Correcting Missing Transverse Energy Using Tracks

F. Golf and Friends…

1. Introduction

Missing Transverse Energy (MET) is an important observable used to establish proper detector performance, observe basic SM signatures and search for new physics. Simple in its definition, a true measurement of MET requires understanding the interaction physics as well as proper detector alignment. Of the basic physics objects, it is the most susceptible to detector calibration and noise problems.

Two main features characterize the MET performance of a detector: resolution and the non-Gaussian tail. Both properties suffer in CMS from

- the large material budget of the tracking system,
- the non-linear response of the calorimeter,
- the strong (3.8T) magnetic field

The last two points are particularly enlightening for they explain the two primary ways that MET is mis-measured. The energy a particle deposits is generally under-measured due to the non-linear response of the calorimeter. This mis-measurement of the magnitude of the deposition is not the whole picture though. The strong magnetic field in CMS bends the trajectories of charged particles, causing energy deposited in the calorimeter to be displaced in phi relative to the particle's initial direction.

Both issues of magnitude and direction have to be addressed in order to obtain an accurate determination of missing E_T. We have developed a method to improve performance in measuring MET by taking advantage of the CMS tracker. At the scale of interest (~1 GeV), the tracker has excellent resolution compared to the calorimeter. The proposal is simple:

Correct the measured missing E_T by replacing, for all well reconstructed tracks, the average (or expected) energy deposition in the calorimeter by the measured momentum in the tracker.

2. The problem: Fake MET in $Z \to \ell\ell$ events

Any search for new physics will begin with the study of SM processes necessary to calibrate the detector and develop experimental techniques. These processes

present significant backgrounds that distort signatures of new physics. Top quarks, for example, due to the large production rate and complex signature, will be a serious background to many BSM searches. More details may be found in our public Physics Analysis Note on the ttbar cross-section[x].

The n-jet distribution (right) for di-leptons plus MET illustrates some of the processes we have been investigating and our current understanding of their relative contributions. The figure to the right shows the n-jet distribution for the di-muon final state after requiring:

- lepton $p_T > 20$ GeV,
- MET > 30 GeV ,
- count jets with $p_T > 30$ GeV (corrected) and $|\eta| < 3$

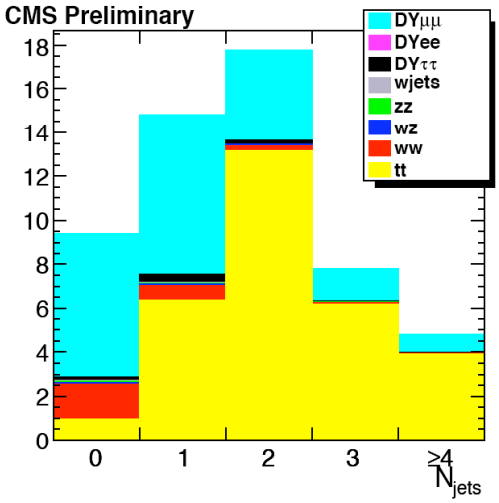

The sizable contribution of $Z \rightarrow \mu\mu$ (cyan) is an indication of the problem CMS faces in measuring MET. At this level, it will be a serious background to any search for new physics whose signature involves missing E_T.

3. Releases & Data Sets Used

This work was done in the context of CMSSW release 1.6.12. The studies were performed on the following samples.

a) Single π

/RelValSinglePiPlus0To100/CMSSW_1_6_7-RelVal-1202415807/GEN-SIM-DIGI-RECO

b) DY→ℓℓ

/DrellYan_ll_40/CMSSW_1_6_7-CSA07-1201885455/RECO

c) W→μυ

/Wmunu/CMSSW_1_6_7-CSA07-1192835438/RECO

d) W→eυ

/Wenu/CMSSW_1_6_7-CSA07-1197047869/RECO

4. Procedure

The goal then is to construct a procedure to correct MET on a track-by-track basis. The logical place to begin is with a sample of single charged pions and so will our discussion. Beginning with this sample, the track and calorimeter information is extracted and a set of filters is used to select a subset of well-reconstructed pions. From this subset, a fixed lookup table of detector response in η, p_T is built.

The following section will discuss the implementation methodology that uses the response function to remove expected depositions in the calorimeter and replace them with values determined from the tracker. Special treatment is given to leptons. The performance evaluation of the track-corrected MET (tcMET) is discussed next. The effects of this correction on Drell-Yan samples with fake missing E_T is compared to events with real MET, such as W production.

A few comments on the correction's effect on missing E_T will follow. The final section will remark briefly on hybrid implementations of track-based MET corrections. Our work tries to address large tails in the Drell-Yan MET distribution, particularly in the 0-jet bin which contains unclustered energy. However, in noisier environments, additional gains may be made by considering alternate response functions such as that used in JPT, jet corrections such as Type-I JES corrections, corrections for neutrals, etc.

5. Derivation of the Response Function (RF)

In order to account for the non-linearity of the calorimeter we need to understand how the detector responds to charged hadrons (leptons receive special treatment and will be discussed later). The goal then is to create a fixed lookup table that allows for a simple determination of the expected detector response based only on track kinematics.

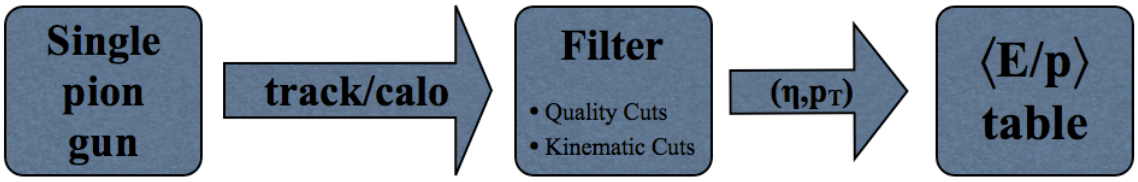

The diagram above shows the conceptual steps involved in the derivation of the response function. Starting with a single particle gun, the track and calorimeter information for each pion is extracted and passed through a set of filters. If the pion passes all cuts, the detector response is calculated and entered into the

appropriate η,p_T bin of a 2d histogram. Finally, the response function is extracted from the histogram. More detail on each of these steps follows below.

a) Selection of pions

The starting point is a sample of single charged pions. Each pion in the sample is passed through a kinematic filter that selects those tracks with:

- $p_T \in (2,100)$ GeV,
- $|\eta| < 2.4$

The requirement on η is imposed by the tracker. The p_T cuts select tracks that are both abundant and well-measured. Those tracks passing the kinematic cuts are then passed through an additional quality filter. We require

- number of hits > 7,
- χ^2 / ndof < 5,
- |d0| < 0.01

The first two are obvious. The third, d0, usually gets distorted to some "large" non-zero value if the estimate of the curvature is in err and hence is a good indicator of badly measured tracks. The three N-1 distributions for our sample of pions are shown below, justifying the choice of cuts made above.

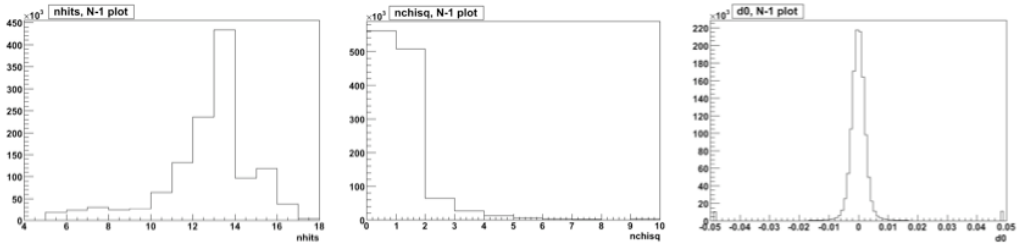

Quality cuts are tight to ensure that only good tracks go into the lookup table and prevent outliers from distorting the response.

b) Definition of the response

The events that pass kinematic and quality filters form the working subset of pions from which the response function is built. For each pion, the momentum, p, measured at the vertex serves as a seed to the analytical propagator which swims the trajectory out to the expected location of entry of the particle into the calorimeter. Using this point as the center, a cone of $\Delta R = 0.5$ is opened and the energy in each calorimeter tower is summed to determine the energy the pion deposited in the detector.

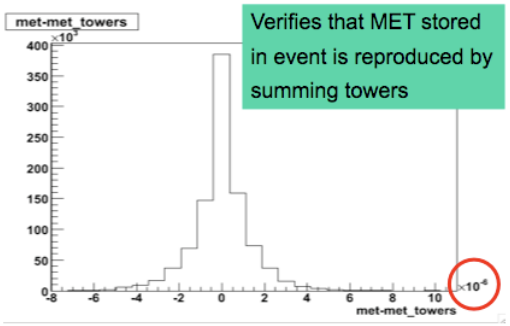

The use of towers rather than rec hits is a deliberate and important one. Missing E_T is calculated as a vector sum of calorimeter towers as verified to the right. Thus, towers are a natural calorimeter segmentation with which to determine detector response. The difference goes beyond geometry though. Rec hits are subject to different thresholds than than towers and thus can give different results. This is particularly important in the 0-jet bin where unclustered energy dominates. Put simply, a naive sum over rec hits in the cone can give a significantly different result than the corresponding sum over towers. This having been said, there are environments in which rec hits may be the more appropriate choice, such as the cores of jets. More on this last point will be discussed in a later section.

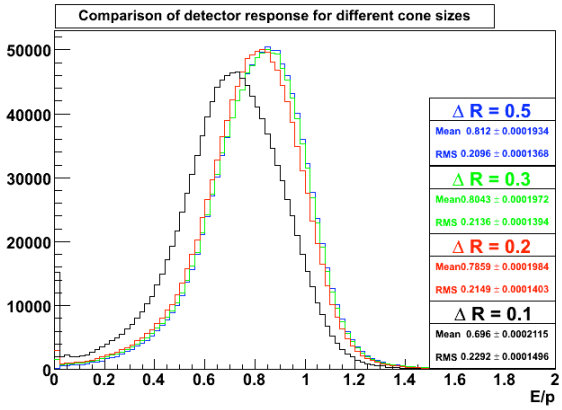

The detector response, E/p, is defined to be the energy deposited in the cone divided by the magnitude of the track momentum measured at the vertex. The histogram to the left shows the detector response for several different cone sizes. One notices right away that a cone of $\Delta R=0.1$ provides a much lower response that do larger ones. As the cone size increases, there is a significant improvement for $\Delta R=0.2$ and then the response improves asymptotically for larger radii. This observation is not surprising for the clean environment of a single particle sample. Given this flexibility, a large cone size of $\Delta R=0.5$ was chosen which is reasonable given the particle source.

c) Response function

The response for each good pion is determined in this way and entered into a two dimensional histogram with variable bins in η, p_T. Once again, the use of variable bin sizes is an important and intentional choice. Bins in eta coincide with towers in the calorimeter, a logical choice given the choice of definition of the response. On the other axis, bins have 1 GeV size at low p_T and increase to 10 GeV at the high end of the spectrum. This last point is important. A larger variation in

detector response is expected at low p_T, thus requiring finer binning, than higher pT, where coarser separation will suffice. Once the histogram is filled, a one dimensional histogram of the response in each bin is created and the mode (i.e. largest bin) is extracted. The extracted value fills the corresponding bin of the response function.

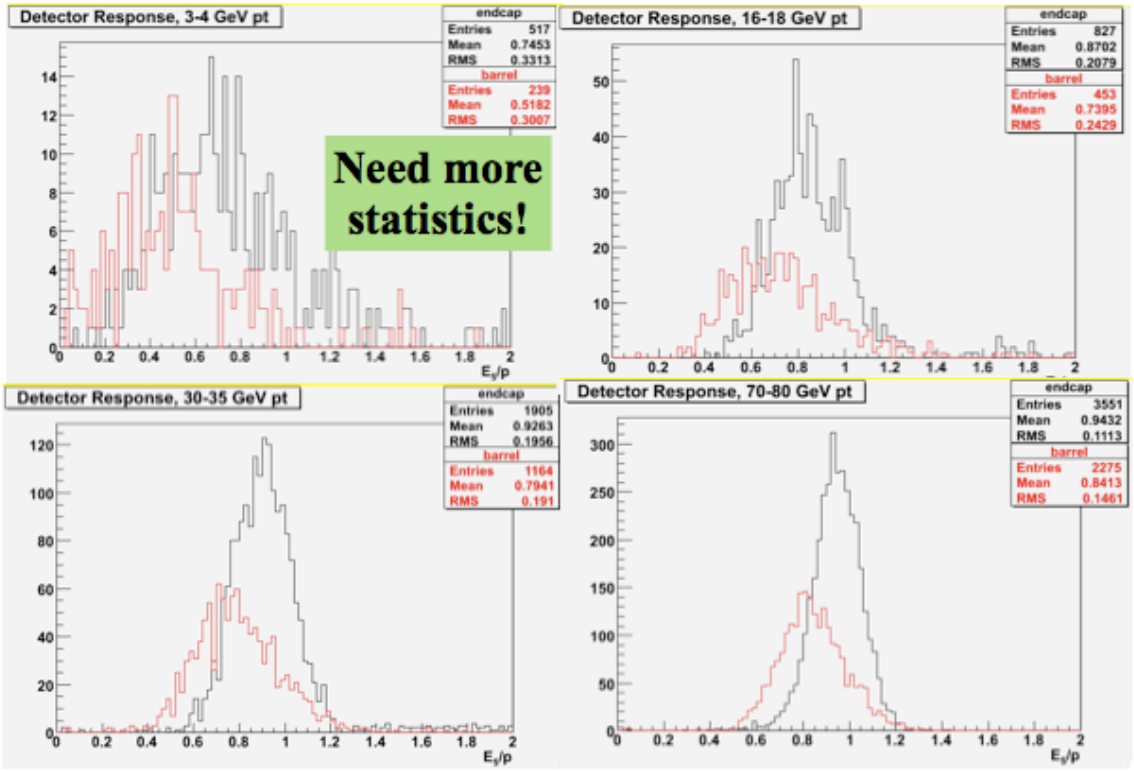

The figures above show detector response in the end cap and barrel for several different p_T intervals. Two features general are immediately manifest:

- for a given p_T interval, response is better in the end cap than in the barrel
- response improves with increasing p_T

Better response in the end cap is likely due to the increased material budget at larger η. Improved response at higher p_T is expected. It is also important to note the large amount of noise as low p_T. This is mostly a result of low statistics. Larger samples for these bins will be produced when the response function is derived in the next release.

Extracting the response using the mode is a common method, but not the only one possible. Several other options were evaluated. The table below compares the mode with two other methods.

	Mean	Gaussian Fit	Mode
p_T=2.5, η=0.05	0.52	0.49	0.27
p_T=2.5, η=2.25	0.75	0.66	0.43
p_T=17, η=0.05	0.74	0.7	0.63
p_T=17, η=2.25	0.87	0.84	0.85
p_T=32.5, η=0.05	0.79	0.79	0.79
p_T=32.5, η=2.25	0.93	0.9	0.87
p_T=75, η=0.05	0.84	0.84	0.77
p_T=75, η=2.25	0.94	0.95	0.93

The response in the first column was found by calculating the mean of entries in each bin. The value in the second column was determined by fitting each bin with a gaussian on the response interval (0.1, 1.2) and extracting the mean. This basically acts like a truncated mean. The response in the third column, the one used in the lookup table, is determined from the mode in each bin. A couple of general features are evident. The mean and gaussian fit yield similar values, except at low p_T and large η. The mode tends to give a slightly lower response than the other methods. This is particularly evident at low p_T.

6. Application of the RF

Track-corrected MET (tcMET) is calculated for an event using:

- caloMET
- muon collection
- electron collection
- track collection
- response function (RF)

For each event, the algorithm begins by identifying and correcting muons. Muons are corrected at the outset by subtracting the p_T of each associated track from the caloMET.

$$\begin{aligned} \text{baseline MET} &= \text{caloMET} - \sum_{\text{muons}} \vec{p}_T, \\ &= -\sum_{\text{towers}} \vec{E}_T - \sum_{\text{muons}} \vec{p}_T \end{aligned}$$

The caloMET corrected for muons will serve as the baseline MET for comparisons throughout this paper. Next, the algorithm identifies "electron-like" objects and skips them - no correction is applied for tracks matched to these objects. This is an appropriate as these objects deposit a large fraction of their energy electromagnetically. Further discussion on electrons follows below.

The remainder of the tracks are candidates for correction using the derived response function. The response function serves as a fixed lookup table that takes as input the kinematics of a track and returns as output the expected energy deposited by the track in the calorimeter. This sequence is represented in the figure below.

Those tracks not matched to leptons are passed through a set of filters cutting on kinematic quantities as track variables, similar to those used in the derivation of the response function. The subset of tracks that survive form the array of "good tracks" for which MET is corrected. The actual correction is implemented by removing the expected energy deposited by each good track in the calorimeter, determined using the response function, and replacing it with the track momentum at the vertex.

$$\begin{aligned} \text{tcMET} &= \text{baseline MET} + \sum_{\text{good tracks}} \langle \vec{E}_T \rangle - \sum_{\text{good tracks}} \vec{p}_T, \\ &= -\sum_{\text{towers}} \vec{E}_T - \sum_{\text{muons}} \vec{p}_T + \sum_{\text{good tracks}} \langle \vec{E}_T \rangle - \sum_{\text{good tracks}} \vec{p}_T \end{aligned}$$

It is important to note that the correction for each good track involves two sets of coordinates. The expected energy deposition for each track is removed from the calorimeter. This location is determined using the vertex track as a seed to the analytical propagator, as was done in the derivation of the response function. The track momentum that replaces it is taken at the vertex. To be explicit, the correction for a single component of MET takes the form:

$$(\mathrm{tcMET})_x = (\text{baseline MET})_x + \sum_{\text{good tracks}} \langle E \rangle \sin\theta_c \cos\phi_c - \sum_{\text{good tracks}} p_T \cos\phi_v$$

where θ_c, ϕ_c are the polar and azimuthal position coordinates of the particle at the calorimeter face and ϕ_v is the azimuthal angle of the track at the vertex.

a) Track quality cuts - how <u>not</u> to generate fake MET

Badly measured tracks can generate fake missing E_T and correcting for such a track can make things worse. Consider a track that, due to a pattern recognition problem in a noisy environment, has a very large estimated momentum. The amount of this over-estimate will translate directly into excess MET if the track is passed to the response function. A set of kinematic and quality cuts is used to filter these mis-measurements from the track collection and avoid generating fake missing E_T.

It is necessary that the same kinematic cuts applied in the derivation of the response function are also used in its application. However, the quality cuts need not take the same values and are loosened here in the interest of inclusivity. The application of the response function defines good tracks to have:

- number of hits > 6,
- χ^2 / ndof < 5,
- |d0| < 0.05

The set of N-1 distributions below are from the Drell-Yan sample, justifying the choice of cuts above.

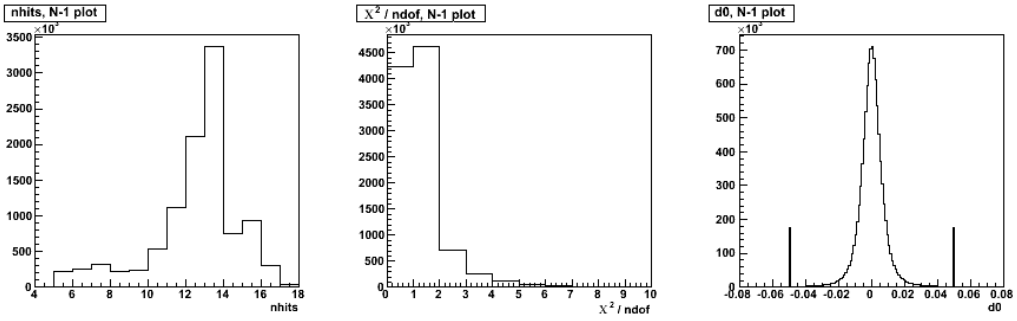

b) Correction for electron-tracks

Applying the response function to an electron-like track, i.e. either an electron or a pion that showered predominantly in the ECAL, will generate fake MET. Thus, it is important to identify and avoid correcting such particles. The discriminating

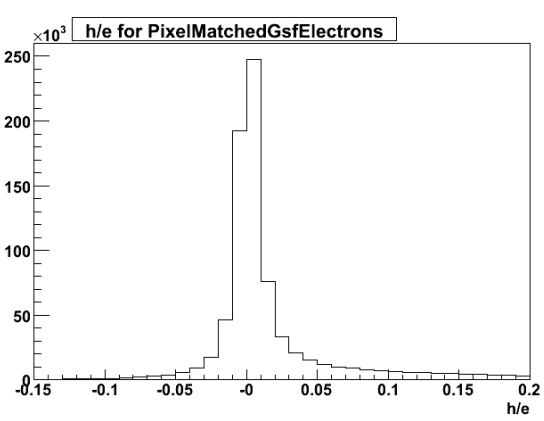

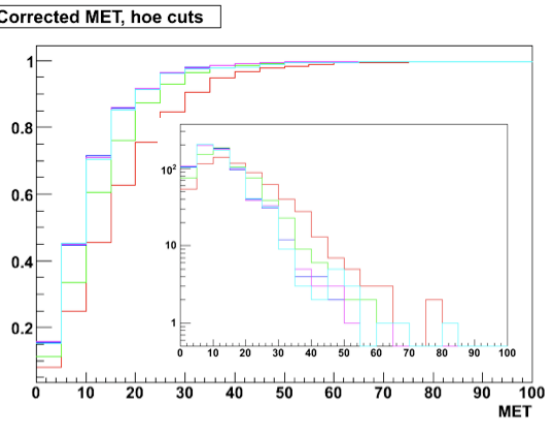

variable h/e, the ratio of hadronic to electromagnetic energy deposited, is shown in the above figure on the left. The efficacy of a cut on h/e was studied on a subset of the Z→ee sample. The figure on the right shows tcMET, where the h/e cut has been varied. The main histogram is a cumulative representation of the fraction of events with tcMET < X. The inset is a log scale plot of tcMET for the same scenarios. The fraction of events with tcMET > 30 for the various cut values is given in the table below.

	Fraction of events with tcMET > 30 GeV
No Cut	0.15
|h/e| < 0.03	0.07
|h/e| < 0.05	0.04
|h/e| < 0.10	0.04
|h/e| < 0.20	0.04

The figures and table clearly demonstrate that some cut on h/e is mandatory. Compared to not discriminating based on the h/e variable even a very loose cut of 0.03 sees a factor of 2 improvment which improves to a factor 4 for a tighter selection. Skipping objects in the PixelMatchGsfElectronCollection with |h/e/ < 0.1 is well motivated by the results above.

7. Performance Evaluation

The goal at the outset was to reduce the sizable tails of the Drell-Yan MET distribution. The correction was tested on a Drell-Yan sample with the metric being the number of events with MET > 30. In addition, the response function was also used to correct a sample of W leptonic decays. The use of a sample with

real missing E_T is an important control. Correcting the Drell-Yan tails with a naive additive or multiplicative factor will eliminate events with high MET, but will also do so for events with real missing E_T. Thus, it is necessary check that any attempt to correct for mis-measured MET does not significantly impact the ability to identify those events for which missing E_T is an important and real signature.

Separate evaluations will be performed for two distinct cases: di-muon and di-electron final states. Consider the former to begin. The baseline missing E_T is defined to be caloMET plus muon corrections. Additionally, a "good muon" will be defined to have:

- $p_T > 20$,
- χ^2 / ndof < 5,
- number of hits > 6,
- |d0| < 0.25

Only Drell-Yan events with exactly two good muons and W decays with at least one good muon in the final state are used for comparison. These requirements ensure the former contains predominantly events with fake MET and the latter with real.

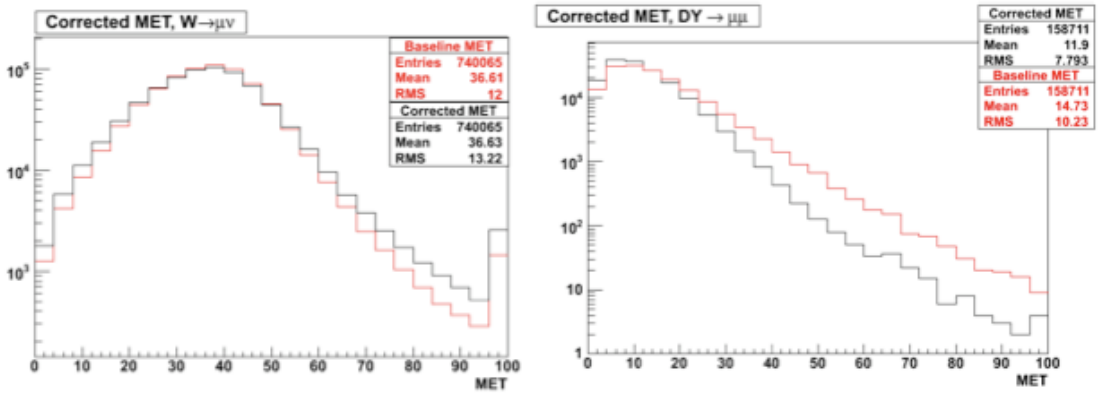

The figures above show baseline MET in red and tcMET in black. The W distribution with real MET on the left shows no significant change after correction. The tail of the Drell-Yan distribution on the right shows a dramatic reduction in the number of events with large missing E_T. The number of events with MET > 30 is reduced by a factor of 2.7. The reduction is enhanced to a factor of 4.7 for MET > 50.

It is also useful to break these results down into n-jet bins. Here, the jet counting parameters are:

- $E_T > 15$ (uncorrected),
- $|\eta| < 3$

Changes in the W sample are less than 5% in any n-jet bin. The table below shows the breakdown for Drell-Yan. Each entry shows the fraction of events with MET > 30 in raw numbers and as a percentage. The final row shows the factor of improvement for tcMET over baseline MET.

	0 jets	1 jet	2 jets	3+ jets
baseline	4006/120177 3%	6278/32161 20%	1859/5517 34%	363/856 42%
tcMET	1687/120177 1%	1981/32161 6%	767/5517 14%	165/856 19%
factor of improvement	2.4	3.2	2.4	2.2

The majority of Drell-Yan events fall into the 0-jet bin. These events are the target of the correction and reduction of the tail by a factor of 2.4 is reason for optimism. Although this correction wasn't conceived with jets in mind, the large number of tracks in these events gives reason to expect similar reductions and the results do not disappoint. The improvement in the 2-jet bin is commensurate with that seen in events with no jets. The 1-jet bin shows even better pruning. This is not entirely unexpected. These events contain one sizable jet which provides additional tracks for correction. The presence of a single large jet that is typically under-measured produces significant asymmetry, providing excellent conditions for correction. The observation of a smaller improvement in the 2-jet bin is also reconciled in this view as the two jets typically have some spatial separation which results in corrections for tracks in one jet partially canceling out corrections for tracks in the other jet. This cancellation becomes more evident in those events with more than 2 jets.

A similar comparison is also done for the electron final state. Only Drell-Yan events with exactly two electrons with $p_T > 20$ are used for comparison. The requirement for W decays is the presence of at least one electron with the same p_T cut.

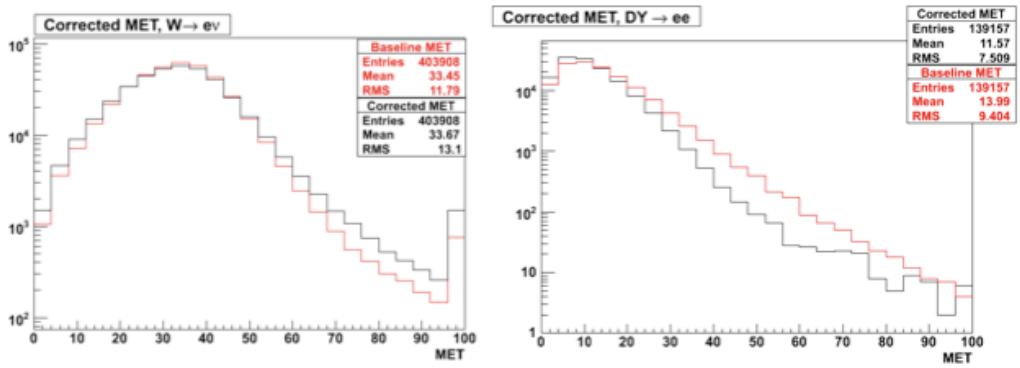

The figures above show baseline MET in red and tcMET in black. The W distribution with real MET on the left shows no significant change after correction. The tail of the Drell-Yan distribution on the right shows a significant reduction in the number of events with large missing E_T. The number of events with MET > 30 is reduced by a factor of 2.7, the same as in the muon final state.

The table below shows the breakdown by n-jet bin for Drell-Yan. The W sample is not changed by more than a couple percent in any bin. In addition to the previously jet counting criteria, additional care must be taken here to remove electrons from the list of jets.

	0 jets	1 jet	2 jets	3+ jets
baseline	3117/108383 3%	4169/26291 16%	1057/3968 27%	205/515 40%
tcMET	1392/108383 1%	1329/26291 5%	378/3968 10%	103/515 20%
factor of improvement	2.2	3.1	2.8	2.0

These results are similar to those seen for the muon final state. The correction for the electron final state is slightly less effective in the 0-jet bin and slightly more effective in the 2-jet bin than that observed for muons. The variation in performance amongst the various bins can be explained by the same reasons given in the muon case above.

8. Missing E_T Resolution

The track-based correction shown above is quite effective at reducing the tail of the MET distribution in Drell-Yan. No effort was made to try to improve resolution. Although reducing the tails is a desirable goal, tcMET may be of limited utility if it significantly worsens MET resolution. A crosscheck was performed by fitting the distribution of each component of missing E_T with a Gaussian from which the the resolution determined using

$$\sigma_{MET} = \sigma \cdot \sqrt{\frac{4-\pi}{2}}$$

as outlined in equation 3 of AN-2007/041. Here, σ is the width of the distribution of one of the components. Here, we take σ to be the average of the widths of the two distributions.

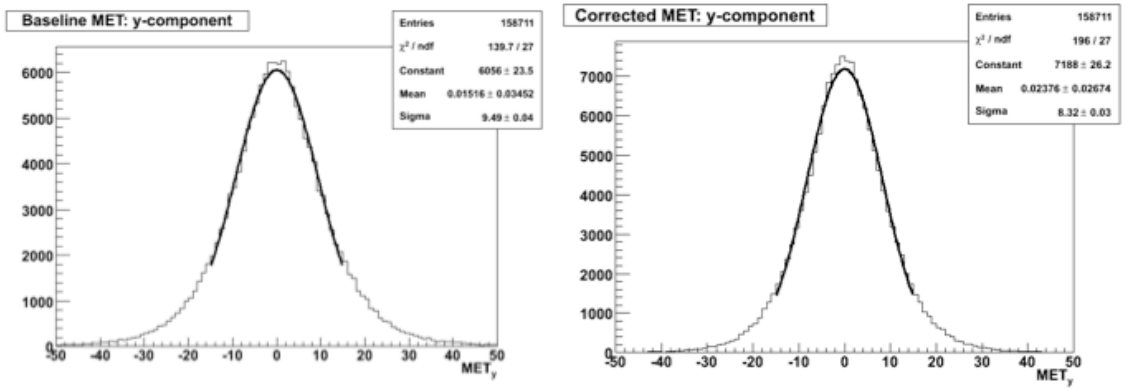

The figures below show the distributions of the y-component of baseline MET (left) and tcMET (right) with Gaussian fits overlaid. The table below shows the width of each component and the MET resolution calculated from these values.

	σ_x	σ_y	σ_{MET}
baseline MET	9.43	9.49	6.2
tcMET	8.29	8.32	5.4

MET resolution gets better after applying the track-based correction! The resolution improves by ~15%. This is an unintended bonus - we set out to reduce the tail and along the way also improved the resolution. The presence of this enhancement validates procedure and indicates that the algorithm developed is a reasonable one.

9. Jet Energy Corrections and MET

Correcting MET for tracks shows large reductions of the tails of the MET distributions for Drell-Yan. In particular, the algorithm corrects for unclustered energy in the 0-jet bin, a feature that does not currently exist in the standard missing E_T corrections. However, type-I JES corrections do correct for clustered energy and thus it is interesting to ask what effect, if any, do these on the Drell-Yan sample.

A comparison was made between baseline MET and baslinet MET + type-I JES corrections. Only events with two good muons or electrons were kept. The standard type-I JES corrections as implemented in CMSSW_1_6_X were used.

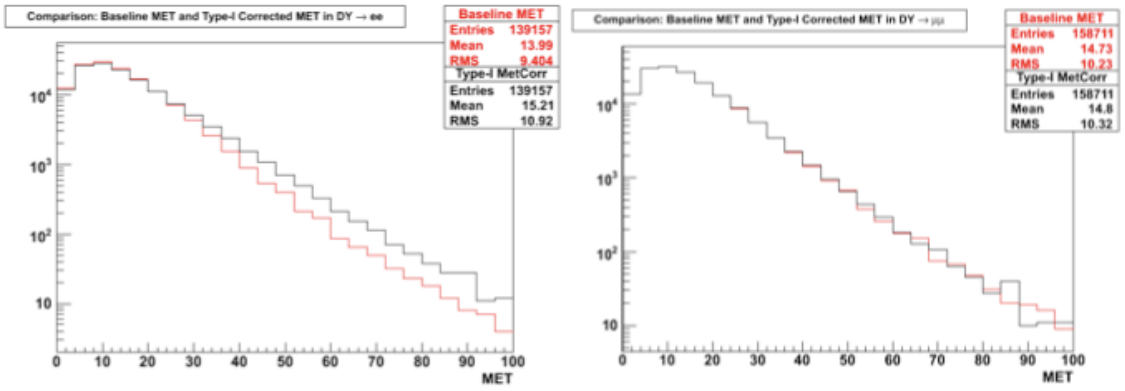

The figures above show the comparison for the electron (left) and muon (right) final states. Baseline MET is shown in red while the addition of the JES corrections is shown in black. No improvement over the baseline is observed in the muon final state. Type-I JES corrections make the tail considerably worse in the electron final state. The suspicion is that the JES implementation is correcting electrons as jets, thus generating fake MET. However, according to the documentation, this should not be the cause as electrons are explicitly removed and there is an additional cut of $|h/e| < 0.1$. This discrepancy has been reported to the MET group and is under investigation. In summary, to reduce the tails of the MET distribution of Drell-Yan, tcMET performs significantly better than the type-I JES corrections.

10.Conclusion