

Disentangling Jet Modification

Jasmine Brewer,^{a,b} Quinn Brodsky^{a,*} and Krishna Rajagopal^{a,b}

^a*Massachusetts Institute of Technology,
77 Massachusetts Avenue, Cambridge, USA*

^b*Center for Theoretical Physics, Massachusetts Institute of Technology,
77 Massachusetts Avenue, Cambridge, USA*

E-mail: jtbrewer@mit.edu, quinnb@mit.edu, krishna@mit.edu

Jet modification in heavy-ion collisions is an important probe of the nature and structure of the quark-gluon plasma (QGP) produced in these collisions and also encodes information about how the wakes that jets excite in a droplet of QGP form and relax. However, in experiment, one cannot know what a particular jet seen in a heavy ion collision would have looked like without quenching, making it difficult to interpret measurements in terms of individual jet modification. The goal of this Monte Carlo study is to gain insight into the modification of jet observables using the hybrid strong/weak coupling model of jet quenching as a test bed. In this Monte Carlo study (but not in experiment) it is possible to watch *the same jet* as it evolves in vacuum or in QGP. We use this ability to disentangle the effects of modification of individual jets in heavy ion collisions vs. the effects of differing selection bias on the distribution of two observables: the fractional energy loss and the groomed ΔR . We find that in the hybrid model the distribution of groomed ΔR appears to be unmodified in a sample of jets selected after quenching, as in heavy ion collisions, and confirm that this lack of modification arises because of a selection bias toward jets that lose only a small fraction of their energy. If instead we select a sample of jets in a way that avoids this bias, and then follow these selected jets as they are quenched, we show that there is, in fact, a substantial modification of the ΔR of individual jets. Furthermore, we show that this jet modification is principally due to the incorporation of particles coming from the wake that the parton shower excites in the plasma as a component of what an experimentalist reconstructs as a jet. The effects we discuss, both those due to selection bias and those due to jet modification coming from the wake in the plasma, are substantial in magnitude, suggesting that our qualitative conclusions are more general than the Monte Carlo study in which we obtain them.

*HardProbes2020
1-6 June 2020
Austin, Texas*

*Speaker

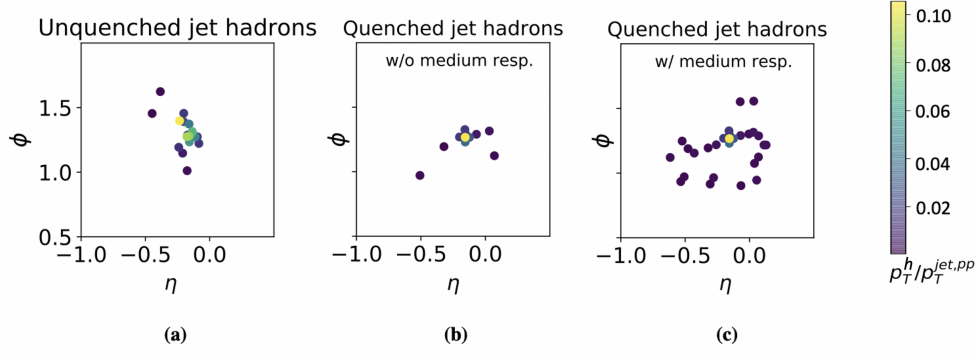


Figure 1: (a) (η, ϕ) distribution of hadrons in an unquenched jet; (b) the same jet, after quenching, leaving out hadrons originating from the response of the medium, e.g. the wake the jet excites; (c) the same jet, after quenching, with medium response. Color indicates hadron p_T as a fraction of the p_T of this unquenched jet.

Introduction: Jet quenching is an important probe of the QGP produced in heavy ion collisions, but quantitative comparisons are difficult because the modification of jets by the plasma necessarily introduces biases in the selection of jets. In particular, since the production rate of jets drops rapidly as a function of jet energy, when one selects jets with a given energy that have propagated through QGP – that have been quenched – one is selecting a sample of jets biased toward those that have lost the least energy. We use a Monte Carlo study to see that this also biases in favor of selecting those jets that are least modified in other respects, and to show that those jets which lose more energy are significantly modified in other respects too. In the hybrid model [1, 2], it is possible to compare a jet as it is before or after quenching. This provides us the opportunity to disentangle jet modification from selection bias by looking at a quenched jet and then matching it to its unquenched version. We can either select a sample of quenched jets (as experimentalists do in heavy ion collisions) and compare them to the same jets as they would have been in the absence of any QGP, e.g. unquenched, or we can select a sample of unquenched jets (as experimentalists do in proton-proton collisions) and watch what would have happened to those same jets if they were quenched. Because selecting quenched jets with a given energy introduces selection bias, the comparisons we obtain via these two different procedures look very different. Hence, drawing conclusions without first disentangling jet modification from selection bias can lead to misconceptions.

Methods: In the hybrid model, we start with a vacuum PYTHIA shower, and embed it in an evolving boost-invariant hydrodynamic medium [1]. The energy loss of partons in the medium is determined by a formula inspired by strongly-coupled energy loss in holography [1]. By turning energy loss off, one can look at the unquenched version of any individual quenched jet in the Monte Carlo ensemble. In this analysis, hadronization effects and the wake that the jet leaves behind in the medium through which it propagates were included [1]. In this Monte Carlo study, we can distinguish between particles in the reconstructed jet coming from the PYTHIA shower and particles produced from medium response. We reconstruct jets using the anti- k_T algorithm with $R = 0.4$. We compare jets in the quenched and unquenched versions of the same event, matching those that are closest in the (η, ϕ) plane and within $\sqrt{\Delta\eta^2 + \Delta\phi^2} \leq 0.4$ of each other. Fig. 1 shows an example.

In this way we can obtain a sample of the same jets before and after quenching. There are two ways to do this, which makes it possible to study how jet selection, and jet modification each impact

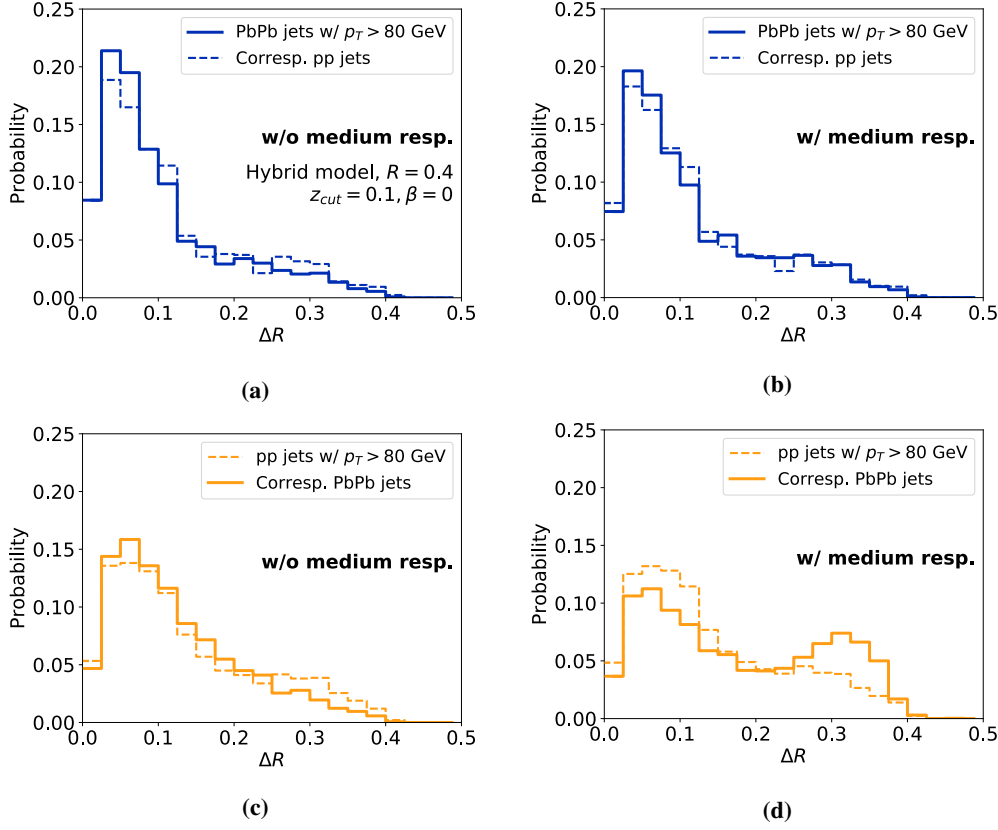


Figure 2: Softdrop ΔR distributions for Quench-then-Select Method (blue; upper panels) and Select-then-Quench Method (orange; lower panels), in both cases without (left panels) and with (right panels) medium response. Upper panels are PbPb jets without (a) and with (b) medium response, selected with $p_T > 80$ GeV (dashed) and the same jets before quenching (solid). Lower panels are pp jets selected with $p_T > 80$ GeV (dashed), and the same jets (solid), without (c) and with (d) medium response.

observables. First, we select quenched (heavy ion) jets whose transverse momentum p_T lies above a threshold cut, which in our Monte Carlo study we can do with or without the particles coming from medium response included in the jets, and then for each quenched jet we use the matching procedure to find the corresponding unquenched (proton-proton) jet. We'll call this the "Quench-then-Select Method". In the alternative "Select-then-Quench Method", we select proton-proton jets whose p_T lies above a threshold and then use the matching procedure to find the heavy ion jet that the pp jet becomes after being quenched according to the hybrid model.

Results: If the effect of the jet selection were small, one might anticipate that these methods would produce similar results since in both cases we compare quenched and unquenched versions of the same jet. However, by plotting various observables, we see that this is not at all the case. We show the Soft Drop ΔR distributions for jets with $z_{cut} = 0.1$ and $\beta = 0$ [3] in Fig. 2, with and without medium response. The plots using the Select-then-Quench method (Figs. 2a and 2b) are similar to what has been seen at the partonic level in previous hybrid model studies of jets selected in heavy ion collisions [2], and seem to indicate that quenching hardly modifies the ΔR distribution. The lower panels in Fig. 2, though point to a completely different conclusion.

Figs. 2a and 2b indeed show that the ΔR distribution for jets selected in PbPb collisions is

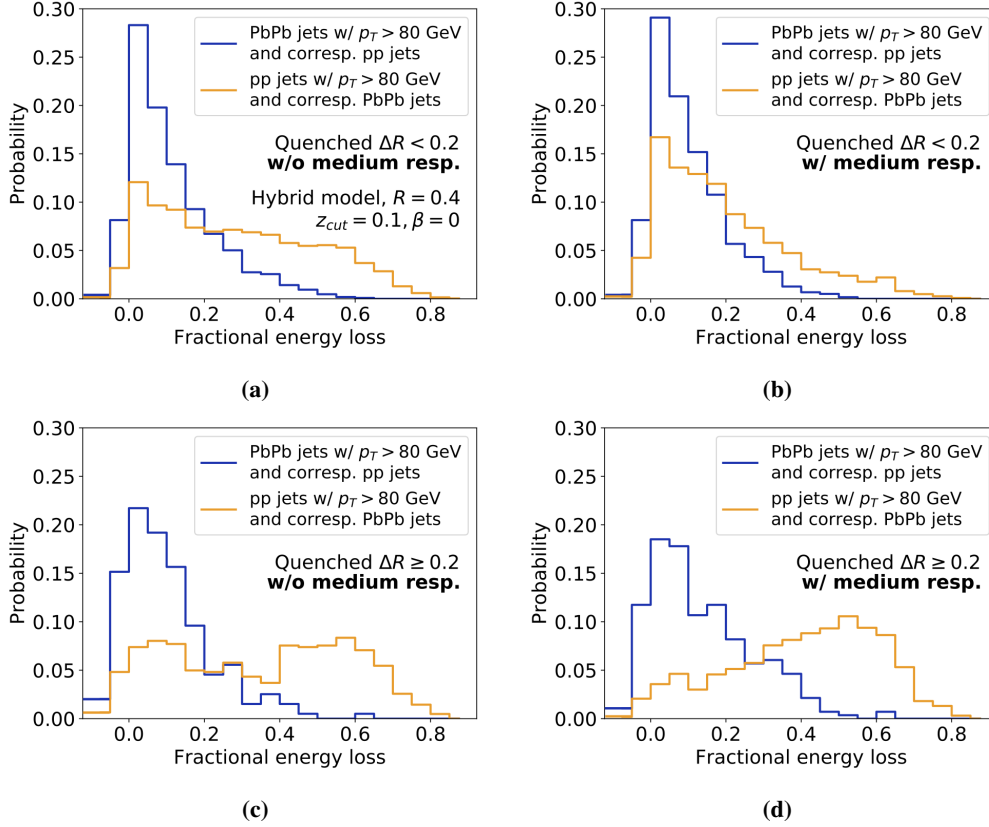


Figure 3: Fractional energy loss for Quench-then-Select (blue) and Select-then-Quench (orange) Methods, without (left panels) and with (right) medium response, for $\Delta R > 0.2$ (upper panels) and $\Delta R \leq 0.2$ (lower).

not much modified by quenching. However, selecting heavy ion jets with $p_T > 80$ GeV yields a sample of jets with a strong bias toward selecting those jets that do not lose much p_T , since the probability to produce a proton-proton jet falls rapidly with p_T . This bias is well known, but we confirm it in Fig. 3. What we are seeing in the upper panels of Fig. 2 is that the ΔR distribution of those jets which lose only a small fraction of their energy is not much modified by quenching. In Figs. 2c and 2d, we use the Select-then-Quench Method, selecting a sample of jets in proton-proton collisions and then following what happens to *those* jets after quenching – regardless of how much energy they lose. By selecting unquenched jets we eliminate the selection bias favoring jets that lose little energy. And, with the selection bias eliminated, we see that quenching actually very substantially modifies the ΔR distribution – when we include particles coming from the response of the medium to the jet. This rather dramatic modification of the ΔR distribution is wiped out in the upper-right panel by selection bias: in the upper-right panel we are looking at jets whose energy is little modified, and so neither is their ΔR . From the lower panels of Fig. 2, we conclude that ΔR is in fact substantially modified by quenching on a jet-by-jet basis, contrary to what one might have been tempted to conclude from the upper panels – in which the effects of jet modification are obscured by the effects of selection bias.

In Fig. 3 we confirm these conclusions by looking at the distribution of the fractional energy loss of jets in the Quench-then-Select Method (blue curves) and the Select-then-Quench Method

(green curves), with and without medium response, for jets with $\Delta R < 0.2$ and $\Delta R > 0.2$. In all cases, selecting quenched jets (Quench-then-Select) yields a sample of jets that have lost less energy than if one selects unquenched jets and then quenches them (Select-then-Quench). We can furthermore compare the upper (small ΔR) and lower (large ΔR) panels of Fig. 3 and see that the enhancement of jets at large ΔR caused by quenching that we see in Fig. 2d is coming from jets that lose a substantial fraction of their energy. These are exactly the jets that are missed in the Quench-then-Select Method due to its strong and intrinsic selection bias, as described earlier.

Conclusion and Discussion: In the hybrid model, quenching modifies the Soft Drop ΔR of individual jets substantially. The jets whose ΔR is substantially modified (see Fig. 2d) are those which lose a large fraction of their energy (see Fig. 3d). Selecting a jet sample using a cut on the jet p_T in heavy ion collisions creates a bias towards jets that lose very little energy. These are the jets whose ΔR is not substantially modified. By selecting a jet sample using a cut on the jet p_T in proton-proton collisions and considering the quenched versions of these jets, we remove the bias toward less modified jets and see that the ΔR of individual jets is substantially modified in the hybrid model. We furthermore see that this modification is to a significant degree coming from the response of the medium to the jet. This indicates that, unlike for jets in vacuum, the ΔR of jets in heavy ion collisions does not give us access in any straightforward way to the first hard splitting during the formation of the jet.

Because the effects of selection bias and jet modification on the observables that we have looked at are so large, we expect our qualitative conclusions to generalize far beyond the model that we have used as a testbed in which to disentangle these effects. That said, what we have done is intrinsically a Monte Carlo study, not something that can be done in the same way in an analysis of experimental data. Our results motivate experimental analyses of jet samples selected in a way that includes no bias toward small fractional energy loss — for example Z+jet samples with selection cuts imposed only on the Z — and which includes those jets that have lost a substantial fraction of their energy. It would be very interesting to look at the ΔR distribution in such a sample.

Acknowledgments: We are grateful to Dani Pablos for providing us with hybrid model samples and much helpful advice and to Yen-Jie Lee, Andrew Lin, Guilherme Milhano, James Mulligan, and Jesse Thaler for helpful conversations. This work was supported by the U.S. Department of Energy, Office of Science, Office of Nuclear Physics grant DE-SC0011090.

References

- [1] J. Casalderrey-Solana, D. C. Gulhan, J. G. Milhano, D. Pablos and K. Rajagopal, JHEP **10**, 019 (2014) [arXiv:1405.3864 [hep-ph]]; JHEP **03**, 053 (2016) [arXiv:1508.00815 [hep-ph]]; JHEP **03**, 135 (2017) [arXiv:1609.05842 [hep-ph]]; Z. Hulcher, D. Pablos and K. Rajagopal, JHEP **03**, 010 (2018) [arXiv:1707.05245 [hep-ph]]; J. Casalderrey-Solana, Z. Hulcher, G. Milhano, D. Pablos and K. Rajagopal, Phys. Rev. C **99**, 051901 (2019) [arXiv:1808.07386 [hep-ph]].
- [2] J. Casalderrey-Solana, G. Milhano, D. Pablos and K. Rajagopal, JHEP **01**, 044 (2020) [arXiv:1907.11248 [hep-ph]]; [arXiv:2002.09193 [hep-ph]].
- [3] A. J. Larkoski, S. Marzani, G. Soyez and J. Thaler, JHEP **05**, 146 (2014) [arXiv:1402.2657 [hep-ph]].