

University of California
Santa Barbara

**A search for supersymmetry with the CMS detector
in the single-lepton final state using the sum of
masses of large-radius jets**

A dissertation submitted in partial satisfaction
of the requirements for the degree

Doctor of Philosophy
in
Physics

by

Ryan Edward Heller

Committee in charge:

Professor David Stuart, Chair
Professor Claudio Campagnari
Professor David Berenstein

June 2018



The Dissertation of Ryan Edward Heller is approved.

Professor Claudio Campagnari

Professor David Berenstein

Professor David Stuart, Committee Chair

June 2018

A search for supersymmetry with the CMS detector in the single-lepton final state
using the sum of masses of large-radius jets

Copyright © 2018

by

Ryan Edward Heller

To the thousands of scientists, engineers, and others whose efforts have rewarded us with the incredible privilege of taking a peek into the universe at 13 TeV.

Acknowledgements

There are so many individuals I would like to thank for all the support I've received along the way, and there is no way I'll be able to do them all justice. I can only try to express the first layers of my gratitude.

Above all, I'd like to thank my advisor, David Stuart, who constantly challenged me with new and interesting projects and ensured I had a diverse foundation that has brought me to where I am today. Thank you especially for your empathy and awareness of our lives as graduate students, and for never adding criticism after we've already deeply felt our own mistakes. I'd also like to thank Jeff Richman and Claudio Campagnari, whose wisdom and guidance have been invaluable. All of you have made UCSB into a truly exceptional environment for becoming a scientist.

The search described in this thesis would not have been possible without the entire RA4-MJ team: postdocs Manuel Franco Sevilla, Jae Hyeok Yoo, Ana Ovcharova, Chris West; and fellow students Rohan Bhandari, Adam Dishaw and honorary member Jack Bradmiller-Feld. Each of you have taught me so much, from CMS basics, good search design, statistics and C++ magic, to delivering a great talk. I am incredibly lucky to be part of such an elite team. Alex and James, I am very happy to know you will be carrying the baton.

I couldn't have survived graduate school without all of the great friends in Santa Barbara who made each day worthwhile. Griffin, Grace, Rohan, Jack, Adam, Jon, Charlotte, Christina, Kevin, Chantal, Kevin, Clint and many, many others: you will always have a special place in my heart.

Last but certainly not least, I thank Ilaria for her constant support, patience and love. I would have never made it this far without the joy and richness you bring to my life. I can't wait to resume our lives together in one place!

Curriculum Vitæ

Ryan Edward Heller

Education

- 2018 Ph.D. in Physics (Expected), University of California, Santa Barbara.
- 2016 M.A. in Physics, University of California, Santa Barbara.
- 2012 B.A. in Physics and Mathematics, University of California, Berkeley.

Selected Publications

CMS Collaboration, “Search for R-parity-violating supersymmetry in proton-proton collisions at $\sqrt{s} = 13$ TeV in events with a single lepton and large jet and bottom quark jet multiplicity,” [arXiv:1712.08920](#). Submitted to Physics Letters B.

CMS Collaboration, “Search for supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV in the single-lepton final state using the sum of masses of large-radius jets,” Phys. Rev. Lett. 119, 151802 (2017)

CMS Collaboration, “Search for supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV in the single-lepton final state using the sum of masses of large-radius jets,” J. High Energy Phys. 08 (2016) 122

ATLAS Collaboration, “Search for microscopic black holes in a like-sign dimuon final state using large track multiplicity with the ATLAS detector,” Phys. Rev. D 88, 072001 (2013)

Abstract

A search for supersymmetry with the CMS detector in the single-lepton final state
using the sum of masses of large-radius jets

by

Ryan Edward Heller

The discovery of the Higgs boson casts new urgency to an old question: the Higgs mass hierarchy problem. Supersymmetry can provide an elegant and natural solution to the hierarchy problem, and would result in many new particles accessible at the TeV scale probed by the LHC.

This dissertation describes a search for a classic natural supersymmetry signature, gluino mediated top squark production. The search is performed in the single-lepton final state, relying on the large missing energy and high jet and b-flavor jet multiplicities to separate the gluino signal from the Standard Model backgrounds. The background measurement is centered around the variable MJ, the sum of masses of large radius jets. The search uses a sample of proton-proton collisions at center of mass energy 13 TeV recorded by the CMS detector, corresponding to an integrated luminosity of 35.9 fb^{-1} . The observed yields in the signal regions are consistent with the expected Standard Model backgrounds, and the results are interpreted in the context of simplified models of gluino pair production. Scenarios with gluino masses up to approximately 1.9 TeV are excluded at 95% confidence level for neutralino masses less than 1 TeV. This negative result joins a substantial body of evidence disfavoring supersymmetry realized near the electroweak scale.

Contents

Curriculum Vitae	vi
Abstract	vii
 Part I Introduction	 1
1 Theoretical context	2
1.1 Hierarchy problem	2
1.2 Supersymmetry	4
 Part II Experimental Apparatus & Event Reconstruction	 9
2 Large Hadron Collider	10
3 CMS detector	13
3.1 Solenoid	15
3.2 Tracker	16
3.3 Electromagnetic Calorimeter	25
3.4 Hadronic Calorimeter	33
3.5 Muon System	46
3.6 Trigger	55
 4 Event Reconstruction	 58
4.1 Tracking	58
4.2 Particle Flow	64
4.3 Jets	69
4.4 Leptons	76
4.5 Missing Energy	81
4.6 Summary of Definitions	83

Part III	The Search for Supersymmetry	84
5	Introduction	85
6	Data samples and simulation	89
6.1	Triggers and data sample	89
6.2	Simulated event samples	93
7	Event Selection	97
7.1	Baseline selection	97
7.2	Suppressing single-lepton backgrounds	101
7.3	Suppressing dilepton backgrounds	103
7.4	Background outlook	105
8	M_J	107
8.1	Large-Radius jets and M_J	107
8.2	Understanding high M_J backgrounds	109
8.3	M_J performance studies	113
9	Background Estimation	116
9.1	Method	116
9.2	Implementation	122
10	Systematic Uncertainties, Background Estimate	125
10.1	Impact of Monte Carlo variations on κ	126
10.2	Understanding sources of M_J - m_T correlation	127
10.3	Control samples for systematic uncertainties	131
10.4	Validation of systematics procedure	137
10.5	Background systematics summary	139
11	Systematic Uncertainties, Signal Model	140
12	Results	144
12.1	Results	144
12.2	Interpretations	152
13	Conclusions	156
	Bibliography	158

Part I

Introduction

Chapter 1

Theoretical context

1.1 Hierarchy problem




Figure 1.1: Distribution of the four-lepton invariant mass for the critical $H \rightarrow ZZ \rightarrow 4\ell$ discovery channel. Points with error bars represent the data, shaded histograms represent the backgrounds, and the unshaded histogram represents the signal expectation for a mass hypothesis of $m_H = 126$ GeV. From Ref. [1].

The discovery of the Higgs boson (Figure 1.1) is the greatest triumph of the LHC ex-

periments to date, and serves as a long-awaited confirmation of the theory of electroweak symmetry breaking, 40 years after its inception. The presence of the Higgs boson at the electroweak scale, though, casts a new urgency to an old problem. The Higgs boson, as an elementary scalar, receives mass contributions from every other particle through corrections to the propagator, like those shown in Figure 1.2.



Figure 1.2: One-loop corrections to the Higgs propagator for fermions (left) and scalars (right).

The contributions from fermions, for example, take the form

$$\Delta m_H^2 = \frac{|\lambda_f|^2}{8\pi^2} \left[-\Lambda_{UV}^2 + m_f^2 \ln(\Lambda_{UV}/m_f) + \dots \right]. \quad (1.1)$$

These contributions contain two problematic terms: a quadratic divergence in the cutoff scale Λ_{UV} , and even worse, sensitivity to the mass scales of all new physics (m_f in this case). Considering all of the known deficiencies of the Standard Model—inability to describe dark matter, dark energy, quantum gravity, and more—it seems irrefutable that there must new physics entering at higher energies between the electroweak and the Planck scale. Each new particle will come with contributions to the Higgs mass similar to 1.1, and asking these contributions to cancel over more than 10 orders of magnitude to leave a Higgs boson at 125 GeV by chance would require extremely lucky fine tuning. This unexplained fine tuning is known as the Higgs mass hierarchy problem.

1.2 Supersymmetry

One way to make sense of the hierarchy problem is to notice that contributions from scalar particles, like that in Figure 1.2 (right), enter with similar structure but opposite sign

$$\Delta m_H^2 = \frac{\lambda_s}{16\pi^2} \left[\Lambda_{UV}^2 - m_f^2 \ln(\Lambda_{UV}/m_s) + \dots \right]. \quad (1.2)$$

This motivates proposing a new symmetry, called a *supersymmetry* [2], that introduces partner particles (“sparticles”) for each particle, with spins offset by 1/2: a scalar for every chiral fermion (“sleptons” and “squarks”) and a gauge fermion for every gauge boson (“gauginos”). In unbroken supersymmetry, the partner particles share the same mass, and the Yukawa couplings are matched so that, for example, $\lambda_s = |\lambda_f|^2$ and

Under truly unbroken supersymmetry, all of the contributions of the form of Eq. 1.1 and Eq. 1.2 would neatly cancel, and new particles would always be introduced in pairs, cleanly resolving the hierarchy problem. However, unbroken supersymmetry would also imply the existence of another 17 particles at easily accessible mass scales—obviously excluded by all of our experimental understanding. To keep supersymmetry alive, we are thus required to consider some kind of broken supersymmetry, sacrificing the matching of either the Yukawa couplings or the partner masses. A common choice is to consider “soft” supersymmetry breaking, where supersymmetry is preserved at high energy scales (and so the matching of Yukawas must be maintained), but the Lagrangian contains low-energy supersymmetry violating terms that can introduce a mass splitting between partners of $O(m_{soft})$, allowing them to have evaded detection so far. In this case, the fermion and scalar contributions only partially cancel, and the residual contributions to

the Higgs mass are of the form:

$$\Delta m_H^2 = m_{\text{soft}}^2 \left[\frac{\lambda}{16\pi^2} \ln(\Lambda_{UV}/m_{\text{soft}}) + \dots \right]. \quad (1.3)$$

In this case, if m_{soft} is on the order of the TeV scale, then the residual fine tuning is much more palatable.

There are a few particles closely related to the Higgs boson that are more strongly constrained to small m_{soft} if supersymmetry is to solve the hierarchy problem with limited tuning. First, since Higgsinos share the Higgs mass parameter, μ , they should be of comparable mass to the Higgs. Second, as top quarks have the largest Yukawa couplings, the residual contribution from the stop-top mass splitting has the largest impact, so $m_{\tilde{t}}$ should be relatively small. Finally, though gluinos do not directly couple to the Higgs, they also make a large contribution at two loops through the top squark. Placing quantitative bounds based on these arguments is somewhat a matter of taste and ultimately boils down to how much fine-tuning can be stomached. However, generally accepted bounds for natural supersymmetry can be roughly summarized as the following:

$$m_{\tilde{H}} \approx O(100 \text{ GeV}), \quad (1.4)$$

$$m_{\tilde{t}} \lesssim 1 \text{ TeV}, \quad (1.5)$$

$$m_{\tilde{g}} \lesssim 2 \text{ TeV}, \quad (1.6)$$

$$(1.7)$$

though the bounds tend to drift with time to accommodate new experimental constraints. Other sparticles can be several TeV, effectively decoupled from LHC phenomenology, without causing too much tension with the Higgs mass. Much more thor-

ough discussion of naturalness can be found in Refs. [3, 4]. The limits from Run I of the LHC already bring some tension, excluding $m_{\tilde{t}} \lesssim 800$ GeV and $m_{\tilde{g}} \lesssim 1300$ GeV, and help set the stage for what to look for at the start of Run II.

Figure 1.3 shows the ratio of LHC parton luminosities at 13 TeV to those at 8 TeV as a function of the mass of pair produced particles. This essentially quantifies the increase in cross section gained from the energy upgrade, which tends to be much more substantial for more massive particles (though, it should be taken with a grain of salt—the cross section to produce a pair 3 TeV particles may have increased by a factor of 100, but it is still quite small.) For gluinos at the edge of Run I exclusion around $m_{\tilde{g}} = 1500$ GeV, the cross section increased by almost a factor of 10, while for stop at $m_{\tilde{t}} = 900$ GeV, the cross section increased by approximately a factor of 5. Meanwhile, the cross section for inclusive $t\bar{t}$ production increased by only a factor of 2. This implies that, especially for gluinos, even a small dataset of a few fb^{-1} can significantly extend the reach beyond the Run I bounds (and indeed, this was the case). At the start of Run II, the most interesting searches to perform are those targeting very high mass processes like the pair production of gluinos, which can breach new territory quickly.

In the case that all of the supersymmetric partners constrained by naturalness do exist at the scales indicated in Eq. 1.7, then one of the first processes to be discovered would certainly be gluino mediated stop production, shown in Figure 1.4. This process benefits from the large cross sections afforded by gluinos, as well as a incredibly distinctive final state containing four top quarks and two neutralinos. The top quarks result in a spectacular display of jets, b-flavor jets, and frequently, isolated leptons. The neutralinos, admixtures of Higgsinos and other neutral electroweakinos, are only weakly interacting. Under the typical assumptions, an additional conserved quantum number (R-parity) requires that the lightest SUSY particle must be stable—typically the lightest neutralino, $\tilde{\chi}_1^0$. As a result, the stable, weakly interacting neutralinos escape invisibly, resulting in a




Figure 1.3: Ratios of LHC parton luminosities for pair production of particles of mass M_X , for gg , $q\bar{q}$, and qg initial states. From Ref. [5].

powerful signature of momentum imbalance (“missing energy”).




Figure 1.4: Feynman diagrams for gluino mediated stop production

The subject of this dissertation is a search for gluino mediated stop production in the single lepton final state using the Compact Muon Solenoid (CMS). Part II describes the CMS detector and how it is used to reconstruct pp collisions delivered by the Large Hadron Collider. Part III describes the development of the search strategy, background

prediction, systematic uncertainties, and the final results of the search over the 2016 dataset corresponding to 35.9 fb^{-1} , the first large dataset taken at $\sqrt{s} = 13 \text{ TeV}$.

Part II

Experimental Apparatus & Event Reconstruction

Chapter 2

Large Hadron Collider

The Large Hadron Collider (LHC) holds many superlatives: it is the highest energy and highest luminosity collider ever built, the largest cryogenic facility at liquid helium temperature, and the largest single machine on Earth. It is located at CERN (le Conseil Européen pour la Recherche Nucléaire) on the border of France and Switzerland, inside the 27 km circumference tunnel that previously housed LEP, the Large Electron-Positron collider. It produces collisions at four different interaction regions, shown in Figure 2.1. The protons are kept traveling in the ring by 1232 dipole magnets that produce magnetic fields of up to 8.3 T and accelerated and maintained at high energy by 400 MHz RF cavities. Thousands more higher order magnets keep the beams collimated and focused. In total, the magnets use approximately 96 t of superfluid helium-4.

The LHC is a proton-proton collider. In general, it would be much more convenient and economical to collide particles with their own antiparticles, whose opposite charge allows sharing of a common magnetic channel for the two counter-rotating beams.[6] This has been a very successful strategy used by most recent colliders: the Super Proton Synchrotron and LEP at CERN, as well as the Tevatron at Fermilab. Unfortunately in this scheme, the production of anti-protons to feed the accelerator becomes the limiting




Figure 2.1: Artistic rendering of the LHC and the detectors at the four interaction points: CMS, ATLAS, LHCb and ALICE. Depth not to scale.

factor in the attainable luminosity, and so in the design of LHC, there was potential for a huge gain if a way could be found to accommodate counter-rotating same-sign proton beams. At the same time, there was a strong financial incentive to fit inside the pre-existing LEP tunnel, which has an internal diameter of just 3.7 m. Ultimately this was achieved with a twin-bore “two-in-one” magnet design, where the opposite-sign dipoles and both beams fit inside a single cold volume, economical with space as well as the cost of cryogenics.[7]

Though its place as the highest energy collider is often emphasized, it shouldn’t be neglected to mention that the twin-bore development has paved the way for the LHC to reach record-setting luminosity as well. During the current period, Run II (2015-2018), the LHC collides protons at $\sqrt{s} = 13 \text{ TeV}$ with a typical luminosity of $O(10000) \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$. This luminosity corresponds to the production of roughly eight top quark pairs per second and one Higgs boson every two seconds. For comparison, at the Teva-

tron during the time of the top quark discovery, the typical luminosity was significantly smaller—approximately $14 \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$ [8], though this is not a value to scoff at. The majority of the improvement can be understood by examining the difference in bunch crossing interval: just 25 ns at the LHC, contrasted with 3500 ns at the early Tevatron, $140\times$ less frequent. The remaining factor of 5 comes from packing more protons in a small bunch area, and results in a $5\times$ greater number pileup interactions in events at the LHC. By the end of its life, the Tevatron reached luminosities of $O(200) \times 10^{30} \text{ cm}^{-2} \text{ s}^{-1}$, thanks mostly to the addition of more bunches, reducing the bunch crossing interval to 396 ns. Each order of magnitude improvement in bunch crossing interval imposes stricter and stricter technological constraints on the detector time response and the DAQ bandwidth, and it likely wouldn't have been possible to take advantage of a high-luminosity same-sign collider until the LHC era. Ultimately, we are very fortunate to be in position both to produce and fully exploit such a remarkable rate of high energy collisions.

Chapter 3

CMS detector

The Compact Muon Solenoid (CMS) is one of two general-purpose high-luminosity particle detectors at the LHC, located at interaction point 5 in Cessy, France. It consists of several concentric layers, each layer functioning more or less independently for specialized purposes. The innermost layer, the tracker (Section 3.2), measures trajectories of charged particles in the magnetic field in order to infer their momentum and the position of the vertices where they originate. The second layer, the electromagnetic calorimeter (Section 3.3), stops electrons and photons to identify them and measure the energy they carry. Hadrons penetrate through to the third active layer, the hadronic calorimeter (Section 3.4), where their energy is measured and absorbed. Finally, the fourth layer is occupied by the muon detectors (Section 3.5), which extend the tracking system to large radius for muons, extremely useful and uniquely penetrating particles which easily pass through all of CMS.

A schematic of the CMS detector is shown in Figure 3.1. The detector and all of the subsystems take the shape of a cylindrical barrel surrounding the LHC beampipe, closed by endcaps on each end. The majority of the detector fits in a cylindrical volume approximately 20 m long with a 15 m diameter. The CMS detector is described in coordinates r ,




Figure 3.1: Cutaway schematic view of the CMS detector.

the distance from the beamline, z , the distance along the beamline from the interaction point, and ϕ , the azimuthal angle. Observed particle directions are typically described in coordinates ϕ and η where η is the pseudorapidity, a remapping of θ , the angle from the beamline, with $\eta \equiv -\ln[\tan(\theta/2)]$. The pseudorapidity η is a more intuitive coordinate for describing relativistic kinematics: particle production is approximately uniform in η (strong preference for low θ), and separations in η are approximately invariant under boosts along the z -axis (forward and central jets have the same width in η). CMS is a hermetic detector, with essentially no visible particles escaping detection out to $\eta = 5$ ($\theta = 0.8^\circ$).

CMS is distinguished from other particle detectors by the substantial investment into its namesake, the solenoid, which is the largest superconducting magnet ever built. This choice reflects an emphasis on tracking measurements, performed by an extremely

sophisticated tracker: the first entirely silicon tracker and the largest silicon detector ever built, with an active area of about 200 m^2 . As a result, momentum measurements from the inner tracker are dominant for most charged particles, except for very high momentum muons, electrons and charged hadrons. This was an aggressive strategy given the cutting-edge technical challenges and risks associated with operating silicon detectors in the harsh LHC environment, but ultimately the great success of the tracker has made this decision very profitable.

The following subsections describe each subdetector of CMS in detail. Much of the information comes from the master CMS reference [9], with supplementary references used as indicated.

3.1 Solenoid

The superconducting solenoid is a remarkable device that forms the central backbone of CMS. It is the largest superconducting magnet ever built, providing an approximately 3.8 T field over a cylindrical region with diameter 6.3 m and length 12.5 m, and stores an energy of 2.3 GJ. The solenoid is also exceptionally thin, to minimize the rate of multiple scattering of muons traversing the magnet, with a thickness of just 31.2 cm and a small mass relative to its size, only 220 t. The combination of the immense field and lightweight structure imply a high strain compared to magnets in other detectors, requiring interesting engineering innovations to accomodate. For instance, the NbTi conductor had to be extruded simultaneously inside an aluminum reinforcement to help bear the strain, as well as serving as the cold mass. A comparison of strain and stored energy for several high energy physics detector magnets is shown in Figure 3.2. The CMS magnet withstands a larger strain and stores a greater energy than all other detector magnets ever made.




Figure 3.2: Energy to mass ratio (measure of mechanical strain) vs stored energy for several detector magnets [9].

The ambitious design of the solenoid yields substantial rewards for the performance of the tracking detectors in CMS. Track curvature is linearly proportional to the field magnitude, and so improvements in the field translate directly into the momentum resolution, easing pressure on the tracker granularity. Additionally, the solenoid extends substantially farther along the z -axis than the tracker, ensuring an extremely uniform, powerful field in the entire relevant tracker volume and paving the way for the dominance of tracking measurements in CMS. Contrast this with ATLAS, whose solenoid produces half the field strength and is slightly shorter than their silicon tracker, resulting in a weak and non-uniform field near the ends that complicates the performance of the tracker away from the most central regions [10].

3.2 Tracker

The inner tracker is the innermost layer of CMS, designed to observe the trajectories of charged particles with minimal disturbance. These trajectories, or “tracks” can be

used both to measure particle momenta, based on their curvature in the magnetic field, as well as to infer the position of interaction vertices, where multiple tracks originate. The successful design of a tracking system requires making compromises among numerous stringent and competing constraints. Trackers should reach as close as possible to the beamline (within a few cm) to provide precise vertexing, but need to maintain sparse channel occupancy for unambiguous assignment of hits from nearby particles in busy environments; together these requirements impose channel segmentation of order 100 μm . Trackers should also extend to large radii from the interaction point, with as many intermediate layers as possible, to ensure accurate and sensitive curvature measurements even for very high momentum, straight tracks, but are constrained by the great cost (in CHF, material, and cooling budgets) of instrumenting larger and larger surfaces with millions of fine-granularity channels. Meanwhile, the tracker sensors and readout electronics need to survive in an extremely intense radiation environment, be read out on ns timescales, and kept at cold and stable temperatures, all while using minimal material to avoid affecting the particle trajectories through scattering or by initiating showers. The only detector technology that can feasibly satisfy all of these constraints, particularly those on granularity, timing, radiation hardness, and mass, are silicon sensors.

In the region with radii below 10 cm, the higher track density requires pixelated sensors to ensure occupancy less than 1%. Section 3.2.1 describes the CMS pixel detector occupying this region. At larger radii, the occupancy can be controlled by using sensors that are narrow in just one dimension, so-called “strips”. The CMS strip tracker, which extends the tracker out to about 1 m, is described in Section 3.2.2. A schematic diagram of the entire tracking system is shown in Figure 3.3.




Figure 3.3: Cross section of the CMS tracker in the r - z plane [11].

3.2.1 Pixel Detector

The pixel detector is the closest detector to the interaction point and is responsible for providing precise impact parameter measurements necessary for efficient vertexing, as well as providing seeds for track reconstruction. It consists of three barrel layers at $r = 4.4$, 7.3 and 10.2 cm, each extending for 53 cm along the z -axis, and two endcap disks at each end, at $z = \pm 34.5$ and ± 46.5 cm, from $r = 6$ to $r = 15$ cm. A schematic of the pixel detector can be seen in Figure 3.4.

Each individual pixel has dimensions of $100 \times 150 \mu\text{m}^2$ in r - ϕ and z , on wafers of thickness $285 \mu\text{m}$. The finite thickness leads to charge sharing among neighboring pixels, which is enhanced by Lorentz drift of charge carriers along the ϕ direction in the magnetic field. Additionally, the forward pixels are tilted by 20° from normal incidence in a turbine geometry to further induce charge-sharing (visible in Figures 3.4 and 3.5.) Since the pixels have analog pulse height readout (as opposed to boolean readout), the charge-sharing is, in fact, a huge asset and enables interpolation between neighboring channels to achieve spatial resolution of order 10 - $15 \mu\text{m}$, $10\times$ more precise than the pixel cell size![11].

In total, the pixel detector contains 66 million individual channels, with an active area of 1.06 m^2 . The coverage as a function of η can be seen in Figure 3.6. In the central region for $\eta < 2$, the efficiency for tracks to yield two or more pixel hits is close to 100%.




Figure 3.4: Schematic of barrel (green) and forward pixels (orange) [9].

3.2.2 Strip Tracker

The CMS silicon strips extend the tracker from $r = 20 \text{ cm}$ to $r = 116 \text{ cm}$. As the track impact parameters are already very precisely constrained by the pixel detector, the primary purpose of the strips are to carefully follow the track curvature in the magnetic field, with a longer lever arm and better sensitivity to small angular deviations from straight paths. Towards this end, the sensors have very fine segmentation in the $r - \phi$ direction, but are quite long along z (10+ cm), hence the name “strips”. The strip geometry dramatically reduces the number of channels necessary to instrument a given area, which leads to crucial savings in cost, material, power consumption (cooling capacity), and readout complexity, but is only feasible at larger radii where the track density and therefore the occupancy is smaller. This comes at the cost of sacrificing measurement




Figure 3.5: Two forward pixel half-disks, with sensors inclined at 20° from normal incidence (left). Half of the barrel pixel detector, surrounding the beampipe. Three horizontal barrel modules are visible immediately below the beampipe (right). [12]

of the z -coordinate (r in the endcap disks), which is partially restored by instrumenting some of the layers in “stereo”: back-to-back mounting of modules with a relative rotation of about six degrees, enabling measurement in the secondary coordinate with resolution of hundreds of microns.

The strip tracker is divided into three different subdetectors, visible in Figure 3.3. The tracker inner barrel and disks (TIB/TID) occupy the region out to $r < 55$ cm and $|z| < 118$ cm and consist of 4 barrel layers supplemented by 3 disks, providing up to 4 measurements in r - ϕ . This region ($|z| < 118$ cm) is surrounded by the tracker outer barrel (TOB), which adds another 6 layers out to $r = 116$ cm. Finally, the tracker endcaps (TEC) add 9 disks, out to ($|z| = 282$ cm), for $|\eta| < 2.5$. The first two layers (rings) of TIB/TOB (TID/TEC) as well as the 5th ring of TEC are instrumented in stereo. Together, all of the strip tracker subdetectors provide 9 or more hits for almost




Figure 3.6: Geometric coverage of the CMS pixel detector (top). Efficiency for ≥ 2 hits as a function of track η (bottom) [9].




Figure 3.7: Number of measurement points as a function of track pseudorapidity η . Solid circles show the total number of measurements, while open squares show the number of stereo 2D hits [9].




Figure 3.8: Photo of the CMS strip tracker [13].

all trajectories between $\eta = 0$ and $\eta = 2.5$. The number of strip hits as a function of η is shown in Figure 3.7.

Occupancy considerations have largely driven the design of the strip sensors. In the inner regions with high track density ($r < 55$ cm), the sensors have extremely narrow pitch, with dimensions approximately 10 cm x 80 μm . At radii from 55-110 cm, the track density is smaller and the dimensions can be increased to approximately 25 cm x 180 μm , which helps maintain a reasonable number of read-out channels despite the large surface area to cover. As the surface area of an individual sensor increases, its capacitance increases, which leads to a reduction in the signal to noise ratio (a given charge deposition results in a pulse with smaller voltage amplitude). To compensate, the wafer thickness increases from 320 μm to 500 μm in the outer layers of the tracker, and the larger path length through the sensor translates directly into larger signals. Thicker sensors are in general more susceptible to bulk radiation damage effects (e.g. growth in leakage current), but this can be tolerated at the larger radii where the fluence is smaller.

3.2.3 Tracker general aspects

The channel occupancy for the tracker is shown in Figure 3.9 and can be used to infer the limiting factors in the design of each tracking detector. The strip detectors have occupancy very close to the limit of 1% necessary for efficient track-finding, implying that the requirements for curvature resolution (momentum resolution) are less stringent and already achieved by accommodating the occupancy constraints. In fact, the position resolution of strip hits, approximately 10-40 μm (relying again on charge sharing interpolation as in the pixels), is typically much finer than the track position itself can be known, because of loss of information due to multiple scattering. In contrast, the pixel sensors, by virtue of segmentation in 2D have occupancy much smaller than necessary

to enable track-finding, and instead the pixel size is strongly constrained by the physics goal of achieving vertex resolution of order 100 μm .

The fantastic performance of the tracker in terms of resolution, readout speed, radiation tolerance and hit efficiency does not come for free. Unfortunately, the tracker material budget, shown in Figure 3.10, grew well beyond the original specifications, driven primarily by the needs for cooling, structural support, and cabling. The substantial amount of material, which reaches almost $2X_0$ near the endcap transition region, causes electrons to bremsstrahlung and many pions to start hadronic showers. In fact, 35% of electrons lose more than 70% of their energy in the tracker! These interactions make it extremely difficult to disentangle the original momenta correctly and place a limit on the final efficiency of the track reconstruction. Future upgrades to the tracker will reduce the material budget, utilizing less dense support structures and two-phase evaporative CO_2 cooling, which brings more cooling power with less mass.




Figure 3.9: Map of typical occupancy across the CMS tracker. Based on measurement at $\sqrt{s} = 7$ TeV with an average of 9 pp interactions per beam crossing. [14]




Figure 3.10: Material budget of the tracker, measured in radiation lengths (left) and interaction lengths (right), for various tracking subdetectors [11].

3.3 Electromagnetic Calorimeter

The next layers of CMS, beyond the tracker, are the calorimeters. The basic principle of calorimetry is to force particles to interact, produce a large cascade, and measure the energy deposition of all of the subsequent particles as they pass through an active material. The energy of the original particle can be inferred from the size of the signal observed, and its direction can be measured based on the position of the shower, depending on the detector granularity. Particles can be identified based on where they are stopped: electromagnetic showers initiated by electrons and photons have much shorter characteristic radiation lengths and are stopped in the electromagnetic calorimeter, described in this section. Hadronic showers can penetrate through much more material, and are only stopped in the hadronic calorimeter, described in Section 3.4.

The electromagnetic calorimeter (ECAL) aims to measure electrons and photons originating from the interaction point. At high energies electrons primarily interact in matter through bremsstrahlung radiation, producing high energy gamma rays. High energy

photons, in turn, produce pairs of electrons and positrons, and so incident photons or electrons can quickly multiply into an electromagnetic shower containing thousands of electrons, positrons, and photons. Electromagnetic showers are relatively well behaved with a predictable size and particle composition, as well as minimal losses of energy to invisible excitations. As result, electromagnetic calorimeters can achieve phenomenal resolution, at the 1% level or better, in stark contrast to measurements of hadronic showers which are more complex and inherently bear a loss of information [15]. For many purposes, such precise resolution is unnecessary, but there is one critical application where it is essential: discovery of the Higgs boson in the diphoton final state. The Higgs boson is unique among heavy resonances with an extremely narrow width of only 4 MeV (in contrast to GeV-scale widths for W and Z bosons and the top quark). This means the observed width of the Higgs resonance, and therefore its prominence above the random background, is entirely dominated by the experimental resolution. It is thus extremely profitable to push the electromagnetic energy resolution as far as possible, and this was a driving force in the design of the ECAL.

Critical for good resolution is minimizing the material in front of the ECAL in order to reduce energy losses in transit. This necessitates placing the ECAL inside the solenoid, and limiting the material used for tracker services. At the same time, the calorimeter must be many radiation lengths deep for good containment of high energy showers, since fluctuations in leakage out the back can seriously degrade the resolution. To achieve shower containment in the compact space inside the solenoid, extremely dense lead tungstate crystals (PbWO_4) were selected and serve as both the absorber and the active material [16]. PbWO_4 has a high density, around 8 g cm^{-3} , and therefore a short radiation length (0.89 cm) and small Moliere radius (2.2 cm), which drives the transverse development of the shower and enables fine granularity. Furthermore, PbWO_4 has other nice properties that make it suitable for use in the LHC: good transparency and radiation

hardness, and a fast scintillation time that ensures 80% of the light is emitted in 25 ns.

Crystals of length 23 cm (22 cm) corresponding to staggering 25.8 (24.7) radiation lengths are used in the ECAL barrel (endcap), rendering fluctuations in the rear leakage essentially irrelevant for the resolution (otherwise a nasty uncertainty, since it wouldn't shrink with increasing energy). The crystal faces in the barrel are $22 \times 22 \text{ mm}^2$, expanding to $26 \times 26 \text{ mm}^2$ at the rear to maintain a pseudo-projective geometry. In the endcaps, the crystals are slightly larger, $28.6 \times 28.6 \text{ mm}^2$ at the front face and $30 \times 30 \text{ mm}^2$ at the rear. When combining adjacent crystals out to a 3×3 or 5×5 grid (superclusters), the resolution degradation due to fluctuations in lateral leakage is about 2% at 1 GeV, and decreases with $\sqrt{E}$.

Ultimately, one of the limiting factors in the crystal resolution turns out to be longitudinal non-uniformity of light collection related to the truncated pyramid shape [17]. To compensate, one face of each barrel crystal is depolished, decreasing light collection for showers that develop early, ensuring more uniform light yield for early and late showers. Ultimately, fluctuations in longitudinal development contribute an uncertainty of just 0.3%, constant with energy.

Photos of the crystals can be seen in Figure 3.11.




Figure 3.11: Barrel crystal with depolished face and attached APD (left). Endcap crystal with VPT (right) [9].

The main disadvantage of lead tungstate is its poor light yield: only $O(100)$ photons are created per MeV of energy deposition, two orders of magnitude fewer than LYSO or NaI crystals. This is not a big problem at LHC energy scales, where energies are measured in GeV, and can be compensated for by using photodetectors with intrinsic gain, mounted directly on the crystals. In the barrel, each crystal is instrumented with a pair of $5 \times 5 \text{ mm}^2$ avalanche photodiodes (APDs), operated at gain 50. The APDs have very good quantum efficiency (75%) at the PbWO_4 scintillation wavelength, 430 nm. In the endcap, where the radiation tolerance requirements are more stringent and the magnetic field is non-uniform, vacuum phototriodes (VPT) are used instead, which are essentially PMTs with a single gain stage. VPTs have a third of the quantum efficiency and a fifth of the gain of the APDs, but compensate with a larger active area of approximately 280 mm^2 .

The barrel crystals are arranged with 360-fold granularity in ϕ and 85-fold granularity in η per half out to $|\eta| < 2.5$, for a total of 62100 channel, a volume of about 8 m^3 and a weight of 67.4 tons. The front faces of the barrel crystals are at a radius of 1.29 m from the beamline. In each endcap, which cover the range $1.479 < |\eta| < 3$, there are 7324 channels split into two half-endcaps called ‘Dee’s’. The front of the endcap crystals are located at $|z| = 315.4 \text{ cm}$. A diagram of the ECAL is shown in Figure 3.12.

Additionally, a preshower detector is installed in front of the ECAL endcap, whose purpose is to identify boosted π^0 decays that can merge to fake an isolated photon. The preshower detector consists of 2 layers of lead absorber, for a total of 3 radiation lengths, interspersed with 2 layers of silicon strip detectors along perpendicular axes, with a pitch of 1.9 mm. This provides position measurements with much better granularity than the PbWO_4 crystals for the 95% of photons that begin to shower in the preshower stage, allowing rejection of π^0 s as well as improved position resolution for electrons and photons.




Figure 3.12: Schematic of the CMS electromagnetic calorimeter [9].




Figure 3.13: Barrel electromagnetic calorimeter (left). Endcap electromagnetic calorimeter (right) [9].

3.3.1 Calibration

Operation of the electromagnetic calorimeter in the actual LHC environment presents additional challenge due to the extreme radiation dose. Irradiation of the crystals produce color centers that absorb light and reduce the transparency [18]. At the ECAL temperature (18 C), some of the damage begins annealing immediately, and the transparency reaches an equilibrium value depending on the instantaneous luminosity, which is constantly changing. A simulation of the transparency cycles over a series of modest 10-hour fills is shown in Figure 3.12, yielding variation of approximately 0.5% over a single fill, even for a relatively central crystal ($\eta = 0.92$).




Figure 3.14: Simulation of transparency evolution for a series of 10 hour fills with initial luminosity of $2 \times 10^{33} \text{ cm}^{-2}\text{s}^{-1}$ [9].

The variation in transparency with time could have been disastrous for achieving percent-level energy resolution. To mitigate this effect, tremendous effort was spent developing a laser-monitoring system that directly injects light to the front of each crystal, and is used to track the transparency channel-by-channel over time with 30-40 minute

granularity, using gaps between beam crossings. The laser measurements also captures any variation in light collection and gain due to small drifts in temperature. The impact of the time-dependent transparency correction factors can be seen in Figure 3.15, where they dramatically improve the observed Z mass width, particularly in the endcap where the dose is much higher.




Figure 3.15: Impact of various calibrations on the observed Z mass width for the barrel (left) and endcap (right) [19]. The intercalibrations include standard calibrations like imposing ϕ symmetry in isolated electron energy spectra and mass constraints from π^0 decays. The dark blue curve shows the effect of the time-dependent transparency monitoring. The decrease in width from the red curve is the key gain from the monitoring; the mean could be corrected trivially with a time-independent correction.

After the laser monitoring corrections, the largest remaining sources of energy uncertainty are the constant 0.3% from fluctuations in longitudinal development, a $3\%/\sqrt{E}$ stochastic term from fluctuations in photostatistics and lateral shower development, and a constant noise term of $O(100 \text{ MeV})$ due to digitization noise and fluctuations in dark current. Additionally, interactions in the material of the tracker make it difficult to improve the energy resolution beyond 1% [20, 21]. Ultimately, thanks to the excellent ECAL performance, the expected Higgs diphoton mass width reaches just above 1%, paving the way for very good discovery sensitivity. A simulation of the width for a 120 GeV Higgs is shown in Figure 3.16.




Figure 3.16: Simulation of Higgs resolution for a 120 GeV Higgs boson, yielding an expected resolution of 1.34 GeV [22].

3.4 Hadronic Calorimeter

The hadronic calorimeter (HCAL) measures energies of hadronic jets. It has a significantly larger kinematic acceptance than any other detector and is therefore a crucial catch-all that enables accurate measurement of the missing energy left behind by invisible particles such as neutrinos or dark matter.

Hadronic showers are inherently more complex and unpredictable than electromagnetic showers, and so their intrinsic resolution is naturally much worse. One example of the difficulties with hadronic showers is that 30% to 50% of energy deposited by hadrons ends up in nuclear breakup or excitation, which are invisible or only produce signals that are $O(\mu s)$ out of time [15, 23]. Fluctuations in this substantial invisible fraction lead to severe broadening of the resolution. Contrast this with electromagnetic showers, which ultimately deposit nearly all of their energy into ionization leading to visible signals. The electromagnetic-hadronic response asymmetry introduces an additional complication: since the production of neutral pions continually siphons energy from the hadronic component into electromagnetic showers, all hadronic showers have a large electromagnetic component, comprising 1/3 of the energy at low energy and rising logarithmically with energy. Because the response is so different between the two components (factor of 1.4 for the CMS HCAL), fluctuations in the electromagnetic fraction also dramatically broaden the resolution.

Despite the challenges associated with hadronic calorimetry, it's important (and very difficult!) to build a calorimeter that maintains resolution close to the intrinsic shower resolution. In CMS, since the ECAL already contains a couple nuclear interaction lengths and will initiate hadronic showers, there is a strong preference to also fit the HCAL inside the solenoid to avoid a large dead zone in a critical region of early shower development. The space constraints imposed by the solenoid, though, make it very difficult

to achieve enough stopping power to reliably contain high energy and late-developing jets. Ultimately, the barrel hadronic calorimeter simply isn't thick enough for optimal performance, though this is partially mitigated by an additional tail-catcher calorimeter outside the solenoid. The barrel and endcap calorimeters are described in 3.4.1, the forward calorimeter in 3.4.2, and the outer tail-catcher calorimeter in 3.4.3.

3.4.1 Barrel and Endcap Calorimeters (HBHE)

Occupying the barrel region out to $|\eta| < 1.3$ and from $r = 1.77\text{ m}$ to $r = 2.95\text{ m}$ is the HB detector, which consists of 16 alternating layers of brass absorber and scintillator tiles, approximately 50 mm and 3.8 mm thick, respectively. At $\eta = 0$, this corresponds to 5.82 interaction lengths, but increases as $1/\sin\theta$ to 10.6 interaction lengths at $|\eta| = 1.3$. The scintillator tiles, made of Kuraray SCSN81, are segmented hermetically into 72 divisions in ϕ , and 16 sectors in η per half barrel. The endcap calorimeters, HE, has a similar composition and extends to $|\eta| = 3$, with another 14 segments in η and the same 72-fold segmentation in ϕ . The HE tile granularity matches HB at low η , with $\delta\eta \times \delta\phi = 0.087 \times 0.087$, expanding to $\delta\eta \times \delta\phi = 0.17 \times 0.17$ at high η . A schematic of the HCAL geometry is shown in Figure 3.17.

The tiles in HB and HE are equipped with 0.94-mm diameter wavelength shifting fibers (Kuraray Y-11) that carry the scintillation light out of the scintillators. Eventually, fibers in the same $\eta - \phi$ position are spliced together and fed into a Hybrid Photodiode (HPD) [24]. HPDs are photodetectors capable of operating in a strong magnetic field, and consist of a photocathode 3.3mm away from a pixelated silicon photodiode. The cathode is held at a potential of -8 kV, and the accelerated photoelectrons impinging on the photodiode are amplified with a gain of 2000. The signals from the HPDs are digitized in specially developed Charge Integrator and Encoder chips (QIEs) that make use of a




Figure 3.17: Map of HBHE tile segmentation. Each color corresponds to a single readout depth, as different tile depths are optically ganged together in the same η - ϕ tower. Most of HB is collapsed to a single depth, while HE mostly preserves two or three depths [9].

clever modified floating point concept to more efficiently use each floating point range, enabling coverage of a larger dynamic range with fewer bits, and ensuring digitization makes a negligible contribution to the resolution [25].

Ultimately, the HBHE energy resolution is dominated by the stochastic sampling fluctuations as well as fluctuations in the shower electromagnetic and invisible fractions, and is approximately $120\%/\sqrt{E}$, plus an additional 5–10% constant term. Here, the impact of the solenoid constraints appears. To fit as much brass as possible in the tight space, the absorber to scintillator ratio is rather large: 50 mm thick absorber interspersed with 3.8 mm tiles. Contrast this to the ATLAS tile calorimeter, which uses 14 mm iron absorber tiles and 3 mm thick scintillators for a much large sampling fraction, and achieves a stochastic resolution term approximately half that of CMS (still somewhat large, due to the inescapable fluctuations in the shower invisible and electromagnetic components.) [10]

A recent measurement of the HB response and resolution is shown in Figure 3.18. Although its resolution is coarse at low energy, the HCAL is highly complementary with the tracker, which has excellent momentum resolution at low p_T but has difficulty with high p_T , straight tracks, where the HCAL resolution is more precise.




Figure 3.18: HB response to 40 GeV to 60 GeV isolated charged hadrons in 2016 pp collisions. HB energy is compared against track momentum, accounting for energy deposited in ECAL [26]. After calibrations, the total resolution for this energy range is approximately 20%.

3.4.2 Forward Calorimeter (HF)

A unique element of the hadronic calorimeter system are the forward calorimeters (HF), which provides sole coverage for CMS in the forward region from $|\eta| = 3$ to 5, critical for accurate accounting of the missing transverse energy. This region close to the beamline experiences extreme radiation doses, close to 100 Mrad/year [27], and so ensuring sufficient radiation tolerance was the main design constraint. Quartz was selected

as the active material for its radiation hardness. Charged particles produce Cherenkov radiation as they pass through the quartz. Each HF consists of cylindrical steel absorber of length 165 cm (about 10 interaction lengths), from $r = 12.5$ cm to $r = 130$ cm, interspersed with quartz fibers. The fibers have a diameter of 600 μm and are placed in a square matrix with 5 mm spacing. In total, HF contains over 1000 km of quartz fibers, and each absorber has a weight of 108 tons.

Quartz calorimetry comes with a number of substantial challenges. Above all, the energy loss for Cherenkov radiation is about 10000 times smaller than that due to ionization in quartz, so the number of photons is relatively small (about 200 photons per cm for $\beta \approx 1$) [28]. Additionally, since the light collection relies on total internal reflection in the fiber at large angles, only a few percent of the light can be transported to the photodetectors. Beyond the inherently low yield, only about 1% of the HF volume is occupied by quartz (compared to almost 10% active volume for HBHE). As consequence of all of this, it takes a roughly 4 GeV energy deposition to yield a single photoelectron, and so the contribution to the resolution from photostatistics is horrendous: for electrons, the HF energy resolution is roughly $200\%/\sqrt{E}$ (contrast with 3% for the equivalent term in ECAL) [29]. Even worse, the Cherenkov mechanism is about 5 times more sensitive to the electromagnetic component of showers than the hadronic component, and so even hadronic jets are primarily observed through their electromagnetic component. Fluctuations in the shower composition thus have an even more dramatic effect on the resolution than in most calorimeters, yielding poor pion energy resolution of approximately $280\%/\sqrt{E}$. Though this may seem a shockingly large value, there is some compensation from the fact that forward jets have large energy, and in particular very large E/E_T . As result, the transverse energy resolution, which is what matters most, is more precise than this value suggests.

Despite its somewhat coarse energy resolution, HF provides critical information as a

catch-all for electrons, photons and jets alike in a region where most detectors would not be able to survive at all. To distinguish between electromagnetic and hadronic-initiated showers, HF uses a clever trick: half of the fibers (“short” fibers) start 22 cm beyond the front face of the absorber. Electromagnetic showers are mostly contained before reaching the short fibers, so by comparing the light yield in the short and long fibers (which extend to the front of the absorber), electrons and photons can be distinguished from jets, crucial for correct energy assignment given the extreme e/h response ratio.

3.4.3 Outer Calorimeter (HO)

Since the HB in the most central region has a depth of only 5.82 interaction lengths, it was considered prudent to build an additional tail-catcher calorimeter outside the solenoid, to ensure good measurement for late-developing showers. The HO consists of an additional layer of scintillator tiles (BC408, 10 mm thick) mounted inside the iron magnet return yoke. The HO tiles are arranged in 5 rings along the z -axis, and roughly correspond to the η - ϕ segmentation of HB. The layout can be seen in Figure 3.19. Taking all of the material in ECAL, HCAL, the solenoid and return yoke, the minimum thickness is extended to 11.8 radiation lengths. Figure 3.20 shows the reduction in the under-measurement tail with the addition of HO.

Originally, HO was read out with the same HPD sensors as HBHE, but they were found to have difficulty operating in the moderate but non-uniform field near the magnet yoke, leading to spurious signals due to discharge across the 8 kV potential. During the LHC long shutdown 1, the HPDs were replaced with silicon photomultipliers (SiPMs) that can tolerate the field [30].




Figure 3.19: Layout of HO trays in 5 rings around the solenoid (left). Correspondence of HO tiles with HB, EB and tracker (right) [9].




Figure 3.20: Comparison of observed HCAL response to a 300 GeV π^- beam with and without HO, scaled so that 300 GeV is unity [31].

3.4.4 HBHE Phase I Upgrade

The HPDs in HBHE also faced several difficulties that would become problematic for maintaining performance long-term. They also experience spurious discharge signals, though less frequent than in HO because of the different field configuration. More troubling, the HPD responses were observed to drift independently in time, attributed to the cumulative effect on the cathodes of repeated discharges [32]. The drift can be seen in Figure 3.21 over a period of two years, without any indication of stabilization. Finally, the scintillator tiles experienced darkening due to radiation damage faster than expected, and with significantly more damage accruing at shallower depths [33]. Any darkening is problematic for photostatistics, but growing differences in depth response introduces longitudinal non-linearity that contributes to the resolution as well. The radiation damage observed for early and later depths are shown in Figure 3.22.




Figure 3.21: Variation in HBHE response as a function of time through Run I, based on LED illumination. [32]

To mitigate these effects, the HBHE HPDs are in the process of being replaced with




Figure 3.22: Response degradation of HE tiles as a function of integrated luminosity, for depth 1 (left) and depth 7 (right), at various values of η . Measurements taken from 2012, 2015 and 2016. The luminosity delivered in 2012 was at $\sqrt{s} = 8$ TeV, and is scaled down by a factor of 1.25 to account for the smaller irradiation per dose relative to 13 TeV collisions [33].

SiPMs. SiPMs operate at lower bias voltage and do not have the stability issues associated with HPDs. SiPMs have better quantum efficiency than HPDs by a factor of three [34], which directly compensates for the scintillator darkening. Additionally, they have much higher gain (3×10^5) and are very compact (few mm) per channel, which together enable a key capability: finer segmentation in the longitudinal readout. In the original HBHE design, all depths in a given η - ϕ position have to be ganged together optically before reaching the HPDs, both to ensure large signals relative to the electronic noise and to reduce the number of HPDs packed in a compact space. The improved SiPM efficiency and gain ensure good signal to noise without combining so many channels, enabling the longitudinal information to be preserved. In the longitudinally segmented scheme, the differential radiation damage can be compensated for simply by weighting signals from darkened layers more heavily. The new segmentation is shown in Figure 3.23.

One wedge of the HE was upgraded in time for 2017 data-taking in order to gain experience installing, commissioning and operating the next generation detector. A key




Figure 3.23: Map of proposed phase 1 HBHE tile segmentation [34].

issue to resolve is whether projections for radiation damage to the SiPMs are accurate in the LHC environment. As SiPMs are irradiated, bulk damage from neutrons increases the rate for spontaneous micro-pixel avalanches, resulting in a growing dark current. If the dark current grows large enough, its fluctuations could contribute to the energy resolution, or worse, the SiPM could heat significantly and enter thermal runaway.

I performed an analysis of the SiPM dark current growth throughout 2017, with the goal of validating the understanding from test beam irradiations and making predictions for the SiPM performance over the lifetime of HB and HE.

The SiPM pedestal currents were measured periodically over 2017 to monitor the growth of the spontaneous avalanche rate. Since each spontaneous micro-pixel avalanche produces the same response as a single photoelectron (SPE), the radiation damage manifests as a simple increase in the dark photoelectron rate. The pedestal current in a single SiPM in the first upgraded HE is shown in Figure 3.24 after various integrated doses. The quantized photoelectron spectrum is clearly visible, and the increased dose

translates into a shift in the mean number of photoelectrons (N_{PE}) observed in a given time window. The single photoelectron gain remains constant with dose as evidenced by the unchanged spacing between consecutive N_{PE} peaks.

The mean SiPM dark current in the upgraded wedge over 2017 is shown in Figure 3.25, and reaches 0.8–1.2 μA after 50 fb^{-1} . This current corresponds to an increase in the mean N_{PE} from near 0 to approximately 1, in a 50 ns window relevant for HCAL pulse reconstruction. Since a single photon corresponds to approximately 30 MeV, fluctuations in the dark current after 50 fb^{-1} contribute noise comparable to 30 MeV. Over the full lifetime of the upgraded HB, which will last for 3000 fb^{-1} , linear extrapolation of this growth predicts an average of 60 dark photoelectrons per 50 ns, which would contribute gaussian fluctuations of approximately 250 MeV (8 photon RMS). Continued darkening of the scintillators will necessitate reducing the photoelectron to energy conversion ratio and give the dark current fluctuations more weight, but it seems unlikely that the dark current noise will ever reach 1 GeV, which would still be significantly less than the full resolution at all energies. Additionally, the maximum dark current at 3000 fb^{-1} will be approximately $60\text{ }\mu\text{A}$, comfortably less than the thermal runaway current ($200\text{ }\mu\text{A}$). These projections make a conservative assumption that there is no annealing; in practice, only approximately half of the SiPM damage will last beyond a few months.




Figure 3.24: Photoelectron charge spectrum from a single SiPM after various exposures in situ in CMS. This SiPM was connected to the scintillator tile at $i\text{Eta} = 29$, $i\text{Phi} = 63$, and depth = 1. The position of the $N_{\text{PE}} = 0$ peak is determined by the QIE pedestal current.




Figure 3.25: Progression of SiPM dark current over 2017. Non-SiPM pedestal currents are subtracted, and the SiPM currents are averaged over all channels of the same size. Growth of the dark current is roughly linear, but due to time-dependent annealing and the complex time structure of LHC operation, there are many breaks of linearity and shifts of slope. Significant recovery is visible over approximately 1 month at the end of the year without appreciable LHC activity. The most central HE channels, where the longitudinal segmentation remains coarser, use larger 3.3 mm SiPMs to accommodate the larger signals from ganged tiles. The larger SiPMs effectively absorb a larger dose at a given luminosity and have more micropixels activated, resulting in a steeper growth in current proportional to the SiPM area [35].

3.5 Muon System

The outermost regions of CMS are occupied by the muon system. Muons are a very powerful handle at the LHC—distinctive and easily measured, and often the signature of an interesting process worth recording. Thanks to their large mass, 200x that of the electron, high energy muons are much less sensitive to deflection by electrons in material they traverse and have very small radiative losses. As result, they can travel much farther through matter (approximately 1 m of rock per GeV!) and penetrate the CMS calorimeters and solenoid with ease. Dedicated detectors outside the solenoid can then easily identify muons and carefully measure their trajectories, which are still relatively faithful representations of their initial momenta.

Muons hold significant branching ratios for the final states of nearly all interesting Standard Model particles observable at the LHC: Higgs, W, and Z bosons, top quarks and b-flavor hadrons, as well as many extensions to the standard model such as higher mass gauge bosons (Z') or supersymmetry (including strongly produced SUSY, the subject of this thesis, but even more important as a signature of decays of electroweak superpartners). For decades, muons have been an essential tool in every single discovery of a new massive elementary particle. For instance, the decay $H \rightarrow ZZ^* \rightarrow 4\mu$ was considered the “golden channel” for Higgs discovery, with a low background and the excellent mass resolution afforded by muons.

The key requirements for the muon system are to enable efficient triggering and identification of muons. Efficient triggering requires reasonably good momentum resolution for low p_T muons, but offline, the muon system can only be expected to compete with the tracker resolution for very high p_T muons, for which the extra lever arm is necessary and the loss of information of multiple scattering is less significant.

Since the muon detectors are at such large radius, they need to cover a daunting

surface area to achieve hermiticity, and counting all layers the muon system reaches an area of about $25\,000\text{ m}^2$. Covering such a large surface necessitates using relatively cheap, reliable technology, but the larger radius allows use of coarser segmentation—measured in mm, not μm , keeping costs lower without reaching dangerous occupancy or sacrificing momentum resolution. The ideal technology for these constraints are gas ionization chambers. In gas chambers, muons ionize gas molecules, and high electric fields cause avalanche amplification of electrons that drift close to the anodes, creating large signals even for small ionization. Typically, the electron drift time enables precise inference of the original location of ionization. There are three muon systems in CMS: drift tubes (DTs) 3.5.1, which occupy the barrel; cathode strip chambers (CSCs) that cover the endcap region; and resistive plate chambers (RPCs), which give redundancy in both the barrel and endcap regions.

3.5.1 Drift tubes

The basic drift cell comprising the drift tubes is shown in Figure 3.26. It is an approximately rectangular tube, $42\text{ mm} \times 13\text{ mm} \times 2.4\text{ m}$ with a $50\text{ }\mu\text{m}$ anode wire suspended in the center. The tube is filled with a mixture of Ar and CO_2 gas, and in the presence of the approximately 2 kV cm^{-1} field, the electron drift velocity towards the anode is $60\text{ }\mu\text{m ns}^{-1}$. Electrons ionized by the passage of a muon can thus take up to 350 ns to reach the anode, and the delay identifies the distance of closest approach very precisely [36]. The spatial uniformity of the field along the path of the muon important to ensure reliable position measurements. This is achieved with a special arrangement of additional electrodes between the anode and the cathode. The equipotential lines inside a drift cell are shown in Figure 3.27, and are quite parallel for the majority of the z-axis range, ensuring electrons produced anywhere along the z-axis will have uniform drift time. In the

end, the single cell resolution is about $250\text{ }\mu\text{m}$. Since the drift tube cells don't measure the coordinate parallel to the wire, each drift tube station contains several layers of cells, with some tubes orientated for r - ϕ (typically 8 layers) and some for z (typically 4 layers). Additionally, for a single drift tube cell, there is a left-right ambiguity. To resolve this, adjacent layers are staggered by half the cell width. Including all 12 layers in a station, the position resolution per station reaches $100\text{ }\mu\text{m}$. There are four layers of drift tube stations placed in the magnet return yoke, and the layout is shown in Figure 3.28.

The drift tubes are an elegant, simple technology that are well suited for the barrel, where the occupancy and the magnetic field are low. In the endcaps, the particle rate is high enough that the 350 ns drift time is too long to ensure low occupancy, and the non-uniform field would be problematic. In this region, it is necessary to use a faster detector that is able to cope with the non-uniform fields: cathode strip chambers.




Figure 3.26: Diagram of drift tube cell [9].

3.5.2 Cathode strip chambers

To handle the significantly higher endcap occupancy, it's necessary to use ionization chambers with much finer wire pitch (and therefore faster signals) and the ability to




Figure 3.27: Diagram of equipotential lines in one half of a drift cell (anode wire on the right, cathode on the left [9].)




Figure 3.28: Diagram of the muon barrel system, with 4 layers of muon detectors (blue, “MB”) interleaved with the magnet yoke (“YB”) [9].

measure both coordinates of every avalanche. The cathode strip chambers achieve this using anode wires with 3.2 mm pitch to measure r , wrapped around cathode strips that measure ϕ . The strips and wires are arranged in a wedge geometry shown in Figure 3.29. The ionization signals from muons form avalanches near the anode wires just like in the drift tubes, but additionally the avalanche forms an image charge in the cathode strips, and interpolation among the analog strip signals yields a precise measurement of the ϕ coordinate. Each CSC chamber contains 7 layers cathode strip panels, with wires wrapped around alternating layers, for a total of 6 wire layers. The resolution combining all layers in each chamber reaches approximately 70–150 μm in the ϕ direction, and the signals are resolved in 25 ns with an ultimate time resolution of 4 ns (for a drift speed of 60 $\mu\text{m ns}^{-1}$, this indicates a z resolution of a couple hundred μm .) Forward muons pass through 3 or 4 CSCs; the endcap layout is shown in Figure 3.30.

3.5.3 Resistive plate chambers

Finally, there is an additional muon system that provides redundancy for both the DTs and the CSCs, particularly for the trigger: the resistive plate chambers (RPCs). Due to the long timescale for signal propagation in DTs (350 ns), there was concern DTs wouldn't be able to unambiguously associate muons with the correct bunch crossing to form an efficient trigger.

The RPCs are designed explicitly for excellent time resolution. In normal gas chambers like the DTs and CSCs, avalanche amplification doesn't begin until the electrons finally drift into the small high-field region in the immediate vicinity of the anode wires. The RPCs instead use *planar* anodes and cathodes which ensure high fields throughout the volume, and avalanche amplification begins immediately wherever the ionization occurs, drastically reducing the time for the signal to develop. Ultimately, the RPCs reach




Figure 3.29: Schematic of CSC wedge, with cathode strips of constant width in ϕ , and a few of the anode wires shown [9].




Figure 3.30: Diagram of the endcap muon detector placement (ME 1–4) [9].

time resolution of $O(1 \text{ ns})$ [37].

A key innovation that made the planar geometry feasible is the use of resistive electrodes (bakelite) that are depleted of charge much faster than they can recover, allowing the avalanche to self-extinguish and reset without an external switching circuit or significant dead time.

In CMS, the RPC stations consist of two 2 mm gaps, each under a potential of 9 kV. There are 6 layers of RPC stations in the barrel augmenting the DTs, and 4 layers of RPCs in the endcaps. The layout is shown in Figure 3.31.

Ultimately, the timing performance of the drift tubes was adequate and the RPCs were not critical to ensure efficient triggering. Still, it is extremely valuable to have redundancy for such an important CMS subsystem.




Figure 3.31: Diagram of the RPC placement (magenta) around the DTs (green) and CSCs (blue) [36].

3.6 Trigger

There is an astounding variation in cross section among different processes in proton-proton collisions, as visible in Figure 3.32. The most interesting events, containing of W or Z bosons, top quarks, or Higgs bosons, are 5 to 10 orders of magnitude more rare than the total proton-proton inelastic cross section (dominated by soft QCD processes), and many anticipated beyond the Standard Model processes are even rarer. Thus obtaining a substantial dataset from proton collisions necessitates the extremely aggressive collision scheme provided by the LHC: colliding bunches every 25 ns, with an intensity such that approximately 50 pp interactions are expected every time! This collision rate is much faster than data acquisition could be sustained, but the vast majority of these events are simply not worth the tape they would require to be stored. The solution is to use a trigger that can quickly filter out and select only the interesting events for storage and full reconstruction offline. CMS achieves this with a two-level design: the Level-1 Trigger (L1), which makes a primitive selection to reduce the rate from 40 MHz to about 100 kHz; and the High-Level Trigger (HLT), which further reduces the rate to several hundred Hz. [38]

The Level-1 Trigger has to analyze every single bunch crossing, at 40 MHz, and distribute a trigger decision within 4 μ s. Achieving this rate and latency performance required a substantial effort to develop a reconstruction pipeline in hardware, entirely parallel to the offline datastream. The L1 trigger is built from programmable FPGAs where possible, and ASICs where flexibility has to be sacrificed for speed. The calorimeter front-end electronics provide *trigger primitives*, essentially per-tower energy sums, which are passed forward to Regional and Global Calorimeter Triggers, which produce a list of electrons/photons, jets, and energy sums. Meanwhile, the muon detectors produce local muon tracks, which are passed to the Global Muon Trigger and compared with




Figure 3.32: Theoretical cross sections for various processes in proton-(anti)proton collisions as a function of center of mass energy. Near $\sqrt{s} = 2$ TeV, cross sections are for proton-antiproton collisions relevant for the Tevatron. Above the discontinuity at $\sqrt{s} = 4$ TeV, the predictions are for proton-proton collisions [5].

output from the Regional Calorimeter Trigger to incorporate isolation information. In the current state, CMS can't perform tracking fast enough to include the inner tracker hits at L1, and as result, electrons and photons can't be distinguished and all low p_T objects are measured with poor resolution. After the Phase II upgrade anticipated for the High Luminosity LHC, though, the strip tracker will contain double-layer modules capable of finding track "stubs" locally, enabling drastically simpler tracking fast enough to use in the L1 trigger. For now, the objects produced by the Global Muon and Global Calorimeter system and passed to the Global Trigger, which compares the list against a menu of up to 128 conditions to decide whether to pass the event onto the next stage. The conditions are typically simple requirements on the number of objects (jets, leptons, etc) passing certain p_T or isolation criteria, or requirements on total or missing energy sums.

If an event satisfies the requires for an L1 trigger, the detector is read out to the HLT, a farm of thousands of CPUs. The HLT is able to perform an analysis close to that used in the full offline reconstruction, but with a number of simplifications for improved speed. For example, considerable care is taken to order the reconstruction in a way to enable rejection of objects as early as possible, before performing computationally intensive tasks. For instance, electron candidates first complete the calorimeter reconstruction, before looking for corresponding pixel hits, before finally performing tracking with the strip information, with significant rejection at every stage [39]. If an event satisfies the requirements for one of the configurable high level triggers, it is saved to mass storage and reconstructed offline.

Chapter 4

Event Reconstruction

This chapter describes how the information from the CMS subdetectors is used to reconstruct events. This includes the reconstruction of important physics objects like jets, electrons, muons, and the missing energy, which are all based on a reconstruction technique known as Particle Flow. All variables used in the analysis are developed along the way, and summarized in Section 4.6.

4.1 Tracking

4.1.1 Track Reconstruction

Developing an efficient and reliable track reconstruction procedure has required a substantial concerted effort by many members of the CMS collaboration. The resulting tracking procedure and its performance is described in great detail in Ref. [11]. This section summarizes general aspects of the track reconstruction.

Track reconstruction is performed by software called the Combinatorial Track Finder (CTF), based on the Kalman filter. The CTF algorithm is applied iteratively. On the first pass, clean, easily measured tracks with high p_T and small impact parameters are

found. Their hits are removed from the collection, and subsequent passes target more difficult tracks—softer, displaced, or with suboptimal seeds. This process continues for approximately six iterations, though the majority of the tracks are found in the first pass.

The precursor to the CTF procedure is to generate the collection of hits, a process known as *local reconstruction*. This involves taking the zero-suppressed signals from the pixels and strips, and grouping adjacent channels in order to form hit clusters. As mentioned in Section 3.2, the analog readout allows the hit position within clusters to be precisely inferred based on the detailed charge-sharing between channels. Pixel hits have a precision of $10\,\mu\text{m}$ in $r\phi$, and $20\text{--}40\,\mu\text{m}$ in z , depending on the angle of incidence—while the Lorentz drift ensures charge sharing in $r\phi$ for all hits, charge sharing in z is only generated by oblique incidence. Strip hit resolution is approximately $10\text{--}40\,\mu\text{m}$, but is typically more precise than the track trajectory can be known, anyway, due to material effects. The hit reconstruction efficiency for good modules in both detectors is greater than 99%.

After the local reconstruction of hits, the first step of CTF is the generation of seeds, the starting points for track forming. Although it wasn't immediately obvious, it turns out to be very advantageous to seed tracking from the innermost hits, in the pixel detector, for a number of reasons. First, despite the extreme track density close to the interaction point, the occupancy of the pixels is actually much lower than the strip occupancy thanks to their fine 2D segmentation (refer to Figure 3.9), and so pixel tracking is inherently cleaner and easier. Additionally, the pixel hits are 3D, providing an additional constraint compared to 2D strip hits (except for stereo layers). Finally, the most compelling reason is that pions are prone to inelastic nuclear interactions in the tracker material, and many do not reach the outer layers intact. Figure 4.1 shows the survival probability after each tracker layer. It is clear that seeding tracks from the outside in would unnecessarily bound the tracking efficiency to less than 90%. For the most pristine

early iterations of the CTF, seeds consist of triplets of pixel hits with an estimated $p_T > 0.8 \text{ GeV}$ and impact parameter consistent with the beam spot. In later iterations, seeds are allowed with looser kinematic constraints, or with some or all of the hits coming from the strips rather than pixels, enabling reconstruction of displaced tracks formed even outside of the pixel volume.




Figure 4.1: Pion survival probability as a function of number of tracker layers traversed [11].

The next step, *track finding*, expands the seeds into full-fledged tracks based on a Kalman filter. The seed trajectories are extrapolated from layer to layer, adding new hits consistent with the track and updating the trajectory at each stage. The extrapolation is broken into many incremental steps, but for computational considerations, most of them use analytical helical paths (uniform magnetic field approximation) adjusted to account for energy losses in the material of each layer. As the reconstruction proceeds outwards, multiple track candidates can be formed from a single seed, and the five with the best

χ^2 are kept at each layer. The process is continued until a track reaches the end of the tracker volume, or until it meets rejection criteria based on the number of missing hits or falling beneath a p_T threshold. Additionally for HLT tracking, tracks can exit the reconstruction early to save time if they reach certain precision thresholds. At the end of the track-finding process, several redundant track candidates may have been found for a single particle. The track candidate collection is then cleaned of duplicate tracks that share hits, keeping the tracks with the best χ^2 .

With the complete set of hits defined for each track candidate, the track parameters are then carefully re-evaluated in a process called *track fitting*. The track fitting stage uses a trajectory propagator that takes into account both the material effects as well as small field non-uniformities. This is a significantly more intensive computation, since the non-uniformities mean the extrapolation must be performed numerically, not analytically, and is done using the standard fourth-order Runge-Kutta method. The fitting procedure significantly reduces bias due to field non-uniformities and improves the track p_T and impact parameter resolutions.

At the end of each iteration of the track finding and fitting procedure, there are a significant number of fake tracks. To ensure high track purity, several selection criteria are applied, particularly minimum requirements on the number of layers with hits and the number of layers with 3D hits, and a limit on how many layers can be traversed with missing hits. Additionally, requirements are imposed on the track χ^2/dof and the impact parameters with respect to both the primary vertex and the beam spot. The thresholds for these cuts are designed to grow more stringent for tracks with hits in fewer layers, which are much more likely to be fake, and become essentially irrelevant for tracks with hits in 10 or more layers.

Figure 4.2 shows the tracking efficiency for muons originating from Z bosons using a tag and probe sample in 2016 data. In the absence of pileup interactions, the efficiency is

above 99%, but falls to 95% by around $N_{PV} = 30$. Pileup interactions can introduce many low- p_T particles that spiral around the tracker volume creating many hits and increasing the track-finding complexity.

Isolated muons are an idealized tracking case. For pions, which are often produced inside dense jets and undergo secondary nuclear interactions before escaping the tracker, the efficiency is degraded and ranges from approximately 85 to 95% in simulated $t\bar{t}$ events. It is degraded at forward η and high p_T , since high p_T tracks are likely to be part of collimated jets with many nearby tracks. Additionally, in realistic jet environments up to 5-10% of tracks can be “fake”, not associated with a primary particle, largely due to confusion from secondary particles produced as pions interact with the tracker material.

Parametrizing the resolution for track impact parameter and p_T is complex, but typical values are approximately 10-40 μm for tracks with $p_T > 10\text{ GeV}$, degrading to 100 μm at 1 GeV; and 1-10% of the p_T , steeply increasing at forward η and above 50-100 GeV.

4.1.2 Primary Vertex Reconstruction

Primary vertex reconstruction attempts to reconstruct the positions of all of the primary pp interactions in the beam spot luminous region, which is approximately 3.5 cm x 10 μm x 10 μm . The track reconstruction described above is based on a vertex collection generated by a simple, fast reconstruction done solely using the pixel detector. This method simply takes all pixel triplet tracks and clusters them into vertices, only separating nearby vertices if they are separated by more than 2 mm.

Then, after the track collection is formed, a more robust primary vertex is performed to improve the ability to resolve nearby vertices. This is done by selecting well-measured tracks consistent with originating in the primary interaction region, and clustering tracks




Figure 4.2: Tracking efficiency in data and simulation for muons from Z boson decays, as a function of η (left) and the number of primary vertices (right). Z bosons decays are identified based on a combined tracker and muon chamber muon (tag) paired with a muon chamber standalone muon (probe), and then a search is performed for tracker-only tracks matching the probe [40].

together based on their z -coordinate. The clustering is performed using a sophisticated method called *deterministic annealing* [41], which is based on analogy with thermodynamics. Tracks are assigned to multiple vertices at the initial stage, and an analogue to the free energy is defined with the track-vertex compatibility χ^2 serving as the energy. As the “temperature” is lowered, tracks are penalized more heavily for their distance from vertices and eventually a finite set of vertices survives to low temperature. The primary vertex reconstruction efficiency based on this procedure is above 99%.

The primary vertex resolution is shown in Figure 4.3, as measured in 2015 and 2016 data. In the 2016 dataset relevant for this analysis, the high instantaneous luminosity resulted in “dynamic inefficiency”, where pixel hits could be lost due to occupancy from previous events. This caused an approximately 10% degradation of the primary vertex resolution. Note that this degradation is mostly independent of the number of vertices in

a single event, depending instead on the *average* number of vertices per bunch crossing.

Of all of the primary interaction vertices found in an event, there is almost always a single vertex representing the interesting hard scatter process. This vertex is referred to as *the* primary vertex in the context of the analysis, and is chosen by ranking the vertices based on the sum p_T^2 of tracks associated with each vertex.




Figure 4.3: Primary vertex resolution in the z -coordinate measured in data from 2015 and 2016, as a function of sum of track p_T associated with the vertex (left) and the number of primary vertices in the event (right). Resolution in the transverse direction is slightly better, due to the smaller pixel size in the transverse dimension [40].

4.2 Particle Flow

With the collection of tracks in hand, the next step is to reconstruct the calorimetric information and proceed to the measurement of jets. It is worth pausing for a moment to consider the best way to measure hadronic energy. Figure 4.4 shows representative energy/momentum resolutions for charged pions from both calorimetry and tracking, as a function of energy/ p_T . As mentioned in Section 3.4, the two techniques are highly

complementary: calorimetric measurement has very coarse resolution at low energy, but the resolution steadily improves with energy. On the other hand, the tracker provides extremely precise measurements of low p_T tracks, but the resolution is degraded for high p_T particles with limited curvature. The ideal reconstruction paradigm then is to leverage both the tracking and calorimetry measurements where they each perform best: measure soft to medium charged pion momenta cleanly using the tracker, before they initiate showers, and only rely on the calorimeter when tracking becomes imprecise and the showers are large enough to be protected against fluctuations. This is one of the core ideas behind Particle Flow [42] [43], the reconstruction strategy employed by CMS: identify and reconstruct particles individually, rather than as aggregate jets, and rely on the subdetector best suited for each particular case. As we will see, this philosophy also provides a number of convenient features for calibration and physics analysis. CMS is the first hadron collider experiment to successfully implement particle flow reconstruction, which was originally developed for the ALEPH detector at LEP.




Figure 4.4: Left: HCAL energy resolution for isolated charged pions as a function of pion energy [44]. Right: Tracker p_T resolution for charged particles (majority pions) in simulated $t\bar{t}$ events [11].

The basic procedure for particle flow reconstruction is to “link” tracks and calorimeter deposits based on spatial coincidence and extrapolation of track trajectories, and then use these “blocks” of linked objects to identify particle candidates and assign energy/momentum. A prerequisite then is the ability to reliably and uniquely identify tracks with calorimeter energy deposits that come from distinct particles. CMS was not designed with this explicit goal in mind, but it turned out to be quite well suited to the task. Key enablers of the linking are efficient track-finding, provided by the CMS tracker, and fine calorimeter granularity to enable resolving energy deposits left by nearby particles. The fine granularity of the ECAL with crystals of $\delta\eta \times \delta\phi = 0.017 \times 0.017$, ensures very good ability to separate adjacent particles—frequently even photons from high p_T π^0 decays can be resolved. The HCAL granularity, on the other hand, is $5\times$ coarser, which makes it much more difficult to distinguish nearby deposits. However, since hadronic showers are intrinsically much broader, there are limits to the resolving power that could be obtained with finer granularity, anyway. What turns out to save the day is the strong magnetic field provided by the solenoid: charged and neutral particles from a single jet are separated by the time they reach the calorimeters. Additionally, the fact that both calorimeters are inside the solenoid and fairly close to the tracker simplifies the propagation of tracks, making linking more reliable.

Energy deposits in the calorimeters are reconstructed with a dedicated clustering algorithm to facilitate particle flow, with a focus on separation of partially-merged showers from nearby particles. The basic idea is to find “seed” cells, defined as cells that are more prominent than their neighbors, and then aggregate all contiguous cells above a zero-suppression threshold to form “topological clusters” around the seeds. If there are two nearby seeds, their clusters can merge into a single topological cluster with multiple seeds. Then for each topological cluster, a fit is performed to N Gaussian energy deposits, where N is the number of seeds. Each Gaussian has free parameters for amplitude

(energy) and position, but a fixed width. This procedure is essentially just finding a reasonable way to divide the energy among N deposits (“clusters”) and estimate their positions.

Tracks, ECAL clusters and HCAL clusters can then be linked into blocks, more or less based on grouping the closest objects together in η and ϕ . Then there is an ordered procedure to deduce the constituent particles in a block. First, as the cleanest case, muon candidates are identified and their tracks are removed from the block. Then, electrons are identified, as tracks associated with ECAL clusters, but with limited HCAL energy. This stage requires special care: electrons frequently bremsstrahlung photons (35% of electrons lose more than 70% of their energy in the tracker!), so a specialized tracking algorithm is necessary, and particle flow attempts to correctly handle energy from bremsstrahlung photons without double-counting. This necessitates also reconstructing isolated photons from track-less ECAL deposits during the same step. Then, the remaining particles are typically hadrons arising from jet fragmentation: charged hadrons ($\pi^\pm$, p , $K^\pm$), photons (from π^0), and long-lived neutral hadrons (K_L^0 or neutrons). The assignment of hadrons is fairly intuitive based on comparison of the linked track momenta and cluster energies:

- If $p \approx E$ (considering uncertainties), assign one charged hadron per track, using the weighted average of both measurements.
- If $p < E$, assign a charged hadron as well as photons and neutral hadrons to account for the remaining energy ($E-p$). Preference is given to assigning as much of the remaining energy as possible to photons (all of the ECAL energy), which are much more common than neutral long-lived hadrons.
- If $p > E$, this can indicate a problem requiring special treatment. Often, this can be attributed to fake high p_T tracks, or to muons that failed identification and survived beyond the first step (to suppress this case, particle flow uses a significantly looser

muon definition than typical muon reconstruction).

This is of course a simplification of the full procedure, and there are many caveats to accommodate special and difficult cases. Luckily, the relatively simple, clean structure of CMS helps minimize the number of special cases, making particle flow feasible in the first place.

A comparison of the jet energy resolution in particle flow and calorimeter-only jets is shown in Figure 4.5. At low and moderate energy, particle flow vastly improves the resolution by relying on the tracker information. For very high energy jets, the resolution converges to the calorimeter resolution.




Figure 4.5: Jet energy resolution compared for particle flow jets and calorimeter-only jets, as a function of p_T , for the barrel (left) and endcap (right). Based on simulation at $\sqrt{s} = 13$ TeV [42].

Note the critical role played by the calorimeter granularity. If each deposit can't be associated to a single particle, then it becomes nearly useless to replace the coarse calorimetry measurements with those from the tracker. For instance, suppose a charged pion deposit is merged with a neutral energy deposit, for a total energy E_{tot} (second case

in the list above). The charged pion can still be reconstructed with excellent precision based on its track momentum, p . Due to the merging, though, the neutral energy now suffers from the full impact of fluctuations on E_{tot} , not on $E_{\text{tot}} - p$, and we have gained almost nothing on the total jet resolution! The 1-to-1 track and deposit correspondence is critical, then, in order to divide and conquer the shower fluctuations.

All future CMS upgrades are now designed with particle flow techniques in mind, particularly focusing on granularity and simplicity of linking. The High Granularity Calorimeter (HGCAL), which will replace both endcap calorimeters, has very fine granularity of $O(1\text{ cm})$, and is essentially a dedicated particle flow calorimeter. Ultimately, it turns out to be much more valuable to focus investment into granularity rather than standalone energy resolution, enabling more advantageous use of the precise tracker measurements and sidestepping shower fluctuations whenever possible.

4.3 Jets

4.3.1 Jet Clustering

From the particle flow event description we can cluster together particle candidates to form jets, and estimate the momenta of particles that initiate showers. Starting with the clustering with the particle flow candidates bring a number of advantages compared to simply clustering calorimeter deposits. The particle candidates represent substantially more accurate momenta measurements, combining all detector information and including response corrections based on particle type. Additionally, charged particles (and most importantly their resolved energy deposits) can be traced back to their original orientations exiting from the hard scatter, rather than relying on their final position on the face of the calorimeter after traversing the magnetic field.

CMS uses the anti- k_T algorithm to cluster jets [45]. The algorithm essentially starts each jet from the highest p_T particles, and sequentially includes nearby soft particles out to a characteristic radius parameter, R (in CMS, $R = 0.4$.) More specifically, for each particle candidate i and j , we define distance measures:

$$d_{ij} = \min\left(\frac{1}{p_{Ti}^2}, \frac{1}{p_{Tj}^2}\right) \frac{\Delta_{ij}^2}{R^2} \quad (4.1)$$

$$d_{iB} = \frac{1}{p_{Ti}^2} \quad (4.2)$$

where Δ_{ij} is the angular separation between particles i and j in the η - ϕ plane. The algorithm proceeds sequentially, at each iteration merging (summing 4-momenta) the two particles with the smallest separation d . This is often the closest neighbor to the hardest particle. When d_{iB} is the smallest measure (it has already eaten its neighbors out to radius R), particle i is removed from the combination list and added to the list of jets.

The jets produced by the anti- k_T algorithm possess nice experimental and theoretical properties. Since they are clustered top-down from the hardest particles, their momenta and shapes are stable against the addition of extra, soft radiation, for example from pileup interactions (“infrared safe”). The sequential combination ensures they are also “colinear safe”—if a high p_T hadron fragments into two nearby particles, it typically won’t affect the resulting jet, since the two fragments will be merged early on. Historically, this has been an issue at hadron colliders where jets were formed by placing fixed cones around the hardest particles, all in one step—in this case, an arbitrary division of the hardest particles completely changes the outcome. Additionally, the dependable, circular shapes from anti- k_T are very amenable to energy calibrations and corrections for pileup.

4.3.2 Jet Energy Corrections

In general, jets have nonlinear response in the calorimeters, and are susceptible to energy from pileup interactions. To neutralize these effects, a series of “jet energy corrections” are applied. Since the rate of pileup interactions varies substantially with time, correcting for pileup energy contributions is a natural first step to remove the time dependence that could otherwise plague the response corrections. Typically in hadron colliders, pileup corrections are performed by subtracting an average energy density from each jet, assuming a perfectly diffuse energy background with no ability to detect localized fluctuations. Particle flow, however, enables an extremely powerful technique to compensate for non-uniformities, known as *charged hadron subtraction*. Charged hadrons that can be unambiguously associated with vertices other than the primary vertex are simply removed from jets. This typically removes 50% of the charged hadrons arising from pileup [46], and substantially reduces the burden placed on the area-based energy subtraction, which is still necessary for neutral pileup and the charged hadrons with ambiguous vertex association. The area correction is based on an event-by-event measurement of the diffuse energy deposition, convolved with a parametrized template of the η dependence of the pileup density. The pileup corrections very effectively remove the variation in the mean jet response, as visible in Figure 4.6 in the transition from the left to center frame.

With the pileup variation removed, a series of corrections are applied to remove nonlinearities and calibrate the jet response. First, the jet response is measured in Monte Carlo simulation, which enables precise control over the calibration sample and is actually found to agree well with the data. Jet energy corrections are applied based on the simulated response, in bins of p_T and η . Then, to complete the calibration, additional “residual” corrections are applied based on data control samples. The first step of the residual calibration, the relative correction, is based on imposing energy balance in dijet

events and seeks to achieve a uniform energy scale as a function of η . The second step, the absolute correction, uses a sample of jets recoiling against well-measured isolated photons and leptonically decaying Z bosons to calibrate the absolute scale, as a function of p_T .




Figure 4.6: Average jet response ($\text{reco } p_T / \text{true } p_T$) as a function of true p_T in bins of mean number of pileup interactions, μ . Left: before any corrections. Center: after pileup corrections. Right: after all corrections. [47]

Finally, for the context of this analysis, a set of quality requirements are imposed to eliminate fake jets that can arise from calorimeter noise or beam halo interactions. Jets must contain at least one charged particle, and at least two constituents in total. Additionally, the neutral hadronic fraction, neutral electromagnetic fraction, and charged electromagnetic fraction must each be less than 99% of the total jet energy, and the charged hadronic component must be non-zero. For all purposes, only jets with $p_T > 30 \text{ GeV}$ and $|\eta| < 2.4$ are considered.

Two event variables are defined directly from the collection of jets passing this selection: N_{jets} , the number of jets; and H_T , the scalar sum of all jet p_T .

4.3.3 Large-radius Jets

This analysis is based around the clustering of large-radius jets (“fatjets”) that contain showers resulting from multiple hard scatter partons. These jets are also clustered using the anti- k_T algorithm, but with a radius parameter $R = 1.4$. In order to capitalize on the official energy and pileup calibrations, the constituents entering the clustering are not the raw particle flow candidates, but instead the already-calibrated standard CMS $R = 0.4$ jets passing our selection (referred to as “skinny jets”). To capture a more complete event description, isolated leptons (described in Section 4.4) are also included as fatjet constituents.

The large-radius jets are used to define the quantity M_J , the scalar sum of the large-radius jet masses:

$$M_J = \sum_{\text{Large-R jets}} m_J \quad (4.3)$$

where the m_J are the masses of the large-radius jets (no selection is applied on large-radius jets in the summation). The motivation for M_J and its use in the analysis is discussed in detail in Part III.

4.3.4 b-flavor Jets

Heavy flavor jets are a interesting handle at the LHC, always present in top decays and in many beyond the Standard Model signatures as well. The ability to identify b-flavor jets is a powerful and important tool, in particular for our gluino analysis, but more generally for understanding any process involving top quarks.

b-flavor hadrons are distinctive for a number of reasons. They have a long lifetime— $c\tau \approx 450 \mu\text{m}$, so with a boost of $O(10)\text{GeV}$ they can propagate for several mm before decaying, which is significantly farther than the vertexing precision of $O(10 \mu\text{m})$. Their

cascade decays through charm hadrons also means b-flavor jets tend to have more tracks than light flavor jets.

Strategies for identifying (“tagging”) b-flavor jets exploit two related signatures of the long lifetime—the presence of secondary vertices, and displaced tracks with significant impact parameters. Figure 4.7 shows a schematic of a b-flavor jet. Track impact parameters are defined by propagating tracks backwards to the point of closest approach to the primary vertex, and are considered positive if the track appears to have been produced downstream along the jet axis from the primary vertex. If there are enough displaced tracks, then a secondary vertex can be reconstructed, which is an even more powerful indicator of a displaced decay, since it requires kinematic consistency of all of the displaced tracks.

Most analyses in CMS contemporary with this one used a multivariate b-tagging algorithm called the Combined Secondary Vertex tagger v2 (CSVv2)[49][50], which tries to make best use of the information in secondary vertices and displaced tracks. It first attempts to use secondary vertex information, particularly the significance of the vertex transverse displacement, the number of tracks associated with the vertex, and the vertex mass. Discrimination between charm and b-flavor secondary vertices is achieved by looking at impact parameter significance of the *least* significant track that raises the vertex mass above the charm threshold, 1.5 GeV (tracks are ordered by impact parameter significance and the mass is computed after adding each track in order.) This particular impact parameter distinguishes cases where a track with fake displacement overlaps with a charm vertex, pushing the mass into b-hadron range.

In the absence of a genuine secondary vertex (typically because not enough tracks are found), the CSVv2 looks for “pseudo-vertices” formed from displaced tracks that aren’t found in strict vertex reconstruction. Failing that, jets without a secondary vertex are still considered based purely on the number of tracks in the jet and their impact




Figure 4.7: Schematic of displaced b-flavor jet decay, indicating the secondary vertex and the impact parameter definition. Positive impact parameters correspond to an angle of less than 90° between the dark green impact parameter vector and the pink jet axis. [48]

parameters.

The efficiency of the CSVv2 algorithm is shown in Figure 4.8 for a variety of working points. This analysis uses the medium working point, corresponding to light flavor fake rate of approximately 1%. Since the analysis was performed, an improved algorithm was developed based on applying deep learning techniques to the same inputs as the CSVv2. The efficiency for the improved algorithm at the same fake rates is also shown in Figure 4.8.

The number of b-tagged jets in an event is one of the primary discriminating variables in this analysis, and is denoted N_b .




Figure 4.8: b-tagging efficiency as a function of jet p_T for a variety of algorithms and working points. Loose, medium and tight correspond to light-flavor fake rates of 10%, 1% and 0.1%, respectively [51].

4.4 Leptons

Leptons identified by particle flow undergo only a rudimentary selection designed to maximize efficiency. As leptons are a distinctive and valuable signature in our gluino

model and our analysis relies on accurate reconstruction of lepton kinematics, we apply additional selection criteria to ensure good purity. The specific criteria are listed in Tables 4.1 and 4.2

4.4.1 Isolation

The most distinctive and interesting leptons produced in pp collisions are those originating in the decay of massive W and Z bosons, arising commonly in our gluino model due to the presence of four top quarks in the final state. Since these leptons are not produced inside jets and receive substantial boost from the energy released in the decay, they are typically isolated, without other radiation nearby. The extremely short lifetimes of the vector bosons also ensure the leptons emerge directly from the primary vertex—they are called “prompt”. To distinguish prompt leptons from non-prompt leptons originating from heavy flavor decays inside jets, requirements are placed on the displacement from

Table 4.1: Muon selection criteria. The parameters d_{xy} and d_z are the transverse and longitudinal impact parameters relative to the primary vertex. $I_{\text{mini}}^{\text{rel}}$, the mini-isolation, is defined in Section 4.4.1. The last three rows are alternatives if the segment compatibility requirement fails.

Property	Requirement
p_T	$> 20 \text{ GeV}$
$ \eta $	< 2.4
d_{xy}	$< 2 \text{ mm}$
d_z	$< 5 \text{ mm}$
$I_{\text{mini}}^{\text{rel}}$	< 0.2
Is PF muon	Yes
Is standalone muon	No
Fraction of tracker hits found	> 0.8
Segment compatibility	> 0.303 , OR the following:
Global track χ^2/N_{dof}	< 3
Tracker-standalone match χ^2/N_{dof}	< 12
Track kink χ^2/N_{dof}	< 20

the primary vertex, and the “isolation”. In general, isolation is typically an energy sum in a cone of radius 0.2 - 0.4 around the lepton, and is usually required to be less than several GeV or a certain fraction of the lepton p_T . In many CMS analyses, a slightly more sophisticated version is used, called “mini-isolation”. Mini-isolation takes advantage of the fact that as jets are more and more boosted, their decay products become more collimated. Rejection of leptons from high p_T heavy flavor jets can thus be achieved with narrower cones, avoiding the risk of needlessly including distant radiation from pileup interactions or other jets. This motivates basing the isolation cone size on the expected opening angle of a two-body particle decay, $\Delta R \approx 2M/p_T$, for b-flavor mesons of mass approximately 5 GeV:

Table 4.2: Electron selection criteria for barrel and endcap electrons (boundary at $\eta_{\text{supercluster}} = 1.479$.) The parameters d_{xy} and d_z are the transverse and longitudinal impact parameters relative to the primary vertex. $I_{\text{mini}}^{\text{rel}}$, the mini-isolation, is defined in Section 4.4.1.

Property	Barrel requirement	Endcap requirement
p_T [GeV]	> 20	> 20
$ \eta_{\text{supercluster}} $	≤ 1.479	$1.479 - 2.5$
d_{xy} [mm]	< 0.118	< 0.739
d_z [mm]	< 3.73	< 6.02
$I_{\text{mini}}^{\text{rel}}$	< 0.1	< 0.1
$\sigma_{i\eta i\eta}$	< 0.101	< 0.0283
$\Delta\eta(\text{supercluster, track})$	< 0.0103	< 0.00733
$\Delta\phi(\text{supercluster, track})$	< 0.0336	< 0.114
$E_{\text{had}}/E_{\text{EM}}$	< 0.0876	< 0.0678
$1/E - 1/p$ [GeV $^{-1}$]	< 0.0174	< 0.0898
Missing tracker hits	≤ 2	≤ 1
Flagged as photon conversion	No	No

$$R^{\text{mini-iso}} = \begin{cases} 0.2, & p_T^\ell \leq 50 \text{ GeV} \\ \frac{10 \text{ GeV}}{p_T^\ell}, & p_T^\ell \in (50 \text{ GeV}, 200 \text{ GeV}) \\ 0.05, & p_T^\ell \geq 200 \text{ GeV}. \end{cases} \quad (4.4)$$

The cone is truncated at 0.05 to ensure good calorimeter granularity, and prevented from growing larger than 0.2 to avoid including neighboring independent jets. Since lepton p_T is highly correlated with the energy of nearby radiation, it is most meaningful to place requirements on the relative mini-isolation, $I_{\text{mini}}^{\text{rel}}$, the surrounding energy divided by the lepton p_T . For muons, $I_{\text{mini}}^{\text{rel}}$ is required to be less than 20% and for electrons, less than 10%.

The combined efficiency for the reconstruction, identification and isolation requirements for leptons as a function of p_T and η is shown in Figure 4.9. The efficiency for electrons is generally significantly worse than for muons, because they are both more fragile and less distinctive—they can be lost or mistaken for photons due to bremsstrahlung, and they can be more easily faked, necessitating stricter reconstruction requirements which comes with an efficiency cost.

4.4.2 Veto Tracks

Since identification of all leptons from W boson decays in $t\bar{t}$ is important for the gluino analysis, we additionally employ a looser lepton selection, defining objects known as *veto tracks*. This lepton definition focuses on efficiency, not purity, and imposes no further identification criteria beyond the criteria used by the particle flow algorithm. The p_T thresholds are also lowered to 10 GeV to expand the kinematic acceptance. Additionally, to target single-prong hadronic decays of τ leptons, a category of isolated




Figure 4.9: Combined efficiency for the reconstruction, identification and isolation requirements for electrons (left) and muons (right), based on leptons in kinematic acceptance from W boson decays in simulated $t\bar{t}$ events [51].

hadronic tracks is defined as well. The hadronic τ isolation calculation considers only charged tracks, not neutral energy, to include all single-prong hadronic decays, many of which also include neutral pions. The track definitions are shown in Table 4.3. The full use of veto tracks, as well as a detailed analysis of leptons that fail identification, is described in Part III.

Table 4.3: Selection requirements for veto tracks.

Property	Leptonic tracks	Hadronic tracks
p_T [GeV]	> 10	> 15
$ \eta $	< 2.5	< 2.5
$I_{\text{mini}}^{\text{rel}}$	< 0.2	< 0.1
d_{xy} [mm]	< 0.5	< 0.5
d_z [mm]	< 0.7	< 0.7

4.5 Missing Energy

Though they cannot be observed directly, the presence of invisible particles such as neutrinos, or dark matter, can be inferred based on a concept known as missing energy. Since the partons initiating the collisions have essentially zero net momentum in the transverse plane, the *vector* sum of the transverse momenta of all of the products of the collision should be zero. In practice, this means we can construct the vector sum of all visible particles in the event and take it to be an estimate of the total transverse momentum carried by invisible particles (modulo resolution effects). We define the missing transverse momentum:

$$\vec{p}_T^{\text{miss}} = - \sum \vec{p}_T \quad (4.5)$$

where the summation is performed over all particle flow jets, including their calibrations, as well as all other unclustered particle flow candidates. Colloquially, the missing transverse momentum is (arguably incorrectly) referred to as the missing energy, E_T^{miss} . This is because historically at hadron colliders, in the absence of particle flow techniques, this quantity would be measured mainly with calorimeter energy deposits.

Note that at a hadron collider, the momentum of the center of mass frame along the z -axis is unknown, and so it is impossible to determine the z -component of the missing energy.

4.5.1 Transverse Mass

Most events with a single isolated lepton and significant E_T^{miss} result from the decay of a W boson into a charged lepton and a neutrino. To identify these events, it is extremely

useful to define a quantity known as the transverse mass, m_T :

$$m_T(\ell, \vec{p}_T^{\text{miss}}) = \sqrt{2p_T^\ell E_T^{\text{miss}} [1 - \cos(\Delta\phi_{\ell, \vec{p}_T^{\text{miss}}})]} \quad (4.6)$$

If the E_T^{miss} in an event arises solely from a neutrino produced in association with the charged lepton, m_T is bounded above by the mass of the parent particle. This can be understood by the fact that the invariant mass of the lepton-neutrino pair must be the parent mass, and the transverse mass is essentially the portion of the invariant mass visible in the transverse plane. When energy released by the decay results in momenta of the decay products along the z -axis, m_T is reduced to less than the full invariant mass.

For events with two leptonically decaying W bosons, such as dileptonic $t\bar{t}$ decays, m_T is no longer bounded by the W mass. If both leptons are reconstructed, though, it is possible to construct a generalization of the transverse mass called the “stransverse mass” that restores the cutoff at m_W [52]:

$$m_{T2}(\ell_a, \ell_b, \vec{p}_T^{\text{miss}}) = \min_{\vec{p}_1 + \vec{p}_2 = \vec{p}_T^{\text{miss}}} \left[\max \{m_T(\vec{p}_{\ell_a}, \vec{p}_1), m_T(\vec{p}_{\ell_b}, \vec{p}_2)\} \right] \quad (4.7)$$

where the ℓ_a and ℓ_b are the two charged leptons, and the minimization is taken over all ways to partition $\vec{p}_T^{\text{miss}}$ into two hypothetical neutrino vectors, $\vec{p}_1$ and $\vec{p}_2$. For the true partitioning, the lepton momenta are paired with their neutrino partners and both m_T expressions are bound by m_W , and therefore the pairwise maximum is also bounded by m_W . Since of course the true partitioning is in the set of all partitions, the minimum of all maxima is bound by m_W as well. An interesting special case is when $\vec{p}_T^{\text{miss}}$ points between the directions of the two charged leptons. Then it is possible to partition the full E_T^{miss} with both $\vec{p}_1$ and $\vec{p}_2$ parallel to the leptons. In this case, both m_T vanish and therefore m_{T2} also vanishes.

4.6 Summary of Definitions

Table 4.4: Summary of variables used in the gluino analysis.

Variable	Definition
N_{lep}	Number of electrons and muons
N_{veto}	Number of veto tracks
N_{jets}	Number of $R = 0.4$ jets, $p_{\text{T}} > 30 \text{ GeV}$
N_{b}	Number of b-tagged jets
$E_{\text{T}}^{\text{miss}}$	Missing energy
H_{T}	Sum of p_{T} of all $R = 0.4$ jets
S_{T}	H_{T} + sum of all lepton p_{T}
M_{J}	Scalar sum of large-radius jet masses
m_{T}	Transverse mass of primary lepton and $E_{\text{T}}^{\text{miss}}$

Part III

The Search for Supersymmetry

Chapter 5

Introduction

As described in Part I, gluino mediated stop production is one of the most important signatures of supersymmetry, if supersymmetry is the solution to the Higgs mass hierarchy problem. This process involves all of the sparticles that are constrained by naturalness to the TeV scale, and requires no others. If these particles do exist, this process would certainly be the first in which they are revealed, thanks to the large cross section of gluino production and the striking final state. Since the Run I exclusions already probed such large gluino masses, few searches stood to gain as much as this one from the increase in LHC center of mass energy to $\sqrt{s} = 13$ TeV. Together, all of these factors made this analysis one of the most promising and exciting searches to perform at the start of Run II.

This search considers both classic gluino mediated stop production shown in Figure 5.1 (left), as well as a simplified model with the cascade decay replaced by a three-body decay of the gluino shown in Figure 5.1 (right) [53, 54, 55, 56]. The direct three-body decay arises for the case of off-shell stop quarks, and is a simpler benchmark model that captures most of the kinematic features in the on-shell model but with a much more approachable parameter space. We use the simplified three-body model (named “T1tttt”

in the SMS framework) as a guide for the design of the analysis, and ultimately also interpret the results for some interesting realizations of the on-shell cascade model (named “T5tttt”). We make a concerted effort to maintain sensitivity to both the traditional “non-compressed” spectra, where the gluino is much more massive than the neutralino, as well as “compressed” models, where the neutralino is closer in mass to the gluino and as result all of the final state objects carry significantly less momentum.




Figure 5.1: Feynman diagrams for gluino mediated stop production (left), and a simplified model with a direct 3-body decay of $\tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$ (right).

Due to the presence of four top quarks in the final state, these processes are incredibly distinctive, with often nine or more jets, four b-flavor jets, and frequently isolated, prompt leptons from the decays of W bosons. Additionally, the invisible neutralinos ($\tilde{\chi}_1^0$) can result in a substantial amount of missing energy. Each of these features alone is quite rare in the Standard Model, and it’s only possible for background events to possess all of them in truly exceptional cases.

The distribution of the number of isolated leptons is shown in Figure 5.2. Given the branching ratio of W bosons into electrons or muons (about 25% including 1-pronged leptonic cascades through τ), the single lepton final state holds the largest branching ratio. Even after accounting for acceptance and reconstruction efficiency, the single

lepton final state still contains a substantial fraction of all events, reaching almost 40%. In general at the LHC, Standard Model processes involving isolated leptons are much rarer than purely hadronic processes. Particularly for the case of E_T^{miss} signatures, leptonic backgrounds are much easier to understand—dominated by small cross section processes with real missing energy carried by neutrinos, without the threat of uncontrolled QCD tails where so-called “fake” E_T^{miss} can arise due to jet mismeasurement. By restricting the analysis to the single lepton final state, we maintain a large branching ratio with a somewhat smaller and simpler background to measure.




Figure 5.2: Distribution of number of leptons in the $\tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$ T1tttt model. The true number of leptons resulting from W boson decays is shown in red, and the number of reconstructed isolated leptons in kinematic acceptance is shown in blue.

One of them main challenges in advancing the search for gluinos stems from the fact that, due to their spectacular signatures, all models with substantial cross sections

(low masses) have already been excluded. The surviving models, which are still quite compelling theoretically, have cross sections on the order of 10 fb or less and live in extreme kinematic regions relative to Standard Model processes. The challenge is not finding a way to tease out a subtle signal over a large background, as gluinos come with plenty of powerful handles. Rather, the challenge is in robustly predicting the Standard Model contributions to the most extreme kinematics tails, well beyond the reach of calculations or reliable simulation. This search is explicitly designed to reduce to a single dominant background component (dileptonic $t\bar{t}$) amenable to a robust data-driven prediction, with extensive validation in data, purely data-driven background systematic uncertainties, and in general only minimal reliance on Monte Carlo simulation. We believe this approach has resulted in an analysis with the potential for a genuinely credible discovery of this critical supersymmetric signature.

Searches for supersymmetry in the single-lepton final state at the LHC have been performed previously by ATLAS and CMS at $\sqrt{s} = 7, 8,$ and 13 TeV [57, 58, 59, 60, 61, 62, 63, 64]. The strategy employed by this search was developed for the first data at $\sqrt{s} = 13$ TeV in 2015 with an integrated luminosity of 2.3fb^{-1} , published in [64]. This dissertation presents the results from an expanded analysis on the 35.9fb^{-1} dataset collected in 2016, the first 13 TeV dataset with substantial discovery potential. This analysis was published in [65], and the discussion in this dissertation draws heavily on internal CMS documentation in [66] and [67].

Chapter 6

Data samples and simulation

6.1 Triggers and data sample

Fortunately, the distinctive final state of our gluino model gives it a number of useful handles for triggering. Traditionally, isolated leptons are the key trigger object for this kind of search, which enable reasonable trigger rates even for relatively low p_T thresholds. For this analysis, specialized hybrid triggers were developed which combine a lepton requirement with a moderate threshold on the hadronic energy, H_T . The leptons are specifically designed to have only the loosest possible isolation requirements (“very very very loose”, or VVVL) to ensure good efficiency and compatibility with the offline variable-cone mini-isolation (defined in Section 4.4). The combination allows significantly lower thresholds than the pure single-component triggers: 15 GeV for leptons, compared to 25-30 GeV for pure lepton triggers (with too-strict isolation criteria); and 400 GeV for H_T , compared to O(1) TeV for pure H_T triggers. Though the non-compressed gluino models are energetic enough for the pure H_T trigger, the reduced threshold cross-triggers significantly extend trigger capability for compressed models, without compromising the ability to use leptons as soft as 20 GeV. Perhaps even more important, the cross-triggers are

based on H_T seeds at L1, saving the otherwise expensive efficiency cost of lepton-finding at L1 without aid of the tracker.

However, as useful and dependable as the lepton triggers are, they still suffer an efficiency of about 5% due to the approximate identification criteria used at HLT. What ultimately proves more effective for this analysis is to trigger on E_T^{miss} . It is a remarkable technical achievement that the E_T^{miss} can be reconstructed at all at L1, and even more remarkable that the E_T^{miss} triggers reach essentially 100% efficiency above 250 GeV (though the nominal threshold is closer to 100 GeV). The arrival of the E_T^{miss} triggers essentially replaced the leptonic triggers, but we still take the logical OR of all of them in order to supplement the efficiency in the low E_T^{miss} region. The full list of trigger paths considered is shown in Table 6.1. The redundancy afforded by so many different triggers is also extremely valuable, since there are frequently bugs that cost efficiency in any single path.

The dataset considered in this analysis corresponds to an integrated luminosity of 35.9 fb^{-1} collected in 2016.

The trigger efficiencies as a function of offline E_T^{miss} are shown in Figures 6.1 (electrons) and 6.2 (muons), for lepton and E_T^{miss} triggers alone, as well as the full combination. The efficiency is about $99.8 \pm 0.1\%$ for events with E_T^{miss} greater than 200 GeV. Additionally, the efficiency as a function of other key analysis variables, N_{jets} , N_{b} and M_J , is shown in Figures 6.3 and 6.4 after a basic selection. The excellent efficiency is maintained in all corners of phase space considered in the analysis.




Figure 6.1: Trigger efficiencies as a function of E_T^{miss} in electron events for the OR of all lepton triggers (left), all E_T^{miss} triggers (center), and all lepton and E_T^{miss} triggers together (right). The efficiencies are measured using a data sample collected using the HLT_PFJet450 and HLT_AK8PFJet450 triggers, and offline requirements of one electron, 3 or more jets, and $S_T > 500$ GeV [67].




Figure 6.2: Combined trigger efficiencies as a function of E_T^{miss} in muon events for the OR of all lepton triggers (left), all E_T^{miss} triggers (center), and all lepton and E_T^{miss} triggers together (right). The efficiencies are measured using a data sample collected using the HLT_PFJet450 and HLT_AK8PFJet450 triggers, and offline requirements of one muon, 3 or more jets, and $S_T > 500$ GeV [67].




Figure 6.3: Combined trigger efficiencies in electron events as a function of N_{jets} (left), N_b (center) and M_J (right). The efficiencies are measured using a data sample collected using the HLT_PFJet450 and HLT_AK8PFJet450 triggers, and offline requirements of one electron, 3 or more jets, and $S_T > 500$ GeV and $E_T^{\text{miss}} > 200$ GeV [67].

Table 6.1: List of HLT paths used to populate all analysis regions. Duplicate triggers with varying thresholds are included to accommodate changes in prescale necessary with variation of instantaneous luminosity.

E_T^{miss} Triggers
HLT_PFMET100_PFMHT100_IDTight, HLT_PFMETNoMu100_PFMHTNoMu100_IDTight
HLT_PFMET110_PFMHT110_IDTight, HLT_PFMETNoMu110_PFMHTNoMu110_IDTight
HLT_PFMET120_PFMHT120_IDTight, HLT_PFMETNoMu120_PFMHTNoMu120_IDTight
Muon Triggers
HLT_Mu15_IsoVVVL_PFHT350
HLT_Mu15_IsoVVVL_PFHT400
HLT_Mu50_IsoVVVL_PFHT400
HLT_IsoMu22_eta2p1, HLT_IsoMu24, HLT_IsoTkMu24, HLT_Mu50
Electron Triggers
HLT_Ele15_IsoVVVL_PFHT350
HLT_Ele15_IsoVVVL_PFHT400
HLT_Ele50_IsoVVVL_PFHT400
HLT_Ele27_eta2p1_WPLoose_Gsf
HLT_Ele25_eta2p1_WPTight_Gsf, HLT_Ele27_WPTight_Gsf
HLT_Ele105_CaloIdVT_GsfTrkIdT, HLT_Ele115_CaloIdVT_GsfTrkIdT




Figure 6.4: Trigger efficiencies in muon events as a function of N_{jets} (left), N_b (center) and M_J (right). The efficiencies are measured using a data sample collected using the HLT_PFJet450 and HLT_AK8PFJet450 triggers, and offline requirements of one electron, 3 or more jets, and $S_T > 500$ GeV and $E_T^{\text{miss}} > 200$ GeV [67].

6.2 Simulated event samples

Though the quantitative background estimate relies only minimally on Monte Carlo simulation, we make extensive use of simulations to study the background processes and develop the analysis strategy. Simulation of the signal, of course, is the only good way to estimate the signal acceptance/efficiency and interpret the results of the search.

Since the events considered in this analysis live in extreme phase space for the Standard Model, it is critical to have a large simulated sample with well-populated tails, even at the expense of accuracy in the bulk regions. For this reason, most of the samples used are generated at leading order, not next-to-leading order, to facilitate generation of a large sample. In particular, CMS produces LO samples with cuts applied after the generation of the hard scatter, saving the costly step of simulating the detector response. This results in extremely useful large samples, for example, $t\bar{t}$ production with a threshold on H_T or neutrino p_T that would be impractical to populate inclusively by brute force. This technique is the only way to keep up with the LHC event rate for relatively common processes like $t\bar{t}$, W , and Z production. It is a testament to the amazing performance of the LHC that it is unfeasible to match, for instance, its inclusive $t\bar{t}$ production

rate even with infinitely-parallelizable simulation offline. This is a potentially a major issue for many analyses and will only be exacerbated with the large rates and huge data samples expected at the High Luminosity LHC.

The full list of background Monte Carlo samples is shown in Table 6.2. Most common processes are generated with MADGRAPH5_AMC@NLO 2.2.2 [68] at LO. Rarer processes like single top production, or $t\bar{t}$ pairs in association with vector bosons, are generated at NLO with MADGRAPH5_AMC@NLO or POWHEG v2 [69, 70, 71]. All samples use the NNPDF 3.0 [72] set of parton distribution functions. Hadronization and fragmentation of partons produced in the hard scatter is performed by PYTHIA 8.205 [73], and the underlying event is based on the CUETP8M1 tune [74]. Simulation of the material interactions and detector response is performed with GEANT4 [75]. Cross sections are taken from the highest order calculation available; for $t\bar{t}$, the cross section is evaluated at next-to-next-to-leading order plus a next-to-next-to-leading logarithmic correction [76].

The signal simulations are generated similarly with MADGRAPH5_AMC@NLO in a 2D grid of the gluino and neutralino masses, for $m_{\tilde{g}}$ from around 600 to 2300 GeV, and $m_{\tilde{\chi}_1^0}$ from 1 to 1600 GeV. Unlike the background samples, which use a full GEANT simulation of the detector, the signal samples use a faster parametrized simulation of the detector response (FastSim) in order to save processing time. Additional scale factors are applied to correct for differences from the full simulation for aspects like the b-tagging and lepton reconstruction efficiencies.

Simulated pileup events are generated by PYTHIA 8.205 and overlaid with the hard scatter events. Since the production of samples is initiated before the average instantaneous luminosity for the year can be known, the distribution of number of pileup interactions in simulation can never perfectly match that of the data. However, because the variables used in the analysis are designed and carefully calibrated to be insensitive to the number of pileup interactions, reweighting of the simulations to match the pileup

distribution in data is not performed, as it would greatly weaken the statistical power of the samples for only a small reduction in systematic mismodeling.

Dataset name	Events	L [fb ⁻¹]
TTJets_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	10,259,872	12.57
TTJets_SingleLeptFromT_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	61,973,977	346.80
TTJets_SingleLeptFromT_genMET-150_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	16,959,409	1,819.14
TTJets_SingleLeptFromTbar_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	60,210,394	336.94
TTJets_SingleLeptFromTbar_genMET-150_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	17,066,856	1,830.66
TTJets_DiLept_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	30,444,678	355.41
TTJets_DiLept_genMET-150_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	9,890,329	1,806.33
TTZToLLNuNu_M-10_TuneCUETP8M1_13TeV-amcatnlo-pythia8	1,992,438	1,708.99
TTZToQQ_TuneCUETP8M1_13TeV-amcatnlo-pythia8	749,400	310.65
TTWJetsToLNu_TuneCUETP8M1_13TeV-amcatnloFXFX-madspin-pythia8	5,280,565	6,838.97
TTWJetsToQQ_TuneCUETP8M1_13TeV-amcatnloFXFX-madspin-pythia8	833,298	547.04
TTTT_TuneCUETP8M1_13TeV-amcatnlo-pythia8	250,000	4,811.39
ST_tW_antitop_5f_NoFullyHadronicDecays_13TeV-powheg_TuneCUETP8M1	11,408,144	586.04
ST_tW_top_5f_NoFullyHadronicDecays_13TeV-powheg_TuneCUETP8M1	11,345,619	582.83
ST_t-channel_antitop_4f_inclusiveDecays_TuneCUETP8M2T4_13TeV-powhegV2-madspin	3,928,063	48.52
ST_t-channel_top_4f_inclusiveDecays_TuneCUETP8M2T4_13TeV-powhegV2-madspin	5,993,676	44.06
ST_s-channel_4f_leptonDecays_13TeV-amcatnlo-pythia8_TuneCUETP8M1	1,000,000	116.20
WJetsToLNu_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	57,026,058	0.93
WJetsToLNu_HT-70To100_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	10,094,300	6.08
WJetsToLNu_HT-100To200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	78,236,266	48.00
WJetsToLNu_HT-200To400_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	38,328,549	87.99
WJetsToLNu_HT-400To600_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	7,759,701	130.93
WJetsToLNu_HT-600To800_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	18,687,480	1,281.68
WJetsToLNu_HT-800To1200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	7,745,467	1,163.65
WJetsToLNu_HT-1200To2500_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	6,290,216	3,911.61
WJetsToLNu_HT-2500ToInf_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	2,388,086	61,368.93
QCD_HT100to200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	80,160,711	0.00
QCD_HT200to300_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	57,580,393	0.03
QCD_HT300to500_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	54,537,903	0.15
QCD_HT500to700_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	62,271,343	2.12
QCD_HT700to1000_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	45,412,780	6.96
QCD_HT1000to1500_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	15,127,293	14.22
QCD_HT1500to2000_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	11,826,702	97.34
QCD_HT2000toInf_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	6,039,005	237.57
DYJetsToLL_M-50_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	49,144,274	8.16
DYJetsToLL_M-50_HT-70to100_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	9,616,188	44.60
DYJetsToLL_M-50_HT-100to200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	10,391,819	60.61
DYJetsToLL_M-50_HT-200to400_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	9,404,362	178.85
DYJetsToLL_M-50_HT-400to600_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	9,950,320	1,471.66
DYJetsToLL_M-50_HT-600to800_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	8,292,957	4,946.62
DYJetsToLL_M-50_HT-800to1200_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	2,668,730	3,210.09
DYJetsToLL_M-50_HT-1200to2500_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	596,079	4,177.73
DYJetsToLL_M-50_HT-2500toInf_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	399,492	125,304.88
WJetsToQQ_HT-600ToInf_TuneCUETP8M1_13TeV-madgraphMLM-pythia8	1,026,587	8.92
WH_HTtoBB_ZToLNu_M125_13TeV_amcatnloFXFX_madspin_pythia8	2,179,687	2,872.12
WWTo2L2Nu_13TeV-powheg	1,999,000	164.15
WWToLNuQQ_13TeV-powheg	1,999,200	39.99
WZTo1L1Nu2Q_13TeV_amcatnloFXFX_madspin_pythia8	23,939,924	735.80
WZTo1L3Nu_13TeV_amcatnloFXFX_madspin_pythia8	1,703,772	170.78
WZTo2L2Q_13TeV_amcatnloFXFX_madspin_pythia8	26,517,272	1,699.59
WZTo3LNu_TuneCUETP8M1_13TeV-powheg-pythia8	1,993,200	449.97
ZH_HTtoBB_ZToNuNu_M125_13TeV_amcatnloFXFX_madspin_pythia8	2,159,477	6,950.97
ZZ_TuneCUETP8M1_13TeV-pythia8	990,064	59.92
ttHTobb_M125_13TeV_powheg_pythia8	3,936,004	3,936.00

Table 6.2: List of simulated event samples along with their equivalent luminosities [67].

Chapter 7

Event Selection

7.1 Baseline selection

Our gluino model results in a final state containing four top quarks and two invisible neutralinos, and therefore we expect substantial E_T^{miss} , many jets (9+ hard partons if only one top decays leptonically!), four of which are b-flavor. To form an inclusive sample as a starting point for the analysis, we define the following *baseline selection*:

$$N_{\text{lep}} = 1, N_{\text{veto}} = 0, S_T > 500 \text{ GeV}, E_T^{\text{miss}} > 200 \text{ GeV}, N_{\text{jets}} \geq 6, N_b \geq 1$$

These requirements bring us to a region in the proximity of the final search selection, which is ultimately delineated by two additional discriminating variables, m_T and M_J . The events in the baseline selection outside the search region share similar composition and kinematic properties, and will form invaluable control samples for the background prediction. Additionally, the requirements on N_{lep} , S_T and E_T^{miss} ensure a uniform trigger of close to 100% over the entire baseline region.

The distributions of S_T , E_T^{miss} , N_{jets} , and N_b in data and simulation are shown in

Figure 7.1. Note that the validity of the background prediction will not depend directly on agreement with simulation in these variables. Further categorization of events is achieved by subdividing or “binning” events in E_T^{miss} , N_{jets} and N_b for a total of 18 ($3 \times 2 \times 3$) categories:

- Three E_T^{miss} bins: $200 \text{ GeV} < E_T^{\text{miss}} \leq 350 \text{ GeV}$, $350 \text{ GeV} < E_T^{\text{miss}} \leq 500 \text{ GeV}$, $E_T^{\text{miss}} > 500 \text{ GeV}$
- Two N_{jets} bins: $6 \leq N_{\text{jets}} \leq 8$, $N_{\text{jets}} \geq 9$
- Three N_b bins: $N_b = 1$, $N_b = 2$, $N_b \geq 3$

A cutflow showing the impact of each successive selection is shown in Table 7.1. After the baseline selection is applied, the dominant background component stems from events containing $t\bar{t}$ production, in particular events with exactly one top quark decaying leptonically. This is unsurprising, given that $t\bar{t}$ is the only common process that at tree level can yield an isolated lepton and several jets, two of which are b-flavor, without the need for additional jets from higher order processes. Events involving W boson production are typically suppressed by the b-tagging requirements, and multijet QCD is suppressed by the lepton and E_T^{miss} requirements.




Figure 7.1: Distributions of key variables for the baseline selection, in data (points) and Monte Carlo (stacked histograms). S_T and E_T^{miss} are shown in the top row; N_{jets} and N_b in the bottom row. Baseline selection is applied, except for the variable plotted. The blue hatched area in the ratio panel indicates the statistical uncertainty of the simulation. The simulations are normalized to the number of events observed in data for each selection, amounting to approximately 80-85% of the expectation based on luminosity. The precise normalization ratio is shown in parentheses on each frame [67].

Table 7.1: Simulated event yields, after the selection criteria are applied successively. The two benchmark signal models correspond to a non-compressed scenario (T1tttt(NC)) with $m_{\tilde{g}} = 1800$ GeV and $m_{\tilde{\chi}_1^0} = 100$ GeV, and a compressed scenario (T1tttt(C)) with $m_{\tilde{g}} = 1400$ GeV and $m_{\tilde{\chi}_1^0} = 1000$ GeV. The requirements above the horizontal line correspond to the baseline selection [67].

$\mathcal{L} = 35.9 \text{ fb}^{-1}$	Other	QCD	$t\bar{t}V$	Single t	W+jets	$t\bar{t}$ (1ℓ)	$t\bar{t}$ (2ℓ)	SM bkg.	T1tttt NC	T1tttt C
$1\ell, S_T > 500 \text{ GeV}, \text{MET} > 200 \text{ GeV}$	8468.1	3623.0	1624.4	9096.5	96033.4	49212.8	11813.9	179872.1	31.0	92.6
Track veto	6783.3	3414.2	1367.6	8103.3	92184.9	45789.0	7605.8	165248.1	28.5	80.2
$N_{\text{jets}} \geq 6$	420.2	200.8	512.1	1149.3	4185.5	11588.4	1435.9	19492.2	25.0	74.5
$N_b \geq 1$	125.9	105.4	411.6	913.4	1007.1	9408.2	1183.0	13154.6	23.7	70.8
$M_J > 250 \text{ GeV}$	107.8	79.4	357.9	749.5	819.4	7943.3	919.1	10976.5	23.6	66.1
$m_T > 140 \text{ GeV}$	9.5	6.7	42.3	43.2	28.2	87.2	369.1	586.2	19.3	39.4
$M_J > 400 \text{ GeV}$	5.5	6.1	20.4	25.7	16.9	51.2	143.8	269.5	18.9	25.1
$N_b \geq 2$	1.8	0.6	8.3	11.1	3.8	25.1	69.4	120.1	14.2	18.8
$\text{MET} > 350 \text{ GeV}$	0.29	0.16	2.30	2.82	0.81	2.36	15.19	23.93	12.50	9.08
$\text{MET} > 500 \text{ GeV}$	0.00	0.00	0.63	0.70	0.27	0.35	3.50	5.40	9.96	3.82
$N_{\text{jets}} \geq 9$	0.03	0.00	0.12	0.06	0.05	0.04	0.39	0.69	3.82	2.61
	Other	QCD	$t\bar{t}V$	Single t	W+jets	$t\bar{t}$ (1ℓ)	$t\bar{t}$ (2ℓ)	SM bkg.	T1tttt NC	T1tttt C

7.2 Suppressing single-lepton backgrounds

The background sample surviving the baseline selection is dominated by events containing production of $t\bar{t}$ pairs with a single leptonic W decay. The composition is summarized in Figure 7.2, left. One of the classic techniques to address this background is to examine the transverse mass, m_T , of the lepton and E_T^{miss} , and remove events where it is kinematically consistent that they arise from a single W boson. Events with additional sources of E_T^{miss} such as neutralinos, or just additional neutrinos, can have m_T significantly larger than the W mass. The distribution of m_T in the baseline selection is shown in Figure 7.3. Requiring m_T greater than 140 GeV dramatically reduces the contribution from semileptonic $t\bar{t}$, leaving a much smaller background from dileptonic $t\bar{t}$ production. The T1tttt signal models, on the other hand, pass the m_T requirement with high efficiency (60-80%), thanks to the additional E_T^{miss} arising from the two neutralinos, especially for non-compressed models. We thus take the high m_T selection to form our search region.

The background composition after the requirement of $m_T > 140$ GeV is shown in Figure 7.2 (right) and the largest share is from dileptonic $t\bar{t}$ events. Since we explicitly veto events with two reconstructed leptons, dileptonic $t\bar{t}$ events can only enter the analysis if one lepton is “lost”, as will be discussed in Section 7.3.




Figure 7.2: Composition of background processes contributing to the baseline selection (right) and the baseline selection with $m_T > 140$ GeV (right).




Figure 7.3: Distribution of m_T in the baseline selection, in data (points) and Monte Carlo (stacked histograms) [65].

7.3 Suppressing dilepton backgrounds

The dominant background process for high m_T selections is the production of $t\bar{t}$ pairs, with both top quarks decaying leptonically. These events can only pass the $N_{\text{lep}} = 1$ requirement if one lepton is “lost”, because it is either too soft ($< 20 \text{ GeV}$), fails a strict identification requirement, overlaps with a jet, or if it is a hadronically decaying τ , which sidestep the lepton reconstruction altogether. The lost lepton composition is shown in Figure 7.4 with large components coming from hadronic τ leptons, leptons out of kinematic acceptance, and failed identification.




Figure 7.4: Composition of lost leptons by mutually exclusive category. “Failing acceptance” indicates the presence of electrons and muons with $p_T < 20 \text{ GeV}$ [67].

Given that the background is dominated by such a specific event topology, it begs the question whether further reduction can be achieved by tailoring a looser, second lepton selection to recover the lost leptons and veto these events. It turns out there are a number of obvious ways to expand the lepton selection. Though the strict identification criteria listed in Section 4.4 are certainly necessary to ensure purity in broader samples of events, they are excessively strict for our narrow corner of phase space at high m_T . Since most events do contain a lost lepton, the presence of an isolated track with an only a loose indication of being a lepton is enough tag leptons with reasonable purity. Isolation itself is

also a fairly strong discriminator between true, prompt leptons and fakes or leptons from heavy flavor decays. For these reasons, we can safely consider isolated “leptonic tracks”, tracks identified as electrons or muons by Particle Flow, which is tuned for maximal efficiency rather than purity. It is also profitable to expand the kinematic acceptance by lowering the p_T threshold from 20 to 10 GeV.

A similar approach is useful for identifying hadronic τ , which are the largest category of lost leptons. Since the high m_T region is enriched with τ leptons, isolated hadronic tracks are good indicators of the presence of single-prong τ decays. To include common single-prong τ decays with additional neutral pions, the hadronic tracks use a special version of mini-isolation that only includes charged particle candidates. Only hadronic tracks with $p_T > 15$ GeV are considered.

The leptonic and hadronic tracks passing these requirements are found in approximately 40-45% of dileptonic $t\bar{t}$ events with lost leptons. However, extra care is required before vetoing these events. It turns out that almost half of the $T1tttt$ signal events also have a lost lepton (since there are 4 W bosons in each event), and so naively vetoing all of the tracks could remove approximately 20% of the signal events as well.

To preserve signal efficiency, there are a couple of good ways to tailor the veto to specific features of dileptonic $t\bar{t}$. The first is simple: recognizing that in $t\bar{t}$ events, the two leptons are always of opposite charge, we veto only tracks of opposite sign from the primary lepton. This removes from consideration one third of the true lost leptons in signal events (two of the non-primary W bosons are opposite sign to the primary lepton, one is same sign), and also reduces the fake rate by half.

The second way to improve the veto takes advantage of the variable m_{T2} , defined in Section 4.5. m_{T2} is a generalization of m_T , which restores the cutoff at the W mass for events with two leptonically decaying W bosons. Distributions of m_{T2} are shown in Figure 7.5. Almost all $t\bar{t}$ events with tracks have $m_{T2} < 80$ GeV, while more than

half of $T1tttt$ events with tracks have larger m_{T2} , because the E_T^{miss} is not produced in association with the two leptons. Only vetoing leptonic (hadronic) tracks with m_{T2} less than 80 (60) GeV saves approximately half of the signal events that would otherwise be vetoed. With both of these improvements, the veto still removes nearly 40% of the dileptonic $t\bar{t}$ background (or 30% of the total background) at the cost of only 6-9% of the signal, substantially increasing the signal to background ratio.

Where the variable N_{veto} is used, it always refers to the number of leptonic and hadronic tracks that satisfy the charge and m_{T2} requirements.




Figure 7.5: Distributions of m_{T2} for events in the baseline selection and with $m_T > 140$ GeV for leptonic tracks (left) and hadronic tracks (right) in signal (red) and dileptonic $t\bar{t}$ (blue). The distributions are normalized to 100% to illustrate only their shapes [67].

7.4 Background outlook

Based on the results developed in this chapter, the background measurement will primarily involve predicting the dileptonic $t\bar{t}$ contribution in an extreme kinematic region with many jets and large E_T^{miss} . Since dileptonic $t\bar{t}$ pairs yield only two jets from the top decay product (or sometimes three, including hadronic τ jets), many of the jets will come from higher order processes with additional radiation. Figure 7.6 shows an example

cartoon of one contributing diagram. To reach the baseline selection with six jets, three or four additional jets are required. For the most sensitive bins, which require nine or more jets, six or seven jets must come from additional radiation!

High order diagrams like these are very distant from the core $t\bar{t}$ processes that our (leading order!) simulations aim to model, and the modeling of the high order QCD processes must be taken *cum grano salis*. There are a number of experimental aspects that are difficult to simulate as well, including how often leptons are lost, and the extreme tail of the E_T^{miss} resolution. For these reasons, it is critical to develop a background measurement deeply grounded in data in order to be confident in the results. Our background estimate will be delivered by an additional powerful discriminating variable, M_J , as will be developed in the next sections.




Figure 7.6: Example diagram contributing to dileptonic $t\bar{t}$ production with $N_{\text{jets}} = 6$.

Chapter 8

M_J

8.1 Large-Radius jets and M_J

The standard, $R = 0.4$ anti- k_T jets used in CMS are intended to reconstruct the momenta of individual partons from the hard scatter. More recently, many techniques involving large-radius clustering have become popular, where the goal is to capture multi-parton information in a single jet (colloquially called “fatjets”). This is extremely useful for identifying hadronically decaying boosted massive objects, like vector bosons or top quarks [77, 78]. These applications emphasize optimization of the jet mass resolution as well as variables that quantify substructure [79], taking advantage of the fact that, for example, boosted top fatjets should contain three prongs.

Recently, large-radius jet techniques have been applied using a completely different approach, to create a generic event-wide discriminating variable for high-multiplicity searches [80, 81, 82]. This variable is known as M_J , and is formed by taking the scalar sum of masses of jets formed by large-radius clustering, as described in Section 4.3:

$$M_J = \sum_{\text{Large-R jets}} m_J \tag{8.1}$$

where the m_J are the masses of each large-radius jet (clustered with $R = 1.4$). The purpose of M_J is not to capture precise mass resonances from real particles, but to allow uncorrelated partons to overlap in a single large-radius jet, resulting in what is called “accidental mass”, where separated partons with high p_T can fake the signature of a massive resonance. This can be understood by examining the formula for the invariant mass of two highly relativistic particles:

$$m(p_1, p_2) = \sqrt{2p_1 p_2 (1 - \cos(\Delta R))} \quad (8.2)$$

If two high momentum partons at large ΔR are interpreted as decay products of a highly boosted parent, the parent mass would be inferred to be extremely large, since its decay is able to establish a large separation. Hence, clustering of neighboring partons at large radius results in fatjets with very large mass. The accidental overlap is enhanced by using extremely large jet radii— $R = 1.4$ in this analysis, compared to $R = 0.8$ typically used for boosted object tagging.

The use of M_J in the analysis fills a role traditionally filled by variables that characterize the energy scale, such as H_T . In fact, M_J is highly correlated with H_T , but combines information about the jet multiplicity and the angular structure of the final state as well.

Distributions of M_J in the T1tttt signal and $t\bar{t}$ are compared in Figure 8.1. The signal events populate high M_J due to a combination of frequent overlap of partons from the top quark decays, as well as collimation of boosted hadronic tops into a single fatjet. The $t\bar{t}$ events, meanwhile, have a softer spectrum. Though $t\bar{t}$ events receive moderate M_J from the intrinsic top mass, they rely on additional radiation to reach large mass, which results in a softer spectrum. The compressed T1tttt benchmark model (dotted histogram) has a significantly softer M_J spectrum than the non-compressed model. This

is easily understood by the fact that all of the gluino decay products are less boosted in the compressed case, and contribute less mass even with wide-angle overlap.




Figure 8.1: Shape comparison of M_J distributions in $t\bar{t}$ events (blue) and $T1tttt$ events (red), for compressed and non-compressed model points (dashed and solid, respectively). The selection is similar to the baseline selection [67].

8.2 Understanding high M_J backgrounds

In this analysis, regions with $M_J > 400$ GeV are considered “high M_J ”, a threshold slightly beyond twice the top quark mass. Substantial effort has been devoted to understanding how $t\bar{t}$ events can reach large values of M_J beyond their intrinsic mass.

It turns out to be cleanly delineated by the presence of extra jets from initial state radiation (ISR). Figure 8.2 shows M_J distributions with and without substantial contributions from ISR. In cases with minimal ISR (left), M_J quickly falls off beyond $2m_t$, even for highly boosted $t\bar{t}$ pairs! Contrast this with H_T —since M_J sums invariant masses instead of p_T , it is insensitive to boosts that don’t introduce additional overlap. On the

other hand, in cases with substantial ISR (right), the M_J distribution grows a tail. In general, isolated ISR jets don't contribute to M_J on their own; mass arising from the parton shower alone is suppressed by a factor of α_s relative to mass from overlap of partons [80]. However, occasionally ISR jets can overlap with each other or with partons from the $t\bar{t}$ decay, resulting in a massive fatjet. Understanding the rate for overlap of with ISR jets is the main task for predicting the tail of the M_J distribution in $t\bar{t}$.

Two characteristic simulated $t\bar{t}$ events with low and high M_J values are shown in Figures 8.3 and 8.4, respectively. Both events contain isolated ISR jets that don't contribute substantial mass to M_J (blue circles). However, in the high M_J event, one ISR jet overlaps with a top quark fatjet, resulting in significantly larger M_J .




Figure 8.2: Comparison of M_J shapes in simulation, for events with limited ISR (left) and a significant ISR contribution (right), classified based on the total ISR p_T , or equivalently the net p_T of the $t\bar{t}$ or $\tilde{g}\tilde{g}$ pair. $t\bar{t}$ is shown in blue, and $T1tttt$ in red. Events are required to have at least one lepton, $E_T^{\text{miss}} > 200$ GeV, and $H_T > 500$ GeV [67].




Figure 8.3: Event display showing the large radius jet clustering in a typical, low M_J $t\bar{t}$ event from simulation. The dots represent the clustering constituents (black for jets and leptons, green for b-tagged jets). Each ring represents a fatjet, with the ring color indicating the fatjet mass based on the z -axis scale. The directions of the top quarks and their decay products are indicated with symbols (in light green and for each top family). This event has two moderate m_J fatjets from each collimated top (red and yellow rings), and two essentially massless fatjets containing isolated jets from ISR (blue rings).




Figure 8.4: Event display showing the large radius jet clustering in a high M_J $t\bar{t}$ event from simulation. The dots represent the clustering constituents (black for jets and leptons, green for b-tagged jets). Each ring represents a fatjet, with the ring color indicating the fatjet mass based on the z -axis scale. The directions of the top quarks and their decay products are indicated with symbols (in light green and for each top family). This event has one moderate m_J fatjet from an isolated collimated top quark (top center, in yellow), and one fatjet with significantly larger mass (bottom left, red) which contains a collimated top quark overlapping with an additional jet from ISR. There are three additional single-constituent ISR fatjets with very low m_J (blue rings).

8.3 M_J performance studies

Since M_J combines information about both the energy scale and the jet multiplicity, there is a significant overlap with classic search variables like H_T and N_{jets} . The precise overlap is studied in detail in [66], but key results are summarized in Figure 8.5, which compares ROC curves of $t\bar{t}$ efficiency vs T1tttt efficiency for M_J , N_{jets} and H_T , as well as boosted decision trees (BDTs) trained to make best use of each pairwise combination of these variables for each particular benchmark model. For the non-compressed model (left), M_J significantly outperforms both H_T and N_{jets} for most of the range. Furthermore, combining M_J and H_T into a single BDT does not improve the performance relative to M_J alone, indicating that the majority of information in H_T is already encapsulated by M_J . However, the BDT combination of N_{jets} and H_T does, in fact, outperform M_J alone in the high purity, low background regime—meaning there is discrimination power in the jet multiplicity beyond what is captured by M_J .




Figure 8.5: Receiver operating characteristic curves (ROC curves) comparing performance of discrimination between T1tttt and $t\bar{t}$ for $M_J, H_T, N_{\text{jets}}$, and BDTs formed from their pairwise combinations. Each curve shows the $t\bar{t}$ efficiency vs the T1tttt efficiency for a simple cut on each variable for continuously varying thresholds. Points closer to the bottom right corner indicate superior performance. Left: non-compressed T1tttt. Right: compressed T1tttt. [66]

A key feature of M_J , though, is that it attains excellent performance without need for

dedicated tuning to each model point, which facilitates understanding of systematics and provides coverage for a variety of models with a single unifying variable. The simplicity also has substantial aesthetic appeal. Given that the performance of M_J and the BDT converge in the high efficiency regime, a successful strategy is to use a loose, efficient M_J requirement driven mostly by the background shape ($M_J > 400 \text{ GeV}$), and additionally categorize events based on N_{jets} , recovering much of the performance gap relative to the BDT.

This strategy also pays off well for compressed models. Based on Figure 8.5 (right), M_J is not an optimal variable for compressed mass points, where N_{jets} actually performs much better. The loose M_J requirement then primarily serves to enable the background estimation (explained in the next chapter), while most of the discrimination power for these models comes from the categorization in N_{jets} and N_{b} .

Figure 8.6 compares performance of M_J for various choices of the clustering radius. In general, larger values of R increase the rate of overlap, and shift M_J distributions to the right in both signal and background. For non-compressed models (left), the performance is not strongly dependent on the radius, though there is a slight preference for 1.2 relative to 1.4 or 0.8. However, for compressed models, where the top quarks are substantially less boosted, there is a strong preference for larger radii, which enable overlap and containment of top quarks even at modest p_{T} . As a compromise of performance across the mass plane, $R = 1.4$ was selected for this analysis.




Figure 8.6: ROC curves comparing discrimination power against $t\bar{t}$ for variations in the clustering radius. Left: non-compressed T1tttt. Right: compressed T1tttt. [67]

Chapter 9

Background Estimation

9.1 Method

The background prediction is centered around the two main discriminating variables of the analysis, M_J and m_T . Events are categorized according to where they lie in the M_J - m_T plane, as shown for the signal and dominant background components in Figure 9.1. The introduction of M_J gives us a powerful tool: we can now split the high m_T region in half, yielding a signal-rich search region at high M_J and a signal-poor control region at low M_J . Similarly, we divide the low m_T regions into two additional control regions, for a total of four regions:

- R1 (CR): $m_T \leq 140$ GeV, $250 \leq M_J \leq 400$ GeV
- R2 (CR): $m_T \leq 140$ GeV, $M_J > 400$ GeV
- R3 (CR): $m_T > 140$ GeV, $250 \leq M_J \leq 400$ GeV
- R4 (SR): $m_T > 140$ GeV, $M_J > 400$ GeV.

We can then use the high m_T control region, R3, to anchor the high m_T normalization,




Figure 9.1: Distribution of simulated single-lepton $t\bar{t}$ events (dark-blue triangles), dilepton $t\bar{t}$ events (light-blue inverted triangles), and $T1tttt(1800,100)$ events (red squares) in the M_J - m_T plane after the baseline selection. Each marker represents one expected event at 35.9fb^{-1} . Overflow events are placed on the edge of the plot. The values of the correlation coefficients ρ for each background process are given in the legend. Region R4 is the nominal signal region, while R1, R2, and R3 serve as control regions. [65]

encapsulating effects including the lost lepton rate, and the high m_T tail for single-neutrino processes. Then, we can use the low m_T regions, R1 and R2, to precisely measure a transfer factor from low to high M_J , and apply this shape at high m_T . This leads to the following prediction relating the expected yields in the four regions:

$$\mu_{R4}^{\text{bkg}} = N_{R2} \times N_{R3} / N_{R1}, \quad (9.1)$$

where N_{Ri} are the yields observed in the three control regions, and μ_{R4}^{bkg} is the best prediction for the rate in the signal region R4. This is an example of a relatively standard technique known as an ABCD prediction.

The validity of this method depends on the ability to translate the M_J shape measured

in the low m_T regions to the high m_T regions. It's not at all obvious that this translation should work. Since hadronic decays funnel more of the top mass into visible energy than leptonic decays do, the semileptonic $t\bar{t}$ events (the dominant contribution at low m_T) have a higher intrinsic mass than the fully leptonic events dominating the high m_T regions. This difference is clear in the M_J shape comparison previously shown in a loose selection and the absence of ISR, in Figure 8.2 (left).

The key enabling factor, though, is that in the ISR dominated regime at high N_{jets} , the M_J shapes for both $t\bar{t}$ components converge. This can be understood by the fact that, once there are enough hard partons in an event, the cost to add an additional parton is fixed, independent of the initial state—more or less just a factor of α_s . To satisfy the baseline requirement of $N_{\text{jets}} \geq 6$, semileptonic $t\bar{t}$ events need two additional jets from ISR, while dileptonic $t\bar{t}$ events need four additional jets. As result, for $N_{\text{jets}} \geq 6$, the difference in number of partons from top decays is essentially washed away, and the N_{jets} spectra match very closely for single and dileptonic $t\bar{t}$ events, as shown in Figure 9.2 (left). Since the rate for accidental overlap with ISR jets is largely driven by the number of jets in the event, once the N_{jets} spectra match, the M_J shapes agree as well. This is visible in Figure 9.2 (right), where the M_J distributions for each component share almost exactly the same slope on the tail from moderate M_J and above. At low M_J , where the ISR jets do not overlap with the top jets, the intrinsic mass difference between the two components remains. Therefore, for this analysis, the region with $M_J < 250 \text{ GeV}$ is excluded from the control regions. Ultimately, the M_J transfer factors match because the transfer from medium to large M_J is driven by properties of the ISR system, common to both $t\bar{t}$ components after moderate requirements on N_{jets} and M_J .

Though the common ISR system removes most of the difference in the M_J spectra for the low and high m_T regions, there can still be residual correlations remaining from the varying $t\bar{t}$ composition or introduced by subdominant background processes. The




Figure 9.2: Comparison of N_{jets} and M_J distributions, normalized to the same area, in simulated $t\bar{t}$ events with two true leptons at high m_T and one true lepton at low m_T , after the baseline selection is applied. The shapes of these distributions are similar. These two contributions are the dominant backgrounds in their respective m_T regions. The dashed vertical line on the right-hand plot indicates the $M_J > 400$ GeV threshold that separates the signal regions from the control samples. The shaded region corresponding to $M_J < 250$ GeV is not used in the background estimation. [67]

effects of small, remaining correlations are removed with factors, κ , that are obtained from Monte Carlo simulation:

$$\kappa = \frac{N_{\text{R4}}^{\text{MC,bkg}}/N_{\text{R3}}^{\text{MC,bkg}}}{N_{\text{R2}}^{\text{MC,bkg}}/N_{\text{R1}}^{\text{MC,bkg}}} \quad (9.2)$$

The κ factors can be understood as removing any “non-closure” observed in simulation: the result of dividing the actual R4 yield in Monte Carlo by the naive ABCD prediction obtained from evaluating Eq. 9.1 in Monte Carlo. The modified ABCD prediction is thus given by:

$$\mu_{\text{R4}}^{\text{bkg}} = \kappa N_{\text{R2}} \times N_{\text{R3}}/N_{\text{R1}}. \quad (9.3)$$

In principle, this prediction could be applied independently in each of the 18 bins of $E_{\text{T}}^{\text{miss}}$, N_{jets} , and N_{b} defined in Section 7.1:

- Three $E_{\text{T}}^{\text{miss}}$ bins: $200 \text{ GeV} < E_{\text{T}}^{\text{miss}} \leq 350 \text{ GeV}$, $350 \text{ GeV} < E_{\text{T}}^{\text{miss}} \leq 500 \text{ GeV}$, $E_{\text{T}}^{\text{miss}} > 500 \text{ GeV}$
- Two N_{jets} bins: $6 \leq N_{\text{jets}} \leq 8$, $N_{\text{jets}} \geq 9$
- Three N_{b} bins: $N_{\text{b}} = 1$, $N_{\text{b}} = 2$, $N_{\text{b}} \geq 3$

However, the fine binning would result in control samples with poor statistics, particularly in bins of R3 with many jets or high $E_{\text{T}}^{\text{miss}}$, where the M_J distribution is shifted to larger values. To improve the control region statistics, we will again exploit the fact that the semileptonic and dilepton $t\bar{t}$ events share the same N_{jets} spectrum (Figure 9.2 left), which ensures the relative normalization of the high and low m_{T} regions is independent of N_{jets} . We study this further by examining $R(m_{\text{T}})$, the ratio of the number of events




Figure 9.3: The ratio $R(m_T)$ of high- m_T (R3+R4) to low- m_T (R1+R2) event yields for the simulated SM background, as a function of N_{jets} and N_b . The baseline selection requires $N_{\text{jets}} \geq 6$. The uncertainties shown are statistical only. [67]

at high m_T to the number of events at low m_T :

$$R(m_T) \equiv \frac{N(m_T > 140 \text{ GeV})}{N(m_T \leq 140 \text{ GeV})}. \quad (9.4)$$

Figure 9.3 shows $R(m_T)$ as a function of N_b and N_{jets} . In particular, thanks to the common N_{jets} and N_b shapes in both semi- and fully leptonic $t\bar{t}$, $R(m_T)$ is essentially independent of these variables beyond the $N_{\text{jets}} = 6$ requirement, especially in the low M_J regions. This motivates merging all of the N_{jets} and N_b categories in R1 and R3 and using a single value of $R(m_T)$ at low M_J for each category of E_T^{miss} . The merging allows the $R(m_T)$ transfer factor to be measured with good precision in data, but also introduces correlation among the N_{jets} and N_b bins in a given E_T^{miss} bin. The κ factors for all 18 bins, taking into account merging of bins in R1 and R3, are shown in Figure 9.2. The κ factors are close to unity, which makes it plausible that they could be well modeled by the Monte Carlo simulation. The modeling of κ is studied in detail in Chapter 10.5.




Figure 9.4: Values of the double-ratio κ in each of the 18 signal bins, calculated using the simulated SM background. The κ factors are close to unity, indicating the small correlation between M_J and m_T . The uncertainties shown are statistical only. [67]

9.2 Implementation

The method described above is implemented with a likelihood function that includes uncertainties on κ , the yields in each control region, accounts for correlations arising from the merged R1 and R3 bins, and correctly incorporates the signal contamination in each control region.

The Standard Model background contribution is described based on estimated mean background rates, μ_{Ri}^{bkg} , in each region Ri , for $i = 1, 2, 3, 4$. The four rates are parametrized in terms of a global rate, μ , transfer factors $R(m_T)$ and $R(M_J)$, and κ :

$$\begin{aligned}
 \mu_{R1}^{\text{bkg}} &= \mu, & \mu_{R2}^{\text{bkg}} &= \mu \cdot R(M_J), \\
 \mu_{R3}^{\text{bkg}} &= \mu \cdot R(m_T), & \mu_{R4}^{\text{bkg}} &= \kappa \cdot \mu \cdot R(M_J) \cdot R(m_T).
 \end{aligned} \tag{9.5}$$

where κ is defined as in Eq. 9.2.

There are a total of 18 copies of $R(M_J)$ and κ for each bin in N_{jets} , N_{b} and $E_{\text{T}}^{\text{miss}}$. Due to the merging of R1 and R3 across N_{jets} and N_{b} , there is only one copy of $R(m_{\text{T}})$ per bin of $E_{\text{T}}^{\text{miss}}$ (three total).

Then, defining $N_{\text{R}i}^{\text{data}}$ as the observed data yield in each region, $\mu_{\text{R}i}^{\text{MC,sig}}$ as the expected signal rate in each region, and r as a relative signal strength parameter, we can write the likelihood function as

$$\mathcal{L} = \mathcal{L}_{\text{ABCD}}^{\text{data}} \cdot \mathcal{L}_{\kappa}^{\text{MC}} \cdot \mathcal{L}_{\text{sig}}^{\text{MC}}, \quad (9.6)$$

$$\mathcal{L}_{\text{ABCD}}^{\text{data}} = \prod_{i=1}^4 \prod_{k=1}^{N_{\text{bins}}(\text{R}i)} \text{Poisson}(N_{\text{R}i,k}^{\text{data}} | \mu_{\text{R}i,k}^{\text{bkg}} + r \cdot \mu_{\text{R}i,k}^{\text{MC,sig}}), \quad (9.7)$$

$$\mathcal{L}_{\kappa}^{\text{MC}} = \prod_{i=1}^4 \prod_{k=1}^{N_{\text{bins}}(\text{R}i)} \text{Poisson}(N_{\text{R}i,k}^{\text{MC,bkg}} | \mu_{\text{R}i,k}^{\text{MC,bkg}}), \quad (9.8)$$

$$\mathcal{L}_{\text{sig}}^{\text{MC}} = \prod_{i=1}^4 \prod_{k=1}^{N_{\text{bins}}(\text{R}i)} \text{Poisson}(N_{\text{R}i,k}^{\text{MC,sig}} | \mu_{\text{R}i,k}^{\text{MC,sig}}). \quad (9.9)$$

The indices k run over each of the bins in $E_{\text{T}}^{\text{miss}}$, N_{jets} and N_{b} ; for R2 and R4, there are 18 bins, while for R1 and R3 there are only 3. The likelihoods $\mathcal{L}_{\text{ABCD}}^{\text{data}}$ accounts for the statistical uncertainty in the observed data yields, while $\mathcal{L}_{\kappa}^{\text{MC}}$ account for the uncertainties arising due to the finite size of the simulated samples.

Systematic uncertainties (described in the next Chapter) are incorporated by including additional log-normal constraints with a nuisance parameters for each set of correlated uncertainties.

This likelihood function is used to perform two distinct background estimates: a *predictive fit*, which enables simpler interpretation of the null hypothesis, and a *global fit*, which can more fully accommodate fluctuations in the observations and estimate the signal strength parameter, r . To perform the predictive fit, observations in R4 are excluded, leaving only one background parameter per constraint. In this case, the best estimates

of the background rates, μ_{Ri}^{bkg} simply converge to the observations, and the estimated background in R4 is just the simple ABCD expectation given by the combination of the three yields and κ given in Equation 9.3. The purpose of the likelihood machinery in this case is simply to handle the propagation of uncertainties correctly.

For the global fit, yields in all four regions are used. In this case, there are four constraints but only three floating parameters per ABCD plane, which allows the estimates of the background rates to be adjusted away from the observed yields, and enables the best estimate of the signal strength to be determined.

Chapter 10

Systematic Uncertainties, Background Estimate

The main assumption of the background prediction is that M_J and m_T are mostly uncorrelated, and that any small correlations are encapsulated by the Monte Carlo correction factors, κ . Understanding the systematic uncertainties in the background measurement then is a matter of quantifying how well the simulation captures these correlations.

We ultimately take a conservative approach and assign systematic uncertainties no smaller than can be bounded by direct probes of the modeling in data. Due to the statistical imprecision inherent to data-driven tests of the simulation, this strategy yields larger uncertainties than what may feel justified based on the robustness of κ and the good agreement in control samples. However, since this analysis probes challenging phase space for leading order Monte Carlo generator, building credibility in case of a discovery requires treating the inputs from simulation with no more confidence than can be proven. In fact, the systematic uncertainties are typically eclipsed by statistical uncertainties stemming from the observed control region yields, so the conservative approach has little impact on the sensitivity. Additionally, since the evaluation of the systematic uncertainties is

limited by the statistical precision of the data, the systematic understanding will be refined as more integrated luminosity is accumulated, remaining subdominant as the purely statistical uncertainties also become better-controlled.

10.1 Impact of Monte Carlo variations on κ

One of the traditional ways for probing the sensitivity to systematic effects is to introduce a physically motivated mismodeling to the simulation and measure the impact on the background prediction. Indeed, the M_J and m_T distributions themselves certainly depend on a variety of aspects which are dubiously well-modeled, like the ISR spectrum and the E_T^{miss} resolution. The key advantage of the background prediction, though, is that κ is a *double ratio*, the ratio of the M_J transfer factors at low and high m_T (or equivalently, the ratio of the m_T transfer factors at low and high M_J). As result, if the M_J shape suffers from the same mismodeling at both low and high m_T , it cancels in the ratio and κ is unaffected. Because most well-motivated variations affect M_J and m_T “globally”, the variations essentially cancel out to first order in κ . Only variations that affect M_J differently at low and high m_T , introducing an M_J - m_T correlation, can have a significant impact on the background prediction.

We have assessed the impact on κ of many well-motivated variations and the results are summarized in Table 10.1. Even for relatively extreme scenarios, κ is very stable, with variation larger than a few percent observed only in bins susceptible to statistical fluctuations. These variations are not used to assign uncertainties, but simply to better understand which effects are important.

There are several challenges we would have to resolve if we were to take the Monte Carlo variations as quantitative measures of the systematic uncertainty. First, due to limited Monte Carlo statistics, the systematic variations are ultimately convolved with

Table 10.1: Summary of mismodeling scenarios studied.

Mismodeling Studied	Variation	Change in κ (%)
ISR p_T	ISR tail weight $\times 0.25$ or $\times 0.5$	0-6
ISR Multiplicity	2σ change in ISR weight	2-4
E_T^{miss} Resolution	Double weight if $[E_T^{\text{miss}}(\text{reco}) - E_T^{\text{miss}}(\text{true})]/E_T^{\text{miss}}(\text{true}) > 0.5$	0-6
b-Tagging Efficiency	3σ variation of scale factors	0-7
QCD Cross section	$\text{QCD} \times 4$	1-7
W Cross section	$W \times 3$	2-8
$t\bar{t}W$ Cross section	$t\bar{t}W \times 3$	0-3
$t\bar{t}Z$ Cross section	$t\bar{t}Z \times 3$	1-11

statistical fluctuations and it is very difficult to distinguish significant systematic shifts from statistical artifacts of the variation procedure. Second, even if a significant systematic shift is identified, it is hard to determine an appropriate weight to assign the effect. Each Monte Carlo variation estimates a quantity that is more or less analogous to a derivative $\partial\kappa/\partial x$ for each variation x —how much κ is affected by a particular modeling parameter, x . Translating this result into an uncertainty on the final prediction requires choosing a corresponding amplitude Δx to pair with the derivative. Sometimes, there are well-motivated choices for Δx , but often the amplitude of variation has to be chosen somewhat arbitrarily. Third, even if a long list of statistically significant effects with well-motivated amplitudes is identified, it is difficult to prove that the list is complete. For these reasons, we avoid the use of Monte Carlo variations and instead turn to data-driven probes of the modeling of κ to assign uncertainties, as described in the next sections.

10.2 Understanding sources of M_J - m_T correlation

The first step to designing a suite of data-driven tests to constrain systematic mismodeling is to better understand different categories of events that pull κ away from unity. It turns out to be useful to categorize events in the search region based on how

they reach $m_T > 140 \text{ GeV}$. There are essentially four categories:

≥ 2 prompt neutrinos: True dilepton events, which are the dominant category and the main focus of the background prediction. This category consists primarily of dilepton $t\bar{t}$ events with a lost lepton or is a hadronically-decaying τ , with smaller contributions from $t\bar{t}W$ production as well as the production of $t\bar{t}Z$ with 1 charged lepton from a top quark and extra neutrinos from a Z decay.

Fake E_T^{miss} : Events for which the true value of m_T calculated with generator-level information is less than 140 GeV , but due to mismeasurement of E_T^{miss} is reconstructed with $m_T > 140 \text{ GeV}$. These events are labeled $m_T^{\text{true}} \leq 140 \text{ GeV}$.

Non-prompt neutrinos: Events for which $m_T^{\text{true}} > 140 \text{ GeV}$ due to additional E_T^{miss} from high p_T non-prompt neutrinos produced in heavy flavor decays.

Off-shell W bosons: Events for which $m_T^{\text{true}} > 140 \text{ GeV}$ due to the presence of an extremely off-shell W boson.

A comparison of M_J shapes for the different categories realized in $t\bar{t}$ is shown in Figure 10.1. The 2ℓ and $\ell\tau_{\text{had}}$ (blue and violet) categories match very well the shape of the low m_T semileptonic $t\bar{t}$ events (black) which dominate the M_J measurement in R1 and R2, as evidenced by the flat ratio for these components. These categories are the ones for which the background measurement is explicitly designed, and if all components matched this well, then κ would be equal to 1. However, the other two components, semileptonic $t\bar{t}$ events with significant fake E_T^{miss} or E_T^{miss} from non-prompt neutrinos, have a significantly harder M_J spectrum than the measurement sample dominated by low m_T semileptonic $t\bar{t}$. Introducing these events to the high m_T sample results in a correlation between M_J and m_T and pushes $\kappa > 1$. The preference for larger values of M_J in events with fake E_T^{miss} can be understood by the fact that the absolute E_T^{miss}

resolution degrades with increasing hadronic energy. Thus, events with high hadronic activity, which tend to have larger values of M_J , also migrate to the high m_T region more frequently. Similarly, the energy of b-flavor jets is highly correlated with the hadronic energy overall, so high p_T non-prompt neutrinos are more likely to be present in high activity, high M_J events. Both of these mechanisms for reaching high m_T depend strongly on the event energy scale and kinematics, unlike the main dilepton background, whose m_T efficiency is relatively insensitive to the energy scale. The extent to which κ is pulled away from 1 is largely determined by the relative size of these secondary contributions.




Figure 10.1: Comparison of M_J distributions for $t\bar{t}$ events for the various high m_T categories, compared against the low m_T $t\bar{t}$ sample which dominates the measurement in R1 and R2. Events are selected after the baseline selection, with the E_T^{miss} requirement relaxed to 100 GeV. [65]

Figure 10.2 shows the event composition over all the analysis bins, extended to $N_{\text{jets}} = 5$ and $100 \text{ GeV} < E_{\text{T}}^{\text{miss}} < 200 \text{ GeV}$. At lower values of $E_{\text{T}}^{\text{miss}}$, the contribution from $E_{\text{T}}^{\text{miss}}$ mismeasurement grows substantially. This is expected since random fluctuations in $E_{\text{T}}^{\text{miss}}$ have greater impact when there is less genuine $E_{\text{T}}^{\text{miss}}$ from neutrinos. Since this component introduces an M_J - m_{T} correlation, we then expect κ to grow larger than 1 as $E_{\text{T}}^{\text{miss}}$ is decreased less than 200 GeV. Indeed, examining κ over an extended range in Figure 10.3, we see clearly that κ grows closer to 1 as $E_{\text{T}}^{\text{miss}}$ is increased and the contribution from fake $E_{\text{T}}^{\text{miss}}$ reduces. κ is also closer to 1 for bins with more jets, which are further into the ISR dominated regime that ensures the M_J shapes agree.

CMS *Simulation Supplementary*
arXiv:1705.04673

13 TeV




Figure 10.2: Relative contribution of each of the four categories in bins of $E_{\text{T}}^{\text{miss}}$ (left to right: $E_{\text{T}}^{\text{miss}}$ 100-150 GeV, $E_{\text{T}}^{\text{miss}}$ 150-200 GeV, $E_{\text{T}}^{\text{miss}}$ 200-350 GeV, $E_{\text{T}}^{\text{miss}}$ 350-500 GeV, $E_{\text{T}}^{\text{miss}}$ 500+ GeV) and bins of N_{jets} (top to bottom: $N_{\text{jets}} = 5$, $6 \leq N_{\text{jets}} \leq 8$, and $N_{\text{jets}} \geq 9$). [65]




Figure 10.3: Values of κ along an extended E_T^{miss} range starting from 100 GeV through the $E_T^{\text{miss}} > 200\text{GeV}$ signal regions. [65]

10.3 Control samples for systematic uncertainties

Armed with the knowledge of the categories relevant for probing the understanding of κ , we define two control samples used to estimate the systematic uncertainties. The first is a dilepton sample with two reconstructed leptons, designed to directly check the central assumption of the analysis, that the M_J shapes for single and dilepton $t\bar{t}$ converge at sufficiently high jet multiplicity. The second is a single lepton control sample with $N_{\text{jets}} = 5$, which shares very similar composition to the main analysis regions and probes the modeling of the subdominant backgrounds. In each control sample, an ABCD prediction is performed analogous to the main analysis, and any non-closure is taken as a systematic uncertainty. In the case of good agreement between the data and the predictions, then the precision of the test is assigned as an uncertainty.

The uncertainties derived from the dilepton test are evaluated primarily as a function of N_{jets} . For the $N_{\text{jets}} = 5$ control sample, the uncertainties are evaluated as a function

of E_T^{miss} . The uncertainty in each signal region bin is then given approximately by:

$$(\sigma^{\text{SR}})_{E_T^{\text{miss}}, N_{\text{jets}}}^2 = (\sigma^{N_{\text{jets}}=5})_{E_T^{\text{miss}}}^2 + (\sigma^{2\ell})_{N_{\text{jets}}}^2 \quad (10.1)$$

10.3.1 Dilepton control sample

To test the fundamental assumption of the method, that the M_J shape for dileptonic $t\bar{t}$ can be measured in a single lepton $t\bar{t}$ sample, we evaluate a modified ABCD prediction in a dilepton sample. To form the dilepton sample, the two “lost lepton” regions from the main analysis, R3 and R4, are replaced with analogous regions where both leptons are reconstructed, D3 and D4, as illustrated in Figure 10.4. These include events with two fully-fledged leptons passing the signal lepton definition, as well as events with one lepton and one loose “veto track”. Events with two full leptons pass a modified version of the baseline selection with the N_{jets} requirement lowered by one to keep the number of fatjet clustering constituents constant, and $N_b = 0$ events included to improve statistics (in the dilepton sample, high $t\bar{t}$ purity is attained without any need for b-tagging). No requirement on m_T is applied to the dilepton events. The lepton + track events joining D3 and D4 are taken from the same selection as R3 and R4, simply inverting the track veto. We then perform the an ABCD in two bins of N_{jets} using R1 and R2 (unmodified), as well as D3 and D4. All N_b and E_T^{miss} regions are combined for this test, but the highest $N_b \geq 3$ and $E_T^{\text{miss}} > 500\text{GeV}$ bins are excluded to limit signal contamination. The results are shown in Table 10.2, including also a lower E_T^{miss} region which is not used to quantitatively. For both regions, note that κ is close to 1 for the dilepton component, and improves with increasing jet multiplicity, as the ISR becomes more dominant. The agreement of the data with the predictions is well within the precision of the test.




Figure 10.4: Diagram of the translation from the signal region ABCD plane (left) to the dilepton control regions (right). R1 and R2 are unchanged for the dilepton test.

Table 10.2: Background yields and signal contamination in the dilepton control sample in bins of N_{jets} [67].

$\mathcal{L} = 36 \text{ fb}^{-1}$	T1tttt(NC)	T1tttt(C)	κ	MC bkg.	Pred.	Obs.	Obs./MC	Signi.
$100 < E_T^{\text{miss}} \leq 200$ (Obs/MC = 0.87 ± 0.01)								
R1: All N_{jets}	0.03	21.02		22449.55		22506	1.00 ± 0.01	
R2: $6 \leq N_{\text{jets}} \leq 8$	0.38	6.91		7744.93		7684	0.99 ± 0.01	
R2: $N_{\text{jets}} \geq 9$	0.35	5.41		1099.01		1060	0.96 ± 0.03	
D3: All N_{jets}	0.04	8.36		3013.45		3064	1.02 ± 0.02	
D4: Low N_{jets}	0.33	2.52	$1.09^{+0.01}_{-0.01}$	1132.93	$1140.00^{+28.79}_{-28.48}$	1142	1.01 ± 0.03	0.1σ
D4: High N_{jets}	0.11	1.62	$1.00^{+0.11}_{-0.11}$	147.14	$143.94^{+16.64}_{-16.41}$	131	0.89 ± 0.12	-0.6σ
$200 < E_T^{\text{miss}} \leq 500$ (Obs/MC = 0.81 ± 0.01)								
R1: All N_{jets}	0.07	10.37		5054.89		5173	1.02 ± 0.01	
R2: $6 \leq N_{\text{jets}} \leq 8$	1.23	6.80		3086.58		2895	0.94 ± 0.02	
R2: $N_{\text{jets}} \geq 9$	0.81	9.02		372.39		325	0.87 ± 0.05	
D3: All N_{jets}	0.15	7.83		809.67		911	1.13 ± 0.04	
D4: Low N_{jets}	2.04	3.69	$1.09^{+0.01}_{-0.01}$	537.31	$554.08^{+23.03}_{-22.91}$	564	1.05 ± 0.04	0.3σ
D4: High N_{jets}	0.79	3.92	$0.99^{+0.02}_{-0.02}$	59.16	$56.77^{+3.97}_{-3.94}$	52	0.88 ± 0.12	-0.5σ

To better evaluate the modeling of κ in particular, we can also form κ^{data} by taking the analogous ratio of yields in R1,R2,D3 and D4 in data, and compare it with κ^{MC} . The observed and simulated κ factors are shown in Figure 10.5. Since the κ are from data and Monte Carlo are in agreement, the precision of the comparison in the $200 \text{ GeV} < E_T^{\text{miss}} < 500 \text{ GeV}$ bin is taken as an uncertainty: 6% for low N_{jets} bins, and 15% for high N_{jets} bins.




Figure 10.5: Comparison of κ observed in data and simulation for the dilepton control sample. The error bars represent the statistical uncertainty from the Monte Carlo (red) combined with the expected uncertainty on the data (black). The two N_{jets} bins in the $200 \text{ GeV} < E_T^{\text{miss}} < 500 \text{ GeV}$ selection are used to assign systematic uncertainties [67].

10.3.2 $N_{\text{jets}} = 5$ control sample


The $N_{\text{jets}} = 5$ control sample is defined by the same selection as the main analysis regions with N_{jets} reduced to five. This control sample shares similar composition to the main analysis, and serves to validate the modeling of the contribution of secondary processes which move κ away from 1. The $N_{\text{jets}} = 5$ test is used to measure an uncertainty as

a function of E_T^{miss} , the most important variable for controlling the secondary mechanisms for reaching high m_T . To avoid signal contamination, the region with $E_T^{\text{miss}} > 500 \text{ GeV}$ is excluded, and instead the uncertainty derived in the $350 \text{ GeV} < E_T^{\text{miss}} \leq 500 \text{ GeV}$ bin is applied to the highest E_T^{miss} bins as well. This procedure is very safe, as we find no realistic scenarios that induce mismodeling that increases with increasing E_T^{miss} . As E_T^{miss} is increased, the composition tends to be further dominated by the well-behaved two prompt neutrino category.

The results of the $N_{\text{jets}} = 5$ test are shown in Table 10.3. Like the dilepton test, the observed yields are consistent with the expectations for all bins. A comparison of the κ factors in data and Monte Carlo is shown in Figure 10.6. κ is modeled remarkably well, even in the lowest E_T^{miss} regions which have significant deviations from $\kappa = 1$. Given the good agreement, the precision of the test is taken as uncertainty for the predictions in the main analysis regions: 16% for bins with $200 \text{ GeV} < E_T^{\text{miss}} \leq 350 \text{ GeV}$ and 40% for bins with $350 \text{ GeV} < E_T^{\text{miss}} \leq 500 \text{ GeV}$ or $E_T^{\text{miss}} > 500 \text{ GeV}$.

Table 10.3: Background yields and signal contamination in the $N_{\text{jets}} = 5$ sample in bins of E_T^{miss} [67].

$\mathcal{L} = 36 \text{ fb}^{-1}$	T1tttt(NC)	T1tttt(C)	κ	MC bkg.	Pred.	Obs.	Obs./MC	Signi.
$100 < E_T^{\text{miss}} \leq 150 \text{ (Obs/MC} = 0.90 \pm 0.01\text{)}$								
R1: All N_{jets}	0.01	0.59		12867.62		12792	0.99 ± 0.01	
R2: $N_{\text{jets}} = 5$	0.03	0.10		1683.07		1697	1.01 ± 0.03	
R3: All N_{jets}	0.00	0.30		1046.26		1095	1.05 ± 0.04	
R4: $N_{\text{jets}} = 5$	0.03	0.08	$1.67^{+0.08}_{-0.08}$	229.05	$243.13^{+15.44}_{-14.99}$	242	1.06 ± 0.08	-0.0σ
$150 < E_T^{\text{miss}} \leq 200 \text{ (Obs/MC} = 0.87 \pm 0.01\text{)}$								
R1: All N_{jets}	0.00	0.57		6019.39		6040	1.00 ± 0.01	
R2: $N_{\text{jets}} = 5$	0.01	0.24		922.55		892	0.97 ± 0.03	
R3: All N_{jets}	0.02	0.44		331.80		351	1.06 ± 0.06	
R4: $N_{\text{jets}} = 5$	0.03	0.04	$1.66^{+0.09}_{-0.09}$	84.26	$85.89^{+7.20}_{-7.22}$	75	0.89 ± 0.11	-0.9σ
$200 < E_T^{\text{miss}} \leq 350 \text{ (Obs/MC} = 0.83 \pm 0.01\text{)}$								
R1: All N_{jets}	0.01	1.08		4744.06		4733	1.00 ± 0.02	
R2: $N_{\text{jets}} = 5$	0.07	0.26		1024.74		994	0.97 ± 0.03	
R3: All N_{jets}	0.05	0.79		255.97		292	1.14 ± 0.07	
R4: $N_{\text{jets}} = 5$	0.23	0.26	$1.34^{+0.06}_{-0.05}$	74.23	$82.33^{+6.64}_{-6.44}$	80	1.08 ± 0.13	-0.2σ
$350 < E_T^{\text{miss}} \leq 500 \text{ (Obs/MC} = 0.81 \pm 0.03\text{)}$								
R1: All N_{jets}	0.03	0.17		525.10		531	1.01 ± 0.04	
R2: $N_{\text{jets}} = 5$	0.05	0.03		233.26		217	0.93 ± 0.06	
R3: All N_{jets}	0.06	0.14		33.55		43	1.28 ± 0.21	
R4: $N_{\text{jets}} = 5$	0.29	0.11	$1.08^{+0.09}_{-0.08}$	16.09	$18.97^{+3.71}_{-3.46}$	17	1.06 ± 0.26	-0.3σ

Figure 10.6: Comparison of κ observed in data and simulation for the $N_{\text{jets}} = 5$ control sample. The error bars represent the statistical uncertainty from the Monte Carlo (red) combined with the expected uncertainty on the data (black). The two E_T^{miss} bins above 200 GeV are used to assign systematic uncertainties.

10.4 Validation of systematics procedure

The systematic procedure described in this chapter is only valid if any mismodeling that affects the signal regions also appears proportionately in the control samples. Extensive studies of variations in Monte Carlo have been performed, like those in Table 10.1, and in all cases the impact on κ is mirrored by a comparable effect in a control region. For most physically motivated variations, though, the change in κ is typically smaller than could be detected with the coarse precision in data.

As an example of a significant variation captured by a control region, we can examine the impact of doubling the rate of events with extremely mismeasured E_T^{miss} (double weight if $E_T^{\text{miss}}(\text{reco}) - E_T^{\text{miss}}(\text{true}) > 0.5 E_T^{\text{miss}}(\text{true})$.) Figure 10.7 compares κ in the default Monte Carlo and in pseudodata based on the same Monte Carlo under this scenario, with uncertainty assigned corresponding to a 35 fb^{-1} sample. Each $\Delta\kappa$ observed in the signal regions is accompanied by a similar shift in the $N_{\text{jets}} = 5$ control regions. Note that this shift is much more extreme than would ever be expected in real data, but gives confidence in the coverage provided by the control regions.




Figure 10.7: Changes in κ caused by doubling the number of mismeasured events in region R4. The $N_{\text{jets}} = 5$ control sample (top) captures the change as seen in the signal region sample (bottom). The error bars represent the statistical uncertainty from the Monte Carlo (red) combined with the expected uncertainty on the data (black) [67].

Table 10.4: Systematic uncertainties assigned in each signal bin, as measured in 35.9 fb^{-1} of data. Note that the uncertainty for the mid- and high- E_T^{miss} bin are the same since they both share the $N_{\text{jets}} = 5$ mid- E_T^{miss} control sample. [67]

Bin	Syst. unc.
$200 < E_T^{\text{miss}} \leq 350 \text{ GeV}, 6 \leq N_{\text{jets}} \leq 8$	17%
$200 < E_T^{\text{miss}} \leq 350 \text{ GeV}, N_{\text{jets}} \geq 9$	23%
$350 < E_T^{\text{miss}} \leq 500 \text{ GeV}, 6 \leq N_{\text{jets}} \leq 8$	41%
$350 < E_T^{\text{miss}} \leq 500 \text{ GeV}, N_{\text{jets}} \geq 9$	44%
$E_T^{\text{miss}} \geq 500 \text{ GeV}, 6 \leq N_{\text{jets}} \leq 8$	41%
$E_T^{\text{miss}} \geq 500 \text{ GeV}, N_{\text{jets}} \geq 9$	44%

10.5 Background systematics summary

Based on the agreement observed in the control samples, uncertainties were derived as a function of N_{jets} (from the dilepton sample) and E_T^{miss} (from the $N_{\text{jets}} = 5$ sample). Uncertainties on a particular N_{jets} bin are treated as completely correlated across all bins of N_b and E_T^{miss} , but independent of uncertainties assigned to other N_{jets} bins. Similarly, uncertainties assigned to E_T^{miss} bins are correlated across N_{jets} and N_b , but independent of other bins in E_T^{miss} . Ignoring the complication of the correlations, we can approximate the total systematic uncertainty for each bin by simply adding the two relevant uncertainties in quadrature. The resulting approximate uncertainties are shown in Table 10.4.

Chapter 11

Systematic Uncertainties, Signal Model

Without the possibility of a dedicated control sample in data, we derive systematic uncertainties for the signal efficiency and acceptance based on variations of the signal Monte Carlo simulation (though many of the variations are motivated by measurements in data as well). The uncertainties fall in to two categories: experimental uncertainties (mostly relevant for the efficiency) and theoretical uncertainties (mostly relevant for the acceptance and cross section).

The most important experimental uncertainties in general are related to the efficiencies for b-tagging and lepton reconstruction. These efficiencies are measured in dedicated control samples in data (e.g. tag and probe with leptons from Z boson decays, and large statistics $t\bar{t}$ samples for b-tagging), and scale factors were developed to reweight events in the simulation to improve the agreement with data. The uncertainty in the signal efficiency due to uncertainty in the individual object efficiencies is assessed simply by varying these scale factors within their own uncertainties. Additionally, since the signal Monte Carlo doesn't use the full simulation of the detector, but only a parametrized "FastSim",

larger uncertainties are warranted, which were assessed based on a comparison of samples (e.g. $t\bar{t}$) that are produced both in full and fast simulation.

Additionally, for the case of compressed models with softer E_T^{miss} spectra, the efficiency is sensitive to the E_T^{miss} resolution. Thus, uncertainty in the jet energy corrections, and in particular the jet energy corrections for FastSim, can have a large impact on the overall signal efficiency.

The uncertainty in the integrated luminosity is 2.6% [83]. To assess the uncertainty related to the impact of pileup, the signal efficiency is compared between signal samples with low and high pileup. This comparison is dominated by statistical fluctuations, so uncertainties of 10% and 15% were assigned to the low $N_{\text{jets}}/E_T^{\text{miss}}$ and high $N_{\text{jets}}/E_T^{\text{miss}}$ bins, respectively, which are found to bound all observed variations. These values were ultimately chosen for internal political expediency and are now considered to be significantly larger than is truly warranted.

The theoretical aspects contributing to the systematic uncertainties in the signal acceptance are the renormalization and factorization scales, as well the modeling of the ISR. Since the uncertainty in the cross section (from QCD scales as well as PDFs) is factorized and reported in the final results separately, the variations in the renormalization and factorization scales are held at constant cross section and used simply to probe the impact of the scales on the event kinematics. As result, their variation a very small effect (less than 1% in essentially all cases). Since the ISR spectrum in simulation is corrected based on a dedicated measurement in a $t\bar{t}$ dominated sample, an uncertainty is also assessed based on variation of the ISR weights within their uncertainty. The ISR is typically not important for non-compressed models, but can be significant for compressed models that rely on additional boost to fall into acceptance.

The typical magnitudes for these uncertainties are summarized in Table 11.1. The detailed uncertainties for the efficiencies in all 18 signal region bins are shown in Table 11.2

Table 11.1: Typical values of the signal-related systematic uncertainties. Uncertainties due to a particular source are treated as fully correlated between bins, while uncertainties due to different sources are treated as uncorrelated. [67]

Source	Fractional uncertainty [%]
Lepton efficiency	3–4
FastSim lepton efficiency	3–4
Trigger efficiency	1
b-tag efficiency	1–5
FastSim b-tag efficiency	1–28
Mistag efficiency	1–2
FastSim mistag efficiency	1–8
FastSim MET	1–13
Jet energy corrections	1–24
QCD scales	< 1
ISR	1–16
Jet ID	1
Pile up	10–15
Integrated luminosity	3

for the compressed and non-compressed benchmark models, as an example. Corresponding tables are generated based on variations performed for each model in the $m_{\tilde{g}}-m_{\tilde{\chi}_1^0}$ plane. Each source of systematic uncertainty is treated as correlated across all bins, but different sources are treated as uncorrelated from each other. The correlations are particularly important for sources of uncertainty like b-tagging, where a change in the efficiency results in a migration among bins much more than it affects the total efficiency.

Table 11.2: Summary of the signal systematic uncertainties. The sign of each uncertainty corresponds to the impact on the yields of the upward variation of the uncertainty source. A negative sign indicates an anti-correlation with respect to bins with a positive sign. Uncertainties due to a particular source are considered fully correlated between bins, while uncertainties due to different sources are considered uncorrelated. [67]

Uncertainty [%]	$200 < E_T^{\text{miss}} \leq 350$				$350 < E_T^{\text{miss}} \leq 500$				$E_T^{\text{miss}} > 500$			
	$6 \leq N_{\text{jets}} \leq 8$		$N_{\text{jets}} \geq 9$		$6 \leq N_{\text{jets}} \leq 8$		$N_{\text{jets}} \geq 9$		$6 \leq N_{\text{jets}} \leq 8$		$N_{\text{jets}} \geq 9$	
	$N_b = 1$	$N_b = 2$	$N_b \geq 3$	$N_b = 1$	$N_b = 2$	$N_b \geq 3$	$N_b = 1$	$N_b = 2$	$N_b \geq 3$	$N_b = 1$	$N_b = 2$	$N_b \geq 3$
Signal model: T1tttt(1800,100)												
Lepton efficiency	3	3	4	3	4	4	4	4	4	3	4	4
FastSim lepton efficiency	4	4	4	3	4	4	4	4	4	4	4	4
Trigger efficiency	1	1	1	1	1	1	1	1	1	1	1	1
b-tag efficiency	-4	<1	4	-5	3	<1	5	4	-4	<1	5	4
FastSim b-tag efficiency	-16	12	5	5	4	2	10	4	-1	2	7	6
Mistag efficiency	-1	<1	1	1	<1	<1	<1	<1	-1	<1	1	1
FastSim mistag efficiency	3	3	3	4	3	3	3	4	3	3	3	3
FastSim MET	-5	-7	1	7	2	-1	1	6	<1	<1	<1	1
Jet energy corrections	-5	2	-7	4	5	8	-4	-2	-1	-1	-2	8
QCD scales	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1
ISR	-2	-4	-3	5	8	4	-2	-3	-2	-3	8	7
Jet ID	1	1	1	1	1	1	1	1	1	1	1	1
Pile up	10	10	10	15	15	15	10	10	15	15	15	15
Integrated luminosity	3	3	3	3	3	3	3	3	3	3	3	3
Signal model: T1tttt(1400,1000)												
Lepton efficiency	3	4	4	3	4	4	3	4	4	3	3	3
FastSim lepton efficiency	3	3	3	3	3	3	3	3	3	3	3	3
Trigger efficiency	1	1	1	1	1	1	1	1	1	1	1	1
b-tag efficiency	-3	<1	3	-5	-1	4	-3	<1	-3	<1	3	3
FastSim b-tag efficiency	-1	-1	1	-1	-1	2	-1	<1	<1	-1	-1	28
Mistag efficiency	-1	<1	1	-1	<1	<1	-1	<1	-1	<1	-1	2
FastSim mistag efficiency	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1	8
FastSim MET	<1	-8	11	-4	4	<1	7	10	-3	-3	4	-3
Jet energy corrections	-7	-6	-4	18	13	8	3	11	-6	8	6	7
QCD scales	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1	<1
ISR	1	-1	<1	11	10	8	5	2	3	6	2	13
Jet ID	1	1	1	1	1	1	1	1	1	1	1	1
Pile up	10	10	10	15	15	15	10	10	15	15	15	15
Integrated luminosity	3	3	3	3	3	3	3	3	3	3	3	3

Chapter 12

Results

With the development of the background measurement and the systematic uncertainties complete, we arrive at the final step: unblinding the signal regions and confronting our predictions with the data. Section 12.1 shows distributions over the complete dataset as well as detailed tables of the predictions, results, and possible signal contributions. Section 12.2 shows the interpretation of the results in the context of our supersymmetric signal model.

12.1 Results

Before diving into the bin-by-bin results, it is useful to get an overview of the data and basic check on the agreement with the background prediction. Figure 12.1 compares M_J shapes observed in the high m_T regions (points) with the shapes taken from the low m_T regions (histograms) for a variety of E_T^{miss} and N_b selections, integrating over N_{jets} bins. These distributions are binned much more finely than used in the actual background prediction or interpretation, but are a clean way to check the basic assumption of M_J - m_T independence in the main analysis regions. In particular, the left column shows the

shapes in the low E_T^{miss} bins, where there are adequate statistics for a detailed comparison, and the shapes do agree well overall. The background prediction could almost be read directly off the plot by simply integrating the histograms in the high M_J regions, except for a key missing a piece: the κ factors. Since the κ factors are slightly larger than 1, including κ would effectively scale up the histograms in the high M_J regions, and scale down the histograms in the central M_J regions. Based on the trends in the ratio plots, the simulated correction factors are indeed necessary and would significantly improve the agreement.

Meanwhile, the right column of Figure 12.1 compares the M_J shapes in the highest E_T^{miss} bins, where a contribution from the signal would be most prominent. Unfortunately, there is no sign of an excess at large M_J , though precise analysis will await the full detailed predictions.

Another overview can be seen in Figure 12.2, which shows 2D distributions of the data in M_J and m_T for the three different E_T^{miss} bins, compared with the expected shapes from simulation, and an rough approximation of the events expected from the non-compressed benchmark model. This view confirms that the data follow reasonable shapes in 2D, and offers a nice handle to identify outlier events. All of the events with extreme values of both m_T and M_J have been examined in detail and do not reveal any unanticipated features.

Table 12.1, 12.2, and 12.3 summarize the observed event yields, the fitted backgrounds, and the expected signal yields for the two T1tttt benchmark model points in all signal bins. The background shown in the table is obtained from the predictive fit, which uses only the yields in regions R1, R2, and R3 and is model independent. The observed yields are consistent with the background predictions in all 18 signal bins within 2 standard deviations, and most bins are consistent within 1 s.d. The R4 bins with $E_T^{\text{miss}} > 500\text{GeV}$ show what may appear to be excess in the observed yields relative to




Figure 12.1: M_J distributions observed in data for $200 < E_T^{\text{miss}} \leq 350$ GeV (left column), $350 < E_T^{\text{miss}} \leq 500$ GeV (center column) and $E_T^{\text{miss}} > 500$ GeV (right column); and $N_b = 1$ (top row) and $N_b \geq 2$ (bottom row) in the 1ℓ data for low and high m_T regions. In each plot, the data at low m_T (histograms) have been normalized to the yield observed at high m_T (points) to facilitate shape comparison. The data are integrated over $N_{\text{jets}} \geq 6$. The two benchmark T1tttt models are shown in the solid and dashed red histograms. [67]

the background predictions. However, the predictions in these bins are all proportional to a single R3 yield, integrated over all bins of N_{jets} and N_{b} , which introduces a substantial correlation. Given the poor precision of this bin (only two events observed!) and accounting for the correlation, the significance of the excess in these six bins is only 1.9 s.d., mostly driven by the low N_{b} bins. In fact, under the background-only global fit, which adjusts the observed yields according to their uncertainties, the best-fit for the R3 yield with $E_{\text{T}}^{\text{miss}} > 500\text{GeV}$ is approximately 5 events, nearly erasing the excess altogether. It seems very plausible that the majority of the 1.9 s.d. deviation is caused by a simple fluctuation from a mean rate of 5 to an observation of 2 events.




Figure 12.2: Two-dimensional distributions for data and simulated event samples integrated over the N_{jets} and $N_b \geq 2$, shown for each E_T^{miss} bin: 200-350 GeV (top), 200-350 GeV (middle) and ≥ 500 GeV (bottom). The black dots are the data; the colored histogram is the total simulated background, normalized to the data; and the red dots are a particular signal sample drawn from the expected distribution for gluino pair production in the T1tttt model with $m_{\tilde{g}} = 1800\text{GeV}$ and $m_{\tilde{\chi}_1^0} = 100\text{GeV}$ for 35.9 fb^{-1} . Overflow events are shown on the edges of the plot. [67]

Table 12.1: Observed and predicted event yields for the signal regions (R4) and background regions (R1–R3) with $E_T^{\text{miss}} \leq 350$ GeV in data (35.9 fb^{-1}). Expected yields for the two SUSY T1tttt benchmark scenarios (1800, 100) and (1400,1000) are also given. The prediction is the result of the predictive fit. The uncertainties on the prediction account for the available statistics in the data control samples, the precision of κ from MC, and the systematic uncertainties assessed from control samples in data. [67]

Bin	T1tttt (1800,100)	T1tttt (1400,1000)	κ	Pred.	Obs.
200 GeV < $E_{\text{T}}^{\text{miss}} \leq 350$ GeV					
R1: all $N_{\text{jets}}, N_{\text{b}}$	0.0	9.1			4761
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 1$	0.1	1.2			1353
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} = 1$	0.1	1.0			127
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 2$	0.3	2.1			910
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} = 2$	0.1	2.0			118
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} \geq 3$	0.3	1.7			214
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} \geq 3$	0.2	3.1			33
R3: all $N_{\text{jets}}, N_{\text{b}}$	0.1	12.5			247
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 1$	0.4	1.9	$1.2 \pm 0.0 \pm 0.2$	84.6 ± 14.3	106
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} = 1$	0.2	1.6	$1.0 \pm 0.1 \pm 0.2$	6.5 ± 1.5	11
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 2$	0.6	3.0	$1.2 \pm 0.0 \pm 0.2$	55.1 ± 9.3	75
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} = 2$	0.3	2.1	$1.2 \pm 0.1 \pm 0.3$	7.6 ± 1.9	11
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} \geq 3$	0.6	2.2	$1.5 \pm 0.1 \pm 0.2$	16.4 ± 3.0	16
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} \geq 3$	0.4	3.1	$1.4 \pm 0.1 \pm 0.3$	2.3 ± 0.7	2

Table 12.2: Observed and predicted event yields for the signal regions (R4) and background regions (R1–R3) with $350 < E_T^{\text{miss}} \leq 500$ GeV in data (35.9 fb^{-1}). Expected yields for the two SUSY T1tttt benchmark scenarios (1800, 100) and (1400, 1000) are also given. The prediction is the result of the predictive fit. The uncertainties on the prediction account for the available statistics in the data control samples, the precision of κ from MC, and the systematic uncertainties assessed from control samples in data. [67]

Bin	T1tttt (1800,100)	T1tttt (1400,1000)	κ	Pred.	Obs.
350 GeV < $E_{\text{T}}^{\text{miss}} \leq 500$ GeV					
R1: all $N_{\text{jets}}, N_{\text{b}}$	0.0	1.0			412
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 1$	0.2	0.4			226
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} = 1$	0.1	0.7			15
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 2$	0.2	0.7			155
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} = 2$	0.1	1.0			25
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} \geq 3$	0.2	0.5			37
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} \geq 3$	0.2	1.0			7
R3: all $N_{\text{jets}}, N_{\text{b}}$	0.1	2.3			32
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 1$	0.7	1.1	$1.0 \pm 0.1 \pm 0.3$	17.4 ± 6.6	25
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} = 1$	0.3	1.0	$1.1 \pm 0.1 \pm 0.4$	1.3 ± 0.6	2
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 2$	0.9	1.3	$1.1 \pm 0.1 \pm 0.4$	13.7 ± 5.3	10
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} = 2$	0.5	1.1	$0.8 \pm 0.1 \pm 0.3$	1.6 ± 0.8	2
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} \geq 3$	0.8	0.9	$1.3 \pm 0.1 \pm 0.4$	3.8 ± 1.6	1
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} \geq 3$	0.7	2.1	$1.2 \pm 0.2 \pm 0.4$	0.6 ± 0.4	0

Table 12.3: Observed and predicted event yields using for the signal regions (R4) and background regions (R1–R3) with $E_T^{\text{miss}} > 500$ GeV in data (35.9 fb^{-1}). Expected yields for the two SUSY T1tttt benchmark scenarios (1800, 100) and (1400,1000) are also given. The prediction is based on the predictive fit. The uncertainties on the prediction account for the available statistics in the data control samples and the precision of κ from MC. [67]

Bin	T1tttt (1800,100)	T1tttt (1400,1000)	κ	Pred.	Obs.
$E_T^{\text{miss}} > 500 \text{ GeV}$					
R1: all $N_{\text{jets}}, N_{\text{b}}$	0.1	0.4			74
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 1$	0.5	0.4			71
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} = 1$	0.2	0.4			8
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 2$	0.6	0.4			32
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} = 2$	0.3	0.7			5
R2: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} \geq 3$	0.5	0.3			10
R2: $N_{\text{jets}} \geq 9, N_{\text{b}} \geq 3$	0.6	0.8			1
R3: all $N_{\text{jets}}, N_{\text{b}}$	0.3	0.5			2
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 1$	2.5	0.6	$1.0 \pm 0.1 \pm 0.3$	1.9 ± 1.5	8
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} = 1$	1.0	0.7	$1.0 \pm 0.2 \pm 0.4$	0.2 ± 0.2	2
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} = 2$	3.6	1.0	$1.0 \pm 0.1 \pm 0.3$	0.9 ± 0.7	4
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} = 2$	1.8	1.2	$1.0 \pm 0.3 \pm 0.3$	0.1 ± 0.1	0
R4: $6 \leq N_{\text{jets}} \leq 8, N_{\text{b}} \geq 3$	3.2	0.4	$1.5 \pm 0.3 \pm 0.5$	0.4 ± 0.4	1
R4: $N_{\text{jets}} \geq 9, N_{\text{b}} \geq 3$	2.3	1.7	$3.1 \pm 1.0 \pm 1.1$	0.1 ± 0.1	0

12.2 Interpretations

Given the absence of any significant excess, the results are interpreted in the context of exclusion limits on the production cross section for T1tttt and T5tttt models as a function of $m_{\tilde{g}}$ and $m_{\tilde{\chi}_1^0}$.

A 95% confidence level (CL) upper limit on the production cross section is estimated using the modified frequentist CL_S method [84, 85, 86], with a one-sided profile likelihood ratio test statistic. For this test, we perform the global fit under the background-only and background-plus-signal (r floating) hypotheses. The statistical uncertainties from data counts in the control regions are modeled by the Poisson terms in the likelihood Eq. (9.9). All systematic uncertainties are multiplicative and are treated as log-normal distributions. Exclusion limits are also estimated for $\pm 1\sigma$ variations on the production cross section based on the NLO+NLL calculation [87].

Figure 12.3 shows the corresponding excluded region at a 95% CL for the T1tttt model in the $m_{\tilde{g}} - m_{\tilde{\chi}_1^0}$ plane with 35.9fb^{-1} of data. At low $\tilde{\chi}_1^0$ mass we exclude gluinos with masses of up to 1900 GeV. The highest limit on the $\tilde{\chi}_1^0$ mass is 1175 GeV, attained for $m_{\tilde{g}}$ of approximately 1750 GeV. The observed limits differ by just over 1σ from the expected limits at maximum across the full scan range. In the “corner” region with moderately compressed spectra, near $m_{\tilde{\chi}_1^0} \approx 1000$ GeV and $m_{\tilde{g}} \approx 1800$ GeV, the extended observed reach relative to the expected limits is largely driven by the small deficit observed in the moderate $E_{\text{T}}^{\text{miss}}$ bins with $N_{\text{b}} \geq 3$.

An interpretation is also performed for the more natural scenario, T5tttt, with gluinos decaying through on-shell stop quarks. Since fully probing the mass space would require an impractical three-dimensional scan in $m_{\tilde{g}}$, $m_{\tilde{t}}$, and $m_{\tilde{\chi}_1^0}$, instead one 2D slice of the 3D space is chosen that captures the most distinguishing features from the simplified three-body decay model. For the each point in the scan in $m_{\tilde{g}}$ and $m_{\tilde{\chi}_1^0}$, $m_{\tilde{t}}$ is chosen so

that $m_{\tilde{t}} - m_{\tilde{\chi}_1^0} = 175 \text{ GeV}$, so that none of the mass energy of the stop quark is converted to kinetic energy in the final state. This choice allows for the most perverse case, where $m_{\tilde{t}} \rightarrow m_t$ and $m_{\tilde{\chi}_1^0} \rightarrow 0$. In this case, the $\tilde{\chi}_1^0$ receives basically none of the $\tilde{t}$ momentum (the $\tilde{\chi}_1^0$ momentum is essentially just $\gamma_{\tilde{t}} m_{\tilde{\chi}_1^0}$, which also approaches 0.) This “stealth” spectrum erases the powerful E_T^{miss} signature and cripples the signal efficiency for the m_T requirement. In fact, the signal acceptance for this case relies almost entirely on events with two leptonic top decays and one lost lepton (which are further suppressed by the track veto!). The exclusions for this choice of T5tttt are compared with the three-body T1tttt exclusions in Figure 12.4.

12.2.1 Aggregate bins

To facilitate reinterpretation of the results in the context of other theoretical models, we provide four aggregated binning options for a simplified analysis in which only a single analysis bin is used for each of the four ABCD regions. For each option, the single bin is created by summing over a subset of the 18 bins in E_T^{miss} , N_{jets} , and N_b . In the low- M_J regions R1 and R3, only the E_T^{miss} requirement is applied, while in the high- M_J regions R2 and R4, all three requirements are applied. This procedure is analogous to the standard background estimation, in which a separate ABCD plane is used for each E_T^{miss} bin (i.e., the E_T^{miss} requirement is applied to all four regions), but, to improve the statistical power of the regions R1 and R3, the N_{jets} and N_b bins are used only in the high- M_J regions R2 and R4. The four options, described in Table 12.4, extend down to $E_T^{\text{miss}} = 200 \text{ GeV}$, $N_{\text{jets}} = 6$, and $N_b = 1$, allowing sensitivity to a variety of potential signal models. The aggregate bins are designed with at most two stringent requirements and one loose requirement, as a compromise between inclusivity and sensitivity.




Figure 12.3: Interpretation of results in the T1tttt model. The colored regions show the upper limits (95% CL) on the production cross section for $pp \rightarrow \tilde{g}\tilde{g}, \tilde{g} \rightarrow t\bar{t}\tilde{\chi}_1^0$ in the $m_{\tilde{g}}-m_{\tilde{\chi}_1^0}$ plane. The curves show the expected and observed limits on the corresponding SUSY particle masses obtained by comparing the excluded cross section with theoretical cross sections. [67]

Table 12.4: Observed events yields and mean background yields from a predictive fit in four aggregate search bins. In all four cases, the predicted yields refer to R4 with the usual requirements of $m_T > 140$ GeV and $M_J > 400$ GeV applied in addition to the baseline selection. Unlike the finely binned approach where all 18 background predictions are found simultaneously, the four aggregate bin predictions here are computed separately and may be highly correlated due to overlapping definitions. [67]

Bin	T1tttt (1800,100)	T1tttt (1400,1000)	κ	Pred.	Obs.
$E_T^{\text{miss}} > 200$ GeV, $N_{\text{jets}} \geq 9$, $N_b \geq 3$	3.4	6.9	1.4 ± 0.3	3.1 ± 0.8	2
$E_T^{\text{miss}} > 350$ GeV, $N_{\text{jets}} \geq 9$, $N_b \geq 2$	5.3	6.2	1.0 ± 0.4	2.7 ± 1.2	2
$E_T^{\text{miss}} > 500$ GeV, $N_{\text{jets}} \geq 6$, $N_b \geq 3$	5.4	2.1	1.7 ± 0.6	0.5 ± 0.4	1
$E_T^{\text{miss}} > 500$ GeV, $N_{\text{jets}} \geq 9$, $N_b \geq 1$	5.1	3.6	1.2 ± 0.4	0.4 ± 0.4	2




Figure 12.4: Excluded region (95% CL), shown in blue, in the $m_{\tilde{g}}-m_{\tilde{\chi}_1^0}$ plane for a model combining T5tttt, gluino pair production, followed by gluino decay to an on-shell top squark. The top squarks decay via the two-body process $\tilde{t} \rightarrow t\tilde{\chi}_1^0$. The neutralino and top squark masses are related by the constraint $m_{\tilde{t}} = m_{\tilde{\chi}_1^0} + 175\text{GeV}$. For comparison, the excluded region (95% CL) from Fig. 12.3 for the T1tttt model, which has three-body gluino decay, is shown in red. The small difference between the two boundary curves shows that the limits for the scenarios with two-body gluino decay have only a weak dependence on the top squark mass. [67]

Chapter 13

Conclusions

The exclusions on gluino mediated stop production presented in Section 12.2 reach $m_{\tilde{g}} = 1.9 \text{ TeV}$ for neutralinos up to $m_{\tilde{\chi}_1^0} = 1 \text{ TeV}$. We are now faced with significant tension between the exclusions and the expectations based on naturalness, that $m_{\tilde{g}}$ should be less than 2 TeV and the lightest neutralino certainly less than 1 TeV , being partially composed of $\mathcal{O}(100 \text{ GeV})$ higgsinos.

There are many good proposals that add complexity or relax assumptions to keep natural supersymmetry alive for a little while longer. The exclusions can be eroded if there are gluino decay channels that reduce the $\tilde{t}$ branching fraction, or if there are additional electroweakinos that siphon energy away from the $\tilde{\chi}_1^0$ through compressed cascade decays. If R-parity conservation is abandoned, we lose our stable neutralino and are faced with a more challenging final state relying purely on visible decay products. It seems unlikely, though, that these caveats could diminish the reach as much as the stealth stop case already considered does, which takes $m_{\tilde{t}} \approx m_t$ and $m_{\tilde{\chi}_1^0} \approx 0$, essentially removing the extra $E_{\text{T}}^{\text{miss}}$ signature altogether. The exclusion obtained for this extremely difficult case reaches $m_{\tilde{g}} = 1.6 \text{ TeV}$, which itself is remarkable and still approaching tension with the expectations from naturalness. Comparable results are obtained for gluino searches

in the context of natural R-parity violating supersymmetry, as well [88].

Even taking these caveats into consideration, this search together with many others like it constitutes a truly meaningful negative result. Though it is still important to continue the search for classic supersymmetry signatures, we are starting to feel a push to abandon the mainstream, long-held belief in electroweak scale supersymmetry and depart from the well-defined path for discovery. This is a bittersweet moment; though we have already exhausted the possibilities for easy discoveries granted by the upgrade in LHC energy, we have the opportunity to discover something perhaps altogether unanticipated and much more revolutionary. Though it's likely to be a long journey ahead, it is truly an exciting time to continue the search for new physics at the LHC.

Bibliography

- [1] CMS Collaboration. Measurement of the properties of a Higgs boson in the four-lepton final state. *Phys. Rev. D*, 89(CMS-HIG-13-002. CMS-HIG-13-002. CERN-PH-EP-2013-220):092007. 45 p, December 2013. Comments: Published in *Phys. Rev. D*.
- [2] Stephen P. Martin. A Supersymmetry Primer. *arXiv:hep-ph/9709356*, 18:1–98, July 1998.
- [3] Michele Papucci, Joshua T. Ruderman, and Andreas Weiler. Natural SUSY Endures. *Journal of High Energy Physics*, 2012(9), September 2012.
- [4] Nathaniel Craig. The State of Supersymmetry after Run I of the LHC. *arXiv:1309.0528 [hep-ex, physics:hep-ph]*, September 2013.
- [5] W. James Stirling. Progress in Parton Distribution Functions and implications for LHC. In *Proceedings, 38th International Symposium on Multiparticle Dynamics (ISMD 2008): Hamburg, Germany, September 15-20, 2008*, pages 46–55, 2009.
- [6] Stephen Myers. THE LARGE HADRON COLLIDER 20082013. *International Journal of Modern Physics A*, 28(25):1330035, October 2013.
- [7] Lyndon Evans and Philip Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001, 2008.
- [8] Claudio Campagnari and Melissa Franklin. The discovery of the top quark. *Rev. Mod. Phys.*, 69(1):75, 1997.
- [9] CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, 2008.
- [10] Daniel Froidevaux and Paris Sphicas. General-Purpose Detectors for the Large Hadron Collider. *Annual Review of Nuclear and Particle Science*, 56(1):375–440, November 2006.
- [11] The CMS Collaboration. Description and performance of track and primary-vertex reconstruction with the CMS tracker. *Journal of Instrumentation*, 9(10):P10009–P10009, October 2014.

- [12] Michael Hoch. CMS pixel detector. CMS Collection., December 2014.
- [13] CMS Collaboration. CMS Tracker showing silicon strip detectors in the barrel module.
- [14] CMS Collaboration. Commissioning and Performance of the CMS Pixel Tracker with Cosmic Ray Muons. *Journal of Instrumentation*, 5(03):T03007–T03007, March 2010.
- [15] David Cockerill. Introduction to Calorimeters, May 2016.
- [16] A.A Annenkov, M.V Korzhik, and P Lecoq. Lead tungstate scintillation material. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 490(1-2):30–50, September 2002.
- [17] D J Graham and C Seez. Simulation of Longitudinal Light Collection Uniformity in PbWO₄ Crystals.
- [18] CMS Collaboration. Performance and operation of the CMS electromagnetic calorimeter. *Journal of Instrumentation*, 5(03):T03010–T03010, March 2010.
- [19] Federico De Guio and the CMS collaboration. Performance of the CMS electromagnetic calorimeter and its role in the hunt for the Higgs boson in the two-photon channel. *Journal of Physics: Conference Series*, 455:012028, August 2013.
- [20] Andrea Benaglia. The CMS ECAL performance with examples. *Journal of Instrumentation*, 9(02):C02008–C02008, February 2014.
- [21] CMS Collaboration. Energy calibration and resolution of the CMS electromagnetic calorimeter in pp collisions at $\sqrt{s} = 7$ TeV. *Journal of Instrumentation*, 8(09):P09009–P09009, September 2013.
- [22] CMS Collaboration. Observation of a new boson with mass near 125 GeV in pp collisions at $\sqrt{s} = 7$ and 8 TeV. *JHEP*, 06:081, 2013.
- [23] Sehwook Lee, Michele Livan, and Richard Wigmans. On the limits of the hadronic energy resolution of calorimeters. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 882:148–157, 2018.
- [24] P Cushman, A Heering, and A Ronzhin. Custom HPD readout for the CMS HCAL. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 442(1-3):289–294, March 2000.
- [25] Tom Zimmerman and James R Hoff. The Design of a Charge Integrating, Modified Floating Point ADC Chip.

- [26] CMS Collaboration. Calibration of hadron calorimeter using isolated charged hadrons. May 2017. CMS-DP-2017-016.
- [27] CMS Collaboration. *The CMS Hadron Calorimeter Project: Technical Design Report*. Technical Design Report CMS. CERN, Geneva, 1997.
- [28] Nural Akchurin and Richard Wigmans. Quartz fibers as active elements in detectors for particle physics. *Review of Scientific Instruments*, 74(6):2955–2972, June 2003.
- [29] A Penzo, Y Onel, and the CMS Collaboration. The CMS-HF quartz fiber calorimeters. *Journal of Physics: Conference Series*, 160:012014, April 2009.
- [30] Artur Lobanov and the CMS Collaboration. The CMS Outer HCAL SiPM Upgrade. *Journal of Physics: Conference Series*, 587:012005, February 2015.
- [31] S. Abdullin and others. The CMS barrel calorimeter response to particle beams from 2-GeV/c to 350-GeV/c. *Eur. Phys. J.*, C60:359–373, 2009. [Erratum: *Eur. Phys. J.*C61,353(2009)].
- [32] J. Mans et al. CMS Technical Design Report for the Phase 1 Upgrade of the Hadron Calorimeter. Technical Report FERMILAB-DESIGN-2012-01, 1151651, September 2012.
- [33] CMS Collaboration. Measurement of Radiation Damage of HCAL Endcap Calorimeter (HE) using Laser data collected in Run 1 and Run 2. CMS Detector Performance Note CMS-DP-2016-052, July 2016.
- [34] Nadja Strobbe. The upgrade of the CMS hadron calorimeter with silicon photomultipliers. *Journal of Instrumentation*, 12(01):C01080–C01080, January 2017.
- [35] CMS Collaboration. Results related to the Phase1 HE upgrade. Technical Report CMS-DP-2018-019, May 2018.
- [36] Ingo Bloch and Jason Gilmore. The CMS Muon System, June 2008.
- [37] Giacomo Luca Bruno. *The RPC Detectors and the Muon System for the CMS Experiment at the LHC*. PhD thesis, Universit degli Studi di Pavia, 2001.
- [38] CMS Collaboration. The CMS trigger system. *Journal of Instrumentation*, 12(01):P01020, 2017.
- [39] CMS Trigger and Data Acquisition Group. The CMS High Level trigger. *Eur. Phys. J.*, C46:605–667, 2006.
- [40] CMS Collaboration. CMS Tracking POG Performance Plots 2016. Technical report.

- [41] K. Rose. Deterministic annealing for clustering, compression, classification, regression, and related optimization problems. *Proceedings of the IEEE*, 86(11):2210–2239, Nov./1998.
- [42] CMS Collaboration. Particle-flow reconstruction and global event description with the CMS detector. *Journal of Instrumentation*, 12(10):P10003, 2017.
- [43] Patrick Janot. Particle-Flow Event Reconstruction from LEP to LHC, February 2011.
- [44] V Daniel Elvira. Measurement of the Pion Energy Response and Resolution in the Hadronic Barrel Calorimeter using CMS HCAL Test Beam 2002 Data.
- [45] Matteo Cacciari, Gavin P Salam, and Gregory Soyez. The anti- k_t jet clustering algorithm. *Journal of High Energy Physics*, 2008(04):063–063, April 2008.
- [46] CMS Collaboration. Jet energy scale and resolution performances with 13 TeV data. June 2016.
- [47] CMS Collaboration. Jet energy scale and resolution in the CMS experiment in pp collisions at 8 TeV. *Journal of Instrumentation*, 12(02):P02014, 2017.
- [48] Nazar Bartosik. B-tagging diagram, June 2016.
- [49] CMS Collaboration. Identification of b-quark jets with the CMS experiment. *Journal of Instrumentation*, 8(04):P04013, 2013.
- [50] Identification of b quark jets at the CMS Experiment in the LHC Run 2. Technical Report CMS-PAS-BTV-15-001, CERN, Geneva, 2016.
- [51] CMS Collaboration. CMS SUSY Results: Objects Efficiency. Technical report, March 2018.
- [52] C. G Lester and D. J Summers. Measuring masses of semi-invisibly decaying particle pairs produced at hadron colliders. *Physics Letters B*, 463(1):99–103, September 1999.
- [53] Nima Arkani-Hamed, Philip Schuster, Natalia Toro, Jesse Thaler, Lian-Tao Wang, Bruce Knuteson, and Stephen Mrenna. MARMOSET: The path from LHC data to the new standard model via on-shell effective theories. 2007.
- [54] J. Alwall, P. Schuster, and N. Toro. Simplified models for a first characterization of new physics at the LHC. *Phys. Rev. D*, 79:075020, 2009.
- [55] Johan Alwall, My-Phuong Le, Mariangela Lisanti, and Jay G. Wacker. Model-independent jets plus missing energy searches. *Phys. Rev. D*, 79:015005, 2009.

- [56] Daniele Alves, Nima Arkani-Hamed, Sanjay Arora, Yang Bai, Matthew Baumgart, Joshua Berger, Matthew Buckley, Bart Butler, Spencer Chang, Hsin-Chia Cheng, Clifford Cheung, R. Sekhar Chivukula, Won Sang Cho, Randy Cotta, Mariarosaria D’Alfonso, and others. Simplified models for LHC new physics searches. *J. Phys. G*, 39:105005, 2012.
- [57] ATLAS Collaboration. Search for supersymmetry in final states with jets, missing transverse momentum and one isolated lepton in $\sqrt{s} = 7$ TeV pp collisions using 1 fb^{-1} of ATLAS data. *Phys. Rev. D*, 85:012006, 2012. [Erratum: 10.1103/PhysRevD.87.099903].
- [58] ATLAS Collaboration. Search for squarks and gluinos in events with isolated leptons, jets and missing transverse momentum at $\sqrt{s} = 8$ TeV with the ATLAS detector. *JHEP*, 04:116, 2015.
- [59] ATLAS Collaboration. Search for pair production of gluinos decaying via stop and sbottom in events with b-jets and large missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev.*, D94(3):032003, 2016.
- [60] ATLAS Collaboration. Search for gluinos in events with an isolated lepton, jets and missing transverse momentum at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Eur. Phys. J. C*, 76:565, 2016.
- [61] CMS Collaboration. Search for supersymmetry in pp collisions at $\sqrt{s} = 7$ TeV in events with a single lepton, jets, and missing transverse momentum. *JHEP*, 08:156, 2011.
- [62] CMS Collaboration. Search for supersymmetry in pp collisions at $\sqrt{s} = 8$ TeV in events with a single lepton, large jet multiplicity, and multiple b-jets. *Phys. Lett. B*, 733:328, 2014.
- [63] CMS Collaboration. Search for supersymmetry in events with one lepton and multiple jets in proton-proton collisions at $\sqrt{s} = 13$ TeV. *Phys. Rev. D*, 95:012011, 2017.
- [64] CMS Collaboration. Search for supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV in the single-lepton final state using the sum of masses of large-radius jets. *JHEP*, 08:122, 2016.
- [65] CMS Collaboration. Search for supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV in the single-lepton final state using the sum of masses of large-radius jets. *Phys. Rev. Lett.*, 119(15):151802, October 2017.
- [66] CMS Collaboration. Prospects for SUSY searches using the sum of the masses of large radius jets. CMS Analysis Note CMS AN-15-002.

- [67] CMS Collaboration. Search for supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV in the single-lepton final state using the sum of masses of large radius jets. CMS Analysis Note CMS AN-16-187.
- [68] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. S. Shao, T. Stelzer, P. Torrielli, and M. Zaro. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations. *JHEP*, 07:079, 2014.
- [69] Simone Alioli, Paolo Nason, Carlo Oleari, and Emanuele Re. NLO single-top production matched with shower in POWHEG: S- and t-channel contributions. *JHEP*, 09:111, 2009. [Erratum: *JHEP* **02** (2010) 011].
- [70] Emanuele Re. Single-top Wt-channel production matched with parton showers using the POWHEG method. *Eur. Phys. J. C*, 71:1547, 2011.
- [71] Tom Melia, Paolo Nason, Raoul Rontsch, and Giulia Zanderighi. W+W-, WZ and ZZ production in the POWHEG BOX. *JHEP*, 11:078, 2011.
- [72] Richard D. Ball and others. Parton distributions for the LHC Run II. *JHEP*, 04:040, 2015.
- [73] Torbjørn Sjstrand, Stefan Ask, Jesper R. Christiansen, Richard Corke, Nishita Desai, Philip Ilten, Stephen Mrenna, Stefan Prestel, Christine O. Rasmussen, and Peter Z. Skands. An Introduction to PYTHIA 8.2. *Comput. Phys. Commun.*, 191:159, 2015.
- [74] CMS Collaboration. Event generator tunes obtained from underlying event and multiparton scattering measurements. *Eur. Phys. J. C*, 76(3):155, 2016.
- [75] S. Agostinelli and others. GEANT4 — a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250, 2003.
- [76] M. Czakon, P. Fiedler, and A. Mitov. The total top quark pair production cross-section at $O(s)$. *Phys. Rev. Lett.*, 110:252004, 2013.
- [77] CMS Collaboration. Identification techniques for highly boosted W bosons that decay into hadrons. *Journal of High Energy Physics*, 2014(12):17, December 2014.
- [78] CMS Collaboration. Boosted Top Jet Tagging at CMS. CMS Physics Analysis Summary CMS-PAS JME-13-007.
- [79] Jesse Thaler and Ken Van Tilburg. Identifying boosted objects with N-subjettiness. *Journal of High Energy Physics*, 2011(3):15, March 2011.
- [80] Anson Hook, Eder Izaguirre, Mariangela Lisanti, and Jay G. Wacker. High multiplicity searches at the LHC using jet masses. *Physical Review D*, 85(5), March 2012.

- [81] Timothy Cohen, Eder Izaguirre, Mariangela Lisanti, and Hou Keong Lou. Jet substructure by accident. *Journal of High Energy Physics*, 2013(3):161, March 2013.
- [82] Sonia El Hedri, Anson Hook, Martin Jankowiak, and Jay G. Wacker. Learning how to count: A high multiplicity search for the LHC. *Journal of High Energy Physics*, 2013(8):136, August 2013.
- [83] CMS Luminosity Measurements for the 2016 Data Taking Period. Technical Report CMS-PAS-LUM-17-001, CERN, Geneva, 2017.
- [84] Thomas Junk. Confidence level computation for combining searches with small statistics. *Nucl. Instrum. Meth. A*, 434:435, 1999.
- [85] Alexander L. Read. Presentation of search results: The CL_s technique. *J. Phys. G*, 28:2693, 2002.
- [86] ATLAS Collaboration, CMS Collaboration, LHC Higgs Combination Group. Procedure for the LHC Higgs boson search combination in Summer 2011. Technical Report CMS-NOTE-2011-005, ATL-PHYS-PUB-2011-11, CERN, 2011.
- [87] Christoph Borschensky, Michael Kramer, Anna Kulesza, Michelangelo Mangano, Sanjay Padhi, Tilman Plehn, and Xavier Portell. Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV. *Eur. Phys. J. C*, 74:3174, 2014.
- [88] CMS Collaboration. Search for R-parity violating supersymmetry in pp collisions at $\sqrt{s} = 13$ TeV using b jets in a final state with a single lepton, many jets, and high sum of large-radius jet masses. Technical Report CMS-SUS-16-040, CERN, Geneva, December 2017.