Universidade do Minho
Escola de Ciências

Filipa Cavaco Reis Peres

**New observables and techniques
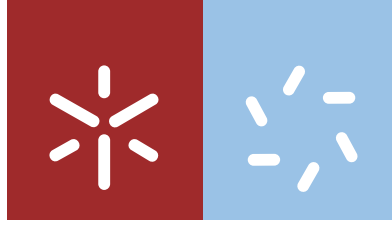for the study of jets in hadron collisions**

outubro de 2019

Universidade do Minho
Escola de Ciências

Filipa Cavaco Reis Peres

**New observables and techniques
for the study of jets in hadron collisions**

Dissertação de Mestrado
Mestrado em Física

Trabalho realizado sob orientação do
Professor Doutor José Guilherme Milhano
e do
Professor Doutor Nuno Filipe Castro

Direitos de autor e condições de utilização do trabalho por terceiros

Este é um trabalho académico que pode ser utilizado por terceiros desde que respeitadas as regras e boas práticas internacionalmente aceites, no que concerne aos direitos de autor e direitos conexos.

Assim, o presente trabalho pode ser utilizado nos termos previstos na licença abaixo indicada. Caso o utilizador necessite de permissão para poder fazer um uso do trabalho em condições não previstas no licenciamento indicado, deverá contactar o autor, através do RepositóriUM da Universidade do Minho.

Acknowledgements

I have been through a lot of things in my life, but finishing my master's degree proved to be one of the hardest. This is not because the courses were of unreachable or insurmountable difficulty but because reconciling academic studies with a full-time job requires extreme organisation, dedication, compromise and sacrifice. There is this quote from a book that has stuck to me as an absolute truth from the moment I read it: "You could not arrive at excellence by the average of these people. Excellence was an individual quest, not a group effort." Indeed, this quote resonated in me and I work and live by it everyday: the individual has to strive for the extraordinary, to have the desire to succeed.

However, as I worked on finishing my master's degree, I have come to regard this quote with more caution. I still think that it is true that only through the individual's desire and ambition can excellence be achieved. Nevertheless, the individuals that we see distinguished for their extraordinary achievements are never alone in their quests. They may be alone in the work itself, but that is not all that matters. In order for these individuals to be able to put in the effort and dedication they need to succeed, other people are there to help them, behind the scenes, walking in the shadows of their success, giving them support and content in helping them succeed often without recognition. This is why I believe the aforementioned quote must be regarded with caution. Yes, the individual must aspire to do his/her best, but there is inevitably a group behind, giving support.

During the full timeframe of my academic endeavours, there were a lot of people supporting, advising and encouraging me. I could not stand to leave those people go unrecognised, and this is one chapter that will bring me particular relief and contentment to write.

First of all, I should start by thanking both of my advisors: Doctor José Guilherme Milhano and Doctor Nuno Filipe Castro. None of the work presented in this master dissertation would have been possible were it not for their guidance, effort, dedication and commitment. They were always supportive, encouraging and inspiring. I can only wish that they have enjoyed working with me as much as I have enjoyed working with them and that we can keep collaborating in the future.

Secondly, I would like to thank Liliana Apolinário for generating the event samples used to perform the studies in this dissertation; and to Miguel Romão and Tiago Vale for their suggestions and assistance in the machine learning part of this work.

On a more personal level, I would like to thank my friends for understanding my detachment and seclusion over the past year. They have been nothing but supportive and considerate and my gratitude is beyond their imagination. I am thankful to all of them but there are some who deserve special mention. To Danya, thank you for being my best friend and the one I can always count on for unconditional loyalty and complete honesty (even if brutal). To Filipe for being in a similar situation to my own and thus understanding exactly the nature of the burden of finishing an academic degree while juggling a full-time job. And (last but not least) to Osvaldo for lightening my nervousness and teaching me "the serenity to embrace that which I cannot change", but also for his support and encouragement; all of these things helped me more than I

can express into words.

I also have a lot to thank to my family. I would like to start with my grandparents. They are far away and I do not get to see them as much as I would like to. I know they think that I am distant and, perhaps, even a bit indifferent to them, but that could not be further from the truth. In reality, I often think about them and how much I owe them. They have done a lot for me over the years, being present in some of my most difficult times, and helping me through most of the difficulties regarding my health. I know that they will never be able to read this text and understand how much they mean to me, but they do. And that had to be written regardless.

I also would like to thank my brother and my father. They are amongst some of the most knowledgeable people I know, with their knowledge spanning a lot of different contexts. It is refreshing to always be able to count on them for fruitful discussions and help on multiple topics.

Finally, there is one individual that should have a full chapter of her own, but I hope I can summarise all I have to say in one paragraph. Though without my advisors none of this work would have been possible, without my mother I am absolutely sure that it also would not have been. I owe her everything, there is nothing that I have accomplished while finishing my master's degree that is mine alone; there is a piece of her in everything, because she walks in the shadows of my successes and I know that I can count on her no matter what. In the tempest she is my lighthouse, my safe haven and my compass, all in one. I can only hope that I find a way of giving back to her as much as she has given me. It is a hard thing to strive for but after delivering this dissertation she can be sure that I will keep working every day to make her happy and proud, and to reciprocate all that she has ever done for me.

Statement of integrity

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the University of Minho.

Novas observáveis e técnicas para o estudo de jatos em colisões de hadrões

Resumo

Jatos consistem em conjuntos colimados de hadrões que resultam da ramificação de partões ener-géticos (quarks e gluões) e sua subsequente hadronização. Por um lado, os jatos providenciam-nos formas exatas de testar a teoria da interação forte. Por outro, são também ferramentas valiosas para etiquetar processos específicos (por exemplo, o decaimento do bosão de Higgs para dois quarks b, medido recentemente) onde desvios relativamente ao Modelo Padrão podem, ou não, ser encontrados.

Para além disso, os jatos são também ferramentas essenciais no contexto de colisões de iões pesados, onde o objetivo principal passa por caracterizar um novo estado da matéria - o plasma de quarks e gluões - que existia no universo primordial (breves instantes após o Big Bang) e no qual quarks e gluões interagem fortemente, sem estarem confinados em hadrões. Este plasma é recriado nestas colisões (PbPb no LHC) e tem um tempo de vida muito curto ($\sim 10\,\mathrm{fm}/c$), pelo que não pode ser estudado com recurso a sondas externas. Partões energéticos produzidos como resultado da colisão e os jatos a que eles dão origem são modificados por interações com o plasma. Isto é extraordinariamente diferente do que acontece noutras colisões onde não é formado um outro estado da matéria (por exemplo, colisões pp). O estudo destas modificações permite-nos extrair informações sobre as propriedades do plasma de quarks e gluões. Infelizmente, algumas das modificações são subtis e difíceis de detetar (jatos que atravessaram o plasma partilham muitas propriedades com jatos que não o fizeram).

Tradicionalmente, procuram-se as melhores observáveis que permitam distinguir entre jatos atenuados (os que atravessaram o plasma) e jatos que se desenvolveram no vácuo (i.e., na aus-ência do meio denso e quente), através da análise de cada observável individualmente. Muito recentemente, no contexto de colisões pp, descobriu-se que técnicas que aprendizagem automática aplicadas a planos de Lund são muito poderosas na distinção entre jatos iniciados por quarks e aqueles que são iniciados por gluões. O plano de Lund consiste num mapa bi-dimensional no qual cada ponto corresponde a uma ramificação que ocorreu durante processo de desenvolvimento de um jato.

O objetivo desta dissertação é migrar e adaptar estas técnicas para o contexto da discrimin-ação entre jatos que interagiram com o plasma de quarks e gluões e aqueles que não o fizeram.

Palavras-chave: aprendizagem automática, atenuação de jatos, jatos, plano cinemático de Lund, plasma de quarks e gluões.

New observables and techniques for the study of jets in hadron collisions

Abstract

Jets are collimated bunches of hadrons that result from the multiple QCD branching of energetic partons (quarks or gluons) and their subsequent hadronisation. On the one hand, jets provide an accurate testing ground for the theory of strong interaction. On the other hand, they are invaluable tools for tagging specific processes (for example, the recently measured decay of the Higgs boson to two b-quarks) where deviations from the Standard Model may, or may not, be found.

Further, jets are essential tools in the context of heavy-ion collisions where the main goal is to characterise a new state of matter - the quark-gluon plasma (QGP) - that existed in the primordial universe (shortly after the Big Bang) and in which quarks and gluons are strongly interacting without being confined into hadrons. This plasma is recreated in such heavy-ion collisions (PbPb at the LHC) and lives for a very short time ($\sim 10\,\mathrm{fm}/c$), such that it cannot be studied using external probes. Energetic partons are produced as part of the overall collision and the jets they give rise to are modified by interaction with the QGP. This is strikingly different from what happens in other collisions where no other state of matter is created (e.g. pp collisions). The study of these modifications allows us to extract information about the properties of the QGP. Unfortunately, some of the modifications are subtle and difficult to detect (jets that traversed the QGP share many properties with jets that did not).

Traditionally, one has searched for the best observables that distinguish between quenched and non-quenched jets by exhaustively analysing one observable at a time. Very recently, it has been found, in the context of pp collisions, that machine learning techniques applied to Lund planes are very powerful in distinguishing between jets initiated by quarks and those initiated by gluons. The Lund plane is a bi-dimensional splitting map where each point corresponds to one branching that took place during the development of a jet.

The aim of this dissertation is to migrate and adapt these techniques to the context of discriminating between jets which have interacted with the QGP and those which have not.

Keywords: jet quenching, jets, kinematical Lund plane, machine learning, quark-gluon plasma.

Contents

List of abbreviations

AI	Artificial intelligence
ANN(s)	Artificial neural network(s)
AUC	Area under the curve
C/A	Cambridge/Aachen
CERN	European Organisation for Nuclear Research
CNN(s)	Convolutional neural network(s)
DIS	Deep inelastic scattering
DL	Deep learning
DLA	Double logarithmic approximation
FSR	Final-state (QCD) radiation
GRU	Gated recurrent unit
IRC	Infrared and collinear
ISR	Initial-state (QCD) radiation
LHC	Large Hadron Collider
LSTM	Long short-term memory
MC	Monte Carlo
ML	Machine learning
NN(s)	Neural network(s)
QCD	Quantum chromodynamics
QED	Quantum electrodynamics
QFT(s)	Quantum field theory(ies)
QGP	Quark-gluon plasma
RHIC	Relativistic Heavy Ion Collider
RNN(s)	Recurrent neural network(s)
ROC	Receiver operating characteristic
SLAC	Stanford Linear Accelerator Centre
SPEAR	Stanford Positron Electron Accelerating Ring
SPS	Super Proton Synchrotron

List of Figures

Chapter 1

Introduction

Particle physics research is at the heart of our understanding of the laws of nature and provides our current picture of the universe as made of basic constituents, known as elementary particles, and governed by four fundamental forces: the electromagnetic, weak, strong and gravitational forces. The Standard Model of particle physics represents one of the triumphs of modern physics as it successfully describes how the elementary particles and three of the fundamental forces are related to each other,[1] thereby allowing to explain almost all experimental results as well as precisely predict a wide variety of phenomena [1]. As such, having seen comprehensive validation from experimental data, the Standard Modern has established itself as the cornerstone of modern particle physics.

The first fundamental interaction to be well understood was the electromagnetic interaction. The physical theory that describes it is known as electrodynamics, and its classical formulation is due to Maxwell well over one hundred years ago (refs. [2–4], dating from 1861 to 1873). The quantum theory of electrodynamics (QED) was perfected by the works of Tomonaga [5], Schwinger [6, 7] and Feynman [8–10] in the late 1940s and provides a covariant formulation for the interaction between charged particles (mediated by photons) which is finite at any order in perturbation theory. QED was thus the first prototype of a quantum field theory (QFT).

Some years later, came into existence the important idea that particle interactions are dictated by local gauge symmetries. Although this was not the view point of the physicists that developed QED, this theory can actually be regarded as a special class of QFT: an Abelian gauge theory based upon the $U(1)$ symmetry group. This beautiful idea that symmetry dictates the form of the interaction was proposed in 1954 by Yang and Mills [11], in an attempt to explain the strong interaction. They constructed a theory based not on the simple one-dimensional Abelian $U(1)$ group, but on the three-dimensional non-Abelian $SU(2)$, hoping that this would become the theory of strong interaction. It took a few time until the Yang-Mills idea started to be applied to the weak interaction. However, Glashow [12], in 1961, and independently Salam and Ward [13], in 1964, proposed a $SU(2) \times U(1)$ group structure to explain this interaction, thus unifying the electromagnetic and weak forces. Later, in 1967, Weinberg [14] and Salam [15] incorporated the Higgs mechanism into the electroweak interaction thus conferring it its modern form. Experimental validation for the electroweak model came in 1973 when neutral currents were discovered at CERN [16] and in 1978 at Stanford Linear Accelerator Center (SLAC) with the discovery of parity violation in those currents [17].

The final piece of the Standard Model is quantum chromodynamics (QCD), the theory of strong interaction. In 1964 the quark model was proposed independently by Gell-Mann [18] and Zweig [19] and presented the idea that hadrons are made of quarks and anti-quarks, allowing

---

[1]A satisfactory quantum formulation of gravity is yet to be provided, thus preventing this fundamental force from being unified with the other three and included into the Standard Model.

to explain the growing diversity of hadrons observed in collider experiments. The quark model got experimental support from an experiment done at SLAC in 1968 under the leadership of Friedman, Kendall, and Taylor [20]. However, no further advancements had been made in describing the strong interaction between these "partons". It was the success of the electroweak theory that restored interest in Yang–Mills theory and propelled further investigation into the strong interaction. In 1973, through many different contributions, QCD acquired its modern form. By that time, the concept that quarks must come in three different colours (red, green and blue) was already well established (refs. [21–23]), and so it was natural to take the gauge symmetry of the strong interaction as a $SU(3)$ group acting on the (three-valued) colour quantum number of the quarks. In that same year, Gross and Wilczek [24] and Politzer [25] independently discovered that a wide class of non-Abelian gauge theories (amongst which the $SU(3)$ non-Abelian gauge theory) have the remarkable property of asymptotic freedom. Furthermore, several independent authors (see refs. [26–28]) proposed that gluons, the mediators of strong interaction, are massless and cannot be detected for the same reason we do not see quarks: coloured particles can never be isolated. Indeed, the property that gluons carry colour charge means that they interact with each other, conducing to the asymptotic freedom of QCD.

In conclusion, the present formulation of the Standard Model would not have been possible were it not for the conjunction of beautiful and brilliant theoretical ideas and models with experimental information and validation provided by collider events. It was thus during the latter half of the twentieth-century and through the works and contributions of many first-class physicists that remarkable developments were made, allowing our current understanding of the universe. For an in-depth description of the making of the Standard Model, reference [29], used as a source of information for writing this overview, is recommended.

In this chapter, each section is built up on top of the knowledge provided by the preceding ones such that the concepts and results presented gradually grow in complexity and relevance to the work developed. The idea behind this structure is to provide any interested reader with the knowledge required for the understanding of this master dissertation and its results. As such, firstly some core fundamental concepts and properties of QCD are introduced, notably its Lagrangian (section 1.1) and its structure, i.e., its logarithmic divergences (section 1.2). Following that, an overview of Monte Carlo (MC) event generators in the context of pp collisions (where no dense QCD matter is formed) is presented in section 1.3, and with the knowledge acquired at that point a description of how jets arise in the events is laid out. The quark-gluon plasma (QGP) and its importance in the understanding of fundamental properties of nature, as well as its formation and evolution in the context of ultra-relativist heavy-ion collisions, are described in section 1.4. Finally, the chapter reaches its climax and is concluded with the introduction of the concept of formation time, a paramount piece in this work. The expression widely used for its calculation is derived and the connection to a proxy formation time using hadron collider observables is established.

In some of these sections, the discussions and descriptions are general overviews focusing on the key concepts required for the understanding of future sections and chapters, rather than in-depth, detailed descriptions. When that is the case, more comprehensive references are usually provided for the interested reader. At any rate, it is our understanding that a more exhaustive description of some points would only serve to increase the bulk of the present text, and would indeed defeat its purpose, which is to bring focus to the results obtained in the developed research.

1.1   Fundamentals of quantum chromodynamics

In the previous introduction we have been presented with the remarkable idea that particle interactions are dictated by local gauge symmetries. Owing to that, the QFTs of the three

fundamental forces described by the Standard Model are gauge theories in which the Lagrangian $\mathcal{L}$ is invariant under a certain Lie group of local transformations. For any study involving modern particle colliders, may it be experimental, phenomenological or theoretical, it is vital to understand some QCD fundamentals. However, before delving into such deep waters, it is advisable to start by having a look at QED, since this is a simpler theory from which we can build upon to understand the much richer theory that is QCD.[2]

QED is the theory of the electromagnetic interaction between electrically charged particles. As we have seen, although is was not originally formulated in the framework of gauge theories, it has since been established that QED is indeed a gauge theory based upon the one-dimensional Abelian $U(1)$ symmetry group. From this symmetry, the Lagrangian of QED can be deduced and is given by:

$$\mathcal{L}_{QED} = \bar{\psi}\left(i\gamma^\mu\partial_\mu - m - e_f\gamma^\mu A_\mu\right)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}\,, \tag{1.1}$$

where ψ represents a single (fermionic) spinor, μ and ν are Lorentz indices, γ^μ are the Pauli matrices, m is the mass of the considered fermion and e_f its electric charge,[3] A_μ is the gauge field and $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the gauge field tensor.[4] From imposing a $U(1)$ symmetry group we see that the gauge boson (i.e., the mediator of the electromagnetic interaction) has to be a massless particle with no electric charge. Indeed, the photon is the gauge boson of QED; it has zero mass and is electrically neutral.

QCD is the theory of strong interaction, i.e., the theory of quarks and gluons and the interactions between them. We can start by thinking of QCD in terms of its similarities and differences with respect to QED [31]. Just as the latter describes the interaction between particles which carry electric charge, the former describes the interaction between particles carrying the "QCD charge", known as colour charge (or simply colour). A major difference in this point is that while there is only one single kind of electric charge, colour comes in three varieties: red, green and blue. The massless gluons are the gauge bosons of QCD, thus playing a similar role to QED's photon. However, while the photon is electrically neutral, gluons are not colour neutral, particularly, they can be thought of as carrying both colour and anti-colour. There are eight different combinations of colour and anti-colour for gluons, meaning that there are eight different gluons. The fact that the gluons themselves carry colour charge means that they can interact directly with one another. This possibility is not available in QED since photons are electrically neutral. Field theories in which the field quanta may interact with each other are called non-Abelian. In QCD, the self-interaction of gluons leads to asymptotic freedom, i.e., the coupling α_s tends to zero for high momentum scales and blows up at small scales (a striking difference with respect to QED, in which the opposite behaviour is observed). In-between the two limits the evolution with scale is rather fast: the so-called running coupling effect.

The fact that there are three different colour charges makes the $SU(3)$ group the natural symmetry group of strong interaction. The fundamental representation of this group is a triplet, making it the natural representation for quarks. That is, quarks are represented by a field $\psi_{(a)}$, carrying a colour label $a = 1\ldots 3$. This quark field corresponds to the product of a vector in colour space with a spinor such that:

$$\psi_{(1)} = |R\rangle \otimes u(p)\,, \qquad \psi_{(2)} = |G\rangle \otimes u(p) \qquad \text{and} \qquad \psi_{(3)} = |B\rangle \otimes u(p)\,,$$

---

[2]The sign conventions used throughout are those by Peskin and Shroeder (ref. [30]), although the notation is slightly different in some places.

[3]The convention is such that $e_f = Q\,|e|$, with Q for example equal to -1, $+2/3$ and $-1/3$ respectively for the electron, up and down quarks.

[4]Note that while a sum over all flavours could have been written explicitly for completeness, we chose not to do so because the QED Lagrangian is often written for a single family.

where

$$|R\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \qquad |G\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix} \qquad \text{and} \qquad |B\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}.$$

The generators of $SU(3)$ are eight linearly independent traceless hermitian 3×3 matrices, traditionally denoted by λ^A, where the index A runs from 1 to 8 :

$$\lambda^1 = \begin{pmatrix} 0 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda^2 = \begin{pmatrix} 0 & -i & 0 \\ i & 0 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda^3 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \lambda^4 = \begin{pmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix},$$

$$\lambda^5 = \begin{pmatrix} 0 & 0 & -i \\ 0 & 0 & 0 \\ i & 0 & 0 \end{pmatrix}, \lambda^6 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{pmatrix}, \lambda^7 = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & -i \\ 0 & i & 0 \end{pmatrix}, \lambda^8 = \frac{1}{\sqrt{3}}\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -2 \end{pmatrix}.$$

These eight generator matrices represent the eight gluons of QCD, which are thus in the adjoint representation of $SU(3)$.

Since we say that $SU(3)$ is the symmetry group of strong interaction, this means that the Lagrangian of QCD must be invariant under local transformations of that group. From this we can obtain its form as:

$$\mathcal{L}_{QCD} = \sum_{flavours} \bar{\psi}_{(a)} \left(i\gamma^\mu \partial_\mu \delta_{ab} - m\delta_{ab} + g_s \gamma^\mu \mathcal{A}_\mu^A t_{ab}^A \right) \psi_{(b)} - \frac{1}{4}\mathcal{F}_{\mu\nu}^A \mathcal{F}^{A,\,\mu\nu}, \tag{1.2}$$

where g_s is the coupling constant which relates to the QCD coupling via $\alpha_s = \frac{g_s^2}{4\pi}$, $t^A \equiv \frac{1}{2}\lambda^A$, $\mathcal{A}_\mu^A$ are the gluon or gauge fields (analogous to A_μ in (1.1)) and $\mathcal{F}_{\mu\nu}^A$ are the gluon field tensors given by:

$$\mathcal{F}_{\mu\nu}^A = \partial_\mu \mathcal{A}_\nu^A - \partial_\nu \mathcal{A}_\mu^A + g_s f^{ABC} \mathcal{A}_\mu^B \mathcal{A}_\nu^C, \tag{1.3}$$

with f^{ABC} the structure constants of $SU(3)$, given by $\left[t^A, t^B\right] = if^{ABC}t^C$.

By comparing the QED and QCD Lagrangians in equations (1.1) and (1.2) respectively, we note that the major difference between the two is the presence of the term $g_s f^{ABC} \mathcal{A}_\mu^B \mathcal{A}_\nu^C$ with two gluon fields. This is the term that is responsible for the gluon self-interaction and, hence, for the asymptotic freedom of QCD.

To each Lagrangian there corresponds a set of Feynman rules, that is, a set of propagators and vertex factors. This section does not aim at deriving such entities from the two Lagrangians (1.1) and (1.2), but, since these rules will be necessary for the calculations performed in the next section, the results are presented in appendix A for consultation whenever required. With these rules identified, it is possible to calculate a desired observable using perturbation theory based upon Feynman diagrammatic techniques.

Perturbation theory relies on the idea of an order-by-order expansion in a small coupling α_s, such that a given observable is obtained by:

$$f = f_0 + f_1\alpha_s + f_2\alpha_s^2 + f_3\alpha_s^3 + ...,$$

where the coefficients f_i can be determined through Feynman diagrammatic techniques and it is generally sufficient to calculate the first order terms with the understanding that the remaining terms will be small.

From the above formulation, it is straightforward to understand that the applicability of perturbation theory rests on the assumption that the coupling is small. Therefore, as a result of asymptotic freedom, the calculation of a given observable using perturbative QCD is only valid at high momentum (or energy) scales, where the coupling α_s satisfies that condition.
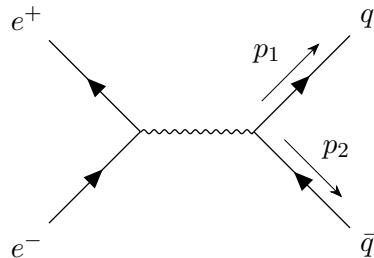
1.2   The soft and collinear limits of QCD

In the previous section, we have introduced the QCD Lagrangian, from which a set of Feynman rules can be identified and employed to calculate observables using perturbation theory based on Feynman diagrams. We have also seen that in order for perturbation theory to be applicable, we need to limit ourselves to the behaviour of QCD at high momentum scales.

As we shall see, jets take up a particularly prominent role in this master dissertation. **Jets** are collimated sprays of hadrons that result from the fragmentation of energetic partons into a cascade of lower-energy quarks and gluons that eventually hadronise. Jet substructure is, at its most fundamental, the study of QCD at high momentum scales and in the near-soft and near-collinear limits [32]. "Soft" means that the emitted gluon carries very little energy compared to the parent parton that emitted it. That is, if we call the fraction of energy carried out by the gluon ζ then we say that $\zeta \ll 1$. "Collinear" means that the gluon is emitted very close in angle to another parton in the event. The soft and collinear approximation is important enough in particle physics in general (and jet physics in particular) that we will repeatedly employ it in our discussions and calculations.
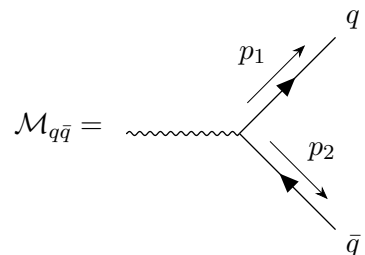
In this section we intend to derive the probability of emission of a soft and collinear gluon from a pair of quark and anti-quark $q\bar{q}$. The calculations presented follow the steps done in ref. [31].

To simplify the calculations let us consider the process of e^+e^- annihilation into hadrons, specifically into a pair $q\bar{q}$:

$$\text{(1.4)}$$

Considering this process has the advantage that only the final state involves QCD. We are interested in the hadronic side of (1.4), so let us write the amplitude of the QED process of a virtual photon decaying to a $q\bar{q}$ pair: $\gamma^* \to q\bar{q}$,

$$\mathcal{M}_{q\bar{q}} = \quad \cdots \quad = \bar{u}(p_1)(-ie_q\gamma^\mu)\delta_{ab}v(p_2)\,, \tag{1.5}$$

where $\bar{u}(p_1)$ and $v(p_2)$ are the spinors for the outgoing quark and anti-quark respectively, with four-momenta p_1 and p_2, $(-ie_q\gamma^\mu)$ is the QED vertex (as given in equation (A.1)) and e_q is the electric charge of the quark. The Dirac-δ imposes that the quark and anti-quark carry the same colour, as required in the process under consideration (QED process/vertex).
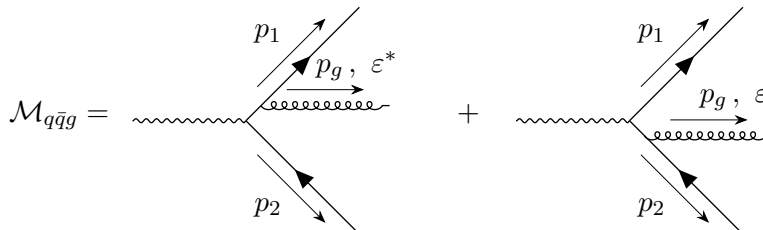
By doing this we have discarded the $e^+e^- \to \gamma^*$ interaction vertex and also the photon propagator, since those are inconsequential for the calculations we are interested in. Usually we need the squared amplitude of matrix element (1.5), summed over all possible (colour) final

states:

$$
\begin{aligned}
|\mathcal{M}_{q\bar{q}}|^2 &= |\bar{u}(p_1)(-ie_q\gamma^\mu)\delta_{ab}v(p_2)|^2 \\
&= \bar{u}(p_1)(-ie_q\gamma^\mu)\delta_{ab}v(p_2)\bar{v}(p_2)ie_q\gamma^\nu\delta_{ba}u(p_1) \\
&= N_C\,|\bar{u}(p_1)ie_q\gamma^\mu v(p_2)|^2\,,
\end{aligned}
\tag{1.6}
$$

where $N_C = 3$ is the number of colour charges in QCD. Note that, when taking the squared modulus of the matrix element, there are two implicit sums over the colour indices a and b, i.e., there are implicit sums over all possible final states ($\delta_{ab}\delta_{ba} = \delta_{aa} = N_C$).

Let us now incorporate into this process the emission of a soft gluon with momentum p_g and polarisation vector ε^*, such that the Feynman diagram of the problem is

$$
\begin{aligned}
&= \bar{u}(p_1)\varepsilon^*_\alpha(p_g)(ig_s\gamma^\alpha t^A_{ab})\left[i\frac{\not{p}_1 + \not{p}_g + m}{(p_1 + p_g)^2 - m^2 + i\epsilon}\right](-ie_q\gamma^\mu\delta_{bc})v(p_2) \\
&+ \bar{u}(p_1)(-ie_q\gamma^\mu\delta_{ab})\left[-i\frac{\not{p}_2 + \not{p}_g + m}{(p_2 + p_g)^2 - m^2 + i\epsilon}\right](ig_s\gamma^\alpha t^A_{bc})\varepsilon^*_\alpha(p_g)v(p_2)\,.
\end{aligned}
$$

To write the equation above, the Feynman rules presented in appendix A were used. We note that there are two different terms that contribute for the amplitude of this process: the first for the emission of the gluon from the quark and the second for the emission of the gluon from the anti-quark. In equation (1.7) m stands for the mass of the quark (and anti-quark) and $\not{p} = \gamma^\mu p_\mu$. In high energy physics, it is common to take particles to be massless ($m = 0$). In that approximation, we can re-write equation (1.7) as

$$
\mathcal{M}_{q\bar{q}g} = i\bar{u}(p_1)\not{\varepsilon}^* g_s t^A_{ac}\frac{\not{p}_1 + \not{p}_g}{(p_1 + p_g)^2}e_q\gamma^\mu v(p_2) - i\bar{u}(p_1)e_q\gamma^\mu\frac{\not{p}_2 + \not{p}_g}{(p_2 + p_g)^2}g_s t^A_{ac}\not{\varepsilon}^* v(p_2)\,.
\tag{1.8}
$$

At this stage it is important to pause and consider several details of the matrix element (1.8) and of the process represented in (1.7). Since we are under the assumption of massless particles, the Dirac equation for the quark yields $\bar{u}(p_1)\not{p}_1 = 0$ while for the anti-quark we can write $\not{p}_2 v(p_2) = 0$. Moreover, we assume that all final-state particles (q, $\bar{q}$ and g) are on-shell, such that: $p_1^2 = p_2^2 = p_g^2 = 0$. Additionally, in each term of the diagram there is a virtual particle, represented by the small internal line (i.e., by a propagator) between the two vertices. Each one of these particles is an off-shell particle with a given virtuality $(p_1 + p_g)^2, (p_2 + p_g)^2 \neq 0$. Finally, we have already stated that we are considering that the emitted gluon is soft, meaning that $p_{\mu,g} \ll p_{\mu,i}$, with p_i the four-momentum of either the final-state quark or anti-quark.

Knowing also that the γ-matrices obey the anti-commutation relation $\{\gamma^\mu, \gamma^\nu\} = \gamma^\mu\gamma^\nu + \gamma^\nu\gamma^\mu = 2g^{\mu\nu}$, where $g^{\mu\nu}$ is the Minkowski metric $(+,-,-,-)$, yields that

$$
\begin{aligned}
\not{A}\not{B} &= \gamma^\mu A_\mu\gamma^\nu B_\nu = \gamma^\mu\gamma^\nu A_\mu B_\nu \\
&= 2g^{\mu\nu}A_\mu B_\nu - \gamma^\nu\gamma^\mu A_\mu B_\nu \\
&= 2A.B - \not{B}\not{A}\,.
\end{aligned}
$$

With all of these considerations in mind, we can re-write each of the terms in equation (1.8), starting with the first:

$$i\bar{u}(p_1)\not{\varepsilon}^* g_s t^A_{ac} \frac{\not{p}_1 + \not{p}_g}{(p_1 + p_g)^2} e_q \gamma^\mu v(p_2) = i\bar{u}(p_1) g_s t^A_{ac} \frac{2\varepsilon^*.(p_1 + p_g) - (\not{p}_1 + \not{p}_g)\not{\varepsilon}^*}{(p_1 + p_g)^2} e_q \gamma^\mu v(p_2)$$

$$\simeq i\bar{u}(p_1) g_s t^A_{ac} \frac{2\varepsilon^*.(p_1 + p_g) - \not{p}_g\not{\varepsilon}^*}{2p_1.p_g} e_q \gamma^\mu v(p_2)$$

$$\simeq \bar{u}(p_1) i e_q \gamma^\mu t^A_{ac} v(p_2) g_s \frac{p_1.\varepsilon^*}{p_1.p_g} \, .$$

The second term of equation (1.8) can be handled in a similar way to the first term and we get:

$$i\bar{u}(p_1) e_q \gamma^\mu \frac{\not{p}_2 + \not{p}_g}{(p_2 + p_g)^2} g_s t^A_{ac} \not{\varepsilon}^* v(p_2) = i\bar{u}(p_1) e_q \gamma^\mu \frac{2(p_2 + p_g).\varepsilon^* - \not{\varepsilon}^*(\not{p}_2 + \not{p}_g)}{(p_2 + p_g)^2} g_s t^A_{ac} v(p_2)$$

$$\simeq i\bar{u}(p_1) e_q \gamma^\mu \frac{2(p_2 + p_g).\varepsilon^* - \not{\varepsilon}^*\not{p}_g}{2p_2.p_g} g_s t^A_{ac} v(p_2)$$

$$\simeq \bar{u}(p_1) i e_q \gamma^\mu t^A_{ac} v(p_2) g_s \frac{p_2.\varepsilon^*}{p_2.p_g} \, .$$

As such, we can finally write the matrix element (1.7) as:

$$\mathcal{M}_{q\bar{q}g} \simeq \bar{u}(p_1) i e_q \gamma^\mu t^A_{ac} v(p_2) g_s \left(\frac{p_1.\varepsilon^*}{p_1.p_g} - \frac{p_2.\varepsilon^*}{p_2.p_g} \right) \, . \tag{1.9}$$

Just as before we need the squared amplitude of the matrix element (1.9), summed over all possible colour and polarisation final states:

$$|\mathcal{M}_{q\bar{q}g}|^2 \simeq \sum_{\text{pol.}} \left| \bar{u}(p_1) i e_q \gamma^\mu t^A_{ac} v(p_2) g_s \left(\frac{p_1.\varepsilon^*}{p_1.p_g} - \frac{p_2.\varepsilon^*}{p_2.p_g} \right) \right|^2 \, ,$$

where we chose to write the sum over all polarisation states explicitly while leaving the sums over repeated final-state colour indices (a, b and A) implicit. Proceeding with the calculations, we can factorise this probability into a product of the sum over all colour final states and the sum over all polarisation final states:

$$|\mathcal{M}_{q\bar{q}g}|^2 \simeq g_s^2 \left| \bar{u}(p_1) i e_q \gamma^\mu t^A_{ac} v(p_2) \right|^2 \sum_{\text{pol.}} \left| \frac{p_1.\varepsilon^*}{p_1.p_g} - \frac{p_2.\varepsilon^*}{p_2.p_g} \right|^2 \, . \tag{1.10}$$

The first factor in the equation above resembles the squared amplitude of the matrix element for $q\bar{q}$ production (equation (1.6)) except that with a generator matrix t^A instead of a Dirac-δ. To perform the sum over all colour states and thus evaluate this first factor it is useful to keep in mind the Fierz identity:

$$t^A_{ab} t^A_{cd} = \frac{1}{2}\delta_{bc}\delta_{ad} - \frac{1}{2N_C}\delta_{ab}\delta_{cd} \, .$$

Using this result and noting that the squared modulus yields $\left| t^A_{ac} \right|^2 = t^A_{ac} t^A_{ca}$, we obtain

$$t^A_{ac} t^A_{ca} = \frac{1}{2} \left(\delta_{cc}\delta_{aa} - \frac{1}{N_C}\delta_{ac}\delta_{ca} \right)$$

$$= \frac{1}{2} \left(N_C^2 - 1 \right) = C_F N_C \, ,$$

which means we get an extra C_F multiplicative factor with respect to $|\mathcal{M}_{q\bar{q}}|^2$. C_F is the Casimir of the fundamental representation and arises in the previous equation because

$$C_F = \frac{N_C^2 - 1}{2N_C} = \frac{4}{3} \, .$$

For the sum over polarisations states it is important to know that,

$$\sum_{\text{pol.}} \varepsilon_\mu \varepsilon_\nu^* = -g_{\mu\nu} \, .$$

Using this result we can develop the sum over all polarisation states such that

$$\sum_{\text{pol.}} \left| \frac{p_1.\varepsilon^*}{p_1.p_g} - \frac{p_2.\varepsilon^*}{p_2.p_g} \right|^2 = \sum_{\text{pol.}} \left(\frac{p_1^\mu \varepsilon_\mu}{p_1.p_g} - \frac{p_2^\mu \varepsilon_\mu}{p_2.p_g} \right) \left(\frac{p_1^\nu \varepsilon_\nu^*}{p_1.p_g} - \frac{p_2^\nu \varepsilon_\nu^*}{p_2.p_g} \right)$$

$$= \left(\frac{p_1^\mu}{p_1.p_g} - \frac{p_2^\mu}{p_2.p_g} \right) \left(\frac{p_1^\nu}{p_1.p_g} - \frac{p_2^\nu}{p_2.p_g} \right) \sum_{\text{pol.}} \varepsilon_\mu \varepsilon_\nu^*$$

$$= \frac{2p_1.p_2}{(p_1.p_g)(p_2.p_g)} \, .$$

Finally, returning to equation (1.10) yields,

$$|\mathcal{M}_{q\bar{q}g}|^2 \simeq g_s^2 C_F |\mathcal{M}_{q\bar{q}}|^2 \frac{2p_1.p_2}{(p_1.p_g)(p_2.p_g)} \, .$$

One main point of this result is that, in the soft limit, the $|\mathcal{M}_{q\bar{q}g}|^2$ squared matrix element factorises into two terms: the $|\mathcal{M}_{q\bar{q}}|^2$ matrix element and a piece with a rather simple dependence on the gluon momentum.

The next piece we need is the phase space of the $q\bar{q}g$ system, $d\Omega_{q\bar{q}g}$, which also factorises into the product of the phase space of the $q\bar{q}$ pair production, $d\Omega_{q\bar{q}}$, with the phase space for gluon emission $d\Omega_{g\,\text{emission}}$: $d\Omega_{q\bar{q}g} = d\Omega_{q\bar{q}} d\Omega_{g\,\text{emission}}$. The Lorentz invariant phase space for the soft-gluon emission corresponds to the phase space available for the gluon, which can be written as

$$d\Omega_{g\,\text{emission}} = \frac{d^3\mathbf{p}_g}{2(2\pi)^3 E_g} \, .$$

where E_g is the energy of the gluon [31]. As such, the phase space of the $q\bar{q}g$ system is given by

$$d\Omega_{q\bar{q}g} = d\Omega_{q\bar{q}} \frac{d^3\mathbf{p}_g}{2(2\pi)^3 E_g}$$

Finally, we can thus say that the full differential cross section for the $q\bar{q}g$ production, $d\sigma_{q\bar{q}g}$, is given by the squared amplitude of the $q\bar{q}$ production matrix element and corresponding phase space, $|\mathcal{M}_{q\bar{q}}|^2 \, d\Omega_{q\bar{q}}$, multiplied by the probability of soft-gluon emission, $d\mathcal{S}$, such that

$$d\sigma_{q\bar{q}g} = |\mathcal{M}_{q\bar{q}g}|^2 \, d\Omega_{q\bar{q}g} = |\mathcal{M}_{q\bar{q}}|^2 \, g_s^2 C_F \frac{2p_1.p_2}{(p_1.p_g)(p_2.p_g)} d\Omega_{q\bar{q}} \frac{d^3\mathbf{p}_g}{2(2\pi)^3 E_g} \equiv |\mathcal{M}_{q\bar{q}}|^2 \, d\Omega_{q\bar{q}} d\mathcal{S} \, .$$

We can thus identify the probability of emission of a soft gluon as,

$$
\begin{aligned}
d\mathcal{S} &= g_s^2 C_F \frac{2p_1.p_2}{(p_1.p_g)(p_2.p_g)} \frac{d^3\mathbf{p}_g}{2(2\pi)^3 E_g} \\
&= \alpha_s C_F \frac{2p_1.p_2}{(p_1.p_g)(p_2.p_g)} \frac{d^3\mathbf{p}_g}{(2\pi)^2 E_g} \\
&= \alpha_s C_F \frac{2p_1.p_2}{(p_1.p_g)(p_2.p_g)} \frac{\mathrm{p}_g^2 d\mathrm{p}_g \, \sin\theta d\theta \, d\phi}{(2\pi)^2 E_g} \\
&= \frac{\alpha_s C_F}{\pi} \frac{p_1.p_2}{(p_1.p_g)(p_2.p_g)} E_g dE_g \, \sin\theta d\theta \frac{d\phi}{2\pi} \,,
\end{aligned}
$$

where θ and ϕ are, respectively, the polar and azimuthal angles of the gluon with respect to the other final-state child resulting from the splitting (either the quark or the anti-quark, in the first and second terms, respectively). We can proceed, with a little bit more algebra, to re-arrange this result in a different fashion. We place ourselves in the rest frame of the $q\bar{q}$ pair such that the quark and anti-quark are back-to-back: $\mathbf{p}_1 = -\mathbf{p}_2$ and $E_1 = \mathrm{p}_1 = \mathrm{p}_2 = E_2$. Under this condition the probability of soft-gluon emission can be written as:

$$
\begin{aligned}
d\mathcal{S} &= \frac{\alpha_s C_F}{\pi} \frac{E_1^2 + \mathrm{p}_1^2}{(E_1 E_g - \mathbf{p}_1 \cdot \mathbf{p}_g)(E_1 E_g + \mathbf{p}_1 \cdot \mathbf{p}_g)} E_g dE_g \, \sin\theta d\theta \frac{d\phi}{2\pi} \\
&= \frac{\alpha_s C_F}{\pi} \frac{2E_1^2}{E_1^2 E_g^2 - (\mathbf{p}_1 \cdot \mathbf{p}_g)^2} E_g dE_g \, \sin\theta d\theta \frac{d\phi}{2\pi} \\
&= \frac{2\alpha_s C_F}{\pi} \frac{E_g dE_g \, \sin\theta d\theta}{E_g^2(1 - \cos^2\theta)} \frac{d\phi}{2\pi} \\
&= \frac{2\alpha_s C_F}{\pi} \frac{dE_g}{E_g} \frac{d\theta}{\sin\theta} \frac{d\phi}{2\pi} \,.
\end{aligned}
$$

It will prove useful to write the probability of soft-gluon emission as a function of a dimensionless quantity: the energy fraction, ζ, carried out by the gluon with respect to that of its parent (emitter),

$$
\zeta = \frac{E_g}{E_{parent}} = \frac{E_g}{E_g + E_1} \,, \qquad \text{(for the first term)}.
$$

Using this dimensionless quantity and further considering that we are within the collinear limit where $\sin(\theta) \simeq \theta$, the probability of emission of a soft and collinear gluon can be written in the final form:

$$
d\mathcal{S} = \frac{2\alpha_s C_F}{\pi} \frac{d\zeta}{\zeta} \frac{d\theta}{\theta} \frac{d\phi}{2\pi} = \frac{2\alpha_s C_F}{\pi} d\left(\ln\zeta\right) d\left(\ln\theta\right) \frac{d\phi}{2\pi} \,. \tag{1.11}
$$

This result presents two different non-integrable logarithmic divergences: the first, when $\zeta \to 0$, is called the soft divergence while the second, when $\theta \to 0$, is the collinear divergence, corresponding to the emission of the gluon with a direction collinear with that of the quark (or anti-quark). For this reason, the soft and collinear limit is usually referred to as the "double-logarithmic regime" or "double-logarithmic approximation" (DLA).

We will return to this expression in chapter 3, where the kinematical Lund planes are introduced and we can use equation (1.11) as motivation for the choice of axes.

It is very important to note that, although derived in the specific context of $e^+e^- \to q\bar{q}$ production, these soft and collinear (logarithmic) divergences are a very general property of QCD and appear whenever a soft gluon is emitted from a quark or from another gluon, regardless of the process originating the original parent [31].

1.3   A brief overview of Monte Carlo event generators

General-purpose Monte Carlo (MC) [33] event generators, such as HERWIG [34], PYTHIA [35] and SHERPA [36], are essential tools in modern high energy physics without which many physics analyses would not be feasible. They provide a detailed description of the final state of high energy particle collisions,[5] such as those at the Tevatron or the LHC, and can therefore be employed in a multitude of ways. Namely, they can be used (i) to estimate signal and background processes, thus allowing to devise data analysis strategies for separating them from one another; (ii) to plan and design new experiments or estimate the feasibility of a new physics study; (iii) to optimise the detector and its triggers to a specific investigation; (iv) as a vehicle to disseminate ideas from theorists to experimentalists, etc. [38].

In this section we present a general overview of the various MC event generator components through which an event is build up. From amongst these, the understanding of the final-state parton shower is of particular importance to our work. Once the description of the flow of an event in a generator has been presented, we take this opportunity to properly introduce the notion of jets and how they arise in hadron collider events.

One important note is due at this stage: the discussions in this section focus only on the high energy physics of pp collisions. Heavy-ion physics is not addressed in here since it involves rather different aspects, which are left for the two final sections (1.4 and 1.5) of this chapter and for section 2.1 of chapter 2.

1.3.1   The structure of an event

Understanding and describing the final states of high energy particle collisions is a formidable theoretical task. Such final states typically involve the production of hundreds of particles, potentially from every species predicted by the Standard Model (and possibly even beyond it) and with momenta spanning several different orders of magnitude. Therefore, the relevant matrix elements are often too laborious to compute beyond a first few orders in perturbation theory and, additionally, in the case of QCD processes, they involve the unsolved, non-perturbative problem of colour confinement and resulting hadronisation [39].

In spite of these difficulties, modern MC event generators are very successful in describing such final states. That success stems from the large amount of work that has been invested in recent years in the development of techniques to help tackle this extremely challenging problem. From among these, one crucial tool is factorisation. It allows us to separate the treatment of many different processes into three main regimes, according to the momentum (or energy) transfer scales involved. At the core lie the hard interactions between pairs of partons, described by matrix elements which can be computed perturbatively given that the momentum transfer scales are high; this is the hard regime or hard scattering. Following that, a perturbative evolution process is responsible for the production of a multitude of final-state partons, through a cascade of emissions that successively lower the momentum scale down to scales around $1\,\mathrm{GeV}$, where QCD is strongly coupled and perturbation theory breaks down. Finally, at low momentum transfer scales, the final-state coloured partons interact non-perturbatively turning into colourless hadrons; this is the soft regime.

Each one of these regimes is suitable for computer simulation using MC methods and the factorisation means that each element can be improved individually if new advances have been made. In the end, when all of these pieces have been put together, we have a MC event generator capable of simulating a wide range of interesting physical processes.

---

[5]It should be noted that the final state produced by MC event generators does not include detector effects, which can be incorporated into this picture resorting to detector simulation programs such as GEANT4 [37]. Since this is outside the scope of this work no further mentions will be made to this particular point.

The steps by which event generators build up the structure of a hadron collision are presented in figure 1.1. We will now outline what happens individually at each of these steps. Reference [39] is recommended for a detailed description, and has served to inspire the general overview presented here.
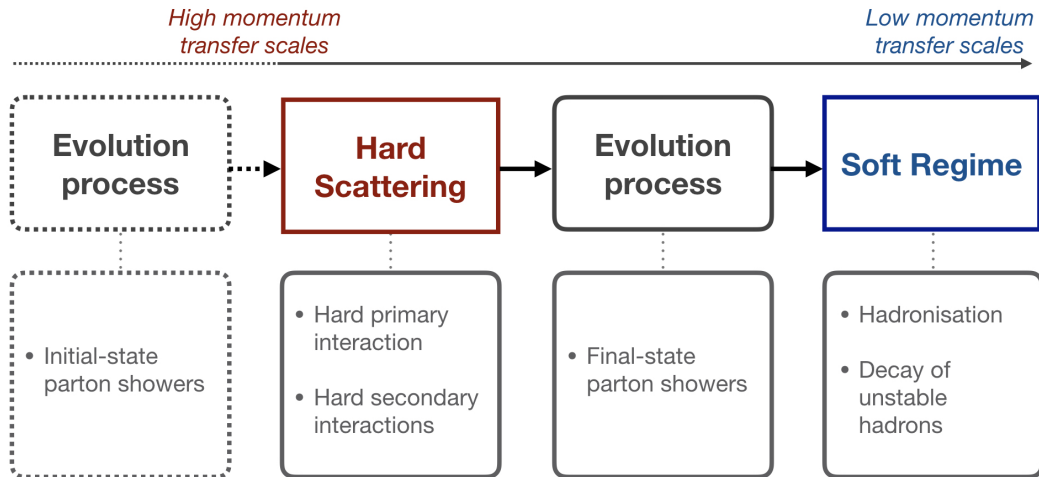
Figure 1.1: Schematic picture of the flow of hard collisions in a MC event generator. The three different regimes are represented, together with a list of the main steps by which the generators build up the structure of an event involving a hard process.

Hard scattering

In a hadron collider, two hadronic beams are accelerated and made to collide with each other. Each hadron is made up of quarks and gluons such that, as a result of a collision of two hadrons (one from each beam), there are interactions between incoming partons. Most of these interactions are soft, i.e., they involve low momentum transfer scales. However, a small fraction of the incoming partons may interact through the transfer of large amounts of momentum, a hard scattering which produces a (relatively) small number of outgoing quarks, leptons and/or gauge bosons at large angles. The processes of interest in hadron collisions generally stem from these rarer hard interactions and, as such, a hard scattering lies at the core of events in MC generators.

Although we are generally interested in a particular kind of event, the variety of $2 \to n$ (hard) processes that can occur as a result of the collision of the incoming partons is substantial. As such, instead of generating typical events and waiting for one of them to be of the desired type, MC event generators allow us to generate events around specific hard subprocesses of interest according to their matrix elements and phase space. Since the momentum transfer scales involved are high, these matrix elements can be computed using perturbation theory.

Thus, from the collision of two hadrons, a primary hard interaction between two partons produces a flow of outgoing primary particles. However, since hadrons are complex bound-states of strongly interacting partons, within the collision of these same two hadrons, further parton pairs (other than the first) may interact via the transfer of large amounts of momentum. This means that besides the primary hard subprocess, one or more secondary hard interactions may occur and produce additional outgoing quarks, leptons and/or gauge bosons of their own. These additional secondary interactions may add significant contributions to any observable besides

that of the primary subprocess. This part of the event structure is known as "underlying event", and should not be confused with pile-up which stems from the interaction of more that one hadron pair at each beam crossing.

Evolution processes

As we have seen, as a result of the hard primary interaction outgoing primary partons can be produced. These outgoing partons can afterwards radiate gluons, which can, in turn, radiate further gluons themselves or produce quark-anti-quark pairs. Thus, the primary outgoing partons initiate a cascade of successive splittings of the form $q \rightarrow qg$, $g \rightarrow gg$ and/or $g \rightarrow q\bar{q}$, responsible for the extraordinarily high multiplicity of final-state partons. This process is simulated by the so called "parton shower algorithm" which consists in a stepwise Markov process. This algorithm probabilistically evaluates whether or not a new parton - emitted by a pre-existing one - should be added to the final state. A key feature of these parton showers is that the successive emissions are explicitly ordered according to a momentum-tranfer-like ordering variable, which is successively lowered at each splitting: $t_1 > t_2 > t_3 > ...$ (t being the ordering variable). New partons are added to the final state as long as the ordering variable associated with their emission is larger than a predefined threshold; when that condition is no longer met the showering is stopped.

Typically, the ordering variable should be a variable in which QCD is naturally ordered, i.e., it should be related to QCD's logarithmic divergences. In vacuum, QCD showering structure is thought to be ordered in such a way that large-angled, hard emissions happen first, followed by gradually softer and smaller-angled emissions. Even so, the choice is by no means obvious or unique and, indeed, different generators have implemented different ordering variables. For instance, Herwig has been formulated with an angular ordered shower [40], while the original Pythia implementation had an ordering in virtuality and the modern Pythia8 generator is ordered in transverse momentum [41].

At this point it is important to point out that as a result of the hard primary interaction some of the outgoing particles will be electrically charged. As such, although we have centred our discussion on QCD showering (since that is the most relevant to the work presented here), subdominant QED radiation should also be appropriately included in this picture.

We have just described the perturbative evolution process through which the primary outgoing partons give rise to a multitude of final-state partons. In the same way, the outgoing partons that originate from secondary interactions also produce successive QCD emissions adding more partons to the final state (underlying event). This branching is formulated exactly as the one for the primary interaction: as a perturbative downwards evolution in a particular ordering variable. As such, together, the primary and secondary processes originate the final-state radiation (FSR) and the parton shower algorithms behind this evolution process are thus called final-state parton showers.

One subtle detail has not yet been mentioned. The accelerated partons within the incoming hadrons can themselves radiate prior to the hard interactions. This gives rise to initial-state radiation (ISR) and the corresponding showering algorithms are known as initial-state parton showers. Their theoretical formulation and implementation is slightly more complicated than final-state parton showers because the complex nature of the incoming hadrons and the confinement of the incoming partons therein come into play. But, in any case, the important thing to take note of is that the final state will have contributions from both FSR and ISR. The latter should be seen as a contamination to what we actually want to observe.

Soft regime

As an event is evolved downwards, the ordering variable is successively lowered until the momentum (energy or virtuality) transfer scale involved reaches a region - around $1\,\mathrm{GeV}$ - where QCD becomes strong interacting ($\alpha_s(1\,\mathrm{GeV}) \sim 1$ [42]) and perturbation theory is no longer applicable. At this stage, the showers must be terminated and we are left with a final state made up of a large multiplicity of partons. However, as we know, in real-life events hadrons (and not partons) are observed.

The transition from the final-state partons into hadrons is known as hadronisation and cannot be calculated from first principles within any calculation technique currently at our disposal. Therefore, models have been introduced to describe the non-perturbative transition of the final-state coloured quarks and gluons into colourless hadrons. Since these models are not derived directly from QCD but rather inspired by it, they are design with free parameters whose values are adjusted until the model's results agree with those of a given experimental dataset (usually one from LEP where the initial state does not involve QCD). It is then said that the model has been "tuned to data". Once that has been done, the model is approximately universal, meaning that it gives a good description of the hadronisation step regardless of the core process that originated the input final-state coloured partons.

Finally, many of the hadrons output by hadronisation are heavy and unstable, never reaching the detector in an actual high energy physics experiment. Therefore, once again models are used to describe the decay of such hadrons into lighter ones, that are long-lived enough to be considered stable at the time scales of the detector.

1.3.2   The arising of jets

It has already been mentioned in section 1.2 that "jets are collimated sprays of hadrons that result from the fragmentation of energetic partons into a cascade of lower-energy quarks and gluons that eventually hadronise". At the time it was presented, this definition might have struck the reader as a bit contrived; however, in light of the parton shower description above it should now be much clearer.

As we have seen, in an event there is an evolution process in which successive emissions of softer partons gradually lower the momentum scale until the non-perturbative region is reached. The structure of QCD, i.e., its soft and collinear divergences (recall eq. (1.11)), means that those emissions are dominated by soft gluon emission and/or collinear parton splittings. As a result, the event's final-state partons will mostly be found in bundles around the original directions of the hard outgoing partons produced in the hard interaction. This distribution remains essentially unchanged by hadronisation and, hence, jets arise as bunches of collimated (i.e. collinear) hadrons grouped together roughly around the direction determined by a hard parent parton resulting from the collision.

1.4   The quark-gluon plasma and heavy-ion collisions

The current generally accepted picture for the origin of the universe was first proposed in 1972 by Geroges Lemaître and is now widely known as "the Big Bang theory" [43]. According to it, the universe has expanded explosively from an extremely hot and dense initial state, continuing to expand and cool to this day.

In its early stages, the primordial universe was filled with an extremely hot and dense form of matter made up of the most elementary building blocks found in nature, but essentially dominated by quarks and gluons [44,45]. In those first moments the temperature was so extreme that these partons were only weakly coupled (due to the asymptotic freedom of QCD). However,

between approximately 10^{-12} to 10^{-6} seconds after the Big Bang (i.e., after the start of the expansion) the universe reached temperatures that, though still incredibly high, were low enough for the quarks and gluons to start interacting strongly (rather than weakly). This astonishingly hot and dense soup of strongly coupled quarks and gluons was the first complex form of matter to be formed in nature and is known as the quark-gluon plasma (QGP). The evolution continued and as lower temperatures were reached the partons started to bind together to form nucleons (protons and neutrons). The transition from the QGP phase to the phase where the universe is made up of protons and neutrons happened via a gradual (smooth) crossover, leaving no "imprint" in the primordial universe. As such, there is no cosmological observation that can be made to directly "observe" this primordial QGP and/or its properties. We know only of its existence through rigorous and complex analytical (lattice) QCD calculations which show the QGP is a phase of the QCD Lagrangian in a specific temperature range.

A deeper study of the QGP offers a very enticing research field with vast potential. Since this form of matter was present in the primordial universe, it offers a tangible connection with those early stages of its evolution and a way of probing its beginnings. Moreover, since QCD governs the formation and evolution of the QGP, its study can provide insights into the properties of diverse multi-scale QCD processes. However, as we have seen, while the QGP was present in the early stages of the universe, it quickly expanded and evolved into a state made up of protons and neutrons, leaving no trace of its existence that could be observed, probed or analysed today. Fortunately, the QGP can be recreated in a collision of ultra-relativistic heavy ions, a fact which has propelled a lot of interest in the field.

In an ultra-relativistic heavy-ion collision, the maximum energy density (or temperature) happens at the moment of impact between two incoming (highly contracted) nuclei. The collision is responsible for multiple interactions (both hard and soft) of several different nucleon pairs while, at the same time, many more nucleons proceed along their route, suffering no such interactions. As a result of the soft interactions, a droplet of highly dense matter is formed. $1\,\mathrm{fm}/c$ after its creation, the energy density of this droplet is about $12\,\mathrm{GeV}/\mathrm{fm}^3$, a value almost 25 times larger than that typically found inside hadrons ($\sim 0.5\,\mathrm{GeV}/\mathrm{fm}^3$). This shows that this droplet cannot be made of a highly packed collection of distinct individual hadrons [46]. On the other hand, at this energy density, the temperature is not quite high enough for the droplet to be a nearly free gas of weakly interacting partons. Hence, it must be made of strongly coupled quarks and gluons that form a collective medium, the QGP, which behaves like a nearly-perfect (i.e., low specific viscosity) liquid [47] that evolves according to the laws of relativistic hydrodynamics.

As time passes, hydrodynamic pressure drives the tiny droplets of QGP to gradually expand in all directions. As this liquid expands and cools, the energy density decreases until it reaches values of the order of those found inside individual hadrons, which happens at a temperature of approximately $T \sim 155\,\mathrm{MeV}$. At that point, the quarks and gluons are bound together, getting confined into hadrons that scatter off one another a few times, before eventually streaming away freely to eventually reach the detectors.
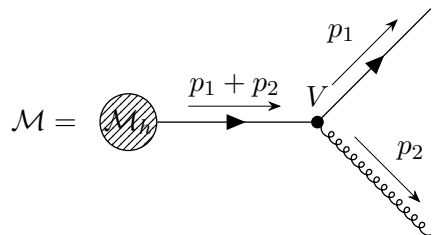
While some of the interactions are responsible for the production of the QGP, a (simultaneous) hard scattering may also produce hard outgoing partons - a process which is rare but of particular interest. These partons will subsequently shower and originate jets, much as we have seen in the previous section for pp collisions. However, now the presence of the QGP adds an additional degree of complexity. As particles from the hard scattering evolve and radiate, they traverse the region where the QGP is being produced and evolved. In particular, in events involving hard partons, the jets they give rise to are modified, with respect to those stemming from pp collisions. These modifications and the phenomena originating them are collectively known as "jet quenching" and contain a wealth of invaluable information that allows us to make grounded inferences about the nature and properties of the QGP. Although some pieces of information

have already been extracted much is still to be learnt and this is a blossoming area of research. Jet quenching is discussed in section 2.1 of chapter 2.

1.5   The concept of formation time

As we have seen, QCD is naturally formulated in momentum space because it is in that space that (i) its structure - i.e., its logarithm divergences - is manifest and (ii) factorisation is formulated. We have also discussed that parton showers produce a cascade of successive emissions or splittings which are explicitly ordered according to a specific ordering variable. All of these discussions were presented in the context of collisions where no dense QCD matter is formed. However, as explained in the previous section, in a collision of heavy ions the QGP is created and some of the partons produced as a result of the hard scattering may interact with it. It is therefore clear that, in this context, a knowledge of which partons interact with the QGP and at which spacetime position is desirable.[6] In order for that to be possible, we need to give a spacetime meaning to the parton shower which was formulated in momentum space (section 1.3).

The relation between momentum space and spacetime is encoded in the uncertainty principle. Recalling that the parton shower evolves the event between two scales (the hard scattering scale and the hadronisation scale) via successive $1 \to 2$ splittings of virtual partons, a natural time scale is the lifetime of a virtual parent to split into two children, with some specific kinematics. Let us consider one such splitting in which the virtual parent particle has four-momentum $q = p_1 + p_2$ and splits into two child particles with four-momenta p_1 and p_2 such that the corresponding Feynman diagram is

$$\mathcal{M} = \quad \raisebox{-1em}{[diagram]} \quad . \tag{1.12}$$

We depict not only the splitting vertex V but also another process with an associated amplitude $\mathcal{M}_h$, responsible for creating the parent particle (either the hard scattering itself or a previous splitting). This is explicitly presented to highlight the fact that the parent is a virtual particle (internal line). From the uncertainty principle, the lifetime of a virtual state is given by the inverse of the square-root of its virtuality q^2,

$$t \sim \frac{1}{\sqrt{q^2}} \,.$$

Since the particles which we are dealing with are boosted particles travelling at very high velocities, that is, they are relativistic particles, the lifetime of the parent particle in a Lorentz frame in which it carries an energy E is dilated by the Lorentz boost factor $\gamma = E/\sqrt{q^2}$. Thus, the lifetime of the virtual parent particle is given by the ratio between its energy and its virtuality:

$$t_F \sim \frac{2E}{q^2} \,. \tag{1.13}$$

This lifetime can be regarded as the characteristic time for the splitting to occur and is usually known as the splitting's "formation time". The factor of two appearing in equation (1.13) is a widely spread and accepted convention for the definition of this quantity [48, 49].

---

[6]Note that it is the QGP itself that provides a physical spacetime reference frame for the formulation of this problem.

Although the development of the shower requires the children to be virtual (off-shell) particles themselves, for the computation of the formation time we will assume that they are on-shell. This is a good approximation because the off-shellness of the child particles contributes to the lifetime of their parent at high orders (i.e., the correction is parametrically small). As such, considering once again massless particles, we can write: $p_1^2 = p_2^2 = 0$. We now assume that the softer child takes an energy fraction ζ while, by energy conservation, the other takes the remaining fraction $(1 - \zeta)$ such that:

- one child particle has $p_1 = (E_1, \mathbf{p}_1) = ((1 - \zeta)E, \mathbf{p}_1)$;

- the other child has $p_2 = (E_2, \mathbf{p}_2) = (\zeta E, \mathbf{p}_2)$.

This allows us to write the virtuality of the parent particle as:

$$\begin{aligned}
q^2 = (p_1 + p_2)^2 &= p_1^2 + 2p_1.p_2 + p_2^2 \\
&= 2p_1.p_2 \\
&= 2\left(E_1 E_2 - \mathbf{p}_1 \cdot \mathbf{p}_2\right) \\
&= 2E_1 E_2 \left(1 - \cos\theta\right) \\
&= 2(1 - \zeta)\zeta E^2 \left(1 - \cos\theta\right) ,
\end{aligned}$$

where θ is the angle between the two child particles. As such, the formation time can be written as

$$t_F \sim \frac{1}{(1 - \zeta)\zeta E \left(1 - \cos\theta\right)} . \tag{1.14}$$

In this expression the kinematics is exact, except for the aforementioned fact that the two children were assumed to be on-shell, when they are actually off-shell particles just like their parent.

In the collinear limit we can write: $1 - \cos\theta \simeq 1 - [1 - \theta^2/2] = \theta^2/2$. Additionally, considering also the soft limit, in which $\zeta \ll 1$, we get $\zeta(1 - \zeta) \simeq \zeta$. Thus, we can write the formation time in its commonly used form [48, 49]

$$t_F \sim \frac{2}{\zeta E \theta^2} . \tag{1.15}$$

One point that is crucial to keep in mind is that the formation time as defined by equation (1.15) does not define the absolute time at which a given splitting has occurred (or a given pair of particles was produced). Instead, it describes the time interval between successive splittings, meaning it is an **incremental time**. This is clear from the way the derivation was performed. Hence, when we calculate this quantity using equation (1.15) together with the internal kinematical information provided by the MC event generator, i.e., from the four-momenta of the particles within the jet, the formation time that we obtain is an incremental quantity.

The formation time of a splitting is responsible for imbuing a spacetime meaning into the cascade of successive splittings from the parton shower, thus allowing to determine whether a certain splitting may or may not have interacted with the QGP and been affected by it. To make things clearer we can consider the simple picture presented in [49], where we take the medium to have a spatial extent L. In that case, the i-th splitting along a certain branch happens at a time that is given by the sum of the (incremental) formation times of all the splittings that happened up to that point in that branch,

$$t_i = t_{F,1} + t_{F,2} + ... + t_{F,i} .$$

If this time is larger than the medium length, that is, if $t_i > L$ the splitting happens outside of the QGP and is therefore not directly affected by its presence. We say that it is a "pure vacuum

splitting". On the other hand, if $t_i < L$ then the splitting happens inside the medium and can potentially be affected by it. This kind of splittings are called "in-medium splittings". It is important to note that not all splittings that happen inside the medium are necessarily affected by it; other effects are at play which determine whether or not a certain splitting may or may not be resolved by the QGP. For a simple and brief discussion on this, the reader is referred to reference [49].

Equation (1.14) presents the formation time written in terms of the energy of the parent particle, the energy fraction carried out by the softer child and the angle between the two children. Instead we could have written this expression using hadron collider variables such as the transverse momentum of the particles with respect to the beam (p_t),[7] their rapidity (y) and their azimuthal angle (ϕ). For that, it helps to note that the four-momentum of a particle i is given by

$$p_i = (E_i,\, p_{i,x},\, p_{i,y},\, p_{i,z}) = p_{i,t} \left(\cosh y_i,\, \cos \phi_i,\, \sin \phi_i,\, \sinh y_i \right).$$

Using this result we can write the virtuality of the parent particle in these variables as

$$\begin{aligned} q^2 &= 2p_1.p_2 \\ &= 2p_{1,t}p_{2,t} \left(\cosh y_1 \cosh y_2 - \sinh y_1 \sinh y_2 - \cos \phi_1 \cos \phi_2 - \sin \phi_1 \sin \phi_2 \right) \\ &= 2p_{1,t}p_{2,t} \left(\cosh (y_1 - y_2) - \cos(\phi_1 - \phi_2) \right) \end{aligned}$$

and its energy as

$$E = E_1 + E_2 = p_{1,t} \cosh y_1 + p_{2,t} \cosh y_2 \,.$$

Thus, making once again use of equation (1.13), the formation time yields

$$t_F \sim \frac{p_{1,t} \cosh y_1 + p_{2,t} \cosh y_2}{p_{1,t}p_{2,t} \left(\cosh (y_1 - y_2) - \cos(\phi_1 - \phi_2) \right)} \,. \tag{1.16}$$

This expression is fully equivalent to that of equation (1.14), even if with a rather more complicated look to it.

We can now focus on relevant limits of equation (1.16), in particular, let us consider that the differences in rapidity and azimuthal angle are sufficiently small so that:

$$\cosh(y_1 - y_2) \simeq 1 + \frac{(y_1 - y_2)^2}{2} \,,$$

$$\cosh(y_1) \simeq \cosh(y_2) \equiv \cosh(y) \,,$$

and,

$$\cos(\phi_1 - \phi_2) \simeq 1 - \frac{(\phi_1 - \phi_2)^2}{2} \,.$$

In this limit, equation (1.16) approximates to,

$$t'_F \sim \frac{2 \left(p_{1,t} + p_{2,t} \right) \cosh y}{p_{1,t}p_{2,t} \Delta R_{12}^2} \,,$$

where $\Delta R_{12}^2 = (y_1 - y_2)^2 + (\phi_1 - \phi_2)^2$. Further considering that, not only the difference in rapidities is small but that the rapidity itself is small $(y \ll 1)$ we have $\cosh (y) \simeq 1$. Moreover, if we take, once more, the soft approximation $(\zeta \ll 1)$, then we can assume that one of the children carries almost all of the momentum such that $p_{1,t} \gg p_{2,t}$ and the formation time can be written as

$$t'_F \sim \frac{2}{p_{2,t} \Delta R_{12}^2} \,. \tag{1.17}$$

---

[7]As is usual in literature, the coordinate system is defined such that the z direction is oriented along the beam.

For convenience, we introduce a dimensionless quantity z defined as the ratio between the transverse momentum carried out by the softer child and the transverse momentum q_t of its parent such that

$$z = \frac{p_{2,t}}{q_t} \,.$$

In analogy to the energy fraction ζ, we will call z the "momentum fraction" (though the term is not fully accurate in this case). Finally, we can write

$$t'_F \sim \frac{2}{z q_t \Delta R_{12}^2} \,. \tag{1.18}$$

The importance of equations (1.17) and (1.18) will be apparent in section 2.2 of chapter 2, when discussing the different reclustering algorithms used throughout this work. For now, it suffices to note that t'_F as given by these two equations is a proxy for the "real" splitting formation time t_F presented exactly in equations (1.14) and (1.16), and approximately in equation (1.15).

Chapter 2

Introduction to jet physics

As previously explained, jets consist of bunches of collimated hadrons which result from the fragmentation of energetic quarks and gluons. They are ubiquitous objects used in the analysis of the experimental measurements, MC parton-shower results, or even in analytical partonic calculations [50]. Jets are not fundamental objects defined by QCD theory and despite tending to be particularly prominent features when one looks at an event display of a collider experiment, that alone is not sufficient to guarantee their broad application throughout. In order for that to be possible, a **jet algorithm** which establishes which particles are projected (or grouped) into a jet in a reproducible and well defined manner is necessary. Besides that, a **recombination scheme**, defining how the four-momenta of the particles should be combined when they are clustered together to form a single object, is also required. The jet algorithm and its parameters together with the recombination scheme form the so-called **jet definition** [50].

In this chapter we will start by making an exposition of the role of jets in heavy-ion physics research. Afterwards, different jet algorithms will be presented and discussed, with particular emphasis on the generalised-k_t family of algorithms, which encompasses all the algorithms used in this master dissertation. Finally, the chapter ends with a brief note on recombination schemes.

2.1 Jets and their role in heavy-ion research - jet quenching

In 1975, an experiment performed at the Stanford Positron Electron Accelerating Ring (SPEAR) at SLAC recorded the observation that the final-state hadrons were not isotropically distributed but instead accumulated around the axis defined by the quark and anti-quark produced as a result of the collision ($e^+e^- \to q\bar{q}$) [51]. This behaviour had been theoretically predicted by Bjorken and Brodsky five years before [52] and constituted the first experimental evidence of jets. The first record of these objects in the context of hadron collisions dates from the 1982 Paris Conference, when the UA2 collaboration announced the observation of very clear hadron jets resulting from $p\bar{p}$ collisions at the Super Proton Synchrotron (SPS) at CERN [53, 54]. The quest for the QGP via the collision of relativistic heavy ions started at the SPS approximately one decade later, in the early 1990s [55–57]. As we have seen in the previous chapter, in such collisions the presence of the QGP induces modifications in the properties of jets relative to the baseline given by simpler collisions where no dense QCD matter is formed. The totality of these modifications together with the phenomena that give rise to them are collectively known as "jet quenching". This term is often used to refer to the suppression of high-p_t hadrons (which necessarily belong to jets). Historically, the first observation of jet quenching as hadron suppression happened in 2001 at the Relativistic Heavy Ion Collider (RHIC) [58] and preceded our ability to reliably reconstruct full jets within the large and fluctuating background of heavy-ion collisions. Indeed, it was not until 2003 that first candidate reconstructed jets were obtained at the STAR and PHENIX experiments
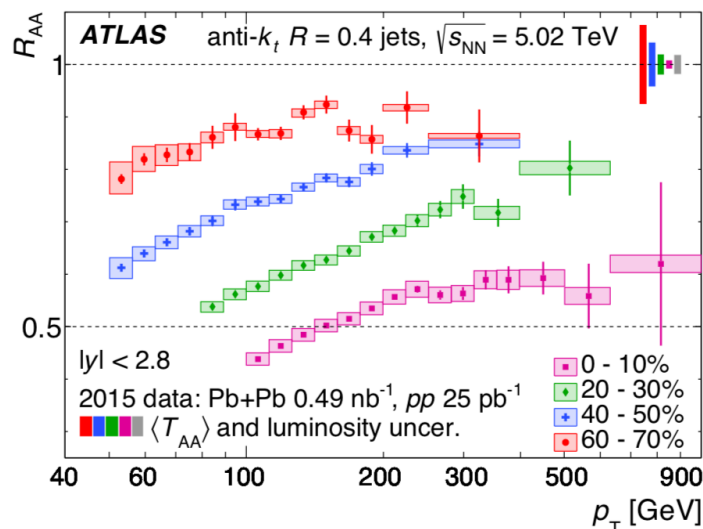
Figure 2.1: The nuclear modification factor R_{AA} as a function of jet transverse momentum p_t, for jets with $|y| < 2.8$ and three different centralities bins [64].

(again) at the RHIC [47]. A full and systematic reconstruction of jets with statistical reliability was only possible in 2010 at the LHC, at which time the phenomenon of jet quenching with fully reconstructed jets was directly observed and confirmed by the ALICE, ATLAS and CMS collaborations [59–62]. This could only be achieved thanks to the much higher energies reached at the LHC ($\sqrt{s_{NN}} = 2.76\,\mathrm{TeV}$, compared with the $\sqrt{s_{NN}} = 200\,\mathrm{GeV}$ at the RHIC) which gave access to an entirely new kinematic range and novel information, and allowed for new and better characterisation of the QCD processes in play.

The showering process through which jets are developed in vacuum is understood and described by perturbative QCD with remarkable (and enviable) precision. However, in the presence of a medium, the scenario is considerably more complex and, despite the efforts of the past two decades, much work is still to be done. As we have seen in section 1.3 of chapter 1, as a result of the hard scattering it is possible (even if rare) that hard partons are produced. These subsequently radiate into a cascade of successive emissions which, in the case of simpler collisions, develops in vacuum. In relativistic heavy-ion collisions, the QGP is simultaneously produced and evolved which means that as the jets develop, i.e., as successive partons are emitted, they may traverse this hot and dense QCD medium and interact with it through many different mechanisms. One of the major challenges in the theoretical description of jet quenching ensues from the diversity of effects that come into play and which are often intertwined, making it extremely hard to account for all of those effects simultaneously in the same analytical calculation.

It has been experimentally well instituted that jets lose energy as they traverse the hot and dense QGP or, in other words, in heavy-ion collisions there is a suppression of high-p_t jets [63]. This suppression is quantified by an observable known as nuclear modification factor, $R_{AA}(p_t)$,

$$R_{AA}(\{p_t, y, b, ...\}) = \frac{dN_{AA}/d\Omega}{\langle T_{AA} \rangle\, dN_{pp}/d\Omega}\,,$$

where N_{AA} and N_{pp} are, respectively, the inclusive yields in nucleus-nucleus and pp collisions, $\langle T_{AA} \rangle$ is the average nuclear overlap function for a given centrality class and $d\Omega$ is simply the relevant phase space. Figure 2.1, taken from reference [64], shows a plot of this observable as a function of the jet's p_t and the suppression ($R_{AA} < 1$) is clear, particularly for central collisions. Parton energy loss is a dominant effect that contributes to the modification of many

jet observables in heavy-ion collisions, when comparing to the results for those observables in the pp baseline.

This sort of observation is clear in showing that there are conspicuous differences between jets stemming from relativistic heavy-ion collisions and those from the pp baseline. These differences are attributed to scattering processes between the partons from the final-state shower and those within the deconfined (yet strongly interacting) medium.[1] One major effect that has been firmly established is that these interactions stimulate soft gluon emissions off the partons from the shower, thus (obviously) causing modifications to the jets' substructure. While a fraction of these emitted soft gluons ends up inside the jet, the remaining land outside of it. The specific proportions will vary depending on the jet definition in place, but a relevant fraction should end up outside of the jet cone, because the induced soft gluon emissions tend to happen at large angles (unlike the collinear nature of the vacuum branching process) [62, 65]. These emissions were early interpreted as a mechanism of radiative energy loss happening in soft modes at large angles [66–70]. This is observed experimentally as a collection of final-state jets with suppressed p_t, together with an abundance of soft quanta ($p_t \lesssim 2\,\mathrm{GeV}$) at large angles.

Another effect that produces modifications in the jet substructure originates in the fact that when shower partons scatter off medium partons they may transfer to them part of their energy and momentum. These partons from the medium may end up as a part of the hadrons that make up the final-state jet hence promoting further alterations to the jets' substructure which cannot be isolated from the contributions from the induce soft gluon radiation. Furthermore, these recoil effects may also push shower partons that would otherwise end up inside the jet cone to be outside of it, thus instituting yet another modification to the substructure of the jets.

Our current phenomenological and theoretical understanding of the dynamics behind the medium-induced modifications was aided by the interpretation of experimental data. The key idea is that the presence of the QGP undoubtedly introduces several different intertwined effects, one of which is the modification of the final-state parton showers. One of the many challenges under work is precisely to reconcile the collinear and the medium-induced radiations in those showers. For that, an improvement of our theoretical understanding of the phenomena at play is paramount. The better our grasp of the effects behind the measured experimental results, the more we can infer about the thermodynamical and transport features of the QGP and, therefore, about the dynamic and collective properties of QCD that govern its creation and evolution [66].

One central question remains to be explicitly clarified: why jets? The answer to this is relatively simple: the QGP created in heavy-ion collisions is very shortly lived ($\sim 10\,\mathrm{fm}/c$) before falling apart into a midst of hadrons. Because of this small lifetime, it is impossible to probe this medium with external probes. Jets are intrinsically hard, multi-scale probes and since they are produced and developed concurrently with the creation and evolution of the QGP itself they are natural probes for it. The importance of jets in the context of collisions of relativistic heavy ions is therefore clear: they provide an unique way of probing the deconfined matter prevalent in the primordial universe. By constructing and calculating different observables (we have seen the example of the R_{AA}), we can bring sensitivity to specific aspects of jet substructure and get a deeper insight into the concrete modifications that are taking place and, hence, a better understanding of the dynamics of the interactions that lead to them. To that end, the use of kinematical Lund planes, which will be introduced in chapter 3, proves particularly effective since they provide an immediate and visual way of identifying which regions of the phase space of emissions suffer enhancements or suppressions in the presence of the medium.

---

[1]Actually, partons from the initial-state shower can also interact with the QGP and the jets they give rise to are modified by it as well. This will be important for further discussions, though in this section we will always refer to the final-state shower.

2.2   Jet algorithms

As explained before, jet algorithms consist in a set of rules which determine which particles are grouped together to form a jet. There are two main categories of jet algorithms: (i) cone algorithms, which constitute a "top-down" approach to jet finding and (ii) sequential recombination algorithms, which can be regarded as "bottom-up" algorithms. A complete and thorough description of the extensive variety of jet algorithms and their history is not provided here, since that is outside the scope of this section (and this work in general). Instead, a very brief discussion on cone algorithms is presented, only for the sake of completeness since they were not used in the analyses made. The main focus is on the generalised-k_t family of the sequential recombination algorithms, from which we take all the algorithms used in our work. It should be pointed out that more sequential recombination algorithms other than those from the generalised-k_t family exist, but their description does not fit into the purpose of this section and, as such, we chose to leave them out. An excellent overview of the main jet algorithms can be found in [50], which was taken as basis for the text presented in this and the following section.

2.2.1   Cone algorithms

As mentioned above, cone algorithms take a "top-down" approach to finding the jets, relying on the idea that QCD showering and hadronisation leave the bulk features of an event's energy flow, specifically the energy flow into a cone, unchanged. The main idea behind this type of algorithms is that all final-state particles within a cone with a given angular reach, determined by a jet radius R, are said to form a jet if the resulting transverse momentum is above a predetermined threshold.

For a better understanding of how these algorithms work, the idea behind iterative cone (IC) algorithms will be described very briefly. In such algorithms, the momentum of a seed particle i sets the initial direction of the cone; we then combine the momenta of all particles j within a circle ("cone") of radius R around i, in azimuthal angle ϕ and rapidity y, i.e., we consider that all particles j such that

$$\Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2 < R^2 \, ,$$

are constituents of the jet and combine their momenta; y_i and ϕ_i are, respectively, the rapidity and azimuthal angle of particle i. The direction of the resulting momenta combination is then used as a new seed direction, and we iterate this procedure until the direction of the resulting cone is stable. In the end, only those jets that are above a certain transverse momentum threshold are considered. Besides IC algorithms, there are also fixed cone (FC) algorithms in which the hardest particle is set as the seed, defining the direction of the cone and there are no further iterations.

Despite producing jets with clean circular boundaries, which is an experimentally favoured feature, one major drawback of most cone algorithms (both IC and FC) is that they are either infrared or collinear (IRC) unsafe [50]. IRC safety means that an observable or algorithm is insensitive to infinitesimally soft or exactly collinear emissions [71]. This is an extremely important feature of any algorithm or observable because only those experimental results obtained under the condition of IRC safety can be compared to analytical predictions from perturbative QCD.

2.2.2   Sequential recombination algorithms: the generalised-k_t family

The second class of algorithms is that of sequential recombination algorithms which, as was previously described, take a "bottom-up" approach to jet finding. These algorithms are rather simple to state: they take the complete set of final-state particles and repeatedly recombine the pair which is closest into a new pseudo-particle, following some specific distance measure. The difference between the various algorithms of this class thus lies in the distance measure used.

Because of the way these algorithms work, sequentially recombining pairs of particles, they go beyond simply finding the jets but also provide a hierarchical clustering sequence for each jet of the event. These clustering sequences are the consequence of the algorithms' attempt to assign a meaningful QCD substructure to the jet and are closely connected to QCD's probabilistic picture of parton showering.

   The fact that these algorithms provide us with this clustering sequence is a key advantage over cone algorithms. In particular, in the context of this master dissertation, the hierarchical ordering of the particle pairings provided by sequential recombination algorithms is indispensable. By accessing this hierarchical tree of successive pairings (or conversely, of successive branchings or splittings) and undoing each pairing we can retrieve the kinematical information associated with each of these splittings and use it to fill the kinematical Lund planes discussed in chapter 3.

   For the analyses performed within the work presented here, we used only sequential recombination algorithms from a specific family of algorithms known as the generalised-k_t family. The distance measures associated with this family of algorithms are defined as [72]:

$$d_{ij} = \min\left(p_{i,t}^{2p},\, p_{j,t}^{2p}\right)\frac{\Delta R_{ij}^2}{R^2}\,, \qquad \Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2\,, \tag{2.1}$$

$$d_{iB} = p_{i,t}^{2p}\,, \tag{2.2}$$

where the labels i and j are used to designate particles (either real final-state particles or pseudo-particles which are already the result of particle recombination) in the event and B refers to the beam, $p_{i,t}$ is the transverse momentum of particle i, R is the parameter of the algorithm which defines the jet's reach and is known as jet radius (in an analogy to R in cone algorithms), and p is the (continuous) parameter which defines this as a family of algorithms. The way the algorithms of this family work based upon these distance measures is as follows:

1. calculate all d_{ij} and d_{iB} for all possible combinations, and determine which one is the minimum;

2. if a distance d_{ij} between two particles is the smallest, then combine those particles into a new pseudo-particle object, and return to step 1;

3. if, on the other hand, a distance d_{iB} is the smallest then take object i and declare it as a final-state jet. Return to step 1 afterwards;

4. repeat this iterative process until there are no more final-state particles left.

From the procedure described above it is quite simple to understand that as long as two particles are within a distance ΔR_{ij} which is smaller than the jet radius R, the distance d_{ij} will always be smaller than both d_{iB} and d_{jB} and, hence, the two particles will be recombined together. Only when there are no more particles within the jet's reach, that is, only when $\Delta R_{ij} > R$ for all pairs $\{i, j\}$ will the remaining objects be declared as final-state jets. Although all algorithms in this family share this common feature, the p parameter determines which particle pairs are perceived as being closest together, inducing different orderings to the clustering sequences associated with each algorithm and also producing different jets.

   We will now proceed to analyse the main algorithms of this family. Three of the algorithms described have seen extensive use in literature both in experimental measurements, analysis of MC event generators and theoretical calculations. These are the Cambridge/Aachen (C/A), the k_t and the anti-k_t algorithms. A fourth algorithm belonging to this family will also be discussed here. It involves the concept of formation time introduced in section 1.5 of chapter 1 and we will thus name it the "t_F algorithm". Though it has not been explored in literature, it was extensively used throughout this work.

The Cambridge/Aachen algorithm

The simplest algorithm of the generalised-k_t family is the Cambridge/Aachen (C/A) algorithm which takes the parameter p to be zero. In this case, the distance measures defined by (2.1) and (2.2) simplify to $d_{ij} = \Delta R_{ij}^2/R^2$ and $d_{iB} = 1$, there being no momentum dependence in this particular algorithm. This corresponds to a very simple situation: this algorithm starts by pairing those particles which are closer in angle (i.e., closer in rapidity and azimuth), thus originating a hierarchical clustering tree which is dictated by an angular ordering. When we decluster this tree, starting by undoing the pair of particles which were last put together, these will present the largest angle between them. Successively moving forward with this declustering will mean the angles will be progressively smaller. This is an extremely powerful feature since it is closely related to the idea of angular ordering associated with QCD parton branching, discussed in section 1.3.

The inclusive longitudinally invariant k_t algorithm

Another algorithm of this family is the longitudinally invariant k_t algorithm, usually simply referred to as k_t algorithm. This algorithm takes $p = 1$, such that: $d_{ij} = \min\left(p_{i,t}^2, p_{j,t}^2\right) \Delta R_{ij}^2/R^2$ and $d_{iB} = p_{i,t}^2$. While the way the particles are clustered in the C/A algorithm is energy (or momentum) independent, for the k_t algorithm that is no longer the case. Indeed, objects which are close in angle still desire to cluster first (such is the tendency of the generalised-k_t family determined by presence of the ratio $\Delta R_{ij}^2/R^2$), but now the k_t algorithm assigns a further preference towards the clustering of particles with smaller transverse momentum, that is, it tends to form jets around softer particles. Once again, this can be related to the QCD parton branching structure which states that large-angled, hard emissions tend to happen earlier on while collinear, soft emissions come at a later stage. As such, we are once again assigning to the jet clustering sequence a hierarchical ordering which can be related to the way we think of the QCD branching process. We are essentially producing a tree which collects what we believe to have been the history of the evolution of the jet in the MC parton shower.

The anti-k_t algorithm

Both the C/A and the k_t algorithms have their roots on e^+e^- collisions (see refs. [73] and [74] respectively). However, the distance measures used in that context differ from the ones that have been presented in here. The C/A algorithm as defined above and currently used in the context of hadron colliders was first discussed in the context of deep inelastic scattering (DIS) experiments [75]. As for the k_t algorithm, its metric as defined in the context of e^+e^- collisions was not invariant under longitudinal boosts, a characteristic which is desirable in hadron collider context. For that reason, a new distance measure was proposed in [76] and, around the same time, in [77], the latter one corresponding to the definition given above. We then understand that although they can now be regarded as members of the same family of algorithms, they were introduced separately through adaptations of algorithms originally designed for e^+e^- experiments. Only later was the idea of a generalised family introduced.

The concept of such a family was first explored in [72] to introduce the anti-k_t jet algorithm. The motivation behind the introduction of this algorithm arose from the fact that the final-state jets output by the C/A and k_t algorithms present irregular boundaries, unlike those of cone algorithms. While this is irrelevant for the analysis of MC event generator results or for theoretical calculations, from the experimentalists' stand point it is an undesirable feature essentially because it complicates detector calibrations. As such, the anti-k_t algorithm was proposed (see

ref. [72]) as a solution to this issue and as simply another algorithm of the generalised-k_t family, which took $p = -1$. This corresponds to using as distance measures:

$$d_{ij} = \min\left(\frac{1}{p_{i,t}^2}\,,\,\frac{1}{p_{j,t}^2}\right)\frac{\Delta R_{ij}^2}{R^2}\,, \qquad \Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2\,, \qquad (2.3)$$

$$d_{iB} = \frac{1}{p_{i,t}^2}\,. \qquad (2.4)$$

The condition of finding the minimum between the inverse of the two squared transverse momenta is equivalent to finding their maximum, that is, equation (2.3) could be alternatively written as

$$d_{ij} = \frac{1}{\max\left(p_{i,t}^2\,,\,p_{j,t}^2\right)}\frac{\Delta R_{ij}^2}{R^2}\,, \qquad \Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2\,. \qquad (2.5)$$

While the objects which are closer in angle still prefer to cluster early, there is now a preference towards harder particles. As a result, the anti-k_t algorithm produces jets which grow around hard particles, moving on through successively softer ones. By being centred around hard particles, the jets that arise from this algorithm present pleasant round boundaries, as is experimentally desired. Besides that, this algorithm also has the advantage that is will not start building jets from particles which are uninteresting or undesirable, thus preventing (or at the very least minimising) the occurrence of the so called "junk jets".

The downsize of this algorithm is that its clustering sequence can not be related to the QCD branching structure and its soft and collinear divergences. We thus fail in our attempt to ascribe a physical meaning to the hierarchical clustering tree. For this reason, the following approach to jet finding has become standard and was thus used in this master dissertation:

1. jets were found using the anti-k_t algorithm, such that each jet has assigned a set of final-state particles and a clustering sequence;

2. we then take the particles within each individual jet and recluster them using either the C/A, k_t or t_F (discussed below) algorithms. By doing so, we get back a clustering sequence which can be related to the structure of QCD;

3. having that new clustering sequence we can proceed by undoing each successive pairing to get the kinematical information associated with each one, and using it to fill in the Lund planes (chapter 3).

The t_F algorithm

All the algorithms discussed hitherto are standards in literature which have seen extensive use in all contexts from experimental, to phenomenological and analytical studies. We will now proceed to propose the use of an algorithm from the generalised-k_t family which has not been explored. Our proposal is to take $p = 1/2$ which translates into the following distance measures:

$$d_{ij} = \min\left(p_{i,t}\,,\,p_{j,t}\right)\frac{\Delta R_{ij}^2}{R^2}\,, \qquad (2.6)$$

$$d_{iB} = p_{i,t}\,. \qquad (2.7)$$

Let us now recall the concept of formation time introduced in section 1.5 of chapter 1. In that section we have arrived at equation (1.17), written in hadron collider variables, and noted

that the quantity t'_F given by it is a formation time which is a proxy for the usually defined formation time given in equations (1.14) and (1.15). If we compare equation (1.17) with (2.6) we quickly identify that the distance measure d_{ij} for this algorithm is related to that formation time: $\min(p_{i,t}, p_{j,t}) \equiv p_{2,t}$ (the transverse momentum of the softer particle), such that

$$d_{ij} = \frac{2}{t'_{F,ij}R^2} \, .$$ (2.8)

Let us now interpret what this means for the clustering sequence produced by this algorithm. In this case, particle pairs with larger formation time are recombined first whereby we obtain a hierarchical tree which is ordered in formation time, such that if we walk back through this tree, undoing each successive pair, the characteristic time between splittings will increase.

In our discussion of the C/A and k_t algorithms, we noted that these produce hierarchical clustering sequences in which the successive splittings are ordered in angle and transverse momentum, respectively. We have also argued that these algorithms possess an extraordinary power associated with the fact that they produce clustering sequences which can be related to the natural QCD parton branching structure by which hard and large-angled splittings happen first. While this is true in vacuum, it is not necessarily the natural ordering of the splittings in the QGP. In fact, in the presence of this hot and dense medium it is more natural to consider an ordering of the splittings according to their formation time [49, 78]. This is why we have chosen to explore this algorithm within the context this master dissertation: we use it to recluster the anti-k_t jets since we believe it to be the algorithm which sets the clustering sequence that best relates to the physical structure of QCD showering in the medium. This idea will be explored further in the upcoming chapters.

2.3   Quick note on recombination schemes

Alongside the jet algorithm used, a recombination scheme must be defined which determines how to combine the four-momenta of the particles which are clustered to form the jet. The most widespread and used algorithm in the E-scheme or 4-vector recombination scheme which operates by simply adding the four-momenta of the two objects (final-state particles or pseudo-particles) which are merged together. Since this is the current recommendation [79], it was the recombination scheme which was used in this dissertation.

Chapter 3

The kinematical Lund plane

Jet substructure studies aim at extracting information from the radiation pattern inside jets. The information gathered from these studies is extremely valuable because it can be employed in a multitude of different investigations. For example, it can be used (i) to determine whether a jet stems from the branching of a hard parton or from the hadronic decay of a boosted Z, W or H boson produced in the hard scattering, (ii) to probe the Standard Model or (iii) to increase sensitivity to new physics studies. In the past years, researches involving jet substructure have been dominated by two main approaches [80]. The first of these consists in engineering (often complex, high-level) observables whose properties can be analytically calculated by perturbative QCD. Those observables can later be experimentally measured and the results compared to the theoretical analytical predictions. The second approach exploits the power of machine learning (ML) algorithms to extract useful information from data (either experimental or from MC parton showers). The inputs for these algorithms have mainly been (η, ϕ) or (y, ϕ) jet images, the jets' clustering sequences obtained with sequential recombination algorithms or a set of (simple) jet substructure observables (see ref. [80] and references therein for more details). Each approach has its advantages and disadvantages. The first requires a laborious effort to be employed into the construction of observables that are sensitive to whichever phenomenon or feature we desire to examine and it is, therefore, straightforward to understand that it involves a high-level processing of the data. At the same time, this approach grants a solid understanding of the physics which is being exploited. The second approach, on the other hand, requires minimum to no (physical) processing: the data can be fed to the ML algorithm, which will extract from it the relevant information for the particular investigation under way. While the performances obtained by this method can be substantially higher than those achieved with manually constructed observables, that comes at the expense of a loss of clarity and physical intuition as to which features are actually being exploited.

Recently, a third approach has been put forward in [80] to *bridge the fault-line between manually constructed observables and machine learning approaches*. That is, the proposed method aims at finding a compromise between the maximum (laborious) processing and solid physical grasp offered by the engineering of (high-level) observables and the minimum processing, low physical understanding provided by ML approaches, while also retaining the remarkable performance achieved by the latter. To that end, the use of kinematical Lund planes [81] is suggested. These provide a visual representation of the radiation pattern inside jets, allowing to easily identify interesting or distinct features of the jet substructure. That (simple) visual identification can then motivate the construction of specific observables which are sensitive to specific regions of those planes. At the same time, the information used to build the Lund planes can also be input into a ML algorithm to extract information. However, contrary to the previously described ML approaches, in this case the visual analysis of the kinematical Lund planes means that the dis-

criminating features detected by the ML algorithm can, in principle, have been clearly identified and understood *a priori*.

In the context of relativistic heavy-ion collisions, the study of jet substructure aims at discriminating between jets which have interacted with the QGP and those which have not or, analogously, aims at identifying which aspects of jet substructure are modified by the presence of the medium. Any of the three approaches presented above can potentially be employed. However, the use of Lund planes has stirred up some interest in the heavy-ion community since the 2017 CERN TH Institute "*Novel tools and observables for jet physics in heavy-ion collisions*" where it was established consensually *that a comprehensive comparison of existing and future jet quenching models can be based on the distribution of hadronic fragments in the kinematical Lund plane* [49].

In the context of this dissertation, we focus on the exhaustive analysis of a myriad of different Lund planes. We exploit the use of different kinematical variables for the $y-$axis, fill the planes in three different ways and make use of clustering sequences from three different jet algorithms. This comprehensive set of studies allows us to draw firm and consistent conclusions onto how jets are modified in the presence of the QGP, when compared to those from pp collisions. These studies are entirely visual in nature, involving the interpretation of the different kinematical planes to understand the effects in play. Within the timeframe set for this dissertation it was, unfortunately, not possible to manually construct observables sensitive to the regions of phase space where stronger modifications were observed. We leave that for future work. However, a ML architecture was explored for confirmation of some of the visually drawn and physically motivated conclusions.

With this chapter we intend on making clear the concept of Lund plane. By the end of it, the reader should be completely comfortable with the idea of what is a kinematical Lund plane and how it is constructed, making it easy to follow the analyses and interpretations of the planes presented in chapters 4 and 5. To that end, in the next three sections, we will motivate and discuss different definitions for the axes of the Lund plane, describe the main features of each one of those definitions, explain the three different ways in which the kinematical planes were filled and examine the matter of normalisation.

3.1 The choice of kinematical variables for the Lund plane

We have seen that a jet is developed through a cascade of successive $1 \rightarrow 2$ parton splittings of the form $q \rightarrow qg$, $g \rightarrow gg$ or $g \rightarrow q\bar{q}$. The Lund plane explores the idea that the phase space of these emissions can be described by two key kinematical variables. As we have emphasised repeatedly throughout this dissertation, the structure of QCD involves two logarithmic divergences: the soft divergence, when $\zeta \rightarrow 0$, and the collinear divergence, when $\theta \rightarrow 0$. We introduced this property promptly in section 1.2 of chapter 1, when we demonstrated that the probability of emission of a soft and collinear gluon is given by equation (1.11). Recalling this equation, it is straightforward to understand that, in the DLA at fixed coupling, the emissions of a soft and collinear gluon are uniformly distributed in a plane with axes defined as $(\ln(\theta), \ln(\zeta))$, with a density given by

$$\rho = \frac{2\alpha_s C_F}{\pi} \,.$$

Hence, the structure of QCD motivates the use of logarithmic scales for the axes of the kinematical Lund plane. Using non-logarithmic scales would mean that emissions would be exponentially spaced throughout the plane, making for an inadequate and unintelligible representation of the phase space.

The structure of QCD obviously also inspires the particular choice of kinematical variables. Specifically, we look for variables that, in the DLA at fixed coupling, grant us an uniform

distribution of emissions throughout the allowed region of the Lund plane. This is motivated by the fact that, in reality, running coupling effects, higher order corrections and hadronisation ruin the simple picture painted by the DLA at fixed coupling, and the Lund planes are anything but uniformly filled (as we will have the chance to observe in chapters 4 and 5). Selecting kinematical variables that ensure that the phase space is uniformly filled according to the DLA at fixed coupling means that any deviations from that uniform behaviour come from the aforementioned effects, allowing us to directly observe the impact that they have on jet substructure.

Typically, the $x-$axis uses the angle of the emission θ, i.e., the angle between the two children resulting from the branching. This explores the collinear (logarithmic) divergence of QCD. The standard choice for this axis is to take $x \equiv \ln\left(1/\theta\right),$ such that collinear emissions are cast away to $+\infty$ and the part of the axis that is actually represented is filled with more interesting (larger-angled) emissions. The $x-$axis will not be filled all the way down to $\ln(1/\theta) = 0$. A kinematical limit is imposed by the jet radius used with the anti-k_t jet-finding algorithm: the angle between the two child partons cannot be larger than the jet radius, since that would cause them to be outside of the jet. This imposes a restriction that $\theta \leqslant R$ and, as a consequence, the region below $\ln\left(1/R\right)$ is kinematically prohibited and will be empty in all the constructed Lund planes.

The definition of the $y-$axis depends on the particular preference of the author and has seen several different variants in literature. However, all of those variants share the characteristic that they take a kinematical variable relatable to the soft divergence of QCD, i.e., to the logarithmic derivative $d\left(\ln\zeta\right)$ in equation (1.11). This serves to guarantee the desired uniform distribution of the emissions in the allowed region of the Lund plane. As a matter of fact, in the DLA and from the point of view of QCD's logarithmic structure, all the kinematical variables used are equivalent. This correspondence is apparent from equation (1.11), by performing a variable transformation of ζ to the desired kinematical variable. We will exemplify this explicitly below for one particular case. That should suffice in making this point clear, given that the remaining cases can be treated identically.

We will now proceed with the presentation of all the kinematical variables exploited for the vertical axis and of the different features of the corresponding Lund planes.

3.1.1 $y \equiv \ln(1/z)$

Usually, the momentum fraction z carried out by the softer child is used instead of the energy fraction ζ. This transformation is harmless, since equation (1.11) could have been derived directly using the momentum sharing fraction z, rather than ζ.

Analogously to the $x-$axis, we can define the vertical axis as $\ln\left(1/z\right)$ such as to cast the soft emissions away to $+\infty$, leaving more relevant emissions to be represented in the visualised plane. Just as we have seen above, in the DLA at fixed coupling, these emissions will be uniformly distributed on the kinematical plane. However, they will not fill the entire region defined by $\ln\left(1/z\right) > 0$. Since we fill the plane using the momentum fraction of the softer child, z is limited such that $z \leqslant 1/2$, which means that there will be no emissions below $\ln(2)$. Other than that, there are no restrictions to the filling of this plane, and it is a semi-finite plane whose appearance should resemble that of figure 3.1. Harder emissions appear at the bottom of the plane, near the horizontal line $y = \ln(2)$ and large-angled emissions populate the left of the diagram, close to the vertical line $x = \ln(1/R)$.

Figure 3.1: Schematic of a Lund plane with the choice of $y-$axis as $\ln{(1/z)}$. The two kinematical limits are represented by the lines $x = \ln(1/R)$ and $y = \ln(2)$. Soft and/or collinear emissions are cast away to $+\infty$ in this plane, just as depicted. The hashed zone is kinematically prohibited and the green region is uniformly filled (in the DLA at fixed coupling).

3.1.2 $y \equiv \ln(z\theta)$

Figure 3.1, presents a schematic of the $(\ln(1/\theta), \ln(1/z))$ Lund plane and it allows us to perceive that the phase space is limited to a semi-finite rectangular region. Some authors have a preference towards triangular Lund planes, i.e., planes in which emissions are bound in a semi-finite triangularly-shaped region and prohibited elsewhere. To achieve such a shape, it is necessary to correlate the two axes, x and y. This can be accomplished by introducing an angular dependence in the kinematical variable chosen for the vertical axis. At the same time, we still desire to retain a variable that is relatable to the soft (logarithmic) divergence of QCD. One simple way to do this, which has been extensively used in literature, is to keep using the momentum sharing fraction but multiply it by the angle of the emission, such that the vertical axis is defined as $\ln{(z\theta)}$. We can show explicitly that this choice of variable is equivalent to the use of $\ln{(1/z)}$ from the point of view of QCD's logarithmic structure and that it also ensures that the plane is uniformly filled by emissions. To do this we take equation (1.11), replace ζ with z (this equation could have been derived in that form anyway) and perform the variable transformation $z \to z\theta$, such that

$$dS = \frac{2\alpha_s C_F}{\pi} d\left(\ln{(z\theta)}\right) d\left(\ln{\theta}\right) \frac{d\phi}{2\pi}$$
$$= \frac{2\alpha_s C_F}{\pi} \left[d\left(\ln{z}\right) + d\left(\ln{\theta}\right)\right] d\left(\ln{\theta}\right) \frac{d\phi}{2\pi}$$
$$\simeq \frac{2\alpha_s C_F}{\pi} d\left(\ln{z}\right) d\left(\ln{\theta}\right) \frac{d\phi}{2\pi} .$$

In the last step we have discarded the term in $\left[d\left(\ln{\theta}\right)\right]^2$ since this is of higher (logarithmic) order in the kinematical variable θ. Thus, we see that in the DLA, the variables z and $z\theta$ are (ordering-wise and logarithmically) equivalent. Moreover, with this kinematical variable (and using a logarithmic scale) the Lund plane will be uniformly filled with emissions. Both of these

Figure 3.2: Schematic of a $(\ln(1/\theta), \ln(z\theta))$ Lund plane. The two kinematical limits are represented by the lines $x = \ln(1/R)$ and $y = \ln(1/2) - \ln(1/\theta)$, the latter corresponding to the condition $z = 1/2$. The phase space of emissions is triangular-shaped and is uniformly filled (in the DLA at fixed coupling).

statements are true for all of the kinematical variables presented subsequently and we will abstain from repeating this explicit demonstration.

In the Lund plane with the vertical axis defined as $\ln(z\theta)$, diagonal lines of slope $m = -1$ mark sets of emissions with the same $z = const.$, i.e., emissions in which the softer child parton carries out the same momentum fraction z. Therefore, in this Lund plane construction, the kinematical limit is established by the diagonal line of slope $m = -1$ for which $z = 1/2$. As such, we are left with an uniformly filled (in the DLA at fixed coupling), semi-finite plane of triangular shape, just as depicted in figure 3.2. There is no lower limitation to this plane. Close to the diagonal kinematical limit, there lie the hard splittings, i.e., splittings with $z \simeq 1/2$. As we successively move away from that line (in the transverse direction), z decreases, that is, the emissions get increasingly softer.

3.1.3 $y \equiv \ln(k_t)$

So far, none of the Lund plane definitions considered the absolute momentum scale involved in the splitting; only the momentum fraction z has been taken into account. We can also choose the kinematical variable using the absolute transverse momentum carried out by the softer child parton, that is, we can replace z with $p_{2,t} = zp_t$, where p_t is the transverse momentum of the parent parton (with respect to the beam). To get a triangularly shaped plane, we can once again introduce the angle of the emission in the definition of the $y-$axis kinematical variable via a multiplication. By doing so, we are left with a Lund plane with the $y-$axis defined as $\ln(k_t) = \ln(zp_t\theta)$, k_t being the transverse momentum of the softer child parton measured with respect to the direction defined by other (harder) emitted parton. Therefore, we have two different transverse momenta: (i) p_t, which refers to the transverse momentum of a given parton with respect to the beam; and (ii) k_t, which is a local property of each splitting (much like z) and will be usually referred to as the "transverse momentum of the splitting or emission". At places we also refer to k_t simply as "transverse momentum", same as for p_t; however, from the context, it should be fairly simple for the reader to understand to which of the two quantities
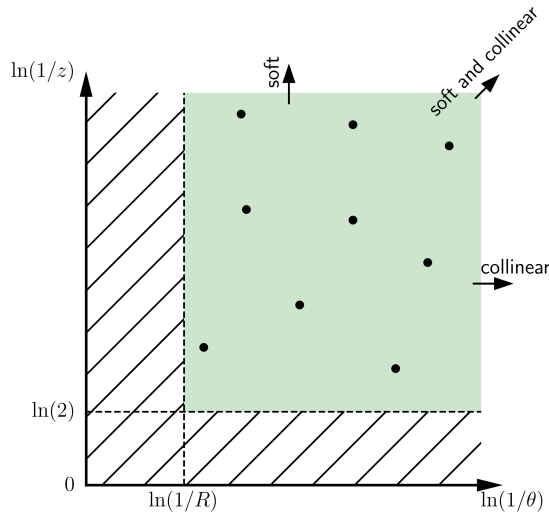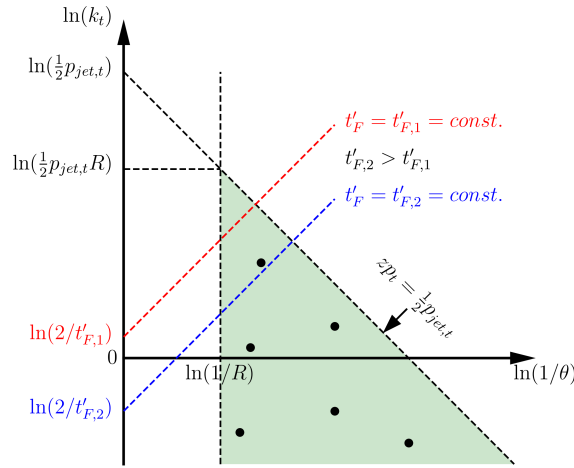
Figure 3.3: Schematic of a Lund plane with the $y-$axis defined as $\ln(k_t)$. The two kinematical limits are represented by the lines $x = \ln(1/R)$ and $y = \ln(\frac{1}{2}p_{jet,t}) - \ln(1/\theta)$. Diagonal lines of slope $m = -1$ correspond to sets of emissions with the same absolute transverse momentum of the softer child parton ($p_{2,t} = zp_t = const.$). On the other hand, horizontal lines plainly stand for splittings with the same transverse momentum k_t. Finally, lines with slope $m = 1$ correspond to splittings with the same formation time t'_F. The intersection of the latter lines with the $y-$axis determines the formation time associated with each one: $b = \ln(2/t'_F)$.

the discussion refers to.

In this case, the diagonals of slope $m = -1$ are not lines of constant z, but rather lines with constant $p_{2,t} = zp_t$, i.e., correspond to splittings in which the respective softer child partons share not the same relative transverse momentum but the same absolute transverse momentum. As such, in this construction, it is impossible to identify if a splitting in one of those lines corresponds to a hard emission with $z \simeq 1/2$ or rather to a softer emission (e.g., $z = 1/10$) from a parent parton with larger transverse momentum p_t. Once again we will have a diagonal kinematical limit, in this case imposed by the condition that $zp_t \leq \frac{1}{2}p_{jet,t}$.

On the other hand, it is plain to see that horizontal lines correspond to sets of splittings which share the same transverse momentum k_t. Also noteworthy about this plane is that unit slope lines (i.e., $m = 1$) define sets of emissions with the same formation time t'_F,

$$
\begin{aligned}
y = \ln(k_t) &= \ln(zp_t\theta) \\
&= \ln(zp_t\theta^2) + \ln(1/\theta) \\
&= \ln(2/t'_F) + x.
\end{aligned}
$$

The interception with the $y-$axis is valued to $b = \ln(2/t'_F)$, meaning that the higher a line intersects that axis the shorter the corresponding formation time.[1] The main features of the $(\ln(1/\theta), \ln(k_t))$ Lund plane are presented in figure 3.3. Close to the kinematical diagonal limit there lie the emissions where the softer splitting child has the largest allowed absolute transverse momentum with respect to the beam ($p_{2,t} \simeq \frac{1}{2}p_{jet,t}$).

---

[1]Note that this feature has been attributed to Lund planes of the form $(\ln(1/\theta), \ln(z\theta))$ [49], but that is misleading.
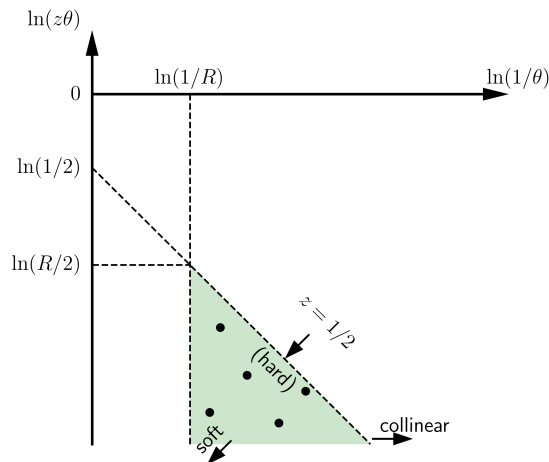
3.1.4 $y \equiv \ln\left(t'_F\right)$
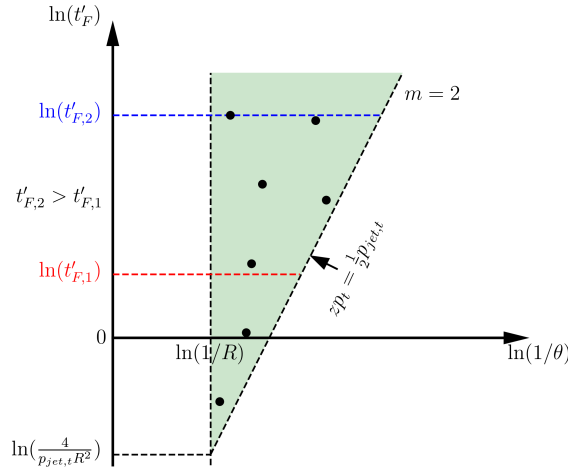
Figure 3.4: Schematic of a $(\ln(1/\theta), \ln(t'_F))$ Lund plane. The two kinematical limits are represented by the lines $x = \ln(1/R)$ and $y = \ln(\frac{4}{p_{jet,t}}) + 2\ln(1/\theta)$ (slope $m = 2$, unlike any of the other constructions). Diagonal lines parallel to the kinematical limit define sets of emissions with the same absolute transverse momentum of the softer child parton ($p_{2,t} = zp_t = const.$), while horizontal lines characterise splittings with the same formation time t'_F.

We have already seen that, in heavy-ion collisions, as the QGP is created and evolved, jets are, simultaneously, produced and developed, interacting with it. For this reason, it would be desirable to be able to ascertain at which spacetime point an interaction between two partons occurs. It is the QGP itself that provides the spacetime reference frame to formulate this problem and, as we have seen in section 1.5 of chapter 1, a variable which should help in describing the evolution of jets within the medium is the formation time. Thus, in our current context of constructing Lund planes with different kinematical variables, the formation time is a natural choice to consider for studying the radiation pattern inside jets from heavy-ion collisions. Furthermore, it is manifest that the formation time t'_F, as defined in equation (1.18), is logarithmically equivalent to any of the previous definitions. We can perform a transformation $z \to t'_F$ and retain the structure of QCD with its logarithmic divergences, the allowed region of the Lund plane will be uniformly filled and the emission probability will be exactly the same. This was our motivation for exploring a novel Lund plane construction not yet seen in literature: a $(\ln(1/\theta), \ln(t'_F))$ kinematical Lund plane. Alternatively, t_F can also be used and indeed we have done so to demonstrate that the kinematical Lund planes that it yields are identical to those using t'_F, thus proving that t'_F is really a good proxy for the formation time t_F (as stated in section 1.5 of chapter 1).

In the $(\ln(1/\theta), \ln(t'_F))$ Lund plane, lines of constant formation time are simply horizontal lines, with larger times corresponding to higher $y-$coordinates. This means that early splittings populate the bottom of the plane, while later ones crowd its top. In this particular kinematical Lund plane, lines of slope $m = 2$ are lines of constant $p_{2,t} = zp_t$, that is, the splittings in those

lines share the absolute transverse momentum of the softer child parton,

$$y = \ln(t'_F) = \ln\left(\frac{2}{zp_t\theta^2}\right)$$

$$= \ln\left(\frac{2}{zp_t}\right) + \ln\left(1/\theta^2\right)$$

$$= \ln\left(\frac{2}{zp_t}\right) + 2\ln\left(1/\theta\right)$$

$$= \ln\left(\frac{2}{zp_t}\right) + 2x\,.$$

The kinematical limit is still given by the condition that $zp_t \leq \frac{1}{2}p_{jet,t}$, which in this case means that the splittings need to be above the line with slope $m = 2$, corresponding to the condition that $zp_t = \frac{1}{2}p_{jet,t}$. Figure 3.4 shows all the main features of this Lund plane, which will take a prominent role in our results. Note that there is no kinematical restriction binding its top, meaning that splittings can, in principle, extend to infinity (as shown).

3.2   Different filling procedures for the kinematical Lund planes

In the previous section, we have seen that the Lund plane is essentially a two-dimensional representation of the phase space of emissions within a jet and that, given our careful choice of kinematical variables, it is uniformly filled in the DLA at fixed coupling. The main features of the different Lund plane constructions were discussed and, as such, the only thing that still requires explaining is how exactly the plane is filled.

The procedure to fill the plane is done jet-by-jet. We have explained in chapter 2 that, given the final-state particles in an event, we find jets using the anti-k_t jet algorithm and afterwards recluster each of the found jets using either the C/A, k_t or t_F (reclustering) algorithms. By doing so, we are left with a clustering sequence which relates (in principle) to the natural structure of QCD. This clustering sequence can be regarded as a history tree of the development of the jet where the two objects that were last put together by the reclustering algorithm should, in principle, correspond to the first branching that occurred in the development of the jet. For that reason, we can walk backwards through this clustering sequence to "see" how the jet developed.

At each branching, we have two child objects with four-momenta p_1 and p_2. As always, we define that p_2 is the four-momentum of the softer child and then calculate the appropriate kinematical variables associated with its emission:

$$\theta \equiv \theta_{12}\,,$$

$$\zeta = \frac{E_2}{E}\,,$$

$$z = \frac{p_{2,t}}{p_t}\,,$$

$$k_t = p_{2,t}\theta = zp_t\theta\,,$$

$$t_F = \frac{2}{E_2\theta^2} = \frac{2}{\zeta E\theta^2}\,,$$

and

$$t'_F = \frac{2}{p_{2,t}\theta^2} = \frac{2}{zp_t\theta^2}\,.$$

where E and p_t are, respectively, the energy and transverse momentum of the parent parton (i.e., the emitter).

This gives us the coordinates of each splitting on each of the previously defined Lund plane constructions. (Note that in our definitions we take always the properties of the softer child parton. This has already been pointed out in the preceding section since the restriction $z \leq 1/2$ determines one of the two kinematical limits.)

The first way in which we have filled the plane consists in undoing only the first splitting in the clustering tree and filling the Lund plane with the coordinates of that emission alone. Therefore, each jet contributes to the Lund plane with one single point (one single splitting or emission). That point should, in principle, correspond to the branching of the original hard parton produced in the hard scattering.

The second way in which we have filled the kinematical Lund plane is to consider only the splittings that happen along the hardest branch. This means that at each splitting we calculate the appropriate kinematical variables, fill the plane with that coordinate point and then move along following the harder splitting child. This produces what is called the "primary Lund plane". We say that we fill the plane only with the "primary emissions", i.e., emissions from the hardest branch.

Finally, we have also produced what we call a "full Lund diagram" in which we consider all the splittings in the clustering tree. By using this configuration, we make use of all the information inside the jet and increase the statistics considerably. The disadvantage is that while with the primary Lund plane we have an explicit sequence or path of emissions along the plane, in the full Lund diagram we lose that unambiguous ordering. This will be a relevant detail for chapter 6.

At this point, it is important to return to the idea that, in the DLA at fixed coupling, we expect an uniform distribution of emissions throughout the constructed kinematical Lund planes. Strictly speaking, that only applies to the latter filling procedure in which all emissions within the jets are considered. Filling the plane with the primary emissions should still yield a behaviour which approximates a uniform distribution; however, it is not expected that the first emissions are uniformly distributed across the plane.

3.3   Comments on event weighting and normalisation

Even though the different events that may result from the hard scattering process are not equally probable, in a MC event generator they can all be produced following an uniform distribution. This means, for instance, that there will be much more events with high-p_t jets than those that will be observed, in reality, in a collider experiment. For that reason, each event has a specific weight, which is a measure of how likely that event actually is. When filling the Lund plane, each emission that is entered into the splitting map should be weighted by the corresponding weight of the event. This will compensate for the excess of certain types of events and will bring the results into concordance with those observed in experimental measurements.

Besides taking care for applying the weight of each event to the corresponding entries in the kinematical Lund plane, normalisation also deserves consideration. There are several different possibilities for the normalisation and, once again, the particular choice depends on the particular preference of the author. Common normalisations are, for instance, normalising the Lund plane to unity or to the total cross-section of the process in question. In our case, we decided to perform a normalisation by the effective number of jets N^{eff},

$$N^{eff} = \sum_i w_i N^i_{jets} \,,$$

where w_i is the weight of event i and N^i_{jets} the number of jets in that same event. The inverse of this constant $(1/N^{eff})$ is used to scale every bin in the 2D Lund plane such that the splitting map is normalised per jet.

The different normalisations cause only a change on the absolute scale of each individual Lund plane, and vacuum planes and medium planes will be rescaled in different ways (since the total effective number of jets in the events will differ from one situation to the other). This calls for attentiveness when interpreting the results of the differences between medium and vacuum kinematical Lund planes, because we are subtracting two self-normalised histograms. Therefore, care should be taken in the conclusions drawn.

Chapter 4

Analysis of $Z + jet$ partonic samples

Now that all the necessary concepts have been introduced, we will turn to the application of kinematical Lund planes in the context of identifying differences between the radiation pattern of jets developed in vacuum and that of those that evolve in the presence of a medium. The events used to perform this study were generated in JEWEL [82], for a $NN \to Z + jet$ hard process, such that samples with a quark or a gluon in the final state of the hard matrix element were produced separately. The vacuum baseline consists in considering pp collisions, while medium events result from the interaction between two nucleons (each within a Pb nucleus), with the products of that interaction travelling across the QGP. The centre of mass energy per nucleon pair used in these samples was $\sqrt{s_{NN}} = 5.02 \, \text{TeV}$ and neither ISR or hadronisation were introduced in the picture.

The event selection criteria used for analysing these samples are as follows:

- only events in which a Z boson with $p_t(Z) > 60.0 \, \text{GeV}$ is present are considered;

- final-state particles with $p_t < 0.5 \, \text{GeV}$ or $|\eta| > 2.5$ are removed from the analysis. (η is the pseudo-rapidity, given by $\eta = -\ln \left[\tan \left(\tilde{\theta}/2 \right) \right]$, where $\tilde{\theta}$ is the angle between the particle's three-momentum and the beam axis)[1];

- jets are found using the anti-k_t algorithm with jet radius $R_1 = 0.4$ and then reclustered with the C/A, k_t and t_F algorithms with $R_2 = 1.0$;

- final-state jets must have $p_t > 30.0 \, \text{GeV}$ and $|\eta| < 2.0$, otherwise they are excluded from the analysis;

- only events with at least one jet are considered.

The results obtained for the different Lund plane constructions and filling procedures detailed in chapter 3 are described and discussed in sections 4.1 to 4.4. A comparison between the two different formation times (t_F and t'_F) is presented in section 4.5 and the explicit paths of splittings in the primary Lund plane are analysed in section 4.6. Finally, the effect of employing a different pseudo-rapidity cut is briefly examined in section 4.7. Because of the large volume of information lying herein, the chapter ends with a summary of the main observations and results (section 4.8).

---

[1]$\tilde{\theta}$ should not be confused with the angle θ of an emission (i.e., the angle between the two children of a splitting), which is used as kinematical variable for the Lund planes.

4.1 Analysing the $(\ln(1/\theta), \ln(1/z))$ Lund planes

4.1.1 Gluon-initiated jets

(a)

(b)

(c)
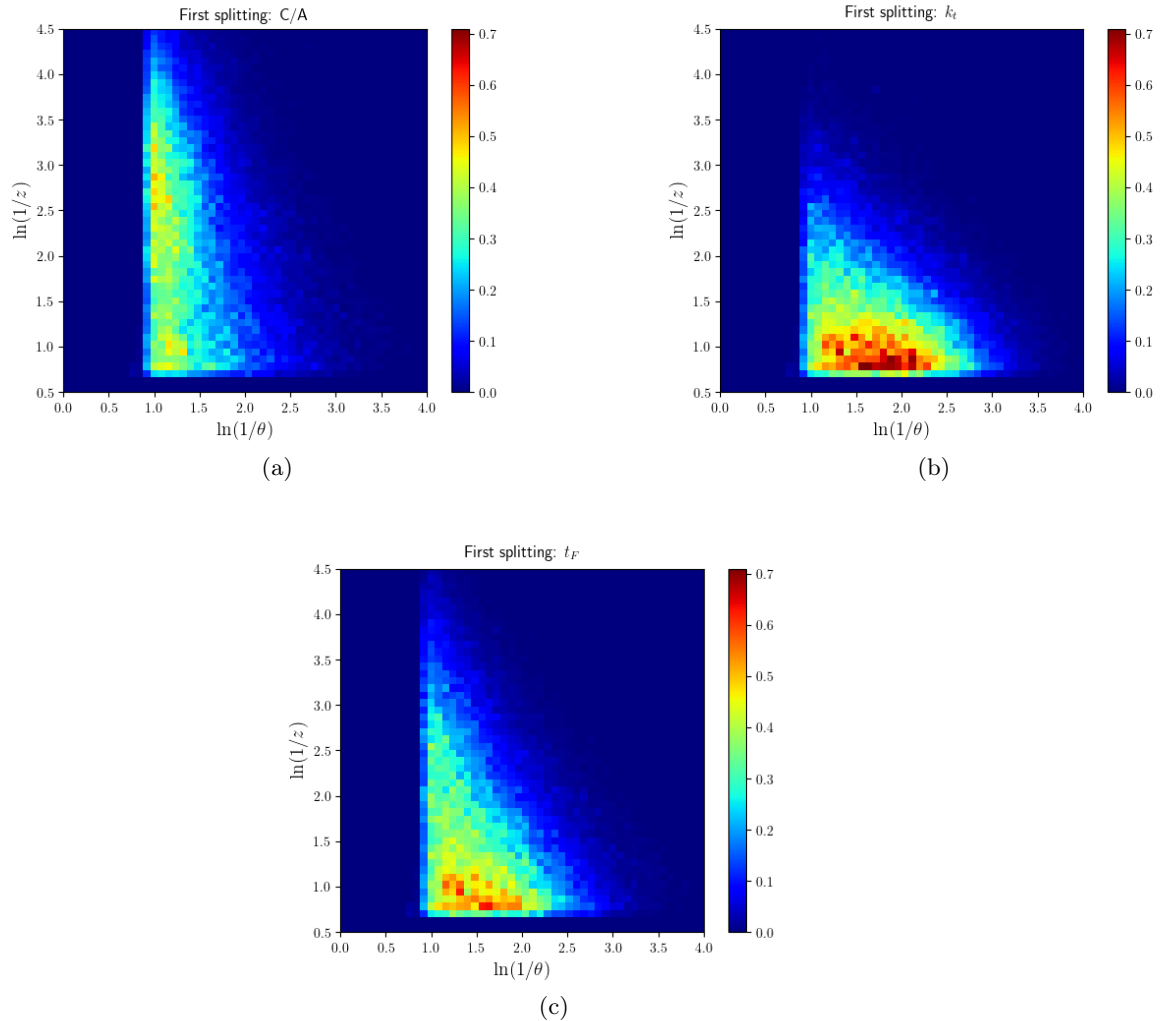
Figure 4.1: Lund planes filled only with the first splitting of gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

Let us start by analysing the results obtained for the samples in which the parton resulting from the hard scattering is a gluon. Before having a look at the kinematical Lund planes it is interesting to note that, the effective number of jets present in the totality of (non-vetoed) vacuum events is 1.6 times larger than the corresponding medium quantity. This is consistent with the well established phenomenon of medium-induced energy loss, which causes a suppression in the number of jets that meet the requirement that $p_t > 30\,\mathrm{GeV}$.

The results for the Lund planes filled only with the first splitting of jets developed in vacuum are presented in figure 4.1. There are several interesting and relevant observations that are immediate from this figure. First of all, we note that the first splitting of C/A-reclustered jets tends to be a large-angled splitting. We see that (first) emissions pool near $\ln(1/\theta) \simeq 1.0 \Leftrightarrow \theta \simeq R_1 = 0.4$ and are roughly uniformly spread vertically along that line, over a region where z
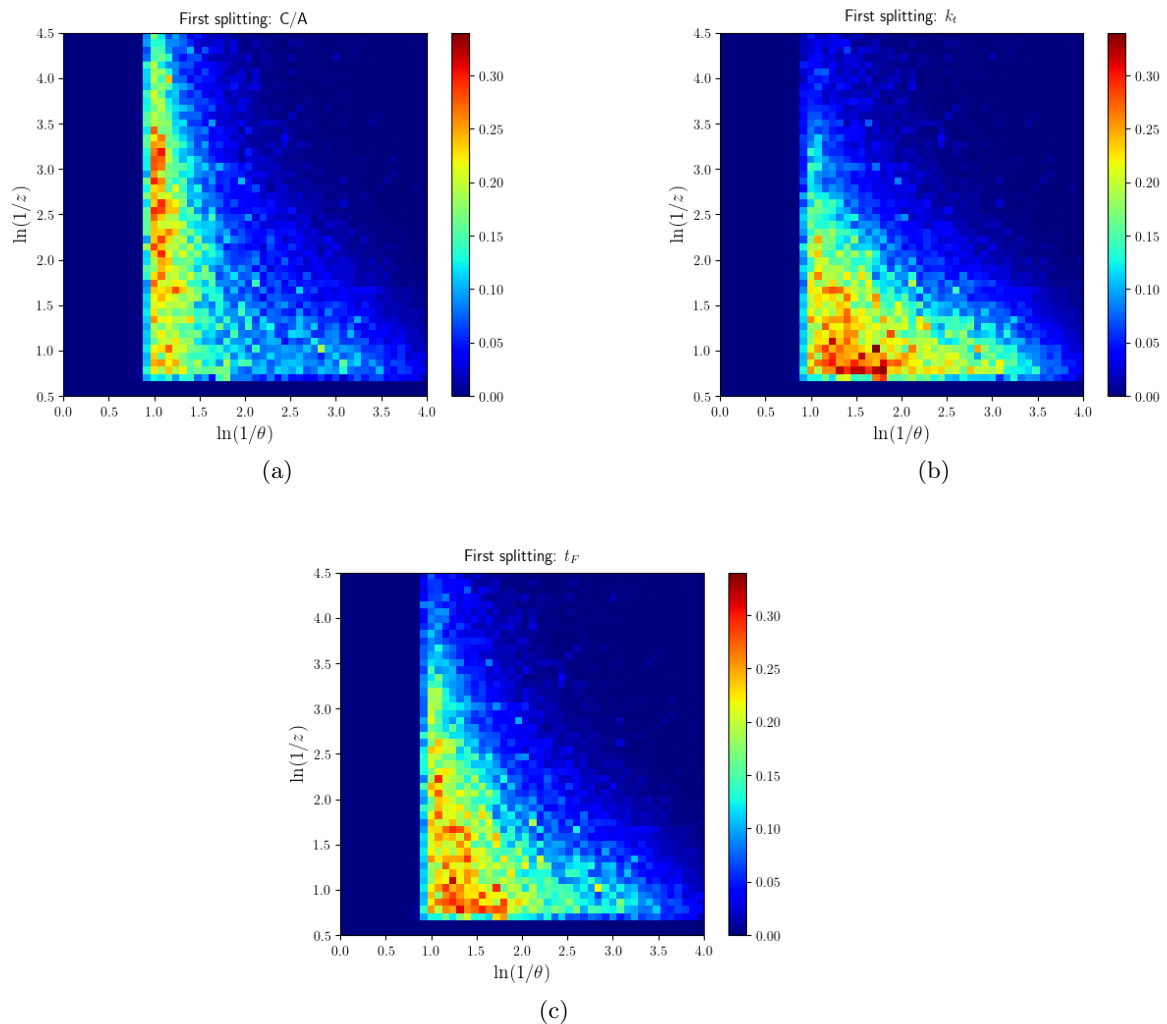
Figure 4.2: Lund planes filled only with the first splitting of gluon-initiated jets, in medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

ranges between 0.03 and $1/2$. That is, it is approximately equally probable for these emissions to be hard emissions or very soft ones. Moreover, we can also see that, for these jets, it is extremely rare for the first emission to be such that $\ln{(1/\theta)} \gtrsim 2.0 \Leftrightarrow \theta \lesssim 0.15$. These observations are all comfortingly consistent with the internal procedure of the C/A algorithm, which produces a clustering sequence that is angular ordered thus identifying as the first splitting the one with larger angle.

On the other hand, the k_t algorithm tends to start by clustering soft particles first, while simultaneously maintaining the preference towards angular proximity (natural to the generalised-k_t family of algorithms). As a result, the first splitting of the clustering tree produced by this algorithm should, in principle, correspond to a hard, (relatively) large-angled emission. Indeed, we observe that using this algorithm yields a Lund plane in which a large portion of (first) emissions is pooled in an area that corresponds to a momentum fraction z approximately between 0.35 and $1/2$, spread over angles θ that range between 0.1 and 0.3. Interestingly, the first emissions are much more densely packed in this case than for C/A-reclustered jets, for which

Figure 4.3: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

we see a much more uniform (spread out) distribution.

Regarding t_F-reclustered jets, we remark that the distribution of (first) emissions appears to be a middle ground between the results from the C/A and k_t algorithms.

The important conclusion here is that the radiation pattern of first emissions is different from one algorithm to another. This follows from the fact that these algorithms reconstruct the history of the shower following different distance measures and consequently produce clustering sequences with different orderings. Another important observation is that since no grooming procedures were employed for "cleaning up" the jet clustering sequences, it may be that many of the first splittings which are being identified are not indeed the ones that would be interesting to isolate. This appears particularly true for C/A-reclustered jets, where the emissions are approximately uniformly distributed in hardness - something that goes against the expected natural QCD ordering, from which we believe hard emissions to happen earlier on. Therefore, it should be clear to the reader that throughout the discussions in this master dissertation the expression "first emission(s)" is meant in the sense of the first as identified by each of the reclustering algorithms.

(a)

(b)

(c)
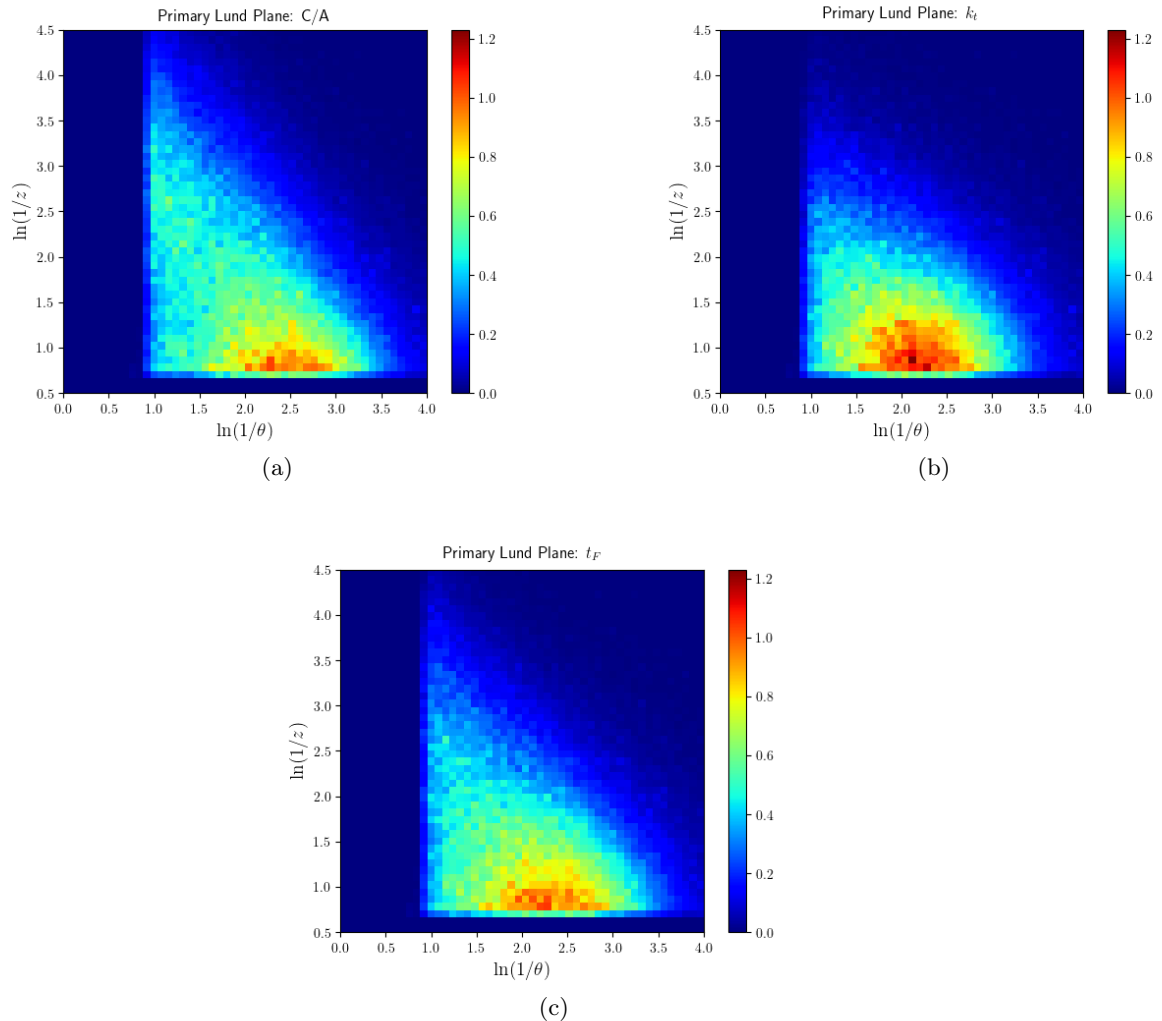
Figure 4.4: Primary Lund planes for gluon-initiated jets developed in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

Hence, conclusions drawn from the corresponding kinematical planes should be approached with caution, and these splitting maps serve essentially as a visual indicator that each algorithm is working as expected.

We can now have a look at the corresponding in-medium results, presented in figure 4.2, where quenching effects come into play and promote modifications to the radiation pattern. The medium appears to promote a dispersion of the emissions throughout the plane.

For the C/A reclustering algorithm that effect is apparent by the spread of emissions beyond $\ln(1/\theta) \gtrsim 2.5$ - a region which is nearly empty in vacuum. However, the first emissions still tend to cluster around $\ln(1/\theta) \simeq 1.0$ and span over the region defined by z ranging from 0.03 to $1/2$.

For the k_t and t_F algorithms, we also see that the emissions spread to smaller-angle regions of the phase space, namely, beyond $\ln(1/\theta) \gtrsim 3.0$, where nearly no emissions are present in the absence of the medium. In the case of these two algorithms, there is also an observable difference to the large-angle emissions around $\ln(1/\theta) \simeq 1.0$, which stretch much more into the soft region than they do in the absence of the QGP. The large splitting concentration observed in vacuum with the k_t algorithm appears to migrate to larger angles and then spread in all (allowed)

(a)

(b)

(c)
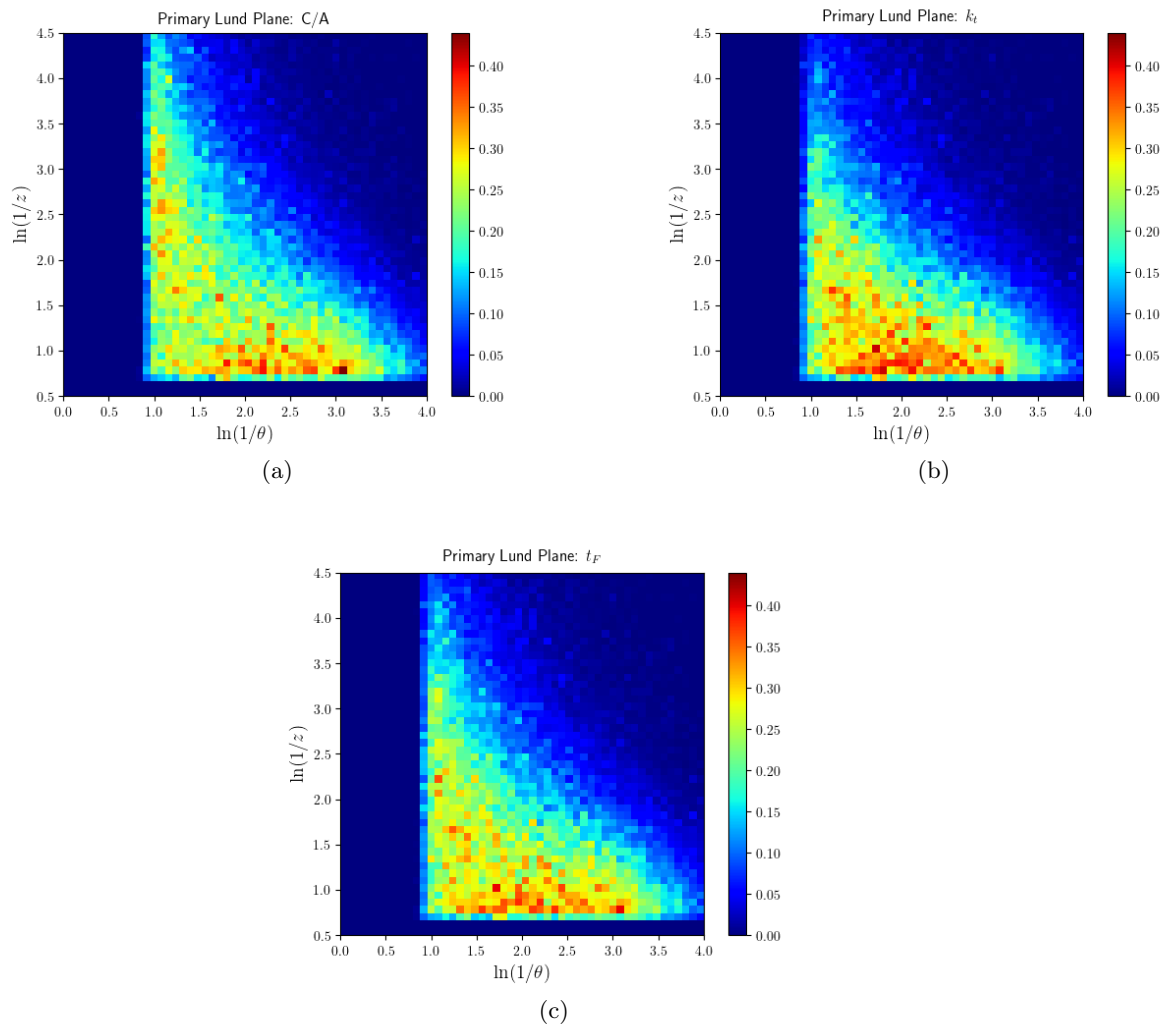
Figure 4.5: Primary Lund planes for in-medium, gluon-initiated jets found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

directions. This is also the case for the (slightly smaller) splitting concentration registered in vacuum with the t_F algorithm.

To corroborate these observations, we can subtract the Lund planes obtained in the presence of the QGP and those produced in vacuum. The results are presented in figure 4.3 from which it is clear that the fraction of splittings that populate the hard, small-angle region of the plane is larger in the medium than in vacuum. Analogously, the fraction of in-medium splittings which populates the soft, large-angle region is also larger than the corresponding vacuum fraction, a fact which is particularly clear for the k_t and t_F algorithms. On the other hand, hard, large-angled regions of the phase space are (comparatively) more densely populated in vacuum than in the medium. These results are consistent with the previously stated observation that the medium spans the emissions throughout the plane.

Let us now analyse the results for the primary Lund planes of both vacuum jets and in-medium jets. First of all, the effective number of splittings entered into the vacuum planes is approximately 2.5 times larger than the effective number of entries in the presence of the medium. This factor is much larger than the factor by which the samples differ in terms of effective number
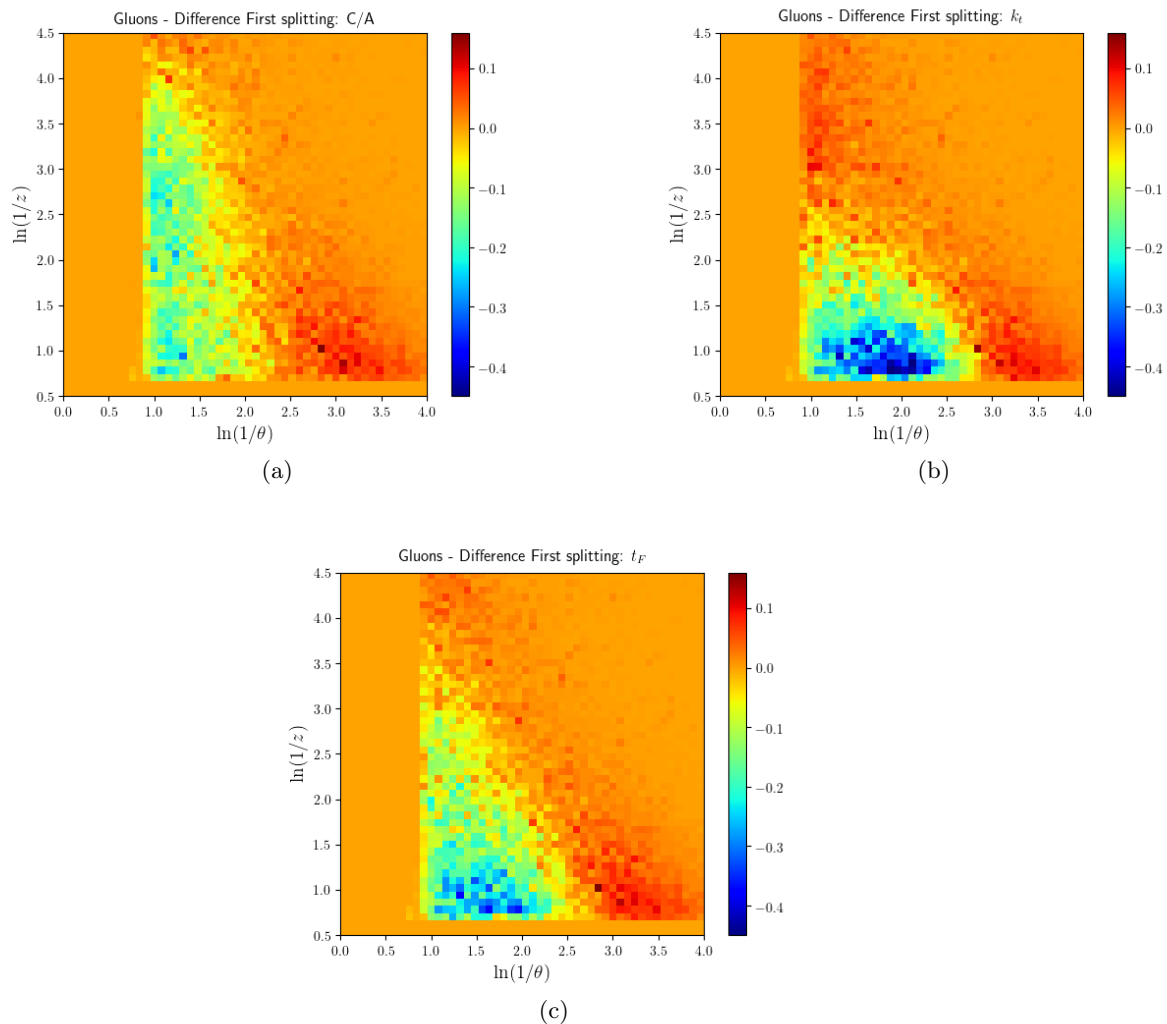
(a)

(b)

(c)
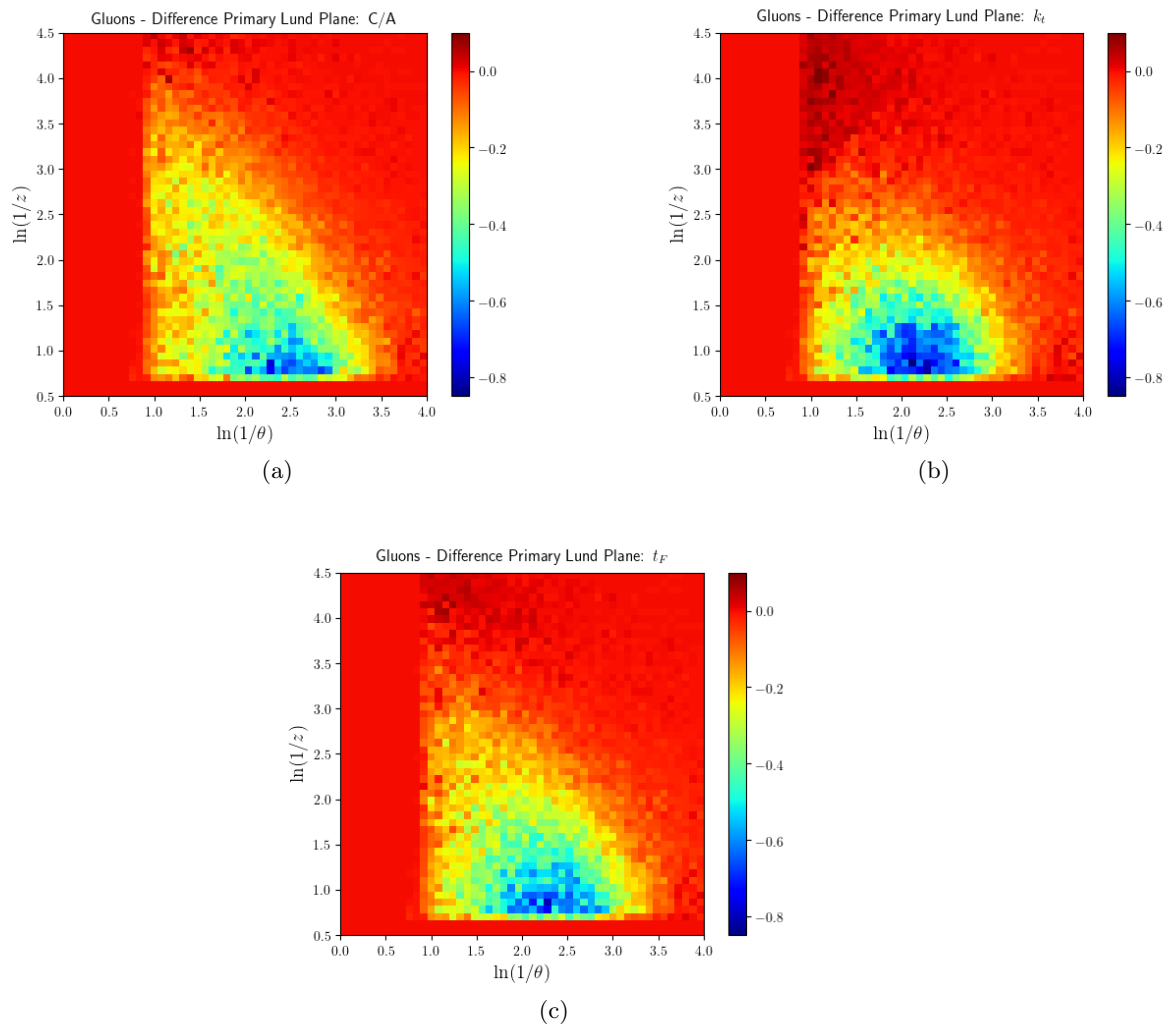
Figure 4.6: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

of jets. This suggests that not only are there fewer jets in the medium case, but these have also fewer emissions within them (at least along the primary branch). This indicates that the QGP causes a fraction of the splittings to be "thrown" outside of the jet cone.

Another interesting feature that deserves some commenting is the fact that the overall number of effective entries in the primary Lund planes corresponding to the different reclustering algorithms is not the same. This indicates that the length of the primary (or hardest) branch, i.e., the number of emissions along that branch, differs from one algorithm to the other. Specifically, C/A-reclustered jets, in vacuum, produce 20% more primary emissions than do k_t-reclustered jets and 10% more than t_F-reclustered ones. In the medium, the differences are slightly attenuated and C/A-reclustered jets yield 15% and 8% more primary emissions than k_t- and t_F-reclustered jets, respectively.

The primary Lund planes for jets developed in vacuum are presented in figure 4.4, for the three different reclustering algorithms. Remarkably, we note that the phase space of primary emissions is triangular-shaped, although there is no kinematical restriction binding the splittings into such

a region (recall figure 3.1). We also note that the plane is not uniformly filled, a fact which stems from the presence of running coupling effects and higher order corrections implemented into the MC parton shower. Finally, the radiation patterns of primary emissions obtained with the three different reclustering algorithms are much more similar amongst themselves than they are for first emissions.

Once again, the t_F algorithm appears to yield an intermediate behaviour between those of the C/A and k_t algorithms. The former has a large fraction of splittings concentrated between $\ln(1/\theta) \simeq 2.0$ and $\ln(1/\theta) \simeq 3.0$, while for the latter there seems to be a slight migration of the emission concentration to the region between $\ln(1/\theta) \simeq 1.5$ and $\ln(1/\theta) \simeq 2.5$. That is, fascinatingly, the k_t algorithm appears to yield a displacement of the splitting pool towards larger angles. Furthermore, those large splitting concentrations appear to be bound closer to the kinematical limit $z = 1/2$ for the C/A algorithm, while extending farther away from it (in the upwards direction) to regions where $z \simeq 0.3$ for the k_t algorithm. It is also apparent that C/A-reclustered jets have a larger fraction of soft, large-angled emissions than do k_t-reclustered ones.

We can now have a look at the primary Lund planes obtained for jets which develop in the presence of the QGP. These results are presented in figure 4.5 from which it is straightforward to understand that the medium (once again) spreads the emissions throughout the plane. In particular, it is clear that the fraction of emissions in the soft, large-angle region is enhanced. This conclusion is consistent with the picture of jet quenching presented in section 2.1 of chapter 2, from which we stated that medium-induced radiation tends to appear in the final state as soft quanta at large angles. Furthermore, the fraction of emissions in the hard, small-angle region is also enhanced by the medium. These conclusions are corroborated by plotting the differences between the radiation pattern of quenched and non-quenched jets (see figure 4.6). The reader should keep in mind that, since the kinematical Lund planes have been normalised by the total effective number of jets, whenever we discuss enhancements (or suppressions) of specific regions of the phase space, in the presence of the QGP, that means that those regions have become more (or less) important per jet, and not in absolute terms.

The results for the full Lund diagram resemble those of the primary Lund planes. For that reason, and as to not overextend this section, those planes are presented in appendix B. However, it is worth mentioning that the total effective number of emissions in vacuum is 2.8 times larger than that for in-medium jets. This is consistent with the result obtained for the primary Lund plane, and reinforces the idea that a lot of the emissions are being "thrown out" of the jet cone by the medium.

4.1.2   Quark-initiated jets

We can now have a look at the behaviour of quark-initiated jets. First of all, it is interesting to note that vacuum samples yield 1.4 times more jets than medium samples. Just as for gluon-initiated jets, this result is in agreement with the well established effect of jet energy loss which happens in the presence of the QGP. More interesting is to note that the suppression in the total effective number of jets is smaller in the present case than it is for gluon-initiated jets. That is, gluon jets lose more energy in the medium. This is consistent with the picture that those jets have (in average) more constituents than quark-initiated ones. And, indeed, for the generated events, gluon jets present approximately 30% and 20% more splittings than do quark jets, respectively for the vacuum and in-medium cases.

In this subsection, we opted towards removing the Lund planes obtained for the t_F algorithm. This was done to reduce the bulk of this part of the discussion and because this algorithm yields an intermediate behaviour between those of the C/A and k_t algorithms, just as we have seen above.
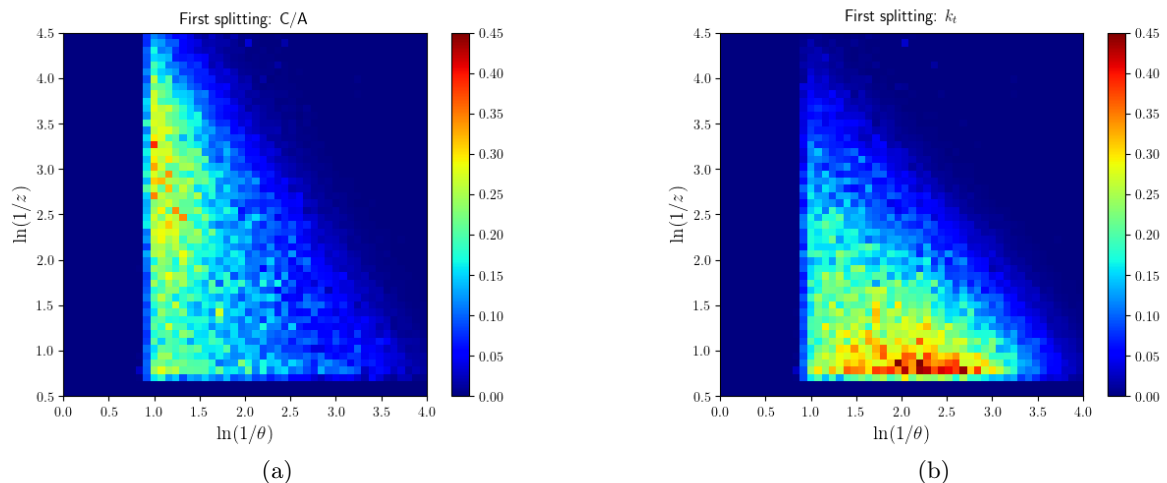
Figure 4.7: Lund planes filled only with the first splitting of quark-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

The results for the kinematical Lund planes filled only with the first splitting of vacuum jets are presented in figure 4.7. Comparing these planes with those obtained for gluon-initiated jets (figure 4.1), we see that the main features observed for those jets remain. The first emissions obtained from declustering the history tree of the C/A algorithm pool themselves along the line $\ln(1/\theta) \simeq 1.0$, stretching vertically from $z \simeq 0.02$ to $z \simeq 1/2$. Also similarly to what we have seen before, the k_t reclustering algorithm produces a large concentration of first emissions in the hard radiation region, specifically where z ranges approximately between 0.4 and $1/2$ and the emission angle θ spans roughly from 0.07 to 0.2.

Besides recording these behavioural similarities, we should also take notice of the differences between quark- and gluon-initiated jets, in vacuum. In the case of the C/A reclustering of quark-initiated jets, we see that the large-angle emissions concentrated along the $\ln(1/\theta) \simeq 1.0$ line stretch more into the soft radiation zone than they do for the gluon case. Furthermore, in the present situation, we see much more first emissions populating the region $\ln(1/\theta) \gtrsim 2.0$, which was nearly empty in the gluon case. As for the reclustering of quark-initiated jets with the k_t algorithm, just as for the C/A algorithm we see that (first) emissions stretch much more into the soft, large-angle region and into the hard, collinear region. The large splitting concentration displayed with this reclustering algorithm, while being a similar feature to gluon-initiated jets also exhibits some differences compared to that case. In particular, it migrates towards smaller angles and it is more concentrated vertically, i.e., it remains closer to the kinematical limit.

Overall, the (first) emissions of quark-initiated jets extend much more throughout the plane, while retaining the distinct features that are observed for gluon-initiated jets.

The impact of the medium over these jets is analogous to that over gluon-initiated ones. Specifically, it destroys a great portion of the structure observed in vacuum and spreads emissions throughout the plane, as can be seen from figure 4.8. For that reason, once again we see that the impact of the medium is such that the fraction of splittings that happen in the soft, large-angle region and in the hard, collinear region is larger in the presence of the medium than it is in vacuum. To help corroborate this observation, we can look at the difference between the radiation pattern of first emissions of quenched and non-quenched jets (figure 4.9).

Figure 4.8: Lund planes filled only with the first splitting of quark-initiated jets, in medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Figure 4.9: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Let us also have a look at the primary Lund planes for quark-initiated jets, in vacuum, depicted in figure 4.10. Once again we can compare these splitting maps with those of gluon-initiated jets (figure 4.4). Just as for the kinematical planes filled only with the first splittings, we notice that the overall features observed for the gluon case remain present while, at the same time, the radiation pattern of quark-initiated jets is more spread out through the plane. In particular, the fraction of emissions in the soft, large-angle region and in the hard, collinear region is larger for these jets than for the gluon case. Once again, the large concentrations of splittings - present both in the C/A and k_t algorithms - migrate slightly towards smaller angles.
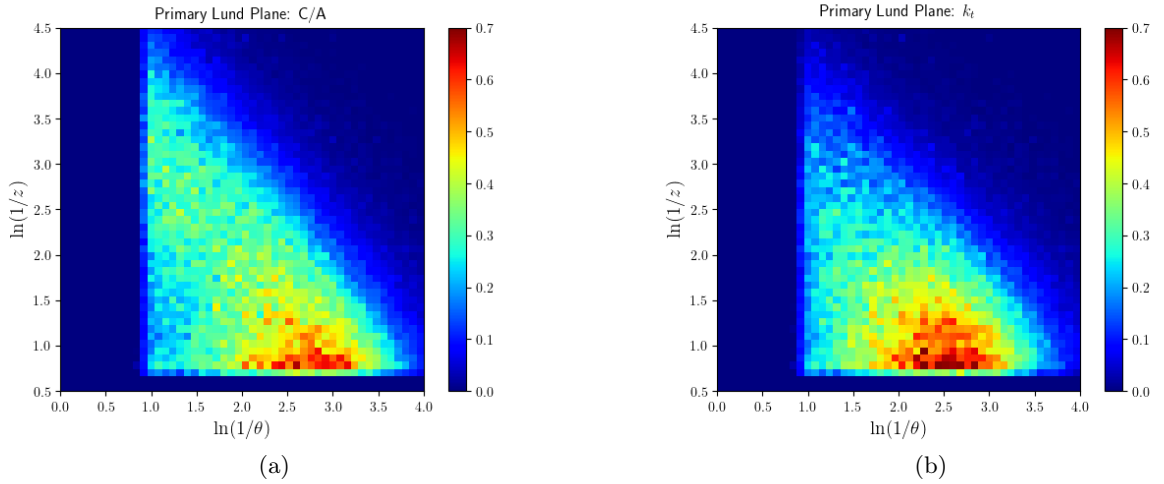
Figure 4.10: Primary Lund planes for quark-initiated jets in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
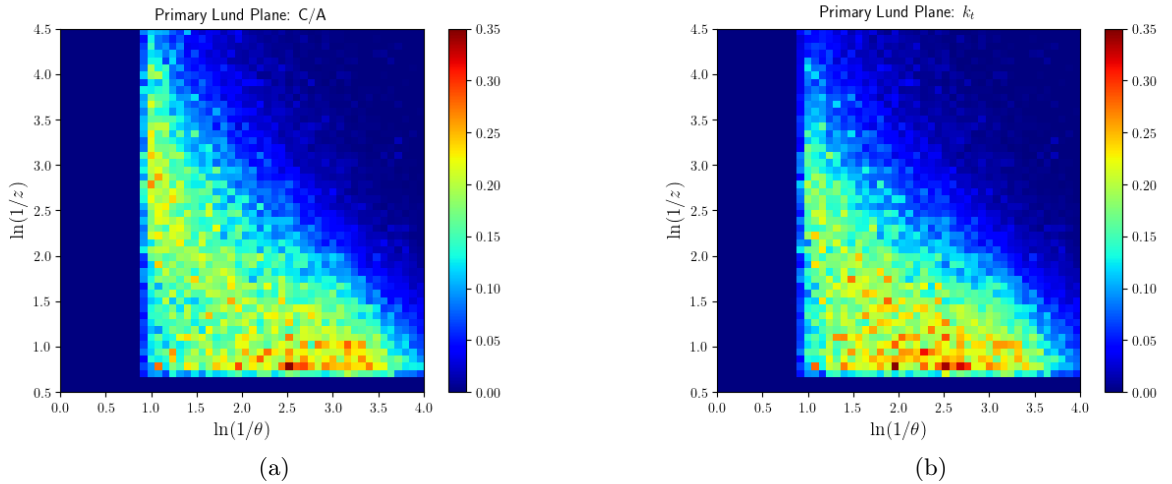
Figure 4.11: Primary Lund planes for in-medium, quark-initiated jets found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

The impact of the medium is the same as noted for gluon-initiated jets, as can be seen by analysing figure 4.11. It visibly destroys much of the observed vacuum structure and disperses the emissions throughout the plane. Once again, it is clear that the fraction of emissions in the soft, large-angle region of the phase space is enhanced in the presence of the medium with respect to vacuum. As we have already noted for gluon-initiated jets, this observation is consistent with the picture that medium-induced radiation manifests as soft quanta at large angles. There is also a migration of the emissions towards the hard, collinear region, though this one is less manifest. These differences can once again be clearly seen by presenting the difference between the radiation patterns in the medium and in vacuum - figure 4.12.

(a)   (b)

Figure 4.12: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Just as for gluon-initiated jets, the overall behaviour and conclusions regarding de full Lund diagram are similar to those associated with the primary Lund plane. For that reason, we chose to put these planes in appendix B, only for consultation of the reader.

This discussion on quark-initiated jets has served to show that the overall features exhibited by the radiation pattern of gluon-initiated jets are present in the quark case, while with some slight migrations and modifications, in particular a larger span of emissions throughout the plane. This is true regardless of the Lund plane construction used and, for that reason, in the following sections we will focus the discussion on gluon-initiated jets. From the behaviour of such jets and with the previous analysis in mind, it should be fairly simple for the reader to understand the expected behaviour for the quark case. The idea behind this omission is simply to reduce the overall bulk size of the discussions and avoid the repetition of the same ideas over and again. Furthermore, with this same motivation in mind, we will abstain from presenting the kinematical Lund planes for the t_F reclustering algorithm, since its behaviour is a middle ground between the other two reclustering algorithms.

4.2   Analysing the $(\ln(1/\theta), \ln(z\theta))$ Lund planes

We can now have a look at the kinematical Lund planes with the vertical axis defined as $\ln(z\theta)$. Since the jets under analysis are exactly the same as in the previous section, the same conclusions as drawn before regarding the influence of the medium should apply, even though the radiation pattern itself will have a different look to it, as a consequence of the use of a different plane definition. Just as before, we can start by analysing the phase space of the first emission within gluon-initiated jets, in vacuum. This is depicted in figure 4.13, which is the analogous of figure 4.1 and, hence, should yield the same information. We note that emissions are constrained into a triangular-shaped region of the plane, as outlined in figure 3.2.

Just as observed with the preceding Lund plane construction, the first emissions obtained from the C/A clustering sequence pool themselves vertically along the line $\ln(1/\theta) \simeq 1.0$, and are extremely rare for $\ln(1/\theta) \gtrsim 2.0$.

Once again, the k_t algorithm yields a much more packed concentration of (first) splittings
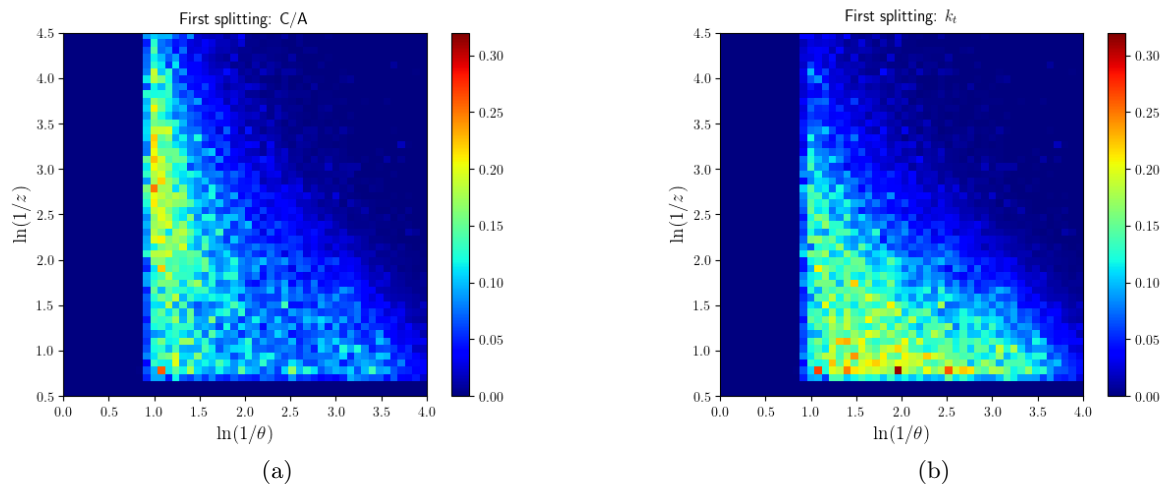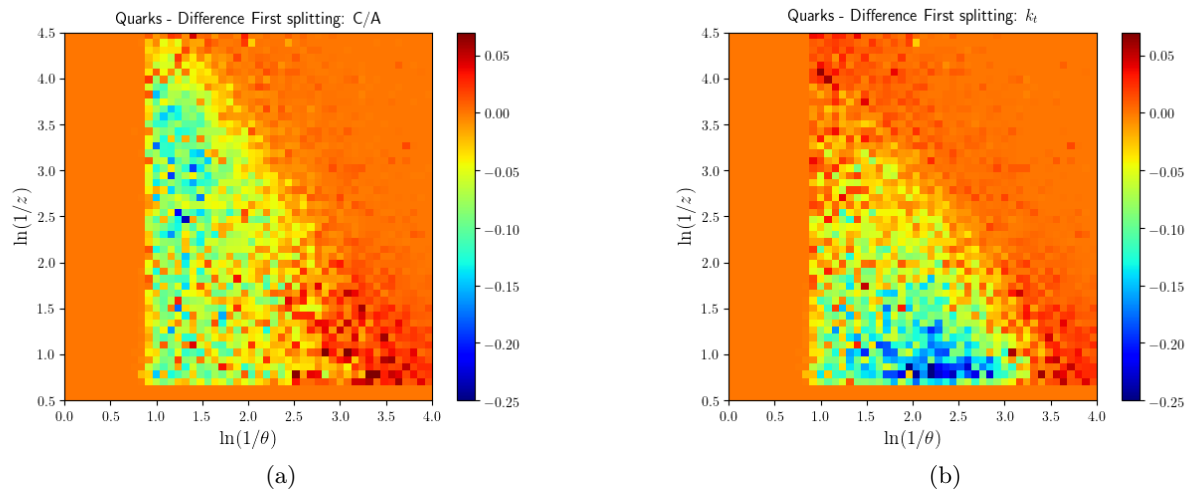
Figure 4.13: Lund planes filled only with the first splitting of gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

than does the C/A algorithm. The large splitting pooling lies near the kinematical diagonal limit thus corresponding, just as before, to hard emissions with z approximately between 0.35 and $1/2$. Moreover, this large splitting concentration is associated with emissions at angles θ roughly between 0.1 and 0.3 (same before).

Essentially, we remark that all the distinct features observed with the $(\ln(1/\theta), \ln(1/z))$ plane definition can now be observed with this construction as well.

Figure 4.14: Lund planes filled only with the first splitting of gluon-initiated jets, in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Recalling our observations from the previous section regarding the impact of the medium over the first emissions, we expect it to spread emissions throughout the plane. In particular, for the C/A reclustering algorithm we anticipate that the fraction of emissions beyond $\ln(1/\theta) \gtrsim 2.5$ is considerably enhanced by the presence of the medium. For the k_t algorithm, we predict that
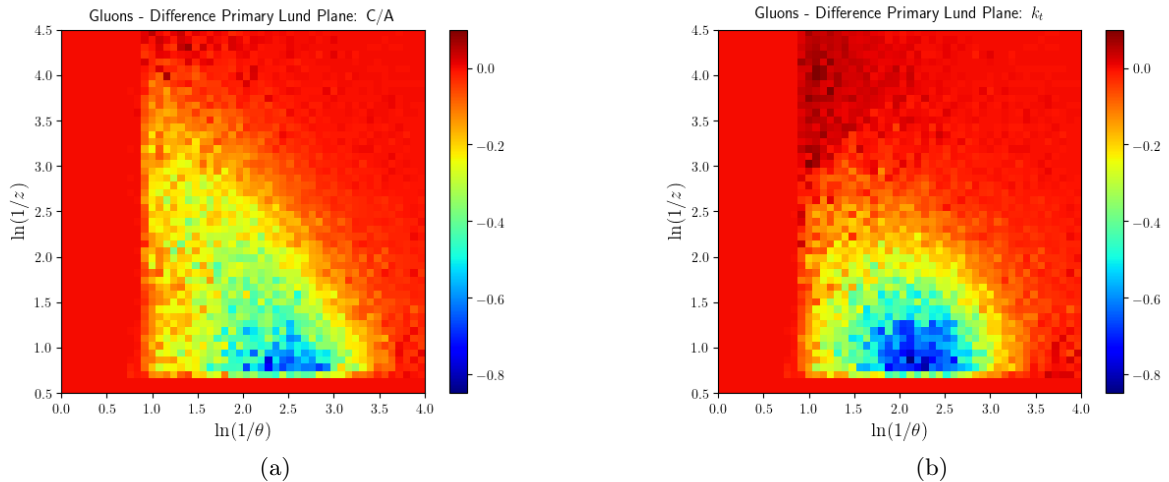
Figure 4.15: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

the fraction of emissions in the soft, large-angle and in the hard, collinear regions of the phase space is larger for jets developed in the medium than for vacuum jets. All of these expectations are corroborated by figures 4.14 and 4.15.

Let us also have a look at the appearance of the primary Lund planes with this definition. These are presented in figure 4.16 for non-quenched, gluon-initiated jets. Inevitably, the radiation patterns of primary emissions of C/A- and k_t-reclustered jets exhibit the exact same features as noted in the $(\ln(1/\theta), \ln(1/z))$ Lund plane construction. Namely, there are large splitting concentrations near the kinematical limit defined by $z = 1/2$. The emission pooling associated with the C/A algorithm lies flatter and closer to the kinematical boundary, while the one obtained with the k_t algorithm grows a little more away from that boundary to regions where $z \simeq 0.3$. Moreover, the C/A splitting concentration lies bounded between $\ln(1/\theta) \simeq 2.0$ and $\ln(1/\theta) \simeq 3.0$, while the one stemming from the k_t algorithm is displaced towards larger angles, within $\ln(1/\theta) \simeq 1.5$ and $\ln(1/\theta) \simeq 2.5$. Finally we also see that the C/A reclustering algorithm yields a larger fraction of emissions in the soft, large-angle region of the plane than does the k_t algorithm. All of these conclusions are the same as the ones that had been drawn in section 4.1.

The impact of the medium is expected to modify the radiation pattern exactly in the same fashion as before. That is, it is expected to destroy a big part of the structure observed in the vacuum primary Lund planes, i.e., to destroy a big part of the splitting concentrations observed near the kinematical limit and disperse the emissions throughout the plane. Particularly, figure 4.17 shows that the fraction of emissions in the soft, large-angle region is enhanced in the presence of the medium, comparatively to the corresponding vacuum fraction. Similarly, the portion of emissions in the hard, collinear region is also larger for quenched jets than it is for jets that evolve in vacuum.

The differences between the radiation pattern, in the primary Lund plane, with and without jet quenching are presented in figure 4.18 and corroborate the picture described above.
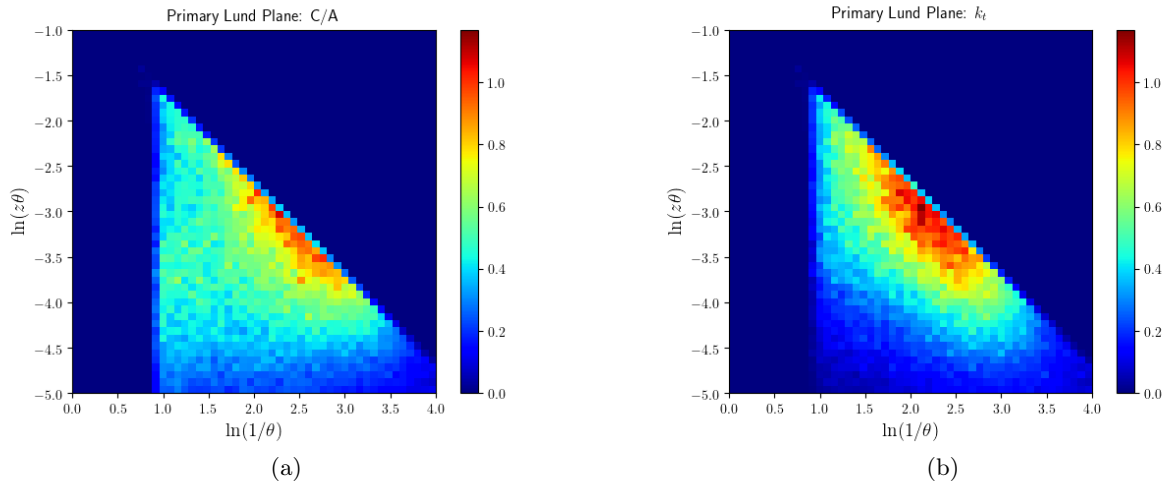
Figure 4.16: Primary Lund planes for gluon-initiated jets, developed in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
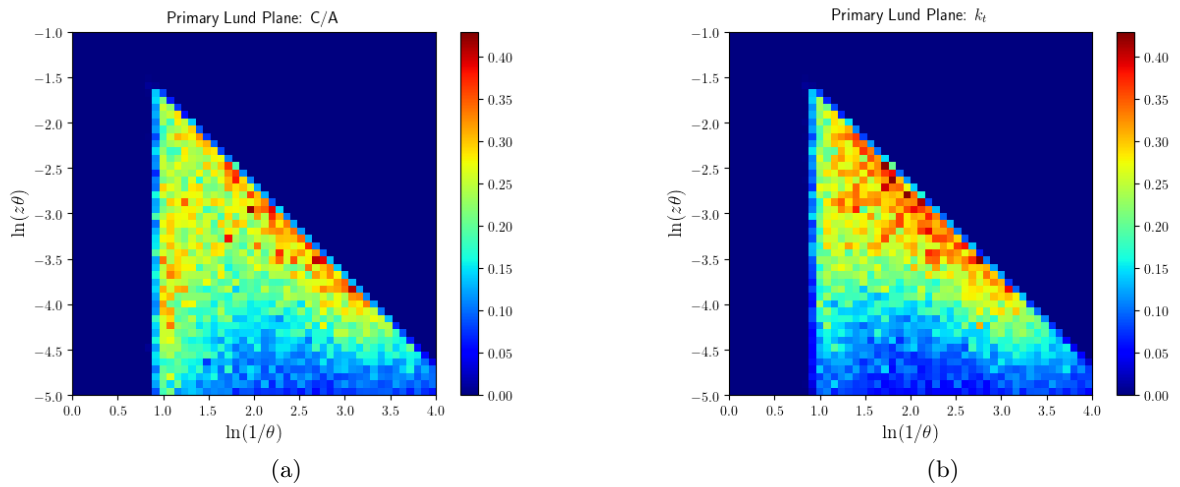
Figure 4.17: Primary Lund planes filled with emissions within gluon-initiated jets, in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
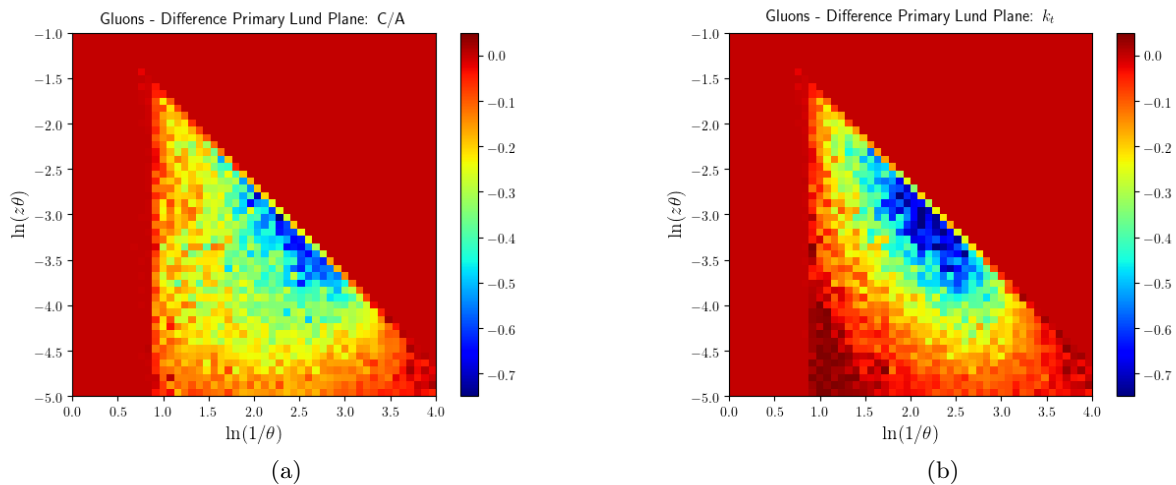
Figure 4.18: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

4.3   Analysing the $(\ln(1/\theta), \ln(k_t))$ Lund planes

We will now have a chance of analysing the kinematical planes defined using $\ln(k_t)$ for the $y-$axis. These planes yield new information with respect to the two previous constructions because they provide us with details about the transverse momentum scale of the splittings, k_t. Just as we have done so far, we start by analysing the Lund planes for the first emission within each jet. These are presented in figure 4.19 for gluon-initiated jets, in vacuum and reclustered with the C/A and k_t algorithms.
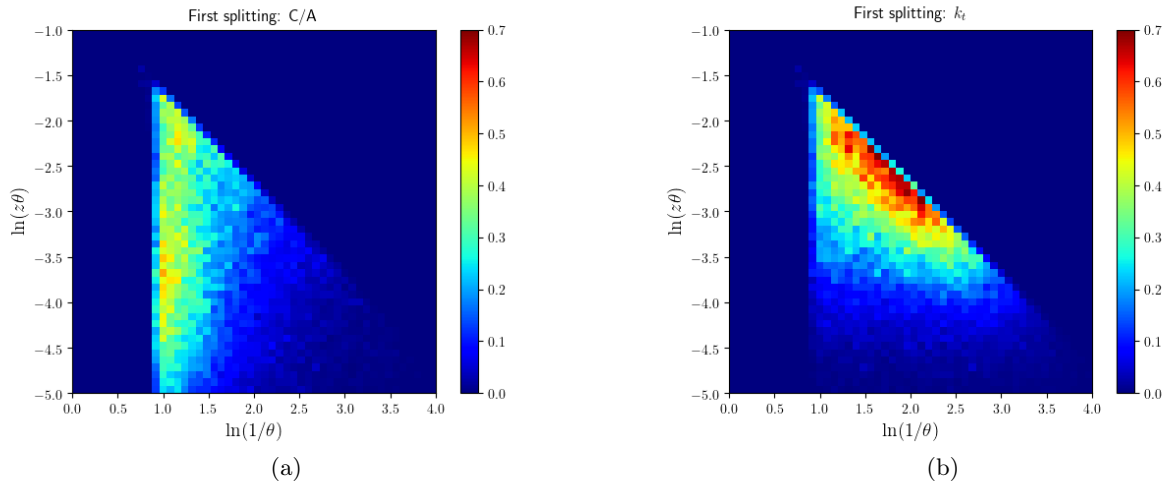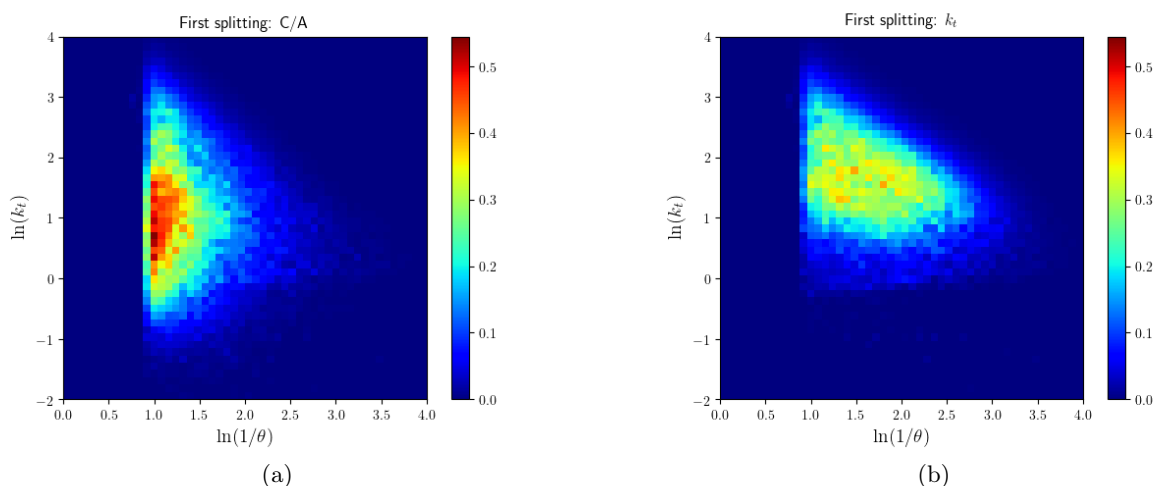

Figure 4.19: Lund planes filled only with the first splitting of gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

One of the most immediate observations is that there appears to be a switch in behaviour between the C/A and k_t algorithms in the sense that now the C/A originates a larger concentra-

tion of (first) emissions in a specific kinematical region, while the k_t algorithm yields a much more uniform distribution. This behavioural reversal is a remarkable observation. It means that when it comes to using the momentum sharing fraction z, the k_t algorithm identifies a large fraction of self-similar first splittings with identical z values, but whose transverse momentum values k_t are distinct amongst themselves, translating into a less packed distribution in the current Lund plane construction. On the other hand, by using the momentum fraction z with the C/A algorithm we see an almost uniform distribution of emissions among different values of z, but which turn out to correspond to splittings with similar k_t values, that crowd together in the present kinematical Lund plane construction. To phrase the previous observations more succinctly, by changing the kinematical variable used in the $y-$axis from $\ln(1/z)$ (or $\ln(z\theta)$) to $\ln(k_t)$, a uniform distribution of first emissions is transformed into a localisation of those same emissions, and vice-versa.

Further analysis of the radiation pattern reconstructed by the C/A algorithm shows that, once more, first emissions pool very tightly around $\ln(1/\theta) \simeq 1.0$ and are rare beyond $\ln(1/\theta) \gtrsim 2.0$. The large splitting concentration that is now observed with this algorithm corresponds to a transverse momentum of the first splitting approximately between $\ln(k_t) \simeq 0.5 \Leftrightarrow k_t \simeq 1.6\,\text{GeV}$ and $\ln(k_t) \simeq 1.5 \Leftrightarrow k_t \simeq 4.5\,\text{GeV}$. Finally, the region below $\ln(k_t) \lesssim -0.5 \Leftrightarrow k_t \lesssim 0.6\,\text{GeV}$ is (nearly) deprived of any emissions.

The analysis of the kinematical Lund plane obtained from the first splitting of k_t-reclustered jets shows that emissions below $\ln(k_t) \lesssim 1.0$ are rare, i.e., the transverse momentum of the first splitting is sparsely below $k_t \simeq 2.7\,\text{GeV}$. The noteworthy observation here is that, overall, the transverse momentum scale of the first splitting is larger for k_t-reclustered jets than for C/A-reclustered ones, as is to be expected given the distance measures of each of these algorithms.

Let us now analyse what is the impact of the medium in the radiation pattern of jets as depicted in this Lund plane construction (figure 4.20).
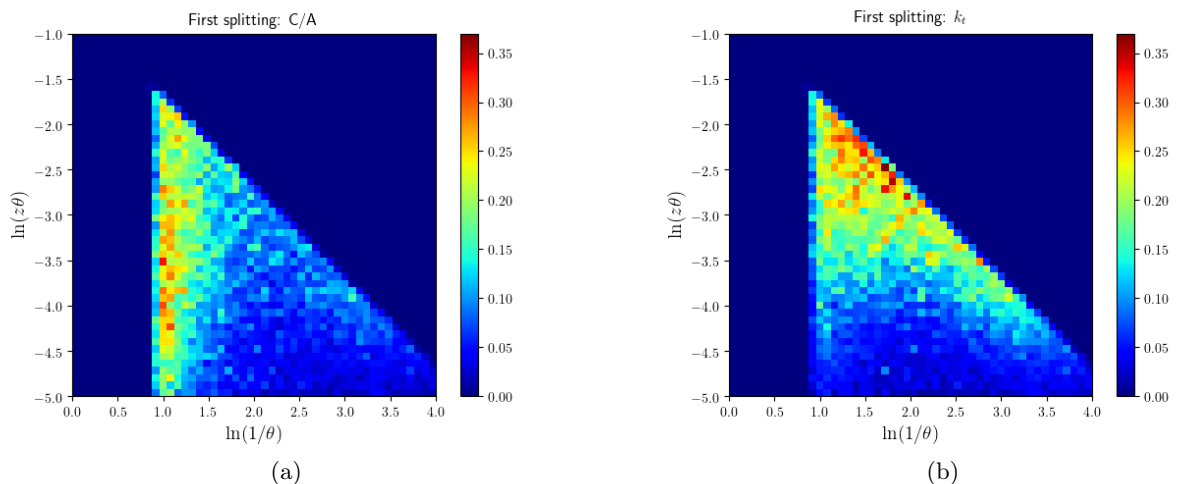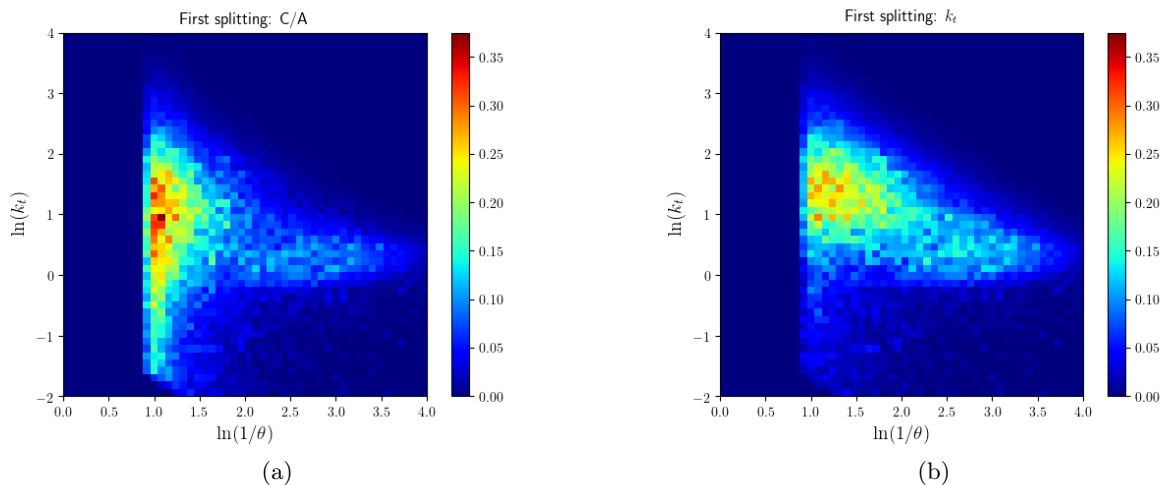

Figure 4.20: Lund planes filled only with the first splitting of gluon-initiated jets, in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

The medium-induced behavioural modifications are similar between the two algorithms. Just as before, the larger emission concentrations tend to be (partially) destroyed and radiation is spread throughout the plane. Particular to this case is the fact that the radiation stretches along two prominent prongs. The first of those lies in the vertical direction, along $\ln(1/\theta) \simeq 1.0$, downwards to regions where $\ln(k_t) \simeq -1.5 \Leftrightarrow k_t \simeq 0.2\,\text{GeV}$. This prong is especially prominent
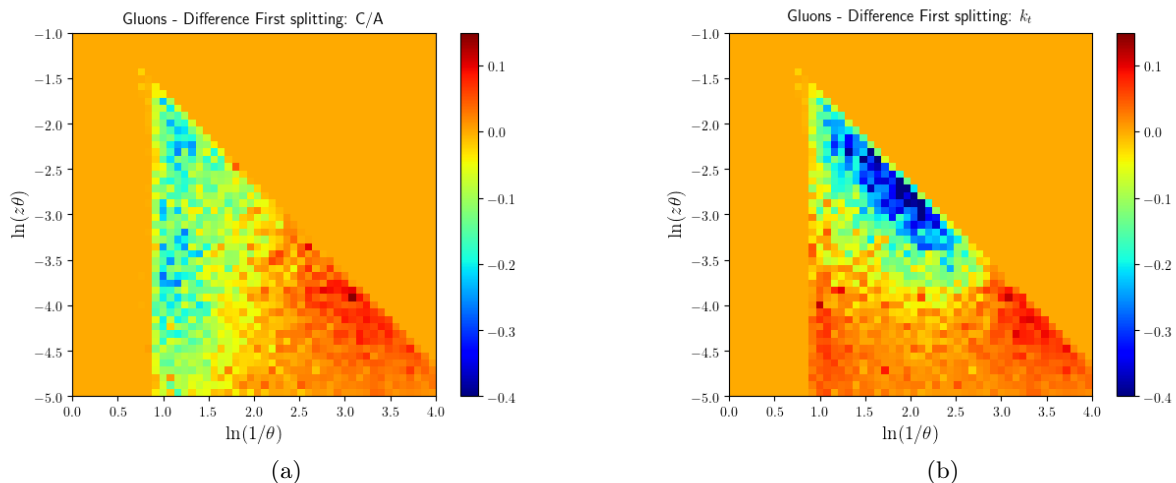
Figure 4.21: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

for the C/A algorithm. The second prong stretches horizontally, approximately along the line $\ln(k_t) \simeq 0.5 \Leftrightarrow k_t \simeq 1.6\,\mathrm{GeV}$, towards the collinear region where $\ln(1/\theta) \simeq 4.0 \Leftrightarrow \theta \simeq 0.02$. This one is a more striking feature for k_t-reclustered jets. These prongs highlight once again the fact that the presence of the medium enhances radiation in the soft, large-angle region of phase space (vertical prong) and also in the hard, collinear region (horizontal prong). They make for particularly conspicuous features when plotting the difference between the radiation pattern of quenched and non-quenched jets, as can be seen in figure 4.21. In particular, it is clear that the vertical prong is more dominant with the C/A algorithm while the horizontal prong is more prominent for k_t-reclustered jets. The enhancement of the soft, large-angle region when using the C/A reclustering algorithm is relevant because that region did not suffer such a large enhancement when using the other two Lund plane constructions. As such, the use of this kinematical definition for the vertical axis appears to have brought some sensitivity towards that region of the phase space of emissions for C/A-reclustered jets. Once again we reinforce the important idea that whenever we mention enhancements (or suppressions) of specific phase space regions in the presence of the QGP, these refer to those regions becoming more (or less) prominent per jet, and not in absolute terms.

The migration of the (first) splittings into the two prongs should correspond to an increase of the formation time of such splittings (recall figure 3.3), something that should be legitimised in the next section by the $(\ln(1/\theta), \ln(t'_F))$ Lund planes.

Given the interesting new behaviours observed for the radiation pattern of the first emissions, we undoubtedly anticipate exciting findings for the primary Lund planes. These are presented in figure 4.22, for gluon-initiated jets, in vacuum. We see that the radiation pattern is remarkably similar between the two algorithms: the overall shape and useful region are identical. In light of that, the seeming attenuation of the primary Lund plane obtained for k_t-reclustered jets is probably due to the overall shorter length of their primary branch, with respect to that of C/A-reclustered ones (notably 20% shorter in vacuum, as stated in section 4.1). For the C/A algorithm, the larger splitting concentration appears to pool between $\ln(k_t) \simeq 0.0 \Leftrightarrow k_t \simeq 1.0\,\mathrm{GeV}$ and $\ln(k_t) \simeq 1.0 \Leftrightarrow k_t \simeq 2.7\,\mathrm{GeV}$, with the angle of emissions spanning approximately between $\ln(1/\theta) \simeq 1.5$ and $\ln(1/\theta) \simeq 3.0$. On the other hand, in the primary Lund plane

for k_t-reclustered jets, the splitting concentration appears to have a dominant fraction between $\ln(k_t) \simeq 0.5 \Leftrightarrow k_t \simeq 1.6\,\text{GeV}$ and $\ln(k_t) \simeq 1.5 \Leftrightarrow k_t \simeq 4.5\,\text{GeV}$ (i.e., at larger transverse momenta than for C/A-reclustered jets) and with an angular range from $\ln(1/\theta) \simeq 1.5$ to $\ln(1/\theta) \simeq 2.5$.

Figure 4.22: Primary Lund planes for gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Figure 4.23: Primary Lund planes for gluon-initiated jets, developed in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

By now, the effect of the medium is remarkably well established. We anticipate that it will (partially) destroy the large pools of splittings set at the centre of the primary Lund planes of vacuum jets, dispersing emissions throughout the plane. In particular, we expect two prongs: one horizontal around $\ln(k_t) \simeq 0.5$ and the other vertical about $\ln(1/\theta) \simeq 1.0$. Indeed these predictions are corroborated by figure 4.23. Once again, the migration of splittings towards these two prongs indicates that the presence of the QGP delays emissions, i.e., enhances the

fraction of splittings with larger formation times.[2]

The differences in the radiation pattern of quenched and non-quenched jets are presented in figure 4.24, corroborating our previous analysis.
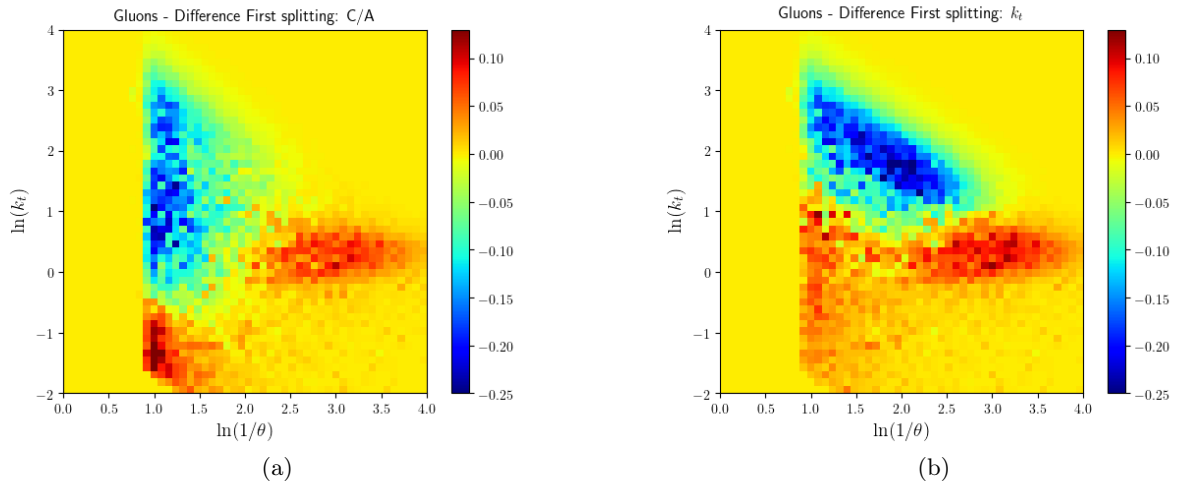

Figure 4.24: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

4.4   Analysing the $(\ln(1/\theta), \ln(t'_F))$ Lund planes

Finally, we can turn to the analysis of the jet radiation pattern in the Lund plane defined using the formation time of the splittings. We will comply with the same structure that has been followed so far and will start by analysing the kinematical Lund planes filled only with the first emission within each gluon-initiated jet, in vacuum. These planes are presented in figure 4.25, for the C/A and k_t reclustering algorithms.

All the conclusions we have drawn so far regarding the first emission of C/A-reclustered jets are valid for this novel Lund plane construction (as is to be expected). In particular, the (first) emissions pool around $\ln(1/\theta) \simeq 1.0$ and are extremely rare beyond $\ln(1/\theta) \gtrsim 2.0$. In the case of the k_t reclustering algorithm, the splittings are much more evenly distributed through the plane (in similarity with the behaviour for the $(\ln(1/\theta), \ln(k_t))$ kinematical Lund plane construction). Also noteworthy is the fact that first emissions are approximately restricted to the region $\ln(t'_F) \lesssim 2.5 \Leftrightarrow t'_F \lesssim 2.5 \, \mathrm{fm}/c$.

We expect the medium to span emissions throughout the plane, specifically, we anticipate that the radiation pattern of quenched jets in this kinematical Lund plane extends upwards, such that there is an enhancement of the fraction of emissions with larger formation times, when compared to the pattern of emissions in vacuum. This is observed in figure 4.26. In particular, we see the formation of an extended vertical prong along the line $\ln(1/\theta) \simeq 1.0$, i.e., large-angled emissions are delayed by the presence of the QGP. Besides that, there is a second prong which extends along a diagonal line of slope $m \simeq 2$. Just as the first prong, this second one, also indicates that the medium delays emissions but further implies that it also enhances the fraction

---

[2]Here it is important to note that we shall use the expression "delay emissions" in the sense of an enhancement of the phase space region associated with larger formation times. To make an actual full correspondence of this observation to a physical time delay of the splittings a more extensive study is needed.

of hard, collinear splittings (as we have already seen in the other Lund plane definitions). Overall, we note that radiation is now rare only beyond $\ln(t'_F) \gtrsim 4.0 \Leftrightarrow t'_F \gtrsim 10.8 \, \text{fm}/c$, a value which is approximately four times larger than the one obtained for vacuum jets.
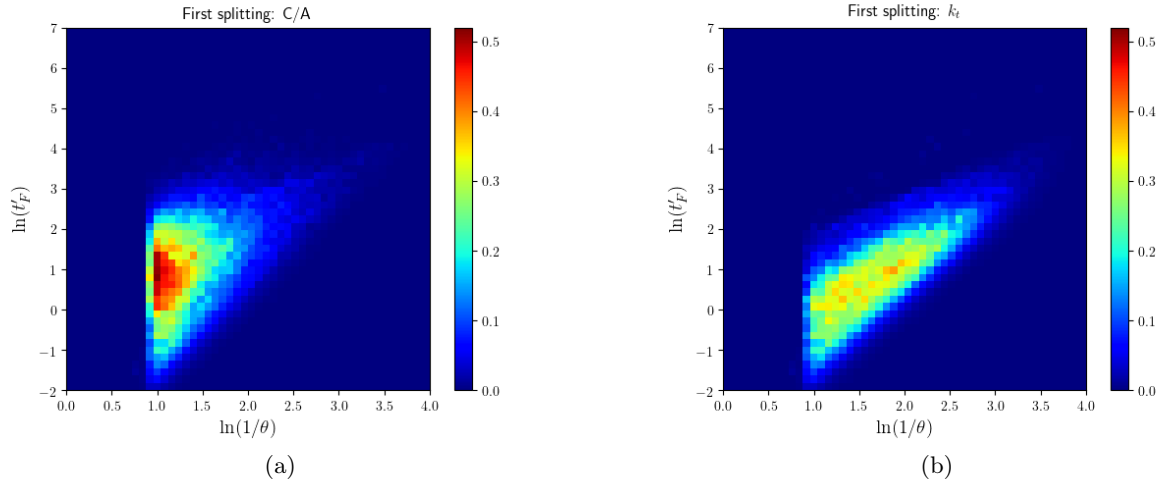
Figure 4.25: Lund planes filled only with the first splitting of gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
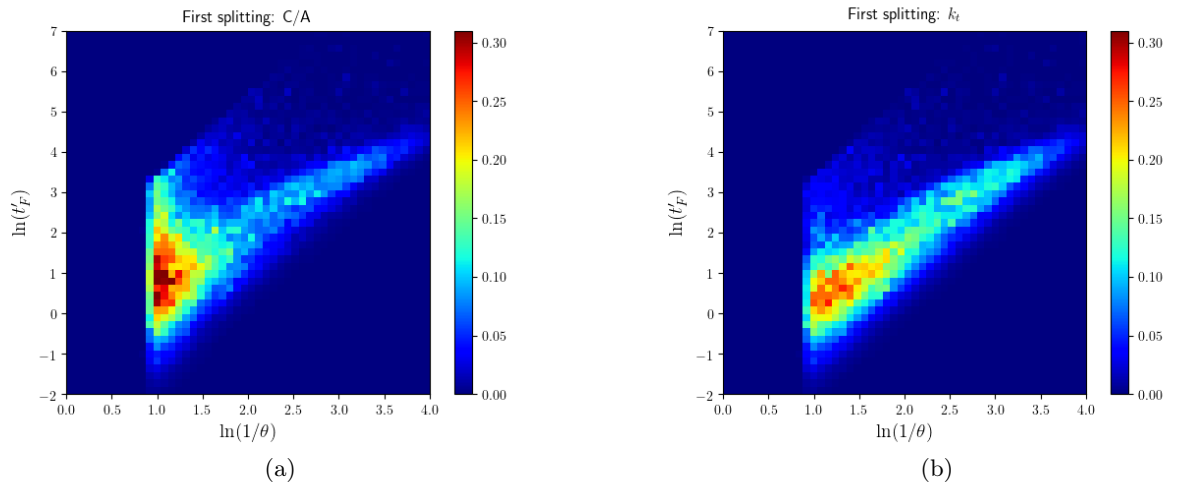
Figure 4.26: Lund planes filled only with the first splitting of gluon-initiated jets, in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Figure 4.27 depicts the difference between the radiation pattern of quenched and non-quenched jets. That the medium promotes late (first) emissions with respect to vacuum radiation is apparent from the Lund planes presented.
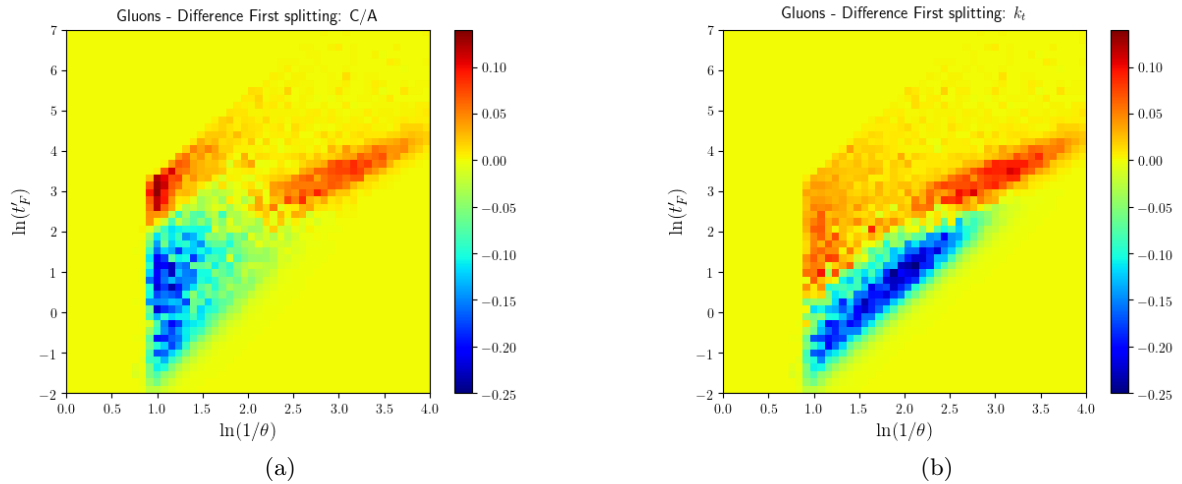
Figure 4.27: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Let us now analyse what happens along the primary branch of gluon-initiated jets. As always, we start by having a look at the picture in vacuum, which is illustrated in figure 4.28. Remarkably, although the emissions can, in principle, extend all the way up through the plane (recall figure 3.4), this does not happen. From the splitting maps presented we see that large-angled radiation (along $\ln(1/\theta) \simeq 1.0$) is bound to short formation times: $\ln(t'_F) \lesssim 2.5 \Leftrightarrow t'_F \lesssim 2.5\,\mathrm{fm}/c$. As we consider the whole set of primary emissions, we see that overall smaller-angled radiation yields larger formation times. There seems to be a strong concentration of splittings oriented diagonally around a line of slope $m \simeq 2$ and indicating a strong correlation between the angle of the emission and its formation time.
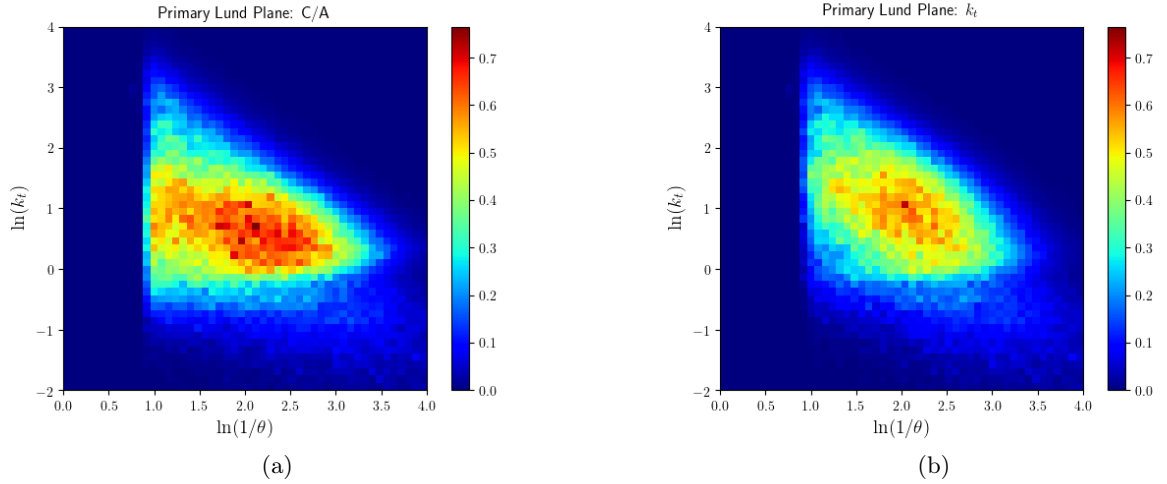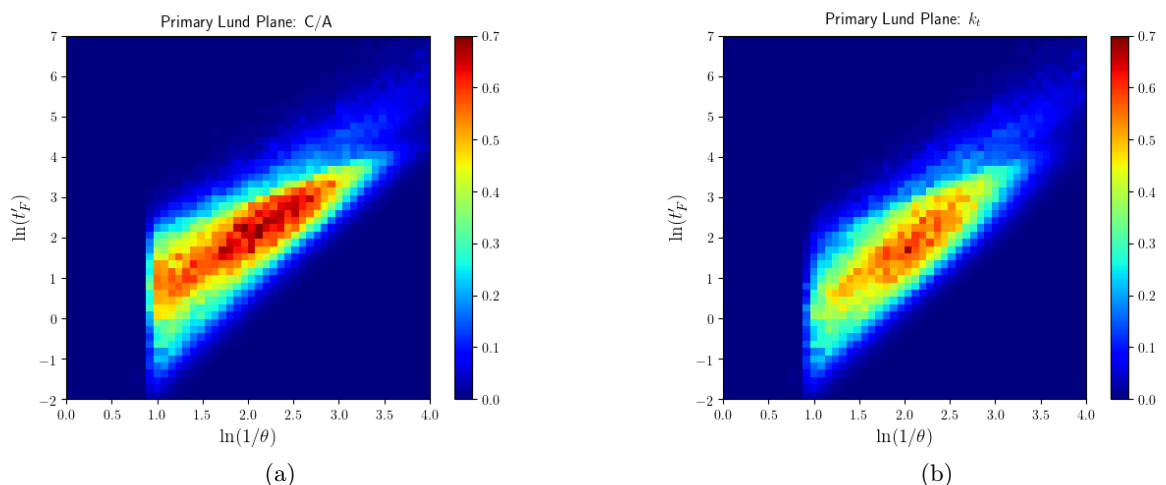

Figure 4.28: Primary Lund planes for gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
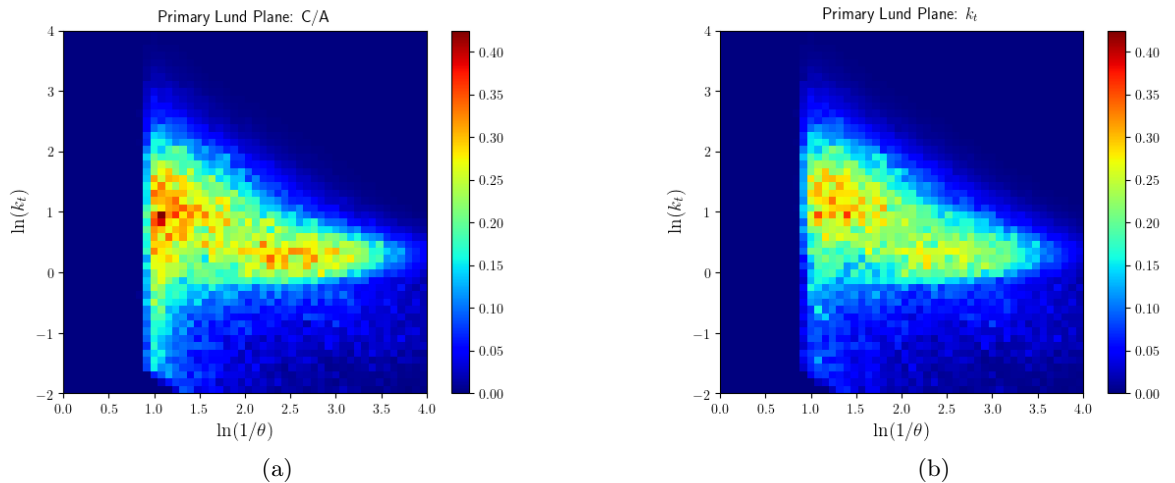
The effect of the medium over the primary emissions is presented in figure 4.29. From

the analysis of the radiation pattern of quenched jets it is particularly clear that the medium undoubtedly promotes delayed emissions (dominantly) at large angles. This is again visible by the prominent vertical prong along the line $\ln(1/\theta) \simeq 1.0$. The diagonal prong is slightly subdued, though there still seems to be a slight increase of the fraction of emissions with larger formation times in the collinear region.

　　The difference between the radiation patterns with and without jet quenching is presented in figure 4.30, for the primary Lund plane of gluon-initiated jets. The enhancement of the regions associated with larger formation times is plain to be seen. Indeed, as it has been stated so far, the medium spreads radiation through the plane. However, that dispersion is not at random: emissions are spanned such that there is a systematic migration to larger formation times.
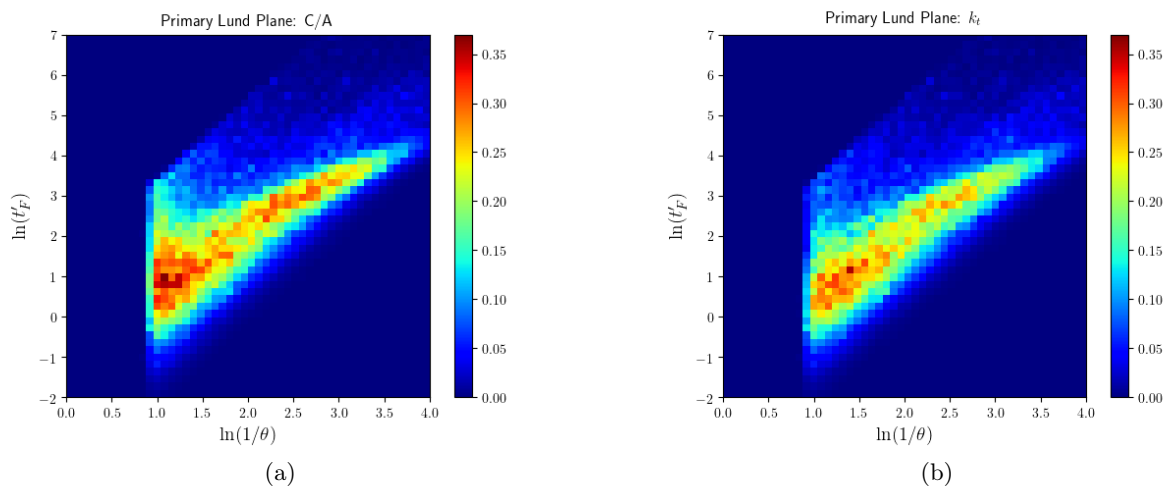
Figure 4.29: Primary Lund planes for gluon-initiated jets, developed in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
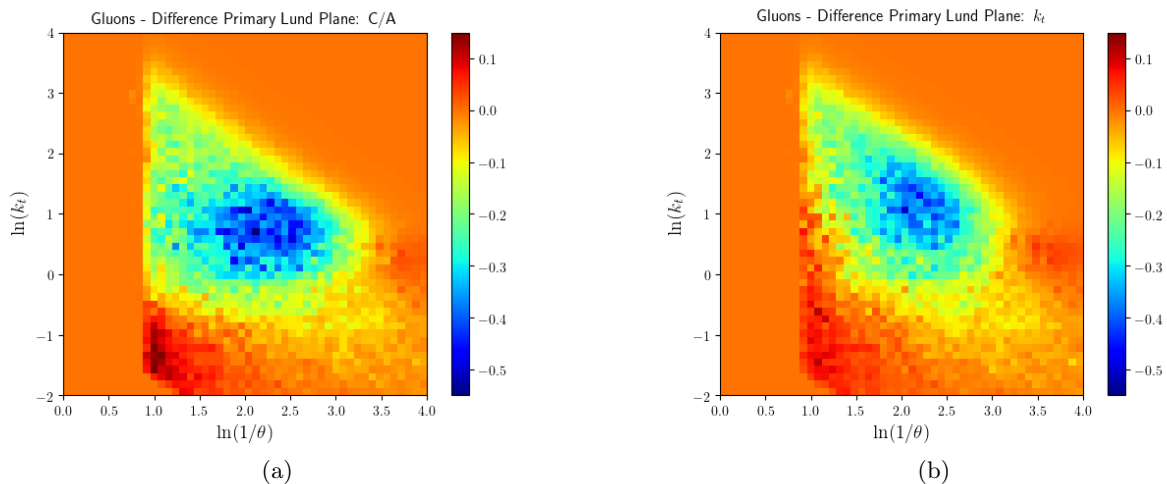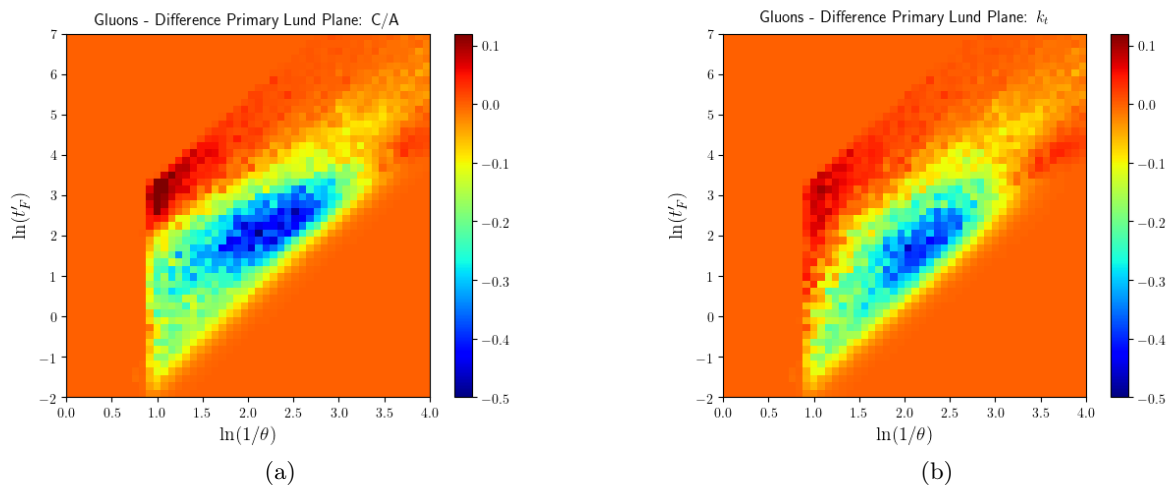

Figure 4.30: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

4.5   Comments on t_F vs t'_F

We had the chance of analysing the novel Lund planes defined by $(\ln(1/\theta), \ln(t'_F))$. In section 1.5, we have seen that the formation time t'_F is a boost invariant quantity which is a proxy for the (often used) formation time t_F (equation (1.15)). We have shown this by deriving equation (1.18) for t'_F, resorting to reasonable approximations and assumptions. To further support this statement that the both formation times can be used, we present in figure 4.31, the primary Lund planes obtained for quenched, gluon-initiated jets using a construction $(\ln(1/\theta), \ln(t_F))$. The comparison between the radiation patterns in this kinematical plane and those in figure 4.29 testifies that the two physical quantities t_F and t'_F are, indeed, equivalent to one another or, in other words, that t'_F is a good proxy for the formation time t_F.
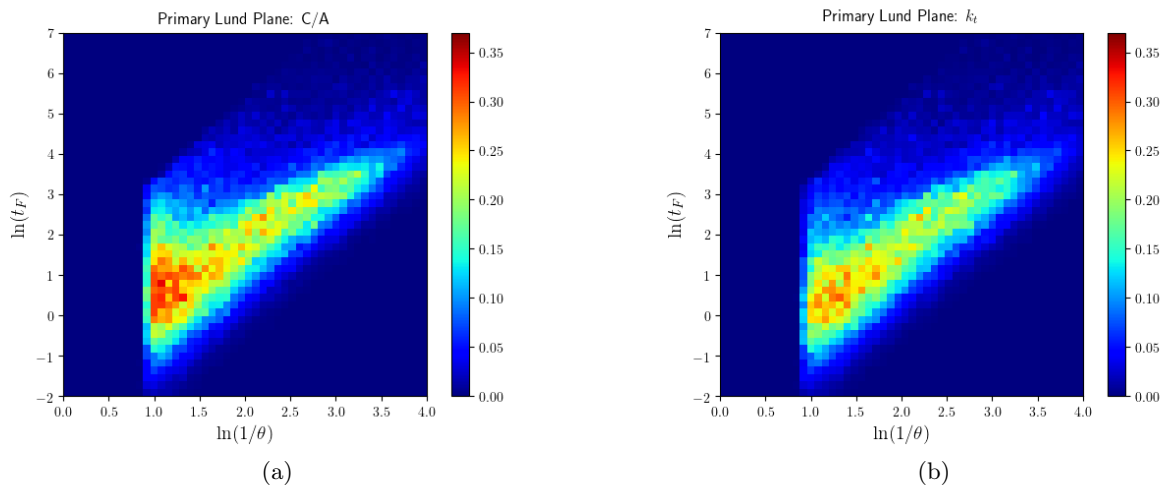
Figure 4.31: Primary Lund planes for gluon-initiated jets, developed in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

4.6   The explicit path of emissions in the primary Lund plane

In this section we present the explicit, average paths of the primary emissions obtained with the same sample of jets as used to fill the kinematical Lund planes in the preceding five sections. To produce these paths we take the coordinates of each primary emission within each jet and average the first emissions, the second emissions, the third emissions, etc. of that sample. Since all the considered jets will have at least one splitting, they all contribute to this procedure with at least one coordinate point - the coordinate point corresponding to the first splitting. Subsequent emissions have a successively smaller number of available data points, i.e., the statistics gets lower as we move along the primary branch. We consider only emissions for which more than 500 coordinate points are at our disposal for the averaging procedure, that is, emissions with the contribution of at least 500 jets. Splittings for which this threshold is not met are not used to build the paths.

The idea behind building these explicit paths of emissions along the primary branch is to show that they are a distinct feature between quenched and non-quenched jets. This is expected given the differences recorded in sections 4.1 to 4.4 between the radiation patterns of vacuum and in-medium jets, but here we decided to explicitly verify that expectation.

We draw the paths of emissions for three of the kinematical Lund plane constructions:

$(\ln(1/\theta), \ln(1/z))$, $(\ln(1/\theta), \ln(k_t))$ and $(\ln(1/\theta), \ln(t'_F))$, since each of these allows us to explore different path features and thus provides us with additional information. For the first case (the $(\ln(1/\theta), \ln(1/z))$ plane), we will consider (once again) both gluon- and quark-initiated jets, as to analyse the repercussions that the nature of the original parton has on the primary emissions' path. Afterwards, we focus on gluon-initiated jets to lighten the comparisons between the different kinematical Lund plane definitions and to reduce the bulk size of this section.

4.6.1   Explicit path in the $(\ln(1/\theta), \ln(1/z))$ primary Lund plane

Following the structure that has been set forth, we start by examining the average paths of primary emissions of the jets in the kinematical Lund plane construction $(\ln(1/\theta), \ln(1/z))$. We will begin with the analysis of gluon-initiated jets and afterwards discuss quark-initiated ones, comparing both situations to establish general differences and similarities. Understanding how the two different jet origins compare to one another allows us to later focus only on the case of gluon-initiated jets.

4.6.1.1   Gluon-initiated jets

The average paths of the splittings along the primary Lund plane for gluon-initiated jets in vacuum are presented in figure 4.32 for the three different reclustering algorithms: C/A, k_t and t_F.

The average sequence of primary emissions for C/A-reclustered jets is, as expected, angular-ordered such that the angle of successive emissions progressively decreases. The (average) first emission is large angled belonging to the region between $\ln(1/\theta) \simeq 1.0$ and $\ln(1/\theta) \simeq 1.5$. This is consistent with the observation of figure 4.1. Furthermore, we see that this first emission corresponds to the softest splitting in the sequence and is followed by successively harder ones, something that deviates from the strict double-logarithmic ordering and might indicate that the C/A algorithm is not properly reconstructing the history of the shower.[3]

On the other hand, jets reclustered with the k_t algorithm exhibit the same angular ordering but are simultaneously ordered such that harder emissions happen first, followed by successively softer ones. As such, this algorithm appears to better reconstruct the development of jets than does the C/A algorithm. Moreover, it is also worth noting that k_t-reclustered jets yield fewer emissions along the primary branch: seven emission data points against the nine observed with the C/A algorithm. This observation is consistent with the fact that the primary Lund plane obtained from the clustering sequence of the k_t algorithm is filled with roughly 20% less splittings than that of C/A-reclustered jets (as was mentioned in section 4.1). Finally, we observe that, on average, the first emission obtained with the k_t algorithm is displaced (with respect to that of the C/A) towards a region associated with harder and smaller-angled splittings, a picture that is, once again, consistent with figure 4.1.

Finally, the t_F algorithm originates a path of emissions that appears to be almost an average between those of the C/A and k_t algorithms. In particular, the first emission is not as large-angled as that yielded by the C/A algorithm but the migration towards smaller angles is shorter than the one registered by the k_t algorithm. Similarly, this first emission is not as hard as the one obtained with the latter but not as soft as the one from the former either. The sequence of splittings for t_F-reclustered jets is nearly horizontally flat, meaning that there is almost no change in the overall splitting hardness along the primary branch. This, once again, appears to lie as a middle ground to the fact that in C/A-reclustered jets the hardness of emissions progressively increases while the opposite behaviour is recorded for k_t-reclustered jets. The fact

---

[3]A similar observation has been made in the workshop https://indico.gsi.de/event/8938/, for which there are as of yet no published results.

that the t_F algorithm yields a compromise between the other two algorithms is also in agreement with our observations from previous sections. It further promotes the idea that the k_t algorithm may be the best choice (amongst the three) to reconstruct the history of the MC parton shower (at least for the vacuum situation).

Figure 4.32: Average path of emissions along the primary Lund plane for gluon-initiated jets, in vacuum, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

We are particularly interested in understanding in what way the presence of the medium and, hence, jet quenching effects, influences these paths of primary emissions. To that end we can have a look at figure 4.33. Globally, the medium shifts the splittings upwards, meaning that, on average, it appears to induce an enhancement of soft radiation. Moreover, the (average) first emission happens at a smaller angle. These observations are common to all three reclustering algorithms and concordant with the results from preceding sections.

Figure 4.33: Average path of emissions along the primary Lund plane for quenched gluon-initiated jets, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

4.6.1.2   Quark-initiated jets

(a)   (b)   (c)

Figure 4.34: Average path of emissions along the primary Lund plane for quark-initiated jets, in vacuum, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

We can now have a look at the paths of primary emissions for quark-initiated jets and compare the results with those obtained for the gluon case. The paths of primary splittings for these jets, in vacuum, and reclustered with the C/A, k_t and t_F algorithms are presented in figure 4.34. The common difference which is observed throughout the three sequences is that the first two emissions of quark-initiated jets are displaced towards smaller angles (with respect to the gluon case) - a feature which is in comforting agreement with the observations done for quark-initiated jets in section 4.1, namely, in figures 4.7 and 4.10.

Another interesting observation here is that the sequence of the k_t algorithm appears to deviate from the the strict double-logarithmic ordering of QCD (after the third emission), something that did not happen in the gluon case. However, this apparent deviation may stem from the reduced statistics of quark-initiated jets, which introduces more fluctuations to the average behaviour of the emissions (especially later ones). To sustain this hypothesis we should calculate the standard deviation associated with each (average) point in the splitting map. Due to the shortage of time, we chose to leave this as an improvement for future work. The reduced statistics of quark-initiated jets is also a further motivation to focus (as we will) on the case of jets in which the original parton is a gluon.

Despite the apparent deviation of the sequence of primary emissions of k_t-reclustered jets from the double-logarithmic ordering of QCD (which, as we stated, is possibly a result of the reduced statistics), this algorithm still appears to exhibit the best behaviour in reconstructing the history of the shower (just as for gluon-initiated jets).

The differences between the path of quenched and non-quenched quark-initiated jets are analogous to the ones observed in gluon jets. Particularly, we see that the first emission is slightly shifted towards smaller angles and that the (average) splittings are overall softer (figure 4.35).

From now on we will focus the analysis on the explicit splitting paths of gluon-initiated jets.

Figure 4.35: Average path of emissions along the primary Lund plane for quenched quark-initiated jets, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

4.6.2 Explicit path in the $(\ln(1/\theta), \ln(k_t))$ primary Lund plane

Analysing the explicit sequence of primary emissions in the $(\ln(1/\theta), \ln(k_t))$ kinematical Lund plane will provide us with additional information that could not be gathered from the previously used construction. In particular, we gain some insight into the evolution of the transverse momentum of the splittings, k_t.

Figure 4.36 presents the paths of primary emissions for gluon-initiated jets in vacuum. Interestingly, from this figure we note the path is approximately linear for all three reclustering algorithms. Moreover, the transverse momentum of the splittings decreases as we move along the primary branch, as we would expect from the natural (vacuum) QCD ordering.
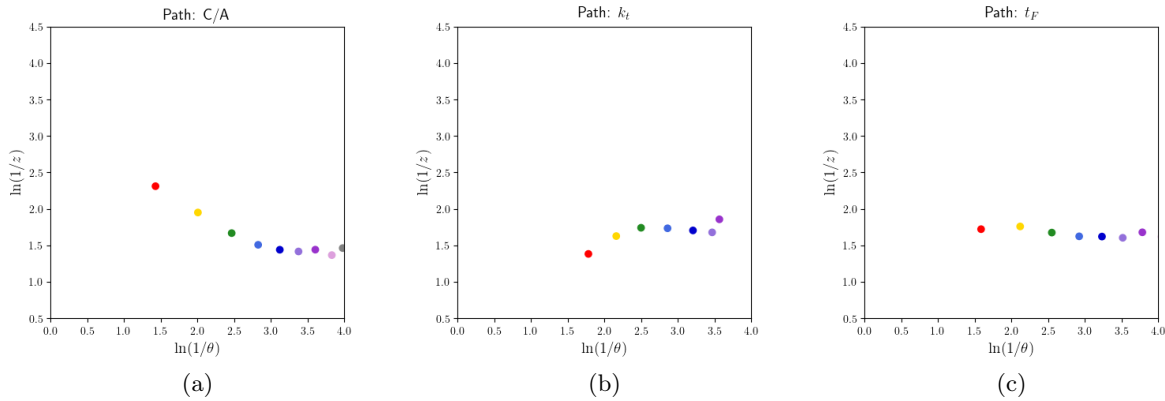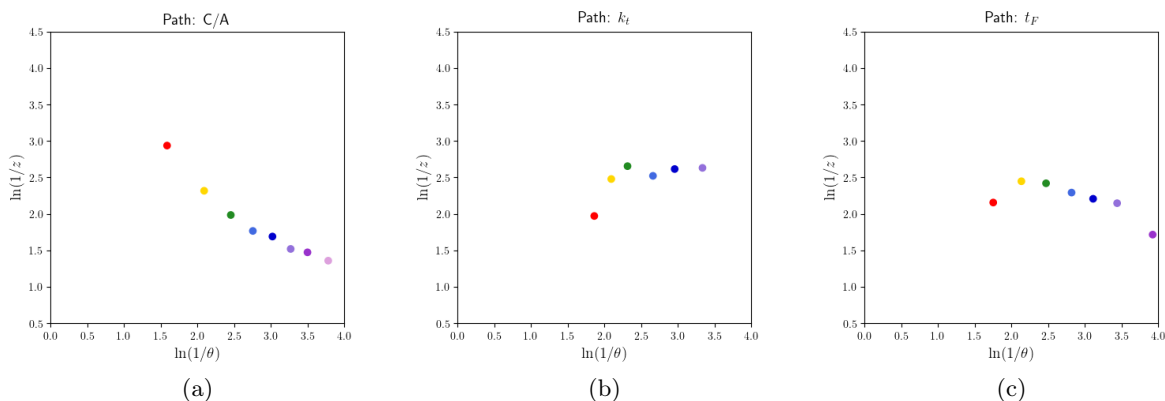

Figure 4.36: Average path of emissions along the primary Lund plane for gluon-initiated jets, in vacuum, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

We notice that the sequence of primary emissions of k_t-reclustered jets yields the (average) first splitting with largest transverse momentum: $k_t \simeq 5.3\,\mathrm{GeV}$, followed by $k_t \simeq 4.6\,\mathrm{GeV}$ for

Figure 4.37: Average path of emissions along the primary Lund plane for gluon-initiated jets, in the presence of the medium, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

t_F-reclustered jets and $k_t \simeq 3.0\,\text{GeV}$ for C/A-reclustered ones. These results are in agreement with our analysis of figure 4.19 in section 4.3. On the other hand, the seventh (and last) primary emission of the k_t clustering sequence has the same transverse momentum value as that same emission in the t_F algorithm - $k_t \simeq 0.2\,\text{GeV}$ - while the ninth and last (average) emission of C/A-reclustered jets has $k_t \simeq 0.4\,\text{GeV}$.

It will be interesting to see how these results are altered by the occurrence of jet quenching effects. The paths of primary emissions in the presence of the medium are depicted in figure 4.37 from which it is apparent that, overall, the medium lowers the transverse momentum scale of the splittings. For instance, the (average) first emissions of C/A- and k_t- reclustered jets have transverse momenta values which are approximately three times smaller than the corresponding vacuum values: respectively, $k_t \simeq 0.9\,\text{GeV}$ and $k_t \simeq 1.8\,\text{GeV}$. This is comfortingly consistent with the observations made in section 4.3.

Finally, these paths also yield information about the formation time, which clearly increases as we progress along the primary branch, meaning that the emissions take progressively longer time to occur. This is true for all algorithms and it is consistent with the behaviour we expect. We shall corroborate this observation in the following subsection, where the $y-$coordinate of the emissions in the kinematical plane is explicitly related to the formation time: $y = \ln(t'_F)$.

4.6.3   Explicit path in the $(\ln(1/\theta), \ln(t'_F))$ primary Lund plane

Similarly to what has been done in the previous two subsections, we will start by analysing the average, explicit paths of primary emissions for gluon-initiated jets which develop in vacuum. Such paths are presented in figure 4.38 for the three different reclustering algorithms. There are several different observations that can be made here. First, we note that the paths are remarkably close to a linear behaviour. Moreover, we confirm that indeed the formation time between successive splittings increases. For the C/A algorithm we note that the (average) first emission takes a total time of $\ln(t'_F) \simeq 1.0 \Leftrightarrow t'_F \simeq 0.6\,\text{fm}/c$ to occur while the time between the eighth and the ninth splittings is much larger: $\ln(t'_F) \simeq 5.7 \Leftrightarrow t'_F \simeq 56.9\,\text{fm}/c$. The progressive delay of emissions is clearly substantial. For k_t-reclustered jets the results are perfectly similar, with the first emission taking approximately $t'_F \simeq 0.4\,\text{fm}/c$ to occur while between the sixth and seventh emissions there is a time span of $t'_F \simeq 55.5\,\text{fm}/c$. Finally, the t_F algorithm yields
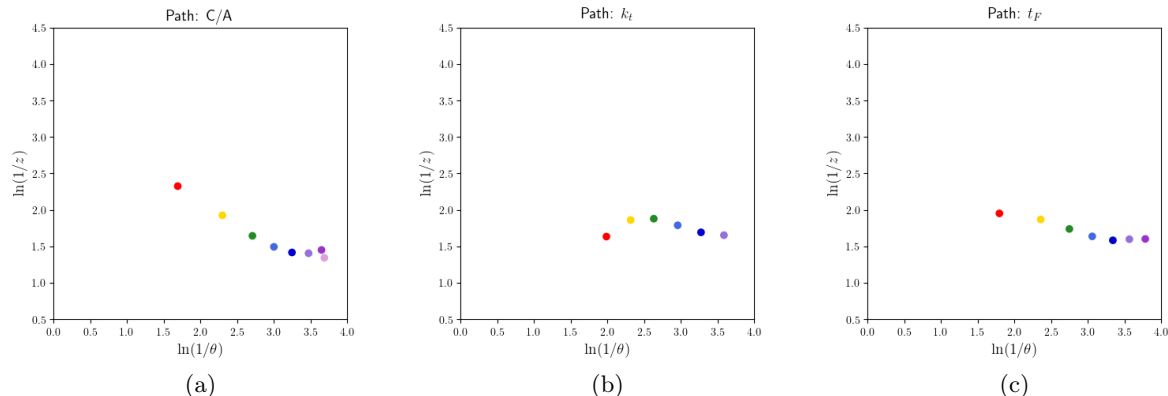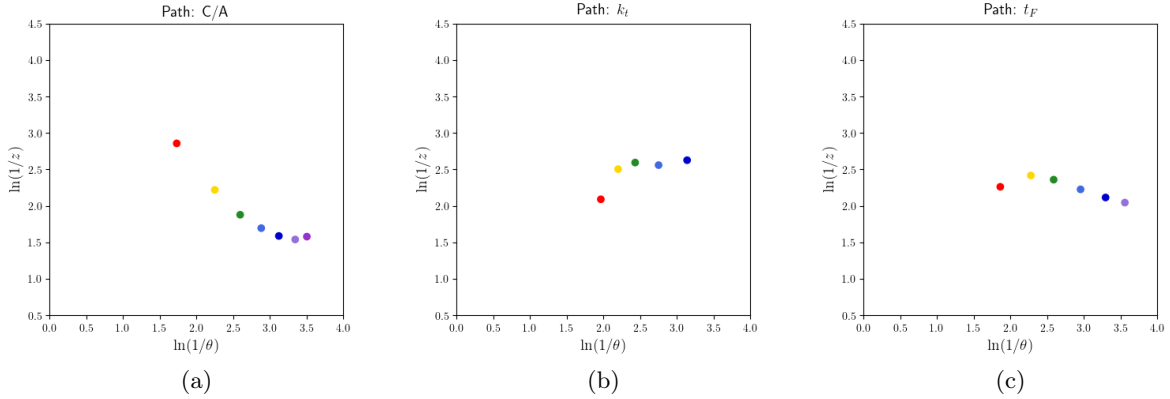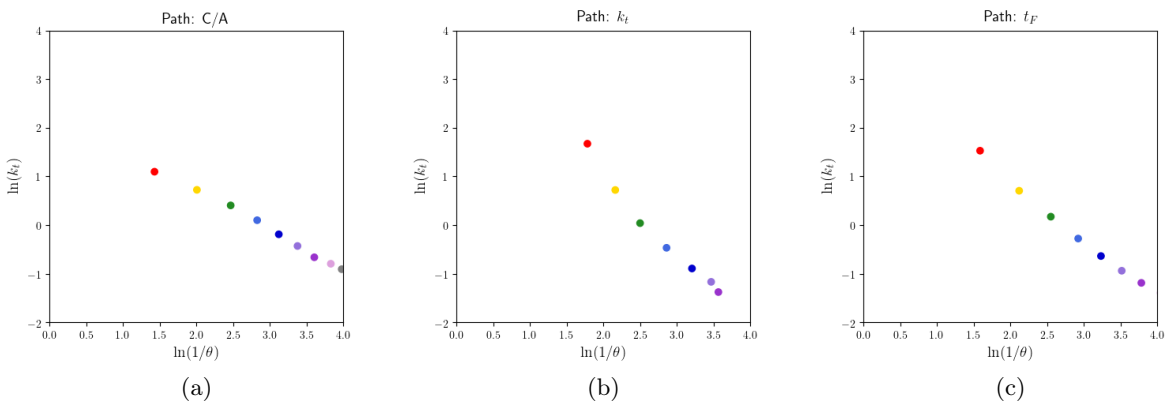
identical values.
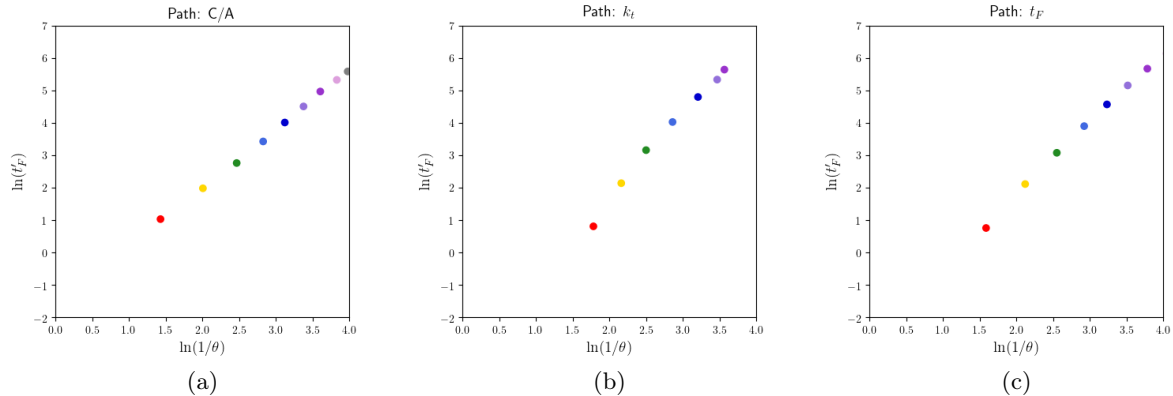
Figure 4.38: Average path of emissions along the primary Lund plane for gluon-initiated jets, in vacuum, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

More interesting now is to analyse how this picture is altered by the presence of the QGP. From our discussion in section 4.4, we expect that the medium is responsible for a delay of emissions, particularly at large angles. This suggests that the average first emissions should yield a (much) larger formation time than the corresponding vacuum value. This is also expected for the later emissions while in lesser scale.

To substantiate these predictions, we can examine the paths of quenched, gluon-initiated jets presented in figure 4.39. Indeed it is immediate to notice that there is an upwards shift of the (average) first splittings, towards larger formation times. In particular, the (average) first emission taken from the C/A clustering sequence has a formation time $t'_F \simeq 2.1 \, \mathrm{fm}/c$, a value which is nearly four times larger than the corresponding result obtained in vacuum. The (average) first emission of quenched, k_t-reclustered jets is also delayed by a factor of (almost) four: $t'_F \simeq 1.4 \, \mathrm{fm}/c$.

The delay of later (average) splittings (i.e., the increase of their respective formation time) can also be spotted in these plots, particularly for the k_t and the t_F algorithms. Overall, our predictions of medium-delayed radiation are easily corroborated by this study.

It is also worth noting that the (average) fifth splitting in the C/A clustering sequence has a formation time $t'_F \simeq 12.7 \, \mathrm{fm}/c$ meaning that this and all posterior emissions should happen outside of the QGP. On the other hand, the (average) fourth splitting of k_t- and t_F-reclustered jets has, respectively, $t'_F \simeq 19.4 \, \mathrm{fm}/c$ and $t'_F \simeq 18.1 \, \mathrm{fm}/c$, such that this and all subsequent splittings (also) occur outside of the medium.
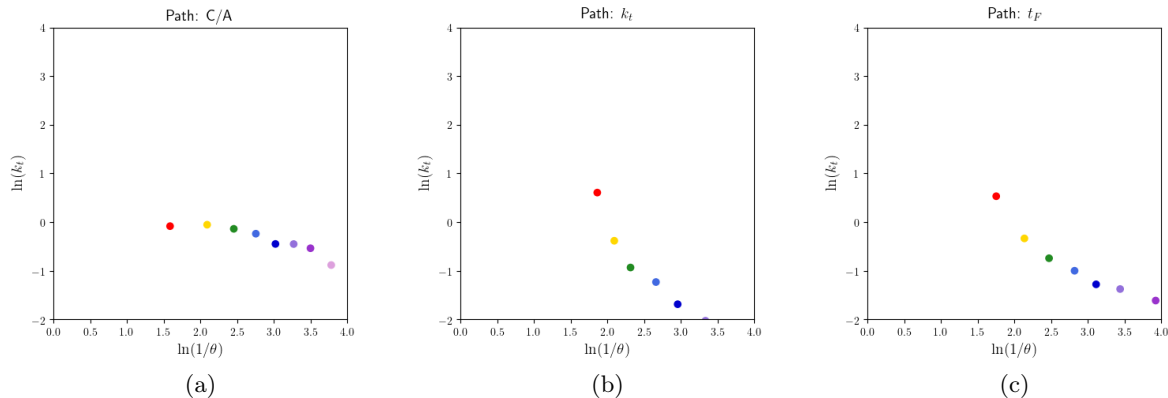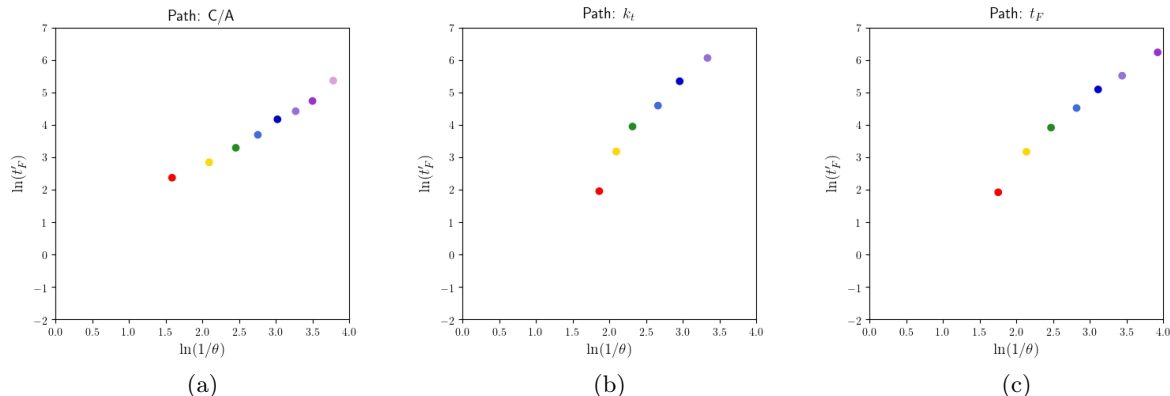
Figure 4.39: Average path of emissions along the primary Lund plane for gluon-initiated jets, in the presence of the medium, reclustered with the: (a) C/A, (b) k_t and (c) t_F algorithms. The ordering of emissions is the following: red (first), yellow (second), green (third), light-blue (fourth), dark-blue (fifth), light-purple (sixth), dark-purple (seventh), pink (eighth) and grey (ninth).

4.7 The effect of the pseudo-rapidity cuts

In the event selection criteria used in all of the preceding analyses there were two different cuts involving the pseudo-rapidity. The first one demanded that the final-state particles had $|\eta| < 2.5$, while the second required that the final-state jets had $|\eta| < 2.0$. We decided to investigate how the radiation pattern presented in the kinematical Lund planes is altered by making more constraining selection criteria. To that end, we changed the cuts to $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets. Using these more restricting values we constructed once more several different kinematical Lund planes. In this section, we chose to present only the results for the Lund plane definition $(\ln(1/\theta), \ln(t'_F))$ and for gluon-initiated jets. The overall conclusions are the same for the other kinematical planes and also for the case of quark-initiated jets.

The Lund planes filled only with the first emissions of gluon jets in vacuum and in the medium are presented, respectively, in figures 4.40 and 4.41, which can be compared to figures 4.25 and 4.26. By doing so, it is fairly straightforward to realise that the use of more constraining pseudo-rapidity cuts yields nearly no difference in the observed radiation patterns. It appears that the previous splitting maps exhibit a smoother behaviour but, other than that, we observe no evident or remarkable differences. Notably, there is no apparent increase in the discrimination between vacuum and in-medium jets. This is clearly observed by examining the kinematical planes with the difference between the radiation patterns with and without quenching effects. These are presented in figure 4.42, to be compared with those depicted in figure 4.27.
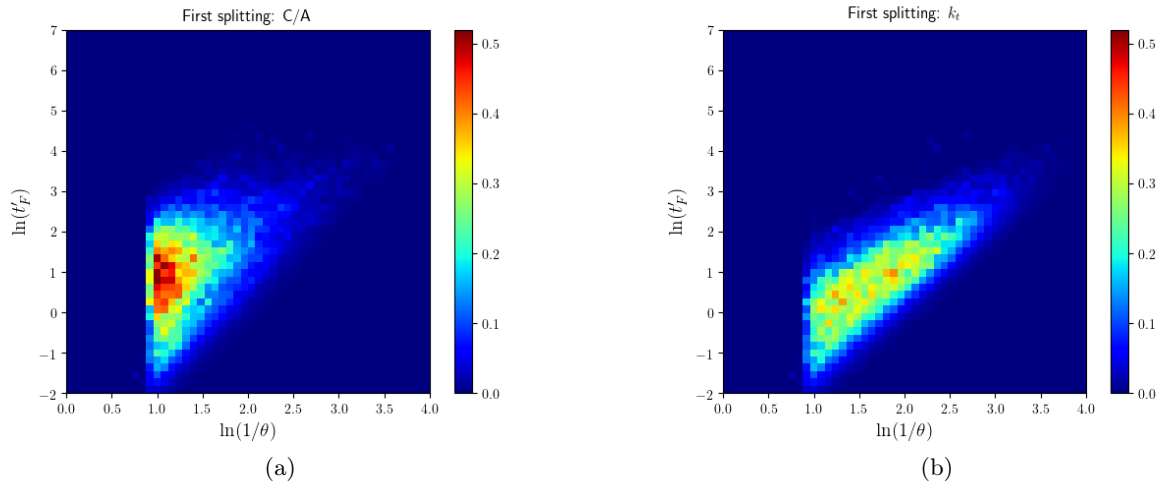
Figure 4.40: Lund planes filled only with the first splitting of gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm. These maps are obtained using more constraining pseudo-rapidity selection criteria, namely: $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets.
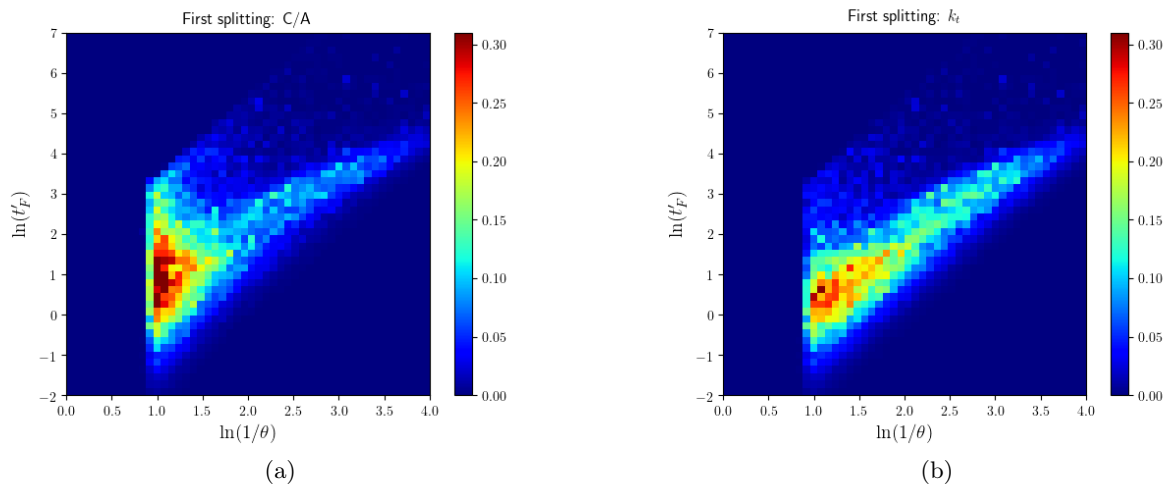
Figure 4.41: Lund planes filled only with the first splitting of gluon-initiated jets, in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm. These maps are obtained using more constraining pseudo-rapidity selection criteria, namely: $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets.

(a)   (b)
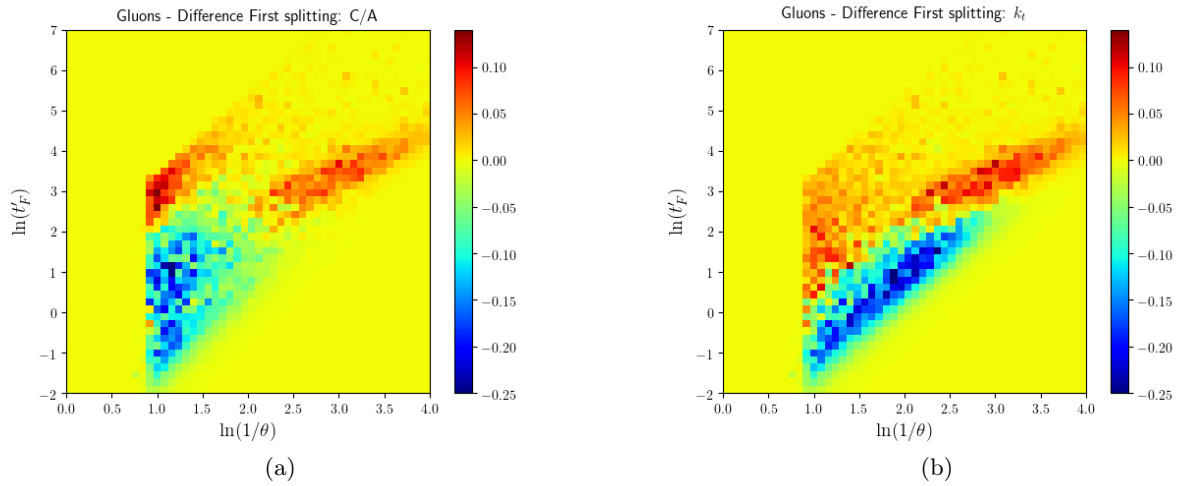
Figure 4.42: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm. These maps are obtained using more constraining pseudo-rapidity selection criteria, namely: $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets.

The same results are true for the radiation pattern in the primary Lund plane. That can be taken from the comparison of figures 4.43, 4.44 and 4.45, with figures 4.28, 4.29 and 4.30 of section 4.4.

(a)   (b)
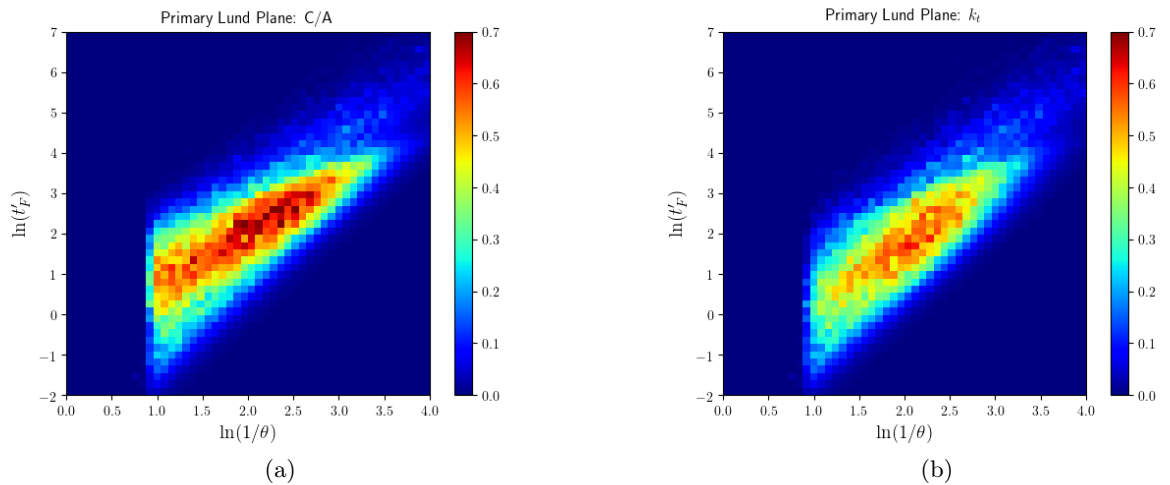
Figure 4.43: Primary Lund planes for gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm. These maps are obtained using more constraining pseudo-rapidity selection criteria, namely: $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets.

(a)   (b)
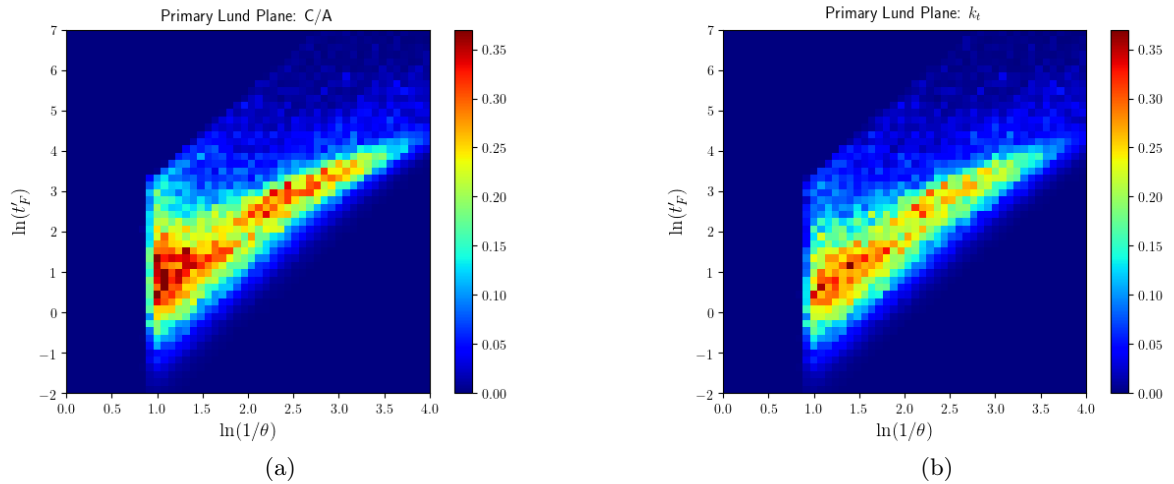
Figure 4.44: Primary Lund planes for gluon-initiated jets, developed in the presence of the medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm. These maps are obtained using more constraining pseudo-rapidity selection criteria, namely: $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets.

(a)   (b)
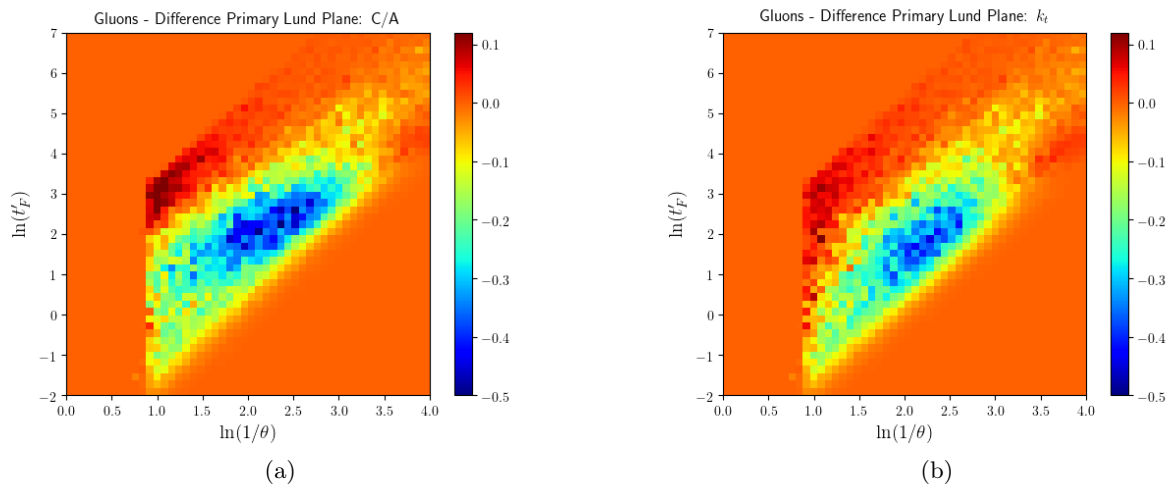
Figure 4.45: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm. These maps are obtained using more constraining pseudo-rapidity selection criteria, namely: $|\eta| < 1.5$ for the final-state particles and $|\eta| < 1.0$ for the final-state jets.

4.8   Summary of the main results

In this chapter we have done a comprehensive analysis of the radiation pattern of jets with and without quenching effects, resorting to different kinematical Lund plane definitions, filling procedures and jet reclustering algorithms. We have also shown that t'_F is a good proxy for t_F and analysed in detail the (average) explicit paths of primary emissions. Finally, we studied the effect of employing different pseudo-rapidity cuts. Since this was an extensive section with a lot

of results, it might be useful at this point to pause and take note of some of the most important findings:

- The effective number of jets available in vacuum is approximately 1.6 and 1.4 times larger than the corresponding in-medium value, respectively for gluon- and quark-initiated jets. This result is consistent with the well established phenomenon of jet energy loss. Furthermore, the fact that the attenuation of quark-initiated jets is slightly smaller than that of gluon-initiated ones corroborates the picture that the former lose less energy in the medium than the latter, as a consequence of the smaller number of emissions in their development.

- The overall number of primary emissions for the totality of vacuum jets is around 2.5 times larger than for quenched jets. This means that the latter have smaller particle multiplicity than the former, suggesting that the medium "throws" a lot of particles outside of the jet cone. Here it is however worth noting that, realistically, this reduction in the overall number of particles within quenched jets should not be as marked as the one obtained with these samples. On the one hand, the inclusion of hadronisation translates in a depletion that is considerably smaller (as we will have a chance to see in the next chapter); on the other hand, some particle content can be recovered by the inclusion of recoil-effects, i.e., by considering jet-induced excitations of the QGP, some of which can end up inside of the jet cone.

- The primary branch is the longest for C/A-reclustered jets and the shortest for k_t-reclustered jets.

- The radiation pattern of quark-initiated jets is less structured and more spread out than that of gluon-initiated jets, while at the same time it retains the distinct features associated with the latter.

- The medium is responsible for destroying part of the structure of the radiation pattern in vacuum and spanning the emissions across the plane. In particular, the relative fractions of large-angled, soft radiation and collinear, hard radiation are enhanced in the medium with respect to the corresponding vacuum fractions. We have also seen that the medium delays emissions, i.e., it enhances the fraction of splittings with larger formation times, particularly in the large-angle region; for instance, on average, the first emission is delayed approximately by a factor of four.

- The k_t algorithm appears to be the best candidate to perform the reclustering of the jets.

- The path of primary emissions along the kinematical Lund planes is a distinct feature between vacuum and in-medium jets. Particularly, the average emissions within quenched jets involve smaller transverse momentum scales and are shifted towards larger formation times with respect to vacuum.

- The pseudo-rapidity selection criteria have almost no influence in the radiation pattern depicted in the kinematical Lund planes.

Chapter 5

Analysis of $Z + jet$ hadronic samples

In the previous chapter we have provided a comprehensive set of studies on the potential of kinematical Lund planes for discriminating between the radiation pattern of quenched and non-quenched jets. These studies were performed on vacuum and medium $Z + jet$ samples at partonic level. This allows us to gain some insight into jet quenching at the fundamental level and without the contribution of effects which are not understood from first principles. However, as we know, the final-state particles of collider experiments are hadrons and not partons. For that reason, we decided to investigate in what way our findings from the previous chapter are altered by the introduction of hadronisation and ISR into the picture. Every other feature of the event generation remained the same and the event selection criteria were also maintained. This way, the results from this chapter can be directly compared to the preceding ones.

Our expectation is that the introduction of hadronisation destroys some of the previously observed structure (or features), i.e., we expect that it smears the radiation patterns obtained at parton level. Even still, we believe that the Lund planes should continue to carry useful information for the discrimination between quenched and non-quenched jets.

Once again, we focus our studies on gluon-initiated jets for which we see, in this case, a suppression of the effective number of in-medium jets by a factor of 1.8 with respect to vacuum, that is, in vacuum there are 1.8 times more jets than in the same amount of in-medium events. This suppression factor is close to the one obtained at partonic level and (once again) evidences the phenomenon of medium-induced energy loss. Furthermore, the total effective number of emissions is 2.0 times larger for vacuum events than for medium events. In this case, the difference between the medium-induced suppression factor of the effective number of jets and of the total effective number of emissions in the events is not as large as in the partonic samples. In fact, the results show that, at parton level, the average effective number of particles per jet for medium events corresponds to 58% of that in the vacuum case, while at hadron level that value is 88%. Furthermore, as it is to be expected, at hadron level the particle multiplicity per jet is much larger than at parton level, specifically, approximately 5.4 and 8.2 times larger, respectively for the vacuum and in-medium cases.

After these introductory findings we can now construct different kinematical Lund planes. We chose to analyse the traditional Lund plane construction $(\ln(1/\theta), \ln(z\theta))$ used in [49] and also our novel kinematical definition $(\ln(1/\theta), \ln(t'_F))$. The results found with each of these constructions can be compared to the findings in sections 4.2 and 4.4, respectively, and also to one another. This is what we set out to do in the following two sections.

5.1   Analysing the $(\ln(1/\theta), \ln(z\theta))$ Lund planes

The kinematical Lund planes filled only with the first splitting of hadronised, gluon-initiated jets reclustered with the C/A and k_t algorithms are presented in figures 5.1 and 5.2, respectively for the vacuum and in-medium cases. The smearing of structure with respect to the parton level situation (figures 4.13 and 4.14) is clear from the analysis of these kinematical planes. In particular, the radiation pattern of quenched jets is essentially featureless.

Figure 5.1: Lund planes filled only with the first splitting of hadronised, gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Figure 5.2: Lund planes filled only with the first splitting of hadronised, gluon-initiated jets, in medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

The considerable smearing which is observed in these splitting maps with respect to the corresponding parton level results complicates the evaluation of the effect of the presence of the medium in the radiation pattern of jets. Based on our observations at parton level, we manage to identify that there still appears to be a dispersion of emissions throughout the plane. Particularly,

73

the fractions of emissions in the soft, large angle region of the phase space (for C/A-reclustered jets) and in the hard, collinear region (for the k_t algorithm) seem to be enhanced in the presence of the medium, with respect to vacuum. However, this is not nearly as apparent as it was before. In fact, plotting the differences between the radiation pattern with and without quenching effects clearly shows that now it is much more difficult to discern between vacuum and in-medium jets (see figure 5.3).

Figure 5.3: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.
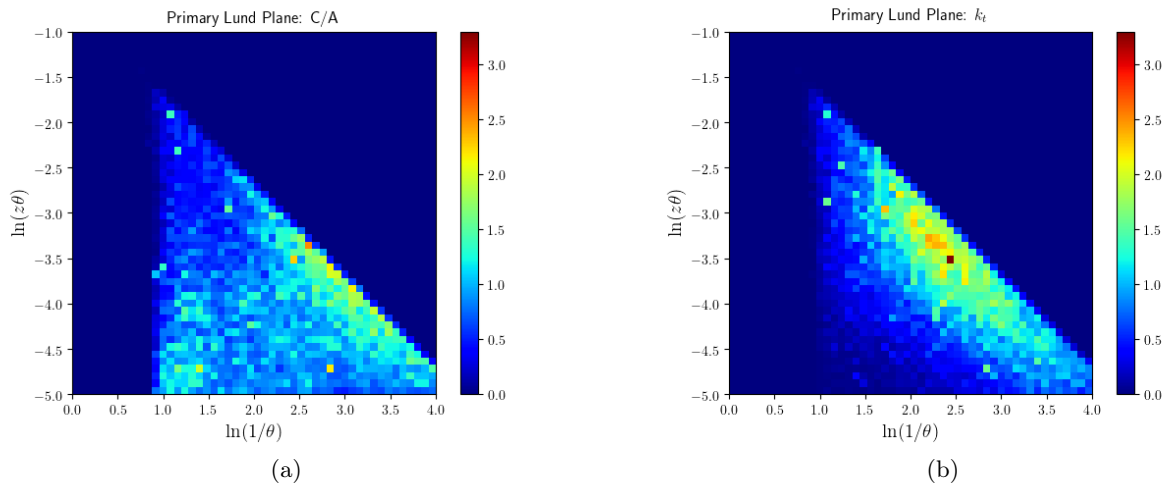
Figure 5.4: Primary Lund planes for hadronised, gluon-initiated jets developed in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

We can also have a look at the radiation pattern of primary emissions, presented in figures 5.4 and 5.5, respectively for the vacuum and medium cases. Once more, there is a significant smearing of structure to the radiation pattern of both quenched and non-quenched jets. However, for the former the results obtained are practically structureless (as we have already noted for the first

emissions). In fact, the radiation pattern of primary emissions of quenched jets is nearly uniform throughout the plane. As a result, it appears that, in the presence of the QGP, the tendency towards larger fractions of softer radiation remains. Nevertheless, most of the discrimination power of the kinematical Lund planes seems to have been lost, as corroborated by figure 5.6.

Figure 5.5: Primary Lund planes for in-medium, hadronised, gluon-initiated jets found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
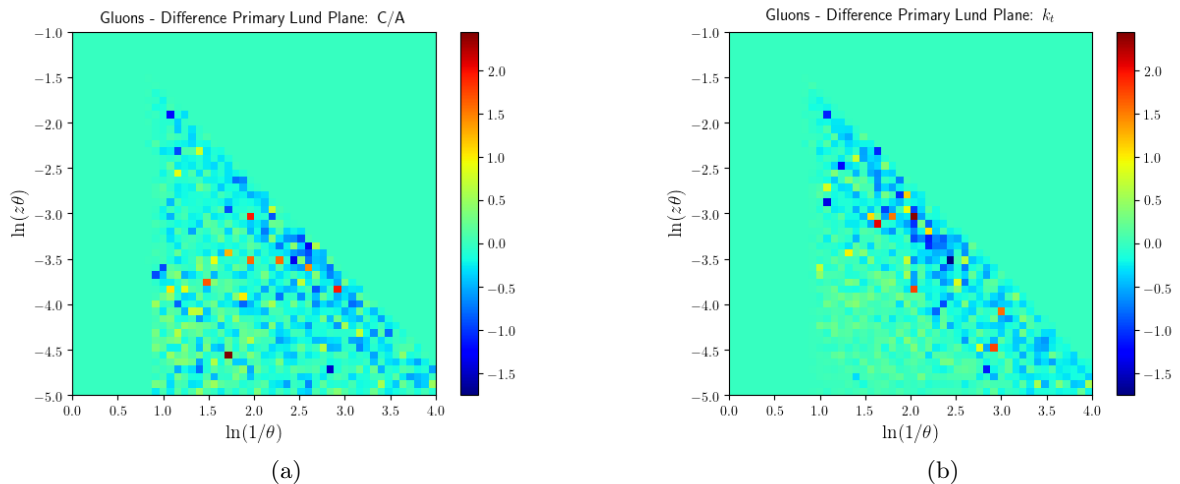
Figure 5.6: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

5.2 Analysing the $(\ln(1/\theta), \ln(t'_F))$ Lund planes

The results from the previous section were not very encouraging since they appear to indicate that the inclusion of hadronisation and ISR spoils the large discrimination power provided by the kinematical Lund planes. Of course, this observation is not benefiting from the absence of

grooming techniques to "clean-up" the jets but even still, given that such procedures were also not applied at parton level, we would like to confirm that this approach to the study of jet quenching is, indeed, promising and worth investing in. Fortunately, the results obtained with our proposed Lund plane definition - $(\ln(1/\theta), \ln(t'_F))$ - are not as discouraging as the ones obtained in the preceding section.

Figures 5.7 and 5.8, show the splitting maps filled only with the first emissions of C/A- and k_t-reclustered jets, respectively in vacuum and in the presence of the QGP. The first reassuring observation which is obvious to make is that these kinematical planes are not as structureless as the ones presented in the previous section. That is particularly true for the C/A algorithm, for which we see a large pooling of first emissions along the vertical line $\ln(1/\theta) \simeq 1.0$, i.e., at large angles, and in the region between $\ln(t'_F) \simeq 2.0 \Leftrightarrow t'_F \simeq 1.5\,\mathrm{fm}/c$ and $\ln(t'_F) \simeq 3.0 \Leftrightarrow t'_F \simeq 4.0\,\mathrm{fm}/c$. The concentration of a large fraction of splittings in this specific region of phase space is present both in vacuum and in the medium, but in the latter there seems to be a slight migration of a dominant fraction towards larger formation times, as we would expect.

Given that the kinematical Lund planes are not as featureless as before, the observation of differences in the radiation pattern of quenched and non-quenched jets is facilitated. Particularly, we see that, in the presence of the medium, emissions are shifted towards larger formation times. In fact, although the two remarkable prongs that were present in the medium radiation patterns at parton level are no longer a perceivable feature, the effect of medium-delayed radiation is still observed. This is manifest in figure 5.9.

(a)   (b)
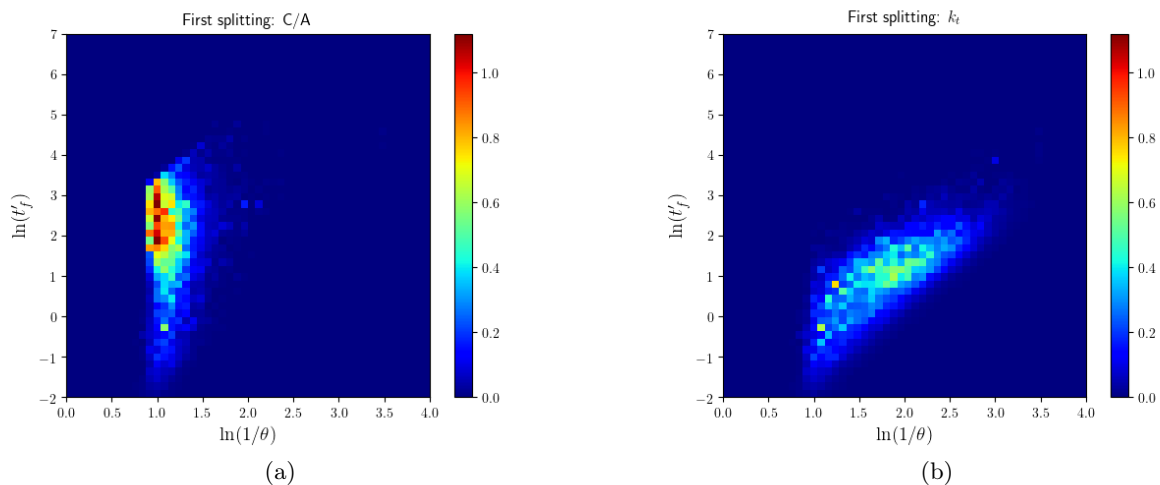
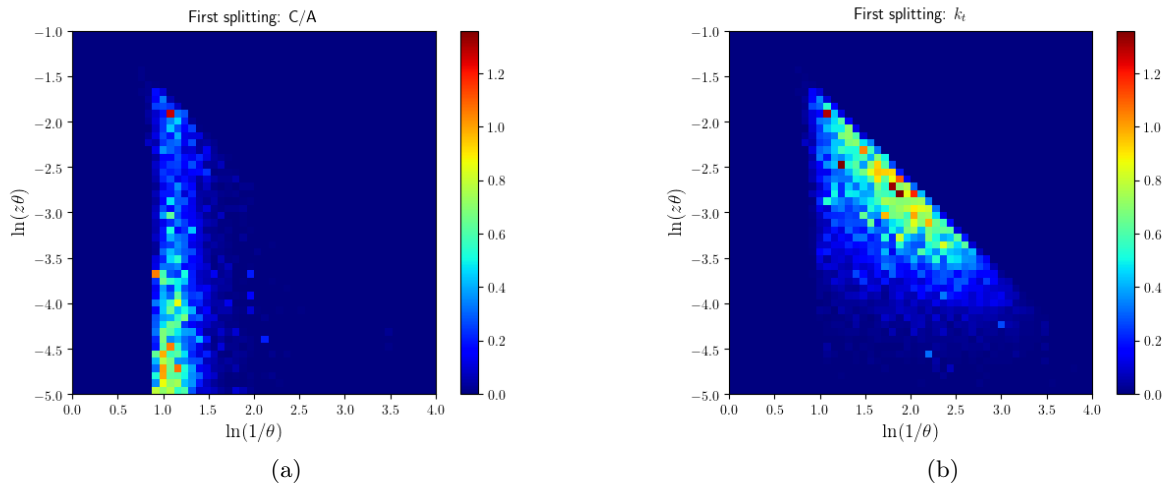Figure 5.7: Lund planes filled only with the first splitting of hadronised, gluon-initiated jets, in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
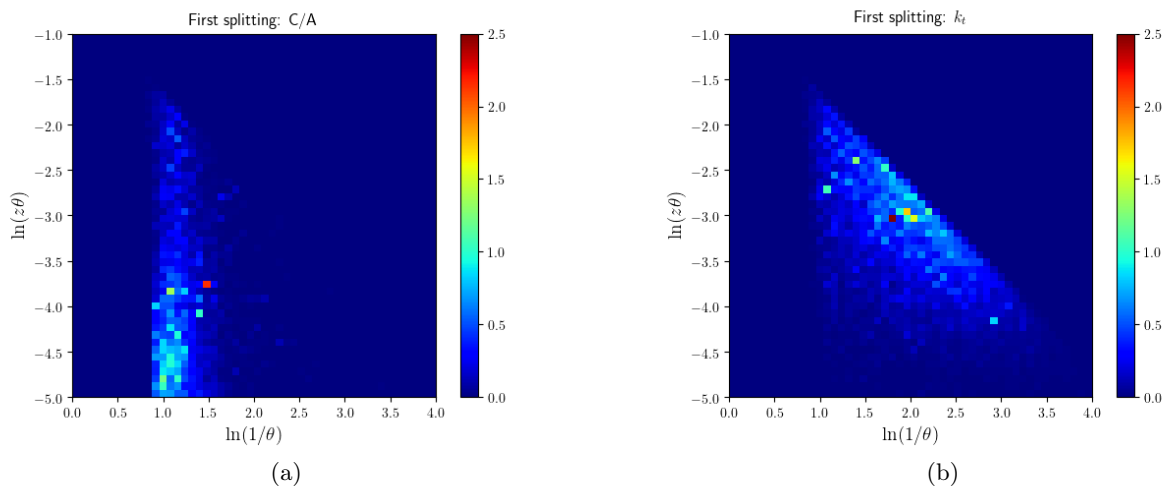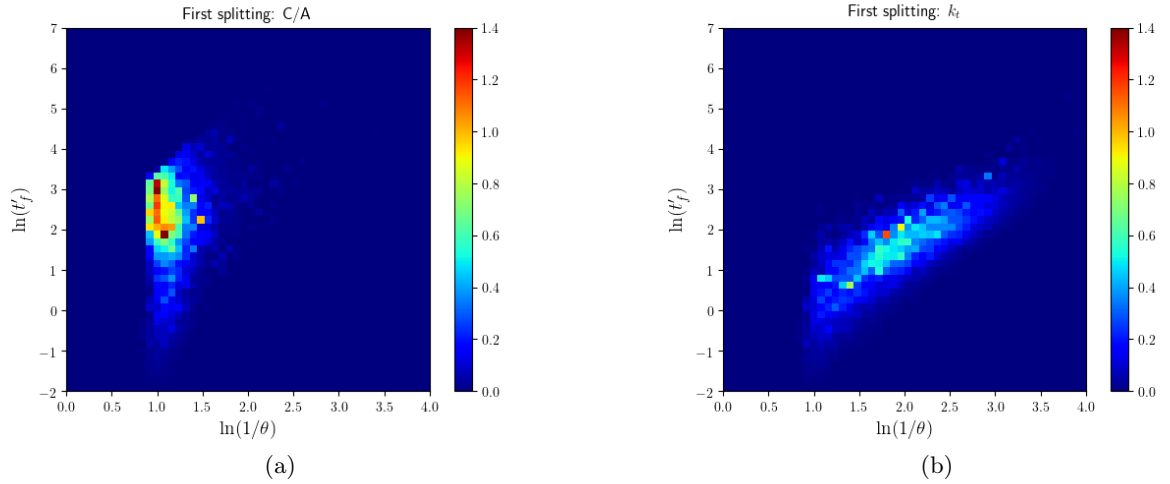
Figure 5.8: Lund planes filled only with the first splitting of hadronised, gluon-initiated jets, in medium, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
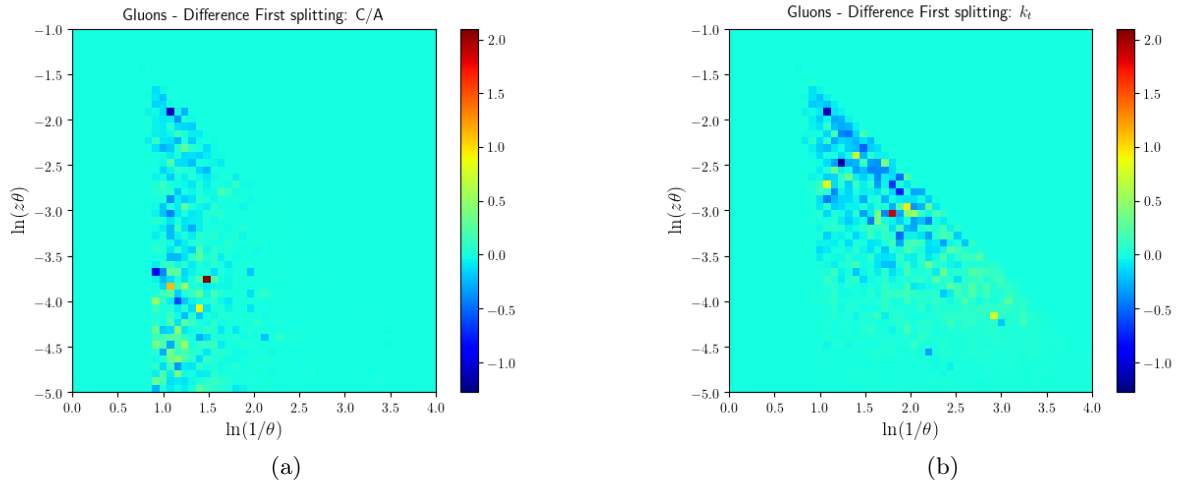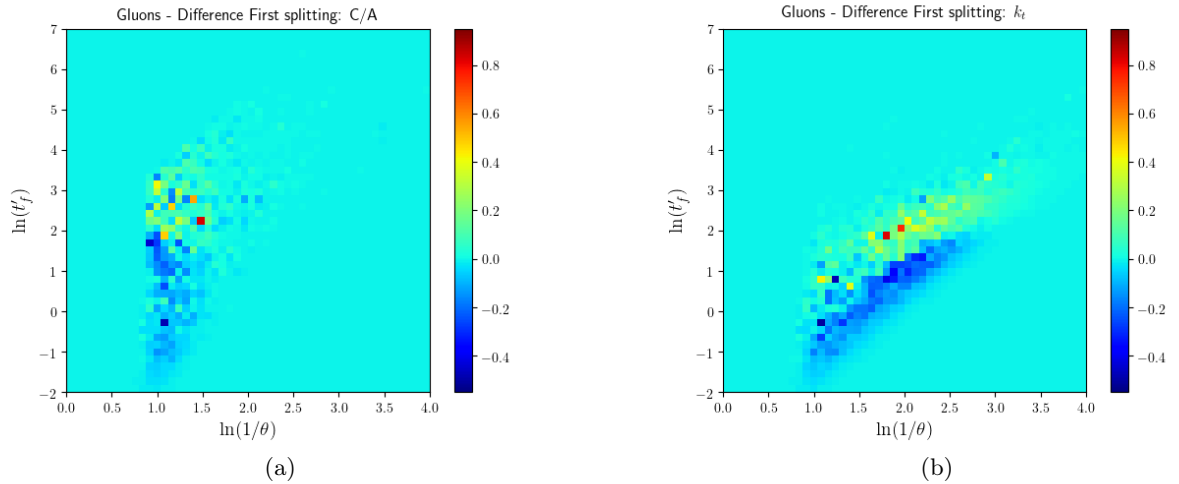

Figure 5.9: Lund planes that register the difference between the radiation pattern of first emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Given the positive results obtained for the radiation pattern of first emissions, we expect that the $(\ln(1/\theta), \ln(t'_F))$ primary Lund plane also exhibits more structure than the one from the preceding section. Attached to that, should also come a good discrimination power between samples with and without quenching effects. The results for the splitting maps of primary emissions are presented in figures 5.10 and 5.11 for vacuum and medium samples, respectively. Indeed, from these figures we see that there is much more structure than for the $(\ln(1/\theta), \ln(z\theta))$ kinematical plane, particularly, the C/A exhibits some interesting concentrations of emissions at large angles and along a diagonal of slope $m \simeq 2$. Similarly, k_t-reclustered jets also present a predilection towards pooling around that same diagonal line.

By comparing the vacuum and medium radiation patterns, we observe that quenched jets once

again appear to present larger splitting fractions at larger formation times, as is corroborated by figure 5.12.
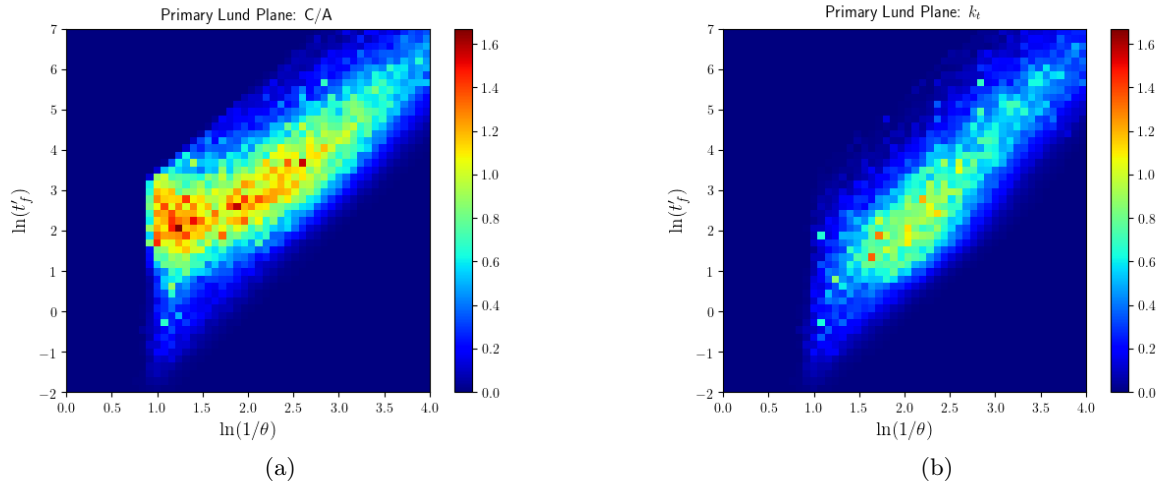
Figure 5.10: Primary Lund planes for hadronised, gluon-initiated jets developed in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
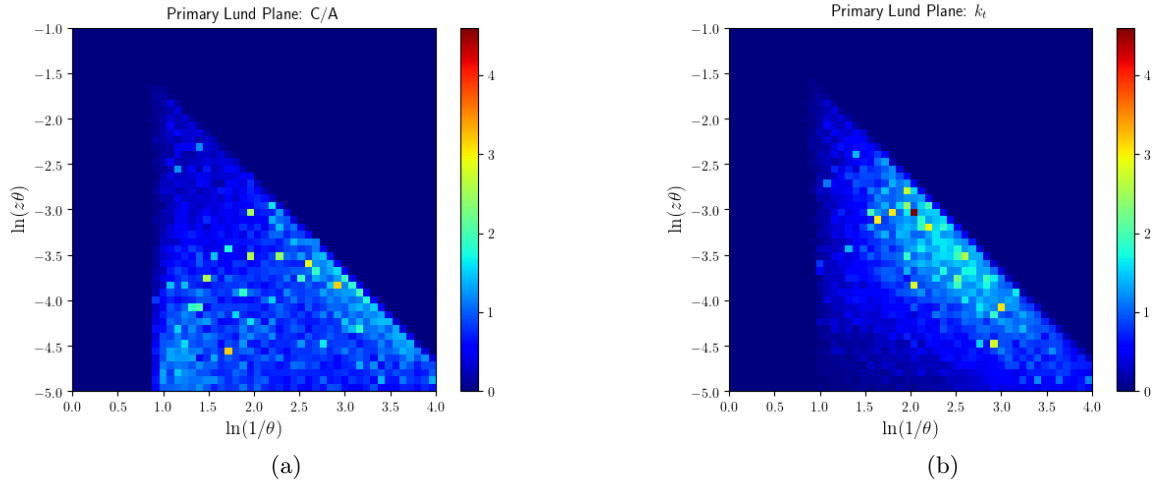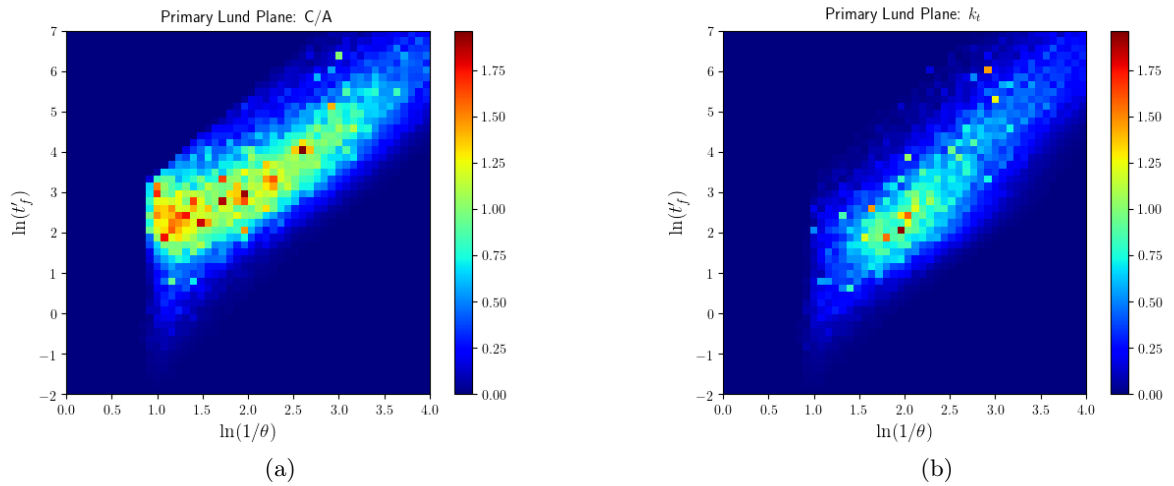

Figure 5.11: Primary Lund planes for in-medium, hadronised, gluon-initiated jets found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.
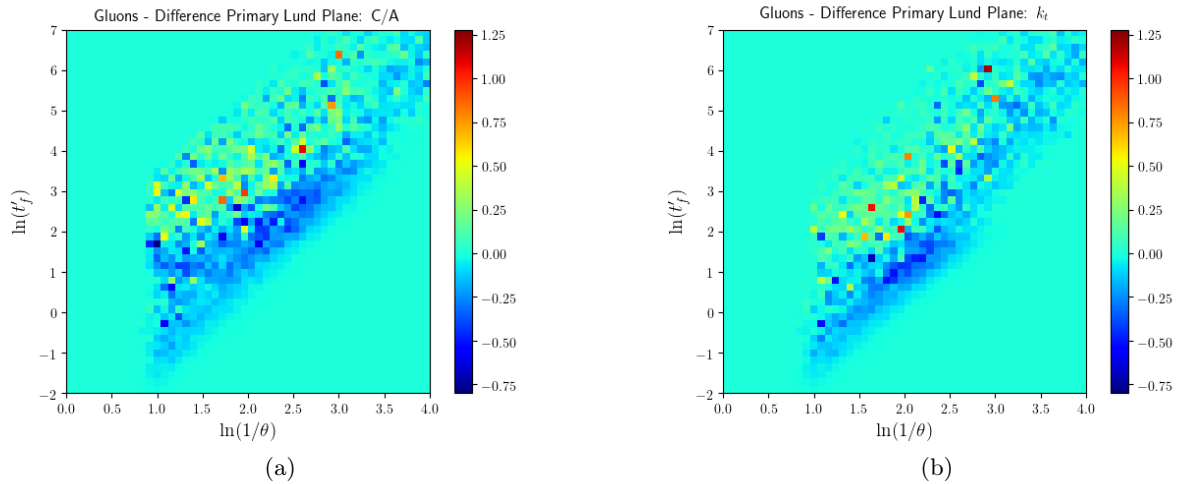
Figure 5.12: Lund planes that register the difference between the radiation pattern of primary emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

The essential conclusion here is that the introduction of hadronisation greatly diminishes the discrimination power provided by the kinematical Lund planes. However, that reduction is much smaller with our proposed Lund plane definition, which retains more structure than the traditional $(\ln(1/\theta), \ln(z\theta))$ construction used in [49]. While this conclusion is drawn entirely based on the visual observation of the radiation patterns with and without quenching effects and their respective differences, that does not make it any less significant. In fact, we have argued that much of the power of the kinematical Lund plane lies in its visual nature, which allows us to easily identify enhancements and/or suppressions in specific regions of phase space and later use those observations to manually craft observables sensitive to those particular regions.

Even still, we desire to corroborate our statement that the $(\ln(1/\theta), \ln(t'_F))$ construction does better at discriminating between vacuum and in-medium jets. This validation was performed resorting to a machine learning (ML) architecture, specifically a recurrent neural network (RNN), and is detailed in chapter 6 of this dissertation.

Chapter 6

Machine learning results on the $Z + jet$ hadronic samples

While artificial intelligence (AI) was once no more than mere science fiction, the astonishing progress registered in the field since its birth in 1956 [83] makes it so that it is nowadays ubiquitous in our daily life. The field of AI deals with programming machines to automate intellectual tasks that would otherwise have to be performed by humans [84]. Initial approaches to this problem consisted in hardcoding explicit rules which encompass the knowledge required by the computer to perform the desired task. While these approaches proved very useful in formal settings with well-defined rules, for instance, for teaching a computer how to play chess, they turned out to be unfruitful in more abstract tasks which are simple to humans, such as object recognition [85]. Thus, researchers realised that, for a large variety of problems, AI systems needed the ability to acquire their own knowledge from an input dataset. This gave rise to a new approach to AI known as machine learning (ML). In this case, the ML algorithm (or model) is given a set of training data and the corresponding answers (or targets). From these, the algorithm learns the appropriate rules, which can afterwards be applied to new sets of data to give back the answers associated with them [84]. In our modern society, ML is behind personalised result refinement in search engines, email spam detection, online advertising and product recommendation, personal assistants such as Siri or Google Now, online fraud detection, etc. It is also used to perform natural language processing, speech, handwriting and object recognition, medical diagnosis, etc.

The large amount of data involved in modern high energy physics experiments motivated physicists to explore ML techniques. These have long since established their worth in the field and are now widely used for (i) analysing experimental results or data from MC parton showers; (ii) reconstructing events and (iii) improving the trigger system.

In the context of this dissertation, we desire to build a ML model that is able to learn distinct features between quenched and non-quenched jets (given our sample of events) and can afterwards correctly classify each jet from a new jet sample as either quenched or non-quenched. In order to provide the reader with the background knowledge required for the understanding of our implemented ML solution, the next section is devoted to making a brief introduction to ML, deep learning (DL) and artificial neural networks (ANNs). The aim is to provide an overview of the most important concepts and not an in-depth description of the broad fields that are ML, DL and ANNs. For that, the reader is referred to [84] and [85]. In section 6.2, we describe our approach to the classification task at hand and present the results obtained.

6.1 A brief overview of machine learning, deep learning and artificial neural networks

We have seen that ML algorithms are able to learn the appropriate rules that convert input data into the respective answers (or targets). In a more formal language, we say that ML models work by automatically performing transformations on the input data with the goal of finding the most useful representation for that data, i.e., the representation that gets the model closer to the expected target [84]. The "learning" in "machine learning" refers to the automated, iterative process of searching for the best representation of the data.

DL is a specific field of ML in which we perform several consecutive (simple) transformations on the data (instead of a single input-to-target transformation), such that we have successive layers of different representations for the data. ANNs are the dominant models in deep learning. They were originally inspired by the neurobiology of the human brain [85] and consist in literal layers stacked on top of each other, each with a certain number of units or neurons. Each layer takes as input the information output by its predecessor, performs a transformation on the received data and passes that information forwards to the subsequent layer. The transformation performed by each layer is encoded into the layer's weights. The process of searching for the optimal representation (i.e., for the optimal weights of the layers) is guided by a feedback signal, which is a measure by which the model is able to evaluate its own performance. This feedback signal is provided by the loss function (also known as objective function) which takes the predictions from the network and checks them against the true targets, yielding a loss score which determines how far the predictions are from the targets. The larger the value of the loss score, the larger the distance between the predictions and the targets and the worse the performance of the model. Since the weights of the NN are initialised at random, the first loss score value obtained will be large. The learning process then consists in using that score as feedback to the NN, so that the weights are iteratively updated in directions that successively reduce the loss. The adjustment of the weights is the job of the optimiser which implements a specific variant of stochastic gradient descent. The training workflow of a generic NN model (with two layers) is depicted in figure 6.1.

There are several different types of layers which can be put together to produce different NN architectures. It is the data processing task at hand that determines the appropriate NN architecture. For instance, to process sequences of data, recurrent neural networks (RNNs) built from recurrent layers such as gated recurrent units (GRUs) or long short-term memory (LSTM) layers are indicated. On the other hand, for dealing with image data processing, convolutional neural networks (CNNs) have proved long ago to be the best ML solution.
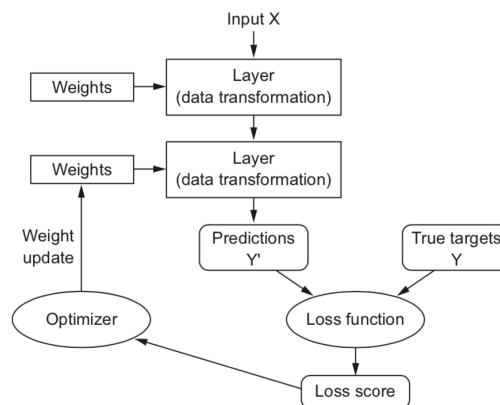

Figure 6.1: Pictogram of the training workflow of a NN, taken from reference [84].

To summarise, the four essential ingredients in the training of an ANN model are: (i) the input data and respective targets; (ii) the (adequate) choice of architecture for the task at hand; (iii) the (appropriate) choice of the loss function to measure the performance of the NN; and (iv) the optimiser to update the weights of the model. The generic description provided above serves the purpose of providing the reader with a general (superficial and oversimplistic) understanding of the overall training process of a NN model and the role played by each of these key ingredients. However, there is still a plethora of ideas, concepts and details that have not been discussed but which are paramount to any ML task. We will address some of them briefly, given their relevance to our own work.

Data partitioning, underfitting, overfitting and early-stopping

In the previous description of the training process of an ANN model, we have seen that we use a loss score to act as a feedback signal so that the optimiser knows in which direction it should update the weights of the model. If we have a given sample of input data (and, of course, the respective targets) we should not use the totality of that sample for the training itself. Instead, we should partition it into three sets of data: (i) a training set; (ii) a validation set; and (iii) an evaluation set. As the name indicates, the training set is used for the training proper, just as described above: the weights of the different layers are initialised, the model is used to generate predictions for the targets of the input training set, the training loss score is computed and the optimiser updates the weights of the ANN such as to minimise that score. This process is iterated a given number of times and each iteration step is called an epoch. Inevitably, given the way this process is implemented, as we move along the successive epochs, the training loss (computed using the training set) gets progressively lower. However, this does not necessarily mean that our model is getting better. In fact, at some point, it will happen that the NN starts to get sensitive to specific features of the training set. This means that, while it performs well on that particular dataset, it will perform poorly on new (unseen) sets. When that happens, the model is said to be overfitting the data.

To establish whether a model is overfitting or not, we use the validation set. This set remains unseen by the training process proper, i.e., it is not used in the search for the optimal weights. At the end of each epoch, this (unseen) set of data is used for computing a loss score of its own - the validation loss score. While the training loss always decreases as we move along the successive epochs, the validation loss will start to increase when the model begins to overfit. At that point, the training should be stopped. Opposite to overfitting lies underfitting which consists in stopping the training before the optimal model has been reached.

To prevent both overfitting and underfitting early-stopping methods can be applied. These methods work by monitoring a given metrics (generally the validation loss) and stopping the training if that metric has not improved over a pre-defined number of epochs. At that point, the weights of the NN are updated to the ones that yielded the best validation loss score over the totality of epochs.

After the training of the NN is completed and its weights set, the model obtained can be evaluated on the third (independent) dataset: the evaluation set.

Activation functions

The transformation carried out by each layer in a NN is essentially a (mathematical) linear tensor transformation. However, having a linear model is too restricting to solve most problems, since it will only perform well on linearly separable data. For more complex datasets, non-linearities are necessary so that the NN can properly learn the features of the data, i.e., so that it can find the best representation for that data.

To give a layer the ability to perform a non-linear transformation, we associate with it a given (non-linear) activation function. This function is applied element-wise to the tensor resulting from the (intrinsic) linear transformation performed by the layer, thus originating different weights and adding non-linearity. An activation function can be incorporated into each layer of the model, allowing it to explore a much richer hypothesis space and achieve better performances.

Some common examples of activation functions are the rectified linear unit ($\text{ReLU}(x) = \max(0, x)$), the sigmoid ($\sigma(x) = 1/(1 + e^x)$) and the hyperbolic tangent ($\tanh(x)$).

Specific considerations on binary classification problems

We are interested in building a ML model which is able to classify jets as quenched or non-quenched. Since there are only two possible classes, this is called a binary classification problem. The number of problems which fit this class is immense and, for that reason, extensive research has allowed for the existence of well established guidelines for tackling such problems. For one, the appropriate loss function is the binary cross-entropy function. Moreover, the last layer of the network should be a dense layer with one single unit and sigmoid activation. This way, an input tensor from the previous layer is converted into a scalar output between 0 and 1. This scalar value is known as the classifier and should be interpreted as the probability that the input data point (in our case a jet) belongs to class 1 (in our case, the class of quenched jets) [84]. The default threshold for the value of the classifier is 0.5, meaning that the input jet will be classified as quenched if the value of the classifier is larger than 0.5, and it will be classified as a vacuum jet otherwise. When testing our model in the evaluation dataset, we generate predictions for each input jet according to the value of the classifier. Afterwards, we can compare those predictions with the real targets. This allows us to determine the binary accuracy of our model, i.e., the probability that the model assigns the correct class to a given data point. Analogously to the loss score, the binary accuracy is a metrics which can be used to evaluate the final performance of the designed NN; however, contrary to the loss, the larger the binary accuracy the better the model.

Finally, we can also evaluate and visualise the performance of a binary classifier by building the so-called receiver operating characteristic (ROC) curve. This curve is built by changing the threshold value of the classifier and computing the fraction of quenched jets which are correctly classified as quenched (true positive rates) and also the fraction of vacuum jets which are incorrectly classified as quenched (false positive rates). A curve of the true positive rates against the false positive rates can then be plotted; this is the ROC curve. If we have a model which is capable of perfectly separating the two classes, i.e., a model such that probability densities of the classifier for the two different classes do not overlap, then the corresponding ROC curve should resemble a step function. On the other hand, with a random (or useless) model the probability densities of the classifier obtained for the two classes completely overlap and the ROC curve corresponds to a unit slope line such that the true positive rates equal the false positive rates. The area under the ROC curve (AUC) is thus another metrics for the performance of the model. It should yield a value between 0.5 and 1, the first value corresponding to a random and useless classifier and the second corresponding to a perfect discriminator. Hence, the higher the AUC the better the model.

6.2   Our machine learning investigation and results

In the previous chapter, we have declared that our proposed kinematical Lund plane construction exploiting the formation time of the splittings yields a better discrimination power between

quenched and non-quenched jets than does the traditional $(\ln(1/\theta), \ln(z\theta))$ plane. This was based purely on the visual examination of the radiation patterns obtained with each of these two constructions. In this dissertation, we decided to go a step further and investigate if inputting the information used to fill these planes into a ML algorithm would allow it to discern patterns in the data and correctly classify the jets as quenched or non-quenched. The idea is identical to the one used in [80] to identify jets stemming from the decay of boosted W hadrons.

As we have seen, along the studies made in this dissertation we have used three different procedures for filling the kinematical Lund planes. While we have mostly abstained from showing the full Lund diagram filled with all the emissions within each jet (except for the ones presented in appendix B), we have argued that it exhibits similar appearance to the primary Lund plane and, hence, should yield identical discrimination power.

To perform our ML investigations, we chose to use as input data the sequence of primary emissions within each jet, either:

$$[[\ln(1/\theta_1), \ln(z_1\theta_1)], [\ln(1/\theta_2), \ln(z_2\theta_2)], ..., [\ln(1/\theta_n), \ln(z_n\theta_n)]]$$

or

$$\left[\left[\ln(1/\theta_1), \ln(t'_{F,1})\right], \left[\ln(1/\theta_2), \ln(t'_{F,2})\right], ..., \left[\ln(1/\theta_n), \ln(t'_{F,n})\right]\right] ,$$

depending on the axis definition which is being explored.

There are several reasons to motivate this choice. First, at parton level we have seen that the primary Lund planes yield a good discrimination between quenched and non-quenched jets. We have also remarked that the explicit (average) path of primary emissions is a distinct feature between the two classes of jets. Furthermore, because we focus on the primary emissions, we have an explicit, sensible sequence of coordinate pairs to input to the ML model.

From the visual discrimination viewpoint, neither the use of only the first emission nor of the totality of emissions lose in comparison to the primary Lund plane. As such, they are potentially also a good choice. We have decided against the use of the first splittings for two main reasons: (i) it would yield a very reduced statistics to perform ML on; and (ii) since no grooming procedures have been employed it is probable that many of the identified first splittings are not the ones worth investigating. On the other hand, the use of the totality of emissions would certainly win over the use of primary emissions alone in terms of the larger statistics it provides. However, using all emissions would spoil the concept of a clear, explicit sequence of emissions and, for that reason, we decided to leave such a study for the future.

Since our input dataset is a collection of ordered sequences (one for each jet) it is natural to use a RNN architecture. We tested several different model architectures using different numbers of stacked recurrent layers and different numbers of units in them. In the end, since the performances achieved by the different models were identical, we settled on the simplest architecture, which has the advantage of being quicker to train. As such, our final NN consists in a single GRU layer with 64 units, followed by a dense layer with a single unit and sigmoid activation.[1]

Our ML was implemented in Keras 2.2.4 [86], using TensorFlow 1.13.1 [87] as backend. For the optimisation we opted towards the use of RMSProp [84, 88] with an initial learning rate of 10^{-2} which is reduced by half if there is no improvement to the validation loss score over the last 5 epochs. We set the network to train on a batch size of 512 for a total of 1000 epochs, but implemented an early-stopping method such that, if there is no improvement to the model for 25 consecutive epochs, the training stops and the model's weights are updated to the ones that yielded the lowest validation loss. As explained in the previous section, this serves to prevent both underfitting and overfitting of the model.

---

[1]Alternatively, a single LSTM layer could have been used instead of the GRU layer since, according to reference [80], that should yield an equivalent performance.

Now that the details of our implemented ML approach have been described, we should get back to the discussion of our input physical sample. We have two distinct samples of events: one for vacuum jets and the other for in-medium jets. Different events within each of these samples present different number of jets. We chose to analyse how those different event types would behave in terms of discrimination between quenched and non-quenched jets.

First, we decided to study how the NN would perform on events with a single jet, for which the available sample is $36\,767$ and $48\,767$ events, respectively for the vacuum and medium samples. Secondly, we considered only events with exactly two jets, for which we have $26\,887$ and $19\,743$ events, respectively for vacuum and medium samples. For these events, we have concatenated the ordered sequences of both leading and sub-leading jets into one single sequence and used this sequence as input for the model. Finally, we considered both events with one and two jets such that we have a total number of events that is $63\,654$ and $68\,510$, respectively for vacuum and medium. To study these events with our NN model, we use the sequence of the one jet in single-jet events while for the remaining events we concatenate the sequences of the two jets.

The use of these three different input data samples allows us to investigate if any of them stands out in terms of discrimination. For each of these three cases, the input dataset was partitioned such that 80% of the jet sequences are used for the training, 10% are used for validation along each epoch and the remaining (independent) 10% are used for evaluating the final model performance.

We have also chosen to input the data sequences in three different ways. First, we used as input the sequence of primary emissions obtained with the C/A algorithm together with the co-ordinate points of those emissions in the traditional Lund plane construction: $(\ln(1/\theta_i), \ln(z_i\theta_i))$. Secondly, we maintained the use of the C/A algorithm but chose to use instead the coordinates of the primary emissions within our proposed kinematical plane: $(\ln(1/\theta_i), \ln(t'_{F,i}))$. The comparison between the results of these two cases allows us to either validate or refute our visually drawn conclusion that our novel plane definition yields a better discrimination between quenched and non-quenched jets. Finally, to evaluate the influence of the reclustering algorithm, we chose to use the sequence of primary emissions obtained from the clustering tree of the t_F algorithm, taking once again the coordinate points of the emissions in our proposed plane.

The results obtained for the validation accuracy along the successive epochs are presented in figure 6.2. These results show that while using the coordinates of emissions in our plane yields a much better discrimination between quenched and non-quenched jets, the use of different reclustering algorithms has virtually no impact in improving discrimination. This is consistent with our visually drawn conclusions from chapter 5. We also note that using either one of the three event samples yields no difference in the overall accuracy of the model, i.e., in the discrimination between quenched and non-quenched jets.

The binary accuracy values obtained on the evaluation dataset are presented in table 6.1, supporting all of the previous conclusions. The reader will note that so far those conclusions are all comparative in nature, that is, the model was used essentially as a tool to corroborate the results visualised in the primary Lund planes. In absolute terms, we notice that we manage to achieve binary accuracies of around 72%, meaning that the model correctly classifies jets 72% of the times. This is a very encouraging result.

(a)

(b)

(c)
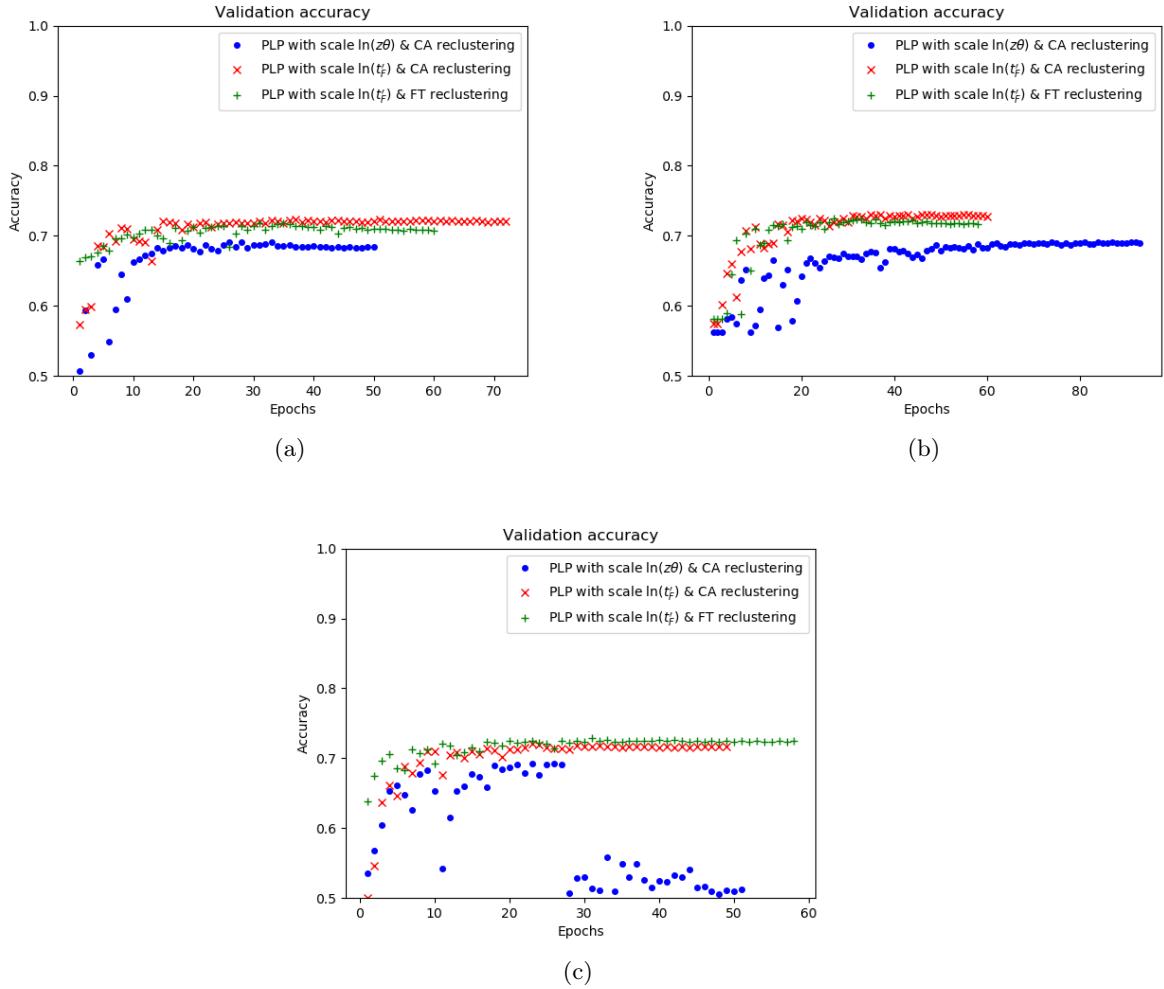
Figure 6.2: This figure presents the binary accuracy achieved by the RNN in the validation set as a function of the successive epochs, when using as input the sequence of primary emissions of events with: (a) one single jet; (b) two jets and (c) either one or two jets. The sequence of primary emissions was input to the RNN as consecutive pairs of coordinate points in the respective kinematical Lund plane construction and for a given reclustering algorithm: particularly, for the $(\ln(\theta), \ln(z\theta))$ kinematical construction with the C/A algorithm (blue dots), and for the new definition $(\ln(\theta), \ln(t'_F))$ together with both the C/A algorithm (red cross) and t_F algorithm (green plus).

	C/A reclustering + $(\ln(1/\theta_i), \ln(z_i\theta_i))$	C/A reclustering + $(\ln(1/\theta_i), \ln(t'_{F,i}))$	t_F reclustering + $(\ln(1/\theta_i), \ln(t'_{F,i}))$
Events with 1 single jet	69.7%	72.8%	72.1%
Events with 2 jets	68.4%	72.4%	72.5%
Events with either 1 or 2 jets	69.1%	72.6%	71.4%

Table 6.1: The table summarises the binary accuracy achieved by the model when applying it to the evaluation set in each of the nine different situations.

We decided to pursue further investigations on the performance of our model for the particular case of events with one single jet, reclustered with the C/A algorithm and using the

86

$(\ln(1/\theta_i), \ln(t'_{F,i}))$ Lund plane coordinates. From table 6.1 it can be seen that for this case we obtained a binary accuracy of 72.8% with the evaluation set and using the default threshold value for the classifier (i.e., 0.5). For this particular instance, we decided to analyse the probability densities of the classifier for quenched and non-quenched jets. These are presented in figure 6.3a. Indeed we confirm that quenched jets tend to have classifier values close to one, but there is still a fraction of 24% of data points which yield values smaller than 0.5. On the other hand, vacuum jets do not exhibit such a peaked probability density close to zero as do quenched jets close to unity, but they still yield a larger fraction of data points to the left of the default threshold than to its right: namely, 69% of vacuum jets return classifier values smaller than 0.5. Conversely to these two values, we have a true positive rate of $1 - 0.24 = 0.76$ and a false positive rate of $1 - 0.69 = 0.31$, for the default threshold value of the classifier.

To evaluate how these values change when we vary the threshold of the classifier we can have a look at the ROC curve plotted in figure 6.3b. The overall look of this curve is very positive and, particularly, the AUC is approximately 0.80, which is a good result. Globally, we conclude, once again, that the results obtained with our model are promising.

(a)   (b)
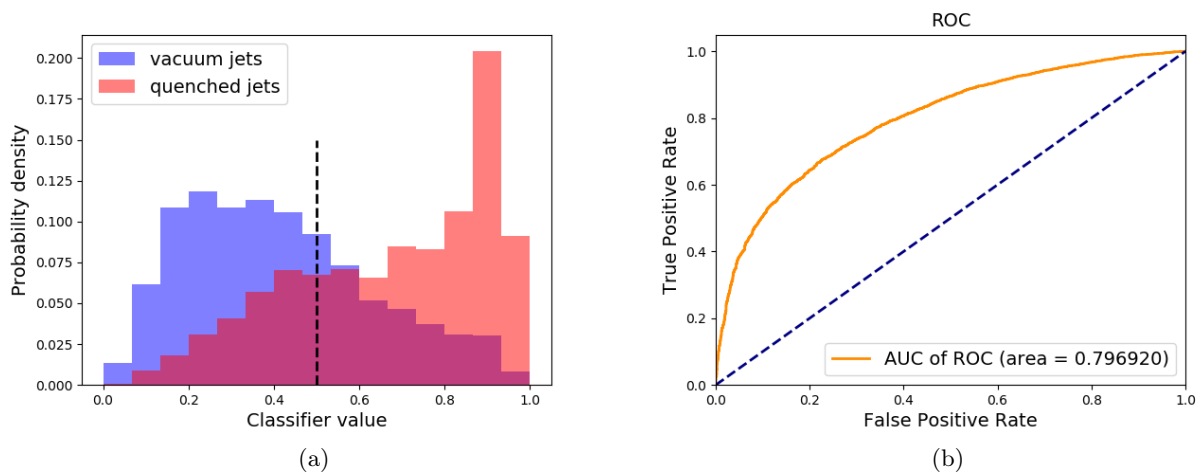
Figure 6.3: This figure presents in (a) the probability density of the classifier values for both quenched (red) and non-quenched (blue) jets and in (b) the corresponding ROC curve and its respective AUC, obtained with our model for single jet events with C/A reclustering and Lund plane coordinates $(\ln(1/\theta_i), \ln(t'_{F,i}))$.

Chapter 7

Concluding remarks

The goal of this master dissertation was to develop a set of tools which allows us to better distinguish between jets which have interacted with the QGP and those which have not. Instead of taking one of the two dichotomous approaches widely established in literature - either the manual crafting of high-level observables or the direct use of a ML model to extract information from jet images (or other input data with low-level physical processing) - we opted towards the adaptation of the approach presented in [80]. That means we resorted to the use of kinematical Lund planes. These present us with a visual representation of the phase space of emissions within jets. Indeed, much of the power of the Lund planes lies in their visual nature, which allows us to easily identify regions of the phase space which exhibit enhancements or suppressions in the presence of the QGP, when compared to the vacuum baseline. By offering that information in such a straightforward fashion, they can, on the one hand, motivate the construction of suitable observables, sensitive to the appropriate regions of the phase space. On the other hand, the data used to build these planes can also be input into a ML algorithm whose architecture should be optimised to secure the highest achievable performance in the task of discriminating between quenched and non-quenched jets. Thus, the building of these planes offers us a good starting ground for further investigations, helping to bridge the two (well established) leading approaches currently in use.

We started our investigations considering only events at parton level (chapter 4). This is a natural starting point since it avoids the inclusion of phenomena which are not calculable directly from perturbative QCD (notably hadronisation). Thus we can study and try to understand the physical effects in play at the most fundamental level, i.e., at the level of quarks and gluons. Our results from this chapter were comprehensive: we explored a variety of kinematical definitions for the $y-$axis of the Lund planes, different fillings procedures and also diverse reclustering algorithms. Amongst our principal conclusions is the observation that these kinematical Lund planes deliver marked differences between the radiation patterns of quenched and non-quenched jets, thus providing a good (visual) discrimination between the two classes of jets. Specifically, we have seen that the QGP spreads radiation throughout the plane, consistently to the soft, large-angle and hard, collinear regions of the phase space. It is also clear that, in the presence of the hot and dense medium, the momentum scale involved in the emissions is decreased with respect to vacuum and that there is an enhancement of the fraction of delayed emissions (i.e., emissions with large formation times). Furthermore, we have seen that the (average) explicit paths of the primary emissions along the splitting maps are a discriminant feature between quenched and non-quenched jets. In particular, these paths show (just as the radiation patterns themselves) the effect of medium-delayed radiation - especially for the large-angled, first emissions which are delayed, on average, by a factor of four. Overall, our conclusions regarding the kinematical Lund planes at parton level are that, following our expectation and desire, they provide a remarkable

discrimination power between quenched and non-quenched jets. Given that observation we expect that high discrimination performances are achieved either by the construction of the appropriate observables (sensitive to the relevant regions of the phase space) or by inputting the information used to build the planes into an adequate ML model.

Following our initial studies at parton level, we decided to investigate in what way our results and conclusions were altered by the inclusion of both hadronisation and ISR into the picture (chapter 5). Although this takes us a step closer to reality, it also adds a degree of uncertainty associated with the fact that hadronisation is modelled and not calculated from first principles of QCD. As a consequence of the inclusion of these two effects, our analysis of the hadronised samples shows a considerable smearing of the radiation patterns; there is a substantial loss of structure to them and the pronounced differences between quenched and non-quenched jets observed at parton level are significantly subdued. Notably, our proposed Lund plane, exploring the use of the logarithm of the formation time of the splittings for the vertical axis, appears to retain more structure and discrimination power than does the tradition $(\ln(1/\theta), \ln(z\theta))$ Lund plane definition.

The final task of this dissertation consisted in the implementation of a RNN model capable of classifying jets as quenched or non-quenched, when input with the paths of primary emissions of the jets along the kinematical Lund plane. We focused this investigation on the comparative study of the performance of our model using different subsets of events, different Lund plane coordinates ($y \equiv \ln(z\theta)$ and $y \equiv \ln(t'_F)$) and reclustering trees obtained with different jet algorithms. Our conclusions are that the use of different event subsets and different reclustering algorithms delivers no perceivable variations in the overall classification performance of the RNN. On the other hand, the kinematical variable chosen for the $y-$axis of the Lund plane has a clear influence on performance. In particular, the results corroborate our visual assessments of the primary Lund planes for the hadronised samples, from which we conclude that our plane proposal yields a better discrimination between quenched and non-quenched jets.

The maximum binary accuracy achieved by our model is around 72%, meaning that the RNN is able to correctly classify a jet as either quenched or non-quenched 72 out of 100 times. Further, we have also seen that while our model does not yield a perfect discrimination between the two classes, the ROC curve looks very good and the AUC has a value close to 0.80. We regard these results as positive and encouraging, specially in light of the fact that these are our first, preliminary benchmark studies.

Of course, the work done is by no means finished. There are still a lot of promising ideas, approaches and avenues of research which need to be pursued in continuation of this work. One notable improvement, which was mentioned repeatedly along the text, is the inclusion of grooming and trimming procedures, to clean the jet structure of uninteresting and irrelevant emissions. The reason to leave the implementation of such techniques out of this dissertation was that our main goal was to establish a benchmark of what is achievable without the inclusion of additional pre-processing techniques. Inevitably, these will introduce into the analysis an additional degree of arbitrariness, which we wanted to avoid at this point. Moreover, it is clear that the inclusion of underlying event and medium recoil effects into these samples is imperative to allow the proper comparison with results from real-life events in collider experiments. Also to that end, the possibility of including the simulation of detector effects should also be analysed. When it comes to ML there are also further avenues of research that remain open. For instance, an analysis of the discrimination power that can be achieved by using only the first splittings or all of the splittings within each jet should be conducted. Further, the study of different architectures (other than a RNN) could also prove informative and beneficial.

All things considered, while the work conducted was comprehensive in many aspects and provides a baseline for future endeavours, much remains to be investigated and explored. In
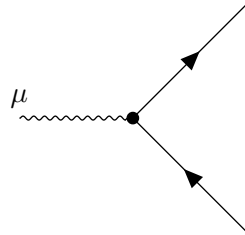
the end, we hope to show that the Lund planes can potentially be used as observables in the discrimination between quenched and non-quenched jets at the experimental level.
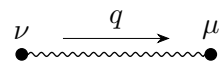
Appendix A

Useful Feynman rules

As explained in chapter 1, a set of Feynman rules can be derived from a given Lagrangian [1]. In this appendix we do not intend to proceed with such derivation and instead limit ourselves to the presentation of various Feynman rules associated with the QED and QCD Lagrangians given in equations (1.1) and (1.2).

In QED, there is a single interaction vertex which comes from the $-\bar{\psi} e_f \gamma^\mu A_\mu \psi$ term of the Lagrangian and can be written as:
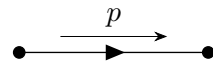
$$\mu \qquad\qquad = -i e_f \gamma^\mu \,. \qquad\qquad (A.1)$$

In a Feynman diagram, internal lines are represented by a propagator. In QED, there are two different propagators: one associated in the (photon) gauge field and the other with the charged fermions. The photon propagator is given by:

$$\nu \xrightarrow{\quad q \quad} \mu \qquad = -i \frac{g_{\mu\nu}}{q^2 + i\epsilon} \,, \qquad\qquad (A.2)$$

where $g_{\mu\nu}$ is the Minkowski metric $(+, -, -, -)$, q^2 is the photon's virtuality (i.e., its four-momentum squared) and $i\epsilon \to i0^+$ to ensure that this propagator is a retarded Green function. On the other hand, the charged fermion propagator is:
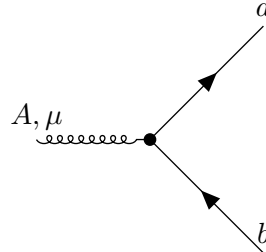
$$\xrightarrow{\quad p \quad} \qquad = i \frac{\not{p} + m}{p^2 - m^2 + i\epsilon} \,, \qquad\qquad (A.3)$$

where $\not{p} = \gamma^\mu p_\mu$.

Also important to know is that external photon lines get polarisation vectors while external fermion lines get spinors such that:

- an incoming (or initial) photon has a polarisation vector $\varepsilon_\mu(p)$ while an outgoing (or final) photon has a polarisation vector $\varepsilon_\mu^*(p)$;

- an incoming or outgoing fermion has a spinor $u(p)$ or $\bar{u}(p)$, respectively, while an incoming (outgoing) anti-fermion has a spinor $\bar{v}(p)$ ($v(p)$).

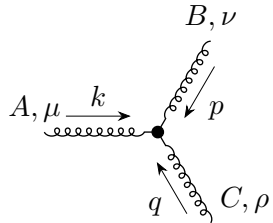As discussed in chapter 1, QCD is a much richer theory than QED and such is reflected in its Lagrangian. In this case, there are three possible interaction vertices. The first vertex comes from the interaction between quarks and gluons (qqg) described by the $\bar{\psi}_a g_s \gamma^\mu t_{ab}^A \mathcal{A}_\mu^A \psi_b$ term in the Lagrangian. This interaction vertex is analogous to that resulting from the interaction between photons and charged fermions in QED (equation (A.1)), and can be presented as:
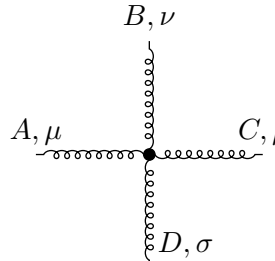

$$= i g_s \gamma^\mu t_{ab}^A. \qquad \text{(A.4)}$$

The second interaction vertex is the triple-gluon vertex, which comes from the cubic term in the gluon field: $g_s f^{ABC} \left(\partial^\mu \mathcal{A}^{A,\nu} \right) \mathcal{A}_\mu^B \mathcal{A}_\nu^C$ such that

$$= g_s f^{ABC} \left[g^{\mu\nu}(k-p)^\rho + g^{\nu\rho}(p-q)^\mu + g^{\rho\mu}(q-k)^\nu \right]. \qquad \text{(A.5)}$$
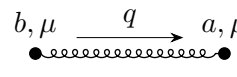
The final interaction vertex is the 4-gluon vertex arising from the quartic term in the gluon field that results from the product of two $g_s f^{ABC} \mathcal{A}_\mu^B \mathcal{A}_\nu^C$ terms and produces a vertex such that:

$$\text{(A.6)}$$

$$= -i g_s^2 \left[f^{ABF} f^{CDF} \left(g^{\mu\rho} g^{\nu\sigma} - g^{\mu\sigma} g^{\nu\rho} \right) + f^{ACF} f^{BDF} \left(g^{\mu\nu} g^{\rho\sigma} - g^{\mu\sigma} g^{\nu\rho} \right) + \right.$$
$$\left. + f^{ADF} f^{BCF} \left(g^{\mu\nu} g^{\rho\sigma} - g^{\mu\rho} g^{\nu\sigma} \right) \right].$$

Note that this vertex is proportional to g_s^2 unlike the other two, making it a higher order term.

Thus, equations (A.4), (A.5) and (A.6) summarise the Feynman rules for the QCD interaction vertices. The only piece of the puzzle which is missing (and is relevant for this work) are the quark and gluon propagators. The gluon propagator is exactly the same as the photon propagator but since gluons carry colour charge, here we explicitly include a Dirac-δ to guarantee conservation of the fundamental colour between the two vertices,

$$= -i \frac{g_{\mu\nu}}{q^2 + i\epsilon} \delta_{ab}, \qquad \text{(A.7)}$$

where the indices a and b are colour indices. Similarly, the quark propagator is a specific case of the fermion propagator written for QCD but which carries colour charge. For that reason, for

the quark propagation we can make an explicit statement of colour conservation of the quark as it propagates between the two vertices,

$$b \xrightarrow{\quad p \quad} a \qquad = i \frac{\not{p} + m}{p^2 - m^2 + i\epsilon} \delta_{ab} \,. \tag{A.8}$$

It is also important to note that, just like for external photon lines, external gluon lines take polarisation such that: an incoming gluon has a polarisation vector $\varepsilon_\mu(p)$ while an outgoing gluon has a polarisation vector $\varepsilon_\mu^*(p)$.

Appendix B

Full Lund diagrams for the $\left(\ln\left(1/\theta\right),\ln\left(1/z\right)\right)$ construction

In order to reduce to the bulk size of section 4.1 of chapter 4 and because the behaviour of the full Lund diagrams closely resembles that of the primary Lund planes, we have decided to remove their presentation from the main text and instead place them in the present appendix, for consultation.

Figures B.1 and B.2 depict the radiation pattern of all emissions of gluon-initiated jets developed, respectively, in vacuum and in the presence of the QGP, for the three jet reclustering algorithms. Indeed it is straightforward to see the similarities to the radiation pattern in the primary Lund planes (figures 4.4 and 4.5). The differences between in-medium jets and vacuum jets are presented in figure B.3.

In the case of quark-initiated jets, following the structure present in the text, we present only the results using the C/A and k_t algorithms, since the t_F algorithm yields and intermediate behaviour between these two. Figures B.4 and B.5 yield the radiation pattern of all emissions within these jets in vacuum and in medium, respectively, while figure B.6, depicts the difference between them.
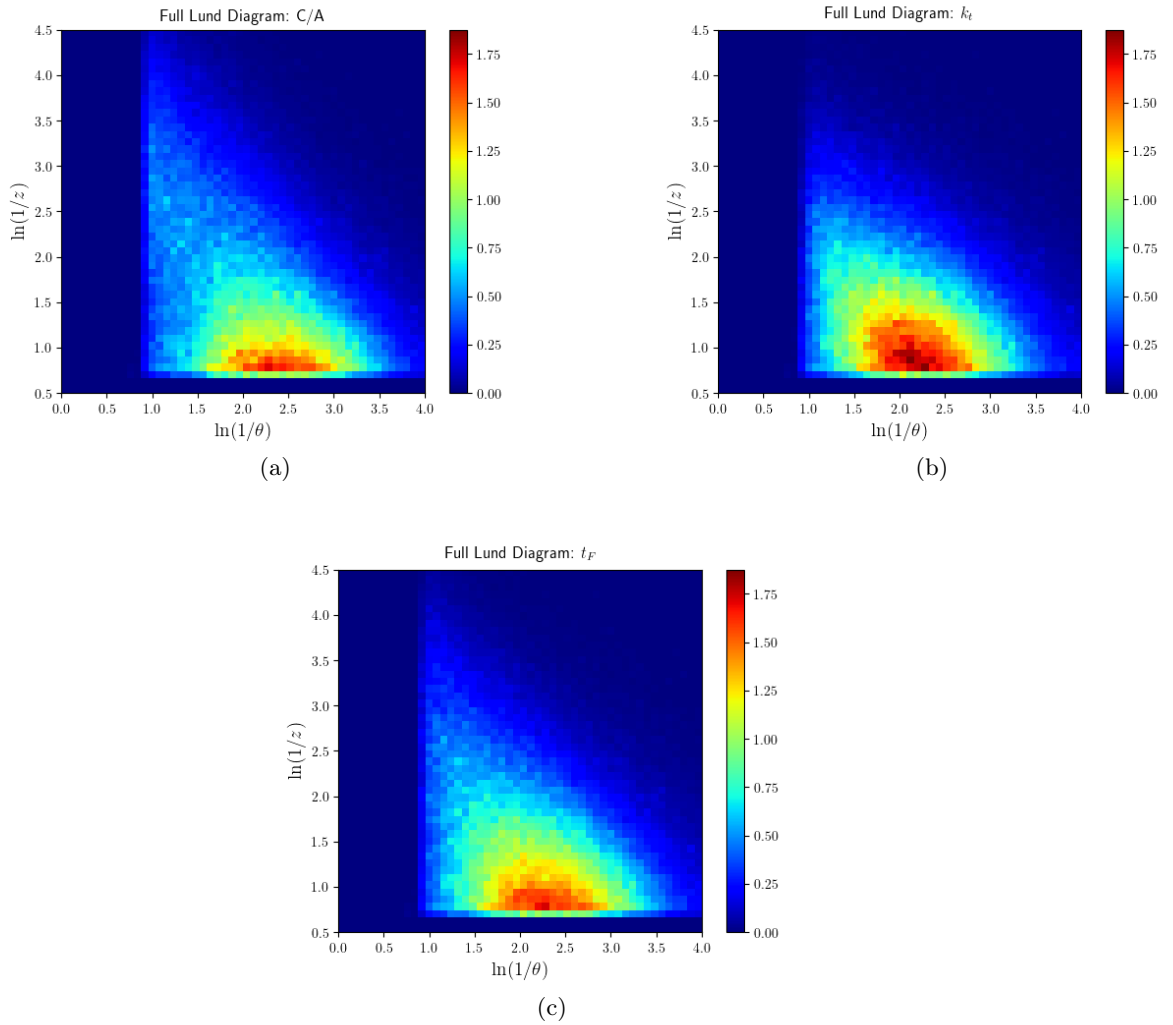
Figure B.1: Full Lund diagram for gluon-initiated jets developed in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

Figure B.2: Full Lund diagram for in-medium, gluon-initiated jets found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

Figure B.3: Lund planes that register the difference between the radiation pattern of all emissions with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm, (b) k_t algorithm and (c) t_F algorithm.

Figure B.4: Full Lund diagram for quark-initiated jets in vacuum, found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

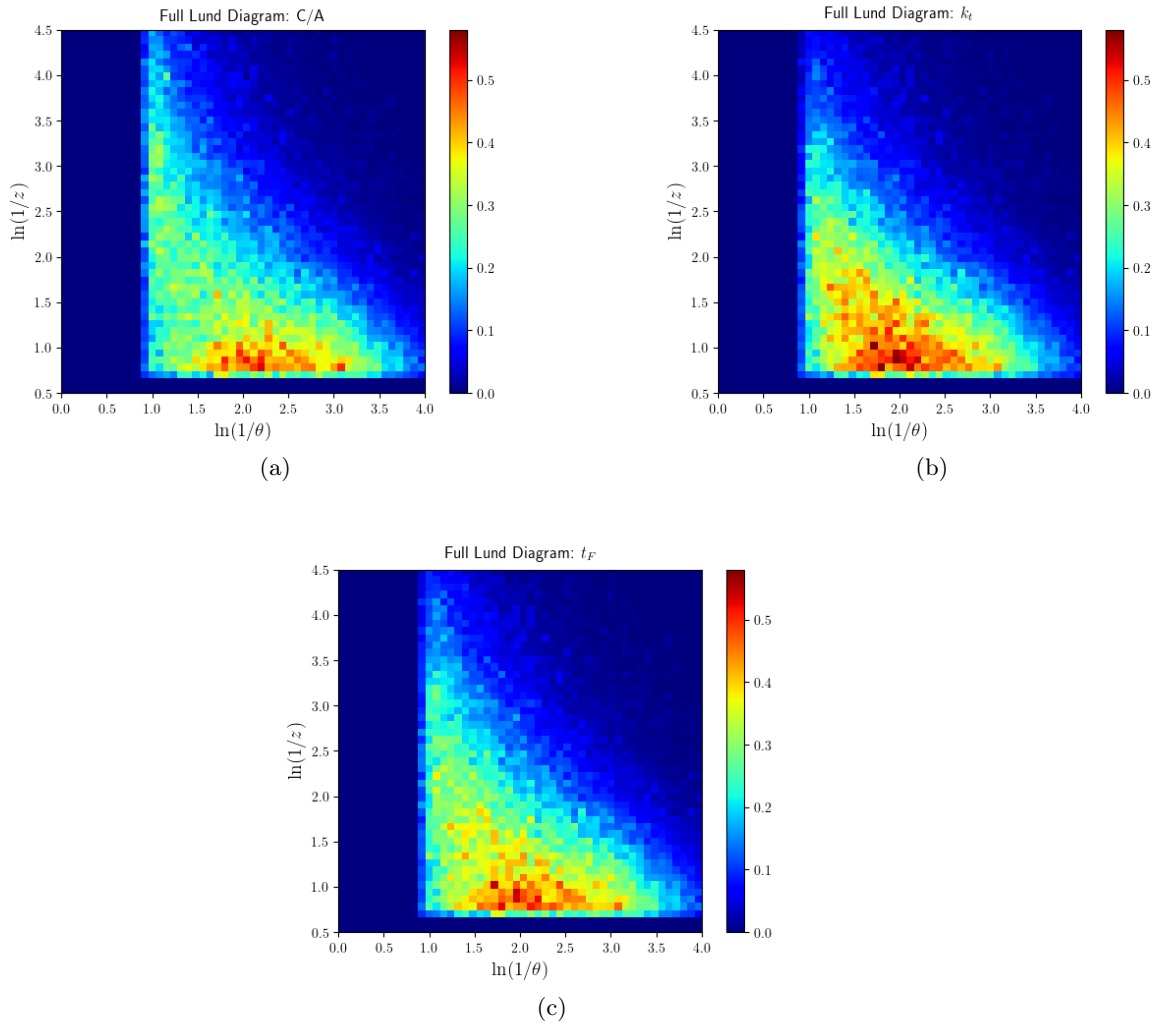

Figure B.5: Full Lund diagram for in-medium, quark-initiated jets found with the anti-k_t algorithm and reclustered with the (a) C/A algorithm and (b) k_t algorithm.

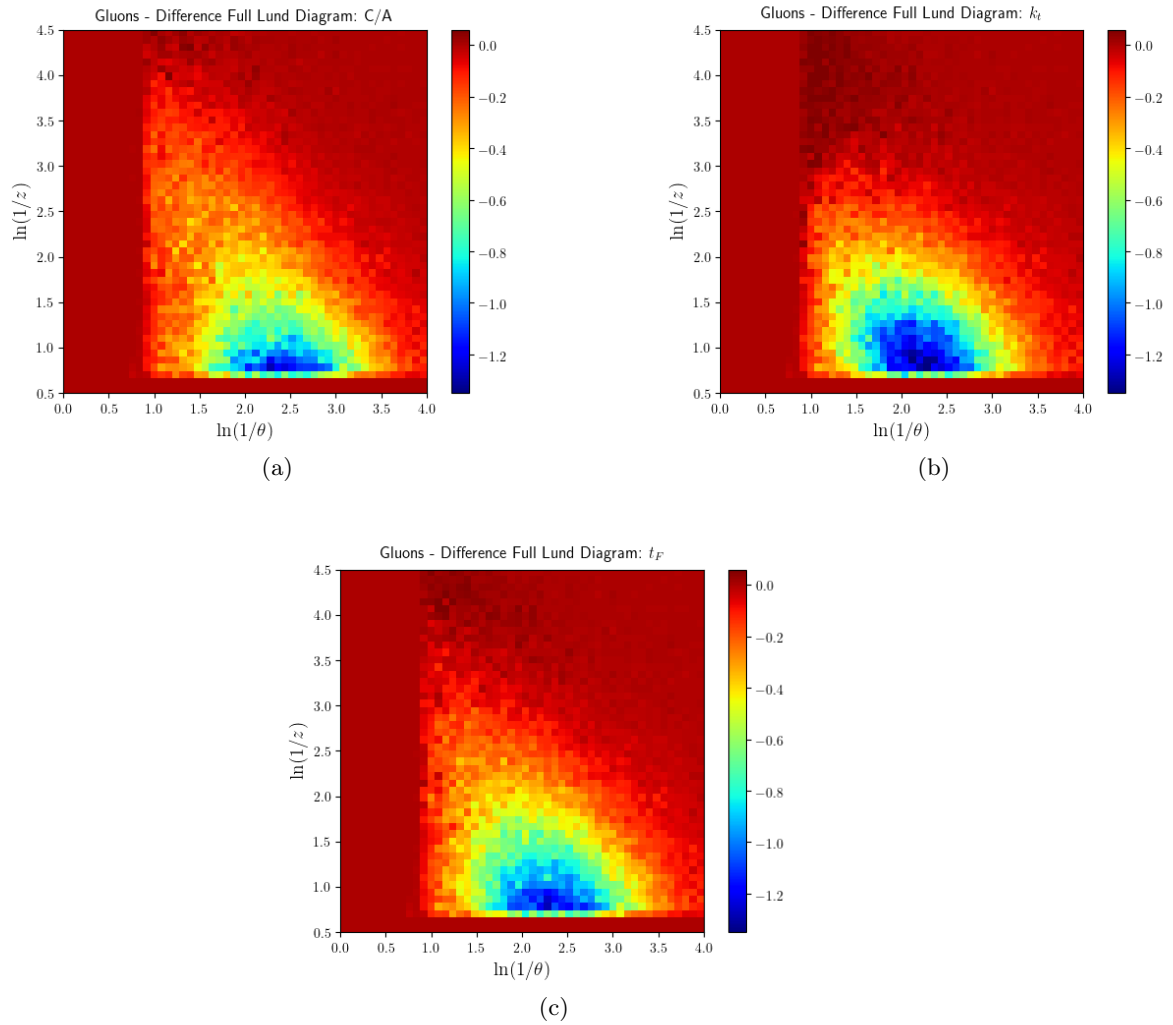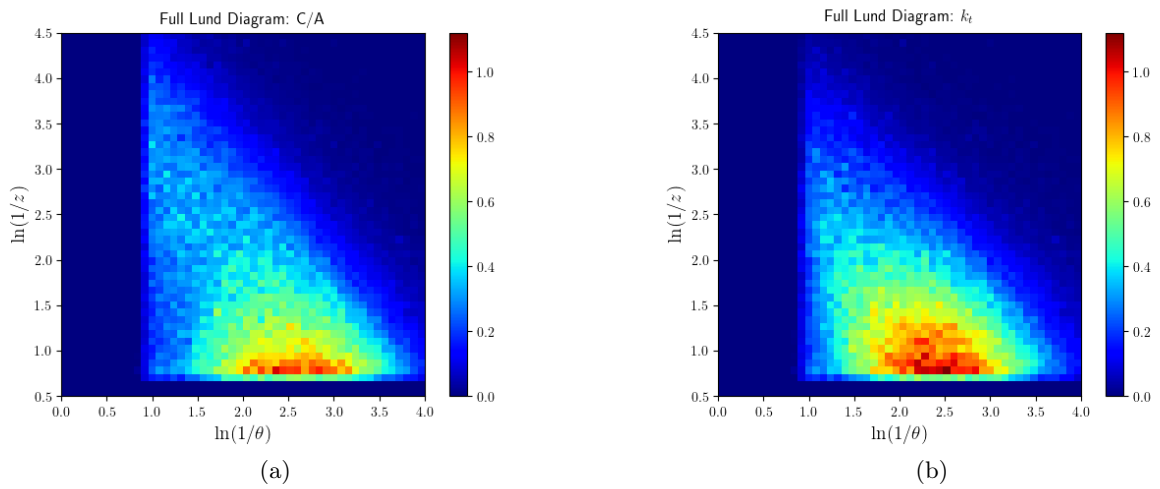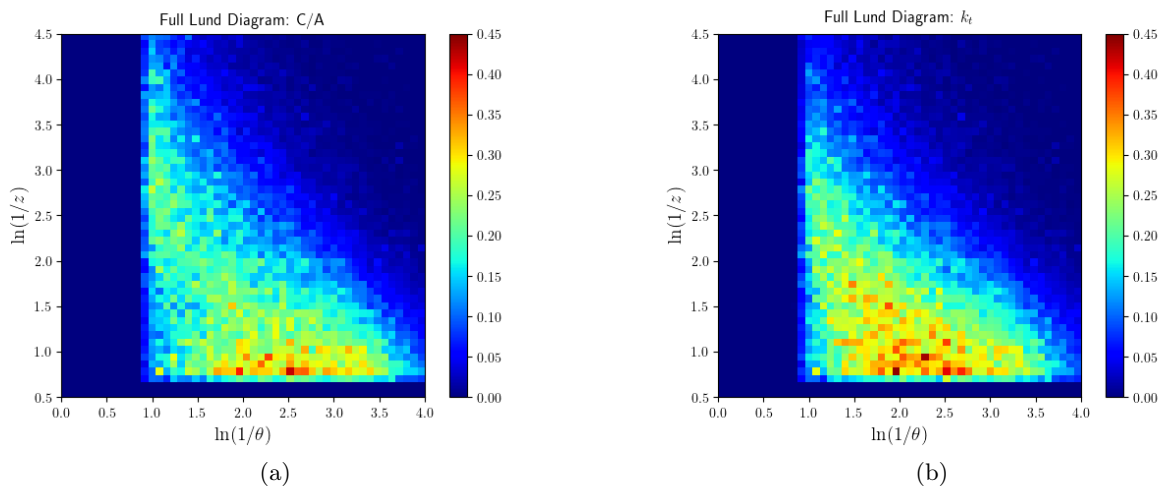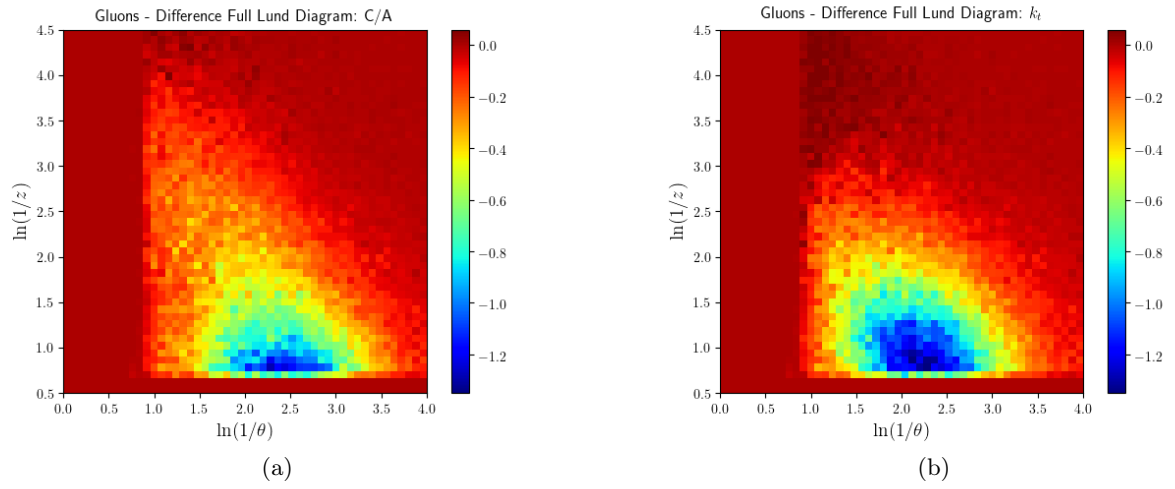Figure B.6: Lund planes that register the difference between the radiation pattern in the full Lund diagram with and without quenching effects, when the jets are reclustered with the (a) C/A algorithm and (b) k_t algorithm.

Bibliography

[1] Matthew D. Schwartz, *Quantum Field Theory and the Standard Model*, Cambridge University Press (2013).

[2] Maxwell, J. C., *On Physical Lines of Force*, Philosophical Magazine **90** (2010) 11.

[3] Maxwell J. C., *VIII. A dynamical theory of the electromagnetic field*, Phil. Trans. R. Soc. **155** (1865) 459.

[4] Maxwell, J. C. , *A Treatise on Electricity and Magnetism* (Cambridge Library Collection - Physical Sciences), Cambridge: Cambridge University Press (2010).

[5] S. Tomonaga, *On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields*, Progress of Theoretical Physics, **1** (1946) 27.

[6] J.S. Schwinger, *On quantum electrodynamics and the magnetic moment of the electron*, Phys. Rev. **73** (1948) 416.

[7] J.S. Schwinger, *Quantum Electrodynamics. I. A Covariant Formulation*, Phys. Rev. **74** (1948) 1439.

[8] R. P. Feynman, *The Theory of Positrons*, Phys. Rev. **76** (1949) 749.

[9] R. P. Feynman, *Space-Time Approach to Quantum Electrodynamics*, Phys. Rev. **76** (1949) 769.

[10] R. P. Feynman, *Mathematical Formulation of the Quantum Theory of Electromagnetic Interaction*, **80** (1950) 440.

[11] C. N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96**, 191 (1954).

[12] S. L. Glashow, *Partial-symmetries of weak interactions*, Nucl. Phys. **22**, 519 (1961).

[13] A. Salam and J. C. Ward, *Electromagnetic and weak interactions*, Phys. Lett. **13**, 168 (1964).

[14] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19**, 1264 (1967).

[15] A. Salam, in N. Svartholm, ed. *Elementary Particle Physics: Relativistic Groups and Analyticity*, Eight Nobel Symposium (Almqvist and Wiksell 1968), p. 367.

[16] F. J. Hasert *et al.*, *Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment*, Phys. Lett. B **46** (1973) 121-138 ; P. Musset *et al.*, *Etat des résultats expérimentaux sur la recherche des courants neutres*, J. Phys. (Paris) **11/12**, T34 (1973).

[17] C. Y. Prescott *et al.*, *Parity Nonconservation in Inelastic Electron Scattering*, Phys. Lett. B **77** (1978) 347 .

[18] M. Gell-Mann, *A schematic model of baryons and mesons*, Phys. Lett. **8**, 214 (1964).

[19] G. Zweig, *An SU_3 model for strong interaction symmetry and its breaking*, CERN preprint TH401 (1964).

[20] E. D. Bloom *et al.*, *High-Energy Inelastic $e-p$ Scattering at 6° and 10°*, Phys. Rev. Lett. **23**, 930 (1969); M. Briedenbach *et al.*, *Observed Behavior of Highly Inelastic Electron-Proton Scattering*, Phys. Rev. Lett. **23**, 935 (1969).

[21] O. W. Greenberg, *Spin and Unitary-Spin Independence in a Paraquark Model of Baryons and Mesons*, Phys. Rev. Lett. **13**, 598 (1964).

[22] M. Y. Han and Y. Nambu, *Three-Triplet Model with Double $SU(3)$ Symmetry*, Phys. Rev. B **139**, 1006 (1965).

[23] W. A. Bardeen, H. Fritzsch, and M. Gell-Mann, in *Scale and Conformal Symmetry in Hadron Physics*, R. Gatto, ed. (Wiley, New York, 1973), p. 139.

[24] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Non-Abelian Gauge Theories*, Phys. Rev. Lett. **30**, 1343 (1973).

[25] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30**, 1346 (1973).

[26] H.Fritzsch, M.Gell-Mann and H.Leutwyler, *Advantages of the color octet gluon picture*, Phys. Lett. B **47** (1973) 365 .

[27] S. Weinberg, *Non-Abelian Gauge Theories of the Strong Interactions*, Phys. Rev. Lett. **31**, 494 (1973).

[28] D. J. Gross and F. Wilczek, *Asymptotically Free Gauge Theories. I*, Phys. Rev. D **8**, 3633 (1973).

[29] `https://cds.cern.ch/record/799984/files/0401010.pdf`, visited on the 23rd of June of 2019, at 14:19.

[30] Michael E. Peskin and Daniel V. Schroeder, *An Introduction to Quantum Field Theory*, Addison-Wesley Publishing Company (1995).

[31] G. P. Salam, *Elements of QCD for hadron colliders*, in *High-energy physics. Proceedings, 17th European School, ESHEP 2009, Bautzen, Germany, June 14-27, 2009* (2010), [arXiv:1011.5131].

[32] A. J. Larkoski, *An Unorthodox Introduction to QCD* (2017), [arXiv:1709.06195v2].

[33] `http://pdg.lbl.gov/2019/reviews/rpp2018-rev-monte-carlo-techniques.pdf`

[34] G. Corcella *et al.*, *HERWIG 6: an event generator for Hadron Emission Reactions With Interfering Gluons (including supersymmetric processes)*, JHEP **01** (2001) 010, [arXiv:hep-ph/0011363v1].

[35] T. Sjöstrand, S. Mrenna and P. Skands, *PYTHIA 6.4 Physics and Manual*, JHEP **05** (2006) 026, [arXiv:hep-ph/0603175v2].

[36] T. Gleisberg *et al.*, *SHERPA 1.a, a proof-of-concept version*, JHEP **02** (2004) 056, [arXiv:hep-ph/0311263].

[37] S. Agostinelli *et al.*, *GEANT4: A Simulation toolkit*, NIM **A506** (2003) 250.

[38] T. Sjöstrand, *Monte Carlo Generators*, High-energy physics. *Proceedings, European School, Aronsborg, Sweden, June 18-July 1, 2006* (2006), [arXiv:hep-ph/0611247].

[39] A. Buckley *et al.*, *General-purpose event generators for LHC physics*, Phys. Rept. **504**, 145-233 (2011), [arXiv:1101.2599v1].

[40] S. Gieseke, P. Stephens, B. Webber, *New formalism for QCD parton showers*, JHEP **12** (2003) 045, [arXiv:hep-ph/0310083].

[41] T. Sjöstrand, P. Z. Skands, *Transverse-momentum-ordered showers and interleaved multiple interactions*, Eur. Phys. J. **C39**, 129–154 (2005), [arXiv:hep-ph/0408302].

[42] `http://pdg.lbl.gov/2019/reviews/rpp2018-rev-qcd.pdf`

[43] G. Lemaître, *Expansion of the universe, A homogeneous universe of constant mass and increasing radius accounting for the radial velocity of extra-galactic nebulae*, MNRAS **91**, 483 (1931) [translated from Annals de la Societe Scientifique de Bruxelles **A47**, 49 (1927)].

[44] `https://home.cern/science/physics/early-universe`, visited on the 23rd of June of 2019 at 1:45p.m..

[45] `https://home.cern/science/physics/heavy-ions-and-quark-gluon-plasma`, visited on the 3rd of June of 2019 at 10:23 a.m..

[46] W. Busza, K. Rajagopal, and W. van der Schee, *Heavy Ion Collisions: The Big Picture, and the Big Questions*, Ann. Rev. Nucl. Part. Sci. **68** (2018) 339–376, [arXiv:1802.04801].

[47] The Brahms, Phenix, Phobos, and Star Collaborations, *First Three Years of Operation of RHIC*, Nucl. Phys. A **757** (2005) 1-283, [arXiv:nucl-ex/0603012].

[48] Yu. L. Dokshitzer, V. A. Khoze, A. H. Mueller and S. I. Troyan, *Basics of Perturbative QCD*, Editions Frontieres (1991).

[49] H. A. Andrews *et al.*, *Novel tools and observables for jet physics in heavy-ion collisions*, [arXiv:1808.03689].

[50] G. P. Salam, *Towards Jetography*, Eur. Phys. J. **C67** (2010) 637-686, [arXiv:0906.1833].

[51] G. Hanson et al. [SLAC-LBL Collaboration], *Evidence for Jet Structure in Hadron Production by e^+e^- Annihilation,* Phys. Rev. Lett. **35**, 1609 (1975).

[52] J. D. Bjorken and S. J. Brodsky, *Statistical Model for Electron-Positron Annihilation into Hadrons*, Phys. Rev. D 1, 1416 (1970).

[53] UA2 Collaboration (Banner, M. *et al.*), *Observation of Very Large Transverse Momentum Jets at the CERN anti-p p Collider* - Phys. Lett. B **118** (1982) 203-210, CERN-EP/82-141.

[54] UA2 Collaboration, Bagnaia, P., Banner, M. *et al.*, *Measurement of production and properties of jets at the CERN $\bar{p}p$ collider*, Z. Phys. C - Particles and Fields **20** (1983) 117-134.

[55] J. Schukraft, *Little bang at big accelerators: heavy ion physics from AGS to LHC*, AIP Conf. Proc. **531** (2000) 3-15.

[56] F. Antinori, A. Billmeier and J. Zaranek, *Introduction to ultrarelativistic heavy-ion physics at the CERN-SPS*, AIP Conf. Proc. **631** (2003) 294.

[57] H. Satz, *The SPS heavy ion programme*, Phys. Rept. **403-404** (2004) 33, [arXiv:hep-ph/0405051].

[58] A. Drees, *First Hints of Jet Quenching at RHIC*, Nucl. Phys. A **698** (2002) 331-340, [arXiv:nucl-ex/0105019].

[59] https://home.cern/news/press-release/cern/lhc-experiments-bring-new-insight-primordial-ur *LHC experiments bring new insight into primordial universe* (Press release) CERN, November 26, 2010, retrieved on the 11th of September 2019, at 6:29p.m..

[60] G. Aad *et al.* [ATLAS Collaboration], *Observation of a Centrality-Dependent Dijet Asymmetry in Lead-Lead Collisions at $\sqrt{s_{NN}} = 2.76$ TeV with the ATLAS Detector at the LHC*, Phys. Rev. Lett. **105** (2010) 252303, [arXiv:1011.6182v2].

[61] B. Abelev *et al.* [ALICE Collaboration], *Measurement of charged jet suppression in Pb-Pb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, JHEP **30** (2014) 013, [arXiv:1311.0633v2].

[62] S. Chatrchyan *et al.* [CMS Collaboration], *Observation and studies of jet quenching in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Rev. C **84** (2011) 024906, [arXiv:1102.1957v2].

[63] V. Khachatryan *et al.* [CMS Collaboration], *Measurement of transverse momentum relative to dijet systems in PbPb and pp collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, JHEP **01** (2016) 006, [arXiv:1509.09029].

[64] M. Aaboud *et al.* [ATLAS Collaboration], *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** (2019) 108 [arXiv:1805.05635v3].

[65] S. Chatrchyan *et al.* [CMS Collaboration], *Jet momentum dependence of jet quenching in PbPb collisions at $\sqrt{s_{NN}} = 2.76$ TeV*, Phys. Lett. B **712** (2011) 176, [arXiv:1202.5022].

[66] Y. Mehtar-Tani, J. G. Milhano, K. Tywoniuk, *Jet physics in heavy-ion collisions*, Int. J. of Mod. Phys. A **28**, 1340013 (2013), [arXiv:1302.2579v2].

[67] D. A. Appel, *Jets as a probe of quark-gluon plasmas*, Phys. Rev. **D33** (1986) 717.

[68] J. Blaizot and L. D. McLerran, *Jets in expanding quark-gluon plasmas*, Phys. Rev. **D34** (1986) 2739.

[69] M. Gyulassy and M. Plumer, *Jet quenching in dense matter*, Phys. Lett. **B243** (1990) 432.

[70] X.-N. Wang and M. Gyulassy, *Gluon shadowing and jet quenching in A+A collisions at $\sqrt{s} = 200A$ GeV*, Phys. Rev. Lett. **68** (1992) 1480.

[71] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and collider physics*, Cambridge University Press (1996).

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

[73] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better jet clustering algorithms*, JHEP **08** (1997) 001, [arXiv:hep-ph/9707323].

[74] S. Catani, Y. L. Dokshitzer, M. Olsson, G. Turnock and B. R. Webber, *New clustering algorithm for multi-jet cross-sections in e^+e^- annihilation*, Phys. Lett. B **269** (1991) 432 .

[75] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross sections in deep-inelastic*, [arXiv:hep-ph/9907280]; M. Wobisch, *Measurement and QCD analysis of jet cross sections in deep-inelastic positron proton collisions at $\sqrt{s} = 300$GeV*, DESY-THESIS-2000-049.

[76] S. Catani, Y. L. Dokshitzer, M. H. Seymour and B. R. Webber, *Longitudinally invariant K(t) clustering algorithms for hadron hadron collisions*, Nucl. Phys. B **406**, 187 (1993).

[77] S. D. Ellis and D. E. Soper, *Successive Combination Jet Algorithm For Hadron Collisions*, Phys. Rev. D **48**, 3160 (1993), [arXiv:hep-ph/9305266].

[78] J. Casalderrey-Solana, J. G. Milhano, and P. Q. Arias, *Out of Medium Fragmentation from Long-Lived Jet Showers*, Phys.Lett. **B710** (2012) 175–181, [arXiv:1111.0310].

[79] G. C. Blazey *et al.*, *Run II Jet Physics: Proceedings of the Run II QCD and Weak Boson Physics Workshop*, (2000), [arXiv:hep-ex/0005012].

[80] F. A. Dreyer, G. P. Salam and G. Soyez, *The Lund Jet Plane*, JHEP **12** (2018) 064, [arXiv:1807.04758].

[81] B. Andersson, G. Gustafson, L. Lonnblad, and U. Pettersson, *Coherence Effects in Deep Inelastic Scattering*, Z. Phys. C **43** (1989) 625.

[82] K. C. Zapp, F. Krauss and U. A. Wiedemann, *A perturbative framework for jet quenching*, JHEP **03** (2013) 080, [arXiv:1212.1599].

[83] J. Mccarthy, M. Minsky, N. Rochester and C.E. Shannon, *A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence*, AI Magazine **27** (2006).

[84] François Chollet, *Deep Learning with Python*, Manning Publications, 1st edition (2017).

[85] Ian Goodfellow, Yoshua Bengio and Aaron Courville, *Deep Learning*, MIT Press (2016).

[86] F. Chollet, "Keras". `https://keras.io`, 2015.

[87] M. Abadi *et al.*, *TensorFlow: Large-scale machine learning on heterogeneous systems*, 2015. Software available from `https://www.tensorflow.org`.

[88] G. Hinton, N. Srivastava, and K. Swersky, *Neural Networks for Machine Learning - Lecture 6a - Overview of mini-batch gradient descent*, 2012. [`http://www.cs.toronto.edu/~tijmen/csc321/slides/lecture_slides_lec6.pdf`]