Available on CMS information server

CMS AN-2008/098

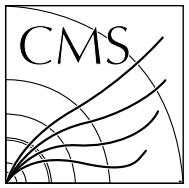

The Compact Muon Solenoid Experiment

Analysis Note

The content of this note is intended for CMS internal use and distribution only

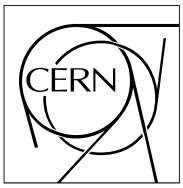

03 November 2008 (v2, 04 November 2008)

Muon Identification in CMS

M. Mulders

CERN, Geneva, Switzerland

I. Bloch, E. James

Fermilab, Batavia, Illinois, USA

A. Everett

Purdue University, West Lafayette, Indiana, USA

D. Barge, C. Campagnari, P. Kalavase, V. Krutelyov, D. Kovalskyi, J. Ribnik

University of California, Santa Barbara, USA

N. Amapane

Università di Torino e Sezione dell'INFN, Torino, Italy

Abstract

We describe the algorithms for muon identification in CMS as of October 2008, and we document their performance.

1 Introduction

In this note we describe the algorithms for muon identification in CMS as of October 2008. This note is complementary to a separate note that describes the algorithms for reconstruction of muon objects at CMS [1]. Here we discuss possible selection criteria for Global Muons and Tracker Muons. Standalone Muons are not addressed in this note. The object of this note is not to define the best possible muon selection criteria, but rather to explore the range of possibilities with the understanding that these issues will need to be re-addressed first with magnet-on cosmics and then collision data.

The selections described in this note are accessed through the member function of the reco::Muon object

```
bool isGood( SelectionType type = AllArbitrated )
```

where the various `SelectionTypes` are discussed in this note. We advocate that silicon-track quality cuts be also applied in analysis, see Section 5.2. For technical reason, these track quality cuts are not included in the `isGood()` function[1)]. The user should then apply track quality requirements externally.

2 Data samples

The main data samples [2] used for these studies consist of Monte Carlo single muons, pions, and kaons of $p_T < 200$ GeV initially generated, simulated, and reconstructed in CMSSW 2.0.6; these events were then partially reprocessed with CMSSW 2.0.11 with special tags:

```
V07-02-08          DataFormats/MuonReco
CMSSW_2_0_11       RecoMuon/Configuration
CMSSW_2_1_0_pre10 RecoMuon/CosmicMuonProducer
CMSSW_2_1_0_pre10 RecoMuon/DetLayers
CMSSW_2_1_0_pre10 RecoMuon/GlobalMuonProducer
CMSSW_2_1_0_pre10 RecoMuon/GlobalTrackFinder
CMSSW_2_1_0_pre10 RecoMuon/GlobalTrackingTools
CMSSW_2_1_0_pre10 RecoMuon/L2MuonProducer
CMSSW_2_1_0_pre10 RecoMuon/L2MuonSeedGenerator
CMSSW_2_1_0_pre10 RecoMuon/L3MuonProducer
CMSSW_2_1_0_pre10 RecoMuon/L3TrackFinder
CMSSW_2_1_0_pre10 RecoMuon/MeasurementDet
V01-04-11_20X      RecoMuon/MuonIdentification
V02-03-00          RecoMuon/MuonSeedGenerator
CMSSW_2_1_0_pre10 RecoMuon/Navigation
CMSSW_2_1_0_pre10 RecoMuon/Records
CMSSW_2_1_0_pre10 RecoMuon/StandAloneMuonProducer
CMSSW_2_1_0_pre10 RecoMuon/StandAloneTrackFinder
CMSSW_2_1_0_pre10 RecoMuon/TrackerSeedGenerator
CMSSW_2_1_0_pre10 RecoMuon/TrackingTools
CMSSW_2_1_0_pre10 RecoMuon/TransientTrackingRecHit
V03-01-10_20X      TrackingTools/TrackAssociator
CMSSW_2_1_0_pre10 TrackingTools/TrackRefitter
CMSSW_2_1_0_pre11 Validation/Configuration
V00-02-09-03       Validation/RecoMuon
V02-10-03-02       Validation/RecoTrack
V01-03-05-27       Configuration/StandardSequences
```

In addition, for the studies described in Section 7, we use a RelVal $\mathrm{t\bar{t}}$ sample reconstructed in CMSSW 2.1.9.

[1)] This has to do with the desire to use the selection function in FWLite. This then forces the code to live in `DataFormats/MuonReco`. One of the recommended track quality requirements is on the impact parameter with respect to the beamspot. The need to access the beamspot information introduces unwanted dependencies in the `DataFormats` package.

3 Definitions

Here we collect some definitions that we will use throughout this note.

- "Fake Muon": any reco::Muon passing whatever cuts we are applying that is reconstructed in single pion or single kaon events. Thus, a muon from a K $\rightarrow \mu$ decay in flight is a Fake Muon.
- "Fake Rate": the ratio of Fake Muons to generated kaons and pions as a function of the p_T or η of the generated kaon or pion.
- "Decays in Flight" vs "Punch-Through": for generated kaon and pion events, if there is a muon in the SIM collection we categorize the event as decay in flight, otherwise as punch-through. Note that we are dealing with single particle events, so this works reasonably well.

Unless specified otherwise, all results are for single particles of $p_T < 200$ GeV and $|\eta| < 2.4$ generated flat in both p_T and η.

4 The need for muon ID

To illustrate the need for muon identification, we show in Figure 1 the probability as a function of p_T that a generated kaon/pion of $|\eta| < 2.4$ results in a reco::Muon. We show this probability separately for Global Muons and Tracker Muons.

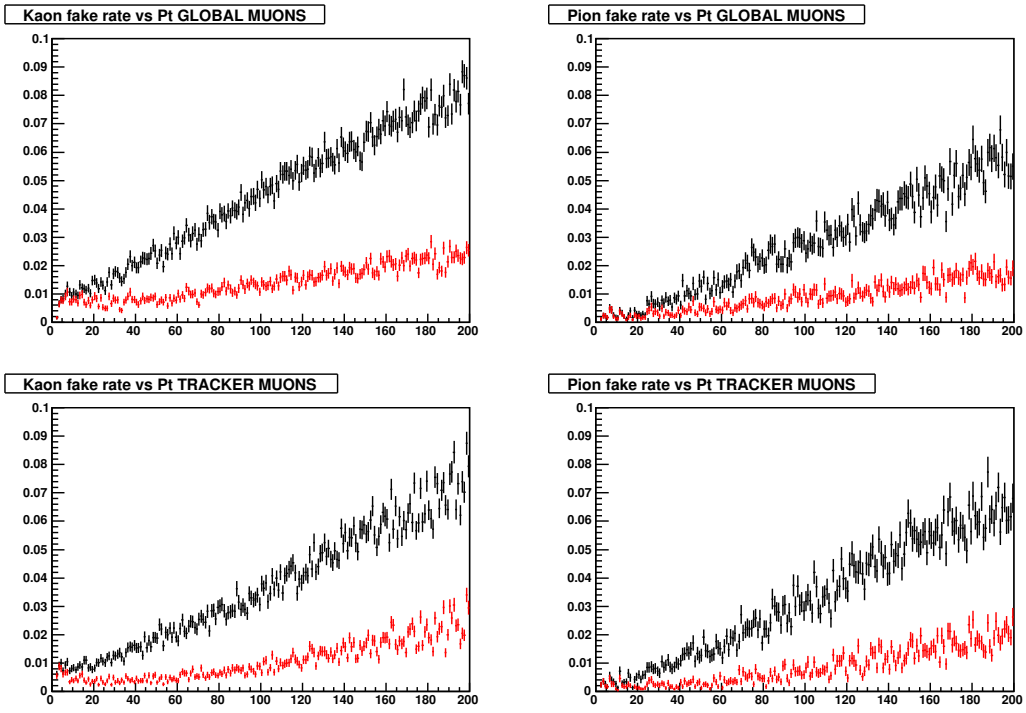

Figure 1: Out-of-the-box Global Muon and Tracker Muon fake rates as a function of p_T. Black histograms: total. Red histograms: decays in flight.

Further, in Figure 2 we show the rate of high-p_T reconstructed Global Muons in fully simulated Pythia QCD Monte Carlo (CMSSW 1.5). We also show fraction of muons muons matched to a GEN muon, *i.e.,* mostly b $\rightarrow \mu$ and c $\rightarrow \mu$[2)]. This fraction is 62% for Global Muons of $p_T > 20$ GeV, and 48% for $p_T > 40$ GeV.

[2)] The matching here is a simple ΔR match.

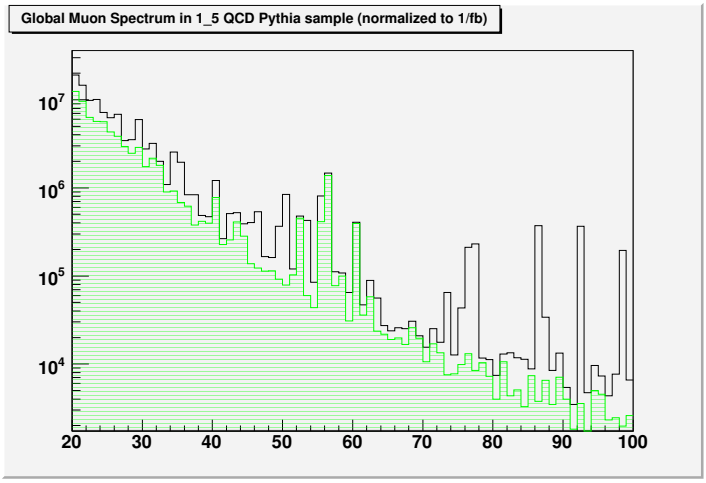

Figure 2: Rates of Global Muons in Pythia QCD Monte Carlo, simulated and reconstructed with CMSSW 1.5. The green histogram shows the subset that are matched to a GEN muon.

5 Global Muon Requirements

5.1 Global Muon χ^2

It is well known that the global muon normalized-χ^2 is a powerful tool to reject both decays in flight and punch-through, see Figure 3. One has to be a bit careful using this variable, since even in Monte Carlo there are long tails in the signal muon distribution.

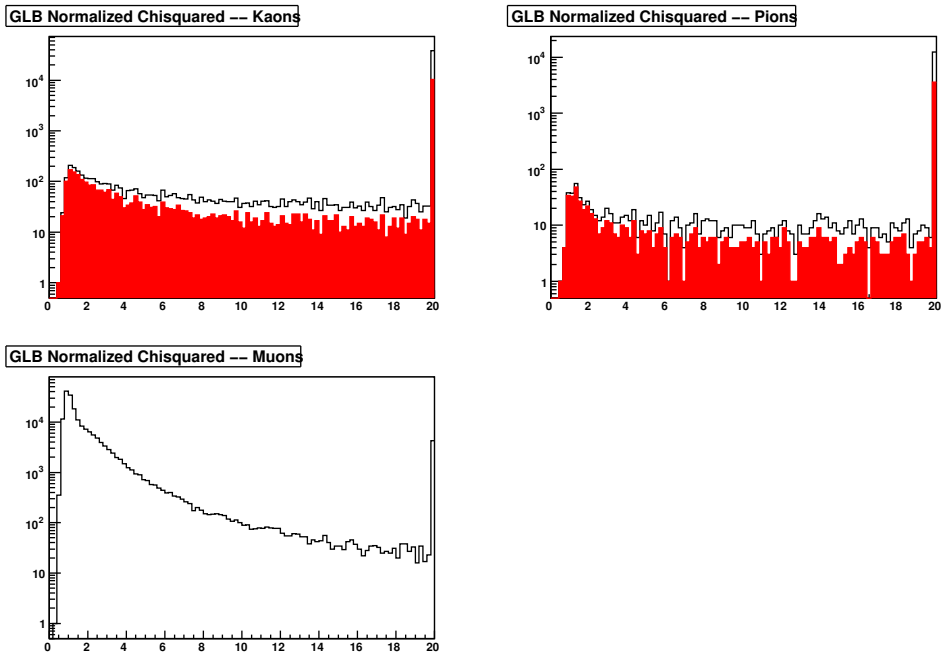

Figure 3: Normalized-χ^2 for Global Muons reconstructed in single kaon, pion, and muon events. The last bin is the overflow bin. For kaons and pions, the open histogram is for all reconstructed Global Muons, the red histogram is for the subset labeled as decays in flight.

The efficiency for a cut on normalized-$\chi^2 < 10$ or < 5 are displayed in Figure 4.
In what follows we apply a normalized-$\chi^2 < 10$ requirement.

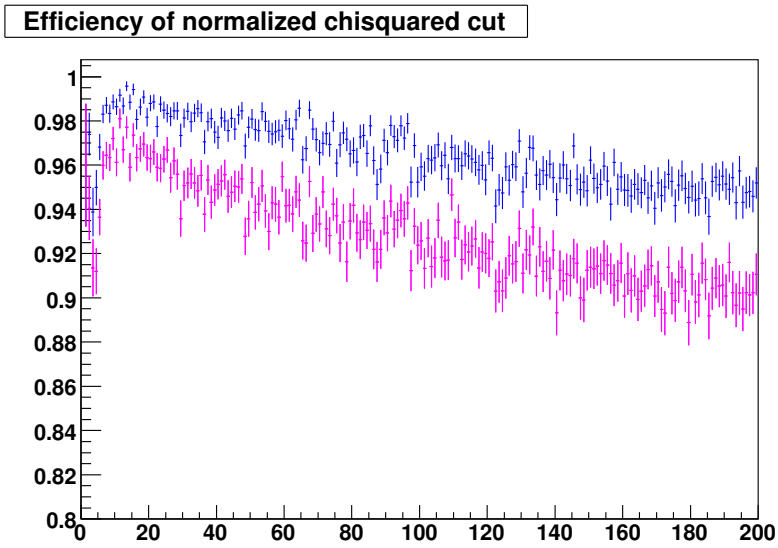

Figure 4: Efficiency for muons of requirement on normalized-$\chi^2 < 5$ (purple) and normalized-$\chi^2 < 10$ (blue).

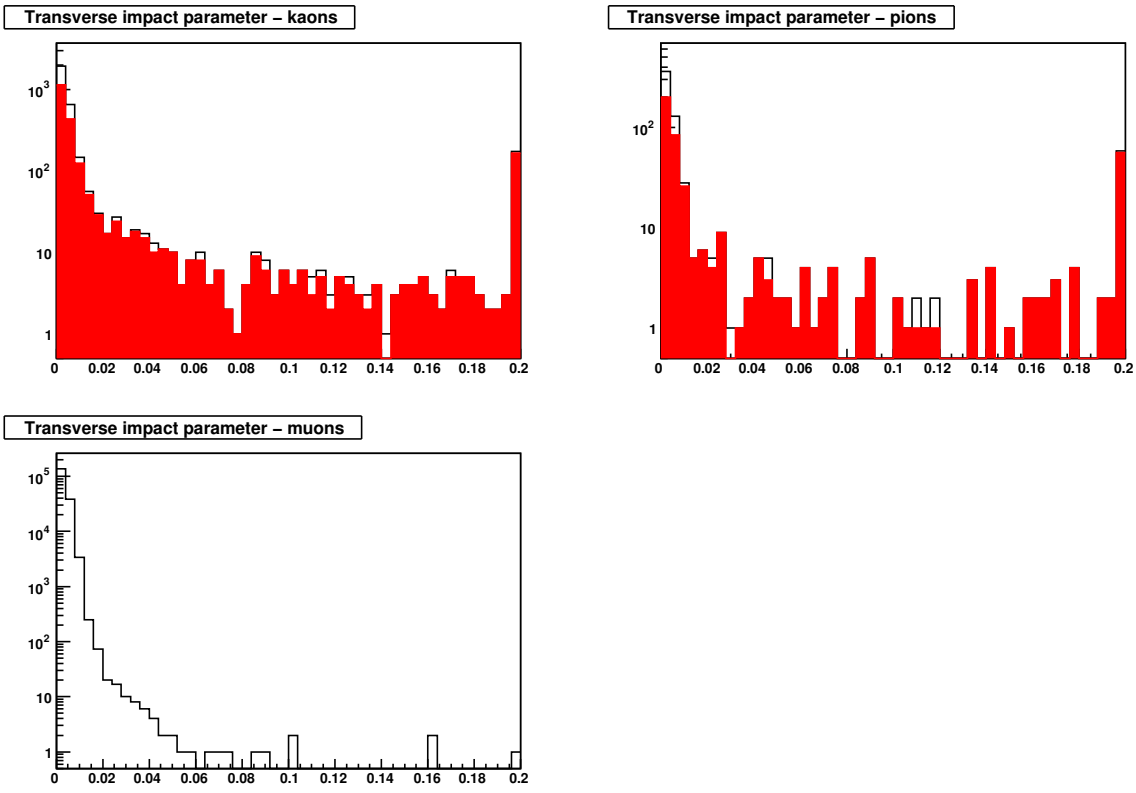

Figure 5: Impact parameter in the XY plane (d_0) for the silicon fit for global muons with global normalized-$\chi^2 < 10$. The last bin is the overflow bin. For kaons and pions, the open histogram is for all reconstructed global muons, the red histogram is for the subset labelled as decays in flight.

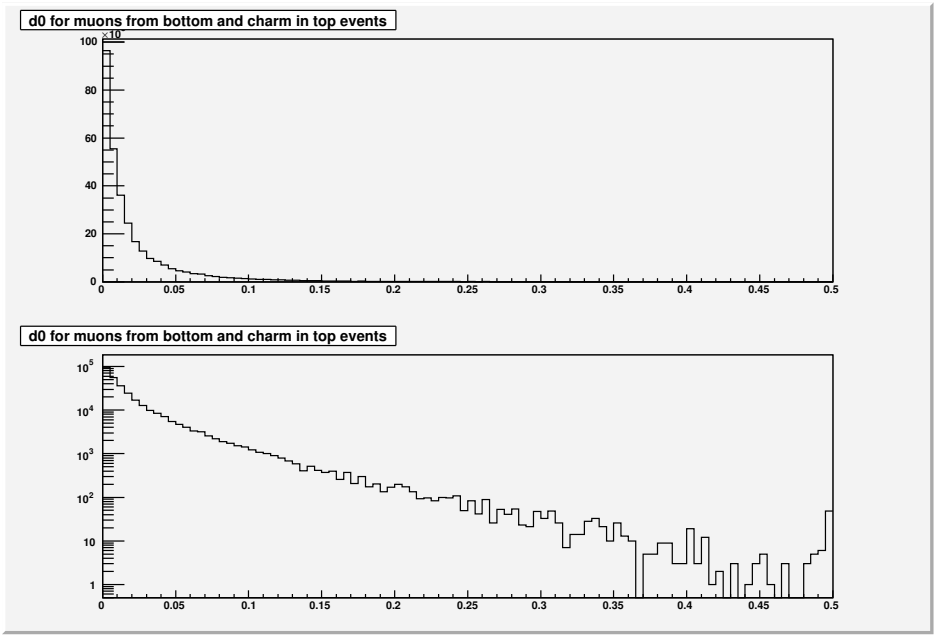

Figure 6: Impact parameter of global muons in MC top events truth matched to GEN muons from bottom or charm decays. This is for a CSA07 Monte Carlo sample, *i.e.*, CMSSW 1.6, reconstructed with 100 pb^{-1} conditions.

5.2 Track quality cuts

Track quality cuts can be used to reject leftover decays in flight. Without doing anything very complicated, *e.g.*, trying to reconstruct kinks in the tracks, there are a few quantities that one should look at:

- Impact parameter of the silicon or global fit.
- Normalized-χ^2 of the silicon fit.
- Number of hits of the silicon fit.

In Figure 5 we show distributions of the XY impact parameter (d_0) for the silicon fit, after having applied a global normalized-χ^2 cut. A loose d_0 cut is very efficient for prompt muons and rejects a significant fraction of decays-in-flight. Note that d_0 here must be calculated from the beam spot or the primary vertex, not from the nominal beamline.

If the cut on d_0 is too tight, muons from b and c-quark decays could be rejected. This is demonstrated in Figure 6. Thus we propose a baseline cut $|d_0| < 2$ mm, which preserves the efficiency for muons from bottom and charm. If so desired, a tighter cut could be used to eliminate these muons, see for example the discussion in [3].

Next, we consider applying a requirement on the normalized-χ^2 of the silicon fit, see Figure 7. This variable does have some discriminating power against decays-in-flight, however the separation between signal and background is not large. Therefore, at this time we choose to not make any requirements on this quantity.

Finally, we examine the distribution of valid hits for the silicon fit, see Figure 8. There is clearly some discriminating power in this variable and we choose a loose requirement $N_{hits} \geq 11$. Note that this cut could in principle be made as a function of η, see Figure 9.

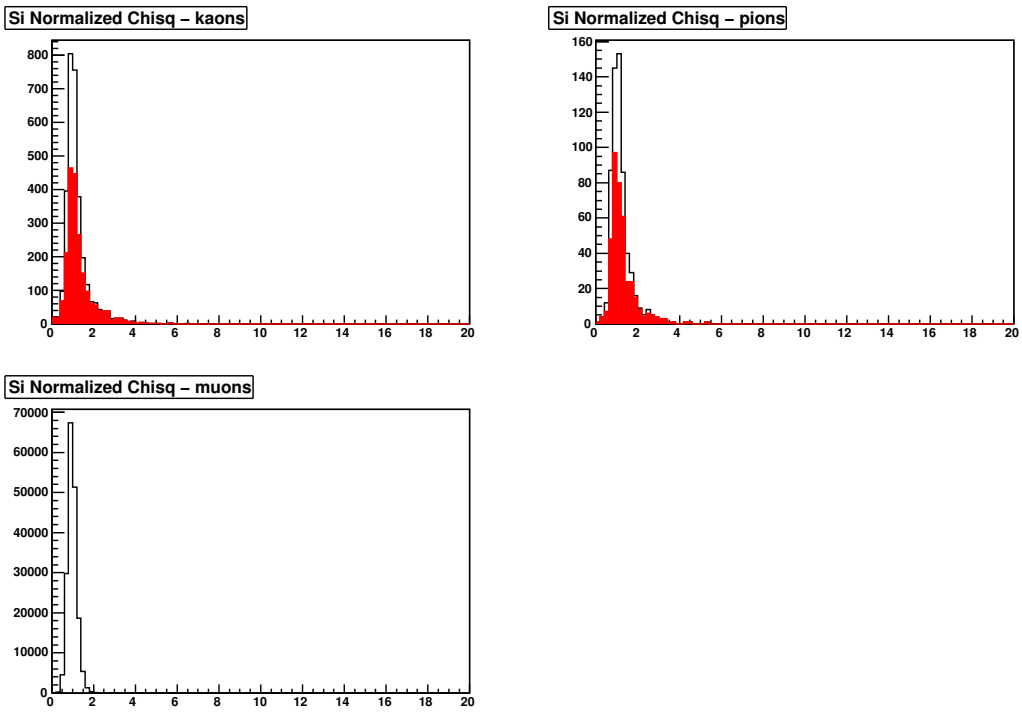

Figure 7: Normalized-χ^2 of the silicon fit for global muons with global normalized-$\chi^2 < 10$ and $|d_0| < 2$ mm. The last bin is the overflow bin. For kaons and pions, the open histogram is for all reconstructed global muons, the red histogram is for the subset labelled as decays in flight.

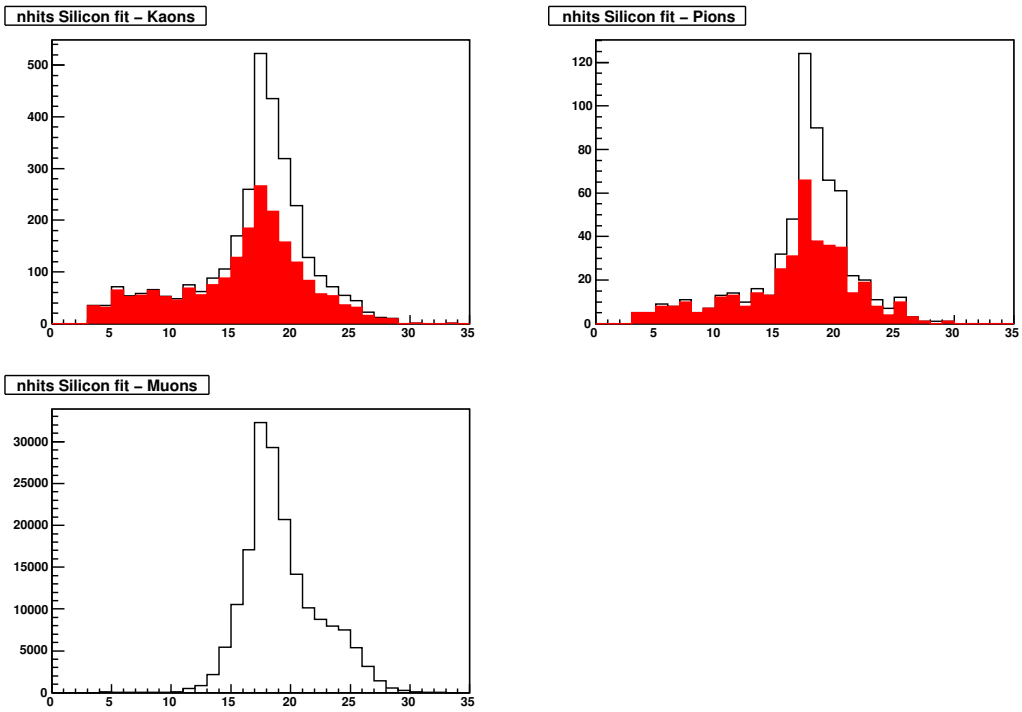

Figure 8: Number of valid hits in the silicon fit for global muons with global normalized-$\chi^2 < 10$ and $|d_0| < 2$ mm. For kaons and pions, the open histogram is for all reconstructed global muons, the red histogram is for the subset labelled as decays in flight.

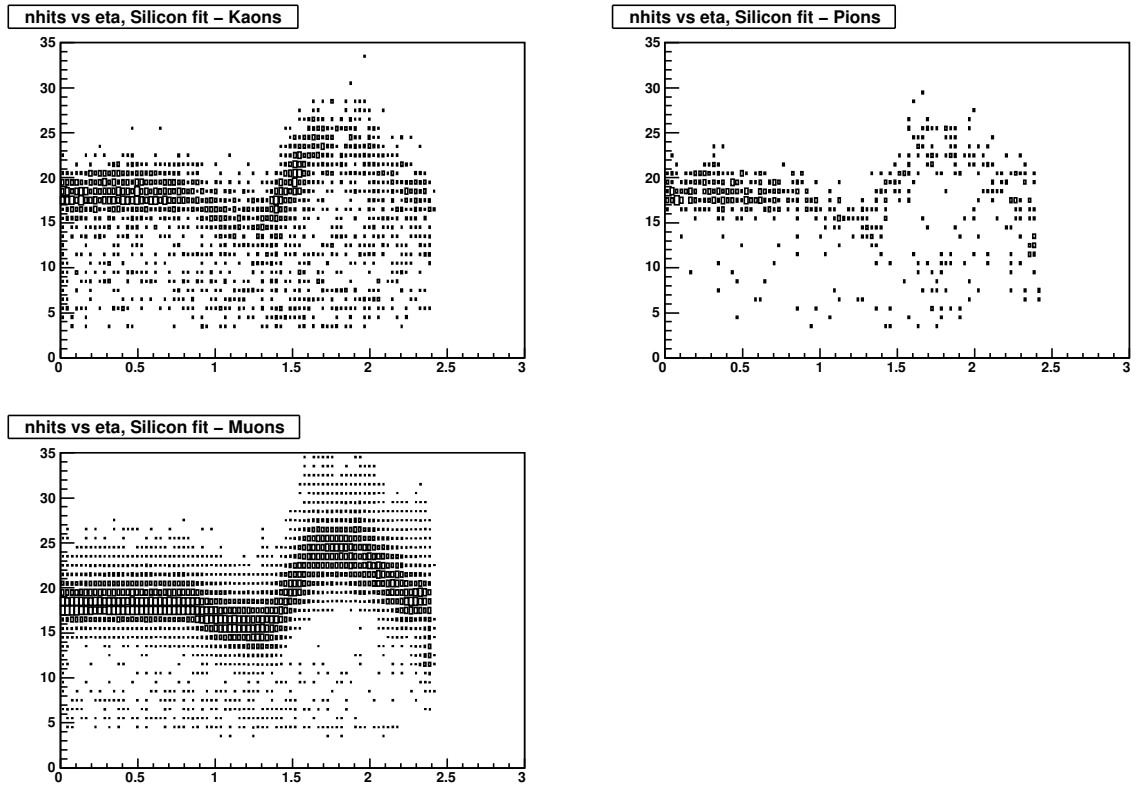

Figure 9: Number of valid hits vs. $|\eta|$ in the silicon fit for global muons with global normalized-$\chi^2 < 10$ and $|d_0| < 2$ mm.

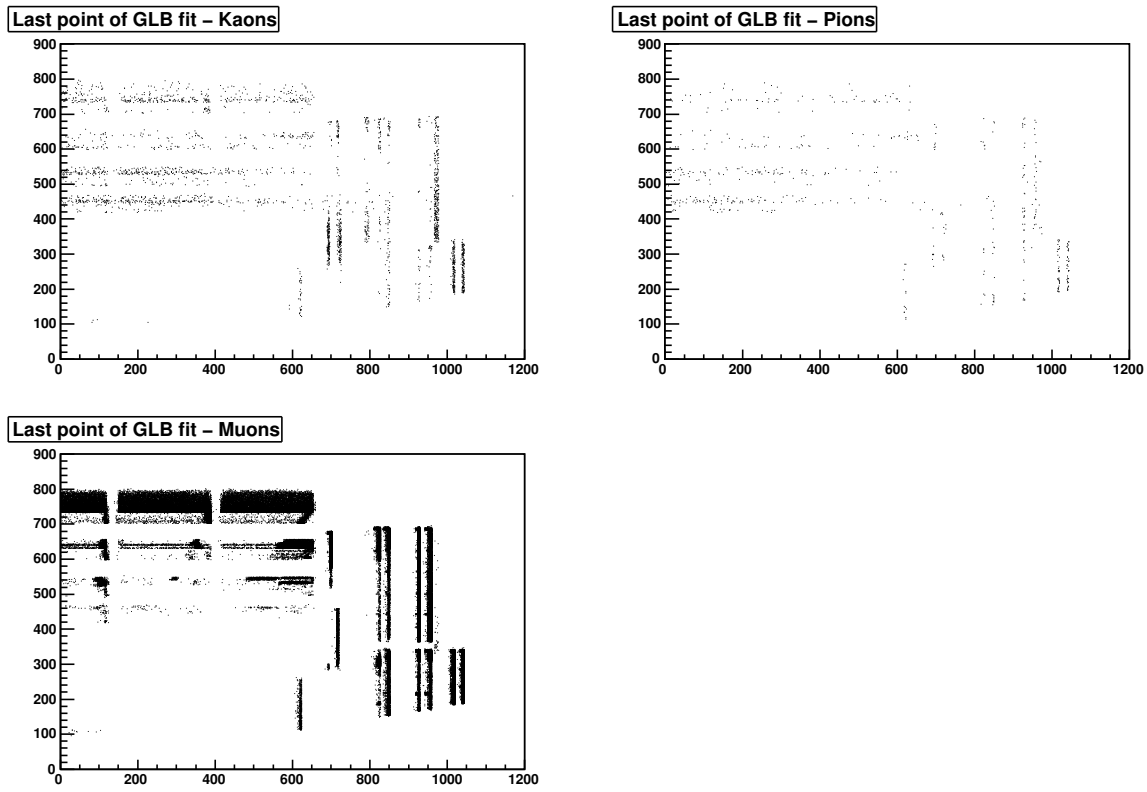

Figure 10: RZ map of the endpoint of the global fit for global muons with global normalized-$\chi^2 < 10$, $|d_0| < 2$ mm, and $N_{hits} \geq 11$.

5.3 Penetration in the muon detector

High p_T muons are expected to penetrate deep into the muon detector. This is illustrated in Figure 10. Very few of the global fits for real muons end in the first layer of the barrel detector, while this happens quite often for kaons and pions, even after all of the requirements. Rejecting these muons results in a 20% reduction in fake rate at very little cost in efficiency. The proposed requirement then is

```
if (  abs(Z) < 660 && r > 400 && r < 480 ) keepMuon=false;
```

where r and Z are the radial and Z coordinates of the end of the global fit; note that this requirement would introduce an inefficiency at low p_T.

One should be able to apply a similar requirement for the first layer of the endcap:

```
if (  abs(Z) > 600 && abs(Z) < 650  && r < 300) keepMuon=false;
if (  abs(Z) > 680 && abs(Z) < 730  && r < 480) keepMuon=false;
```

However, in the data set used for this study many global fits for real muons also end in the first layer of the endcap muon system. This problem has since been resolved, and the solution will be included in CMSSW 2.2.0.

5.4 Summary of requirements for global muons

The global muon requirements are summarized in Table 1. We define a selection, `GlobalMuonPromptTight`, which only includes the requirement in the first line of the Table, global-normalized-$\chi^2 < 10$. The track quality requirements are not included for the reasons discussed in Section 1. The last two requirements are not implemented. This is because the penetration requirement in the barrel would introduce an inefficiency at low p_T; the penetration requirement in the endcap cannot be applied until fixes to the global muon reconstruction are included (already done for CMSSW 2.2.0). It is very easy for the user to implement the penetration requirements externally, if so desired, see Section 5.3.

Requirement	Comment
Global normalized-$\chi^2 < 10$	Watch out for long tails in real data
$\lvert d_0 \rvert < 2$ mm	From beamspot; could be tighter; careful about b-quarks
$N_{hits} \geq 11$	Could be made as function of η
Reject fits ending in 1st barrel station	Watch out for low p_T
Reject fits ending in 1st endcap station	Must wait for fix in CMSSW 2.2.0

Table 1: Summary of baseline global muon requirements. Only the first requirement is implemented in `GlobalMuonPromptTight`.

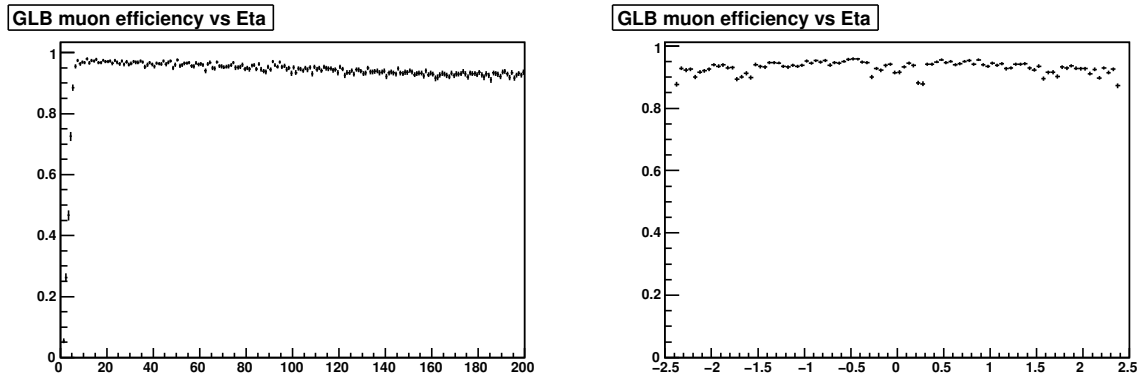

Figure 11: Muon efficiency as a function of p_T and η for `GlobalMuonPromptTight` including the track quality requirements.

5.5 Global Muon fake rates and efficiencies

The fake rates and efficiencies using the `GlobalMuonPromptTight` requirement plus the track quality requirements (first three lines of Table 1) are displayed in Figures 11 and 12. For completeness, we also show the same plots, but adding the requirement on the fourth line of Table 1 to the `GlobalMuonPromptTight` selection, Figures 13 and 14. Note that the efficiency is essentially unchanged, but the fake rate in the barrel is reduced somewhat.

It is also interesting to again separate the background components into decays-in-flight and punch-through. This is shown in Figure 15.

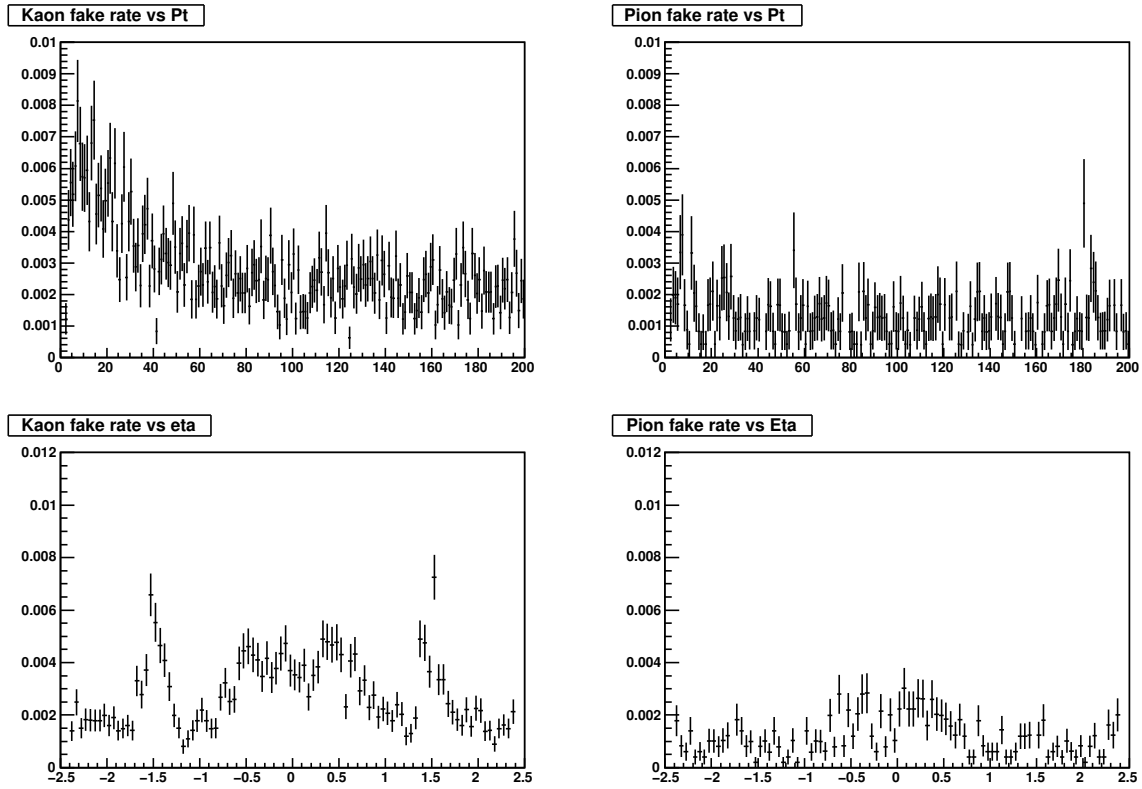

Figure 12: Fake rates for pions and kaons as a function of p_T and η for for `GlobalMuonPromptTight` including the track quality requirements.

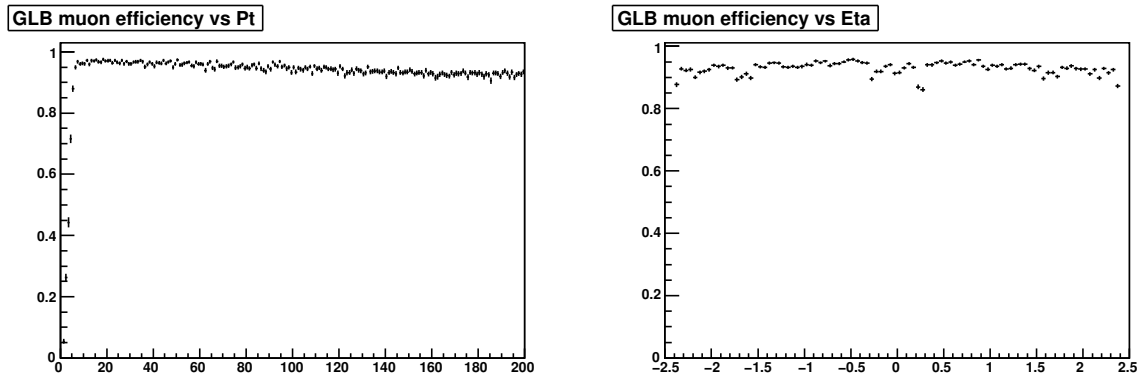

Figure 13: Muon efficiency as a function of p_T and η for the `GlobalMuonPromptTight` selection including the track quality requirements plus the requirement on the 4th line of Table 1.

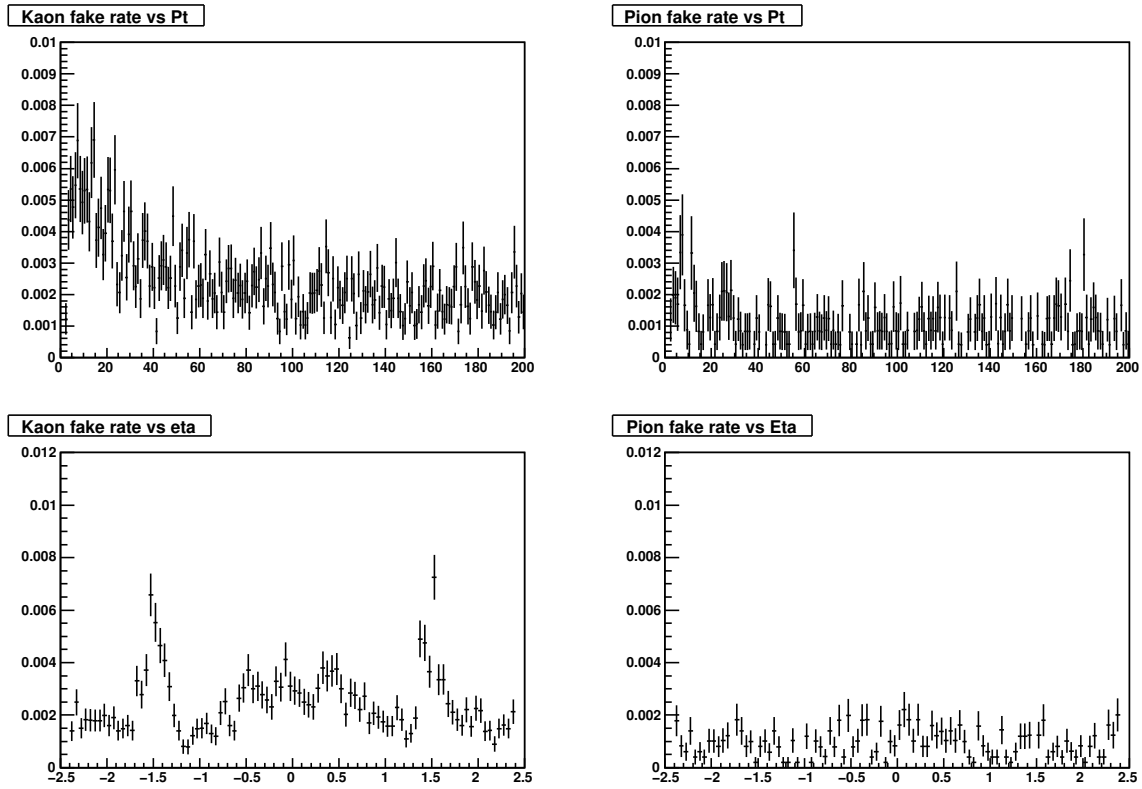

Figure 14: Fake rates for pions and kaons as a function of p_T and η for the `GlobalMuonPromptTight` selection, including the track selection requirements, plus the requirement on the 4th line of Table 1.

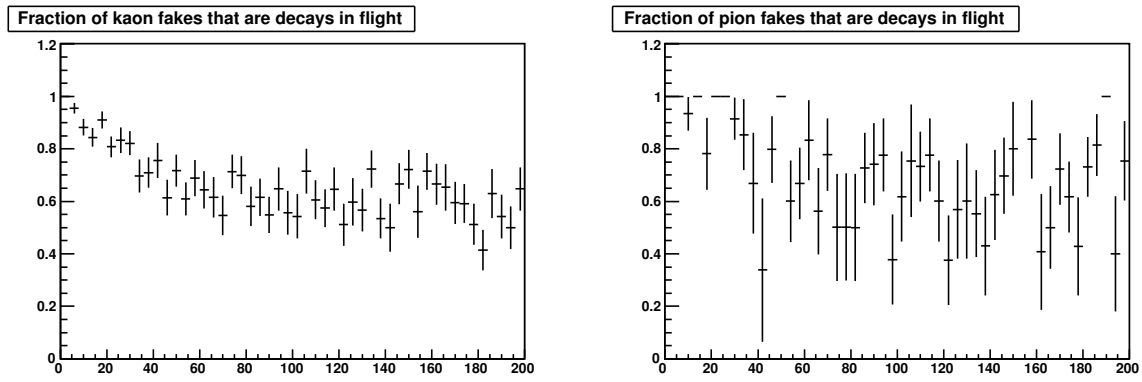

Figure 15: Fractions of pion and kaon fakes that are decays-in-flight as a function of p_T for for `GlobalMuonPromptTight` selection, including the track quality requirements, plus the requirement on the 4th line of Table 1.

6 Tracker Muon requirements

We have developed two approaches to Tracker Muon identification. One approach is cut-based and the other approach constructs a continous variable (aka "compatibility") that can be used to identify muons. In both cases the first step in the algorithm is to perform "arbitration". The concept of arbitration is described in Reference [1], and is briefly discussed below.

6.1 Arbitration

Tracker Muons are built by associating segments in the muon chambers with the silicon track. At reconstruction time a given segment can be associated to more than one silicon track, *i.e.*, reconstructed Tracker Muons are allowed to share segments. Arbitration is the pattern recognition process which assigns each segment uniquely to a single Tracker Muon.

The default arbitration algorithm works as follows: suppose that one segment is associated with more than one Tracker Muon. For each Tracker Muon the quantity $\Delta R^2 = \Delta X^2 + \Delta Y^2$ is calculated, where ΔX and ΔY are the distances between the extrapolated silicon track and the segment in local X and Y coordinates. The segment is uniquely associated to the Tracker Muon for which ΔR is smallest.

6.2 Muon Compatibility Algorithm

The idea of this selection is to construct a continous variable to quantify the compatibility of a Tracker Muon object with the muon hypothesis. In practice, two such variables are constructed, one based on calorimeter information and one based on muon chamber information. The two variables are then combined into a single variable.

6.2.1 Calorimeter Compatibility

The minimum-ionizing signature of muons in the calorimeter can be used to help identify those reconstructed tracks originating from muons. The calorimeter compatibility variable is based on the measured energies in towers along the extrapolated track direction as stored in the muon object [1]. Separate energy measurements are available for the electromagnetic and hadronic sections of the calorimeter as well as for the outer hadronic calorimeter (HO) which is an additional scintillator layer located just outside of the superconducting solenoid in front of the inner muon chambers. Readout thresholds play an important role in how often we can observe minimum-ionizing calorimeter signals. Note that in the endcap region, there are numerous cases where the algorithm returns an electromagnetic energy of identically zero from towers traversed by the extrapolated track.

Based on the measured energies in each calorimeter section, we determine a compatibility value between zero and one that attempts to describe how consistent these energies are with respect to what one expects for a muon. The compatibility value is obtained from a three-dimensional likelihood function of the form

$$\frac{P_S(x) \cdot P_S(y) \cdot P_S(z)}{P_S(x) \cdot P_S(y) \cdot P_S(z) + P_B(x) \cdot P_B(y) \cdot P_B(z)} \, , \tag{1}$$

where P_S and P_B are the signal and background probabilities as functions of the measured energies in the electromagnetic calorimeter (x), the hadronic calorimeter (y), and the HO scintillator layer (z). The signal and background probability distributions are obtained from simulated samples of single-muons and pions, respectively. The distributions are constructed as two-dimensional templates giving the relative probabilities for different values of the measured energies as a function of the input track momentum. We also account for potential differences in calorimeter response as a function of pseudo-rapidity by using a total of nine templates (three for EM response, five for hadronic response, and one for the HO response). The three EM calorimeter templates cover the η ranges from -2.5 to -1.479, -1.479 to 1.479, and 1.479 to 2.5; the five hadronic calorimeter templates cover eta ranges from -2.5 to -1.27, -1.27 to -1.1, -1.1 to 1.1, 1.1 to 1.27, and 1.27 to 2.5; and the HO template covers the full psuedorapidity range of the system ($|\eta| < 1.28$).

In terms of track momentum, the binning of the templates is done in 1 GeV steps from 1 to 5 GeV, 2 GeV steps from 5 to 15 GeV, 3 GeV steps from 15 to 30 GeV, 5 GeV steps from 30 to 60 GeV, 10 GeV steps from 60 to 100 GeV, 20 GeV steps from 100 to 200 GeV, 50 GeV steps from 200 to 500 GeV, 100 GeV steps from 500 to 1000 GeV, and 250 GeV steps from 1000 to 2000 GeV. In cases where the input track momentum is below 1 GeV, we assume that the track will be unable to propogate through the calorimeter and simply assign a calorimeter compatibility value of 0.5. In cases where a measured track momentum is more than 2000 GeV, we detemine a compatibility value based on the modeled probability distributions in the highest track momentum bin.

The binning in templates in terms of the deposited energy is handled differently for each of the three independent calorimeter measurements. For the measured energy in the EM calorimeter, if a returned value is exactly zero, we manually convert the value to -5 GeV. This change is made so that we can separately treat cases for which no calorimeter information is available due to the applied readout thresholds. After applying this change, we bin events into the templates as a function of the measured EM energy based on the following bin structure: -10.0 to -1.0 GeV, -1.0 to 0.0 GeV, 0.0 to 0.165 GeV, 0.165 to 0.21 GeV, 0.21 to 0.25 GeV, in bins of 0.025 GeV from 0.25 to 0.50 GeV, in bins of 0.100 GeV from 0.50 to 1.00 GeV, in bins of 0.200 GeV from 1.00 to 1.40 GeV, and 1.4 to 500 GeV. The lowest bin is primarily for the entries which are exactly zero indicating that the corresponding tower(s) has not been readout. The next-to-lowest bin whose range also lies below zero corresponds to actual readout values that fluctuate below zero based on the applied pedestal. In cases where the returned value of the associated EM energy for a track is either below -10.0 GeV or above 500 GeV, $P_S(x)$ and $P_B(x)$ are simply set to 1.0.

In the case of the returned value for the HAD energy associated with a track, we apply prior to binning an η-dependent correction designed to flatten the muon energy response. In essence this corrects for the $\sin\theta$ dependence on the path length for a muon passing through a specific calorimeter tower. For the region between -1.1 and 1.1 in η, the correction is simply e' = e $\sin\theta$. A slightly more complicated function is used to flatten response in the region of overlap between the central and plug calorimeter modules ($1.1 < |\eta| < 1.27$). In this region the correction factor used is $e' = e\ 1.8/(-2.2\ \eta + 5.5)$. In the forward module ($1.27 < |\eta| < 2.50$), no path length correction is required, but we do scale the observed mip energy to match that observed in the barrel ($e' = 0.818\ e$). After making these corrections, the measured energy from the calorimeter is binned into the templates based on the following bin structure: -5.0 to 0.0 GeV, in bins of 0.15 GeV from 0.0 to 3.9 GeV, 3.9 to 4.1 GeV, in bins of 0.30 from 4.1 to 5.0 GeV, in bins of 0.50 GeV from 5.0 to 6.0 GeV, 6.0 to 7.0 GeV, 7.0 to 8.5 GeV, 8.5 to 10.0 GeV, 10.0 to 1000.0 GeV. Similar to the case of the EM deposits, measured energy deposition outside the range of our binning result in values of 1.0 being assigned to both the $P_S(y)$ and $P_B(y)$ inputs to the likelihood.

No corrections are made to the returned HO energy. HO energies are binned into the corresponding template as follows: -5.0 to -0.6 GeV, in bins of 0.2 GeV from -0.6 to 1.0 GeV, in bins of 0.5 GeV from 1.0 to 2.0 GeV, and 2.0 to 1000 GeV. Once again energies that lie outside the range of our bins result in values of 1.0 being assigned to both $P_S(z)$ and $P_B(z)$.

The signal and background templates are based on one million event simulated single muon and pion samples, respectively. The single particle event samples are generated flat in momentum over the following momentum ranges: 1 to 10 GeV, 10 to 60 GeV, 60 to 300 GeV, and 300 to 2000 GeV. In each momentum range a total of 250,000 events are generated with equal contributions of positively and negatively charged particles. Each generated event contributes entries into up to three corresponding templates based on the measured track P_T and the observed EM, HCAL,and HO energies using the energy corrections and binning structures described above. Once the templates have been filled, entries in each of the 1-dimensional momentum bins are normalized such that the entries within each bin sum to one. In the rare cases where a particular 2-D bin is found to contain no entries, we assign a weight corresponding to what one would expect for a single entry. Figure 16 contains normalized distributions of the calorimeter-based compatibility value for muon candidates reconstructed from simulated samples of single muons and kaons. The secondary peak in the kaon distribution near a compatibility value of one originates from real muons produced in in-flight decays.

6.2.2 Muon Segment Compatibility

We use information on which crossed stations have matched segments (in addition to the quality of the matches between the extrapolated track and the associated segments) to construct a second muon compatibility variable.

In cases where an extrapolated track passes through a muon station but no associated segment is found, the muon object stores the distance between the extrapolated track position and the closest chamber edge. The value is provided in cm and is negative if the actual extrapolated track position is within a physical chamber boundry. In addition, the algorithm calculates the expected uncertainty in the extrapolated position at each station due to cumulative multiple scattering effects. The value of the distance to the closest chamber edge divided by this multiple scattering uncertainty is used as a measure of the probability for the trajectory of a given track to deflect either from inside to outside of a chamber boundry or equivalently from outside to inside. When an extrapolated track does not pass within a reasonable distance of a given muon station, the algorithm simply returns a null pointer. A null pointer is also returned for cases where a low momentum muon is expected to range out in the detector material before reaching a specific muon station.

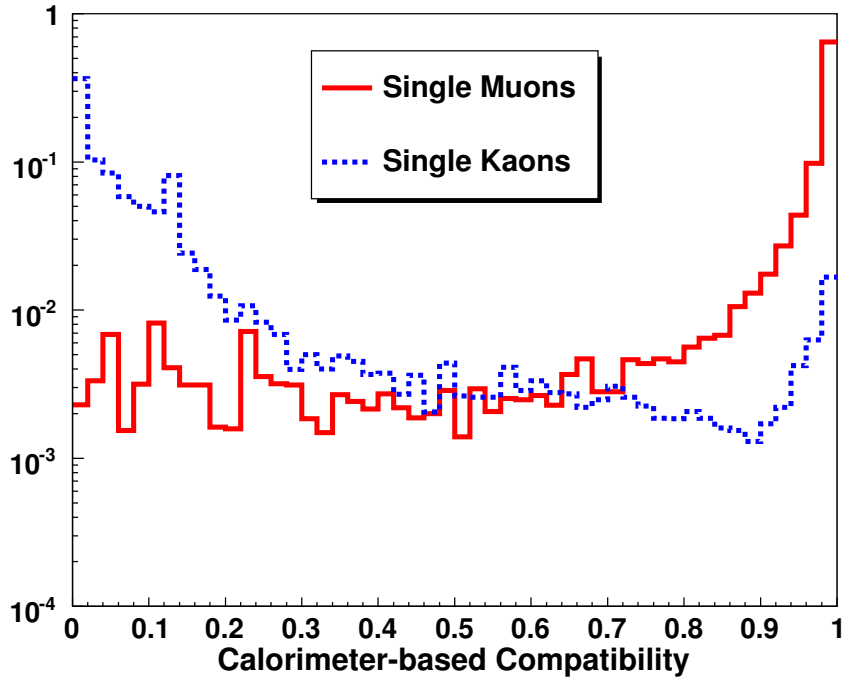

Figure 16: Normalized calorimeter-based compatibility distributions for reconstructed muons in simulated single muon (red) and single kaon (blue) samples. The samples were generated flat in p_T from 0 to 500 GeV/c (muons) and 2 to 200 GeV/c (kaons) and flat in pseudo-rapidity from -2.4 to 2.4.

In order to determine a compatibility variable, we first assign a weight to each of the muon stations that the extrapolated track is expected to either cross or pass within a short distance. If there are no such stations, the compatibility value is simply set to 0.5. If there is one or more such stations, then the compatibility value is at the most basic level determined from the fraction of these stations containing a reconstructed segment that can be associated with the extrapolated track. Because hadronic punchthrough is more likely to produce segments in the inner stations, segments in the outer layers are considered to be stronger evidence for the muon hypothesis. Therefore, the associated segments found in the outer stations are given a higher weight in the determination of the compatibility value. In the case of only one expected station, the station in question makes a full (100%) contribution to the compatibility. In the case of two expected stations, the inner and outer stations contribute with weights of 33% and 67% respectively. For three expected stations the weighted contributions of each (from inner to outer) are 22%, 33%, and 44%. Similarly, in the case of four expected stations, the most possible, the weighted contributions using the same ordering are 10%, 20%, 30%, and 40%. At first order, the compatibility is detemined by multiplying the pre- determined weight for each station by zero or one (one if an associated segment is found) and then summing the resulting values.

We improve upon this first order calculation via two additional corrections applied on a station by station basis. The first correction is used to reduce the contribution to the compatibility value coming from layers in which the match between the associated segment and extrapolated track positions lies outside the one sigma expectation. In particular, we calculate a value

$$\Delta P^2 = \Delta X^2/\sigma_X^2 + \Delta Y^2/\sigma_Y^2$$

where ΔX and ΔY are the differences in position between the extrapolated track and associated segment in the local X and Y coordinates and σ_X and σ_Y are the estimated uncertainties in those quantities. If the resulting value of ΔP is less than one, then no penalty is applied. In cases where ΔP is greater than one, the contribution from the station in question is reduced by a factor of $1/\Delta P^{0.25}$. The factor of 0.25 in the re-weighting factor has no physical significance and was chosen based on studies designed to optimize the separation power of the compatibility value. Due to the fact that uncertainties on the extraploated track position are found to be somewhat underestimated for high-P_T tracks, we utilize the actual difference between the positions of the associated segment and extrapolated track $\Delta R^2 = \Delta X^2 + \Delta Y^2$ rather than ΔP to penalize a small subset of high p_T tracks for which $\Delta R < 3$

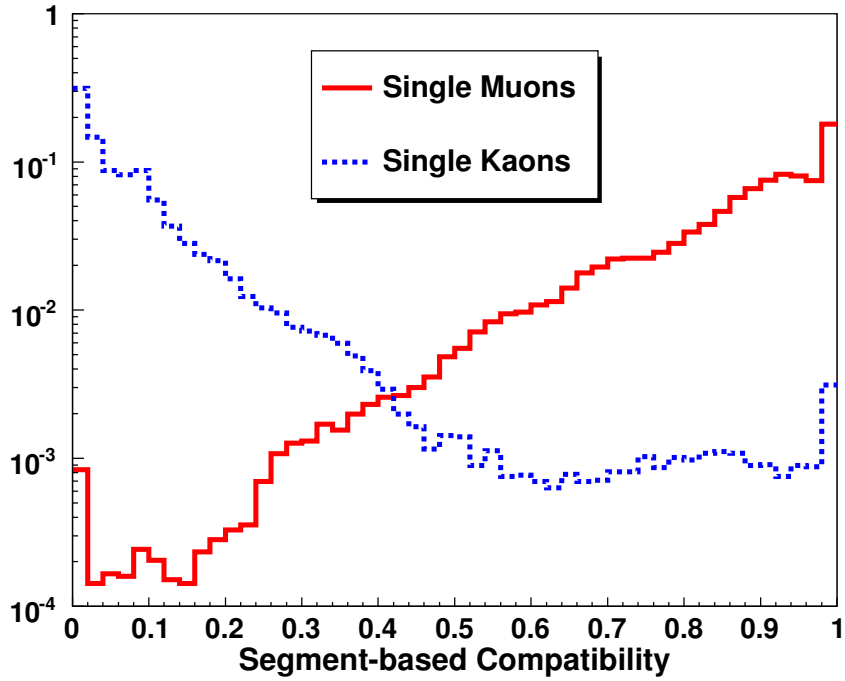

Figure 17: Normalized segment-based compatibility distributions for reconstructed muons in simulated single muon (red) and single kaon (blue) samples. The samples were generated flat in p_T from 0 to 500 GeV/c (muons) and 2 to 200 GeV/c (kaons) and flat in pseudo-rapidity from -2.4 to 2.4.

cm and $\Delta P > 3$. For these special cases, the contribution from the station in question is reduced by a factor of $1/\Delta P^{0.25}$. Using different re-weighting factors in these special cases helps to maintain a consistent seperability between muons and pions over the entire p_T range.

A second set of station by station corrections is applied for stations that are crossed (or nearly crossed) by the extrapolated track but for which no associated segment is found. In these cases, the net contribution to the compatibility value from a station with no associated segment can be increased above zero if there is a non-negligible probability for the track to have passed outside the active chamber boundaries. These corrections are based solely on the returned distance between the extrapolated track position and the closest chamber edge. In cases where the extrapolated track position is more than 10 cm inside of the chamber boundaries no correction is applied (the station contributes nothing to the compatibility value). Conversely, if the extrapolated track position lies less than 10 cm from or outside of the chamber boundaries, then a non-zero contribution from the station is allowed. The actual value of this contribution is given by the following parameterization:

$$c = 0.5 \times 0.5 \times [\mathrm{ERF}(d/6) + 1]$$

where d is the distance between the extrapolated track position and the nearest chamber boundary. Note that d is negative for positions inside the chamber boundaries and positive otherwise. The correction factor for an extrapolated track that points directly at a chamber boundary (d = 0) is 0.25. Note that the nominal expectation is that 50% of such tracks would in fact pass through the active chamber region. The assignment of a smaller 25% contribution to the compatibility value from such a station is once again based on studies designed to optimize seperability. Note also that the maximum contribution coming from a station with no associated segment is 0.5 for the cases of extrapolated tracks that pass well outside the chamber boundaries (c = 0.495 for d = 10 cm). Conversely, in the cases where the extrapolated track position is well within the chamber boundaries, the resulting contribution from the station is negligible (c = 5e-03 for $d = -10$ cm). Normalized distributions of the segment-based compatibility value are shown in Fig. 17 for muon candidates reconstructed from simulated samples of single muons and kaons.

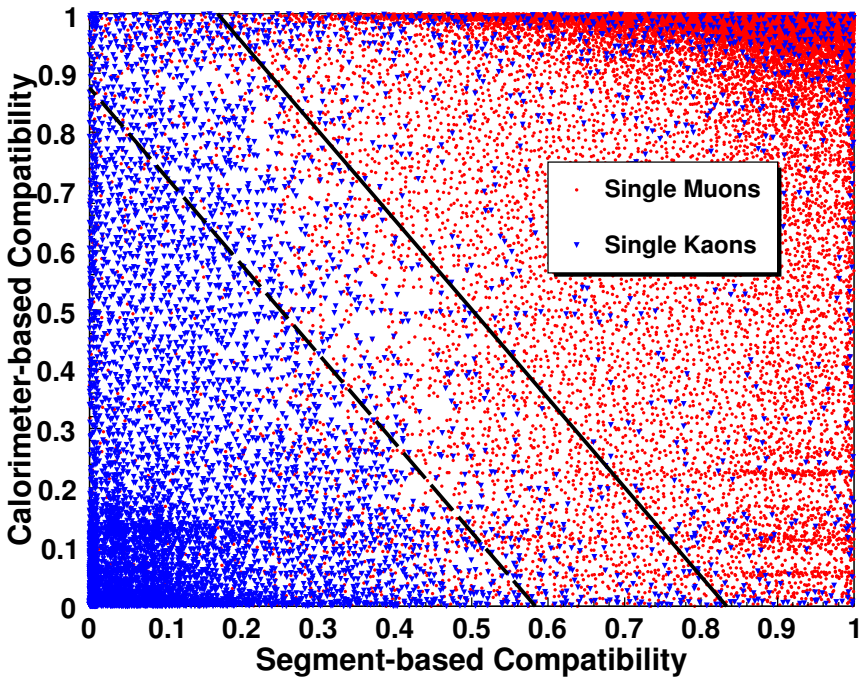

Figure 18: Calorimeter-based versus segment-based compatibility values for reconstructed muons in simulated single muon (red) and single kaon (blue) samples. The samples were generated flat in p_T from 0 to 500 GeV/c (muons) and 2 to 200 GeV/c (kaons) and flat in pseudo-rapidity from -2.4 to 2.4. The `TM2DCompatibilityTight` and `TM2DCompatibilityLoose` criteria select muon candidates to the right of the solid and dashed lines, respectively.

6.2.3 Compatibility-based Muon Selection

A simple selection of muon candidates based on the calorimeter and segment-based compatibility values described above can be made in a straight-forward manner. Each set of chosen cut values provides a specific trade-off with respect to the contribution of fake muons within a sample versus muon selection efficiency. Higher cut values result in less fakes but an associated decrease in selection efficiency. The values of both the calorimeter and segment-based compatibility can be obtained directly from using the following methods embedded within the standard CMS muon object: `reco::Muon::caloCompatibility()` and `reco::Muon::segmentCompatibility()`.

We also define two standard requirements:

- `TM2DCompatibilityTight`, for which the criteria is that 0.8 times the calorimeter compatibility + 1.2 times the segment compatibility is > 1.0.
- `TM2DCompatibilityLoose`, for which the criteria is that 0.8 times the calorimeter compatibility + 1.2 times the segment compatibility is > 0.7.

A two-dimensional distribution of the calorimeter-based versus segment-based compatibility values for muon candidates reconstructed in simulated samples of single muons and kaons is shown in Fig. 18. The clumping of muon events into horizontal line structures at the bottom right corner of the plot originates from high-momentum muons that radiate a significant fraction of their energy within the calorimeter.

6.3 Cut-based requirements

Unlike the compatibility requirements, these algorithms do not demand that the energy in the towers crossed by the muons be compatible with minimum ionizing. Of course it would be possible to add a requirement on `CaloCompatibility` if so desired. We have defined two families of requirements, `TMLastStation` and `TMOneStation`. In both cases we have defined "loose" and "tight" requirements, *e.g.*, `TMLastStationLoose`,

TMLastStationTight, etc. The TMOneStation set is supposed to complement the TMLastStation set at low p_T. We have defined a combined set TMLastStationOptimizedLowPt, which substitutes TMOneStation to TMLastStation at low p_T in the barrel.

6.3.1 TMLastStation algorithms

These algorithms are based on the fact that muons penetrate through the whole muon detector, whereas hadrons tend to stop earlier, and that the segments must be spatially matched to the extrapolated track. Starting from segments passing the arbitration procedure described in Section 6.1, these very simple requirements are imposed:

1. There have to be at least **two** segments well matched to the extrapolated track. The meaning of well matched is the following:
   - $|\Delta X| <$ Max($3\sigma_X$, 3 cm) (TMLastStationLoose and TMLastStationTight)
   - $|\Delta Y| <$ Max($3\sigma_Y$, 3 cm) (TMLastStationTight)

   where ΔX and ΔY are the distances between the extrapolated silicon track and the segment in local X and Y coordinates, and σ_X and σ_Y are the uncertainites in ΔX and ΔY. These uncertainties are the quadrature sums of the uncertainties in the extrapolated track coordinates, including the effects of multiple scattering and dE/dX, and the uncertainties in the position of the segment.
2. One of the segments passing the requirements above must belong to the last station crossed by the extrapolated silicon track. A silicon track is defined to be "crossing a station" if it extrapolates within the fiducial volume of the station and is "well inside" the station, *i.e.*, the nearest chamber edge must be at least 3σ of multiple scattering **and** at least 3 cm away from the nearest chamber edge.

6.3.2 TMOneStation algorithms

The TMLastStation algorithms require that there be two matched segments in the muon detector. This introduces an inefficiency at low p_T, as discussed in Section 6.5. The TMOneStation algorithms are meant to be used in conjuction with the TMLastStation algorithms to recover efficiency at low p_T in the central region by requiring a single matched segment. The exact requirements are

- At least one segment with $|\Delta X| <$ Max($3\sigma_X$, 3 cm) (TMOneStationLoose and TMOneStationTight)
- At least one segment with $|\Delta Y| <$ Max($3\sigma_Y$, 3 cm) (TMOneStationTight)

6.3.3 TMLastStationOptimizedLowPt algorithms

These requirements combine TMLastStation and TMOneStation to improve efficiency at low p_T. The requirements are

- Use TMOneStationLoose or TMOneStationTight for $|\eta| < 1.2$ and $p_T < 8$ GeV.
- Use TMLastStationLoose or TMLastStationTight at higher η and/or higher p_T.

See Section 6.5 for further discussion.

6.4 Performance

We now examine the performance of the Tracker Muon selection algorithms for $p_T < 200$ GeV and $|\eta| < 2.4$. In Figures 19 and 20 we show the fake rates, and in Figure 21 we show the efficiencies.

A few comments can be made:

- The TM2DCompatibility requirements are more efficient than the TMLastStation requirement, but the fake rate is higher.
- There is little difference in performance between "loose" and "tight".

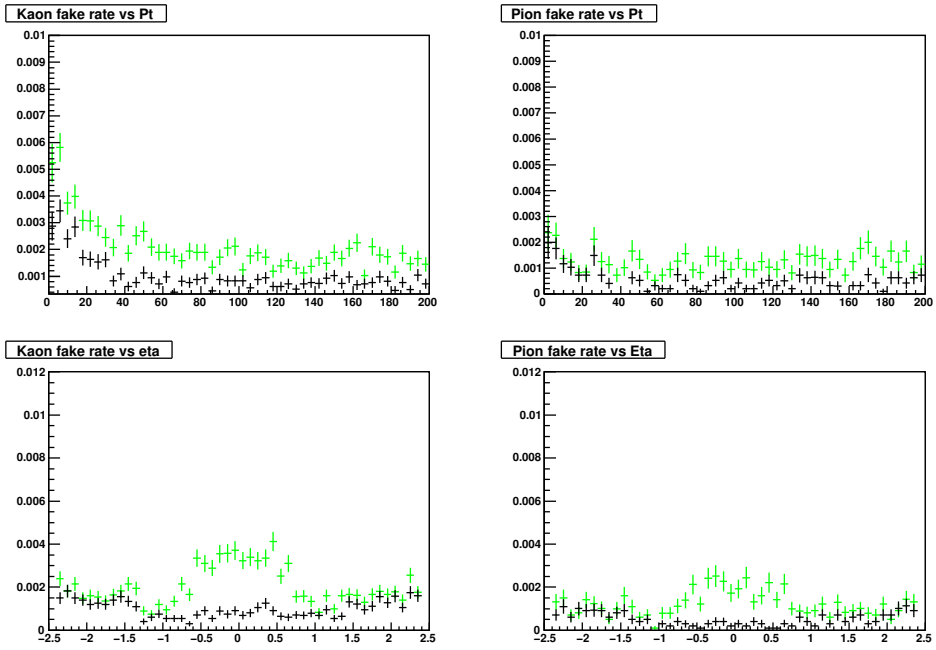

Figure 19: Fake rates for pions and kaons as a function of p_T and η for `TMLastStationLoose` (black) and `TM2DCompatibilityLoose` (green). These include track quality cuts.

- The fake rates and efficiencies are generally a bit better than for Global Muons (compare Figures 12, 19, and 20; also Figures 11 and 21).
- There are localized dips in efficiency for the `TMLastStation` family of requirements in the barrel. This will be investigated and hopefully remedied in the future.

6.5 Low p_T region

At low p_T in the barrel region muons do not penetrate deep into the muon detectors. The `TMLastStation` algorithms require that there be two matched segemts, and this introduces an inefficiency that can be in principle recovered. An example of the efficiency loss is shown in Figure 22.

As discussed in Section 6.3.2, the `TMOneStation` requirements are designed to recover the loss of efficiency of the `TMLastStation` requirements at low p_T, keeping the fake rate reasonably low. This is illustrated in Figures 23 and 24.

6.5.1 Efficiency-optimized cut-based selection at low p_T

Based on the results of Figures 23 and 24, an optimized cut-based selection at low p_T optimized for high efficiency would be

- Use `TMOneStationLoose` or `TMOneStationTight` for $|\eta| < 1.2$ and $p_T < 8$ GeV.
- Use `TMLastStationLoose` or `TMLastStationTight` at higher η and/or higher p_T.

To (hopefully) make it easier for the user, `TMLastStationOptimizedLowPt` bits have been defined to correspond to this (trivial) combination.

6.6 Comparison with Global Muons at low p_T and low η

For completeness, we include here a brief comparison of the performance of Global Muons and Tracker Muons at low p_T in the barrel region, see Figure 25. This is the region where the reconstruction of Global Muons is most challenging and where the gain from Tracker Muons is most important. As expected, there is a significant gain in efficiency from Tracker Muon reconstruction in this region.

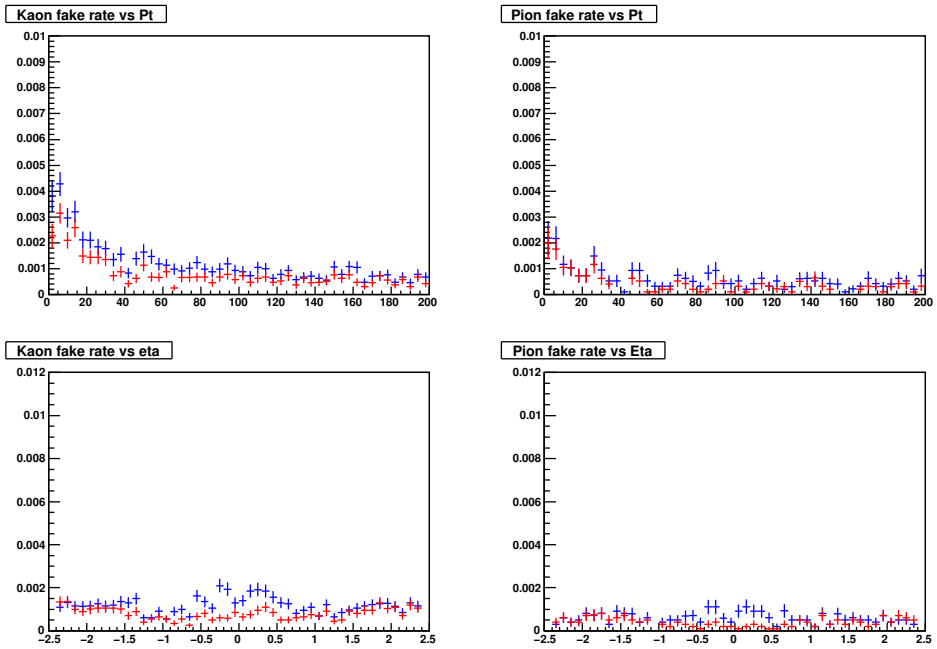

Figure 20: Fake rates for pions and kaons as a function of p_T and η for `TMLastStationTight` (red) and `TM2DCompatibilityTight` (blue). These include track quality cuts.

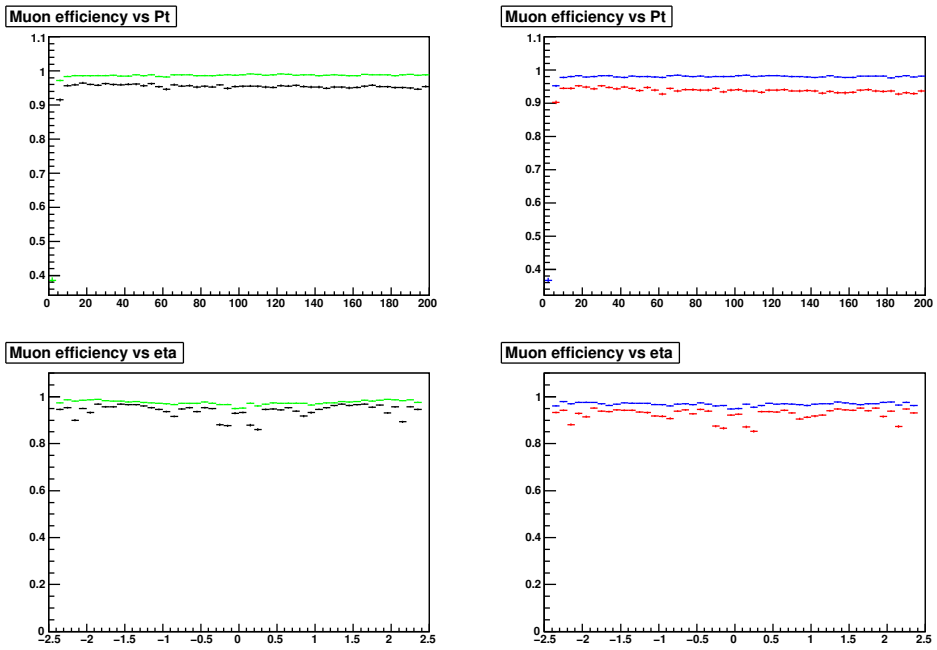

Figure 21: Efficiency for muons as a function of p_T and η for `TMLastStationLoose` (black), `TM2DCompatibilityLoose` (green), `TMLastStationTight` (red), `TM2DCompatibilityTight` (blue). These include track quality cuts.

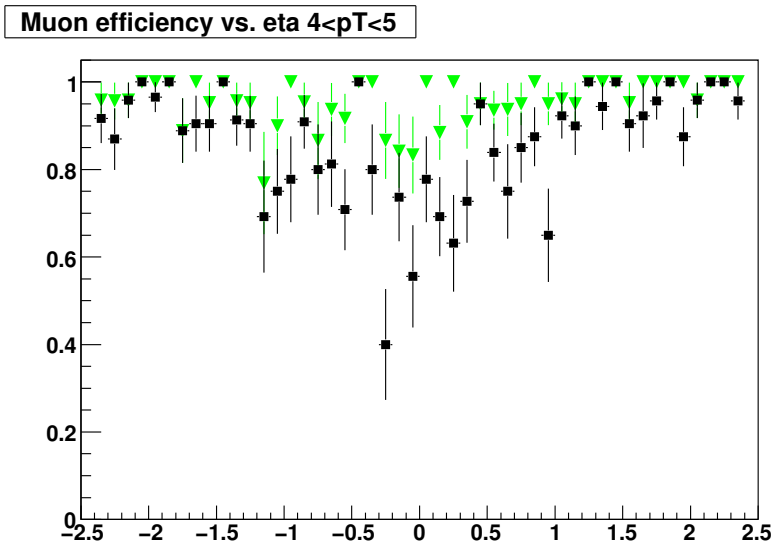

Figure 22: Efficiency for muons as a function of η with $4 < p_T < 5$ GeV for `TMLastStationLoose` (black squares) and `TM2DCompatibilityLoose` (green triangles). These include track quality cuts.

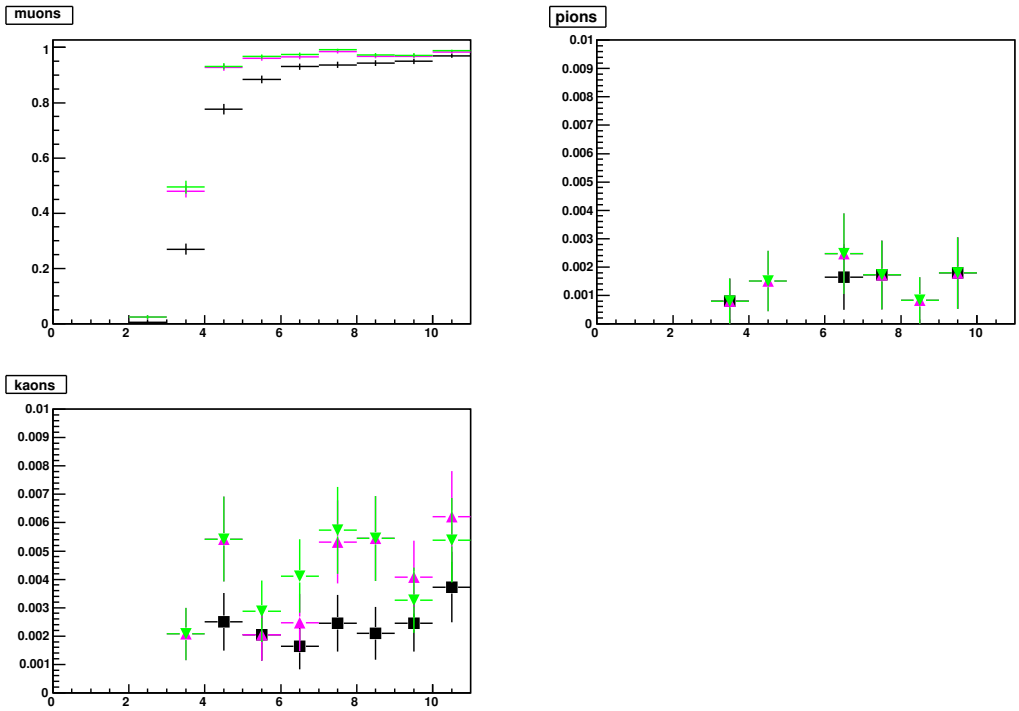

Figure 23: Muon efficiency and pion/kaon fake rates for $|\eta| < 1.2$ for `TMLastStationLoose` (black), `TM2DCompatibilityLoose` (green), and `TMOneStationLoose` (purple). These include track quality cuts.

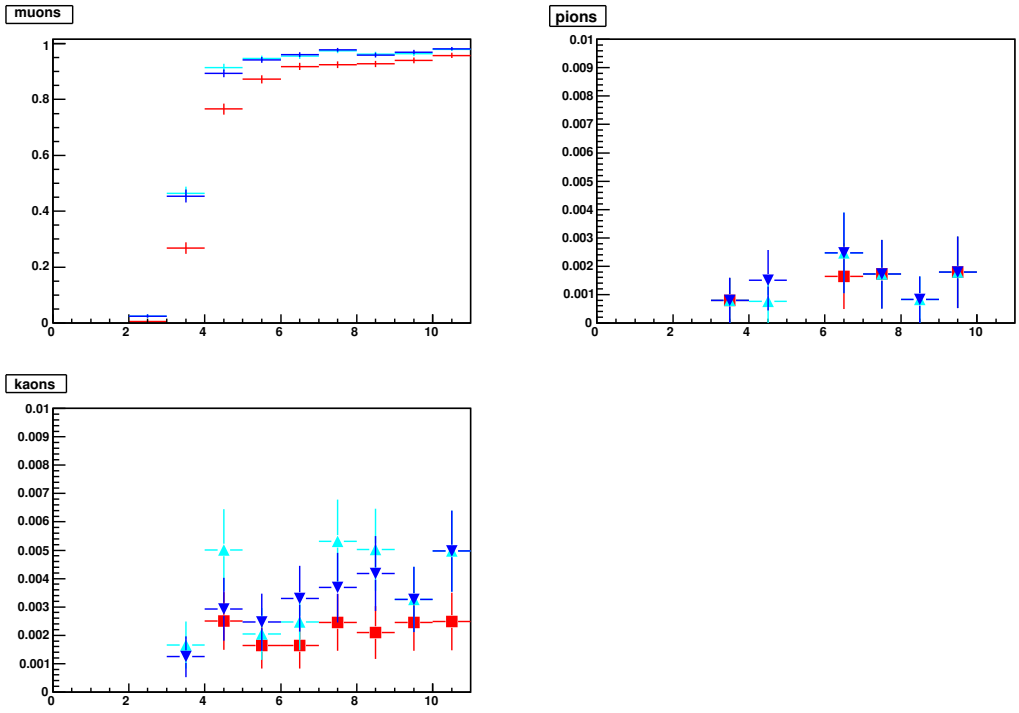

Figure 24: Muon efficiency and pion/kaon fake rates for $|\eta| < 1.2$ for `TMLastStationTight` (red), `TM2DCompatibilityTight` (dark blue), and `TMOneStationTight` (baby blue). These include track quality cuts.

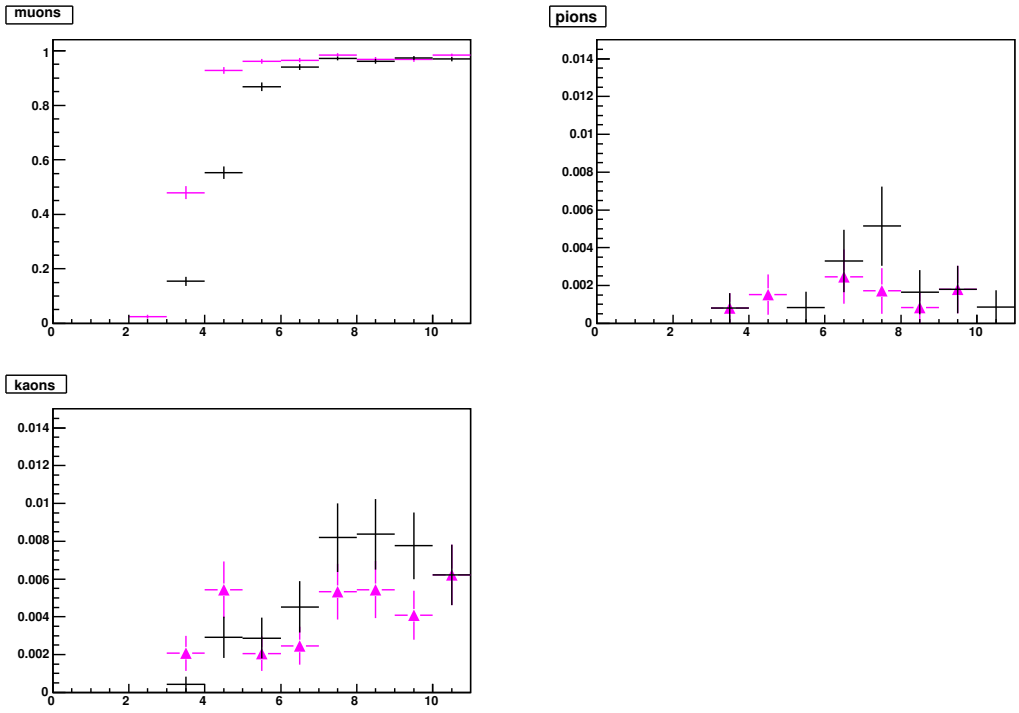

Figure 25: Muon efficiency and pion/kaon fake rates for $|\eta| < 1.2$ for `TMOneStationLoose` (purple) and Global Muons with global normalized-$\chi^2 < 10$ (black). Track quality cuts are included.

7 Efficiency in complex events

The results discussed until this point are based on single particle events. One worries that the efficiency may be lower in real events, particularly for muons inside jets. Possible sources of efficiency loss are:

- Silicon tracking inefficiency inside jets;
- Pattern recognition confusion inside jets which could cause a loss of efficiency of the track quality cuts discussed in Section 5.2;
- In the case of Global Muons, the matching between the standalone track and the silicon track could fail, *e.g.* the standalone track could be matched to a track different from the muon track within a jet;
- Similarly, in the case of Tracker Muons, the arbitration algorithm (see Section 6.1) could assign some of the segments to a a track different from the muon track within a jet;
- Additional particles within a jet occasionally deposit energy in the cells traversed by the muon. This could cause a loss of efficiency for the `TM2Dcompatibility` requirements, since they use the minimum-ionizing signature as one of their inputs.

We use a CMSSW 2.1.9 RelVal $\mathrm{t\bar{t}}$ sample to study these issues.

7.1 Procedure for efficiency test in $\mathrm{t\bar{t}}$ events

The idea is to compare the efficiency of the various selection requirements described in this note as obtained for single muon events and for muons from bottom and charm decays in $\mathrm{t\bar{t}}$ events. The determination of the efficiency in single muon events is straightforward, and it is described in Sections 5 and 6. The procedure in $\mathrm{t\bar{t}}$ events is a bit more complex and is described below:

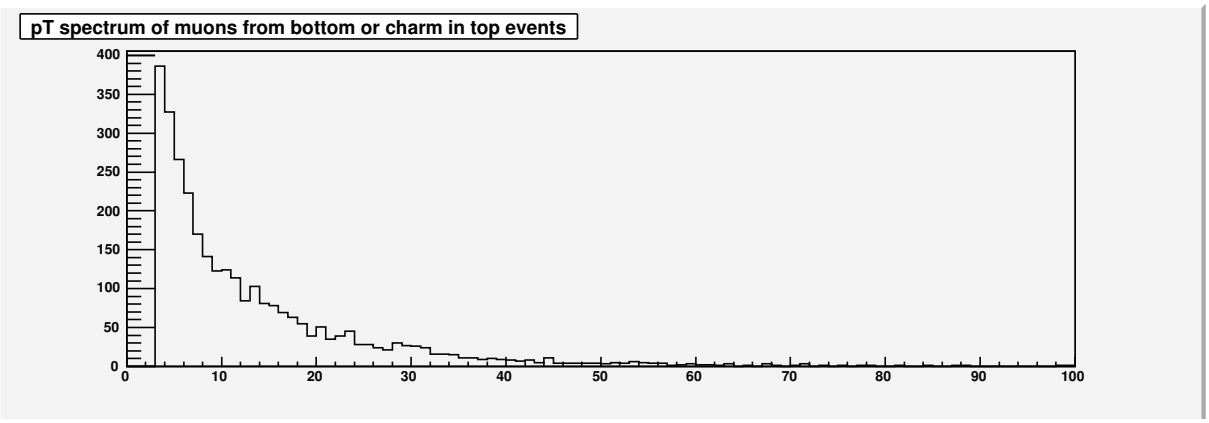

Figure 26: p_T spectrum of GEN muons from bottom and charm decays in $\mathrm{t\bar{t}}$ events ($p_T > 3$ GeV and $|\eta| < 2.4$).

- We loop over the GEN collection and we select muons from bottom and charm decays of $p_T > 3$ GeV and $|\eta| < 2.4$. There are 3047 such muons in this sample of 10050 events. Their p_T distribution is shown in Figure 26.
- We then look for silicon tracks within $\Delta R < 0.015$ of the GEN muon. This GEN muon$-$silicon track matching[3)] only fails for 2 out of 3047 muons[4)]. This shows that the tracking efficiency is very high. If more than one track is found, we only consder the one with the smallest ΔR.
- Next, we look for reco::Muons containing this silicon track, and we apply the various selection requirements to these muons.

[3)] This matching could have been done by hits; for our purposes the simple ΔR matching is sufficient.

[4)] In both cases a silicon track is found within ΔR of 0.025.

- To compare with the results from the single particle gun, we assign to each GEN muons from b- and c-decays in $\mathrm{t\bar{t}}$ events an efficiency as a function of p_T and $|\eta|$ as obtained from the single muon gun events. These efficiencies are computed separately for all sets of identification requirements that we have considered, *e.g.*, `TMLastStationLoose`, `TM2DcompatibilityLoose`, etc.

  We add up the GEN muons, weighted by their efficiency, and compare this "prediction" with the number of observed reco::Muons passing the requirements. Any difference (within statistics) between prediction and observation can then be ascribed to differences in identification efficiencies between muons in muon gun events and $\mathrm{t\bar{t}}$ events. The efficiencies are computed in bins of $\Delta p_T = 1$ GeV and $\Delta|\eta| = 0.2$; there are about 80 muons in the single muon gun sample per bin, so the statistics are not great, but they are adequate.

7.2 Results

The comparison described in Section 7.1 between the observed and predicted number of reco::Muons in $\mathrm{t\bar{t}}$ events matched to GEN muons from bottom and charm decays is shown in Table 2.

The uncertainties in Table 2. are calculated as follows:

- For the "observed" column, the uncertainty is just the binomial uncertainty for finding N reco::Muons passing the requirements out of 3047 GEN muons, (*e.g.*, for the 1st row, N=2504).
- For the "predicted" column, the uncertainty is derived from the statistical uncertainty in the efficiency as a function of p_T and $|\eta|$ as derived from the single muon sample. Each $p_T - |\eta|$ efficiency-bin has its own binomial uncertainty; this uncertainty is convoluted with the $p_T - |\eta|$ distribution of the the 3047 GEN muons to arrive at the final uncertainty. Note that this procedure is not totally correct, because binomial uncertainties are asymmetric, particularly for efficiencies near 100%, where the nominal binomial uncertainties go to zero. This results in an underestimate of the uncertainties. We have artificially set the efficiency uncertainties for the bins with 100% efficiency to sqrt(p*(1-p)/n), where p=(n-1)/n and n is the number of single muons in the bin in question (typically, n=80); this has a very modest impact on the calculated uncertainty, which implies that the uncertainty estimate is stable.

Requirement	Observed muons	Predicted muons	Difference
`GlobalMuonPromptTight`	2504 ± 21	2573 ± 8	$-2.7 \pm 0.9\%$
`TM2DcompatibilityLoose`	2846 ± 14	2843 ± 8	$+0.1 \pm 0.6\%$
`TMLastStationLoose`	2634 ± 19	2664 ± 9	$-1.1 \pm 0.8\%$
`TMOneStationLoose`	2829 ± 14	2825 ± 8	$+0.1 \pm 0.6\%$
`TMLastStationOptimizedLowPtLoose`	2749 ± 16	2774 ± 8	$-0.9 \pm 0.6\%$
`TM2DcompatibilityTight`	2688 ± 18	2800 ± 8	$-4.0 \pm 0.7\%$
`TMLastStationTight`	2604 ± 19	2629 ± 9	$-1.0 \pm 0.8\%$
`TMOneStationTight`	2812 ± 15	2798 ± 8	$-0.5 \pm 0.6\%$
`TMLastStationOptimizedLowPtTight`	2717 ± 17	2736 ± 8	$-0.7 \pm 0.7\%$

Table 2: Comparison of predicted and observed muons in $\mathrm{t\bar{t}}$ events. See text for details.

Note also that, for a fair comparison, the impact parameter requirement was removed from the track quality cuts in Table 2. This is because muons in single muon events are generated from a single point, whereas muons from $\mathrm{b} \to \mu$ and $\mathrm{c} \to \mu$ are affected by the heavy flavor lifetimes (this is a 0.6% effect for the nominal 2 mm requirement).

The differences in efficiencies between single muon gun and busy $\mathrm{t\bar{t}}$ events are quite small. This gives us confidence in our ability to use these requirements in the LHC environment.

8 Selection Code, Tags, Final Comments

As mentioned in Section 1, the selections described in this note can be accessed via the following member function of the reco::Muon object:

```
bool isGood( SelectionType type = AllArbitrated )
```

They correspond to V07_02_11 of `DataFormats/MuonReco`. Note that this tag can be used at the analysis level. It is backward compatible with events reconstructed with CMSSW 2.0 and later.

The recommended track quality requirements, $|d_0| < 2$ mm and $N_{hits} \geq 11$ (for the silicon fit) have to be applied externally, for the reasons mentioned in Section 1. Note that d_0 must be calculated from the beamspot or the primary vertex.

The selections described in this note are:

1. `GlobalMuonPromptTight`
2. `TM2DCompatibilityLoose` and `TM2DCompatibilityTight`
3. `TMLastStationLoose` and `TMLastStationTight`
4. `TMOneStationLoose` and `TMOneStationTight`
5. `TMLastStationOptimizedLowPtLoose` and `TMLastStationOptimizedLowPtTight`

Final comments are summarized below.

- All the selections described here operate on Tracker Muon information, except for `GlobalMuonPromptTight`. Note however that the great majority of Global Muons are also reconstructed as Tracker Muons. Therefore, it is also possible to apply Tracker Muon-based selections for muons selected as Global Muons.
- The `GlobalMuonPromptTight` selection can be augmented by simple "penetration" requirements, as discussed in Section 5.3, with the associated caveats also discussed there. These additional requirements would decrease the fake rate by $\approx 20 - 30\%$.
- The Tracker Muon reconstruction is more efficient at low p_T in the barrel.
- The Tracker Muon-based selections tend to have slightly better fake rates and efficiencies than the `GlobalMuonPromptTight` selection.
- The differences in performance between the `Loose` and `Tight` versions of the Tracker Muon requirements are small.
- The `TM2DCompatibility` selections include requirements on consistency with minimum ionizing energy deposition. This can introduce a small inefficiency for muons in jets. The other selections do not consider the calorimeter energy deposition; separate requirements on `CaloCompatibility` could easily be added, however these were not considered in this note.
- The `TMOneStation` requirements are the loosest ones, and should probably only be used in the barrel at low p_T or if backgrounds are not a concern.
- The `TMLastStation` requirements are generally tighter than the `TM2DCompatibility` requirements.
- The `TMLastStationOptimizedLowPt` requirements combine `TMOneStation` at low p_T in the barrel ($p_T < 8$ and $|\eta| < 1.2$) with `TMLastStation` at higher p_T and/or higher $|\eta|$.
- The `TMLastStation` requirements show localized dips in efficiency in the barrel. This will be fixed in a future release.
- The performance of these selection criteria has been mostly studied and tuned on single particle events. However, it appears that the performance in busy events *e.g.*, $\mathrm{t\bar{t}}$ is not affected significantly, at least for what concerns the efficiency.

References

[1] G. Abbiendi *et al.*, "Muon Reconstruction in the CMS detector", CMS AN-2008/097 (in preparation).

[2] `http://aeverett.web.cern.ch/aeverett/MuonHLTPages/noteDataSets.html`

[3] L. Bauerdick *et al.*, "Data-driven background determination in μ+jets", CMS AN-2008/040.