagopal, Andrey Sadofyev, and Wilke van der Schee

Some of the most important probes of the quark-gluon plasma (QGP) produced in heavy ion collisions come from the analysis of how the shape and energy of jets are modified by passage through QGP. We model an ensemble of back-to-back dijets to gain a qualitative understanding of how the shapes of the individual jets and the asymmetry in the energy of the pairs of jets are modified by passage through an expanding droplet of strongly coupled plasma, as modeled in a holographic gauge theory. We do so by constructing an ensemble of strings in the gravitational description of the gauge theory. We model QCD jets in vacuum using strings whose endpoints move “downward” into the gravitational bulk spacetime with some fixed small angle that represents the opening angle (ratio of jet mass to jet energy) that the QCD jet would have in vacuum. Such strings must be moving through the gravitational bulk at (close to) the speed of light; they must be (close to) null. This condition does not specify the energy distribution along the string, meaning that it does not specify the shape of the jet being modeled. We study the dynamics of strings that are initially not null and show that strings with a wide range of initial conditions rapidly accelerate and become null and, as they do, develop a similar distribution of their energy density. We use this distribution of the energy density along the string, choose an ensemble of strings whose opening angles and energies are distributed as in perturbative QCD, and show that we can then fix one model parameter such that the mean jet shape in our ensemble matches that measured in p-p collisions reasonably well. We send our strings through the plasma, choosing the second model parameter to get a reasonable suppression in the number of jets, and study how the mean jet shape and the dijet asymmetry are modified, comparing both to data from LHC heavy ion collisions.

String-theory-based predictions for nonhydrodynamic collective modes in strongly interacting Fermi gases [53]

In collaboration with Hans Bantilan, Takaaki Ishii, William Lewis, and Paul Romatschke

Very different strongly interacting quantum systems such as Fermi gases, quark-gluon plasmas formed in high energy ion collisions and black holes studied theoretically in string theory are known to exhibit quantitatively similar damping of hydrodynamic modes. It is not known if such similarities extend beyond the hydrodynamic limit. Do non-hydrodynamic collective modes in Fermi gases with strong interactions also match those from string theory calculations? In order to answer this question, we use calculations based on string theory to make predictions for novel types of modes outside the hydrodynamic regime in trapped Fermi gases. These predictions are amenable to direct testing with current state-of-the-art cold atom experiments.

Chapter 3

Adiabatic Hydrodynamization

The work reported in this chapter is based on work done in collaboration with Li Yan and Yi Yin that is currently under review [49]. The abstract can be found in Section 2.1.

Hydrodynamics describes the real-time dynamics of a broad class of interacting many-body systems in the long time and long wavelength limit. In this limit, most degrees of freedom become irrelevant since they relax on short time scales. The surviving slow dynamical variables, or “hydrodynamic modes”, are those associated with conserved densities such as the energy density. Hydrodynamic modelling has seen remarkable success at describing varied and non-trivial results of heavy-ion collision experiments (see Ref. [54] for a concise review). This in turn raises the important question of how the system approaches a state dominated by hydrodynamic modes, namely how “hydrodynamization” occurs in the aftermath of a heavy-ion collision (cf. [20, 55, 56] for a recent review).

In this Chapter, we theorize a new scenario for the process of hydrodynamization with the following defining attribute: *during the interval $\tau_{\text{Redu}} < \tau < \tau_{\text{Hydro}}$, the bulk evolution is governed by a set of slow modes that are “pre-hydrodynamic” in the sense that they are distinct from hydrodynamic modes but evolve gradually into them around the time τ_{Hydro} .* As a premise of this picture, we assume the emergence of a time scale $\tau_{\text{Redu}} < \tau_{\text{Hydro}}$ around which the degrees of freedom required to describe the bulk properties of the system are reduced (see more below).

The pre-hydrodynamic modes in the preceding scenario are the modes with the slowest rate of change at each instant in the pre-hydrodynamic evolution, under the assumption that they remain gapped from faster modes. They are closely analogous to the instantaneous ground states of a time-dependent Hamiltonian in quantum mechanics. Since near thermal equilibrium the hydrodynamic modes are the slowest modes, the pre-hydrodynamic modes are a natural off-equilibrium generalization of the hydrodynamic modes. If a time-dependent and gapped quantum-mechanical system is prepared in its ground state, it will remain in the instantaneous ground state under adiabatic evolution of the Hamiltonian. We will thus refer to situations where the approach to hydrodynamics is governed first by the evolution of pre-hydrodynamic modes as “adiabatic hydrodynamization”.

We will illustrate adiabatic hydrodynamization in a kinetic description of weakly-coupled Bjorken-expanding plasma. We explicitly identify the pre-hydrodynamic modes as the instantaneous ground state modes of a non-Hermitian matrix describing the evolution of bulk quantities from the kinetic equation with a class of collision integrals. Physically, these modes represent the angular distribution in momentum space of the gluons that carry most of the energy of the system. We then demonstrate the emergence of τ_{Redu} induced by the fast longitudinal expansion and show that τ_{Redu} is parametrically smaller than τ_{C} . This is due to the separation of scales between the initial time τ_I when the kinetic description becomes applicable and the typical collision time τ_{C} for QGP in the weak-coupling regime. Because of the hierarchy $\tau_{\text{Redu}} \ll \tau_{\text{C}} \ll \tau_{\text{Hydro}}$, the pre-history of hydrodynamics is (almost) the history of pre-hydrodynamic modes within adiabatic hydrodynamization.

An important implication of adiabatic hydrodynamization is that the macroscopic properties of the medium during the pre-hydrodynamic stage are insensitive to both the initial conditions and the details of the expansion history, and instead are determined predominantly by the features of the pre-hydrodynamic modes. In particular, the most important quantity characterizing the bulk evolution of a plasma undergoing

Bjorken expansion is the percentage rate of change of the energy density

$$g(y) \equiv -\partial_y \log \epsilon, \quad (3.1)$$

where $y \equiv \log(\tau/\tau_I)$ plays the role of a time variable. We shall show that $g(y)$ is related to the eigenvalue $\mathcal{E}_0(y)$ of the pre-hydrodynamic mode if adiabatic hydrodynamization applies, namely

$$g(y) \approx \mathcal{E}_0(y). \quad (3.2)$$

We consider the extensively-studied relaxation time approximation (RTA) of the kinetic equation [57, 58, 6, 59, 60, 7] and confirm quantitatively that Eq. (3.2) holds, demonstrating that hydrodynamization in this model is an example of adiabatic hydrodynamization.

Because of the expansion, the criterion for the dominance of pre-hydrodynamic modes is not that the excited states have decayed, but rather that transitions to the excited states are suppressed. In the absence of better terminology, throughout this manuscript we will use “adiabaticity” as a synonym for the suppression of these transitions. This is consistent with the modern use of this terminology in quantum mechanics (c.f. Ref. [61]). We will show that the regime where this generalized notion of adiabaticity may not apply is parametrically narrow according to the scenario of bottom-up thermalization for weakly-coupled QGP [62]. Although our analysis relies on the smallness of α_s , we hope that many qualitative features of adiabatic hydrodynamization may nonetheless be present in the QGP created in heavy-ion collisions.

There are extensive studies on the formulation of far-from-equilibrium hydrodynamics to describe the pre-hydrodynamic stage of heavy-ion collisions [3, 35, 63, 64, 60, 65, 66]. The key premise of this paradigm is that hydrodynamic modes dominate the bulk evolution, and consequently that hydrodynamics is applicable, even when the system is far from equilibrium [3]. The difference between this paradigm and adiabatic hydrodynamization is that the dominant slow modes for systems undergoing

adiabatic hydrodynamization are pre-hydrodynamic mode(s), which can be qualitatively distinct from hydrodynamic modes. The slow modes of a system generally depend on the state of the medium under consideration, and therefore it is unsurprising that the slow modes in a far-from-equilibrium system are generically different from the hydrodynamic modes. We show that RTA kinetic theory is an example where pre-hydrodynamic and hydrodynamic modes are qualitatively different, and the bulk evolution of the system is dominated by the pre-hydrodynamic modes.

In the modern view, hydrodynamics is a macroscopic effective theory in which hydrodynamic modes are the relevant low energy degrees of freedom. In cases where the relevant degrees of freedom are actually pre-hydrodynamic modes, there is no guarantee that hydrodynamics or its simple generalizations will describe the system. This is not in contradiction to the recent result [64, 67] that some non-trivial generalizations of hydrodynamics like an improved version of Israel-Stewart theory [68] and anisotropic and third-order hydrodynamics describe the bulk evolution of several simplified kinetic theory models even beginning at τ_{Redu} . Rather, since these models include significant contributions from non-hydrodynamic modes, we emphasize that this observation alone does not imply that hydrodynamic modes dominate the evolution. It is worth exploring the applicability of these models in more general settings, however we hope that the identification of pre-hydrodynamic modes as a relevant slow degree of freedom may motivate the future construction of an effective theory of “pre-hydrodynamics”.

3.1 Identification of pre-hydrodynamic mode(s)

We consider a Bjorken-expanding medium of massless particles described by the kinetic equation

$$\frac{\partial}{\partial \tau} f(p_z, p_\perp; \tau) = -\frac{p_z}{\tau} \frac{\partial}{\partial p_z} f(p_z, p_\perp; \tau) - \hat{C}[f], \quad (3.3)$$

where $f(p_z, p_\perp; \tau)$ is the single particle distribution, p_\perp and p_z are the transverse and longitudinal momentum, and \hat{C} is the collision integral. Because of the symmetry, the only relevant hydrodynamic mode is the energy density ϵ . To more directly study the evolution of the energy density, we will focus on the momentum-weighted distribution function

$$F_\epsilon(\cos\theta; \tau) \equiv \frac{1}{2\pi^2} \int_0^\infty dp p^3 f(p_z, p_\perp; \tau), \quad (3.4)$$

where $p = \sqrt{p_\perp^2 + p_z^2}$ and $\theta = \tan^{-1}(p_z/p_\perp)$. Because the angular integration of $F_\epsilon(\cos\theta; \tau)$ is the energy density, $F_\epsilon(\cos\theta; \tau)$ describes the angular distribution of the particles that carry most of the energy. The p^3 -weighted moment of Eq. (3.3) gives the evolution equation for F_ϵ :

$$\begin{aligned} \tau \frac{\partial}{\partial \tau} F_\epsilon(\cos\theta; \tau) = & - \left[-4 \cos^2 \theta + \sin^2 \theta \cos \theta \frac{\partial}{\partial \cos \theta} \right] \\ & \times F_\epsilon(\cos\theta; \tau) - \frac{\tau}{2\pi^2} \int_0^\infty dp p^3 \hat{C}[f]. \end{aligned} \quad (3.5)$$

Following Ref. [69], we assume that $f(p_z, p_\perp; \tau)$ is symmetric under $p_z \rightarrow -p_z$ and expand $F_\epsilon(\cos\theta; \tau)$ in a basis of the Legendre polynomials P_{2n} :

$$F_\epsilon(\cos\theta; \tau) = \epsilon(\tau) + \sum_{n=1} \frac{4n+1}{2} \mathcal{L}_n(\tau) P_{2n}(\cos\theta). \quad (3.6)$$

Eq. (3.6) maps $F_\epsilon(\cos\theta; \tau)$ to an infinite-dimensional vector $\psi = (\epsilon, \mathcal{L}_1, \mathcal{L}_2, \dots)$. We therefore have the correspondence

$$F_\epsilon(\cos\theta) \leftrightarrow \psi = (\epsilon, \mathcal{L}_1, \mathcal{L}_2, \dots). \quad (3.7)$$

Note $p_L = \frac{1}{3}(\epsilon + 2\mathcal{L}_1)$. Since F_ϵ will become isotropic and approach ϵ in the hydrodynamic limit, the hydrodynamic mode corresponds to the vector

$$\phi_0^H = (\epsilon, 0, 0, \dots). \quad (3.8)$$

The problem of hydrodynamization is therefore reduced to studying how ψ evolves

into ϕ_0^H .

In the following discussion, we shall limit ourselves to the class of collision integrals for which Eq. (3.5) can be recast into the form

$$\partial_y \psi = -\mathcal{H}(y) \psi, \quad (3.9)$$

where \mathcal{H} is a non-Hermitian matrix and $y = \log(\tau/\tau_I)$ as we introduced earlier. This is satisfied for any collision integral that is linear in $F_\epsilon(\cos\theta; \tau)$. Eq. (3.9) has the structure of the time-dependent Schrödinger equation in quantum mechanics. The explicit expression for the matrix \mathcal{H} for RTA kinetic theory will be given in the subsequent section.

Throughout this work, we will study the instantaneous eigenmodes $\phi_n(y)$ of $\mathcal{H}(y)$. For clarity we will order them by the real part of their corresponding eigenvalues, e.g. $\text{Re}\mathcal{E}_0 < \text{Re}\mathcal{E}_1 \leq \dots$. Of particular importance is the ground state mode $\phi_0(y)$, which has the lowest damping rate of all of the eigenmodes. In the hydrodynamic limit $\tau \geq \tau_{\text{Hydro}}$, $\phi_0(y)$ will evolve into ϕ_0^H since the conserved densities are the zero-modes of any collision kernel. At times $\tau < \tau_{\text{Hydro}}$, we identify $\phi_0(y)$ as the “pre-hydrodynamic mode” since it is an ancestor to the hydrodynamic mode ϕ_0^H .

3.2 Emergent dominance of pre-hydrodynamic modes

at $\tau \ll \tau_C$

To illustrate the reduction in the degrees of freedom at early times, we study the behavior of ψ for $\tau \ll \tau_C$. In this case, ψ is determined by $\partial_y \psi = -\mathcal{H}_F \psi$ where \mathcal{H}_F is obtained from Eq. (3.5) by neglecting the collision integral (the explicit expression can be obtained from the $\tau \rightarrow 0$ limit of Eq. (3.13)). To solve, we expand ψ in eigenstates of \mathcal{H}_F as $\psi(\tau) = \sum_{n=0} \beta_n(\tau) \phi_n^F$. It is easy to show that $\beta_n(\tau) = \beta_n(\tau_I) \exp(-\mathcal{E}_n y)$ for all n . Therefore contributions from the “excited” modes $\phi_{n>0}^F$ become unimportant

after some emergent time scale

$$\tau_{\text{Redu}} \gtrsim \tau_I \left| \left(\frac{\beta_n(\tau_I)}{\beta_0(\tau_I)} \right)^{1/(\mathcal{E}_{n>0}^F - \mathcal{E}_0^F)} \right|. \quad (3.10)$$

The bulk evolution of the system around τ_{Redu} is then dominated by the ground state mode ϕ_0^F . Related observations have also been made in Refs. [64, 70].

For the description of heavy-ion collisions in the framework of perturbative QCD, τ_I is of the order of Q_s^{-1} , where $Q_s \gg \Lambda_{\text{QCD}}$ is the saturation scale (c.f. Refs. [71, 72, 73, 74]). Meanwhile, a parametric estimate of τ_C can be deduced from the collision integral, $\tau_C Q_s \sim \alpha_s^{-x}$ with exponent $x > 0$ (c.f. Ref. [62]). This hierarchy guarantees the existence of a time scale τ_{Redu} that is parametrically smaller than $\tau_{\text{Hydro}} \geq \tau_C$.

To appreciate the physics underlying the dominance of ϕ_0^F around τ_{Redu} , we compare the explicit expression $\phi_0^F = \epsilon(1, P_2(0), P_4(0) \dots)$ [66] with the definition in Eq. (3.6). It is then transparent that ϕ_0^F corresponds to an angular distribution function $F_\epsilon(\cos\theta; \tau)$ that is sharply peaked at $\theta = \pi/2$. For such a distribution, typical values of p_z are much smaller than those of p_\perp , meaning the longitudinal expansion drives arbitrary initial conditions to a highly anisotropic distribution in momentum space.

The analysis above shows that the longitudinal expansion together with the intrinsic hierarchy $\tau_I \ll \tau_C$ in weakly coupled QCD prepares the system in the instantaneous ground state ϕ_0^F . Since ϕ_0^F depends on \mathcal{H}_F but not on the initial conditions, the bulk evolution around τ_{Redu} becomes insensitive to the details of the initial conditions. The latter has been observed in previous studies of kinetic theory [3, 34, 66], though its connection to the dominance of the mode ϕ_0^F has not been elucidated before.

3.3 Implications of the dominance of pre-hydrodynamic modes

We now explore the implications of the adiabatic evolution of $\mathcal{H}(y)$ after τ_{Redu} . We begin by expanding ψ in terms of the instantaneous eigenmodes $\phi_n(y)$ of $\mathcal{H}(y)$, $\psi(y) =$

$\sum_{n=0} \alpha_n(y) \phi_n(y)$. While in general $\alpha_{n>0}$ can be the same order of magnitude as α_0 , under adiabatic evolution $|\alpha_0| \gg |\alpha_{n>0}|$ and consequently

$$\psi(y) \sim \phi_0(y). \quad (3.11)$$

Eq. (3.11) can be viewed as the definition of adiabatic hydrodynamization.

We emphasize that this dominance of the pre-hydrodynamic mode ϕ_0 indicates that the bulk properties of the pre-equilibrium medium can be related to this mode and its eigenvalue. For example, let us focus on the percentage rate of change of the energy density in Eq. (3.1). Since the zeroth component of ψ is ϵ , it follows from Eqns. (3.9) and (3.11) that $-\epsilon g(y)$ is given by the zeroth component of $\mathcal{H}\psi$, i.e. Eq. (3.2), even though $g(y)$ in general can depend on all modes $\phi_n(y)$. This non-trivial relation is a consequence of the adiabatic evolution. In the next section we will quantitatively test this result in the relaxation time approximation to determine the extent to which adiabatic hydrodynamization applies.

3.4 RTA as an example of adiabatic hydrodynamization

The collision integral under the relaxation-time approximation (RTA) is

$$\hat{C}[f] = \frac{f(p_z, p_\perp; \tau) - f_{\text{eq}}(p/T)}{\tau_C(y)}, \quad (3.12)$$

where τ_C is a function of y . Substituting Eqns. (3.12) and (3.6) into Eq. (3.5) gives [69]

$$\partial_y \mathcal{L}_n = -[a_n \mathcal{L}_n + b_n \mathcal{L}_{n-1} + c_n \mathcal{L}_{n+1}] - \lambda(1 - \delta_{n0}) \mathcal{L}_n, \quad (3.13)$$

where $\lambda \equiv \tau/\tau_C$. Explicit expressions for a_n, b_n , and c_n are given in Ref. [69], for example $(a_0, b_1, c_0) = (4/3, 8/15, 2/3)$. From Eq. (3.13), the evolution of ψ has the

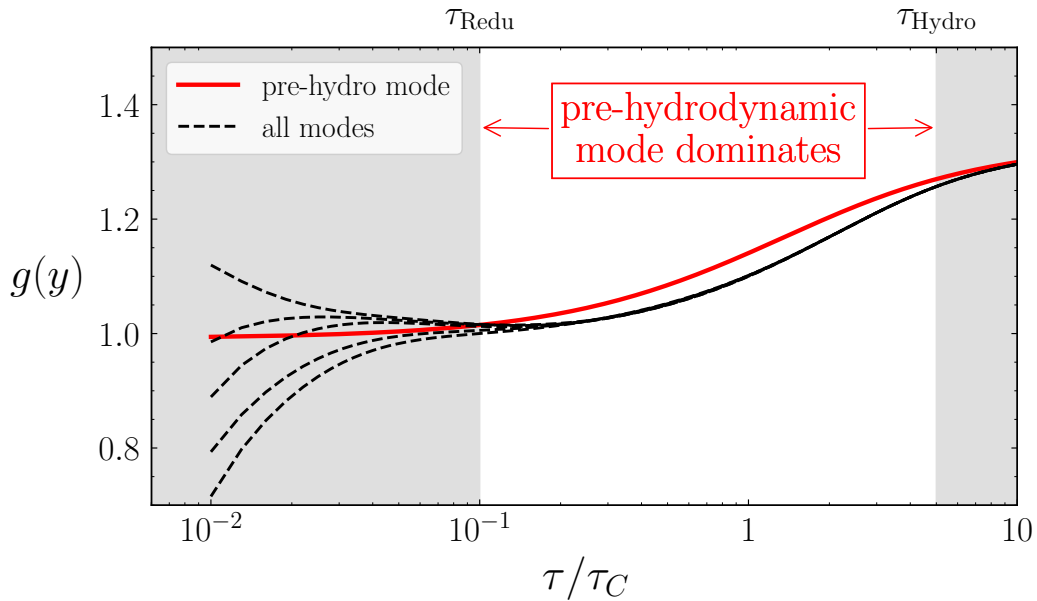


Figure 3-1: We demonstrate the dominance of pre-hydrodynamic modes, and hence adiabatic hydrodynamization, for a Bjorken-expanding plasma in the relaxation time approximation. The red solid curve shows $\mathcal{E}_0(\lambda)$, which is the contribution to the percentage rate of change of the energy density coming from the pre-hydrodynamic mode (c.f. Eq. (3.2)). Black curves are $g(\lambda)$ obtained from solving Eq. (3.3) with constant τ_C for different initial conditions. After τ_{Redu} , they collapse onto the RTA attractor as obtained in Refs. [6, 3, 7]. The left and right shaded regions indicate $\tau \leq \tau_{\text{Redu}}$ and $\tau \geq \tau_{\text{Hydro}}$, respectively.




Figure 3-2: We show the fractional difference between \mathcal{E}_0 and g , which measures the relative importance of contributions from the ground state (pre-hydrodynamic) and excited modes. Red and dashed blue curves show this difference for representative initial conditions with constant and conformal τ_C , respectively. The fact that this quantity is small indicates the dominance of pre-hydrodynamic modes (i.e. adiabaticity) during the interval $\tau_{\text{Redu}} \leq \tau < \tau_{\text{Hydro}}$.

form Eq. (3.9) with

$$\mathcal{H}(y) = \mathcal{H}_{\text{RTA}}(\lambda) \equiv \mathcal{H}_F + \lambda(y) \mathcal{H}_1, \quad (3.14)$$

where the elements of $\mathcal{H}_F, \mathcal{H}_1$ can be read from Eq. (3.13).

From $\mathcal{H}_{\text{RTA}}(\lambda)$ we compute the pre-hydrodynamic modes $\phi_0(\lambda)$ and their energies $\mathcal{E}_0(\lambda)$ for each λ . We note that the minimum gap $\Delta\mathcal{E}_{\text{min}}(\lambda) \equiv \text{Re}(\mathcal{E}_1(\lambda) - \mathcal{E}_0(\lambda))$ is order one for $\lambda \ll 1$ and becomes linear in λ for $\lambda \gg 1$ ¹. For all values of λ , $\phi_0(\lambda)$ is gapped from the excited modes. It is easy to check that ϕ_0^H is the ground state of \mathcal{H}_1 but not that of \mathcal{H}_{RTA} . Since \mathcal{H}_{RTA} evolves in time, the components of $\phi_0(y)$ are different from those of ϕ_0^H for any finite y , exemplifying the distinction between pre-hydrodynamic and hydrodynamic modes.

The solid red curve in Fig. 3-1 shows $\mathcal{E}_0(\lambda)$, which is the contribution to $g(\lambda)$ from the pre-hydrodynamic modes only². For comparison, we also determine $g(\tau)$ by solving the kinetic equation numerically. Following Ref. [7], we use the parametrization $\tau_C \propto \epsilon^{-\Delta/4}$. Solutions to Eq. (3.3) with constant τ_C ($\Delta = 0$) and different initial conditions satisfying $\tau_I \ll \tau_C$ are shown in dashed black in Fig. 3-1. In practice, we solve the coupled differential equations in Eq. (3.13), for which initial conditions on p_L/ϵ are fixed by the initial values of the moments \mathcal{L}_n . The resulting g collapses to a common curve at times much earlier than τ_{Hydro} , which is the well-known ‘‘attractor’’ behavior of Bjorken-expanding RTA kinetic theory. Remarkably, $\mathcal{E}_0(\lambda)$ is close to $g(\tau)$, indicating that the bulk evolution before the hydrodynamic regime is indeed dominated by the evolution of the pre-hydrodynamic modes.

Since the RTA attractor function $g(\tau)$ has already been obtained by many authors [6, 3, 7], what is our purpose of studying this function? Our goal is to demonstrate that the main contribution to this function comes from pre-hydrodynamic modes. We emphasize that the attractor behavior of $g(\lambda)$ alone does not tell us

¹If we were using τ instead of y as our temporal variable, the minimum gap is of the order $1/\tau$ for $\lambda \ll 1$ and of the order $1/\tau_C$ for $\lambda \gg 1$. The latter agrees with Ref. [58] although one has to keep in mind that the minimum gap also evolves in time.

²In practice, we truncate Eq. (3.13) at $n = 9$ so that \mathcal{H}_{RTA} is reduced to a 10×10 non-Hermitian matrix. We have checked that the results shown are not sensitive to the truncation.

whether one mode or many modes are important for the subsequent evolution. In the language of quantum mechanics, the attractor behavior only indicates that the system is in its instantaneous ground state around τ_{Redu} . The system remains in its instantaneous ground state here due to a qualitatively different reason, namely the suppression of transitions to the excited states.

To further demonstrate that the relative importance of contributions from excited modes are suppressed compared to those of the pre-hydrodynamic mode, we show the fractional difference $\delta \equiv |g - \mathcal{E}_0|/g$ as a function of λ in Fig. 3-2. Results for constant τ_C ($\Delta = 0$) and conformal τ_C ($\Delta = 1$) are shown in red and dashed blue, respectively. The evolution is more adiabatic the smaller δ is. δ is small both when $\lambda \ll 1$ and $\lambda \gg 1$ and reaches a maximum of 0.045 at intermediate λ . This indicates that at least 95% of the contribution to g between τ_{Redu} and τ_{Hydro} is from pre-hydrodynamic modes. We emphasize that δ is small even when the Knudsen number $1/\lambda$ is large. In fact, the contributions from the excited modes $\phi_{n>0}$ can be accounted for systematically by expanding in δ , generalizing the method developed in Refs. [75, 61]. Including leading-order contributions from the excited states to $g(\lambda)$ makes the adiabatic result in Fig. 3-1 essentially indistinguishable from the RTA attractor. This will be reported in upcoming work.

3.5 Adiabaticity in the rapidly-expanding QGP

Why does adiabaticity also apply to the violent expansion of the QGP in the early stages of the evolution? In essence, “adiabaticity” only requires that the transition to excited states is suppressed. For example, consider a time-dependent Hamiltonian in quantum mechanics $H(t) = H_0 + \tilde{\lambda}(t)H_1$, where H_0, H_1 are time-independent and $\tilde{\lambda}(t)$ is a monotonic function of time t . The transition rate from the instantaneous ground state $|0, t\rangle$ to instantaneous excited states $|n, t\rangle$ is given by [61]

$$\frac{\partial_t \log \tilde{\lambda}}{\Delta E_n} \langle 0, t | \tilde{\lambda}(t) H_1 | n, t \rangle. \quad (3.15)$$

Therefore “adiabaticity” can arise either due to the smallness of the rate of change of the Hamiltonian $\partial_t \log \tilde{\lambda}$ compared to the energy gap ΔE (slow-quench adiabaticity), or due to the time-dependent part of the Hamiltonian $\langle 0, t | \tilde{\lambda} H_1 | n, t \rangle$ being small in amplitude (fast-quench adiabaticity), see Ref [61] for examples of the applicability of adiabaticity to quantum phase transitions under fast quenches.

We have generalized the aforementioned quantum mechanical expression to a system described by Eq. (3.14). While the slow-quench adiabaticity applies at late times as one might expect, we also find that fast-quench adiabaticity applies at early times because λ is small. To see why this must be so on physical grounds, we recall that ϕ_0^F at very early times represents an angular distribution function $F_\epsilon(\cos \theta; \tau)$ where typical values of p_z are much smaller than those of p_\perp . On the other hand, the excited states at early times have typical values of p_z that are comparable to p_\perp . A “transition” from the ground state to an excited state would therefore require either multiple scatterings or one rare hard scattering among gluons, the probability of which is suppressed when $\tau \ll \tau_C$.

Since our discussion above does not rely on the details of the collision integral, we expect that adiabaticity is a generic feature of both early- and late-time limits for the expanding weakly-coupled QGP. In particular, consider the standard bottom-up thermalization scenario [62]. Following the discussion above, we expect that adiabaticity applies during the stage $\tau_{\text{Redu}} \leq \tau \leq \alpha_s^{-5/2} Q_s^{-1}$ and $\tau \geq \alpha_s^{-13/5} Q_s^{-1}$. In the former stage, $F_\epsilon(\cos \theta; \tau)$ represents the angular distribution of hard gluons (with typical energy Q_s) that rarely collide with one another. In the later stage, $F_\epsilon(\cos \theta; \tau)$ represents the angular distribution of soft gluons (with typical energy T) that are already in thermal equilibrium. Adiabaticity may break down during the transition stage $\alpha_s^{-5/2} \leq Q_s \tau \leq \alpha_s^{-13/5}$ when the numbers of both soft and hard gluons are changing rapidly, however this interval is parametrically narrow compared to other stages.

3.6 Outlook

While our analysis is based on a weakly-coupled kinetic description of the QGP, we anticipate that the concept of pre-hydrodynamic modes and the realization of adiabatic hydrodynamization is relevant more broadly. It would be interesting to explore adiabatic hydrodynamization for the QGP at strong coupling [76], and in table-top experiments [77]. As a first step towards this exploration, it may be necessary to develop a more general method to identify pre-hydrodynamic modes from the pole structure of off-equilibrium correlation functions.

3.7 Ongoing work

An advantage of the present formulation is that there appears to be a small “adiabaticity” parameter that suppresses contributions from the excited states in the evolution. This parameter can be estimated by comparing the fractional deviation of the full evolution compared to the evolution of only the ground state mode, shown in Fig. 3-2, which is seen to be small over the full evolution of the system. This makes it possible to formulate a perturbative expansion in the adiabatic parameter, which we have found to describe the full dynamics exceptionally well at linear order. In contrast, far from equilibrium where gradients are large, the hydrodynamic gradient expansion may require many (or in principle, infinite) orders to describe the full evolution. It is possible that an effective theory based on these pre-hydrodynamic modes may be better behaved than the one based on hydrodynamic modes in this far-from-equilibrium regime.

In this analysis, we have focused on Bjorken flow, which possesses a very high degree of symmetry that is not realized in more realistic situations. One consequence of this symmetry is that there is only one hydrodynamic mode, the energy density ϵ . In this case, there is a single pre-hydrodynamic mode which evolves into this hydrodynamic mode in the hydrodynamic limit. In more general situations, there may be more hydrodynamic and pre-hydrodynamic modes, and their interplay can

be more complex.

The method developed here can be straightforwardly extended to analyze a more general kinetic equation including transverse expansion and non-zero spatial gradients by considering the expansion of $F_\epsilon(\cos\theta; \tau)$ in terms of spherical harmonics $Y_{l,m,s}$. With transverse expansion but without spatial gradients, there is again an emergent time scale for the dominance of a reduced number of pre-hydrodynamic modes that are gapped from other modes. Unlike in the case of Bjorken expansion, however, not all hydrodynamic modes evolve from pre-hydrodynamic modes. This yields a more rich interplay between pre-hydrodynamic and hydrodynamic modes dominating the evolution of the system in different stages. For spatial gradients that are small compared to the energy gap, we find that the adiabatic hydrodynamization scenario discussed here still applies. These generalizations will be discussed in upcoming work.

Chapter 4

A Quantile Procedure for Sorting Out Quenched Jets

The work reported in this chapter is based on work done in collaboration with Guilherme Milhano and Jesse Thaler and published in Physical Review Letters [50]. The abstract can be found in Section 2.1.

The deconfined phase of QCD matter, the quark-gluon plasma, was first discovered in collisions of heavy nuclei at the Relativistic Heavy Ion Collider [12, 13, 14, 15, 16] and confirmed at the Large Hadron Collider [17, 18, 19]. As in high-energy proton-proton collisions, heavy-ion collisions produce collimated sprays of particles, called jets, from highly energetic scatterings of quarks and gluons. The observation of “jet quenching”—a strong suppression and modification of jets in heavy-ion collisions [18, 19, 30]—ushered in a new era of studying the properties of the quark-gluon plasma by measuring its effect on jets [78, 79, 80, 81, 31, 82, 83, 84, 85, 86, 87, 88, 89, 90].

A central issue in interpreting jet quenching measurements is that medium-induced modifications necessarily affect how jets are identified experimentally. Current methods compare proton-proton and heavy-ion jets of the same final (reconstructed) transverse momentum p_T and, as such, inevitably suffer from significant biases from the migration of jets from higher to lower p_T due to medium-induced energy loss (see [91, 92]). While these methods have been very successful in qualitatively demonstrating the phenomena of jet quenching, quantitative studies often necessitate interpreting

the data through theoretical models which include migration effects. Ideally, one would like to isolate samples of jets in proton-proton and heavy-ion collisions which were statistically equivalent when they were produced, differing only by the effects of the plasma.

In this Chapter, we propose a novel data-driven strategy for comparing heavy-ion (AA) jet measurements to proton-proton (pp) baselines which mitigates, to a large extent, the effect of p_T migration. The famous jet ratio R_{AA} compares the effective cross-section for jets in proton-proton and heavy-ion collisions with the same reconstructed p_T :

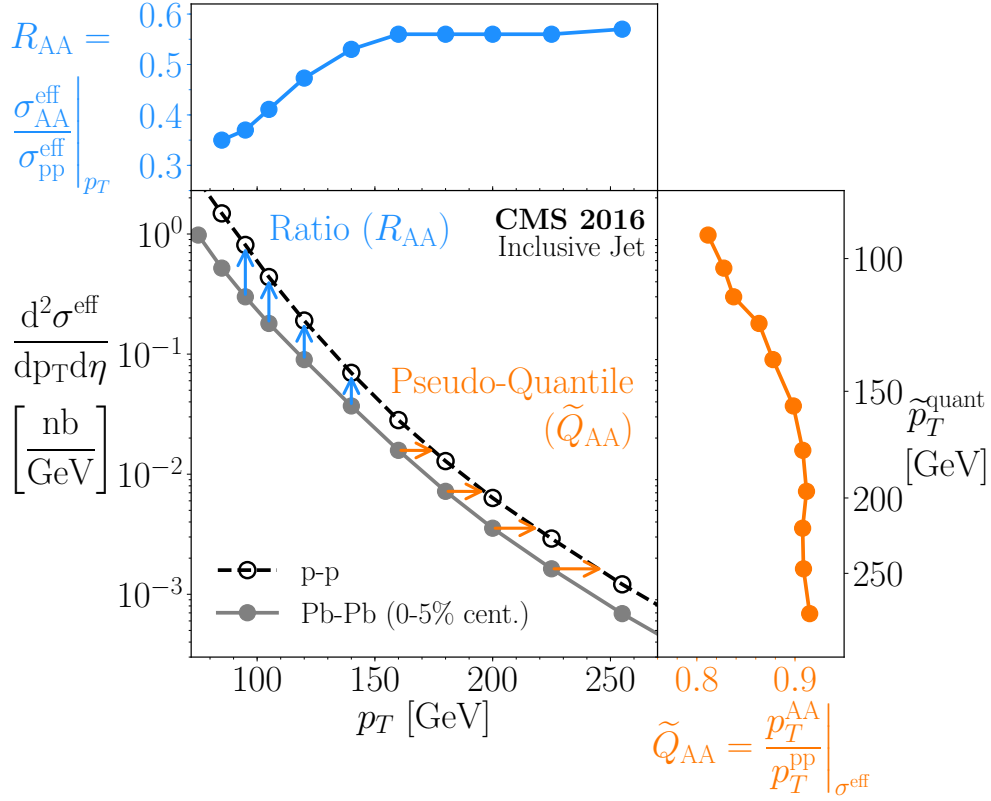
$$R_{AA} = \left. \frac{\sigma_{AA}^{\text{eff}}}{\sigma_{pp}^{\text{eff}}} \right|_{p_T}, \quad (4.1)$$

as illustrated in blue in Fig. 4-1a. Here, we introduce a “quantile” procedure, which divides jet samples sorted by p_T into quantiles of equal probability. Our new proposed observable for heavy-ion collisions is the p_T ratio between heavy-ion and proton-proton jets in the same quantile:


$$Q_{AA} = \left. \frac{p_T^{AA}}{p_T^{pp}} \right|_{\Sigma^{\text{eff}}}, \quad (4.2)$$

as illustrated in red in Fig. 4-1b, where $1 - Q_{AA}$ is a proxy for the average fractional jet energy loss. (Q_{AA} is not related to Q_{pA} used by ALICE [96]). Although R_{AA} can be obtained from Q_{AA} if the proton-proton jet spectrum is known, we will see that the physics interpretation of R_{AA} and Q_{AA} can be quite different. Fig. 4-1a additionally shows the pseudo-quantile \tilde{Q}_{AA} , which is related to the observable S_{loss} introduced by PHENIX for single hadrons [97, 98, 99].

To give an intuitive understanding of Eq. (4.2), consider a simplified scenario where medium-induced energy loss is monotonic in the p_T of the initial unquenched jet. In that case, the n^{th} highest energy jet in a heavy-ion sample is a modified version of the n^{th} highest energy jet in the corresponding proton-proton sample. Thus, in this simplified picture of energy loss, we can obtain a sample of heavy-ion jets that is statistically equivalent to its proton-proton counterpart by selecting jets with the



(a)



(b)

Figure 4-1: (continued) Illustration comparing the ratio and quantile procedures. (a) The inclusive jet p_T spectra measured by CMS [93], for a jet radius of $R = 0.4$. The standard jet ratio R_{AA} (blue) compares heavy-ion and proton-proton jet cross-sections vertically at the same reconstructed jet p_T . (b) The jet p_T cumulative cross-sections extracted from JEWEL [94, 95]. The quantile procedure Q_{AA} (red) compares heavy-ion and proton-proton jet p_T thresholds horizontally at the same cumulative cross-section. From this, one can map each p_T^{AA} (base of red arrows) into the p_T of proton-proton jets in the same quantile, p_T^{quant} (tip of red arrows). For completeness, we also show the pseudo-quantile \tilde{Q}_{AA} (orange, with corresponding $\tilde{p}_T^{\text{quant}}$) defined on the cross-section and pseudo-ratio \tilde{R}_{AA} (purple) defined on the cumulative cross-section. Though we will not explore the use of \tilde{Q}_{AA} or \tilde{R}_{AA} in the present study, we note that in JEWEL, the values of p_T^{quant} and $\tilde{p}_T^{\text{quant}}$ differ by only a few percent.

same (upper) cumulative effective cross-section:

$$\Sigma^{\text{eff}}(p_T^{\text{min}}) = \int_{p_T^{\text{min}}}^{\infty} dp_T \frac{d\sigma^{\text{eff}}}{dp_T}. \quad (4.3)$$

Note that for comparison to proton-proton cross-sections, heavy-ion cross-sections must be rescaled by the average number of nucleon-nucleon collisions $\langle N_{\text{coll}} \rangle$: $\sigma_{\text{pp}}^{\text{eff}} = \sigma_{\text{pp}}$, $\sigma_{\text{AA}}^{\text{eff}} = \sigma_{\text{AA}} / \langle N_{\text{coll}} \rangle$. Of course, energy loss is not strictly monotonic in p_T , since other properties of a jet and of the jet-medium interaction influence its energy loss and cause jets with the same initial p_T to lose different fractions of their energy. Below, we will quantify the usefulness of this quantile picture in the context of a realistic event generator where significant non-monotonicities are indeed present.

Due to the steeply-falling jet production spectrum ($\sigma \sim p_T^{-6}$), jets within a given range in reconstructed heavy-ion p_T are dominated by those which were least modified (see e.g. [100]). Addressing this issue requires comparing jets that had the same p_T when they were initially produced. In rarer events where an energetic γ or Z boson is produced back-to-back with a jet, the unmodified boson energy approximates the initial energy of the recoiling jet [82, 101]. In general jet events, however, the jet energy before medium effects cannot be measured.

A key result of this work is that the quantile picture also provides a natural proxy for the unmodified jet p_T that is observable in general jet events. Given a heavy-ion jet with reconstructed momentum p_T^{AA} , we can define p_T^{quant} implicitly as the momentum

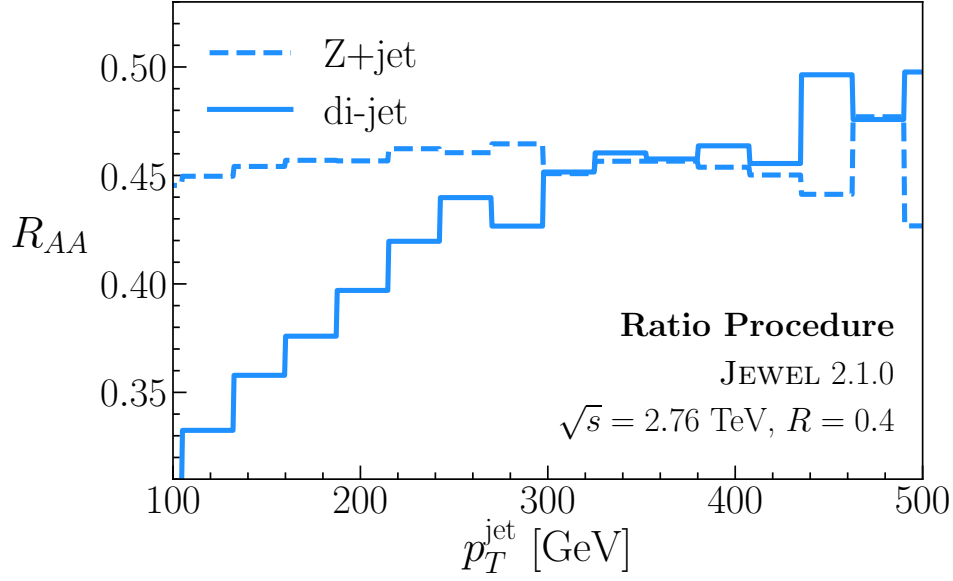
of a proton-proton jet with the same (upper) cumulative cross-section:

$$\Sigma_{\text{pp}}^{\text{eff}}(p_T^{\text{quant}}) \equiv \Sigma_{\text{AA}}^{\text{eff}}(p_T^{\text{AA}}). \quad (4.4)$$


In this quantile picture, p_T^{quant} is viewed as the initial jet p_T prior to medium effects. The mapping from p_T^{AA} to p_T^{quant} is illustrated by the red arrows in Fig. 4-1b, with $p_T^{\text{AA}} = p_T^{\text{quant}} Q_{\text{AA}}(p_T^{\text{quant}})$. Intriguingly, we will show that p_T^{quant} approximates the p_T of a heavy-ion jet before quenching with comparable fidelity to the unmodified boson energy p_T^Z available only in rarer Z+jet events. In particular, comparing properties of proton-proton and heavy-ion jet samples with the same p_T^{quant} may substantially enhance the sensitivity of modification observables by targeting jets that were more strongly modified.

For the remainder of this work, we consider samples of Z+jet and di-jet events in the heavy-ion Monte Carlo event generator JEWEL 2.1.0 [94, 95], based on vacuum jet production in PYTHIA 6 [102]. For each process, we generate 2 million each of proton-proton and head-on (0 – 10% centrality) heavy-ion events at 2.76 TeV and reconstruct anti- k_t jets using FASTJET 3.3.0 [103, 104] with radius parameter $R = 0.4$ and pseudorapidity $|\eta| < 2$. We include initial state radiation but do not include medium recoils, since medium response is not expected to have a significant effect on Eq. (4.3) at the values of p_T^{min} considered here. For Z+jet events we identify the Z from its decay to muons and consider the leading recoiling jet, and for di-jet events we consider the two highest- p_T jets. We consider Z+jet instead of γ +jet events to avoid introducing additional cuts to isolate prompt photons which could bias the validation. The default heavy-ion background in JEWEL is a Bjorken expanding medium with initial peak temperature $T_i = 485$ MeV and formation time $\tau_i = 0.6$ fm, consistent with the parameters used to fit data at 2.76 TeV in more realistic hydrodynamic simulations [95, 105].

Using these Z+jet and di-jet samples from JEWEL, Fig. 4-2a shows the standard R_{AA} (also called I_{AA} for Z+jet) and Fig. 4-2b shows the p_T ratio Q_{AA} . Although the R_{AA} for Z+jet and di-jet events have significantly different p_T -dependence, it is



(a)



(b)

Figure 4-2: Distributions of (a) R_{AA} as a function of p_T^{jet} and (b) Q_{AA} as a function of p_T^{quant} , for the Z+jet (dashed) and di-jet (solid) samples in JEWEL. Although R_{AA} and Q_{AA} are derived from the same underlying jet p_T spectra, they provide different and complementary information. For example, the p_T dependence of R_{AA} is very different for Z+jet and di-jet events in JEWEL, while the average fractional p_T loss $1 - Q_{AA}$ is similar. Note that R_{AA} requires binning of the data, while Q_{AA} , which is based on the cumulative cross-section, can be plotted unbinned.

interesting that the average fractional energy loss of jets is very similar, as quantified by $1 - Q_{AA}$. This might be surprising since Z+jet and di-jet events have different fractions of quark and gluon jets, though Ref. [106] suggests that quark and gluon jets may experience similar energy loss in JEWEL; whether this is borne out in data is an open question. Regardless, it is clear that R_{AA} and Q_{AA} offer complementary probes of the jet quenching phenomenon and are therefore both interesting observables in their own right. The quantile procedure also shows that the highest- p_T jets lose a small fraction of their energy on average ($(1 - Q_{AA}) \sim 5\%$), even though R_{AA} is far below one. This result can be compared to other methods for extracting the average energy loss from data, for example Ref. [107].

We now turn to validating the interpretation of p_T^{quant} as a proxy for the initial p_T of a heavy-ion jet before quenching by the medium. In Z+jet events, p_T^Z can be used as a baseline for the (approximate) initial p_T of the leading recoiling jet, since the Z boson does not interact with the quark-gluon plasma. For a given value of p_T^Z , there is a distribution of recoil jet momenta whose mean is shown in the upper panel of Fig. 4-3a. Even in proton-proton collisions, the recoiling jet p_T is systematically lower on average than p_T^Z due to out-of-cone radiation and events with multiple jets. In heavy-ion collisions, it is even lower due to energy loss. Intriguingly, the mean value of p_T^{quant} (red) is much more comparable to that of p_T^{pp} (dashed black) than p_T^{AA} (blue) is, indicating that p_T^{quant} is a good proxy for the initial jet p_T . On the other hand, the standard deviation of p_T^{quant} , shown in the lower panel of Fig. 4-3a, is higher than that of p_T^{pp} due to energy loss fluctuations. These cannot be undone by the quantile procedure, which can only give a perfect reconstruction of the distribution of p_T^{pp} in the case of strictly monotonic energy loss.

We emphasize that the distribution in Fig. 4-3a is physically observable and could be used to validate the quantile procedure in experimental data. Crucially, quantile matching can also provide a baseline for the initial jet p_T in general jet events. To validate this in di-jet events at the generator level, we use the p_T of the partons from the initial hard matrix element in JEWEL, p_T^{MC} , as an (unphysical and unobservable) baseline for the initial jet p_T (see [108]). We consider the two highest- p_T jets and





Figure 4-3: Mean of the jet p_T distribution compared to a baseline initial p_T (top), along with the corresponding standard deviation (bottom). Shown are (a) Z+jet events where the baseline is the physically observable p_T of the recoiling Z boson and (b) di-jet events where the baseline is the unphysical and unobservable p_T^{MC} of the initial hard scattering obtained from JEWEL. The reconstructed jet p_T for proton-proton and heavy-ion jets are shown in dashed black and blue, respectively. The p_T^{quant} of the heavy-ion sample, shown in red, more closely matches the initial jet p_T than the reconstructed heavy-ion p_T does.

match each jet with the p_T^{MC} that minimizes $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$ between the jet and the parton. Each of the two jets then enters independently in Fig. 4-3b, which demonstrates the correlation of the jet p_T to p_T^{MC} for proton-proton and heavy-ion jets, with the results of the quantile procedure in red. Fig. 4-3b is the only figure in this work that involves an unobservable quantity, and it shows remarkably similar features to Fig. 4-3a which can be measured experimentally.

It might be surprising that the curves in Fig. 4-3 are fairly flat as a function of the baseline initial p_T . This can be understood, however, from a minimal model in which the final energy of a jet is obtained from its initial energy via gaussian smearing. Consider the probability distribution

$$p(p_T^{\text{AA}}|p_T^{\text{in}}) = \int dp_T^{\text{PP}} \mathcal{N}(p_T^{\text{AA}}|\tilde{\mu}_2 p_T^{\text{PP}}, \tilde{\sigma}_2 p_T^{\text{PP}}) \mathcal{N}(p_T^{\text{PP}}|\tilde{\mu}_1 p_T^{\text{in}}, \tilde{\sigma}_1 p_T^{\text{in}}). \quad (4.5)$$

Here, $\mathcal{N}(x|\mu, \sigma)$ is a normal distribution in the variable x with mean μ and standard deviation σ , and $\tilde{\mu}_{1,2}$ and $\tilde{\sigma}_{1,2}$ are dimensionless constants. Eq. (4.5) describes the probabilistic relation between the seed-parton momentum p_T^{in} (interpreted as p_T^Z or p_T^{MC}) and the quenched momentum p_T^{AA} via two stages of gaussian smearing: first from p_T^{in} to the unquenched jet momentum p_T^{PP} , and then from p_T^{PP} to the quenched momentum p_T^{AA} . Integrating over intermediate values of p_T^{PP} gives $p(p_T^{\text{AA}}|p_T^{\text{in}})$, the probability of p_T^{AA} for fixed p_T^{in} . This is an example of a model in which the average energy loss is monotonic in p_T , since $\mu_2 = \tilde{\mu}_2 p_T^{\text{PP}}$ is a monotonic function of p_T^{PP} , but energy loss is not monotonic in p_T jet-by-jet since $\tilde{\sigma}_2 \neq 0$.

The mean and standard deviation of the distribution in Eq. (4.5) can be calculated analytically (see [109]):

$$\begin{aligned} \langle p_T^{\text{AA}}/p_T^{\text{in}} \rangle &= \tilde{\mu}_1 \tilde{\mu}_2, \\ \sigma(p_T^{\text{AA}}/p_T^{\text{in}}) &= \sqrt{\tilde{\mu}_1^2 \tilde{\sigma}_2^2 + \tilde{\mu}_2^2 \tilde{\sigma}_1^2 + \tilde{\sigma}_1^2 \tilde{\sigma}_2^2}, \end{aligned} \quad (4.6)$$

though the resulting distribution is not generally gaussian. These can be compared to the upper and lower panels, respectively, of Fig. 4-3. The fact that Eq. (4.6) has no p_T^{in} -dependence is consistent with the fact that the curves in Fig. 4-3 are approximately

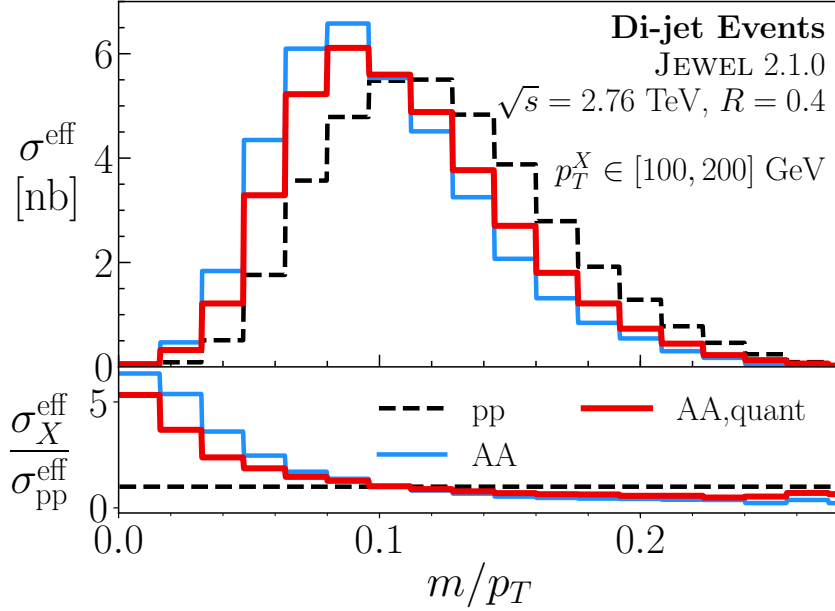


Figure 4-4: Distribution of m/p_T for proton-proton (dashed black) and heavy-ion (blue) jets in di-jet events with reconstructed $p_T \in [100, 200]$ GeV. Heavy-ion jets with $p_T^{\text{quant}} \in [100, 200]$ GeV, corresponding to $p_T^{\text{AA}} \in [80, 173]$ GeV, are in red. The heavy-ion result is normalized to match the proton-proton baseline but the quantile result has the correct normalization by construction. Partially compensating for p_T migration via the quantile procedure shifts m/p_T towards being less modified.

flat. To the extent that this model is semi-realistic, Eq. (4.6) and a measurement of Fig. 4-3a would provide an estimate of the average energy loss and the size of energy loss fluctuations. Taking approximate values from Fig. 4-3a at $p_T^Z = 300$ GeV of $\langle p_T^{\text{PP}}/p_T^Z \rangle \equiv \tilde{\mu}_1 \approx 0.87$, $\sigma(p_T^{\text{PP}}/p_T^Z) \equiv \tilde{\sigma}_1 \approx 0.2 \tilde{\mu}_1 = 0.17$, $\langle p_T^{\text{AA}}/p_T^Z \rangle \approx 0.74$, and $\sigma(p_T^{\text{AA}}/p_T^Z) \approx 0.24 \langle p_T^{\text{AA}}/p_T^Z \rangle = 0.18$, Eq. (4.6) yields $\tilde{\mu}_2 \approx 0.85$ and $\tilde{\sigma}_2 \approx 0.12$. It is satisfying that this extracted $\tilde{\mu}_2$ value is comparable to Q_{AA} in Fig. 4-2b, which is a more direct proxy for fractional energy loss.

As a final application in this Chapter, we demonstrate how the quantile procedure can be used to characterize the effects of p_T migration via an example jet substructure observable, the dimensionless ratio of the jet mass to its reconstructed p_T , m/p_T . Fig. 4-4 shows distributions of m/p_T for proton-proton and heavy-ion jets in a range of reconstructed p_T in dashed black and blue, respectively. Heavy-ion jets with that range of p_T^{quant} are those in the same quantile as the proton-proton baseline, and m/p_T

for that sample is shown in red. For the purpose of this example, we define m/p_T from the reconstructed jet mass and p_T , such that the effect of the quantile procedure is only to change the p_T range of jets in the selection. Using the quantile procedure to (partially) account for the migration of jets to lower p_T , the red distribution shifts toward m/p_T being less modified. We note that the jet mass is known to have significant corrections from medium response [110, 111] so this should be taken only as an illustrative example.

In conclusion, we introduced a new strategy for comparing heavy-ion jets to a baseline of proton-proton jets in the same quantile when sorted by p_T . As shown in Fig. 4-2, our new Q_{AA} observable is based on the same jet p_T spectra as R_{AA} but exposes different and complementary information. As shown in Fig. 4-3, our new p_T^{quant} observable is closely correlated with the initial p_T a heavy-ion jet had before energy loss to the plasma. Thus, the quantile procedure provides a data-driven way to study the modification of quenched jets and minimize the effects of sample migration. Experimental tests in Z +jet or γ +jet can validate the effectiveness of p_T^{quant} as a proxy for the initial p_T of a heavy-ion jet. If these tests are successful, the quantile procedure can then be used to re-analyze measurements of jet modification observables in general jet events with an aim toward characterizing and minimizing p_T migration effects and thus compare jet samples that were born alike. The measurement of Q_{AA} will provide information on the functional form of the average energy loss which would further constrain theoretical models. It can also be used to measure differences in average energy loss between quark- and gluon-dominated jet samples. Measurements of Q_{AA} with jet grooming [112, 113, 114, 115, 116] may also elucidate, for example, how energy is lost by the hard core of a jet compared to the diffuse periphery. It would also be interesting to study the application of this procedure to understanding energy loss fluctuations. Finally, Fig. 4-1 shows two additional observables—the pseudo-ratio \tilde{R}_{AA} and pseudo-quantile \tilde{Q}_{AA} —which may be relevant for experimental applications.

Chapter 5

Separating Quark and Gluon Jets

The work reported in this section is based on work in preparation in collaboration with Jesse Thaler and Andrew Patrick Turner. The abstract can be found in Section 2.1.

High-energy collisions between large nuclei at the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC) are a critical laboratory for studying the deconfined phase of QCD matter, the quark–gluon plasma, created in these collisions. Collimated sprays of high-momentum hadrons, called jets, are produced copiously in these collisions and provide an important probe of the quark–gluon plasma they pass through.

A long-standing question is how the quark–gluon plasma resolves the color charge of high-energy QCD partons [117, 118, 119, 120, 121, 122]. Since jets can originate from either a quark or gluon, and subsequently carry information about their respective total color charge, it is crucial to understand differences in the energy loss and modification of these two categories of jets. Unfortunately, accessing independent information about quark and gluon jets experimentally is very challenging because all jet measurements involve a mixture of contributions from both.

In this Chapter, we demonstrate a data-driven method to estimate both the quark and gluon jet fractions and their separate substructure modification in heavy-ion collisions. Our method is based on a statistical technique called topic modeling, which was pioneered for applications to quark and gluon jet separation in proton–proton collisions in Refs. [123, 124]. We present a proof-of-concept that an extension

of that technique can be used to extract differences in the modification of quark and gluon jets in heavy-ion collisions with the statistics anticipated in Run 4 of the LHC. This is a critical step toward a model-independent determination of quark and gluon jet modification in heavy-ion collisions, which would have dramatic consequences for understanding the microscopic structure of the quark–gluon plasma.

A similar type of analysis was recently performed in Ref. [125], which used a measurement of the jet charge and templates for the jet charge distributions of quark and gluon jets to extract the gluon fraction in proton–proton and heavy-ion collisions. In that study, the same Monte Carlo (MC) distributions were used as templates in both proton–proton and heavy-ion collisions, which makes the implicit assumption that the jet charge distributions of quark and gluon jets are unmodified by the quark–gluon plasma. Here, we present a method that does not require templates and does not assume that substructure observables are unmodified by the plasma, allowing for simultaneous estimates of the modification of quark and gluon jet fractions and of their distributions.

Our method is based on a statistical technique called DEMIX [126] that separates a pair of mixed probability distributions into two common underlying base distributions. This method was demonstrated in Refs. [123, 124] as a way to obtain excellent proxies for quark and gluon jets in proton–proton collisions. Consider two probability distributions $p_1(x), p_2(x)$ for a jet observable x that are a distinct mixture of the same two underlying base probability distributions $b_1(x), b_2(x)$. Namely, we can express the mixture distributions as $p_j(x) = f_j b_1(x) + (1 - f_j) b_2(x)$, for distinct fractions f_j . This expression of the p_j is always ambiguous, however, since there are infinitely many ways to mix the base distributions $b_i(x)$ among themselves to obtain new distributions $\tilde{b}_i(x)$ from which the mixture distributions can be expressed as $p_j(x) = \tilde{f}_j \tilde{b}_1(x) + (1 - \tilde{f}_j) \tilde{b}_2(x)$ with new fractions \tilde{f}_j . The idea behind DEMIX is to resolve this ambiguity by further requiring that the base distributions are *mutually irreducible* [127]. Qualitatively, this means that neither base distribution contains any component of the other; a precise definition will follow shortly. We refer to the mutually irreducible base distributions as *topics*, for their relation to the broader field

of topic modeling established in Ref. [123]. See Refs. [128, 129, 130] for other uses of topic modeling techniques in collider physics.

The notion of “quark- and gluon-initiated jets” is not well-defined at the hadron level. Even at the level of a MC generator, where parton information from the hard process is available, there is still an ambiguity about how to associate final-state jets with their initiating parton. Therefore, the quark and gluon topics we discuss in this Chapter do not correspond directly to any parton-level intuition about quark- and gluon-initiated jets [131]. Instead, these topics correspond to the *operational definition* of jets introduced in Ref. [124], which defines the quark and gluon categories as the mutually irreducible (i.e., maximally separable) distributions underlying a pair of jet samples. To minimize potential confusions, we will often use the language of quark-like and gluon-like (or “quark” and “gluon” in quotes) to refer to this operational definition.

Since the base distributions extracted from a jet observable x using DEMIX are mutually irreducible, they can only agree with the MC quark and gluon jet distributions of x if those are also mutually irreducible. It was argued in Ref. [123] that quark–gluon mutual irreducibility is approximately satisfied for the constituent multiplicity (number of hadrons reconstructed as part of a jet) of groomed and ungroomed jets and n_{SD} [132] in proton–proton collisions, though not for other common jet observables like jet mass. This stems from counting observables having exact quark–gluon mutual irreducibility in the high-energy limit [132]. Ref. [124] further showed that constituent multiplicity is a nearly optimal classifier to separate operationally defined quark and gluon jets (see Ref. [133] for further developments). We thus focus on constituent multiplicity for extracting quark and gluon fractions in our case study.

The algorithm to extract the mutually irreducible base distributions is straightforward. Base distributions are computed from the mixture distributions via

$$\begin{aligned} b_1(x) &= \frac{p_1(x) - \kappa_{12} p_2(x)}{1 - \kappa_{12}}, \\ b_2(x) &= \frac{p_2(x) - \kappa_{21} p_1(x)}{1 - \kappa_{21}}, \end{aligned} \tag{5.1}$$

with

$$\kappa_{ij} = \inf_x \frac{p_i(x)}{p_j(x)}. \quad (5.2)$$

The reducibility factor κ_{ij} is the maximum fraction of $p_j(x)$ that can be subtracted from $p_i(x)$ such that the resulting function remains positive for every x , and can thus be normalized to yield a proper probability distribution. Mutual irreducibility of the p_i is precisely the condition that $\kappa_{12} = \kappa_{21} = 0$. The κ_{ij} are directly related to the mixture fractions:

$$\kappa_{12} = \frac{1 - f_1}{1 - f_2}, \quad \kappa_{21} = \frac{f_2}{f_1}. \quad (5.3)$$

For two analytically known mixtures of two base distributions, it is essentially trivial to compute κ_{ij} using eq. (5.2) and then to extract the mutually irreducible underlying distributions using eq. (5.1). As an example, two Gaussian distributions with different mean and the same standard deviation are mutually irreducible, so any two convex mixtures built from them can be demixed exactly.

When dealing with finite-sampled distributions, however, one encounters substantial technical difficulties using eq. (5.2) directly. A histogram of samples from a probability distribution $p(x)$ has a finite, discretized range of histogram bins $\{x_k\}$ at which $p(x)$ is estimated from the finite-statistics sampled distribution, $\hat{p}(x_k)$. We need a method of defining the reducibility factors $\hat{\kappa}_{ij}$ for a pair of sampled histograms. Naively, the infimum of eq. (5.2) becomes a minimum of the ratio of the histograms over $\{x_k\}$; simply taking the minimum, however, is very sensitive to statistical fluctuations. A more robust approach, introduced in Ref. [124], is to define $\hat{\kappa}_{ij}$ to be the ratio of histograms in the bin for which the ratio plus its uncertainty is minimized. This method turns out to be insufficient to deal with the much more limited statistics we aim to utilize in this work, particularly because $\hat{\kappa}_{ij}$ is typically extracted at the low-statistics end points of the distributions.

One way to address the issue of limited statistics is by using fitting to leverage information about the interior of the distribution, where the statistics are better, to put additional constraints on the tails. We note that the histograms shown later in figs. 5-1a and 5-1d are exceptionally well-described by a simultaneous fit to two

distinct sums of a pair of skew-normal distributions $\text{SN}(x; \mu, \sigma, s)$. That is, they are well-described by the form

$$f_N(x; \boldsymbol{\alpha}_i, \boldsymbol{\theta}) = \sum_{k=1}^N \alpha_{i,k} \text{SN}(x; \mu_k, \sigma_k, s_k) \quad (5.4)$$

with $N = 2$. Here $\boldsymbol{\theta} = (\mu_1, \sigma_1, s_1, \dots, \mu_N, \sigma_N, s_N)$, and $\boldsymbol{\alpha}_i = (\alpha_{i,1}, \dots, \alpha_{i,N})$ contains $N - 1$ independent fractions, with the N th fraction constrained by $\sum_{k=1}^N \alpha_{i,k} = 1$, for jet samples $i = 1, 2$ (dijet and γ +jet in figs. 5-1a and 5-1d). For further generality of the functional form, we consider $N = 4$ and simultaneously fit the two input distributions to $f_4(x; \boldsymbol{\alpha}_1, \boldsymbol{\theta})$ and $f_4(x; \boldsymbol{\alpha}_2, \boldsymbol{\theta})$, respectively, with 18 fit parameters $\boldsymbol{\alpha}_1$, $\boldsymbol{\alpha}_2$, and $\boldsymbol{\theta}$. To estimate the uncertainty on such fits, we use the Markov Chain Monte Carlo (MCMC) ensemble sampler `emcee` [134] to do posterior estimation using the likelihood function [135]

$$\ln \frac{C}{p} = \sum_{i,j} n_i \left[f(x_{i,j}; \boldsymbol{\alpha}_i, \boldsymbol{\theta}) - y_{i,j} + y_{i,j} \ln \frac{y_{i,j}}{f(x_{i,j}; \boldsymbol{\alpha}_i, \boldsymbol{\theta})} \right]. \quad (5.5)$$

Here, j indexes the histogram bins of jet sample i , with the j th bin having constituent multiplicity $x_{i,j}$ and probability density $y_{i,j}$, and the i th sample having total count n_i . This form assumes that the number of counts in each histogram bin, $n_{i,j} = n_i y_{i,j}$, are independently Poisson-distributed around the value $f(x_{i,j}; \boldsymbol{\alpha}_i, \boldsymbol{\theta})$, and estimates distributions of the parameters $\boldsymbol{\alpha}_i, \boldsymbol{\theta}$ for which the observed data is most likely. Following Refs. [135, 136], the likelihood function p in eq. (5.5) is rescaled by a fit-independent constant C that cancels a $\ln(n_{i,j}!)$ that arises when taking the log of the Poisson probability distribution. We take a uniform prior on the parameters $\boldsymbol{\theta}$ and $\boldsymbol{\alpha}_i$ in the range $\mu_k \in [0, 50]$, $\sigma_k \in [1, 15]$, $s_k \in [-20, 20]$, and $\alpha_{i,k} \in [0, 1]$, and we start the MCMC walkers in a Gaussian ball of standard deviation 10% around the least-squares fit parameters. We use the distribution of fits to obtain distributions of $\hat{\kappa}_{ij}$ via eq. (5.2). To combat finite statistics effects, we compute the infimum in eq. (5.2) as a minimum of the MCMC walkers over a reduced range. We consider only the range for which at least one input histogram is non-zero. For each reducibility factor, we further identify

whether the minimum will occur on the left or right side of the range, and truncate the opposite tail at the outermost bin that has at least 10 data points for each input histogram. The distribution of $\hat{\kappa}_{ij}$ is used to compute a distribution of fractions, and its mean and standard deviation are used as the value and uncertainty of $\hat{\kappa}_{ij}$ used to extract the topics.

The samples for our proof-of-concept study come from the heavy-ion MC event generator JEWEL 2.1.0 [94, 95], based on vacuum jet production in PYTHIA 6.4.25 [102]. We consider two mixed distributions coming from photon-jet ($\gamma + \text{jet}$) production and dijet production. For each process, we generate proton-proton and 0% to 10% centrality heavy-ion events at 5.02 TeV and reconstruct anti- k_t jets using FASTJET 3.3.0 [103, 104] with radius parameter $R = 0.4$ within the pseudorapidity range $|\eta| < 1$. We include initial-state radiation, but do not include medium recoil effects. (There is no underlying event model in JEWEL, but we verified that similar results can be obtained after aggressively grooming jets using the Soft Drop algorithm [116] with $z_{\text{cut}} = 0.5$ and $\beta = 1.5$ as in Ref. [137].) For $\gamma + \text{jet}$ events, we consider the recoiling jet with the highest transverse momentum (p_T), and for dijet events, we consider the two highest- p_T jets. In the case of heavy-ion collisions, we downsample our JEWEL events to mimic the statistics that will be available with the anticipated luminosity $\int \mathcal{L} dt = 13 \text{ nb}^{-1}$ after Run 4 [138].

To match the statistics available in Run 4 of the LHC, we had to downsample our heavy-ion $\gamma + \text{jet}$ events. The equivalent statistics of our dijet sample are already less than those achievable in Run 4, but this substantially reduces the statistics of our $\gamma + \text{jet}$ sample. The expected integrated luminosity for heavy-ion collisions after Run 4 is $\mathcal{L}_{AA,\text{tot}} = \int \mathcal{L} dt = 13 \text{ nb}^{-1}$. We should note, however, that the luminosity we are interested in is multiplied by a factor of $A^2 = 208^2$ for lead-lead collisions, $\mathcal{L}_{nn,\text{tot}} = 208^2 \mathcal{L}_{AA,\text{tot}}$ [139]. With the total cross-section σ for a process, this gives the number of events $N = \sigma \mathcal{L}_{nn,\text{tot}}$. Since our JEWEL sample is based on weighted events, we perform a standard unweighting procedure to simulate the statistics possible after Run 4. If we have N_w weighted events, the unweighting procedure probabilistically chooses some subset of these events to act as unweighted events. To obtain an expected number

of kept events N , the j th event is (independently) kept with probability $p_j = N \frac{w_j}{w_{\text{tot}}}$, where $w_{\text{tot}} = \sum_j w_j$. Note that this puts a constraint on the required size of N_w to be able to unweight it to a sample of size N , namely that $1 \geq p_{\text{max}} = N \frac{w_{\text{max}}}{w_{\text{tot}}}$, with w_{max} being the maximum value of w_j . We unweight our proton–proton sample per p_T bin by fixing $p_{\text{max}} = 1$ in each bin and downsampling to a sample of expected size N . After this unweighting we have 63603 and 42466 jets in dijet and γ +jet samples, respectively, with $p_T \in [100, 120]$ GeV, 25046 and 18488 with $p_T \in [120, 140]$ GeV, and 13707 and 9260 with $p_T \in [140, 160]$ GeV. For clarity, we downsample our heavy-ion results uniformly across all p_T bins so that the resulting statistics correspond to fixed luminosities of 13 and 37 nb⁻¹. Our 13 nb⁻¹ γ +jet sample has 6453, 2825, and 1436 jets in p_T ranges [100, 120] GeV, [120, 140] GeV, and [140, 160] GeV, respectively. Our dijet sample has lower luminosity with 71098, 32318, and 15747 jets, respectively, in the same p_T bins.

Starting with proton–proton collisions in the top row of fig. 5-1, we show the distributions of jet constituent multiplicity for γ + jet and dijet samples (fig. 5-1a) and the “quark-like” and “gluon-like” topics extracted from these distributions via the data-driven method described above (fig. 5-1b). The corresponding heavy-ion results are shown in figs. 5-1d and 5-1e, keeping in mind that the proton–proton and heavy-ion analyses are completely independent. The extracted topics are in good agreement with the distributions of constituent multiplicity for quark- and gluon-initiated jets as defined at the MC level.

Furthermore, we can use eq. (5.3) to extract the topic fractions, i.e., the proportions of the topics in the original input distributions. Figures 5-1c and 5-1f show the extracted fraction of the gluon-like topic in the γ + jet and dijet samples as a function of jet p_T . The gluon topic fractions are marginally higher than the MC-level fraction of gluon-initiated jets in proton–proton collisions, and more dramatically higher in heavy-ion collisions. In interpreting these results, however, one has to be mindful both of the inherent ambiguity in using MC-level information to label jets, and the effect of limited statistics.

It is important to remember that MC-level information is inherently ambigu-

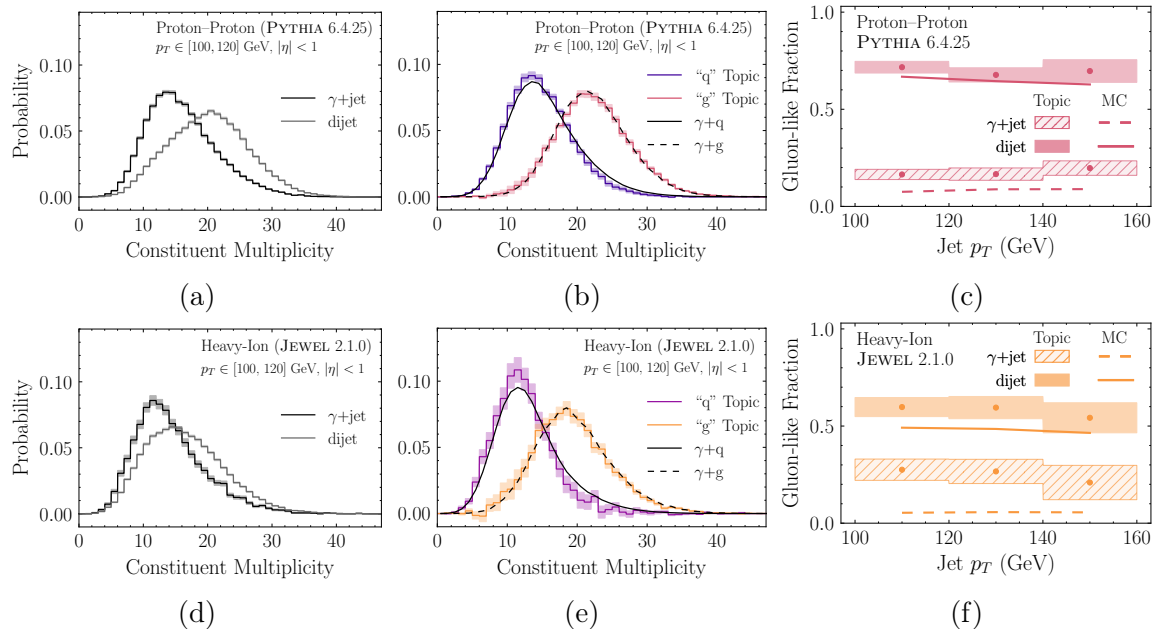


Figure 5-1: Extracting quark-like and gluon-like jet topics from (top row) proton-proton collisions and (bottom row) heavy-ion collisions, as generated by PYTHIA and JEWEL, respectively. (a) Normalized distributions of jet constituent multiplicity for the $\gamma + \text{jet}$ and dijet samples in proton-proton collisions. (b) The two underlying topics extracted from these distributions using the DEMIX method (colorful bands), compared to the self-normalized MC-level definition of quark- and gluon-initiated jets (black) for those jets within $\Delta R = 0.4$ of an initiating parton. (c) Fractions of the gluon-like topic in the $\gamma + \text{jet}$ and dijet samples, compared to the MC-level quark- and gluon-initiated jet fractions. The corresponding results for heavy-ion collisions are shown in (d)–(f). The heavy-ion $\gamma + \text{jet}$ statistics are chosen to match those available in Run 4 at the LHC. Possible reasons for the higher gluon-like topic fractions compared to the MC label fractions are provided in the text.

ous. Here we defined the MC quark- and gluon-initiated jet distributions ($\gamma + \text{quark}$ and $\gamma + \text{gluon}$ in fig. 5-1) to include only those from jets whose axis is within $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = 0.4$ from a quark or gluon in the initial hard scattering matrix element. This requirement is only satisfied, however, by 87% and 90% of the jets in our heavy-ion sample of $\gamma + \text{jet}$ and dijets, respectively, with $p_T \in [100, 120]$ GeV. In the analogous proton-proton samples, it is satisfied for 94% and 95% of the jets in $\gamma + \text{jet}$ and dijet samples, respectively. (These percentages can be further increased by requiring the two jets, or the photon and jet, be nearly back-to-back in azimuthal angle.) Particularly for our heavy-ion results, this implies that the $\gamma + \text{jet}$ and dijet samples are not direct combinations of the MC-level quark and gluon jet distributions.

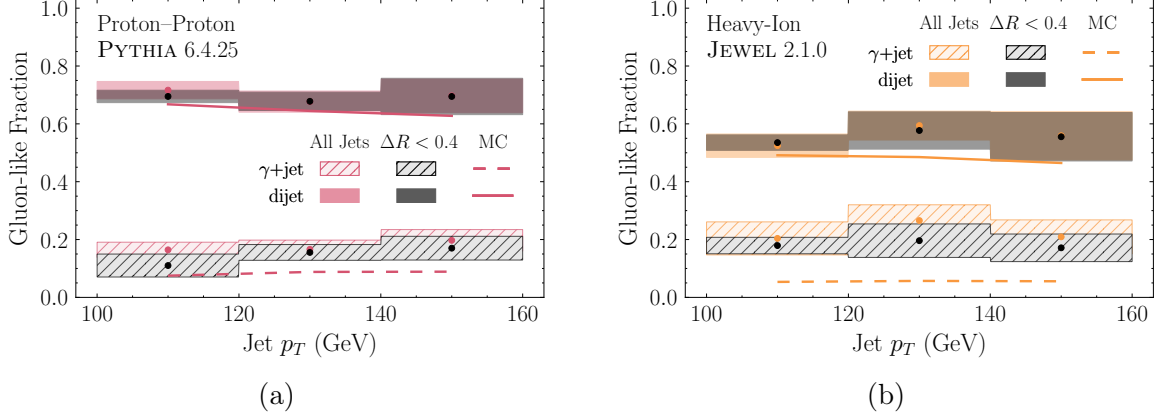


Figure 5-2: Gluon-like topic fractions for (a) proton–proton collisions and (b) heavy-ion collisions, as a function of jet p_T , for the γ +jet (hatched) and dijet (solid) samples. Colorful bands are the extracted gluon-like topic fractions from the full γ + jet and dijet distributions, as already shown in figs. 5-1c and 5-1f. Black bands show for comparison the gluon-like topic fractions extracted from (unphysically) restricted γ + jet and dijet samples including only jets whose axis is within $\Delta R = 0.4$ of a parton in the hard scattering matrix element.

Thus, one might wonder whether part of the mismatch between the gluon-like fractions seen in fig. 5-1f could be due to the fundamental ambiguity in jet flavor labeling. In fig. 5-2, we study this by (unphysically) restricting the γ + jet and dijet distributions to those jets whose axis is within $\Delta R = 0.4$ of an initiating parton, and then using these restricted samples to perform the topic extraction. This has relatively little impact on the gluon-like topic fraction, though it tends to make it somewhat more similar to the MC-level gluon-initiated jet fraction.

In addition to MC label ambiguities, limited statistics drive the extraction of κ from eq. (5.2) into the interior of the distribution where the true minimum is not yet achieved, and can thus artificially enhance the extracted gluon-like topic fraction. While the proton–proton results are relatively robust to statistical effects, the heavy-ion results are impacted by the lower statistics expected for LHC Run 4. This particularly affects the γ + jet sample, where the statistics only enable a rough estimate of the constituent multiplicity distribution. In fig. 5-3, we show the result of topic extraction after increasing the γ + jet sample to have a factor of about 2.8 more events than expected in LHC Run 4. With limited statistics, the method tends to overestimate the gluon-like fraction, though the results are consistent within uncer-

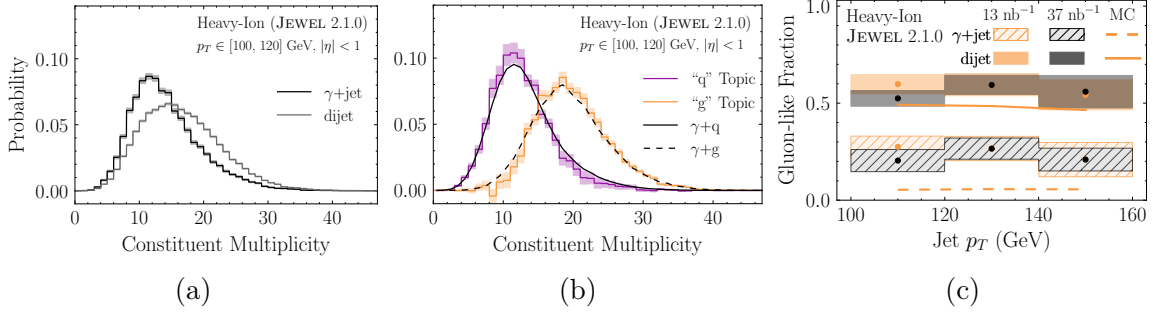


Figure 5-3: Same as the bottom row of fig. 5-1, but using a factor of about 2.8 higher $\gamma + \text{jet}$ statistics than expected after LHC Run 4. With more events, the agreement of the distribution of the quark-like topic (purple) relative to the MC-level definition (black) is somewhat improved compared to fig. 5-1e, though the gluon-like topic fraction remains high compared to the MC label fraction.

tainties. With the higher statistics, the gluon-like topic fractions remain higher than the MC-level fractions, though there may be remaining limited statistics effects even in the higher statistics sample.

Even accounting for these issues, we find a persistently larger gluon-like fraction compared to the MC labeling, at least in the context of JEWEL. One possible explanation for this effect is that a “quark-initiated” jet may become more gluon-like through gluon radiation, an effect which may be enhanced by medium-induced gluon radiation in heavy-ion collisions. For methods like this one, as well as for the method in Ref. [125], this would result in a larger fraction of jets being classified as gluon jets. It is also possible that constituent multiplicity, though apparently nearly mutually irreducible in proton–proton collisions, may be less mutually irreducible in the presence of medium effects, so alternative observables (perhaps from machine learning [124]) might be required. Understanding these issues will be important for interpreting eventual LHC Run 4 data, but we emphasize that the operational definition used to define the quark-like and gluon-like topics is independent of its interpretation.

As a final proof-of-concept, in fig. 5-4 we show the modification of the jet constituent multiplicity distributions for the quark-like (fig. 5-4a) and gluon-like (fig. 5-4b) jet topics as a function of p_T . To our knowledge, this represents the first fully data-driven method to separate the modification of a jet observable for “quark” and “gluon” jets. Though we show here the modification of the constituent multiplicity dis-

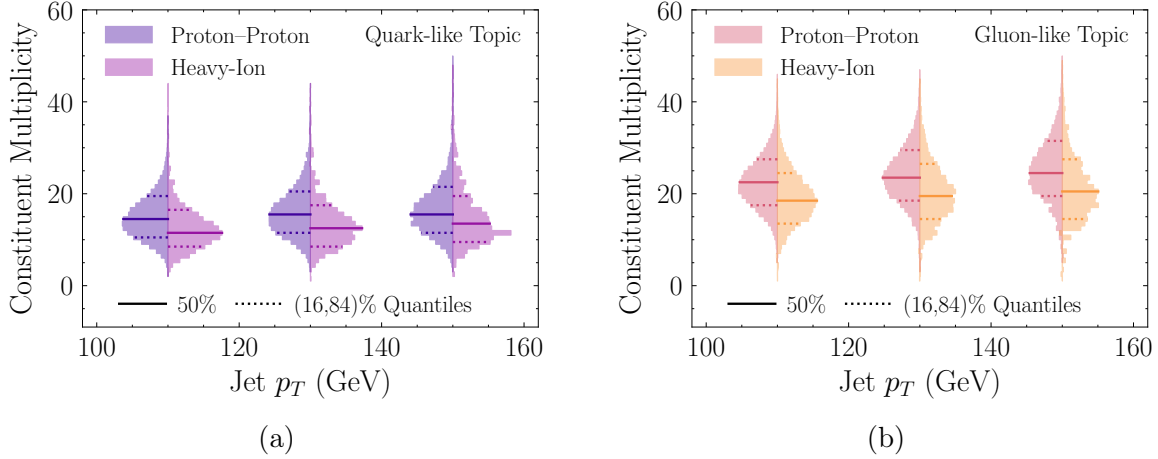


Figure 5-4: Constituent multiplicity distributions for (a) the quark-like topic and (b) the gluon-like topic as a function of jet p_T for p_T bins [100, 120] GeV, [120, 140] GeV, and [140, 160] GeV. Each violin plot has results for both (left side) proton–proton and (right side) heavy-ion collisions, and the change between the two sides illustrates the modification of the constituent multiplicity distribution for the corresponding topic. Horizontal lines indicate the median (solid) and 16% and 84% quantiles (dashed) of the multiplicity distributions.

tribution for clarity, we emphasize that once the topic fractions have been extracted, they can be used to extract separate quark and gluon distributions for any jet observable. Since both jet observable distributions and the quark and gluon fractions may change between proton–proton and heavy-ion collisions, it is substantially simpler to interpret the separate modification of quark and gluon topics compared to, e.g., the modification of the dijet distribution. The fact that the constituent multiplicities of heavy-ion jets simulated in JEWEL are lower compared to proton–proton jets with the same p_T may be due, at least in part, to a selection bias favoring those (typically lower multiplicity) jets that lost least energy, as discussed in Chapter 1. Though not shown here, the method of Chapter 4 (Ref. [50]) could be used to match the proton–proton and heavy-ion jet p_T quantiles and further clarify the interpretation.

In summary, we have illustrated a data-driven method to extract quark-like and gluon-like topic fractions and distributions in proton–proton and heavy-ion collisions. Using JEWEL samples of comparable statistics to those anticipated in Run 4 of the LHC, we have shown that these topics have a similar qualitative interpretation to the (physically ambiguous) definition of the quark and gluon jets at parton level available

from MC generators. We have further shown, as an example, the modification of the constituent multiplicity in heavy-ion collisions separately for quark and gluon jet topics. This study offers an exciting proof-of-concept demonstration of the power of the topic modeling to interpret future heavy-ion collision data, though more quantitative studies will of course be necessary to understand the feasibility of this analysis and associated systematic uncertainties.

We leave a detailed study of the impact of underlying event and background subtraction, medium response, and experimental inefficiencies as important future work. Constituent multiplicity may be particularly difficult to measure in heavy-ion collisions due to the large and fluctuating underlying event and the associated uncertainties arising from background subtraction. Additional studies will be necessary to understand how under- or over-subtracted background and particles from medium response may impact the topic extraction shown here. In the presence of these effects, it may prove more clean to use the constituent multiplicity of jets groomed with the Soft Drop algorithm for this analysis. Alternatively, it may be possible to consider the multiplicity of hadrons above some p_T threshold to remove the contribution from medium response. In addition, some of the interpretations suggested here can be studied in detail in other models. For example, the model in Ref. [140] does not include medium-induced radiation, and could thus be used to assess whether medium-induced gluon radiation could be responsible for the larger deviation of the gluon topic fraction from the MC label fraction in heavy-ion compared to proton–proton collisions observed in JEWEL. In any model with information about the initiating parton flavor, it is further possible to study the mutual irreducibility of quark- and gluon-initiated jets in any jet observable. However, it is important to keep in mind that the topics themselves can be extracted from data and are independent of any direct parton-level interpretation which may additionally arise if the parton-level definitions of quark and gluon jets are mutually irreducible.

Chapter 6

Rapidity Dependence of Cumulants

The work reported in this chapter is based on work done in collaboration with Swagato Mukherjee, Krishna Rajagopal, and Yi Yin and published as an Editors Suggestion in Physical Review C Rapid Communications [51]. The abstract can be found in Section 2.1.

A central goal of heavy ion collision experiments is to map the QCD phase diagram as a function of temperature T and baryon chemical potential μ_B [141, 45, 33]. At zero μ_B , the phase diagram features a continuous crossover from quark-gluon plasma (QGP) to ordinary hadronic matter as a function of decreasing T [142, 143, 144, 40, 145]. Increasing μ_B corresponds to doping the QGP with an excess of quarks over antiquarks. As in condensed matter physics, after the discovery of a new form of strongly correlated matter, here the strongly coupled liquid called QGP, a more complete understanding requires mapping its phase diagram as a function of T and doping. It is an open question whether the crossover at zero doping becomes a sharp first order phase transition as the doping μ_B is increased beyond some critical point [146, 33]. At nonzero μ_B where lattice calculations become extremely difficult [147, 148], there are no first-principles theoretical calculations which provide reliable guidance as to the whether there is a critical point in the phase diagram of QCD, or its location if it does exist [149, 150, 41, 151]. Model calculations suggest the existence of a critical point, but disagree wildly on its location in the (μ_B, T) plane [41, 151]. Reducing the beam energy increases the μ_B of the QGP produced in a heavy ion collision [43, 152, 153, 33]

(principally because lower energy collisions make less entropy but also because they deposit more of their baryon number in the plasma) but it also reduces the temperatures achieved. So, these experiments can scan the crossover (and potentially critical) regime of the phase diagram out to some value of μ_B corresponding to the lowest energy collisions that reach the crossover (critical) temperature [141, 45]. If a critical point is located in the regime that is within reach, it may be detected experimentally. The search for a critical point in the phase diagram of QCD at the Relativistic Heavy Ion Collider (RHIC) is currently underway, with collisions at energies ranging from $\sqrt{s} = 200$ AGeV down to $\sqrt{s} = 7.7$ AGeV, producing fireballs that freeze out with chemical potentials in the range $25 \text{ MeV} \lesssim \mu_B \lesssim 400 \text{ MeV}$ [152, 153]. This exploration will be extended to higher μ_B , albeit at lower temperatures, via collisions at lower \sqrt{s} at the NICA facility [154], now under construction, and fixed target collisions at RHIC and at the future FAIR facility [155]. Phase I of the RHIC beam energy scan (BES-I) was completed in 2014, with no signs of a critical point for $\mu_B < 200$ MeV and with tantalizing but inconclusive results at larger μ_B , in collisions with $19.6 \text{ AGeV} \geq \sqrt{s} \geq 7.7 \text{ AGeV}$ [156, 157, 48, 141, 45, 146]. Phase 2 of the scan (BES-II), to begin in 2019 [48, 146], will focus on this regime with increased luminosity and consequently much higher statistics. One of the improvements planned before BES-II is an upgrade of the inner Time Projection Chamber (iTTPC) at STAR, which will extend its rapidity acceptance for protons from $|y| < 0.5$ in BES-I to $|y| < 0.8$ in BES-II [158].

The energy of a heavy ion collision sets the initial T and μ_B of the QGP which is created, with lower energy collisions being more baryon-rich. The QGP then follows a trajectory in the (T, μ_B) plane as it expands and cools. If there is a critical point in the QCD phase diagram within the range of μ_B which is accessible in the Beam Energy Scan (BES), then at some collision energies the fireball produced may pass through or near the critical region, while at higher (lower) collision energies the fireball produced will pass the critical point on the low (high) μ_B side. The theoretical challenge is to describe the unique signatures of this scenario which would be observable in data from the BES.

A critical point in a thermodynamic system is characterized by an enhanced correlation length. Although the correlation length itself is not observable, because the critical order parameter σ couples to all hadrons the n 'th cumulant moments κ_n of the event-by-event distribution of the measured multiplicity N of various particle species, for example $\kappa_4[N] = \langle(\delta N)^4\rangle - 3\langle(\delta N)^2\rangle^2$, scale with powers of the correlation length ξ near a critical point [47]. Protons couple more strongly to σ than pions or kaons, making cumulants of the proton multiplicity good observables with which to look for critical fluctuations [159, 160]. The non-Gaussian cumulants $\kappa_3[N] = \langle(\delta N)^3\rangle \sim \xi^{9/2}$ and $\kappa_4[N] \sim \xi^7$ scale with higher powers of the correlation length than the Gaussian cumulants and are therefore more sensitive to critical behavior [47, 160]. Furthermore, an analysis that is valid for any critical point in the same (3d Ising) universality class has shown that κ_4 will also change sign near a QCD critical point [5, 161]. Non-monotonic behavior and a sign change of the fourth cumulant as a function of the beam energy are characteristic signatures of the presence of a critical point which can be searched for in the RHIC BES.

The dependence of fluctuation measures on the total rapidity acceptance has been studied before [162, 163, 164] upon assuming that μ_B , and hence ξ and κ_4 , do not depend on rapidity. However, the baryon density does depend non-trivially on rapidity at RHIC BES energies (see e.g. Refs. [165, 166, 4]). Even at top RHIC and LHC energies, the baryon density is significant at very large rapidity, around two units of rapidity below that of the incident beams, although building instrumentation to make the required measurements to explore the hot baryon-rich matter produced in this regime presents considerable challenges [33, 165]. Since the correlation length and fluctuations become very large near a critical point, the rapidity dependence of the baryon density gives rise to a strong, non-trivial rapidity dependence of the cumulants near a critical point which was not incorporated in previous work. Furthermore, since each rapidity is associated with a different value of μ_B and therefore probes a different part of the critical region on the phase diagram, we shall see that integrating over the full rapidity acceptance averages out interesting features in the cumulants which are characteristic of critical behavior. Instead, we propose that binning the

cumulants in rapidity gives a more crisp picture of the critical regime, and demonstrate that the rapidity dependence of these binned cumulants near mid-rapidity will change qualitatively if a critical point is passed in the BES. We therefore propose this observable as a complementary means by which to observe the presence of a critical point at the BES. In particular, it provides a new, and distinctive, signature by which to determine whether downward steps in the collision energy take us past a critical point in the phase diagram.

Rapidity-dependence of μ_B : In this Rapid Communication, we illustrate the effect that the rapidity-dependence of μ_B at RHIC BES energies has on the rapidity-dependence of cumulants, and propose using this toward discovering (or ruling out) a critical point using RHIC BES data.

Since the baryon density at freezeout is symmetric in the spacetime rapidity y_s for symmetric heavy-ion collisions, for small $|y_s|$ the deviation from boost invariance takes the form

$$\mu_B(y_s) \sim \mu_{B,0} + \alpha y_s^2, \quad (6.1)$$

with $\mu_{B,0}$ and α constants that depend on the beam energy \sqrt{s} . We shall use this form for illustrative purposes, noting of course that it cannot be relied upon at large $|y_s|$. As we have discussed, the basis of the BES is that downward steps in \sqrt{s} yield upward steps in $\mu_{B,0}$. For illustrative purposes, we shall pick three values of $\mu_{B,0}$ within the BES range, and see what happens if these steps were to happen to take us past a possible critical point. The value of α has been measured in SPS collisions with $\sqrt{s} = 17.3$ AGeV, where $\alpha = 50$ MeV [166]. At this (and all higher, and some lower) beam energies, $\alpha > 0$ because the baryon number density is peaked at roughly two units of rapidity below the beam rapidity, meaning that it is less at $y_s = 0$ than at larger $|y_s|$ [43, 33]. In AGS collisions with $\sqrt{s} = 5.5$ AGeV, though, the beam rapidity is low enough that the baryon number density is peaked at $y_s = 0$, and $\alpha < 0$ [43]. The \sqrt{s} at which α changes from positive to negative is not known, but is likely near the lower end of the BES range. Ultimately, measurements of the ratios of the mean particle number distribution for different species from the RHIC BES should be used

to measure how μ_B at freezeout depends on y_s at each BES collision energy, and hence to determine the value of α at each energy. For illustrative purposes here, we shall investigate the consequences of choosing $\alpha = 50$ MeV at each of our three values of $\mu_{B,0}$ as well as checking how things change if we choose $\alpha = -50$ MeV instead at our largest value of $\mu_{B,0}$.

Cumulants in the critical regime: Order parameter fluctuations near a critical point induce fluctuations in the event-by-event particle multiplicities. Throughout this work, we will consider the cumulants of protons, as these are expected to be most sensitive to critical fluctuations [160]. From previous work [47, 160, 5, 163], the contribution to the fourth cumulant of the proton multiplicity distribution coming from critical fluctuations (denoted by the subscript σ) takes the form

$$\kappa_4[N]_\sigma = \int_{\mathbf{x}} K_4 \xi^7 T^2 \left(g \int_{\mathbf{p}} \frac{\chi_{\mathbf{p}}}{\gamma_{\mathbf{p}}} \right)^4, \quad (6.2)$$

where the \mathbf{x} -integral is a spacetime integral over the freezeout hypersurface, where T , μ_B and consequently $\xi(\mu_B, T)$ and $K_4(\mu_B, T)$ (proportional to the kurtosis of the event-by-event distribution of the fluctuating order parameter, see below) can take on different values at different points on the freezeout hypersurface, where g is the σ -proton-proton coupling which we set to the same benchmark value $g = 7$ used in Ref. [160], where the \mathbf{p} -integral is a momentum-space integral over the protons at the point \mathbf{x} , where $\chi_{\mathbf{p}} = f_{\mathbf{p}}(1 - f_{\mathbf{p}})/T$ if we assume local equilibrium, with $f_{\mathbf{p}}$ the Fermi-Dirac distribution boosted by the radial flow velocity at the point \mathbf{x} , and where $\gamma_{\mathbf{p}} = \sqrt{\mathbf{p}^2 + m^2}/m$, with m the proton mass. We note that this freezeout prescription allows us to convert the spacetime rapidity dependence of μ_B into momentum-space rapidity dependence of $\kappa_4[N]_\sigma$. (See eq. (6.3) below for details.) We follow Ref. [163] and use a blast wave model to obtain the radial flow velocity and freezeout hypersurface, taking the freezeout curve in the (T, μ_B) plane from the fit to experimental data found in Ref. [152, 153]. Following Ref. [163], we shall make the approximation $\chi_{\mathbf{p}} \approx f_{\mathbf{p}}/T$ and use the Boltzmann distribution for $f_{\mathbf{p}}$, allowing us to do some of the integrals analytically.

The shape of the dependence of K_4 and ξ on μ_B and T are governed by universal properties of critical fluctuations. A critical point in the QCD phase diagram, if it exists, is known to be in the same universality class as the 3d Ising model [167, 168, 169, 44, 46]. The mapping of the Ising variables (r, h) onto the QCD variables (μ_B, T) is not universal, but for illustrative purposes we employ the widely-used assumption [170] that the Ising r axis (and hence the line of first order transitions) is parallel to the QCD μ_B axis, and the Ising h axis is parallel to the QCD T axis. For illustrative purposes, we shall place a hypothetical QCD critical point at $\mu_B = 260$ MeV, $T = 160$ MeV. The 3d Ising universality then determines K_4 at some point away from the critical point in terms of the direction in which that point lies in the (μ_B, T) plane, and ξ in terms of this angle, the distance away from the critical point, and one non-universal parameter whose choice determines the contour on the phase diagram where $\xi = 1$ fm, as illustrated in Fig. 6-1. Because once ξ is less than 1 fm the magnitude of $\kappa_4 \propto \xi^7$ is negligible, for simplicity we set $\xi = 0$ outside the critical regime. (For details, see Refs. [171, 160, 5, 172].)

Following Ref. [163], we cast the momentum integration in terms of the momentum-space rapidity y and transverse momentum p_\perp , which are measured in experiment:

$$\int_{\mathbf{p}} \frac{1}{\gamma_{\mathbf{p}}} \rightarrow \frac{2m}{(2\pi)^3} \int_{y_c - \Delta y/2}^{y_c + \Delta y/2} dy \int_{p_{\min}}^{p_{\max}} p_\perp dp_\perp \int_0^{2\pi} d\psi. \quad (6.3)$$

We have introduced a finite acceptance in both rapidity and transverse momentum. We will keep $p_{\min} = 0.4$ GeV and $p_{\max} = 2$ GeV throughout. We shall compute κ_4 using two different kinds of rapidity cuts, either varying Δy with $y_c = 0$, in which case $|y| < y_{\max} \equiv \Delta y/2$, or varying y_c with fixed bin width Δy .

To simplify the interpretation of our results, we shall show the critical contribution to the cumulants normalized by the average number of protons, $\omega_{4,\sigma} \equiv \kappa_4[N]_\sigma / \langle N \rangle$. This cumulant ratio has the advantage that if the background (noncritical) contribution were Poisson-distributed it would contribute $\omega_{4,\sigma} = 1$, meaning that our results in Figures 6-1 and 6-2 should be interpreted as critical contributions to be added to a background of order 1.

Results and conclusions: In this Section, we demonstrate that the rapidity-dependence of μ_B makes the rapidity dependence of cumulants sensitive to critical fluctuations in a way that yields distinctive, qualitative, observable consequences. In Fig. 6-1 we first compute the dependence of the cumulant ratio $\omega_{4,\sigma}$ on the total rapidity acceptance y_{\max} . This dependence was studied previously in Ref. [163] upon assuming that μ_B itself is constant in rapidity; we find striking consequences of the rapidity-dependence of μ_B . Next, motivated by the expanded rapidity coverage that the STAR iTPC upgrade will bring, we compute $\omega_{4,\sigma}$ for bins in rapidity, something that has not been considered previously. We find that the rapidity dependence of the cumulant ratio is a sensitive and interesting probe of critical behavior.

In Fig. 6-1 we consider a hypothetical set of scenarios motivated by the possibility that there may be a QCD critical point within the energy range to be explored by the RHIC BES. We imagine a critical point at $\mu_B^c = 260$ MeV, and in the first three rows of the figure we consider heavy ion collisions with three decreasing values of the beam energy such that freezeout at mid-rapidity occurs at $\mu_{B,0} = 200, 230, 240$ MeV. In all three rows, we choose $\alpha = 50$ MeV, corresponding to the measured value from SPS collisions with $\sqrt{s} = 17.3$ GeV and $\mu_{B,0} = 237$ MeV. Because a real critical point may lie at larger μ_B^c than this, where α may become negative, in the fourth row we flip the sign of α . The right column of Figure 6-1 shows $\omega_{4,\sigma}$ binned in rapidity bins of width $\Delta y = 0.4$ centered around $y = \pm y_c$, an observable which to our knowledge has not been considered before. This is a more sensitive observable to the unique features of critical behavior than the dependence on the total rapidity acceptance in the center column because it isolates contributions coming from more similar values of μ_B , and the correlation length and other features of the critical regime are sensitive to μ_B near μ_B^c . We see many interesting qualitative features in the rapidity dependence of $\omega_{4,\sigma}$. For example, if $\mu_{B,0}$ is in the red region, where $\omega_{4,\sigma}$ is negative and relatively small in magnitude, larger and positive contributions to $\omega_{4,\sigma}$ can be found at larger rapidity. This can be seen in the middle panels of the first and second rows, but it is much more striking in the right panels, indicating the value of binning in rapidity. On the other hand, if $\mu_{B,0}$ lies in the blue region, in the right column the largest

value of $\omega_{4,\sigma}$ is obtained for the bin centered at $y = 0$, with $\omega_{4,\sigma}$ decreasing with increasing rapidity while staying positive if $\alpha > 0$ as in the third row or decreasing with increasing rapidity while becoming negative if $\alpha < 0$ as in the fourth row. Both the sign change and the non-monotonic behavior in $\omega_{4,\sigma}$, as a function of the rapidity acceptance in the center panel of Fig. 6-1 and even more so as a function of the rapidity bin in the right panel of Fig. 6-1 are new results of this work. They arise from the rapidity dependence of μ_B at freezeout in collisions at RHIC BES energies, and provide distinctive signatures if decreasing the beam energy in this scan takes $\mu_{B,0}$ past a critical point.

To complement Fig. 6-1, Fig. 6-2 shows the cumulant ratio $\omega_{4,\sigma}$ binned in rapidity for a fixed beam energy (fixed $\mu_{B,0}$) as the location of the critical point μ_B^c is changed. There are several features of binning the cumulants in rapidity which we believe will make doing so an important way to probe the critical region, if a critical point is discovered in the RHIC BES. First, $\omega_{4,\sigma}$ increases with $|y_c|$ if freezeout at mid-rapidity occurs at a $\mu_{B,0}$ that is well below μ_B^c , in the red region, whereas it will decrease with $|y_c|$ if $\mu_{B,0}$ is closer to or larger than μ_B^c , in the blue region. This remains true even if α changes sign, as demonstrated in the bottom row of Fig. 6-1. Furthermore, a sign change in $\omega_{4,\sigma}$ as a function of y will be easier to see upon binning in $|y_c|$ since not doing so, as in the middle panels, can obscure it by mixing data from different regions in rapidity. Even in cases where the sign change in $\omega_{4,\sigma}$ is visible in the middle column, as in the first and second rows, it happens at a lower rapidity in the right panel than in the middle panel, making it more feasible to observe at STAR via binning in $|y_c|$.

We conclude that the rapidity-dependence of μ_B at RHIC energies may result in qualitative signatures of critical fluctuations manifest in the rapidity-dependence of the cumulant ratio $\omega_{4,\sigma}$. Complementary to scanning the phase diagram by taking steps in beam energy, the rapidity dependence of μ_B provides additional scans of small regions of the phase diagram. We have seen that non-monotonicity and a sign change of the critical contribution to $\omega_{4,\sigma}$ as a function of rapidity will arise if the BES includes energies on both sides of a critical point. Binning the cumulants in

rapidity provides a sensitive probe of these effects. Signatures of critical behavior in the \sqrt{s} -dependence of $\omega_{4,\sigma}$ can therefore be cross-checked by looking for qualitative changes in the rapidity-dependence of $\omega_{4,\sigma}$ between beam energies on either side of the critical point. We have made arbitrary choices at many points, for illustrative purposes. A future quantitative study should include investigation of changes to these choices, as well as an investigation of consequences of various relevant physical effects which we have neglected in this exploratory study. These include the consequences of the variation of the baryon density across the fireball at a given spacetime rapidity originating from fluctuations in baryon stopping, for example as in Ref. [4]. We note, however, that determining the value of α from experimental data as in Ref. [166] as we propose, will incorporate the most important such consequence. The value of α obtained in this way is an average over many events with fixed \sqrt{s} , meaning that it will be important in future work to assess how κ_4 is influenced by event-by-event fluctuations in baryon stopping over and above their effect on the value of α . Future studies should also include an analysis of the quantitative effects of non-equilibrium dynamics, in particular critical slowing down, on the growth of the correlation length and hence on the values of \sqrt{s} or y_c at which the qualitative features that we have found occur [170, 173, 174, 175]. The effects of baryon number conservation as well as fluctuations in baryon stopping on the cumulants have been studied [176, 177] and should be included in a quantitative model.

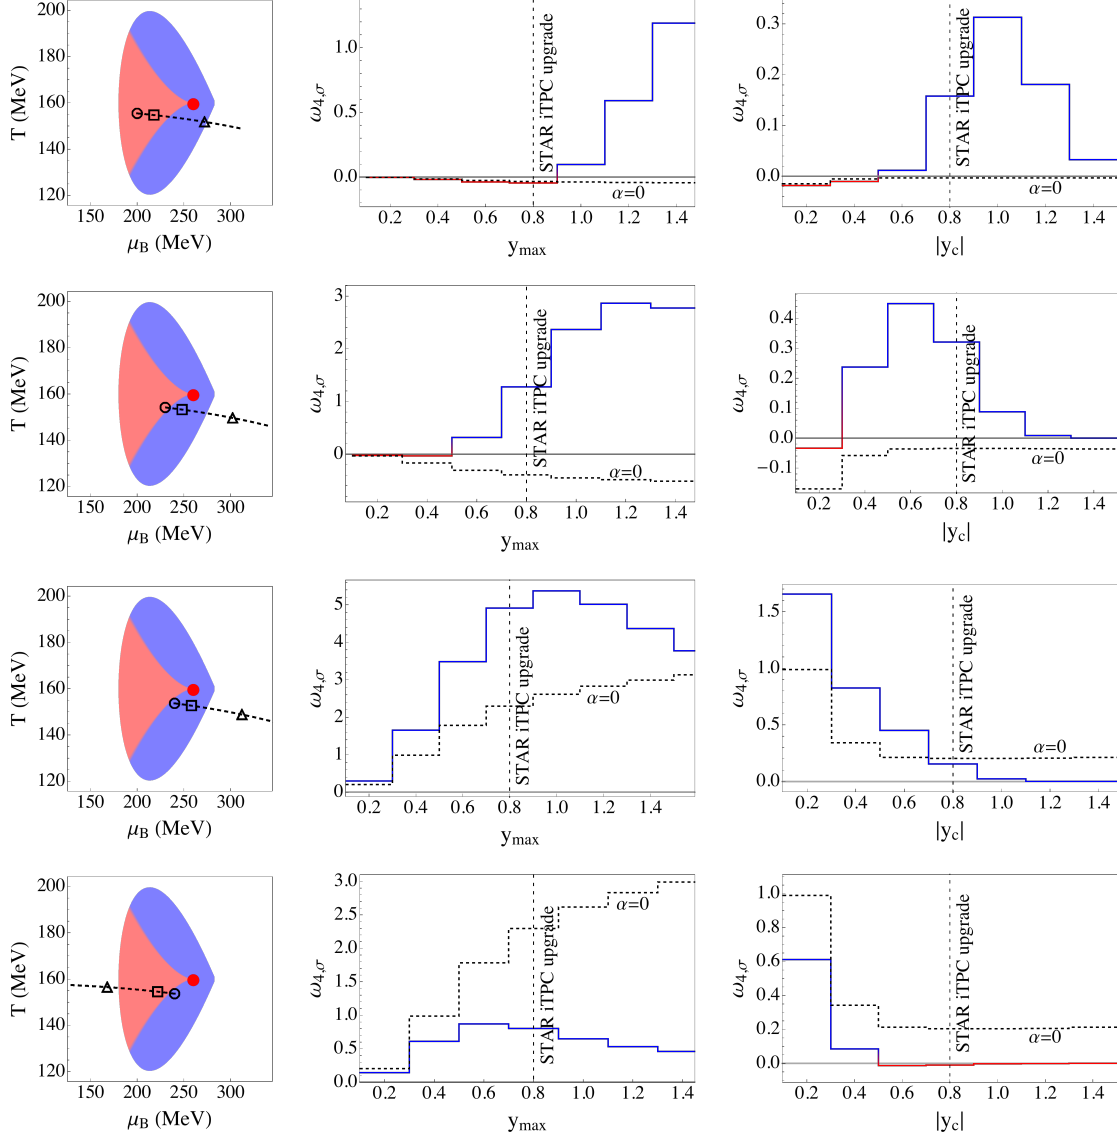


Figure 6-1: In the left column we see that we have assumed the existence of a critical point (red dot) at $(\mu_B^c, T) = (260, 160)$ MeV whose critical region, bounded by the contour where $\xi = 1$ fm, is colored red and blue. The colors denote the sign of $\omega_{4,\sigma}$, with $\omega_{4,\sigma} > 0$ in blue and $\omega_{4,\sigma} < 0$ in red. Different rows correspond to different assumptions for where on the phase diagram a heavy ion collision freezes out, *cf.* collisions with varying beam energy. The black circles show where freezeout occurs at mid-rapidity, from top to bottom with $\mu_{B,0} = 200, 230, 240, 240$ MeV. The black dashed curves show how the freezeout conditions change with increasing spacetime rapidity, with the circle, square, and triangle indicating freezeout at $y_s = 0, 0.6,$ and $1.2,$ respectively.

Figure 6-1: (continued) In the top three rows, we have chosen $\alpha = 50$ MeV (see eq. 6.1) while for the bottom row we have chosen $\alpha = -50$ MeV. The middle column shows how $\omega_{4,\sigma}$ computed for a rapidity acceptance $|y| < y_{\max}$ depends on y_{\max} . The right column shows how $\omega_{4,\sigma}$ computed in a pair of bins with width $\Delta y = 0.4$ centered at $\pm y_c$ depends on y_c . The results in the middle column sum over a wide range of rapidities (with $|y|$ between 0 and y_{\max}) which freezeout with a range of μ_B , meaning that features from the left column are more directly visible in the right column than in the middle. In both the center and right columns, the black dotted lines show $\omega_{4,\sigma}$ with $\alpha = 0$, *i.e.* what would have been obtained if $\mu_B = \mu_{B,0}$, denoted by the black circles in the left column, everywhere. The results shown in the right and middle column should not be taken as quantitative predictions since they depend on the many assumptions that we made for illustrative purposes; they are illustrative of qualitative features to be expected in the rapidity-dependence of cumulants if steps in beam energy take us past a critical point.




Figure 6-2: We illustrate the behavior of $\omega_{4,\sigma}$ when the freezeout conditions are as in the first row of Figure 6-1 with $\mu_{B,0} = 200$ MeV and $\alpha = 50$ MeV, but where here we let the location of the critical point range from $\mu_B^c = 190$ MeV to $\mu_B^c = 300$ MeV. For each value of μ_B^c , namely for each horizontal slice across the figure, color indicates the value of $\omega_{4,\sigma}$ as a function of $|y_c|$, with $\Delta y = 0.4$ fixed as in Fig. 6-1. The slice indicated by the green dashed line corresponds to the top-right panel of Fig. 6-1. For $\mu_B^c < \mu_{B,0}$ ($\mu_B^c > \mu_{B,0}$), $\omega_{4,\sigma}$ decreases (increases) with increasing $|y_c|$.

Chapter 7

The Road Ahead

Over the last few decades, heavy-ion collisions have proven to be an important avenue for studying the many-body physics of Quantum Chromodynamics, with exciting and unexpected results around every corner. The physics of quark–gluon plasma is particularly rich and necessitates a multifaceted approach to understanding its properties. In this thesis, we have presented several avenues of recent progress on understanding this exotic material from its fluid-like behavior, quenching of high-energy particles, and its phase diagram. However, this is of course far from the whole story and there is substantial work to be done in each of these directions.

It remains to be understood why hydrodynamics appears to provide a quantitative description of the highly non-equilibrium quark–gluon plasma produced in heavy-ion collisions. The issue is even more pronounced in smaller systems like proton–ion and high-multiplicity proton–proton collisions, where flow-like correlations have also been observed and described by hydrodynamics despite the fact that these systems are even further from equilibrium. Though hydrodynamic attractors have drawn much recent theoretical attention, their study has so far been restricted to relatively simple systems and their implications for the emergence of hydrodynamic behavior in quark–gluon plasma has yet to be quantified.

At the same time, measurements of jet modification in heavy-ion collisions have become increasingly precise and differential. Regardless, the complexity of modelling jet energy loss in heavy-ion collisions means it remains difficult to use measurements

to discriminate between models with different energy loss mechanisms. Addressing this will likely require concerted effort, both on improving models and developing observables and analysis techniques that enhance the sensitivity of measurements to salient features of models. Quantitative statements about modification effects are especially crucial in small collision systems, where flow measurements suggest that quark–gluon plasma may be produced but no jet modification has been observed. A convincing discovery of quark–gluon plasma in these systems will presumably require observing its effects both in flow measurements and in jet quenching.

A hypothetical critical point in the phase diagram of QCD is still being searched for in the Beam Energy Scan experiment, and it remains crucial to identify observables that are sensitive to the passage of quark–gluon plasma through a critical point. A major problem is that the quark–gluon plasma transiting the phase diagram is far from equilibrium, and therefore necessitates understanding signatures of criticality for a system that is not in equilibrium.

Finally, some of the big emerging questions in the field are how these directions are intertwined. For example, it is unknown to what extent hard processes like jets participate in the process of equilibration and thermalization of quark–gluon plasma in heavy-ion collisions. In principle jet energy loss and equilibration are dictated by the same types of QCD processes, though on different scales, so it may be possible to understand them in a unified framework. In addition, if the quark–gluon plasma is near equilibrium at the time of freezeout, it may retain relatively little information about its far-from-equilibrium structure. The wake of high-energy processes like jets in the quark–gluon plasma may thus provide important phenomenological access to the far-from-equilibrium response and the intermediate stages of equilibration.

The experimental discovery of the quark–gluon plasma was relatively recent, and so we are only at the beginning of an exciting program of studying high temperature QCD through the phenomenology of heavy-ion collisions. In addition to the continued data taking for the Beam Energy Scan II at STAR, there will be sPHENIX [178] at RHIC which will be focused on physics of intermediate-energy jets, which may be a crucial laboratory for studying the interplay of hard processes and equilibration.

At the LHC, there is planned heavy-ion running with higher luminosity through Run 4, and proposals for programs at the high-luminosity LHC and a Future Circular Collider [138]. There has also recently been approval to build an electron–ion collider at Brookhaven National Laboratory to study the structure of nuclei through deep inelastic scattering [179]. With a wealth of upcoming measurements and new theoretical advancements, we look forward to exciting insights on nuclear physics and high-temperature QCD in the coming decades.

Bibliography

- [1] *Beam Energy Scan White Paper: Studying the Phase Diagram of QCD Matter at RHIC*. https://drupal.star.bnl.gov/STAR/files/BES_WPII_ver6.9_Cover.pdf.
- [2] Liam Keegan, Alekski Kurkela, Paul Romatschke, Wilke van der Schee, and Yan Zhu. Weak and strong coupling equilibration in nonabelian gauge theories. *JHEP*, 04:031, 2016.
- [3] Paul Romatschke. Relativistic Fluid Dynamics Far From Local Equilibrium. *Phys. Rev. Lett.*, 120(1):012301, 2018.
- [4] Chun Shen and Björn Schenke. Dynamical initial state model for relativistic heavy-ion collisions. *Phys. Rev. C*, 97(2):024907, 2018.
- [5] M.A. Stephanov. On the sign of kurtosis near the QCD critical point. *Phys. Rev. Lett.*, 107:052301, 2011.
- [6] Michal P. Heller, Alekski Kurkela, Michal Spalinski, and Viktor Svensson. Hydrodynamization in kinetic theory: Transient modes and the gradient expansion. *Phys. Rev.*, D97(9):091503, 2018.
- [7] Michal P. Heller and Viktor Svensson. How does relativistic kinetic theory remember about initial conditions? *Phys. Rev.*, D98(5):054016, 2018.
- [8] David J. Gross and Frank Wilczek. Ultraviolet Behavior of Nonabelian Gauge Theories. *Phys. Rev. Lett.*, 30:1343–1346, 1973.
- [9] H.David Politzer. Reliable Perturbative Results for Strong Interactions? *Phys. Rev. Lett.*, 30:1346–1349, 1973.
- [10] S. Durr et al. Ab-Initio Determination of Light Hadron Masses. *Science*, 322:1224–1227, 2008.
- [11] P.W. Anderson. More Is Different. *Science*, 177(4047):393–396, 1972.
- [12] C. Adler et al. Identified particle elliptic flow in Au + Au collisions at $\sqrt{s_{NN}} = 130$ GeV. *Phys. Rev. Lett.*, 87:182301, 2001.
- [13] I. Arsene et al. Quark gluon plasma and color glass condensate at RHIC? The Perspective from the BRAHMS experiment. *Nucl. Phys.*, A757:1–27, 2005.

- [14] B. B. Back et al. The PHOBOS perspective on discoveries at RHIC. *Nucl. Phys.*, A757:28–101, 2005.
- [15] K. Adcox et al. Formation of dense partonic matter in relativistic nucleus-nucleus collisions at RHIC: Experimental evaluation by the PHENIX collaboration. *Nucl. Phys.*, A757:184–283, 2005.
- [16] John Adams et al. Experimental and theoretical challenges in the search for the quark gluon plasma: The STAR Collaboration’s critical assessment of the evidence from RHIC collisions. *Nucl. Phys.*, A757:102–183, 2005.
- [17] K Aamodt et al. Elliptic flow of charged particles in Pb-Pb collisions at 2.76 TeV. *Phys. Rev. Lett.*, 105:252302, 2010.
- [18] Georges Aad et al. Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.77$ TeV with the ATLAS Detector at the LHC. *Phys. Rev. Lett.*, 105:252303, 2010.
- [19] Serguei Chatrchyan et al. Observation and studies of jet quenching in PbPb collisions at nucleon-nucleon center-of-mass energy = 2.76 TeV. *Phys. Rev.*, C84:024906, 2011.
- [20] Paul Romatschke and Ulrike Romatschke. *Relativistic Fluid Dynamics In and Out of Equilibrium*. Cambridge Monographs on Mathematical Physics. Cambridge University Press, 2019.
- [21] B. Alver, M. Baker, C. Loizides, and P. Steinberg. The PHOBOS Glauber Monte Carlo. 2008.
- [22] J. Adams et al. Azimuthal anisotropy at RHIC: The First and fourth harmonics. *Phys. Rev. Lett.*, 92:062301, 2004.
- [23] Paul Romatschke and Ulrike Romatschke. Viscosity Information from Relativistic Nuclear Collisions: How Perfect is the Fluid Observed at RHIC? *Phys. Rev. Lett.*, 99:172301, 2007.
- [24] Bjorn Schenke, Sangyong Jeon, and Charles Gale. Elliptic and triangular flow in event-by-event (3+1)D viscous hydrodynamics. *Phys. Rev. Lett.*, 106:042301, 2011.
- [25] Thomas Schäfer and Derek Teaney. Nearly Perfect Fluidity: From Cold Atomic Gases to Hot Quark Gluon Plasmas. *Rept. Prog. Phys.*, 72:126001, 2009.
- [26] P. Kovtun, Dan T. Son, and Andrei O. Starinets. Viscosity in strongly interacting quantum field theories from black hole physics. *Phys. Rev. Lett.*, 94:111601, 2005.
- [27] Serguei Chatrchyan et al. Study of high-pT charged particle suppression in PbPb compared to pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Eur. Phys. J. C*, 72:1945, 2012.

- [28] Betty Abelev et al. Centrality Dependence of Charged Particle Production at Large Transverse Momentum in Pb–Pb Collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. *Phys. Lett. B*, 720:52–62, 2013.
- [29] Georges Aad et al. Measurement of charged-particle spectra in Pb+Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV with the ATLAS detector at the LHC. *JHEP*, 09:050, 2015.
- [30] Jaroslav Adam et al. Measurement of jet suppression in central Pb-Pb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. *Phys. Lett.*, B746:1–14, 2015.
- [31] Serguei Chatrchyan et al. Jet momentum dependence of jet quenching in PbPb collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. *Phys. Lett. B*, 712:176–197, 2012.
- [32] Megan Connors, Christine Nattrass, Rosi Reed, and Sevil Salur. Jet measurements in heavy ion physics. *Rev. Mod. Phys.*, 90:025005, 2018.
- [33] Wit Busza, Krishna Rajagopal, and Wilke van der Schee. Heavy Ion Collisions: The Big Picture, and the Big Questions. *Ann. Rev. Nucl. Part. Sci.*, 68:339–376, 2018.
- [34] Aleksı Kurkela, Urs Achim Wiedemann, and Bin Wu. Flow in AA and pA as an interplay of fluid-like and non-fluid like excitations. *Eur. Phys. J. C*, 79(11):965, 2019.
- [35] Michal P. Heller and Michal Spalinski. Hydrodynamics Beyond the Gradient Expansion: Resurgence and Resummation. *Phys. Rev. Lett.*, 115(7):072501, 2015.
- [36] Serguei Chatrchyan et al. Modification of Jet Shapes in PbPb Collisions at $\sqrt{s_{\text{NN}}} = 2.76$ TeV. *Phys. Lett. B*, 730:243–263, 2014.
- [37] Krishna Rajagopal, Andrey V. Sadofyev, and Wilke van der Schee. Evolution of the jet opening angle distribution in holographic plasma. *Phys. Rev. Lett.*, 116(21):211603, 2016.
- [38] Jasmine Brewer, Krishna Rajagopal, Andrey Sadofyev, and Wilke Van Der Schee. Evolution of the Mean Jet Shape and Dijet Asymmetry Distribution of an Ensemble of Holographic Jets in Strongly Coupled Plasma. *JHEP*, 02:015, 2018.
- [39] Jasmine Brewer, Krishna Rajagopal, Andrey Sadofyev, and Wilke van der Schee. Holographic Jet Shapes and their Evolution in Strongly Coupled Plasma. *Nucl. Phys. A*, 967:508–511, 2017.
- [40] A. Bazavov et al. Equation of state in (2+1)-flavor QCD. *Phys. Rev.*, D90:094503, 2014.
- [41] Mikhail A. Stephanov. QCD phase diagram and the critical point. *Prog. Theor. Phys. Suppl.*, 153:139–156, 2004.

- [42] Chun Shen and Björn Schenke. Dynamical initialization and hydrodynamic modeling of relativistic heavy-ion collisions. *Nucl. Phys. A*, 982:411–414, 2019.
- [43] I.G. Bearden et al. Nuclear stopping in Au + Au collisions at $\sqrt{s_{\text{NN}}} = 200$ GeV. *Phys. Rev. Lett.*, 93:102301, 2004.
- [44] Misha A. Stephanov, K. Rajagopal, and Edward V. Shuryak. Signatures of the tricritical point in QCD. *Phys. Rev. Lett.*, 81:4816–4819, 1998.
- [45] Ani Aprahamian et al. Reaching for the horizon: The 2015 long range plan for nuclear science. 2015.
- [46] Misha A. Stephanov, K. Rajagopal, and Edward V. Shuryak. Event-by-event fluctuations in heavy ion collisions and the QCD critical point. *Phys. Rev. D*, 60:114028, 1999.
- [47] M.A. Stephanov. Non-Gaussian fluctuations near the QCD critical point. *Phys. Rev. Lett.*, 102:032301, 2009.
- [48] Xiaofeng Luo. Exploring the QCD Phase Structure with Beam Energy Scan in Heavy-ion Collisions. *Nucl. Phys. A*, 956:75–82, 2016.
- [49] Jasmine Brewer, Li Yan, and Yi Yin. Adiabatic hydrodynamization in rapidly-expanding quark-gluon plasma. 2019.
- [50] Jasmine Brewer, José Guilherme Milhano, and Jesse Thaler. Sorting out quenched jets. *Phys. Rev. Lett.*, 122(22):222301, 2019.
- [51] Jasmine Brewer, Swagato Mukherjee, Krishna Rajagopal, and Yi Yin. Searching for the QCD critical point via the rapidity dependence of cumulants. *Phys. Rev. C*, 98(6):061901, 2018.
- [52] Jasmine Brewer, Andrey Sadofyev, and Wilke van der Schee. Jet shape modifications in holographic dijet systems. 2018.
- [53] H. Bantilan, J.T. Brewer, T. Ishii, W.E. Lewis, and P. Romatschke. String-theory-based predictions for nonhydrodynamic collective modes in strongly interacting Fermi gases. *Phys. Rev. A*, 94(3):033621, 2016.
- [54] Ulrich W. Heinz. Towards the Little Bang Standard Model. *J. Phys. Conf. Ser.*, 455:012044, 2013.
- [55] Michal P. Heller. Holography, Hydrodynamization and Heavy-Ion Collisions. *Acta Phys. Polon.*, B47:2581, 2016.
- [56] Wojciech Florkowski, Michal P. Heller, and Michal Spalinski. New theories of relativistic hydrodynamics in the LHC era. *Rept. Prog. Phys.*, 81(4):046001, 2018.

- [57] Wojciech Florkowski, Radoslaw Ryblewski, and Michael Strickland. Testing viscous and anisotropic hydrodynamics in an exactly solvable case. *Phys. Rev.*, C88:024903, 2013.
- [58] Paul Romatschke. Retarded correlators in kinetic theory: branch cuts, poles and hydrodynamic onset transitions. *Eur. Phys. J.*, C76(6):352, 2016.
- [59] Gabriel S. Denicol and Jorge Noronha. Divergence of the Chapman-Enskog expansion in relativistic kinetic theory. 2016.
- [60] Gabriel S. Denicol and Jorge Noronha. Analytical attractor and the divergence of the slow-roll expansion in relativistic hydrodynamics. *Phys. Rev.*, D97(5):056021, 2018.
- [61] C. De Grandi and A. Polkovnikov. Adiabatic perturbation theory: From Landau-Zener problem to quenching through a quantum critical point. *Lecture Notes in Physics*, page 75–114, 2010.
- [62] R. Baier, Alfred H. Mueller, D. Schiff, and D. T. Son. 'Bottom up' thermalization in heavy ion collisions. *Phys. Lett.*, B502:51–58, 2001.
- [63] Michael Strickland, Jorge Noronha, and Gabriel Denicol. Anisotropic nonequilibrium hydrodynamic attractor. *Phys. Rev.*, D97(3):036020, 2018.
- [64] Sunil Jaiswal, Chandrodoy Chattopadhyay, Amaresh Jaiswal, Subrata Pal, and Ulrich Heinz. Exact solutions and attractors of higher-order viscous fluid dynamics for Bjorken flow. *Phys. Rev.*, C100(3):034901, 2019.
- [65] Alireza Behtash, C. N. Cruz-Camacho, Syo Kamata, and M. Martinez. Non-perturbative rheological behavior of a far-from-equilibrium expanding plasma. *Phys. Lett.*, B797:134914, 2019.
- [66] Jean-Paul Blaizot and Li Yan. Emergence of hydrodynamical behavior in expanding ultra-relativistic plasmas. *Annals Phys.*, 412:167993, 2020.
- [67] Chandrodoy Chattopadhyay, Ulrich Heinz, Subrata Pal, and Gojko Vujanovic. Higher order and anisotropic hydrodynamics for Bjorken and Gubser flows. *Phys. Rev.*, C97(6):064909, 2018.
- [68] G. S. Denicol, H. Niemi, E. Molnar, and D. H. Rischke. Derivation of transient relativistic fluid dynamics from the Boltzmann equation. *Phys. Rev.*, D85:114047, 2012. [Erratum: *Phys. Rev.* D91, no. 3, 039902 (2015)].
- [69] Jean-Paul Blaizot and Li Yan. Fluid dynamics of out of equilibrium boost invariant plasmas. *Phys. Lett.*, B780:283–286, 2018.
- [70] Aleksi Kurkela, Wilke van der Schee, Urs Achim Wiedemann, and Bin Wu. Early- and Late-Time Behavior of Attractors in Heavy-Ion Collisions. *Phys. Rev. Lett.*, 124(10):102301, 2020.

- [71] J. P. Blaizot and Alfred H. Mueller. The Early Stage of Ultrarelativistic Heavy Ion Collisions. *Nucl. Phys.*, B289:847–860, 1987.
- [72] Larry D. McLerran and Raju Venugopalan. Computing quark and gluon distribution functions for very large nuclei. *Phys. Rev.*, D49:2233–2241, 1994.
- [73] Larry D. McLerran and Raju Venugopalan. Green’s functions in the color field of a large nucleus. *Phys. Rev.*, D50:2225–2233, 1994.
- [74] Jamal Jalilian-Marian, Alex Kovner, Larry D. McLerran, and Heribert Weigert. The Intrinsic glue distribution at very small x. *Phys. Rev.*, D55:5414–5428, 1997.
- [75] Gustavo Rigolin, Gerardo Ortiz, and Victor Hugo Ponce. Beyond the quantum adiabatic approximation: Adiabatic perturbation theory. *Physical Review A*, 78(5):052508, 2008.
- [76] Hans Bantilan, Takaaki Ishii, and Paul Romatschke. Holographic Heavy-Ion Collisions: Analytic Solutions with Longitudinal Flow, Elliptic Flow and Vorticity. *Phys. Lett.*, B785:201–206, 2018.
- [77] Jasmine Brewer and Paul Romatschke. Nonhydrodynamic Transport in Trapped Unitary Fermi Gases. *Phys. Rev. Lett.*, 115(19):190404, 2015.
- [78] David A. Appel. Jets as a Probe of Quark - Gluon Plasmas. *Phys. Rev.*, D33:717, 1986.
- [79] J. P. Blaizot and Larry D. McLerran. Jets in Expanding Quark - Gluon Plasmas. *Phys. Rev.*, D34:2739, 1986.
- [80] Miklos Gyulassy and Michael Plumer. Jet Quenching in Dense Matter. *Phys. Lett.*, B243:432–438, 1990.
- [81] Xin-Nian Wang and Miklos Gyulassy. Gluon shadowing and jet quenching in A + A collisions at $\sqrt{s_{NN}} = 200$ GeV. *Phys. Rev. Lett.*, 68:1480–1483, 1992.
- [82] Serguei Chatrchyan et al. Studies of jet quenching using isolated-photon+jet correlations in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Lett.*, B718:773–794, 2013.
- [83] 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,no.2,029903(2015)].
- [84] Serguei Chatrchyan et al. Measurement of jet fragmentation in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Rev.*, C90(2):024908, 2014.
- [85] Vardan Khachatryan et al. Measurement of transverse momentum relative to dijet systems in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 01:006, 2016.

- [86] Morad Aaboud et al. Measurement of jet p_T correlations in Pb+Pb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS detector. *Phys. Lett.*, B774:379–402, 2017.
- [87] S. Acharya et al. First measurement of jet mass in Pb–Pb and p–Pb collisions at the LHC. *Phys. Lett.*, B776:249–264, 2018.
- [88] Morad Aaboud et al. Measurement of the nuclear modification factor for inclusive jets in Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector. *Phys. Lett. B*, 790:108–128, 2019.
- [89] Morad Aaboud et al. Measurement of jet fragmentation in Pb+Pb and pp collisions at $\sqrt{s_{NN}} = 5.02$ TeV with the ATLAS detector. *Phys. Rev.*, C98(2):024908, 2018.
- [90] Shreyasi Acharya et al. Medium modification of the shape of small-radius jets in central Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *JHEP*, 10:139, 2018.
- [91] Jorge Casalderrey-Solana and Carlos A. Salgado. Introductory lectures on jet quenching in heavy ion collisions. *Acta Phys. Polon.*, B38:3731–3794, 2007.
- [92] David d’Enterria. Jet quenching. *Landolt-Bornstein*, 23:471, 2010.
- [93] Vardan Khachatryan et al. Measurement of inclusive jet cross sections in pp and PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV. *Phys. Rev.*, C96(1):015202, 2017.
- [94] Korinna C. Zapp. JEWEL 2.0.0: directions for use. *Eur. Phys. J.*, C74(2):2762, 2014.
- [95] Raghav Kunnawalkam Elayavalli and Korinna Christine Zapp. Simulating V+jet processes in heavy ion collisions with JEWEL. *Eur. Phys. J.*, C76(12):695, 2016.
- [96] Alberica Toia. ALICE Measurements in p-Pb Collisions: Charged Particle Multiplicity, Centrality Determination and implications for Binary Scaling. *Nucl. Phys.*, A926:78–84, 2014.
- [97] S. S. Adler et al. A Detailed Study of High-p(T) Neutral Pion Suppression and Azimuthal Anisotropy in Au+Au Collisions at $\sqrt{s_{NN}} = 200$ GeV. *Phys. Rev.*, C76:034904, 2007.
- [98] A. Adare et al. Neutral pion production with respect to centrality and reaction plane in Au+Au collisions at $\sqrt{s_{NN}} = 200$ GeV. *Phys. Rev. C*, 87(3):034911, 2013.
- [99] A. Adare et al. Scaling properties of fractional momentum loss of high- p_T hadrons in nucleus-nucleus collisions at $\sqrt{s_{NN}}$ from 62.4 GeV to 2.76 TeV. *Phys. Rev.*, C93(2):024911, 2016.

- [100] Harry Arthur Andrews et al. Novel tools and observables for jet physics in heavy-ion collisions. *J. Phys. G*, 47(6):065102, 2020.
- [101] Albert M Sirunyan et al. Study of Jet Quenching with $Z + \text{jet}$ Correlations in Pb-Pb and pp Collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. *Phys. Rev. Lett.*, 119(8):082301, 2017.
- [102] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. PYTHIA 6.4 Physics and Manual. *JHEP*, 05:026, 2006.
- [103] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. The anti- k_t jet clustering algorithm. *JHEP*, 04:063, 2008.
- [104] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. FastJet User Manual. *Eur. Phys. J.*, C72:1896, 2012.
- [105] Chun Shen and Ulrich Heinz. Collision Energy Dependence of Viscous Hydrodynamic Flow in Relativistic Heavy-Ion Collisions. *Phys. Rev.*, C85:054902, 2012. [Erratum: *Phys. Rev.*C86,049903(2012)].
- [106] Liliana Apolinário, João Barata, and Guilherme Milhano. Quark and Gluon Jet Energy Loss. *PoS, HardProbes2018*:082, 2018.
- [107] Yayun He, Long-Gang Pang, and Xin-Nian Wang. Bayesian extraction of jet energy loss distributions in heavy-ion collisions. *Phys. Rev. Lett.*, 122(25):252302, 2019.
- [108] José Guilherme Milhano and Korinna Christine Zapp. Origins of the di-jet asymmetry in heavy ion collisions. *Eur. Phys. J.*, C76(5):288, 2016.
- [109] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- [110] 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.
- [111] Chanwook Park, Sangyong Jeon, and Charles Gale. Jet modification with medium recoil in quark-gluon plasma. *Nucl. Phys. A*, 982:643–646, 2019.
- [112] Jonathan M. Butterworth, Adam R. Davison, Mathieu Rubin, and Gavin P. Salam. Jet substructure as a new Higgs search channel at the LHC. *Phys. Rev. Lett.*, 100:242001, 2008.
- [113] Stephen D. Ellis, Christopher K. Vermilion, and Jonathan R. Walsh. Recombination Algorithms and Jet Substructure: Pruning as a Tool for Heavy Particle Searches. *Phys. Rev.*, D81:094023, 2010.
- [114] David Krohn, Jesse Thaler, and Lian-Tao Wang. Jet Trimming. *JHEP*, 02:084, 2010.

- [115] Mrinal Dasgupta, Alessandro Fregoso, Simone Marzani, and Gavin P. Salam. Towards an understanding of jet substructure. *JHEP*, 09:029, 2013.
- [116] Andrew J. Larkoski, Simone Marzani, Gregory Soyez, and Jesse Thaler. Soft Drop. *JHEP*, 05:146, 2014.
- [117] Simon Caron-Huot. O(g) plasma effects in jet quenching. *Phys. Rev. D*, 79:065039, 2009.
- [118] Martin Spousta and Brian Cole. Interpreting single jet measurements in Pb + Pb collisions at the LHC. *Eur. Phys. J. C*, 76(2):50, 2016.
- [119] Yang-Ting Chien and Raghav Kunnawalkam Elayavalli. Probing heavy ion collisions using quark and gluon jet substructure. 2018.
- [120] Yacine Mehtar-Tani and Soeren Schlichting. Universal quark to gluon ratio in medium-induced parton cascade. *JHEP*, 09:144, 2018.
- [121] 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.
- [122] Liliana Apolinário, João Barata, and Guilherme Milhano. On the breaking of Casimir scaling in jet quenching. *Eur. Phys. J. C*, 80(6):586, 2020.
- [123] Eric M. Metodiev and Jesse Thaler. Jet Topics: Disentangling Quarks and Gluons at Colliders. *Phys. Rev. Lett.*, 120(24):241602, 2018.
- [124] Patrick T. Komiske, Eric M. Metodiev, and Jesse Thaler. An operational definition of quark and gluon jets. *JHEP*, 11:059, 2018.
- [125] Albert M Sirunyan et al. Measurement of quark- and gluon-like jet fractions using jet charge in PbPb and pp collisions at 5.02 TeV. 4 2020.
- [126] Julian Katz-Samuels, Gilles Blanchard, and Clayton Scott. Decontamination of mutual contamination models. *Journal of Machine Learning Research*, 20(41):1–57, 2019.
- [127] Gilles Blanchard, Marek Flaska, Gregory Handy, Sara Pozzi, and Clayton Scott. Classification with asymmetric label noise: Consistency and maximal denoising. *Electron. J. Statist.*, 10(2):2780–2824, 2016.
- [128] Barry M. Dillon, Darius A. Faroughy, and Jernej F. Kamenik. Uncovering latent jet substructure. *Phys. Rev. D*, 100(5):056002, 2019.
- [129] Ezequiel Alvarez, Federico Lamagna, and Manuel Szewc. Topic Model for four-top at the LHC. *JHEP*, 01:049, 2020.
- [130] Barry M. Dillon, Darius A. Faroughy, Jernej F. Kamenik, and Manuel Szewc. Learning the latent structure of collider events. 5 2020.

- [131] Philippe Gras, Stefan Höche, Deepak Kar, Andrew Larkoski, Leif Lönnblad, Simon Plätzer, Andrzej Siódmok, Peter Skands, Gregory Soyez, and Jesse Thaler. Systematics of quark/gluon tagging. *JHEP*, 07:091, 2017.
- [132] Christopher Frye, Andrew J. Larkoski, Jesse Thaler, and Kevin Zhou. Casimir Meets Poisson: Improved Quark/Gluon Discrimination with Counting Observables. *JHEP*, 09:083, 2017.
- [133] Andrew J. Larkoski and Eric M. Metodiev. A Theory of Quark vs. Gluon Discrimination. *JHEP*, 10:014, 2019.
- [134] Daniel Foreman-Mackey, David W. Hogg, Dustin Lang, and Jonathan Goodman. emcee: The MCMC Hammer. *Publ. Astron. Soc. Pac.*, 125:306–312, 2013.
- [135] Steve Baker and Robert D. Cousins. Clarification of the Use of Chi Square and Likelihood Functions in Fits to Histograms. *Nucl. Instrum. Meth.*, 221:437–442, 1984.
- [136] M. Tanabashi et al. Review of Particle Physics. *Phys. Rev. D*, 98(3):030001, 2018.
- [137] Albert M Sirunyan et al. Measurement of the groomed jet mass in PbPb and pp collisions at $\sqrt{s_{\text{NN}}} = 5.02$ TeV. *JHEP*, 10:161, 2018.
- [138] Z. Citron et al. *Report from Working Group 5: Future physics opportunities for high-density QCD at the LHC with heavy-ion and proton beams*, volume 7, pages 1159–1410. 12 2019.
- [139] G. Roland, K. Safarik, and P. Steinberg. Heavy-ion collisions at the LHC. *Prog. Part. Nucl. Phys.*, 77:70–127, 2014.
- [140] 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: *JHEP* 09, 175 (2015)].
- [141] Yasuyuki Akiba et al. The Hot QCD White Paper: Exploring the Phases of QCD at RHIC and the LHC. 2015.
- [142] Frithjof Karsch. Lattice results on QCD thermodynamics. *Nucl. Phys.*, A698:199–208, 2002.
- [143] Y. Aoki, G. Endrodi, Z. Fodor, S. D. Katz, and K. K. Szabo. The Order of the quantum chromodynamics transition predicted by the standard model of particle physics. *Nature*, 443:675–678, 2006.
- [144] Szabocs Borsanyi, Zoltan Fodor, Christian Hoelbling, Sandor D. Katz, Stefan Krieg, and Kalman K. Szabo. Full result for the QCD equation of state with 2+1 flavors. *Phys. Lett.*, B730:99–104, 2014.

- [145] Tanmoy Bhattacharya et al. QCD Phase Transition with Chiral Quarks and Physical Quark Masses. *Phys. Rev. Lett.*, 113(8):082001, 2014.
- [146] Xiaofeng Luo and Nu Xu. Search for the QCD Critical Point with Fluctuations of Conserved Quantities in Relativistic Heavy-Ion Collisions at RHIC : An Overview. *Nucl. Sci. Tech.*, 28(8):112, 2017.
- [147] Philippe de Forcrand. Simulating QCD at finite density. *PoS*, LAT2009:010, 2009.
- [148] Heng-Tong Ding, Frithjof Karsch, and Swagato Mukherjee. Thermodynamics of strong-interaction matter from Lattice QCD. *Int. J. Mod. Phys.*, E24(10):1530007, 2015.
- [149] Krishna Rajagopal. Mapping the QCD phase diagram. *Nucl. Phys.*, A661:150–161, 1999.
- [150] Krishna Rajagopal and Frank Wilczek. The Condensed matter physics of QCD. In M. Shifman and Boris Ioffe, editors, *At the frontier of particle physics. Handbook of QCD. Vol. 1-3*, volume 3, pages 2061–2151. 2000.
- [151] M. A. Stephanov. QCD phase diagram: An Overview. *PoS*, LAT2006:024, 2006.
- [152] J. Cleymans, H. Oeschler, K. Redlich, and S. Wheaton. Comparison of chemical freeze-out criteria in heavy-ion collisions. *Phys. Rev.*, C73:034905, 2006.
- [153] 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.
- [154] N. S. Gerakisev et al. The Nuclotron-based Ion Collider Facility Project. The Physics Programme for the Multi-Purpose Detector. *J. Phys. Conf. Ser.*, 1023(1):012030, 2018.
- [155] T. Ablyazimov et al. Challenges in QCD matter physics –The scientific programme of the Compressed Baryonic Matter experiment at FAIR. *Eur. Phys. J.*, A53(3):60, 2017.
- [156] L. Adamczyk et al. Energy Dependence of Moments of Net-proton Multiplicity Distributions at RHIC. *Phys. Rev. Lett.*, 112:032302, 2014.
- [157] Xiaofeng Luo. Energy Dependence of Moments of Net-Proton and Net-Charge Multiplicity Distributions at STAR. *PoS*, CPOD2014:019, 2015.
- [158] Yaping Wang. STAR upgrade program and future physics. *J. Phys. Conf. Ser.*, 535:012022, 2014.

- [159] Y. Hatta and M. A. Stephanov. Proton number fluctuation as a signal of the QCD critical endpoint. *Phys. Rev. Lett.*, 91:102003, 2003. [Erratum: *Phys. Rev. Lett.*91,129901(2003)].
- [160] Christiana Athanasiou, Krishna Rajagopal, and Misha Stephanov. Using Higher Moments of Fluctuations and their Ratios in the Search for the QCD Critical Point. *Phys. Rev.*, D82:074008, 2010.
- [161] M. A. Stephanov. QCD critical point and event-by-event fluctuations. *J. Phys.*, G38:124147, 2011.
- [162] J. Steinheimer, M. Bleicher, H. Petersen, S. Schramm, H. Stoecker, and D. Zschiesche. (3+1)-dimensional hydrodynamic expansion with a critical point from realistic initial conditions. *Phys. Rev.*, C77:034901, 2008.
- [163] Bo Ling and Mikhail A. Stephanov. Acceptance dependence of fluctuation measures near the QCD critical point. *Phys. Rev.*, C93(3):034915, 2016.
- [164] Adam Bzdak, Volker Koch, and Nils Strodthoff. Cumulants and correlation functions versus the QCD phase diagram. *Phys. Rev.*, C95(5):054906, 2017.
- [165] Ming Li and Joseph I. Kapusta. High Baryon Densities in Heavy Ion Collisions at Energies Attainable at the BNL Relativistic Heavy Ion Collider and the CERN Large Hadron Collider. *Phys. Rev.*, C95(1):011901, 2017.
- [166] F. Becattini, J. Cleymans, and J. Strumpfer. Rapidity variation of thermal parameters at SPS and RHIC. *PoS*, CPOD07:012, 2007.
- [167] Krishna Rajagopal and Frank Wilczek. Static and dynamic critical phenomena at a second order QCD phase transition. *Nucl. Phys.*, B399:395–425, 1993.
- [168] Juergen Berges and Krishna Rajagopal. Color superconductivity and chiral symmetry restoration at nonzero baryon density and temperature. *Nucl. Phys.*, B538:215–232, 1999.
- [169] Adam Miklos Halasz, A. D. Jackson, R. E. Shrock, Misha A. Stephanov, and J. J. M. Verbaarschot. On the phase diagram of QCD. *Phys. Rev.*, D58:096007, 1998.
- [170] Boris Berdnikov and Krishna Rajagopal. Slowing out-of-equilibrium near the QCD critical point. *Phys. Rev.*, D61:105017, 2000.
- [171] J. Zinn-Justin. *Quantum Field Theory and Critical Phenomena*, Clarendon Press, 1996.
- [172] Paolo Parotto. Parametrized Equation of State for QCD from 3D Ising Model. *PoS*, CPOD2017:036, 2018.

- [173] Swagato Mukherjee, Raju Venugopalan, and Yi Yin. Real time evolution of non-Gaussian cumulants in the QCD critical regime. *Phys. Rev.*, C92(3):034912, 2015.
- [174] Swagato Mukherjee, Raju Venugopalan, and Yi Yin. Universal off-equilibrium scaling of critical cumulants in the QCD phase diagram. *Phys. Rev. Lett.*, 117(22):222301, 2016.
- [175] Miki Sakaida, Masayuki Asakawa, Hirotsugu Fujii, and Masakiyo Kitazawa. Dynamical evolution of critical fluctuations and its observation in heavy ion collisions. *Phys. Rev.*, C95(6):064905, 2017.
- [176] Adam Bzdak, Volker Koch, and Vladimir Skokov. Correlated stopping, proton clusters and higher order proton cumulants. *Eur. Phys. J.*, C77(5):288, 2017.
- [177] Adam Bzdak and Volker Koch. Rapidity dependence of proton cumulants and correlation functions. *Phys. Rev.*, C96(5):054905, 2017.
- [178] A. Adare et al. An Upgrade Proposal from the PHENIX Collaboration. 2015.
- [179] 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.