TECHNISCHE
UNIVERSITÄT
WIEN
Vienna University of Technology

DIPLOMARBEIT

Exploring a Novel Signature of New Physics at LHC - A Cut-and-Count Analysis Searching for Soft Displaced Vertices with Missing Energy

zur Erlangung des akademischen Grades

Diplom-Ingenieur/in

im Rahmen des Studiums

Technische Physik

eingereicht von

Felix Lang, BSc
Matrikelnummer 01428291

Ausgeführt am HEPHY - Institut für Hochenergiephysik
der Österreichischen Akademie der Wissenschaften

Betreuung
Betreuer/in: Privatdoz. Mag.phil. Mag. Dr.rer.nat. Manfred Jeitler

Wien, 04.09.2024 _____ _____
(Unterschrift Verfasser/in) (Unterschrift Betreuer/in)

Eidesstattliche Erklärung

Hiermit erkläre ich an Eides statt, dass ich die vorliegende Arbeit eigenständig und ausschließlich unter Verwendung der im Quellen- und Literaturverzeichnis aufgeführten Werke angefertigt habe.

Wien, 04.09.2024

_____

(Unterschrift Verfasser/in)

Abstract

This thesis presents a cut-based, data-driven analysis using the ABCD method and focuses on a novel event signature within the CMS experiment. The explored signature relies on initial-state radiation (ISR), very soft tracks forming a common, measurably displaced vertex, and missing transverse momentum (MET) due to invisible objects in the event. This method is being applied by the Soft Displaced Vertices group in an effort to bridge a gap that exists between previous studies related to long-lived particles (LLPs) and displaced vertices (DV). It falls into the intermediate regime between direct detection searches and prompt decay searches. The study's findings and the production workflow serve as a benchmark for the Soft Displaced Vertices analysis, which is exploring advanced machine learning techniques to improve vertex finding and background-signal separation capabilities.

We describe the steps that have been involved in this study from Monte Carlo (MC) event generation to the statistical methods that have been employed to obtain production cross section limits and 2D exclusion plots. The use of the ABCD method is well-motivated and has been justified in this study by calculating the closure for both data and MC background for the chosen regions. Two signal models, the bino-stop co-annihilation model and the C1N2 model were selected, and samples with varying mass parameters have been generated. In addition to calculating limits based on the 2018 data of the CMS detector, we have given an estimate for the expected improvements from Run III by calculating limits based on background contributions from two centrally produced MC datasets: The Zjets and Wjets datasets which consist of events where Z bosons and W bosons are produced in conjunction with jets.

Our analysis has demonstrated that the large background caused by the inclusion of very soft tracks (down to 0.5 GeV), without requiring any Standard Model objects in the event signature, can be effectively suppressed through a combination of preselection criteria and various well-tuned quality, kinematic, and displacement cuts while retaining a significant number of signal events.

Results from 2018 data show that we are able to achieve great sensitivity for both tested models across a broad parameter space of long-lived particle (LLP) mass and mass splittings. The projected results for Run III of the LHC, based on MC background samples at an integrated luminosity of 300 fb^{-1}, indicate significant improvements, which are supported by the CMS detector's anticipated enhanced tracking capabilities for Run III. Based on our findings, an implementation of a similar study for Run III is expected to achieve even higher sensitivity and potentially exclude even smaller mass gaps.

Contents

1 Introduction

This chapter introduces the experimental designs relevant to this thesis. It also offers a brief overview of the Standard Model (SM) of particle physics, highlighting its limitations. These discussions will set the stage for exploring concepts beyond the SM, forming the foundation of the motivation for this study.

1.1 LHC

The Large Hadron Collider (LHC) is the world's largest particle accelerator. Constructed over a decade from 1998 to 2008 by the European Organization for Nuclear Research (CERN), it is located near Geneva, Switzerland, with significant portions extending into France. The LHC is mainly used as a proton-proton collider and utilizes a system of preaccelerators, many of which have already been instrumental in pushing the boundaries of our fundamental understanding of the universe. An overview of the CERN acceleration complex, which is instrumental for the functionality of the LHC and experiments around it, is depicted in Fig.1.
Since 2020, the source of the proton beams is LINAC4. In this linear accelerator, H^- ions, consisting of a hydrogen atom with an additional electron, are accelerated to an initial energy of 160 MeV. When entering the Proton Synchrotron Booster (PSB), they are stripped of their electrons, leaving only the protons in the accelerator. After being boosted in the Proton Synchrotron (SP) to around 26 GeV, they enter the Super Proton Synchrotron (SPS), where they are further accelerated to 450 GeV before the protons are injected into the LHC.

With a circumference of about 27 km, the LHC is largely buried deep underground. It employs state-of-the-art technology to propel particles to velocities approaching the speed of light. Around it, 1232 dipole magnets are situated, which are used to bend the trajectories of the proton beams going in both directions. An additional 474 quadrupole magnets focus the beams to maximize the number of collisions. The acceleration of the proton beams takes place in the 16 radio frequency cavities. When operating at full capacity, particles in the LHC are being boosted to energies of 6.8 GeV. This allows for an enormous amount of energy to be released in each collision. For Run III, which already started in 2022 and is expected to run until 2025, center of mass energies of $\sqrt{s} = 13.6$ TeV can be achieved. This tops the previous record from Run II (2016-2018) which was set at collision energies of 13 TeV. [1]

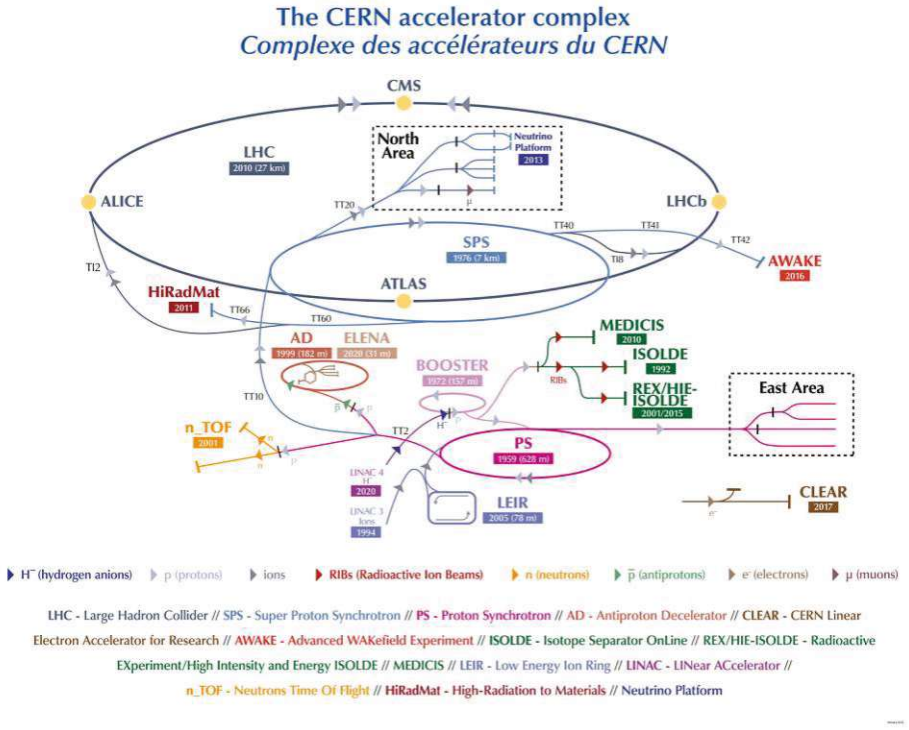

Figure 1: The CERN accelerator complex and experiments located around it. [2]

The energy is, however, not the only important factor when searching for new physics processes or doing precision measurements. It is also crucial to maximize the amount of recorded collisions (Events). The likelihood of any process to occur in a high-energy physics experiment is given by the cross-section of that process. In the search for rare physics processes, this becomes increasingly more important. The relation between the event rate of a given process and its cross-section is given by the following formula:

$$\frac{\mathrm{dN}}{\mathrm{dt}} = \mathcal{L} \cdot \sigma \tag{1}$$

where $\mathcal{L}$ is called the luminosity.

The luminosity is proportional to the number of particles (in this case protons) per bunch (n_1, n_2), the frequency at which they traverse the accelerator ring f, as well as the spread of the bunch in x and y direction (σ_x, σ_y) [3]

$$\mathcal{L} = f \cdot \frac{n_1 n_2}{4\pi\sigma_x\sigma_y} \tag{2}$$

By integrating (1) with respect to time, we get an expression for the total amount of events associated with a given cross-section σ

$$N = L \cdot \sigma \tag{3}$$

where L is the integrated luminosity.

This tells us that the total number of occurrences for a process with a given cross-section is proportional to the integrated luminosity L. Luminosity is precisely measured at the LHC. During Run II, the integrated luminosity for the 2018 data-taking period was recorded at 59.8 fb^{-1}. The expected total integrated luminosity after Run III is projected to be 300 fb^{-1}.

There are numerous experiments located around the LHC, with the four largest ones being ALICE, ATLAS, CMS, and LHCb. This study utilizes data that has been collected at the CMS experiment during Run II and also aims to give prospects for the upcoming data processing efforts for RUN III at CMS.
A summary of the operation principles and detector concepts of the experiment will be provided in the upcoming sections.

1.2 CMS Experiment

The Compact Muon Solenoid (CMS) Experiment is one of the largest experiments at the Large Hadron Collider (LHC) at CERN [4]. It is an international collaboration which was built to study the interactions of subatomic particles at very high-energy scales and in pursuit of explanations for some of the unanswered questions in physics. To accomplish this, the CMS collaboration employs advanced technology and a variety of techniques to measure particle properties more precisely. It is a general-purpose particle detector and is located in an underground hall in the accelerator ring near Cessy in France. One of its original purposes was the search for the Higgs boson, which was announced as discovered on 4 July 2012. It has since been used to investigate a wide range of physics, including further study of the properties of the Higgs boson, the search for evidence of supersymmetry, more generally, physics beyond the standard model, and the study of heavy ion collisions. Its name is referencing its design. In comparison to other experiments at the LHC, such as ALICE [5] and ATLAS [6], it is fairly compact with a cylindrical shape, 21m long, 15m in diameter, and about 14,000 tonnes in weight [4]. Due to its design, it is capable of measuring muon tracks particularly well using a number of different detector types. The reference to solenoid in its name comes from its strong solenoid magnet around which the detector is built. It is 13m long and 6m in diameter, and the cooled superconducting niobium-titanium coil provides a flux

density of up to 4 Tesla. The internal structure of the CMS detector is depicted in Fig.2. The different detector concepts will be described in more detail in the next section.

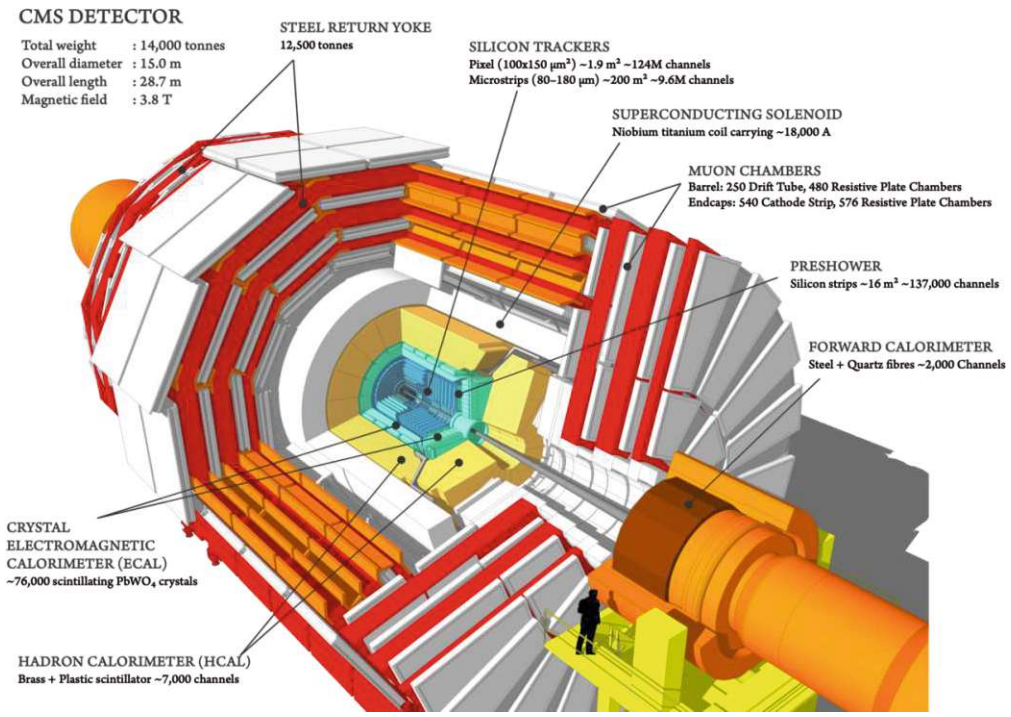

Figure 2: A cutaway diagram of the CMS detector [7]. The cylindrical shape gives the detector its name; the massive muon chambers form the outer layers of the detector and provide the experiment with the unique ability to measure muons very precisely.

1.2.1 Detector Concepts

A crucial part of reconstructing an event is the measurement of the momenta of particles that were produced in the collision. This is done in the CMS experiment by the CMS tracker, which can reconstruct the paths of high-energy muons, electrons, and charged hadrons. The curvature of the tracks through the tracker is used to infer the momentum of the particle. The CMS tracker consists of two different parts: the silicon pixel tracker and the silicon strip tracker. It achieves exceptional accuracy through unparalleled complexity, incorporating approximately 134 million individual electronic readout channels [8, 9]. Tracking is an important concept not just for this analysis. In the CMS detector, tracking is an iterative process where the reconstructed tracks are refitted in each iteration over the detector data. The iterations start by

using seeds, which form the initial track hypothesis and can be based on different parts of the detector. A significant part of the tracking iterations rely on seeds from the pixel detector, which provides the highest spatial resolution [10].

Silicon pixel tracker: The heart of the CMS experiment is the silicon pixel tracker. It consists of multiple modules arranged in a concentric layout. A schematic of one such module is shown in Fig.3. In its pure form, silicon is a semiconductor, meaning that when charged particles pass through the crystals, they will produce electron-hole pairs. An applied voltage is used to collect those charges, which are then further amplified and read out. Signals are stored in the readout chips' memory and undergo a first level of processing before being transmitted to the central data storage using a network of 40,000 fiber optic links. The pixel tracker contains a total number of 124 million pixels, each only being 100 μm by 150 μm in size.

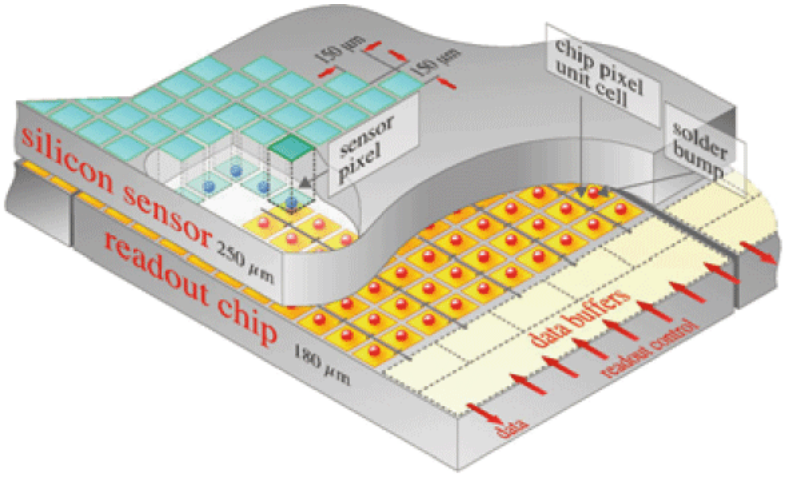

Figure 3: Schematic depiction of a silicon pixel module. Each module consists of a number of tiny silicon elements which can be read out separately from each other by the readout chip. This provides outstanding spatial resolution, a much important feat in any high-energy physics experiment. [11]

Since the pixel detector is the innermost component, it is also subject to the highest amount of radiation and consequently suffers the most radiation damage. For that reason, many of the modules have been replaced during Long Shutdown 2 (LS2), the break between Run II and Run III.

Silicon strip tracker: The second layer of the tracker is formed by the silicon strip detector. It is very similar in functionality to the pixel detector, the

only difference being that the individual silicon elements consist of strips with a length of a couple of centimeters. There are over 15,000 individual modules containing a total amount of about 10 million pixel strips.

Another crucial factor when trying to understand what occurred in the collision is the information about the energy of the various involved particles. Two different types of calorimeters, the electronic calorimeter (ECAL) and the hadronic calorimeter (HCAL) are specifically designed to tackle that task by measuring the energy deposits of passing particles [12, 13].

Electromagnetic calorimeter: The ECAL is the inner part of the two calorimeters. It is designed to precisely measure the energy of incoming photons and electrons and consists of nearly 80,000 lead-tungstate crystals. The structure comprises three distinct sections: the barrel and two endcaps. Lead-tungstate is a high-density transparent scintillator. This means that when photons and electrons pass through the material, they deposit some of their energy, leading to excitations in the material and producing well-defined scintillation light in the process. Photomultipliers on the back of the crystals pick up these signals and amplify them before they are being processed. The ECAL also contains a pre-shower detector, which is used to distinguish between spatially close pairs of photons, which can be produced in neutral kaon decay, and single high-energy photons.

Hadronic calorimeter: The second part of calorimeters is the hadronic calorimeter which, as the name suggests, is able to detect most hadrons. It is a sampling calorimeter by design, meaning that it consists of alternating layers of passive absorbers and active detectors. In the case of CMS, the absorber layers are made of brass, while the active detectors are mainly plastic scintillators. When high-energy hadrons hit the absorber, a large amount of secondary particles is produced, which in turn are able to produce even more particles. This effect is seen as a shower of particles, which can be detected by the layer of scintillators. Since only a fraction of the total energy of the particles is deposited in each absorber, many layers are needed in order to fully infer the total amount of energy of the initial particle. The HCAL is designed to be hermetic and capture, to every extent possible, all particles except muons and neutrinos. To achieve this, there are four different sections of the hadronic calorimeter. The barrel part inside the coil (HB), which itself consists of 36 wedges, each weighing about 26 tonnes, a barrel part outside of the coil (HO), as well as endcap (HE) and forward (HF) sections. This leaves no gaps and ensures that no energy leaks are possible.

One of the most important parts of the CMS detector is the detection of muons. Muons are about 200 times heavier than electrons which leads to a much smaller energy loss from bremsstrahlung in the material. Since they also do not interact strongly, they can penetrate both of the calorimeters without losing much of their energy at all. They are the only particles left that can be directly detected outside the hadronic calorimeter. Muon chambers are arranged in multiple layers, with the return yoke between them. The return yoke acts as an additional filter, only allowing muons and weakly interacting particles like neutrinos to pass through it. The muon chambers utilize four different detector concepts, all working in parallel, allowing for a very robust detection process [14].

Muon chambers: In the barrel region, drift tubes (DTs) and square-shaped resistive plate chambers (RPCs) are used to measure the tracks of the muons. The structure of a DT can be seen in Fig.4. There are 250 drift tube chambers consisting of 8 to 12 aluminum layers, each with up to 90 tubes. Each tube contains a wire within a gas volume. When muons pass through the chamber, they ionize the gas. The produced charges will drift along the electric field to the anode, where they can be detected as a charge pulse. The layers of tubes are arranged in a way that allows us to obtain complete 3D information about the muon track. RPCs function in a similar manner, but instead of tubes, they are made of an anode and a cathode plate with gas between them. The charges produced by the ionization process again drift to the anode, where the signals are picked up by metallic strips. By using the time delay it is possible to accurately determine the muon tracks. The 610 resistive plate chambers provide a very high time resolution of just one nanosecond.
In addition to trapezoidal RPCs, the endcap region also houses 540 cathode strip chambers (CSCs) and 72 gas electron multipliers (GEMs). CSCs consist of a net of copper cathode strips and anode wires, with gas between them. This provides good spatial resolution in the endcap, where the magnetic field is uneven. GEMs were introduced for Run III of the LHC and are installed in the forward regions to supplement the existing muon detectors. Like the rest of the concepts, they are gas detectors that utilize state-of-the-art detection technology to enhance the resolution of muon measurements [15].

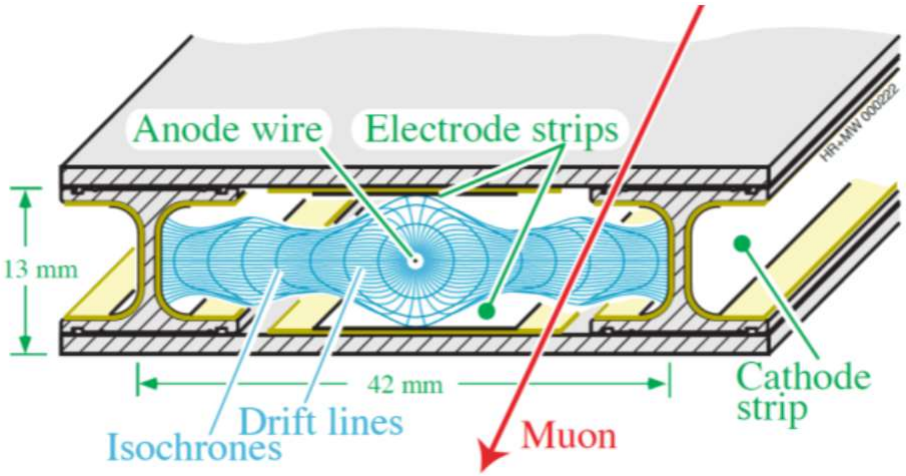

Figure 4: Illustration of the process in a drift tube [16]. The red line represents a muon passing through the drift tube. The charges created by the ionization process drift along the drift lines toward the cathode and anode strips. The isochrones form the surfaces from which charges would have taken the same time to reach the electrodes. These timing differences between incoming charges at the electrodes can be used to calculate the trajectory of the muon through the chamber.

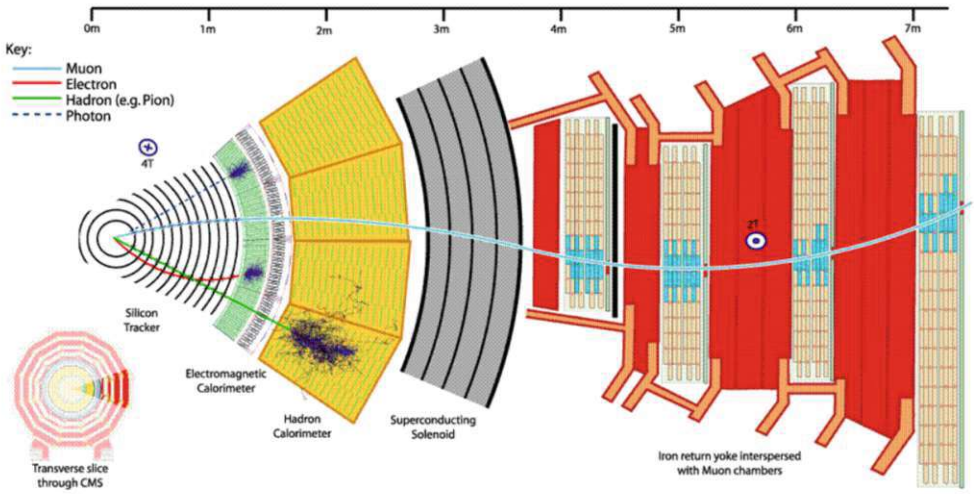

Figure 5: Illustration of the different detection processes taking place in the CMS-detector [17]. The lines represent the trajectories of the muons (blue), electrons (red), uncharged hadrons (green), and photons (dotted blue) and end at the detector element where they are stopped.

Neutrinos will go through the detector structure unhindered and can only be inferred from the missing transverse momentum in the reconstructed event. The trajectory of different subatomic particles is shown in Fig.5.

1.2.2 Trigger System

At its peak, the CMS detector records about a billion proton-proton interactions every second, with minimal time between each of them. The actual frequency of collisions of different proton bunches is around 40 MHz. The time resolution of the detectors is therefore of utmost importance and all the electronic channels have to be in perfect sync to ensure that signals correspond to the same events. Storing all of the data for all of those collisions would also not only be impractical but impossible, since the amount of storage needed for that task would be far too great. To solve this problem, CMS uses a trigger system to decide which events to discard and which events are worth keeping for further investigation. This happens in a two-step process in the experiment. Since the amount of data racks up at an enormous pace, it is necessary to have a system that is able to make these decisions in a very short amount of time. This is achieved by the so-called Level-1 trigger (L1), which is a fully automatic hardware trigger [18]. It is an electronics system specifically designed for this task that takes into consideration the information from the calorimeters as well as the muon detectors and is able to decide whether to discard or keep an event in a matter of 3.8μs. This already reduces the event rate down to 100 kHz. The second part of the trigger system is implemented by the so-called high-level trigger (HLT), which utilizes about 30,000 PC cores and is able to reduce the rate further to only 1 kHz. The HLT is composed of a sequence of reconstruction and filtering modules that implement various algorithms [19]. The events that pass the HLT are finally stored for further reconstruction and analysis.

1.3 The Standard Model of Particle Physics

The standard model is a well-established theory of particle physics. It incorporates three of the four acting forces in the universe, the strong interaction, the weak interaction, and the electromagnetic interaction, and only excludes gravity, which is well described by Einstein's theory of general relativity. The last missing piece of the standard model was the Higgs boson, which was discovered at the LHC in 2012 after a decades-long hunt for its detection.
The standard model is one of the most rigorously tested theories in physics and is able to predict certain results with astonishing accuracy. It is a quantum field theory that is the culmination of many advances in physics during the 20th century. It builds on the concept of gauge theories, in which the Lagrangian of a system, which is used to describe its dynamics, is invariant under certain

transformations, so-called symmetry groups. When requiring not only a global symmetry but also a local symmetry, gauge fields emerge, which are linked to the underlying gauge bosons of the theory. The earliest adaption of such a theory was quantum electrodynamics, which was developed during the 1930s and incorporates a U(1) symmetry. This idea was later generalized to non-abelian groups by Yang and Mills during the 1950s. The SU(3) group is an example of such a group and forms the underlying symmetry group of quantum chromodynamics (QCD). The non-abelian nature of this symmetry group also explains the self-interaction of gluons, the gauge bosons of QCD. A better understanding of the weak interaction came only after the unification with the electromagnetic interaction by Glashow, Weinberg, and Salam, since massive gauge bosons cannot exist in an unbroken symmetry. In the standard model, the electroweak interaction can be described by a $SU(2) \times U(1)$ symmetry. The four massless gauge bosons, the W_1, W_2, W_3 and B boson couple to the weak isospin and the weak hypercharge. Spontaneous symmetry breaking via the Higgs mechanism occurs when the energy falls below the unification scale, and the Higgs field obtains a non-zero vacuum expectation value. This gives rise to mixing states from the precious W_1, W_2, W_3 and B bosons, which combine to the massive $W^{\pm}$ and Z^0 bosons as well as the still massless photon. The weak interaction is also special in the sense that it exhibits parity violation, with maximal parity violation being observed for processes that involve the $W^{\pm}$ bosons. Just as gauge bosons can be seen as excitations of their corresponding gauge field, excitations of the Higgs field can be observed as a Higgs boson [20].
The standard model, as seen in Fig.6, is completed by the fermionic fields which are the quarks and leptons that come in three different generations. The fermions couple to the bosonic fields; this is called being charged under a certain field. All fermions are charged under the weak interaction and, therefore, also obtain their rest mass through the Yukawa coupling to the Higgs field. Only the quarks carry color charge, which couples them to the gluon field. At the same time, all elementary particles apart from the neutrinos are also electrically charged and therefore interact via the electromagnetic interaction mediated by the photon.

The standard model has had tremendous success in predicting experimental results. Not only did it predict several particles before they had been discovered, like the Higgs boson and the top quark. It was also able to accurately predict the masses of the $W^{\pm}$ and Z^0 bosons and has been used to calculate the anomalous magnetic dipole moment of the electron to a precision of one part in a billion in accordance with experiments. Despite its tremendous success, there are also still many open questions that are not explained by the SM, making it ultimately apparent that it is still an incomplete theory.

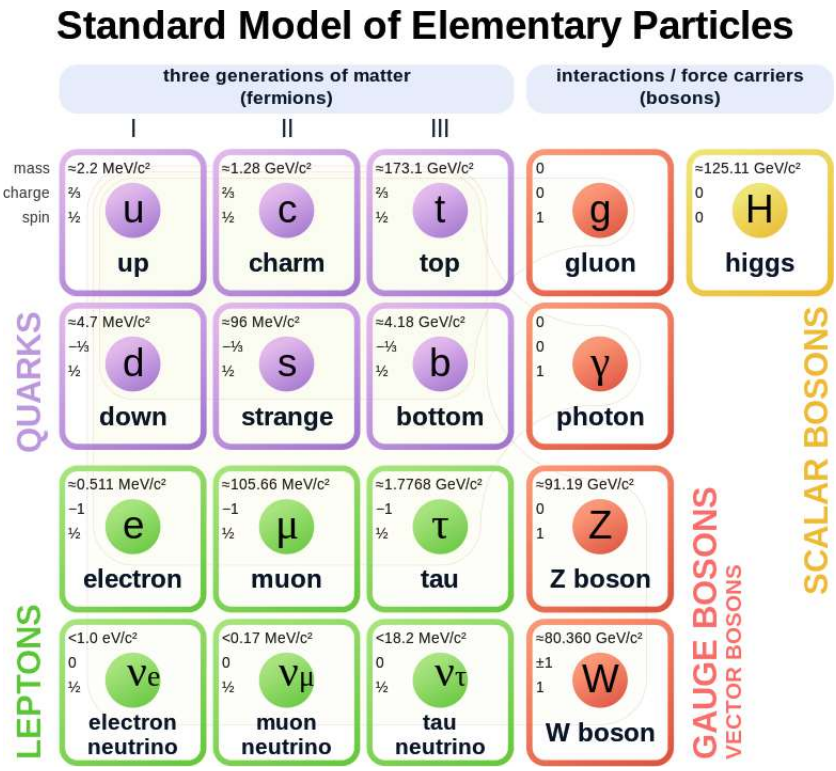

Figure 6: The standard model of particle physics [21]. Particles are structured into the force mediating gauge bosons and the fermions, which are further grouped into three generations that are set apart by their masses. This leads to distinct properties and signatures in high-energy physics experiments that allow their identification. The Higgs boson stands out as a non-mediating scalar boson.

1.4 Physics Beyond the Standard Model

There are some experiments and observations that hint at the existence of physics beyond the standard model. A couple of examples include:

- Neutrinos in the standard model are inherently massless, and there is no understood mechanism that could give rise to neutrino masses. Still, neutrino oscillation experiments have shown that neutrinos are able to change their flavor state over long distances. This implies that there is a mass difference in the mass eigenstates of the neutrino. Although neutrino mass is expected to be extremely small, lighter than 1 eV, this is not yet understood.

- In our observation of rotations of distant galaxies, we have established that their rotation speeds do not follow the expected behavior [22]. The rotation

speed in galaxies is expected to decrease with distances further away from their center. All of our observations, however, point to it being approximately constant, independent of the distance from the galaxy center. This can be explained by the inclusion of new forms of matter called dark matter (DM), which, however, cannot be explained by the standard model. It is, therefore, a strong hint at particles or physics that is not included in the SM.

- This ties into another problem, which is that the standard model does not incorporate gravity and thus cannot make any meaningful predictions at distances where gravity becomes more relevant than the other forces. Gravitational effects can, of course, be described with great accuracy by general relativity. However, it has been a much sought-after goal to combine these two exceptional theories, which is still yet to be accomplished.

- The muon g-2 experiment at Fermilab is designed to precisely measure the anomalous magnetic dipole moment of the muon and has produced some results that seemingly contradict the calculations from the standard model with a confidence level of nearly 5σ [23]. More precise measurements are expected to yield results in the future, but so far, this might already be taken as a hint at new processes giving rise to this phenomenon.

- There is no known process that can explain the overwhelming dominance of matter over antimatter in the universe, which is expected to have been produced in equal quantities. Indirect and direct CP violations (charge conjugation + parity inversion) have been observed in various experiments, but no mechanism to explain this imbalance has yet been found.

- There are also some other arguments to be made about the fine-tuning of the different parameters of the standard model, which is not really understood, as well as the mass hierarchy problem. The mass of the Higgs boson, for example, could be expected to be significantly higher unless some cancellation of the quantum correction terms occurs.

There are many theories, some of which are nearly as old as the standard model itself, which intend to offer answers to some of those open questions. One of them is supersymmetry, which also is the basis of many searches for new particles in physics experiments all around the world, including at CMS.

2 Supersymmetry

The motivation for supersymmetry can be easily explained by the mass hierarchy problem. All massive particles are coupled to the Higgs field; their rest mass is dependent on the Yukawa coupling. It is, therefore, expected that the heaviest particles are coupled the strongest to the Higgs field. At the same time, the mass of the Higgs boson is subject to correction terms caused by Feynman diagrams involving loops of these particles. These corrections would naturally lead to a Higgs mass on the same scale as the heaviest particles. Since new physics is expected to emerge at higher energies, the mass of the Higgs boson would be expected to be of similar magnitude.
The correction term caused by such a fermionic loop can be calculated to be:

$$\Delta m_{\mathrm{H}}^2 = -\frac{|\lambda_f|^2}{8\pi^2}[\Lambda_{\mathrm{UV}}^2 + ...] \tag{4}$$

where λ_{f} corresponds to the Yukawa coupling of the fermion to the Higgs field and Λ_{UV} represents the cutoff scale which associates to the energy scale until which the theory is valid and beyond which new behavior is expected to emerge.

In a similar way loops which are caused by bosonic particles also contribute corrections to the Higgs mass and can be calculated in higher order to:

$$\Delta m_{\mathrm{H}}^2 = \frac{\lambda_S^2}{16\pi^2}[\Lambda_{\mathrm{UV}}^2 + 2\mathrm{m}_S^2 \ln \frac{\Lambda_{\mathrm{UV}}}{\mathrm{m}_S} + ...] \tag{5}$$

where this time λ_{S}, m_S represent the Yukawa coupling and the mass of boson and Λ_{UV} is again the cutoff scale [3].

It can be seen that if two complex scalars existed in addition to all fermions with the same Yukawa coupling then the quadratic quantum corrections would cancel each other exactly thus providing an elegant fix to the mass hierarchy problem.

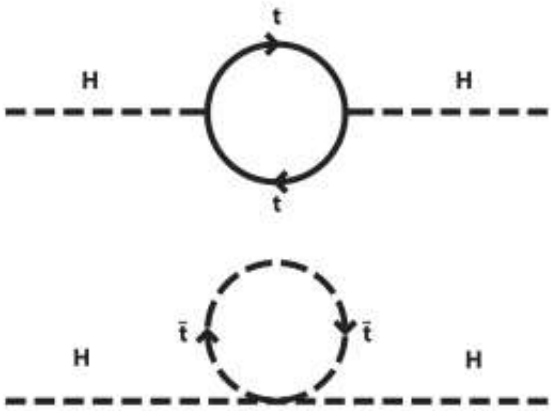

Figure 7: Processes leading to the cancellation of the quadratic mass renormalization for the Higgs boson caused by supersymmetry. The mass corrections caused by the fermionic loop of the top quark get canceled by the loop of its supersymmetric partner, the top squark. [24]

Supersymmetry introduced a new symmetry to the theory, which relates bosonic and fermionic states. In the simplest formulation of supersymmetry, the minimal supersymmetric standard model (MSSM), each particle of the standard model receives a so-called superpartner, with the superpartners of the fermions being bosons and vice versa [25]. Since the couplings of the superpartners to the Higgs field are the same, in an unbroken symmetry, each particle and its superpartner would have the same mass, which has not been observed in any experiments.
There are some theories like Soft SUSY breaking, in which the cancellation of the mass correction terms for the Higgs boson can be preserved even after spontaneous symmetry breaking [26]. Finally, the hope is that the energy at which supersymmetry is spontaneously broken is accessible at the LHC, and new particles could be found on those scales. The supersymmetric partner of the top quark, the top squark, is of particular interest since it is the lightest bosonic SUSY particle in many models. Likewise, the supersymmetric partners of the SM bosons are also of interest. Before SUSY and electroweak symmetry breaking there are two neutral gauginos, the bino $\tilde{B}^0$ and the neutral wino $\tilde{W}^0$ as well as two charged winos $\tilde{W}^+$ and $\tilde{W}^-$. They emerge as the superpartners of the W and Z bosons from the SM. The Higgs field is also extended into a pair of two neutral and two charged higgsinos. After the symmetry breaking, the neutral gauginos and the two neutral higgsinos mix to form the four neutralino mass eigenstates $\tilde{\chi}_i^0$. In a similar way, the remaining two charged gauginos and the charged higgsinos form two different mass eigenstates which can be either positively or negatively

charged and are called charginos $\tilde{\chi}_i^{\pm}$.
This solves an additional problem of the SM because the lightest supersymmetric particle (LSP) would, in this case, be a weakly interacting massive particle (WIMP) which would make it an excellent dark matter candidate [27]. If the mass hierarchy of these supersymmetric particles were to be configured in such a way that the mass difference between LSP and the second-to-lightest supersymmetric particle is reasonably small, then the second-to-lightest supersymmetric particle would be a long-lived particle (LLP). If such an LLP were to be produced in the CMS experiment, then it would experience significant displacement before decaying. This would lead to a so-called displaced vertex (DV), which is something that could be observed in experimental data and is also the basis of this analysis.

Another compelling argument for supersymmetry is made by so-called co-annihilation models, which could explain the scarcity of dark matter (DM) in the early universe [28]. An example of such a co-annihilation model is the bino-stop model, where the LSP is dominated by its bino portion, when the masses of bino and stop lie reasonably close to each other. The result is a so-called compressed model. The small mass gap between those two particles has the effect of enhancing the interaction rates, including scattering and annihilation processes between bino and stop. This enhanced annihilation rate could explain the composition of matter in the early universe.

3 Motivation

The CMS experiment has been actively searching for physics beyond the standard model for over 15 years since its establishment in 2008. Despite these efforts, no evidence of new particles predicted by SUSY or other BSM models has been observed thus far.

During the first two runs, Run I (2009-2013) and Run II (2015-2018), an analysis technique called mono-X search proved to be very popular. Mono-X searches rely on the fact that a visible SM object, as well as any missing transverse energy, is recoiled against the initial state radiation (ISR), which is produced by the protons before the actual collision and consists of high-energy charged particles. Particles are considered boosted by the initial state radiation because any momentum of the ISR particles has to be compensated by the different physical objects, recoiling off of it. This ISR boost experienced by the visible and invisible objects makes it easier to detect any potential new particles, since the reconstruction and detection efficiencies are limited in the low-energy regime. The most notable and also most sensitive analysis in that regard so far has been the mono-jet analysis in which the visible SM object in the event signature is a jet, formed by the hadronization of quarks and gluons produced in the collision [29]. Other analyses have tried to increase the sensitivity for models with compressed spectra by also including different SM objects such as soft leptons [30].

The Soft Displaced Vertices analysis [31, 32] attempts to bridge a gap between previous approaches, which is caused by the fact that in compressed spectra models, SUSY particles might evade detection, since the SM particles produced in their decay at the secondary vertex are evading detection due to them not originating from the primary vertex (PV) or simply being too soft. This study implements a cut-and-count analysis looking for a specific event signature, which includes ISR, very soft tracks down to 0.5 GeV, which form a secondary vertex, and a large part of that energy being carried away by invisible objects. This event signature is shown in Fig.8.
The SM objects that are created in the decays of BSM particles are so soft that they are not reconstructed in the detector. Therefore, we can only rely on the tracks that have been obtained from the tracker to reconstruct any secondary vertices. This imposes some big challenges, since the reconstruction from tracks with very low p_T is a challenging feat and also leads to a very large background caused by the very loose restrictions of the event signature. Even though the initial background is overwhelmingly large, we will show in this study that we are able to suppress the background to a manageable level whilst also retaining enough signal by using well-tuned cuts. This allows us to achieve good sensitivity

to the tested models in a large region of the parameter space. All the steps that have been taken in this analysis, from the data generation to the data processing, will be described in the upcoming sections.

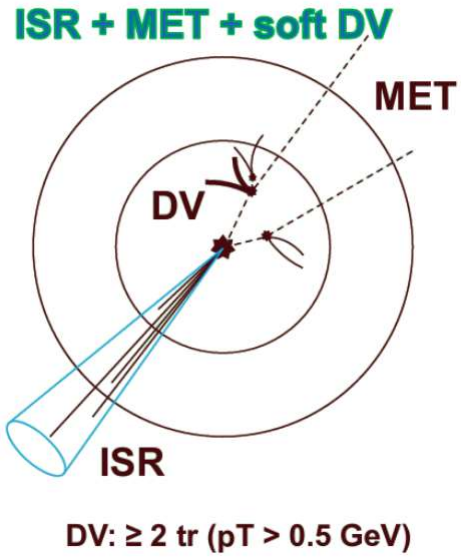

Figure 8: Event signature of the analysis. Initial state radiation leads to a boosting of the event. Displaced vertices, which are being reconstructed from tracks with as little as 0.5 GeV, are the basis of the analysis. We also look for missing transverse momentum in events, which would be expected in the presence of new particles that do not interact with regular matter. [31]

4 Event generation

The first crucial step of many analyses involves acquiring data to serve as a foundation for testing the methods that are being employed in the analysis process. Even for data-driven studies, it can often be beneficial to include and test your analysis against Monte Carlo (MC) generated samples. Monte Carlo event generation is a powerful tool used in particle physics to simulate the expected behavior of particles and interactions based on theoretical models. MC simulations are used to generate synthetic events that mimic the characteristics of real particle collisions. They implement all the crucial steps from the production of particles in the event to their decay along different decay branches and interaction of these particles with the detector.
General purpose Monte Carlo generators like PYTHIA [33] and HERWIG [34] provides a way to model these high-energy collisions. They employ perturbative quantum field theory (QFT) techniques to model these interactions. Perturbation

theory relies on the fact that contributions to an effect can be ordered in powers of a parameter less than one so that contributions of higher order can be neglected without changing the result drastically. This only works for quantum chromodynamics in a limited time/energy frame, since the coupling strength is approximately one at low energies. This means that processes with an arbitrary amount of vertices in their Feynman diagram will not be suppressed to ones with a lower amount of vertices. These generators describe various aspects, starting from the physics at very short distances (below one femtometer), where perturbation theory can be applied, to hadronization, which is the formation of hadrons from their constituent particles. For the event generation of the different signal models gridpacks have been utilized. These are sets of precomputed components of the simulated collision, like the matrix elements for different scattering and production processes, as well as the parton distribution functions (PDFs), and are generated by a matrix element generator, in this case, MADGRAPH [35].
Beyond the standard model (BSM) processes can be included in the event production by utilizing specific gridpacks and specifying the decay tables for those particles. The sample generation can be conveniently done in CMS software (CMSSW) which is a framework for data processing and data analysis by the CMS collaboration. It also seamlessly integrates PYTHIA, which has been used to simulate the decay of the long lived particles in this study.
The data processing follows a flow of different data tiers, depicted in Fig.9. This first step of calculating the events taking place during the initial stages of the collision at very short distances, and the hadronization, is done in the first data tier called GEN.

Figure 9: Flowchart depicting the different steps and data tiers during MC event production. Different formats also include additional information about the generated events, such as generator information or information about detector responses.

The second step is to simulate the energy deposit of the produced particles in the different detector modules, with the resulting data tier being called SIM. The digitization step is used to calculate the detector response to those events of interest but it also adds pile-up information on top of it. Real events in the CMS detector do not consist of a single proton-proton collision. Since many of the protons in the crossing bunches collide with each other, many different collisions happen. These could be collisions that are not head-on and where only a fraction of the available energy is converted into new particles. These interactions are called pile-up events and are mostly a nuisance but do have to

be simulated regardless. During the actual data taking, as many as 50 pile-up events occur for each relevant interaction in the detector. The digitization process simulates the response of the trigger system. This means that both the L1 trigger as well as the HLT are being simulated in this step. This will also result in certain trigger paths being active or inactive for each event, which can be used by the analysis further down the line. After these first steps, the MC-generated data will be in a format similar to the data from the actual experiment. It can be converted to the same datatype (RAW), which can then be used for the T0 offline reconstruction (RECO). The reconstruction process aims to identify a range of physics objects such as tracks, vertices, jets, electrons, muons, etc. and utilizes particle flow algorithms (PF) [10]. The information that is being stored also contains information about reconstructed hits and clusters. In our case, secondary vertices are reconstructed using the inclusive vertex finder (IVF), which is a feature of CMSSW containing a number of tunable variables to adapt the vertex finding algorithm to the needs of different analyses. Its parameters have been adjusted to allow for a looser and more inclusive reconstruction of vertices than the standard settings. The advantages of using simulated events even for a data-driven analysis are significant. For simulated events, the GEN level information is readily available, which allows the testing of different approaches during the analysis. For example, the efficiency of matching the tracks to specific vertices, which is a non-trivial task, can be investigated for different procedures. There are standard data formats used after reconstruction known as AOD, MiniAOD, and NanoAOD. Each aims to condense available information by excluding less relevant data, thereby reducing file sizes for efficient storage and data analysis. However, certain analyses, including the Soft Displaced Vertices analysis, require variables that are not included in these standard formats.

The modularity of CMSSW allows the definition of custom modules, which can be used during event generation or data processing. These modules are, for example, filters, which can be applied to filter events based on certain requirements, or table producers, which can be used to add custom tables to store additional data in different formats. Usually, the AOD data stays the same over the course of an analysis, while multiple reprocessing efforts occur once more insight is gained into the analysis and different variables that could prove helpful. The reprocessing of these custom MiniAODs takes more time and should thus be limited to including a number of different variables, while the reprocessing of custom NanoAODs is quicker and more flexible, which allows it to be done more often.

4.1 Signal processes

There are two different benchmark signal models that have been investigated in this study. The first is the bino-stop co-annihilation model. Fig.10 shows the Feynman diagram of the four-body stop decay, in which the stop is a long-lived particle (LLP) for certain mass splittings from the lightest supersymmetric particle (LSP), which is, in this case, dominated by its bino content. This four-body decay of the top squark is the only decay branch that has been chosen for this analysis, meaning all stops produced in the signal generation will decay in such a process. The second possible decay channel, which involves the decay of stops into a charm quark and the lightest neutralino, depends on a flavour-violating model parameter and is assumed to be suppressed in our case.

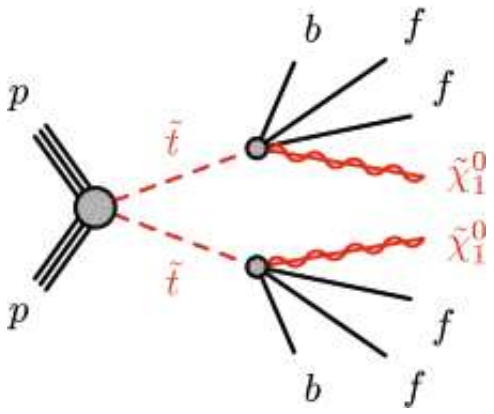

Figure 10: Feynman diagram of the four-body stop decay, which is the primary decay mode that has been implemented in this analysis. A top squark decays into the lightest neutralino, two fermions, and a b quark. Due to the softness of the SM constituents in the decay, they avoid being reconstructed as SM objects in the detector, which is why only the tracks can be relied on to reconstruct the secondary vertex. [36]

The second benchmark signal model is the electroweak C1N2 model, where a chargino and neutralino are produced in the event. In the process, the chargino decays promptly while the neutralino becomes an LLP if the mass splitting to the LSP is only a fraction of its mass. The decay chain that has been implemented for the signal generation of this model can be seen in Fig.11.
In this case, the chargino $\tilde{\chi}_1^{\pm}$ and the neutralino $\tilde{\chi}_2^0$ are assumed to be wino-like, while $\tilde{\chi}_1^0$ is dominated by its bino portion. This constellation leads to the observed mass spectrum and provides for the LLP feature of second lightest neutralino $\tilde{\chi}_2^0$.

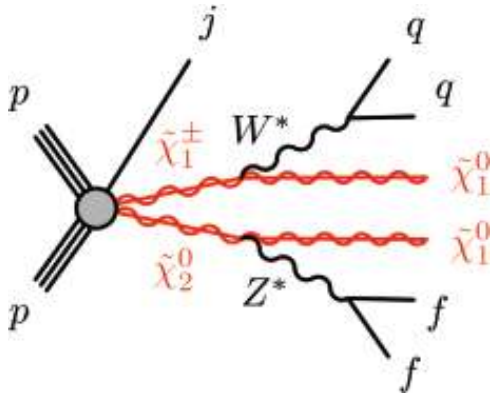

Figure 11: Feynman diagram of the implemented C1N2 production model and ultimate decay of the BSM particles. The chargino decays promptly into the lightest neutralino and an off-shell W boson which will further decay into a pair of quarks. The long-lived second lightest neutralino decays into the LSP and an off-shell Z boson, which decays further into two fermions. [36]

In both scenarios, the small mass splitting between LLP and LSP results in the LLP decay being visibly displaced from the primary vertex. It is the aim of this analysis to demonstrate that we can efficiently detect such vertices and improve the limits on production cross-sections for these processes.

A number of different signal samples have been generated: three different mass points have been chosen, at 600, 1000, and 1400 GeV LLP mass as well as four different mass splittings to the LSP at 12, 15, 20 and 25 GeV. The mass splitting directly correlates to the proper lifetime of the LLP, with the proper lifetime being 200, 20, 2, and 0.2 mm for these different mass splittings respectively.
The inclusion of different mass points and mass splittings is important to understand the effects of the different model parameters on the obtained results. For example, a higher LLP mass, which will correlate to a higher fraction of the vertex energy being carried away by an invisible object will naturally change the reconstruction efficiency of the secondary vertex. Different models might also exhibit different kinematic properties which will also likely result in them being affected differently by some of the cuts. It is important to understand if certain selection criteria that are reliable to suppress the background yield in relation to the signal yield for one model are also suitable for different mass points. A smaller mass gap between LLP and LSP will lead to a larger distance between PV and SV for example, which could provide a limiting factor in our choice of displacement cuts.

4.2 Background processes

For most of the known background processes, the generation of MC samples is conducted centrally, and the data is available worldwide. The only thing that needs to be done is the reprocessing to the custom data formats to include any non-standard variables necessary for the specific analysis. There are two different background processes which have been taken into account for this thesis. The first one is the Zjets dataset, which consists of the production of Z bosons and jets in the event where the Z bosons decay via the process $Z^0 \rightarrow \nu\nu$ to a pair of neutrinos. The second background process comes from the Wjets dataset, where, likewise, W bosons are created in conjunction with jets, and the W bosons decay to a lepton and a neutrino $W^{\pm} \rightarrow l\nu$. These processes have been chosen, since they have a large production cross-section, which makes them an abundant source of background with very high event yields. In addition to that, both of the decays involve the production of neutrinos and jets. Since the neutrinos will not be detected in the experiment, these processes lead to the event signature of jets+MET, which is something that is expected from our BSM signature and also makes them the most dominant background in the mono-jet analysis [29].

5 Implementation

The analysis has been conducted in two distinct phases. We examined both data collected during the 2018 data-taking period at the CMS experiment, as well as MC-generated samples from the Zjets and Wjets datasets. Limits have been obtained using the MC-generated signal samples for the two different signal models with varying model parameters.

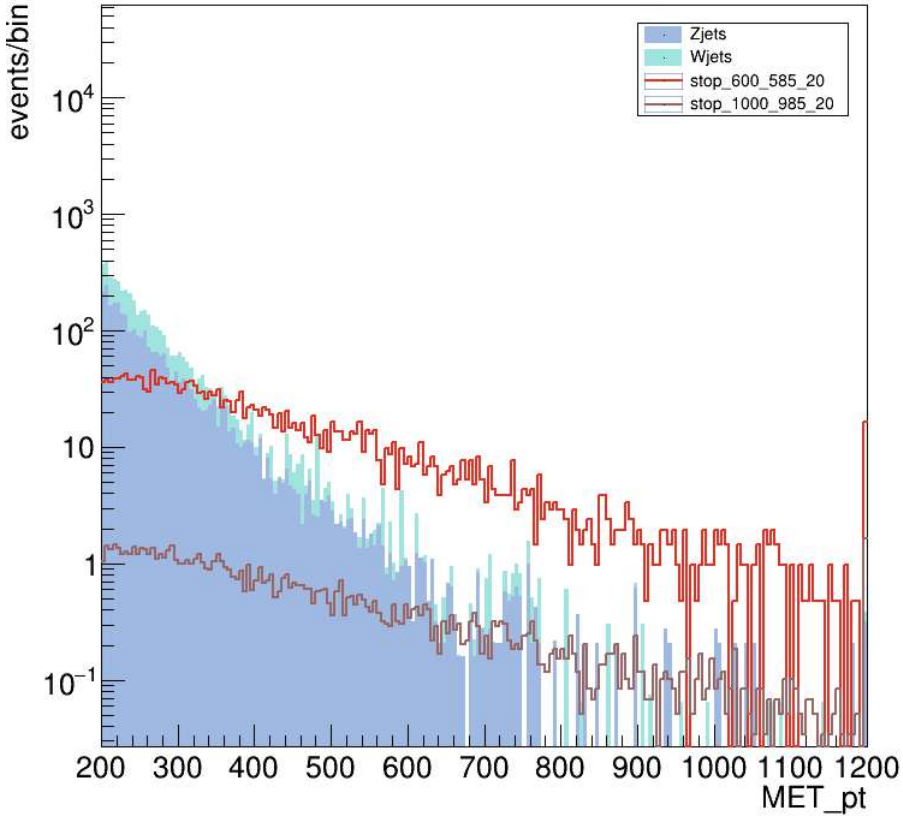

Figure 12: Distribution of E_T^{miss} for the two background processes as well as two signal mass points for the stop production after applying all cuts. The background processes have been stacked on top of each other. The cuts include one that rejects all vertices that have an $L_{xy}^{significance}$ of 20 or lower. The $L_{xy}^{significance}$ is defined by $L_{xy}^{significance} = \frac{L_{xy}}{\sigma(L_{xy})}$ and is the ratio of the transverse displacement of the vertex and the uncertainty of it. Both MC background and signals have been scaled to the same integrated luminosity (300 fb^{-1}). The event yields are displayed on the y-axis, the overflow has been moved to the last bin.

The data has not been unblinded for this thesis. Instead, we used an Asimov dataset to obtain limits. The dataset has been produced by predicting the amount of data in the signal region through the event yield in the control regions. For the sake of consistency, the same procedure has been applied to the MC-generated background. The aim of this study is to establish the limits that can be extracted from the 2018 data using our developed cut-and-count method and then extend these results using MC background to give a projection for future similar analysis with Run III data. Run II at the CMS detector amounted to a total integrated luminosity of 59.8 fb^{-1} [37]. The projection for Run III will be given for an integrated luminosity of 300 fb^{-1} [38]. In the next sections, we will delve into the methods that have been applied to achieve these results, but before explaining the principles of cut-and-count, we first want to showcase its powers.

Fig.12 shows the E_T^{miss} distribution for the background and two selected signal samples after applying all other cuts. From the plot, it is clearly recognizable that beyond a certain value of E_T^{miss}, the signal yields surpass the MC background. This makes it apparent that for statistical analysis, where the limits we obtain will greatly depend on the ratio between signal and background events, a well-placed cut in E_T^{miss} will enhance our sensitivity significantly. The threshold of this cut is dependent on the model parameters like the mass of the LLP and the mass splitting to the LSP. This is easily explainable due to the mass dependence of the production cross-section of both the top squark and the second lightest neutralino, leading to a decrease in produced BSM particles. The shape of the curves of the two different signal models can be seen two be quite similar to each other, which means that both samples are affected by the cuts in the same way. This makes sense if one thinks of the fact that the kinematic properties of the decay remain similar no matter the mass of the LLP.

5.1 Cut & Count analysis

Cut-and-count refers to an analysis type in which the data is subjected to different cuts in the parameter space in order to enhance the signal-to-background ratio and ensure the quality of the data. Events which do not satisfy certain cuts are removed from the dataset.
After applying all the cuts of the analysis the remaining event yields form the basis of additional computations like the calculations of limits for production cross-sections and the production of exclusion charts for signal models. In this study we can classify the different cuts into three categories depending on the variables on which they are based on:

- Event level cuts: These are cuts to variables that are unique per event. An example of such a variable would be the missing transverse momentum (E_T^{miss}). Events that do not satisfy such a criterion will be immediately discarded by the cut. This means that all of the vertices that are associated with this event will be removed from the dataset.

- Vertex level cuts: Each event can contain a number of displaced vertices. Vertex level cuts refer to cuts that are based on variables that are associated with an individual vertex, for example, the displacement of the vertex in relation to the primary vertex in the transverse plane, d_{xy}. Events that do not contain a vertex that passes through all the vertex-level cuts are discarded. Out of all the vertices that are associated with a single event only the one with the highest $L_{xy}^{significance}$ is being considered in the final analysis. The vertex multiplicity after applying all the cuts is relatively low, which can be seen in tables Tab.5 and Tab.6.

- Track level cuts: Tracks that form a common displaced vertex are linked to that vertex via a look-up table allowing unique identification. A single track cannot be counted for multiple different displaced vertices. Track level cuts act on variables that are specific to each track, for example, the number of hits in the tracker that are associated with that track or the momentum of that track. Each vertex requires at least two tracks that satisfy all the track level cuts in order to be kept.

5.2 Discriminating variables

All the cuts that have been applied in this cut-and-count analysis, excluding the ones which are used to construct our signal and control regions, will be listed in this section.
To describe the cuts in a comprehensible way the standard variable names of

the NanoAOD format will be provided. For the purpose of the analysis, some additional variables have been calculated and included in a custom version of the NanoAOD data format. These are variables relating to the tracks and the vertices, and they will be indicated by the "Track" and "Vertex" prefix in front of the variable names.
All the cuts can be classified by the variables on which they act into either event level, vertex level and track level cuts as mentioned before. In addition to that, they serve different purposes: some of them can be seen as quality cuts, for example, to ensure a proper agreement between the displaced vertices and the tracks that they contain. Other cuts, like the angles between missing transverse momentum and the momentum of individual tracks or vertices as a whole, can be classified as kinematic handles. The third and last category are cuts on variables that contain information about the displacement of the tracks and vertices which also play a big role in this study.

Variable	Description	Value
MET_{pt}	missing transverse momentum	> 200
$Jet_{pt}[0]$	transverse momentum of the leading jet	> 100
HLT_PFMET120_PFMHT120_IDTight	HLT trigger path	True
Flag_METFilters	set of noise filters	True
$Jet_{chHEF}[0]$	charged energy fraction of the leading jet	> 0.1
$Jet_{neHEF}[0]$	neutral energy fraction of the leading jet	< 0.8
Electron selection	combination of requirements for electron objects	True
Muon selection	combination of requirements for muon objects	True
Tau selection	combination of requirements for tau objects	True
Photon selection	combination of requirements for photon objects	True

Table 1: Table containing the preselection cuts of this study. The preselection is the first set of cuts that are being applied to the samples and serves the purpose of cleansing the data of events that might not be relevant to this study. All of the cuts that are shown in this table are event-level cuts.

Electron selection		
Electron_cutBased	cut-based electron identification	>= 1
Electron_η	pseudorapidity of the electron object	< 2.5
Electron_pt	transverse momentum of the electron object	> 10
Muon selection		
Muon_looseId	loose muon identification	True
Muon_η	pseudorapidity of the muon object	< 2.4
Muon_pt	transverse momentum of the muon object	> 10
Tau selection		
Tau_idDecayModeOldDMs	tau identification using decay mode finding	True
Tau_η	pseudorapidity of the tau object	< 2.3
Tau_pt	transverse momentum of the tau object	> 18
Photon selection		
Photon_cutBased	cut-based photon identification	>= 1
Photon_η	pseudorapidity of the photon object	< 2.5
Photon_pt	transverse momentum of the photon object	> 15

Table 2: Table containing the cuts for electron, muon, tau and photon objects used for the preselection. We have decided not to include events that contain leptons and photons that fall into a specific parameter space due to a loss in efficiency to correctly reconstruct them below a certain momentum and above a certain pseudorapidity. The cuts for those objects are all structured very similarly: The correct identification of those objects is ensured by using dedicated variables. In addition to that all events with those objects are rejected where the objects fall in a specific window of pseudorapidity and transverse momentum with only the values varying slightly from object to object.

Track level cuts		
abs(Track_dxy/Track_dxyError)	absolute value of the significance of the transverse impact parameter of the track	> 4
abs(Track_dxy/Track_dz)	fraction of transverse and longitudinal impact parameter of the track	> 0.25
Track_pfRelIso03_all	relative isolation of the track	< 5
Track_normalizedChi2	goodness of fit criterium for the track	< 5
Track_numberOfValidHits	number of hits by the track in the tracker	> 13
Track_ptError/Track_pt	significance of the track transverse momentum	< 0.015
acos(cos(Track_phi-Jet_phi[0]	collinearity of the track and the leading jet	> 1
acos(cos(Track_phi-MET_phi)	collinearity of the track and the missing energy	< 1.5

Table 3: Table containing the track-related cuts, including quality, kinematic handles and displacement handles. The first cut can be seen as a displacement cut on the track to select displaced tracks which also display sufficiently high purity. The following five cuts aim to increase the quality of the selected tracks by imposing various constraints on them. The last two cuts confine the spatial relations between the tracks and the missing energy fraction as well as the primary jet of the event.

Vertex level cuts		
Vertex_pAngle	pointing angle of the vertex	> 0.2
Vertex_newgoodtr	number of tracks that satisfy all track criteria to form the vertex	>= 2
acos(cos(Vertex_L_phi-Jet_phi[0]	collinearity of the vector between SV and PV with the leading jet	> 1
acos(cos(Vertex_L_phi-MET_phi	collinearity of the vector between SV and PV with the missing energy	< 1.5
Vertex_ndof	number of degrees of freedom of the vertex	> 1
(for MC signal only) Vertex_matchedLLPnDau_bydau	number of BSM particles matched to the vertex	> 1

Table 4: Table containing the cuts that relate to vertex level variables. Similar to the cuts on the tracks, some of them are more related to ensure the quality of the retained vertices, like the cut requiring a set number of good tracks to be associated with each vertex. Others are focused on the kinematics of the vertex and again confine it spatially in relation to the missing energy and the leading jet of the event.

The preselection is an important step of any analysis and provides the initial step of filtering the data from events that are irrelevant or exhibit poor quality. It is a crucial step to ensure that the subsequent steps of the analysis are being performed on a more manageable and cleaner subset of data to improve the accuracy and efficiency of the analysis. A table of all the cuts that relate to the preselection of this analysis can be seen in Tab.1. We will now provide a more detailed description as to the motivation of their inclusion and their effects.

- MET_pt $>$ 200: This variable corresponds to the missing transverse momentum of the event (E_T^{miss}). The cut at 200 has been chosen because at this value the BSM trigger roughly reaches its efficiency plateau.
- Jet_pt[0] $>$ 100: The recoil condition for the leading jet has been put at 100 GeV to ensure sufficient boosting of the SM and BSM objects. The leading jet is the jet with the highest transverse momentum of the event.
- HLT_PFMET120_PFMHT120_IDTight: This is a high-level MET trigger based on missing transverse momentum and hadronic transverse momentum, which is the transverse momentum of hadronic activity detected in the event.

- Flag_METFilters: This is a noise filter based on hadron barrel and hadron endcap noise filters, which flags events with an unusual level of noise or interference in the HBHE detector.

- Jet_chHEF[0] $>$ 0.1: Cut related to the charged hadronic energy fraction of the leading jet. The charged hadronic energy fraction is the amount of energy carried by charged hadrons in the jet.

- Jet_neHEF[0] $<$ 0.8: This cut relates to the energy fraction attributed to neutral hadrons in the leading jet. Like the other cut on the leading jet energy fraction, it is used to reject anomalous events that might originate from detector noise.

- Muon_looseId && Muon_eta $<$ 2.4 && Muon_pt $>$ 10: Below certain values of p_T and in regions with higher pseudorapidities the efficiencies to reconstruct physics objects drops significantly, which is why these cuts have been included. For muons, the muon loose ID property has been picked to veto muon objects. This is the loosest requirement to identify a muon.

- Photon_cutBased >= 1 && Photon_eta $<$ 2.5 && Photon_pt $>$ 15: In addition to the kinematic handles, the Photon cutBased property has been chosen to select appropriate events. The cutBased variable contains four different values: (0:fail, 1:loose, 2:medium, 3:tight). Again the loosest identification has been chosen.

- Tau_idDecayModeOldDMs && Tau_eta $<$ 2.3 && Tau_pt $>$ 18: This cut has been chosen to veto tau objects. The Tau idDecayModeOldDMs variable has been chosen to identify tau particles, and additional requirements on η and p_T are imposed on them.

- Electron_cutBased >= 1 && Electron_eta $<$ 2.5 && Electron_pt $>$ 10: Electrons are selected similar to photons, also utilizing the cutBased variable and requiring the loosest selection. The η cut has also been set at the same value but lower transverse momenta are allowed for electrons.

Since this analysis relies solely on the tracks for the vertex reconstruction without any additional information from reconstructed objects it is imperative to ensure a high quality of the tracks while also making sure to suppress the vast majority of tracks that are related to vertices from SM processes. A full list of all the track-related cuts can be seen in Tab.3, and we will describe the effects and motivations of all of the cuts in the upcoming segment.

- abs(Track_dxy)/Track_dxyError > 4: The Track_dxy is the transverse impact parameter of the track. It is defined by the point of closest approach to the primary vertex in the transverse plane of the track. The track is reconstructed within a certain accuracy, which is given by the Track_dxyError. The ratio of those two variables is called the $d_{xy}^{significance}$ of the track. Choosing a cut value for the $d_{xy}^{significance}$ ensures that the track has been constructed with sufficient accuracy.

- abs(Track_dxy/Track_dz) > 0.25: Similar to the Track_dxy being the transverse impact parameter, Track_dz is the longitudinal impact parameter of the track. The fraction of d_{xy} and d_z is a valuable cut to reject tracks that originate from pile-up events. These pile-up tracks tend to have a smaller d_{xy} value in conjunction with usually large d_z values. Signal tracks, on the other hand, can also have a large Track_dz but are associated with high values of Track_dxy.

- Track_pfRelIso03_all < 5: The relative isolation of the track is a measure of the energy deposited in a cone around the track, normalized to the energy of the track. The cone size is given by the qualifier "03". It represents the opening angle of the cone described by $\Delta\phi^2 + \Delta\eta^2 \leq 0.3$. There are different isolation variables that correspond to the energy deposited by neutral and charged particles in the cone. The isolation used for this study takes both of them into account and is calculated from the particle flow algorithm. An isolation cut has proven to be a useful way to discriminate background while retaining signal.

- Track_normalizedChi2 < 5: This is a track quality indicator. The normalizedChi2 variable measures the goodness of fit for the track relative to the detector hits. A low normalizedChi2 value indicates a good agreement between the reconstructed track and the hits on the detector.

- Track_numberOfValidHits > 13: The numberOfValidHits variable of a track refers to the number of hits in the silicon tracker that are associated with the track and have been consequently used to fit the track. A larger number ensures that the track is more likely to be better reconstructed while also serving as a cut to reject detector noise.

- (Track_ptError/Track_pt) < 0.015: The ratio between the error of the track's p_T and its p_T represents the track's relative p_T uncertainty. It describes how well the track's p_T has been measured and is therefore a good track quality indicator.

- acos(cos(Track_phi − Jet_phi[0])) > 1: The collinearity between the track and the leading jet in the longitudinal direction is used to ensure spatial separation between the selected tracks and the leading jet. This is physically motivated, since we are interested in objects which are far away from the jet. In order to suppress any tracks which might have originated from the hadronic activity of the jet a cut rejecting lower values has been applied.

- acos(cos(Track_phi − MET_phi)) < 1.5: In a similar way, an upper cut on the collinearity between the track and the missing energy ensures that the track is spatially close to the region in the detector where missing energy has been observed and we expect BSM particles to potentially be present.

Vertices are the key part of this analysis. Therefore we have to ensure that they are of high quality and are also kinematically confined to regions in the parameter space that are of interest to our analysis. We will now provide a more in-depth look at the variables that have been chosen to accommodate these tasks while a more concise and full list of all cut variables and cut values is shown in Tab.4.

- Vertex_matchedLLPnDau_bydau > 1: This cut is only relevant for the simulation of the signal samples. It ensures that at least two particles generated by the particle generator have been matched to the secondary vertex.

- Vertex_pAngle > 0.2: The pointing angle (pAngle) is a very strong variable to cut on, to suppress background vertices while simultaneously retaining signal vertices. It is defined by the angle difference between the displacement vector from the PV to the vertex and the momentum vector of the vertex. SM decays tend to have a small pAngle, since most of the daughter particles are visible. In contrast, decays associated with BSM particles have a large portion of their momentum carried away by the invisible neutralino, resulting in a larger pAngle. This cut allows us to remove many of the vertices that originate from SM decays of heavy hadrons.

- Vertex_newgoodtr >= 2: While the track cuts as described above are used to define so-called good tracks, we have to ensure that only vertices associated with those tracks are considered in the analysis. This cut is the minimal requirement for each vertex to consist of at least two tracks which are compatible with all of the track level cuts.

- acos(cos(Vertex_L_phi − Jet_phi[0])) > 1: The cut on the azimuthal angle between the displacement vector from the PV to SV and the leading jet is motivated in the same way as the corresponding track level cut. It is useful to ensure that the vertex is not related to the hadronic activity from the leading jet.

- acos(cos(Vertex_L_phi − MET_phi)) < 1.5: Likewise, limiting the angle between the displacement vector and the missing energy selects tracks that are more likely to have originated from events that are compatible with the signal we are searching for.

- Vertex_ndof > 1: The number of degrees of freedom of the vertex comes from the tracks that compose it. During the fitting, all of the tracks obtain weights, which correspond to the quality of their fit. The sum of all the track weights of a given vertex corresponds to the ndof variable, which provides insight into the quality of the reconstructed vertex. A cut on this variable ensures that vertices associated with badly reconstructed tracks are discarded.

After discussing all the relevant cuts for the analysis we can now take a closer look at the effects they have on signal and background samples. The following tables (Tab.5 and Tab.6) show the cut efficiencies for a specific signal mass point and the background samples consisting of the Zjets + Wjets datasets. The cut efficiencies refer to the fraction of vertices that are retained after each cut in relation to the number of vertices when no cuts are applied. All of the cuts have been applied cumulatively, meaning that all of the other cuts from previous rows have been applied as well.
From those tables, it is apparent that we start out with an overwhelming amount of background vertices, which we are able to effectively reduce with our combination of cuts to a level that is similar to the amount of signal vertices in the case of the chosen signal mass point.
We have refrained from creating a similar table for 2018 data in order to not accidentally unblind it.
In addition to showing the overall effect of the different cuts on the samples we would like to present some examples of different variables which are being used to discriminate background against signal. These upcoming figures (Fig.13, Fig.14, Fig.15) are so-called n-1 plots in which all of the cuts of the analysis have been applied except for one. The distribution for signal and background in that variable is plotted to visualize the effect of choosing a specific cutting value. All of these variables exhibit a similar feat, which is that in a certain parameter region, the background dominates over the signal, but also, there exist regions for which the opposite is the case. Choosing cut values to include only those regions is the main way of enhancing our signal-to-background ratio in this cut-and-count analysis.

applied cut (cumulative)	events	vertices	cut efficiency
no cuts	127	629	-
preselection	86	437	0.6946
abs(Track_dxySig) > 4	69	106	0.1681
abs(Track_dxy/Track_dz) > 0.25	59	86	0.1361
Track_pfRelIso03_all < 5	56	76	0.1214
$(\text{Track\_ptSig})^{-1} < 0.015$	55	76	0.1214
Track_numberOfValidHits > 13	47	60	0.0960
Track_normalizedChi2 < 5	47	60	0.0960
acos(cos(Track_phi-Jet_phi[0])) > 1	39	47	0.0758
acos(cos(Track_phi-MET_phi)) < 1.5	31	36	0.0567
acos(cos(Vtx_L_phi-MET_phi)) < 1.5	29	32	0.0514
acos(cos(Vtx_L_phi-Jet_phi[0])) > 1	29	32	0.0513
Vtx_pAngle > 0.2	24	26	0.0414
Vtx_ndof > 1	18	19	0.0304
Vtx_LxySig > 20	17	18	0.0283
MET_pt > 650	3	3	0.0055

Table 5: Cut efficiencies for MC signal using the stop production model with a mass of M=1000 GeV, mass splitting of ΔM=15 GeV and 20 mm proper lifetime. The sample has been scaled to an integrated luminosity of 59.8 fb^{-1}. Even though we start off with just a small number of signal vertices, we are able to retain many of them throughout the whole process.

applied cut (cumulative)	events	vertices	cut efficiency
no cuts	8 426 444	23 761 843	-
preselection	1 379 263	4 545 499	0.19129
abs(Track_dxySig) > 4	667 863	1 279 765	0.05386
abs(Track_dxy/Track_dz) > 0.25	378 471	545 319	0.02295
Track_pfRelIso03_all < 5	193 285	226 992	0.00955
$(\text{Track\_ptSig})^{-1} < 0.015$	193 221	226 914	0.00955
Track_numberOfValidHits > 13	92 469	100 564	0.00423
Track_normalizedChi2 < 5	92 468	100 563	0.00423
acos(cos(Track_phi-Jet_phi[0])) > 1	53 112	56 068	0.00236
acos(cos(Track_phi-MET_phi)) < 1.5	33 164	34 483	0.00145
acos(cos(Vtx_L_phi-MET_phi)) < 1.5	32 079	33 276	0.00140
acos(cos(Vtx_L_phi-Jet_phi[0])) > 1	32 066	33 263	0.00140
Vtx_pAngle > 0.2	12 438	12 654	0.00053
Vtx_ndof > 1	1 303	1 309	0.00006
Vtx_LxySig > 20	890	892	0.00004
MET_pt > 650	4	4	< 0.00001

Table 6: Cut efficiencies for MC background consisting of the Zjets and Wjets datasets scaled to 59.8 fb^{-1}. The much lower cut efficiencies compared to the previous table allow us to discard many more of the background vertices compared to signal vertices, leading to a manageable amount of background after all of the cuts have been applied.

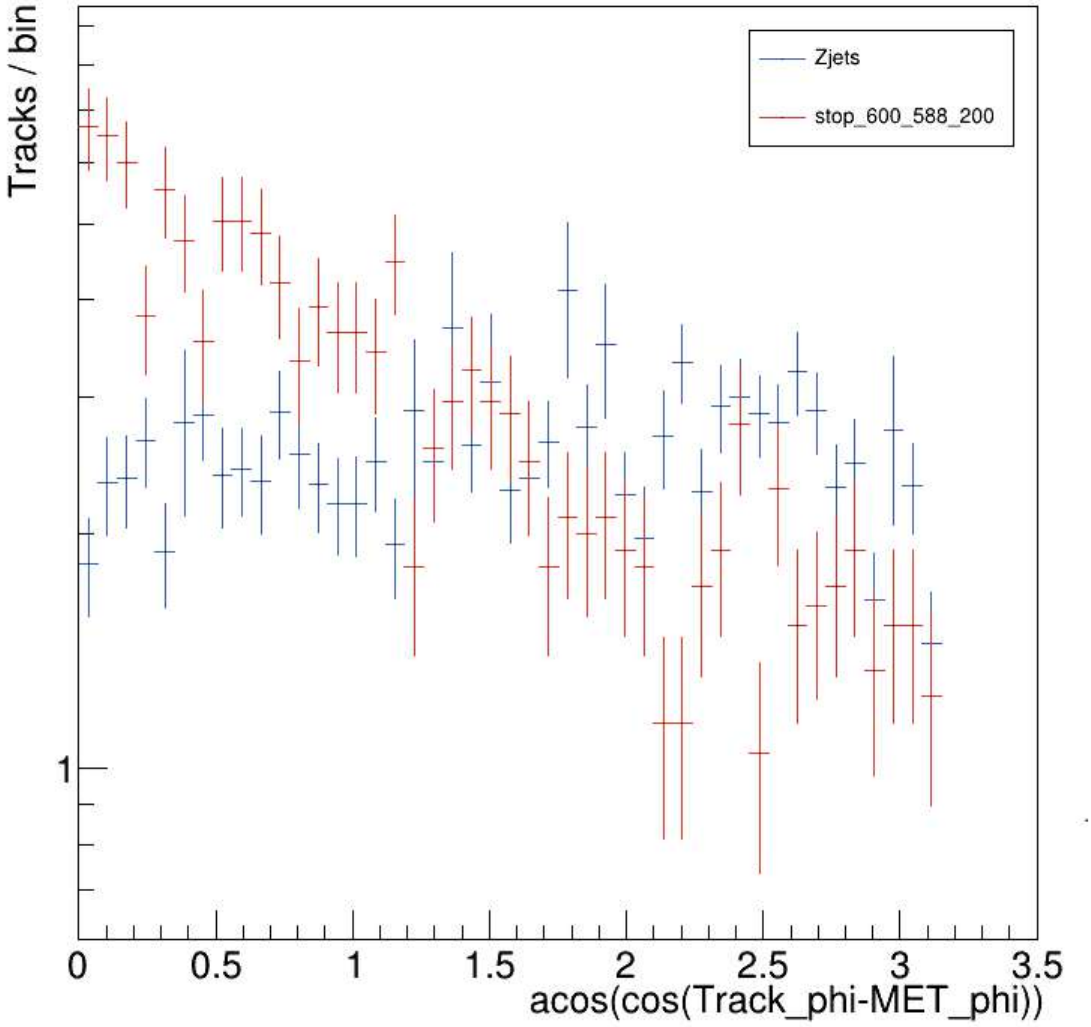

Figure 13: Distribution of the collinearity of the tracks and the missing energy after all other cuts have been applied. The cut which has been chosen for this analysis is discarding values larger than 1.5. This is at the breaking point where tracks originating from background processes become predominant over signal tracks. The shape of the background and signal curves are easily explained by the independence of the background processes with regard to the spatial distribution of missing energy, and our event signature favoring signal tracks to be spatially aligned with the missing energy of the event. Both background and signal have been scaled to an integrated luminosity of 59.8 fb^{-1}.

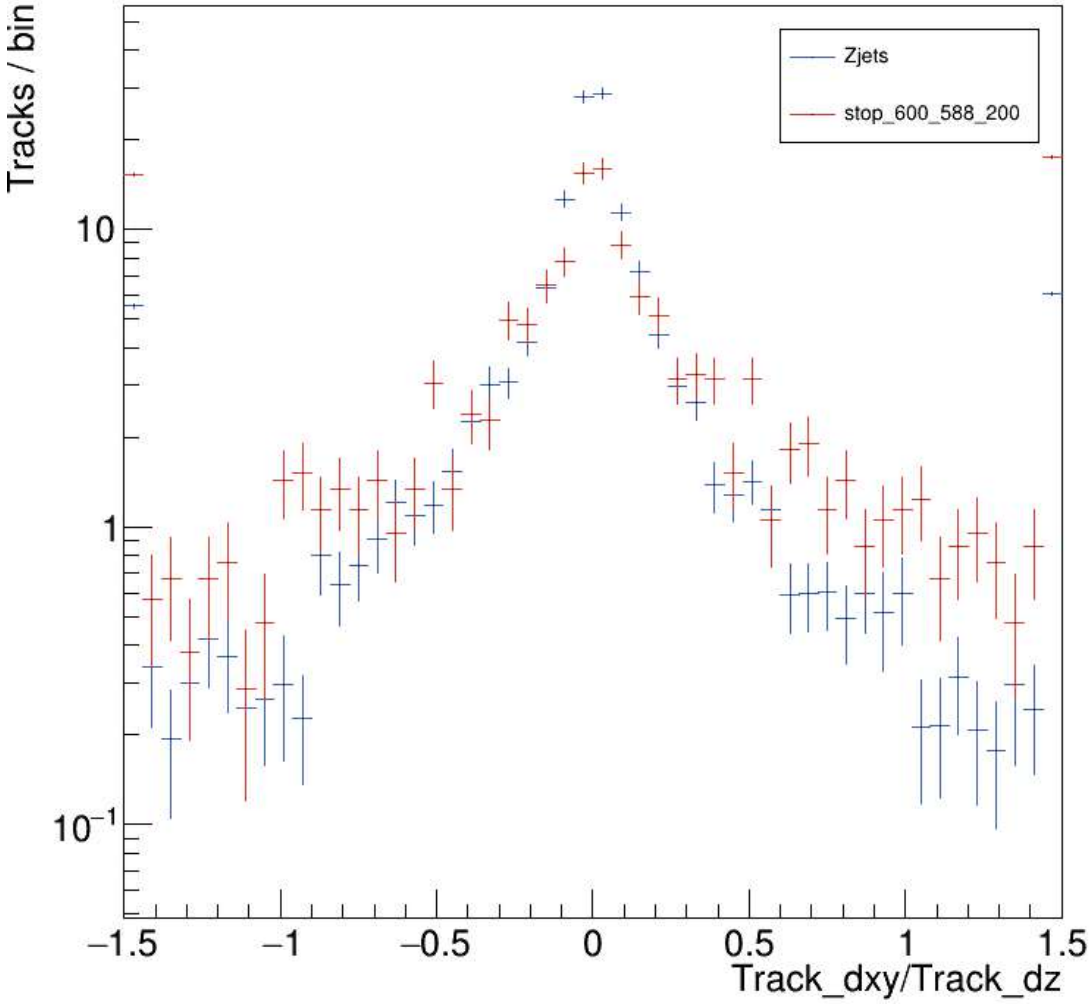

Figure 14: Distribution of Track_dxy over Track_dz, a powerful combination to discriminate against background while also retaining signal. Both signal and background peak at a value of zero, but the background distribution is more steep, leading to the signal becoming dominant over the background after a certain value. The cut for this variable has been set at 0.25 which leads to an enhancement of our signal to background ratio. It is also used to reject pile-up tracks. Both underflow and overflow have been moved to the first and last bin respectively. Background and signal have both been scaled to 59.8 fb^{-1}.

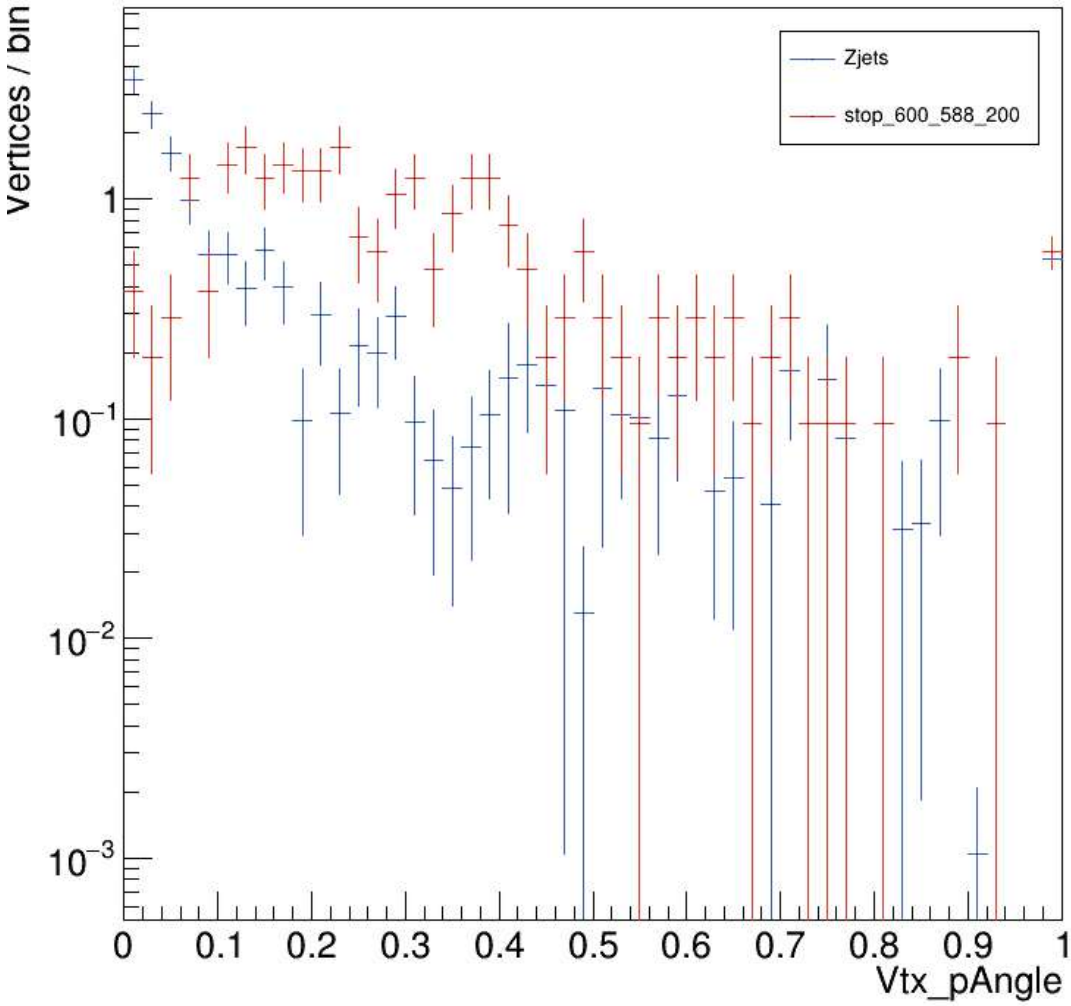

Figure 15: Distribution of the vertex pAngle after all other cuts have been applied. It is observed that background and signal vertices behave differently at low pAngle values. Background vertices peak around 0, while signal vertices peak at higher values. Leveraging this distinction enables us to significantly enhance the signal-to-background vertex ratio. A cut of 0.2 has been applied on this variable to suppress the background while simultaneously keeping a significant amount of signal. The overflow has been moved to the last bin. Background and signal have both been scaled to 59.8 fb^{-1}.

5.3 Background estimation

A widely used background estimation technique in high-energy physics for data-driven studies is the ABCD method [39]. In principle, it works by selecting a set of two different statistically independent variables or classifiers to construct four different regions. Three of those regions are so-called control regions (CRs) and are meant to be located in a parameter space where background processes dominate over the signal processes. The fourth region is called the signal region (SR) and is expected to favor signal over background events in order to achieve maximum sensitivity. The event yields from the experimental data in the control regions can be used to calculate a transfer factor to the signal region which is used to predict the number of background events in the signal region. This expected value can be compared to the actual observation in the SR in order to obtain cross-section limits and exclusion charts.
The predicted yield in the SR is given by

$$N_A = \frac{N_B N_C}{N_D} \tag{6}$$

where the regions B, C, and D are considered the control regions, region A corresponds to the signal region, and N_i is the event yield in each of those bins.
The ABCD regions are usually defined after all of the other cuts have been applied. Naturally, not all the available variables are equally useful in predicting the amount of background in the signal region. The variables used to define these regions need to be approximately uncorrelated for the prediction given by (6) to be accurate. This cannot be tested in the intended CRs/SR, since this would unblind the data. Instead, a different set of regions is being defined, which are so-called validation regions (VRs). Validation regions are typically derived by applying a similar selection procedure but inverting one of the selection criteria. In this study, the validation regions (VRs) are defined by inverting the requirement for the number of tracks that satisfy the track selection per vertex. Whereas vertices in the control regions (CRs) and signal regions (SRs) must have at least two of these "good" tracks, vertices will only qualify for the validation regions if they have one or fewer of these tracks. The exact definition of all the ABCD regions will be given in the upcoming section. In the validation regions the ABCD method can be tested for data and MC simulation by comparing the observed to the expected event yields. Closure can be confirmed for those regions if the predicted and observed values are within the uncertainties of each other. In that case, the variables are assumed to be approximately uncorrelated, and the ABCD method has been validated and can be applied.

5.4 Control/Signal and Validation regions

Three control regions and a signal region have been designed for this analysis (Fig.16). The first of the variables, that has been chosen to define the ABCD regions is the $L_{xy}^{significance}$, which is the ratio of the transverse displacement of the vertex and the uncertainty of it. The second variable is the missing transverse momentum of the event E_T^{miss}. As discussed previously, the four validation regions are obtained by inverting the cut on the number of good tracks per vertex. It should also be noted again that for 2018 data only the real event yields for the VRs and the CRs have been used. Since this thesis has been conducted in a way to avoid unblinding the data in the signal region, we have used an Asimov data set to obtain results from the experimental data. An Asimov dataset represents an idealized scenario in which the observed data perfectly fits the expected values for a given hypothesis [40]. In our case we can use the number of expected events in the signal region which we have obtained from applying the background estimation (Section 5.3) as the observed event yield in the signal region. This allows us to keep the data blinded while still making meaningful predictions and is a common technique in statistical analysis.

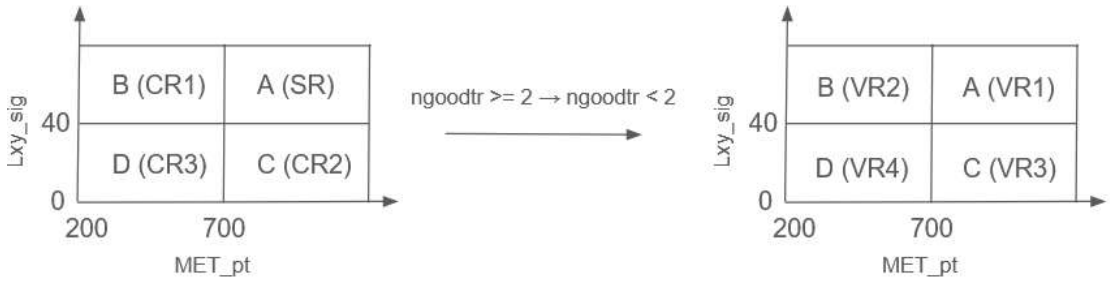

Figure 16: Definitions for the control regions (CRs), signal region (SR) and validation regions (VRs). After applying all of the cuts that have been described previously, a $L_{xy}^{significance}$ boundary of 40 and E_T^{miss} value of 700 have been chosen to split the parameter space into the three control regions and a signal region (left). The validation regions are obtained by applying the same cuts but inverting the requirement for the number of good tracks (tracks that pass all track-related cuts). This leads to another independent set of regions in the parameter space which are split into the four validation regions by the same $L_{xy}^{significance}$ and E_T^{miss} boundaries (right).

The boundaries of the signal region in E_T^{miss} and $L_{xy}^{significance}$ have been tested and optimized for the different signal mass points. During the earlier evaluation stages of the analysis a cut value of 650 for E_T^{miss} and 20 for $L_{xy}^{significance}$ was being used in the definition of the ABCD regions. Based on evaluations presented in Tab.7 a

E_T^{miss} boundary of 700 and a $L_{xy}^{significance}$ boundary of 40 has been chosen to represent the signal region and control regions. These boundaries have been optimized with the 1000 GeV LLP mass and 15/20 GeV mass splittings in mind, since these are the most crucial mass points in extending the exclusion limits. It has however been verified that an increase in $L_{xy}^{significance}$ boundary does not impact the sensitivity for longer lived samples in a negative way. The benefits of greatly increasing the sensitivity for those important mass points far outweigh the downsides.

LxySig \ MET_pt	400	450	500	550	600	650	700	750	800	850	900	950	1000
15	2.357	1.560	1.159	1.081	1.027	0.963	0.921	0.983	1.039	1.182	1.287	1.373	1.412
20	2.217	1.559	1.161	1.144	1.016	0.989	0.895	0.963	1.053	1.149	1.24	1.389	1.402
25	2.022	1.524	1.153	1.126	1.066	0.997	0.966	1.032	1.081	1.169	1.237	1.398	1.382
30	1.889	1.384	1.104	1.094	0.999	0.995	0.966	0.995	1.084	1.192	1.207	1.388	1.394
35	1.744	1.330	1.035	0.954	0.986	0.983	0.937	0.957	1.059	1.087	1.192	1.352	1.451
40	1.601	1.259	1.018	0.985	0.961	0.967	0.845	0.974	1.08	1.164	1.198	1.278	1.372
45	1.530	1.241	0.975	0.989	0.971	0.937	0.928	0.957	1.059	1.121	1.222	1.325	1.411
50	1.436	1.194	0.920	0.973	0.921	0.895	0.894	0.965	1.020	1.110	1.185	1.336	1.392
55	1.392	1.167	0.955	1.000	0.915	0.933	0.901	0.883	1.038	1.135	1.190	1.295	1.414
60	1.343	1.142	0.932	1.016	0.911	0.938	0.889	0.987	1.055	1.158	1.204	1.322	1.415

Table 7: Example table, showing different scale factors on the theoretical stop production cross-sections obtained using different signal region definitions in E_T^{miss} and $L_{xy}^{significance}$. The meaning of these scale factors is the following: The signal yields in the different regions of the ABCD method depend on the production cross-section of a specific signal process. Based on the event yields from data or MC background we can conclude how these production cross-sections would have to be scaled in order for the signal model to be excluded by 95% CL. The scale factors in this table describe this relationship. This table has been obtained using the MC background and the stop signal samples with a mass of 1000 GeV and a mass splitting of 15 GeV both scaled to 300 fb^{-1}. It can be seen that the choice of E_T^{miss} is of great importance; a value of 700 has been chosen as a boundary between control regions and signal region. The boundary for $L_{xy}^{significance}$ has been chosen at a value of 40, increasing the sensitivity further.

5.5 Closure

The closure for the regions defined in Section 5.4 has been tested for data from the CMS experiment during the 2018 data-taking period of Run II, with an integrated luminosity of 59.8 fb^{-1} and Monte Carlo simulated background from the Zjets and Wjets datasets. Closure is ensured if the predicted value by the ABCD method and the real observation lie within the uncertainties of each other. In the case of non-closure, which can happen if the two chosen variables are correlated with each other, closure can be restored by applying a residual correction also derived from simulations. This is not preferred and only works for small corrections thus limiting the usefulness of the ABCD method in certain scenarios but is not necessary for this analysis. It can be seen from Tab.8 that for both 2018 data and our MC simulation, closure is ensured in all cases.
A different observation that can be made from this table is that data and MC event yields differ quite substantially in the validation regions. This can be explained by our choice of MC background contributions. The Zjets and Wjets datasets provide us with an accurate description of the experimental data in our chosen Signal and Control regions. This has been tested and can be seen from Tab.9 and Tab.10. In the validation regions, however, other background processes with very different signatures that have not been included in this study also become more important. The deviation from data and MC background is, therefore, expected to increase in those regions. For the sake of closure this is not important and we can see that closure is ensured for MC in both cases of CRs/SR and VRs. For data this closure has only been tested in the validation regions.

	B	C	D	A (real)	A (pred.)
Data 2018 (VR)	40219 ± 201	499 ± 22	346799 ± 589	56 ± 7	58 ± 3
MC Zjets + Wjets (VR)	10222.0 ± 54.6	286.5 ± 3.5	86342.3 ± 178.3	35.1 ± 1.1	33.9 ± 0.7
MC Zjets + Wjets (CRs/SR)	639.1 ± 10.3	3.1 ± 0.4	926.3 ± 18.4	2.4 ± 0.2	2.1 ± 0.3

Table 8: Closure has been tested for both 2018 data and MC simulation. For data this was done in the validation regions, for MC it has been tested with the background scaled to 59.8 fb^{-1} in the validation regions and the MC background scaled to 300 fb^{-1} in the control regions and signal region. The regions follow the description which has been given in Fig.16. Closure can be confirmed in all of the tested scenarios.

5.6 Statistical methods

Combine is a software tool used for statistical analysis of data obtained from the CMS experiment [41, 42]. It was originally developed by the Higgs Physics Analysis Group (PAG) but has since been adapted and is widely used within the CMS community to obtain exclusion limits. It is integrated into CMSSW, making it a versatile and quick-to-use tool. Combine implements a number of different statistical methods and can be used for both binned and unbinned, shape and non-shape-based analysis. The nature of our ABCD method makes this analysis a binned and non-shape-based analysis, since we are only interested in the total event yields in each region and not the specific distribution of the events inside each bin. The binning corresponds to the different regions in the parameter space defined by the analysis and is, in our case, determined by our control and signal regions. The Combine tool uses datacards as inputs, which are a compact version of the results and statistical models used in the analysis. They can range from simple one-bin experiments to very complex analyses involving multiple different control and signal regions, taking into account various sources of systematic uncertainties. An example of a datacard used in this analysis, consisting of our four regions defined by the ABCD method, can be seen in Fig.17. The most important parts of these datacards are the number of observed events in the different regions as well as the rates and sources of uncertainty for each of the processes. The rates reflect to the number of expected events from a given process in each bin, the number of observed events corresponds to an Asimov dataset.

In the first step of the analysis, where we calculate limits based on 2018 data, the observations correspond to the number of events from this data-taking period after all cuts have been applied. The signal region is an exception. Here the observation has been calculated from Eq.(6) in order to not unblind the data. The expected and observed yields are the same in the control regions. The relationship between the expected yields in the control regions and signal region has been included as a rate parameter which will be fit by the model.
In the case where limits have been calculated from the MC-generated background, both the observed and expected number of events have been obtained by applying all cuts to these datasets. For consistency, the observed yield in the signal region has also been calculated from Eq.(6). The estimation for the number of predicted events in the signal region is again based on the same expression and included as a rate parameter.
In both cases, the expected signal yields have been obtained by applying all of the cuts to the MC-generated signal samples for a specific signal model. The meaning of rate parameters and the statistical methods implemented by the Combine tool will be explained in the upcoming paragraphs.

```
imax 4  number of channels
jmax 1 number of processes -1
kmax *  number of nuisance parameters (sources of systematical uncertainties)
-------
bin             A       B       C       D
observation     5       887     4       677
-------
bin             A       B       C       D       A       B       C       D
process         bkg     bkg     bkg     bkg     sig     sig     sig     sig
process         1       1       1       1       0       0       0       0
rate            1.0     1.0     1.0     1.0     15.93   115.74  0.29    2.00
-------
SigSys lnN              -       -       -       -       1.2     1.2     1.2     1.2
SigSysA lnN             -       -       -       -       1.035   -       -       -
SigSysB lnN             -       -       -       -       -       1.015   -       -
SigSysC lnN             -       -       -       -       -       -       1.577   -
SigSysD lnN             -       -       -       -       -       -       -       1.202
BkgSys lnN              1.2     -       -       -       -       -       -       -
BkgSysA lnN             1.083   -       -       -       -       -       -       -
BkgSysB lnN             -       1.016   -       -       -       -       -       -
BkgSysC lnN             -       -       1.104   -       -       -       -       -
BkgSysD lnN             -       -       -       1.023   -       -       -       -

alpha rateParam A bkg (@0*@1/@2) beta,gamma,delta
beta rateParam B bkg 886.6679625714766
gamma rateParam C bkg 3.559188775636136
delta rateParam D bkg 676.6063875981844
```

Figure 17: Example of a datacard used in the analysis. The number of channels refers to the number of different regions in which event yields have been counted, which is, in our case, four (three control regions and the signal region). We split the contributions in those regions into contributions from the background and from the signal. The observations have been taken from MC background/data. The rates correspond to the number of expected events in each region and have been set to one for the background contributions. The actual expected values from the background are defined as rate parameters at the bottom of the datacard. This is done to allow the expected yields from the background processes in the control regions to be a model parameter that will be fitted by the Combine tool and is important for our background estimation in the signal region. The uncertainties have been calculated for all processes and regions independently and an additional 20% systematic uncertainty has been applied to all of the processes.

Combine uses likelihoods for statistical analysis. A likelihood refers to the probability of obtaining observed data under the assumption of a specific underlying model $\boldsymbol{\theta}$ which can be represented as $\mathcal{L}(\boldsymbol{\theta}|\text{data})$.
The model $\boldsymbol{\theta}$ is defined by a variety of properties: In addition to the parameters of the model, consisting of parameters of interest $\boldsymbol{\mu}$ and nuisance parameters $\boldsymbol{\nu}$, it also encapsulates the relationship between these variables.

For a counting experiment like the one we are conducting, the likelihood can be constructed from a simple poissonian distribution in each analysis bin [41]:

$$\mathcal{L}_{\text{bin}}(\boldsymbol{\mu}, \boldsymbol{\nu}) = \text{Poiss}(\text{n}_{\text{obs}}, \text{n}_{\text{exp}}(\boldsymbol{\mu}, \boldsymbol{\nu})) \tag{7}$$

where $\text{Poiss}(\text{n}_{\text{obs}}, \text{n}_{\text{exp}}(\boldsymbol{\mu}, \boldsymbol{\nu}))$ is a poissonian distribution with mean $\text{n}_{\text{exp}}(\boldsymbol{\mu}, \boldsymbol{\nu})$, evaluated at n_{obs}.

The total likelihood can be obtained by multiplying the likelihood functions for all bins.
It can be seen that in this case the model dependence only stems from the number of expected events in each bin. In the case of this analysis there is only one parameter of interest in our model: it is the signal strength, a parameter that scales the signal event contribution in each bin. This can be explained by the following example: if our datacard contains an observation of 10 events in bin A and a rate of 10 for our signal process in bin A, a signal strength of $\mu = 1$ would mean that in this model all 10 observed events are signal events, while a signal strength of $\mu = 0.5$ would mean that only half of the observed events are signal events. We can later relate the limits for the signal strength directly to our scale factors for the production cross-section and can exclude cross-sections with scale factors higher than the signal strength.
All of the other parameters of our model, like the statistical and systematic uncertainties and the rate parameters for the background processes, which act in a similar way as the signal strength but relate to the background processes instead, are nuisance parameters. This is because we are only interested in the post-fit values of the signal strength to make statements about the signal production cross-sections.
The fitting of the likelihood function to the observed data follows the maximum likelihood approach. This means that the likelihood is evaluated over a large area of the parameter space $(\boldsymbol{\mu}, \boldsymbol{\nu})$ and the best fit is obtained by:

$$(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}}) = \underset{\boldsymbol{\mu}, \boldsymbol{\nu}}{\text{argmax}}\ \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\nu}) \tag{8}$$

In order to test different models with varying parameters against each other, a test statistic is needed. A test statistic is a mathematical function that maps a model and the observed data to a real value $\text{x} \in \mathbb{R}$. Not only is the test statistic useful for comparing different models, but in conjunction with the distribution of the test statistic, we are even able to obtain p-values for our model hypothesis. The Combine tool allows for the use of different test statistics, the one that has been used for this analysis is the following:

$$t_\mu = -2 \ln \left(\frac{\mathcal{L}(\boldsymbol{\mu})}{\mathcal{L}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\nu}})} \right) \tag{9}$$

Here, the numerator corresponds to the profile likelihood, where $\boldsymbol{\mu}$ is a freely floating parameter and the nuisance parameters $\boldsymbol{\nu}$ have been profiled out for a specific value of $\boldsymbol{\mu}$. The profile likelihood is an important concept to reduce the complexity of the likelihood function to only the parameters of interest. As already discussed, the only parameter of interest in our case is the signal strength, and the profile likelihood can be obtained from:

$$\mathcal{L}(\boldsymbol{\mu}) = \max_{\boldsymbol{\nu}} \mathcal{L}(\boldsymbol{\mu}, \boldsymbol{\nu}) \tag{10}$$

The denominator of Eq.(9) corresponds to the likelihood obtained from using the best-fit parameters obtained from Eq.(8).

This allows us to associate a test statistic to different models which are set apart by their model parameters. In order to obtain the p-values associated with the values of the test statistic, we require the distribution of the test statistic given a specific model. Obtaining this distribution is not an easy task, and there are two different approaches implemented in the Combine tool that have been used in this analysis to acquire results.
The first one is called "AsymptoticLimits" and can be used as a quick check to obtain an estimate of the observed and expected limits. Its accuracy is limited by the assumption that the event yields are not too small and the systematic uncertainties are not a large factor on which the result depends. The "AsymptoticLimits" method is based on Wilks theorem, which states that the distribution of the test statistic described in Eq.(9) follows a chi-squared distribution in the asymptotic limit, where the sample size approaches infinity.
The number of degrees of freedom, which is the parameter that defines the chi-squared distribution, is then the number of parameters of interest and is one in our case. The p-value associated with this model can simply be obtained by evaluating this chi-squared distribution at the value of the test statistic for that model. This approximation can be used to calculate two p values p_b and p_μ. Here, p_b corresponds to the p-value under the background-only hypothesis, which can be obtained from the test statistic in Eq.(9) by simply setting the signal strength $\mu = 0$ and p_μ refers to the p-value for the signal + background hypothesis with a signal strength of μ.
A simple way to construct confidence intervals would be by examining the p-values p_μ. For any given confidence level CL, we can set a limit for the signal strength μ by finding the smallest value of μ that satisfies $\mu = 1 - \mathrm{CL}$. This will ensure that all larger values of μ will have a smaller p-value associated with them and we can

exclude them with confidence, since $\mu < 1 - \mathrm{CL}$.
It is, however, common in high-energy physics experiments to construct exclusions differently. The CLs criterion utilizes both p-values $\mathrm{p_b}$ and p_μ. The Combine tool takes this into account and calculates excluded values of μ by requiring:

$$\frac{\mathrm{p}_\mu}{1 - \mathrm{p_b}} < (1 - \mathrm{CL}) \tag{11}$$

The reason for this approach lies in the unique challenges of particle physics experiments. By incorporating the uncertainties of both signal and background models, a more robust approach that is less sensitive to fluctuations and takes systematic uncertainties caused by the complex detectors better into consideration can be ensured.

The "AsymptoticLimits" method has been used to give quick estimates in this analysis. The main purpose has, however, been to calculate rough estimations for starting values of the parameters of more complex and accurate models. The one that has been used in this analysis is Combine's "HybridNew" method.
The "HybridNew" method is a more sophisticated approach to calculating limits on the signal strength. It can be used to compute fully frequentistic limits (LHC-style), which is the preferred approach, or hybrid bayesian-frequentistic limits, which has previously been the preferred approach at other particle experiments like LEP or Tevatron. Depending on the input model, these methods can be both time and resource-intensive but usually yield better and more accurate limits.
The "HybridNew" method achieves this by using toys for the limit calculations. Toys are randomly generated datasets that are based on the model and distribution of the parameters. Each of those toy datasets is subjected to the same statistical analysis procedures, like fitting the models to the data and computing the test statistic. The test statistic for the LHC-style method is equal to the test statistic of the "AsymptoticLimits" method and is defined by Eq.(9). By generating a large amount of toys, it is possible to compute the full distribution of the test statistic for any given model. Having access to this distribution allows us to omit Wilks theorem and we are no longer limited by the assumption of the asymptotic limit. Consequently, the "HybridNew" method yields much more accurate results, especially when dealing with low event yields. After obtaining the p-values for different parameter models, the confidence intervals are constructed in the same way as the "AsymptoticLimits" method, also taking the CLs criterion described in Eq.(11) into account. The speed and robustness of calculating limits with this method can be highly influenced by the starting points of the parameters. It is therefore recommended to use it in conjunction with a hint method like "AsymptoticLimits" in order to achieve maximum performance.

6 Results

In this chapter, we present the results from our cut-and-count analysis applied to two distinct benchmark signal models: the bino-stop model and the C1N2 production model. These models are detailed in Chapter 4.1. The objective of this thesis was to establish and validate a cut-and-count analysis utilizing the event signature ISR + MET + very soft tracks (down to 0.5 GeV), which converge at a common vertex as previously described. We selected these benchmark models to demonstrate the efficacy of this approach.
Although the presented signal models have been delved into in previous studies, the event signature is novel and not yet explored. The cut-and-count analysis conducted in this thesis is also meant as a benchmark for the still ongoing Soft Displaced Vertex analysis which is trying to extract the maximum of this novel approach by applying sophisticated machine learning techniques.
The idea of using very soft tracks to reconstruct displaced vertices without the need of any identified SM objects in the detector signature aims to push the CMS experiment to its technical limitations in its resolution and is expected to be applicable to a variety of different BSM scenarios which contain long-lived particles, like compressed model SUSY spectra and dark sector models.

The results of this study will be presented in two different types of key plots. Three different LLP masses (600 GeV, 1000 GeV, 1400 GeV) and four different mass splittings (12 GeV, 15 GeV, 20 GeV, 25 GeV) have been simulated to obtain them. The results are also split into two parts, the first is obtained by applying the ABCD method to 2018 data, which amounts to a total integrated luminosity of 59.8 fb^{-1}. The simulations involving this have been conducted in a way that avoids unblinding the experimental data. Instead the Asimov dataset has been used, for which both the expected and observed event yields in the signal region are estimated by the yields in the control regions.
Additionally, we have run the analysis framework constructed in the process of this thesis on the MC background samples, consisting of the Zjets and Wjets datasets scaled to an integrated luminosity of 300 fb^{-1}. This is done to project the improvement that will be gained from including the full Run III data, which will result in a total integrated luminosity of 300 fb^{-1} once it has finished. For consistency, instead of using the real event yields in the signal region as an observation, the same procedure of using the Asimov dataset has also been applied to the MC background.

	median	68% band	95% band
Zjets + Wjets	0.286	0.401	0.598
		0.169	0.119
2018 Data	0.330	0.488	0.747
		0.256	0.159

Table 9: Comparison between the scale factors for the theoretical stop production cross-sections obtained from MC background and from 2018 data. 68% and 95% values have been separated into upper (top) and lower (bottom) values for the bands. The signal sample that has been used for this table is the stop with M($\tilde{t}$)=600 GeV and $\Delta M(\tilde{t} - \tilde{\chi}_1^0)$=12 Gev, c_τ=200 mm

	median	68% band	95% band
Zjets + Wjets	18.909	26.556	36.370
		13.373	10.278
2018 Data	25.182	35.238	51.215
		18.011	11.861

Table 10: Comparison between the scale factors for the theoretical C1N2 production cross-sections obtained from MC background and from 2018 data. 68% and 95% values have been separated into upper (top) and lower (bottom) values for the bands. The signal sample that has been used for this table is the C1N2 production model with an LLP mass of M($\tilde{\chi}_2^0$)=600 GeV and $\Delta M(\tilde{\chi}_2^0 - \tilde{\chi}_1^0)$=12 Gev, c_τ=200 mm

For both data and MC, the predicted values are close to the real values, which have been tested and confirmed by checking the closure for both of them (Section 5.5). To emphasize the agreement between data and MC background we show two tables (Tab.9) and Tab.10) that illustrate the obtained cross-section limits

from using both data and MC background scaled to 59.8 fb^{-1} for a specific point in the parameter space. It can be seen that for both the bino-stop and the C1N2 model, the values from data and MC background are in very good agreement with each other. The expected signal yields in all cases are obtained from the signal MC production as described in Section 4 and are appropriately scaled to the same integrated luminosity.

All of the data, MC background, and the various MC-generated signal models underwent the identical cuts as detailed in Chapter 5.1. The resulting event yields, along with their associated uncertainties, were utilized as inputs for the Combine datacards. Additionally, a 20% systematic uncertainty was applied uniformly across all event yields. Using the statistical methods detailed in Section 5.6 we derived scale parameters for the production cross-sections. These scale parameters are the basis of the limit and exclusion calculations. They describe how the production cross-section of a particular process needs to be scaled in order for the results to be in accordance with the background-only hypothesis. This means that any value for the production cross-section that is larger than its scaled value would be excluded at 95% confidence level. These scaled production cross-sections are, therefore, the limits for the production cross-section obtained by the analysis. They have been calculated for all of the different signal models with varying model parameters (the mass of the LLP and the mass splitting to the LSP) by applying the steps of the whole analysis framework to the different signal samples.
The results of applying the cut-based analysis designed in this thesis to 2018 data have been summarized in Fig.18 and Fig.20 for the bino-stop model and the C1N2 model respectively.
Fig.19 and Fig.21 show the production cross-section limits obtained from the projection to 300 fb^{-1} using MC background and signal.
These plots can be understood in the following way: Regions in the parameter space where the scale parameters are less than one, correspond to areas in those plots underneath the theoretical production cross-section curves. In those areas, it can be concluded from our results that the theoretical production cross-sections are too high to be in accordance with the observation, and our analysis is said to be sensitive. In regions where the obtained scale parameters are larger than one, the analysis is not sensitive and no statements can be made based on our results. The theoretical cross-section curves depicted in these figures have been taken from literature, and correspond to the same values that have been used for signal event generation.

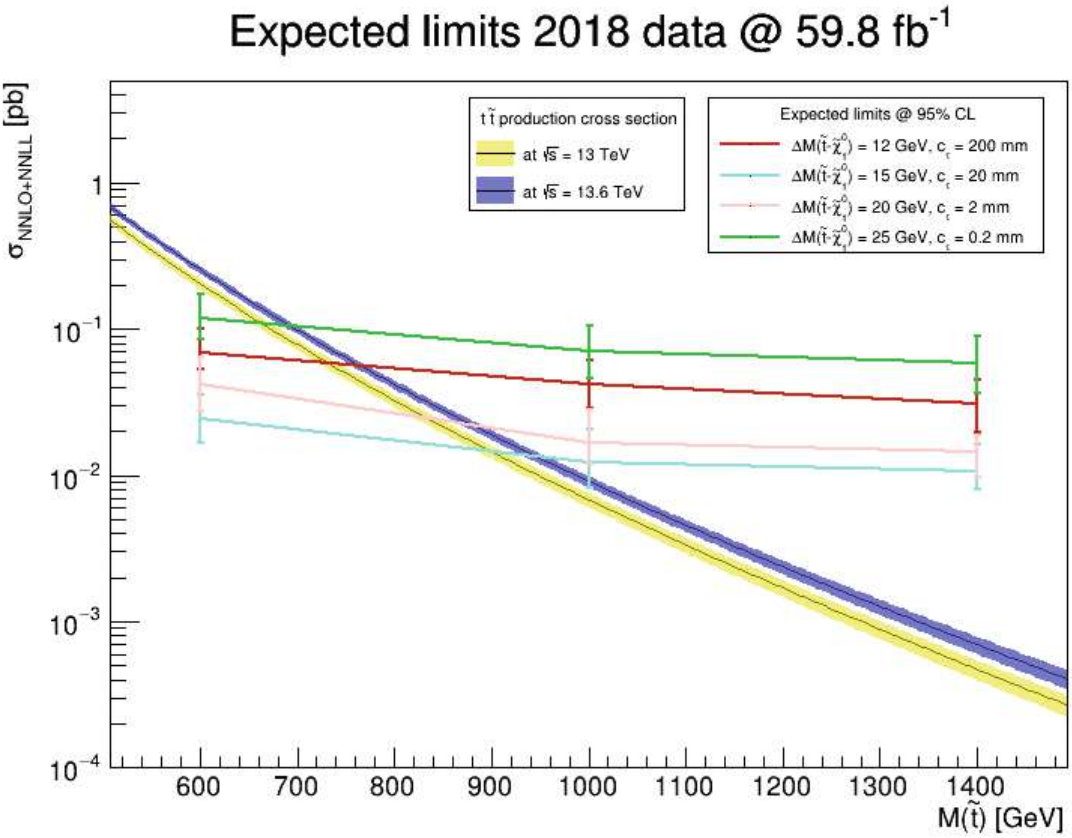

Figure 18: Limits for the stop production cross-section. The x-axis depicts the mass of the top squark. On the y-axis the production cross-section is plotted. The blue and yellow lines represent the theoretical stop production cross-sections at center of mass energies $\sqrt{s} = 13$ TeV (Run II) and $\sqrt{s} = 13.6$ TeV (Run III), respectively. The values for those cross-sections have been obtained from NNLO (next to next to leading order) + NNLL (next to next to leading logarithmic) calculations. The shaded areas represent the statistical uncertainties of the theoretical values. The rest of the colored lines show the limits on those cross-sections by the analysis obtained from 2018 data (59.8 fb^{-1}) and using signal samples with varying model parameters. The uncertainties are indicated by the error bars and correspond to the 1σ band.

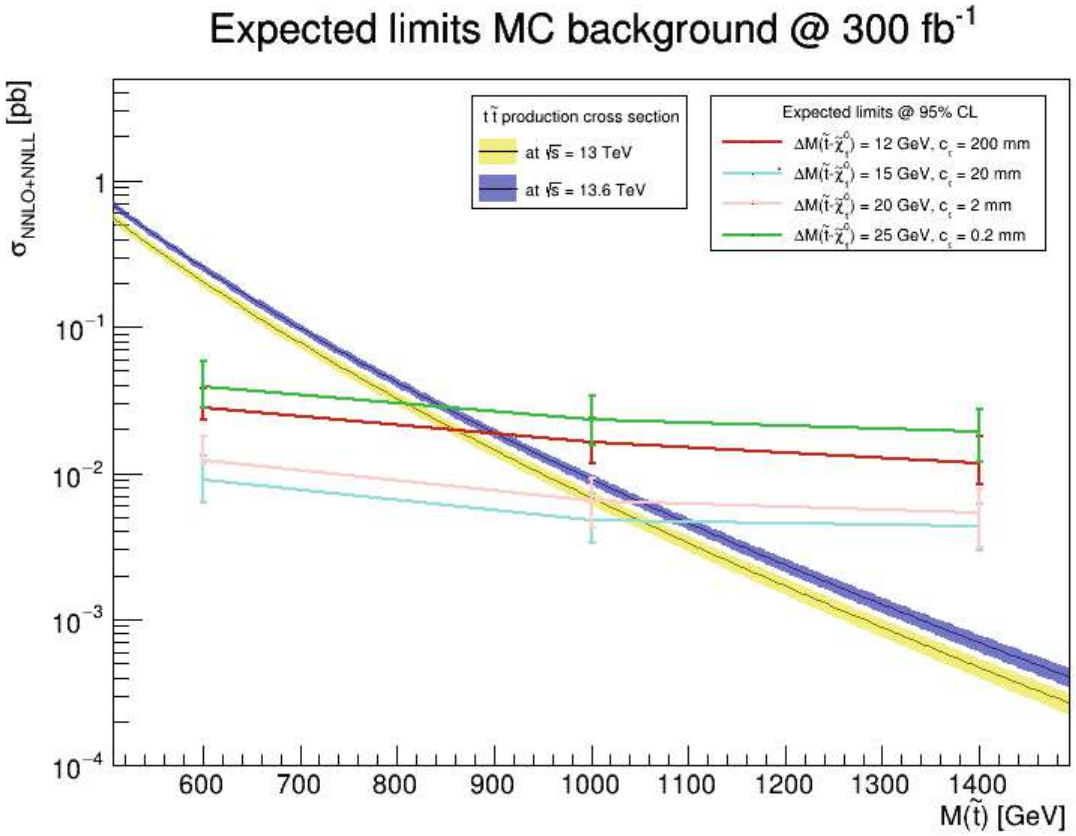

Figure 19: Limits for the stop production cross-section. The x-axis depicts the mass of the top squark. On the y-axis the production cross-section is plotted. The blue and yellow lines represent the theoretical stop production cross-sections at center of mass energies $\sqrt{s} = 13$ TeV (Run II) and $\sqrt{s} = 13.6$ TeV (Run III,) respectively. The values for those cross-sections have been obtained from NNLO (next to next to leading order) + NNLL (next to next to leading logarithmic) calculations. The shaded areas represent the statistical uncertainties of the theoretical values. The rest of the colored lines show the limits on those cross-sections by the analysis obtained from the MC background, which has been scaled to 300 fb^{-1} using signal samples with varying model parameters. The uncertainties are indicated by the error bars and correspond to the 1σ band.

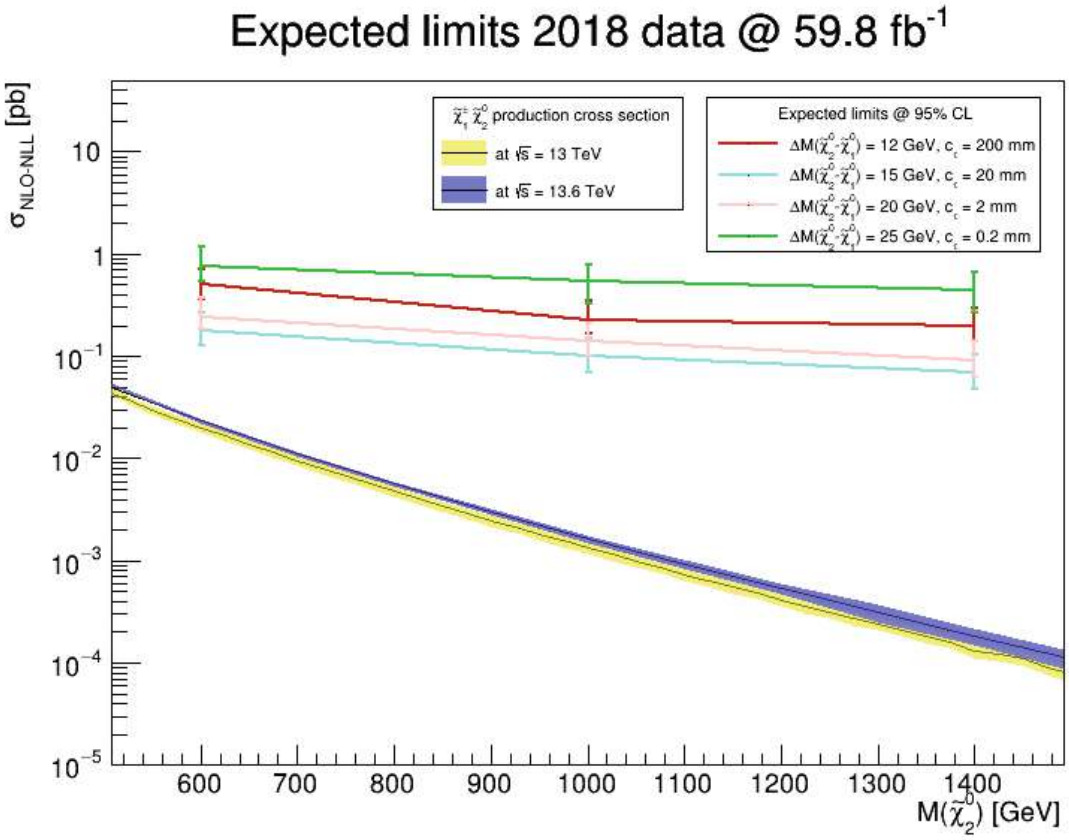

Figure 20: Limits for the C1N2 production cross-section. On the x-axis the mass of the second lightest neutralino is applied. The y-axis depicts the production cross-section. The blue and yellow lines represent the theoretical C1N2 production cross-sections at center of mass energies $\sqrt{s} = 13$ TeV (Run II) and $\sqrt{s} = 13.6$ TeV (Run III), respectively. The values for those cross-sections have been obtained from NLO (next to next to leading order) + NLL (next to leading logarithmic) calculations. The shaded areas represent the statistical uncertainties of the theoretical values. The rest of the colored lines show the limits on those cross-sections by the analysis obtained from 2018 data (59.8 fb^{-1}) and using signal samples with varying model parameters. The uncertainties are indicated by the error bars and correspond to the 1σ band.

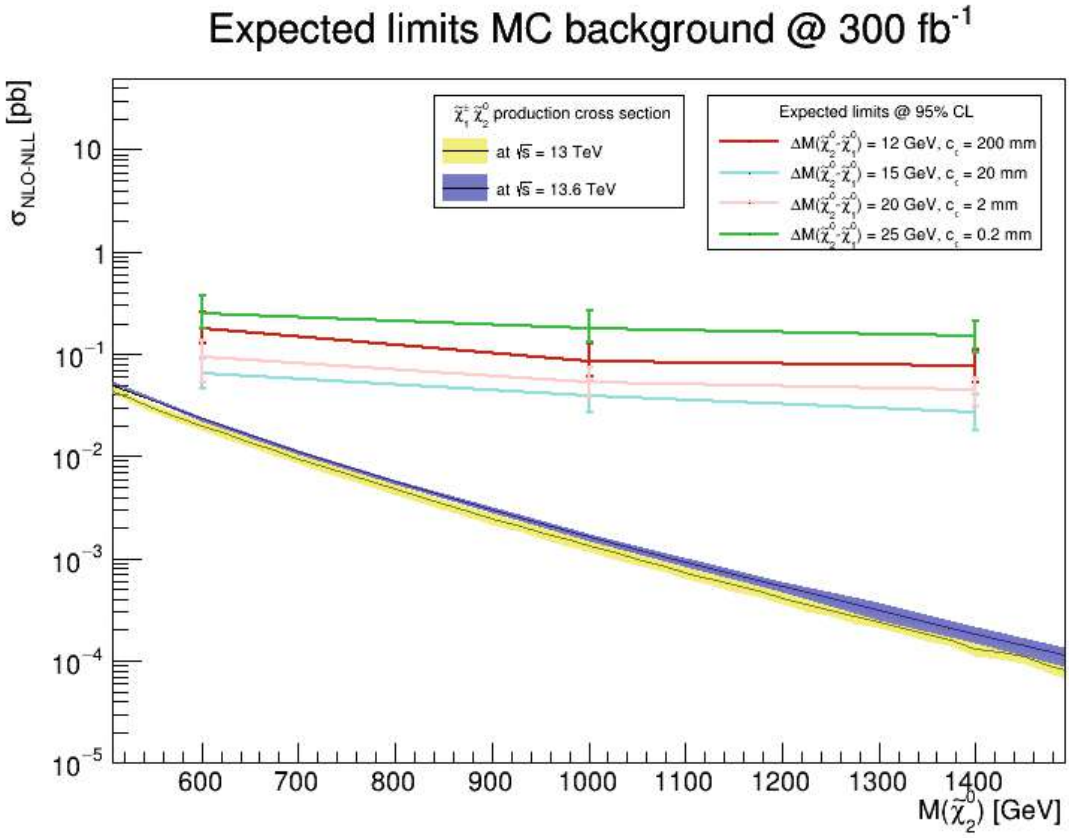

Figure 21: Limits for the C1N2 production cross-section. On the x-axis the mass of the second lightest neutralino is applied. The y-axis depicts the production cross-section. The blue and yellow lines represent the theoretical C1N2 production cross-sections at center of mass energies $\sqrt{s} = 13$ TeV (Run II) and $\sqrt{s} = 13.6$ TeV (Run III), respectively. The values for those cross-sections have been obtained from NLO (next to next to leading order) + NLL (next to leading logarithmic) calculations. The shaded areas represent the statistical uncertainties of the theoretical values. The rest of the colored lines show the limits on those cross-sections by the analysis obtained from the MC background, which has been scaled to 300 fb^{-1} using signal samples with varying model parameters. The uncertainties are indicated by the error bars and correspond to the 1σ band.

From these initial results we can construct another set of plots to gain further insight into the regions of the parameter space which can be excluded from data and simulation. The intersection points between the theoretical production cross-section curves and the analysis-derived limits delineate the boundary between excluded and not excluded regions in the parameter space. These intersection points were obtained by linear interpolation of the analysis limits. Calculating the

intersections for all of the different signal models gives us a number of different points defining the border of exclusion. These points have then been interpolated using spline interpolation to obtain Fig.22 and Fig.23. These plots depict the excluded areas in the parameter space for stop production and the C1N2 model, respectively. To highlight the improvements that can be expected from the increase in integrated luminosity in Run III both exclusion from 2018 data and the MC projection have been plotted in the same figure.

The difference between the two different signal models is expected and can easily be explained. The production cross-section for C1N2 production is noticeably smaller than for stop production. In addition to that, the C1N2 model involves the production of a promptly decaying chargino, leading to only one vertex in the final state (Fig.11), while the four-body stop decay contains two displaced vertices in its final state (Fig.10). The result is a considerable loss of sensitivity for the C1N2 model compared to the stop production model.

Despite this, these results show that this analysis, which utilizes a novel and not yet explored event signature can achieve good sensitivity to both of the chosen benchmark models over a large parameter space of LLP mass and mass splittings. Even though those models are not novel, these results are able to provide a proof of concept, and the framework is expected to be applicable to a variety of different BSM scenarios which include long-lived particles.

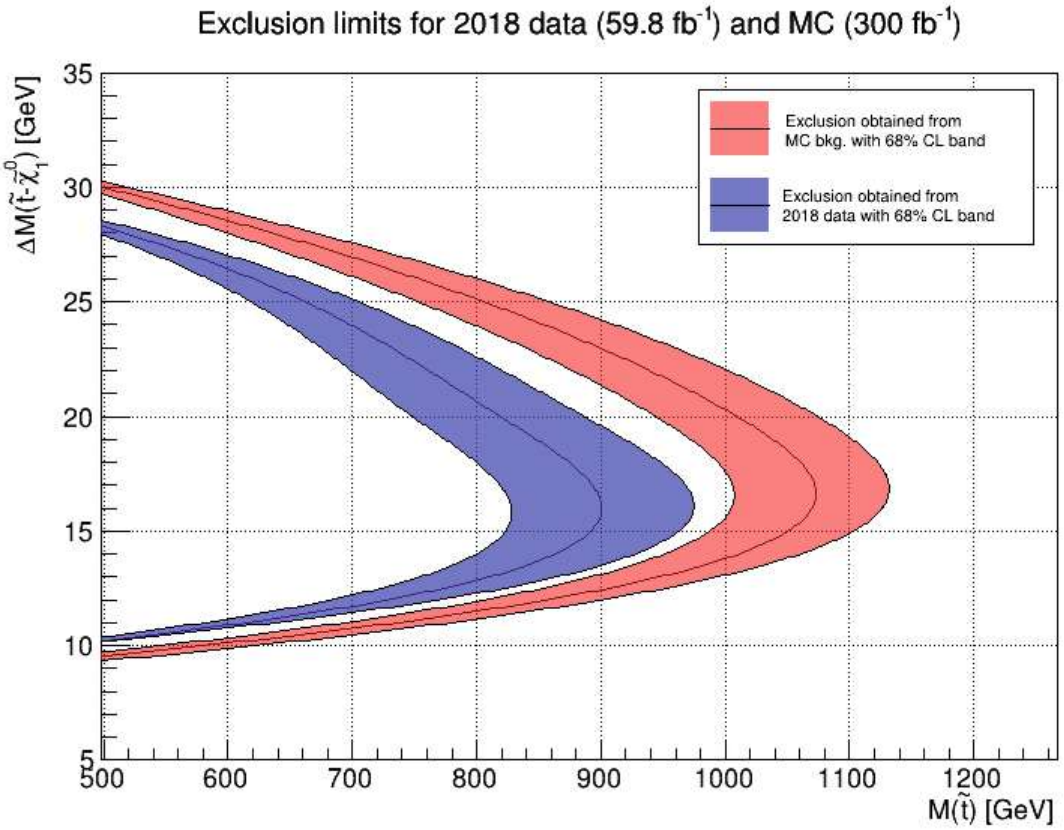

Figure 22: Exclusion chart for stop production. The x-axis shows the mass of the LLP, which, in this case, is the supersymmetric partner of the top quark (stop). On the y-axis, the mass splitting to the LSP (in this model, the lightest neutralino) is plotted. The blue line represents the exclusion obtained from 2018 data (59.8 fb^{-1}), the red line has been obtained from the MC background (300 fb^{-1}). The shaded areas represent the 68% CL level band of the results. The curves have been obtained using cubic spline interpolation, areas to the left of the curves can be excluded by the analysis.

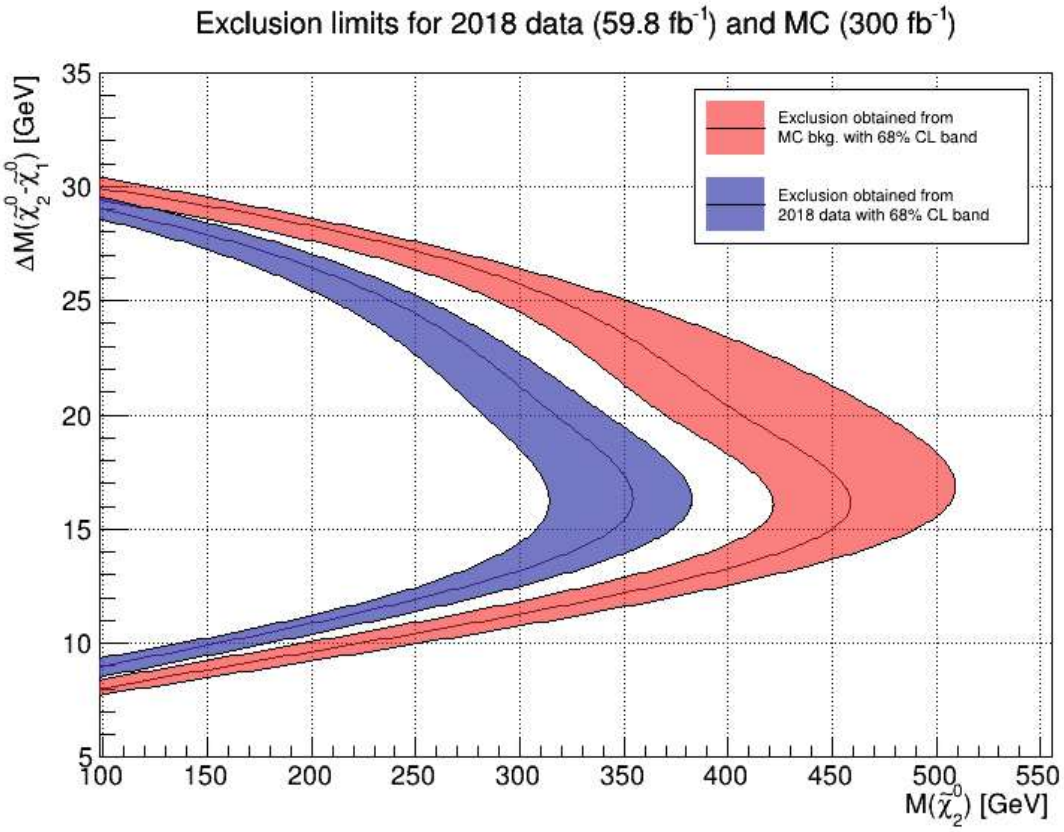

Figure 23: Exclusion chart for C1N2 production. The x-axis shows the mass of the LLP, which in this case is the second lightest neutralino. On the y-axis, the mass splitting to the LSP (in this model, the lightest neutralino) is plotted. The blue line represents the exclusion obtained from 2018 data (59.8 fb^{-1}), the red line has been obtained from the MC background (300 fb^{-1}). The shaded areas represent the 68% CL level band of the results. The curves have been obtained using cubic spline interpolation, areas to the left of the curves can be excluded by the analysis.

7 Outlook

The aim of this study was to establish a cut-and-count analysis based primarily on 2018 data from the CMS experiment while also providing an outlook on the improvements anticipated for Run III at the LHC. Our projections for Run III are primarily based on the increased integrated luminosity compared to Run II. However, despite the promising results presented in this thesis, Run III will bring enhancements beyond just increased luminosity.

A significant improvement for this analysis may come from the changes made to the CMS tracking process [43]. As previously discussed, tracking in CMS is an iterative procedure. During Run II, a process with twelve main iterations, each targeting different track topologies, was employed. This iterative tracking was predominantly seeded by the silicon pixel detector, the innermost module of the detector, which is suboptimal for investigating displaced vertices with very large displacements. In such cases, other iterations of the tracking process are responsible for the efficient reconstruction of displaced tracks and, consequently, the vertices that they originate from.
For Run III, the tracking process has been significantly revised. A new algorithm featuring only six tracking iterations has been developed. This will not only provide much better computational performance, but the inclusion of a new pixel-less iteration seeded by silicon strip triplets will also increase the tracking efficiencies for tracks with larger displacements [44]. This improvement is expected to greatly benefit this analysis by providing a better resolution and quality of reconstructed tracks with large displacements.

In addition to the improvements made by the CMS collaboration, the analysis limits could be further enhanced by incorporating more sophisticated methods, such as machine learning techniques. The variables used in this study to distinguish signal from background vertices provide valuable insights into their relationship. However, this separation could be significantly improved by utilizing multivariate (MVA) discriminants instead of simple cuts on individual variables. ML techniques can exploit correlations between different variables in the parameter space and incorporate additional variables that may initially seem poor discriminators. Additionally, ML could enhance vertex and track-finding capabilities within the workflow. All of these options are currently being explored by the Soft Displaced Vertices work group [31, 32].

8 Conclusion

In this thesis, the framework and workflow of a cut-based and data-driven analysis focusing on a novel event signature have been presented. The explored signature relies on ISR, very soft tracks that form a common vertex that is measurably displaced, and most of the energy being carried away by invisible objects, resulting in missing transverse energy. By utilizing this novel method, the Soft Displaced Vertex analysis is meant to demonstrate the technical capabilities of the CMS detector and attempts to bridge a gap between searches that only rely on the presence of missing transverse momentum and jets and searches that complement this event signature by including additional identified SM objects. This analysis also falls in the gap between direct detection searches, where the proper lifetime of the BSM object is on the order of meters or larger, and prompt searches, where they are expected to decay within a few mm or less.

We have provided an insight into the steps that are involved in conducting such a study starting from the MC event generation. Two signal models have been picked and samples with different mass parameters have been generated for the purpose of this analysis. These benchmark signal models are the bino-stop co-annihilation model and the C1N2 model. Two centrally produced MC datasets, the Zjets and Wjets datasets have been included to describe the background of the analysis.
We have shown that the challenging approach of only relying on very soft tracks down to 0.5 GeV to recover displaced vertices and not requiring any SM objects to be identified in the event signature leads to the inclusion of a large number of background tracks and vertices. Using a combination of preselection and a number of different quality, kinematic, and displacement handles, we were able to reduce this incredibly large background to a manageable level whilst also maintaining a good number of signal events. The statistical methods to obtain limits for the production cross-sections have been laid out and are based on the Combine tool, which is integrated in CMSSW. The motivation and justification for using the ABCD method with the chosen variables have been given by confirming the closure for those regions for both the 2018 data and MC simulation. Overall, the results obtained from the 2018 data show that we are able to achieve good sensitivity to both of the tested models over a large parameter space of LLP mass and mass splittings. We were also able to give predictions about the improvements that can be expected from Run III of the LHC. This has been done by calculating the cross-section limits and exclusion charts based on our MC background samples scaled to 300 fb^{-1} and has been justified by comparing MC obtained limits to limits from data at comparable values of integrated luminosity.
These projections show that for Run III great improvements can be expected.

This is accommodated by the fact that various improvements have been made to the CMS detector and its tracking process allowing for higher resolution and better quality of tracks with larger displacements. These factors are expected to increase our sensitivity even further allowing us to possibly exclude smaller mass gaps in the future.

The results obtained by this study are also expected to be further used as a benchmark for the Soft Displaced Vertices group [31, 32] which is working on the implementation of various sophisticated machine learning techniques to allow for better vertex finding and cleaner separation of background and signal vertices.

List of Tables

List of Figures

References

[1] O. Aberle et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010. URL: https://cds.cern.ch/record/2749422 (cit. on p. 1).

[2] Ewa Lopienska. "The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022". In: (2022). General Photo. URL: https://cds.cern.ch/record/2800984 (cit. on p. 2).

[3] Tanja Rommerskirchen. "Supersymmetry with multijet events at CMS : an experimental strategy". In: 2010. URL: https://api.semanticscholar.org/CorpusID:150327734 (cit. on pp. 2, 13).

[4] CERN. *CMS Collaboration*. URL: https://home.cern/science/experiments/cms (cit. on p. 3).

[5] CERN. *ALICE Collaboration*. URL: https://home.cern/science/experiments/alice (cit. on p. 3).

[6] CERN. *ATLAS Collaboration*. URL: https://home.cern/science/experiments/atlas (cit. on p. 3).

[7] Tai Sakuma. "Cutaway diagrams of CMS detector". In: (2019). URL: https://cds.cern.ch/record/2665537 (cit. on p. 4).

[8] W. Adam et al. "The CMS Phase-1 pixel detector upgrade". In: *Journal of Instrumentation* 16.02 (Feb. 2021), P02027. DOI: 10.1088/1748-0221/16/02/P02027. URL: https://dx.doi.org/10.1088/1748-0221/16/02/P02027 (cit. on p. 4).

[9] M. M. Angarano. "The silicon strip tracker for CMS". In: *Nucl. Instrum. Meth. A* 501 (2003). Ed. by R. Horisberger, D. Kotlinski, and W. Erdmann, pp. 93–99. DOI: 10.1016/S0168-9002(02)02016-8 (cit. on p. 4).

[10] A.M. Sirunyan et al. "Particle-flow reconstruction and global event description with the CMS detector". In: *Journal of Instrumentation* 12.10 (Oct. 2017), P10003–P10003. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/10/p10003. URL: http://dx.doi.org/10.1088/1748-0221/12/10/P10003 (cit. on pp. 5, 19).

[11] CERN. *CMS Collaboration*. URL: https://cms.cern/detector/identifying-tracks/silicon-pixels (cit. on p. 5).

[12] *The Phase-2 Upgrade of the CMS Barrel Calorimeters*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2283187 (cit. on p. 6).

[13] *The Phase-2 Upgrade of the CMS Endcap Calorimeter*. Tech. rep. Geneva: CERN, 2017. DOI: 10.17181/CERN.IV8M.1JY2. URL: https://cds.cern.ch/record/2293646 (cit. on p. 6).

[14] *The Phase-2 Upgrade of the CMS Muon Detectors*. Tech. rep. This is the final version, approved by the LHCC. Geneva: CERN, 2017. URL: https://cds.cern.ch/record/2283189 (cit. on p. 7).

[15] A Colaleo et al. *CMS Technical Design Report for the Muon Endcap GEM Upgrade*. Tech. rep. 2015. URL: https://cds.cern.ch/record/2021453 (cit. on p. 7).

[16] CERN. *CMS Collaboration*. URL: https://cms.cern/detector/detecting-muons/muon-drift-tubes (cit. on p. 8).

[17] David Barney. "CMS Detector Slice". CMS Collection. 2016. URL: https://cds.cern.ch/record/2120661 (cit. on p. 8).

[18] *The Phase-2 Upgrade of the CMS Level-1 Trigger*. Tech. rep. Final version. Geneva: CERN, 2020. URL: https://cds.cern.ch/record/2714892 (cit. on p. 9).

[19] CMS Collaboration. *The Phase-2 Upgrade of the CMS Data Acquisition and High Level Trigger*. Tech. rep. This is the final version of the document, approved by the LHCC. Geneva: CERN, 2021. URL: https://cds.cern.ch/record/2759072 (cit. on p. 9).

[20] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013. DOI: https://doi.org/10.1017/CBO9781139525367 (cit. on p. 10).

[21] Wikimedia Commons. *File:Standard Model of Elementary Particles.svg — Wikimedia Commons, the free media repository*. [Online; accessed 18-June-2024]. 2024. URL: https://commons.wikimedia.org/w/index.php?title=File:Standard_Model_of_Elementary_Particles.svg&oldid=884075516 (cit. on p. 11).

[22] Vera C. Rubin. "The Rotation of Spiral Galaxies". In: *Science* 220.4604 (1983), pp. 1339–1344. DOI: 10.1126/science.220.4604.1339. eprint: https://www.science.org/doi/pdf/10.1126/science.220.4604.1339. URL: https://www.science.org/doi/abs/10.1126/science.220.4604.1339 (cit. on p. 11).

[23] Graziano Venanzoni. *New results from the Muon g-2 Experiment*. 2023. arXiv: 2311.08282 (cit. on p. 12).

[24] Wikimedia Commons. *File:Hqmc-vector.svg — Wikimedia Commons, the free media repository*. [Online; accessed 18-June-2024]. 2020. URL: https://commons.wikimedia.org/w/index.php?title=File:Hqmc-vector.svg&oldid=508869423 (cit. on p. 14).

[25] Adel Bilal. *Introduction to Supersymmetry*. 2001. arXiv: hep-th/0101055 (cit. on p. 14).

[26] D CHUNG et al. "The soft supersymmetry-breaking Lagrangian: theory and applications". In: *Physics Reports* 407.1–3 (Feb. 2005), pp. 1–203. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2004.08.032. URL: http://dx.doi.org/10.1016/j.physrep.2004.08.032 (cit. on p. 14).

[27] John Ellis and Keith A. Olive. *Supersymmetric Dark Matter Candidates*. 2010. arXiv: 1001.3651 (cit. on p. 15).

[28] Chih-Ting Lu, Van Que Tran, and Yue-Lin Sming Tsai. *Confronting dark matter co-annihilation of Inert two Higgs Doublet Model with a compressed mass spectrum*. 2020. arXiv: 1912.08875 (cit. on p. 15).

[29] A. Tumasyan et al. "Search for new particles in events with energetic jets and large missing transverse momentum in proton-proton collisions at $\sqrt{s}$ = 13 TeV". In: *Journal of High Energy Physics* 2021.11 (Nov. 2021). ISSN: 1029-8479. DOI: 10.1007/jhep11(2021)153. URL: http://dx.doi.org/10.1007/JHEP11(2021)153 (cit. on pp. 16, 22).

[30] *Search for supersymmetry in events with at least one soft lepton, low jet multiplicity, and missing transverse momentum in proton-proton collisions at $\sqrt{s} = 13$ TeV*. Tech. rep. Geneva: CERN, 2016. URL: http://cds.cern.ch/record/2226549 (cit. on p. 16).

[31] Ivan Mikulec. 2023 (cit. on pp. 16, 17, 58, 60).

[32] Ang Li, Ali Kaan Güven, and Lisa Benato. *HEPHY-Institute for High Energy Physics Vienna*. URL: https://www.oeaw.ac.at/hephy/home (cit. on pp. 16, 58, 60).

[33] Christian Bierlich et al. *A comprehensive guide to the physics and usage of PYTHIA 8.3*. 2022. arXiv: 2203.11601 (cit. on p. 17).

[34] M. Bahr et al. "Herwig++ Physics and Manual". In: *Eur. Phys. J. C* 58 (2008), pp. 639–707. DOI: 10.1140/epjc/s10052-008-0798-9. arXiv: 0803.0883 [hep-ph] (cit. on p. 17).

[35] J. Alwall et al. "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations". In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph] (cit. on p. 18).

[36] CERN. URL: https://feynman.docs.cern.ch/ (cit. on pp. 20, 21).

[37] *CMS luminosity measurement for the 2018 data-taking period at* $\sqrt{s}$ = 13 TeV. Tech. rep. Geneva: CERN, 2019. URL: https://cds.cern.ch/record/2676164 (cit. on p. 24).

[38] *Luminosity measurement in proton-proton collisions at* 13.6 TeV *in 2022 at CMS*. Tech. rep. Geneva: CERN, 2024. URL: https://cds.cern.ch/record/2890833 (cit. on p. 24).

[39] William Buttinger. "Background Estimation with the ABCD Method Featuring the TRooFit Toolkit". In: 2018. URL: https://api.semanticscholar.org/CorpusID:235806829 (cit. on p. 39).

[40] Ofer Vitells. URL: https://indico.cern.ch/event/74940/contributions/2088584/attachments/1047729/1493442/Wald-Asimov.pdf (cit. on p. 40).

[41] Aram Hayrapetyan et al. "The CMS statistical analysis and combination tool: COMBINE". Submitted to *Comput. Softw. Big Sci.* 2024. arXiv: 2404.06614 [physics.data-an] (cit. on pp. 43, 45).

[42] CMS Collaboration. *The CMS statistical analysis and combination tool: COMBINE*. 2024. arXiv: 2404.06614 (cit. on p. 43).

[43] *Performance of Run-3 HLT Track Reconstruction*. Tech. rep. Geneva: CERN, 2022. URL: https://cds.cern.ch/record/2814111 (cit. on p. 58).

[44] *Performance of Run 3 track reconstruction with the mkFit algorithm*. Tech. rep. Geneva: CERN, 2022. URL: http://cds.cern.ch/record/2814000 (cit. on p. 58).

Appendix

2018 data + stop @ 59.8 fb^{-1}						
sample	600 GeV		1000 GeV		1400 GeV	
ΔM = 12 GeV	0.0677	0.1000	0.0412	0.0617	0.0310	0.0442
		0.0524		0.0292		0.0214
ΔM = 15 GeV	0.0231	0.0378	0.0122	0.0204	0.0106	0.0167
		0.0161		0.0082		0.0069
ΔM = 20 GeV	0.0451	0.0650	0.0165	0.0281	0.0144	0.0189
		0.0283		0.0119		0.0097
ΔM = 25 GeV	0.1183	0.1731	0.0663	0.1078	0.0573	0.0900
		0.0841		0.0462		0.0389

Table 11: stop production cross section limits in pb, obtained from using 2018 data and stop samples scaled to 59.8 fb^{-1}. Values correspond to the mean and the upper/lower 68% CL. (seen in Fig. 18)

ZjetsWjets + stop @ 300 fb^{-1}						
sample	600 GeV		1000 GeV		1400 GeV	
ΔM = 12 GeV	0.0279	0.0383	0.0163	0.0237	0.0118	0.0179
		0.0230		0.0115		0.0085
ΔM = 15 GeV	0.0089	0.0131	0.0048	0.0073	0.0044	0.0062
		0.0064		0.0034		0.0030
ΔM = 20 GeV	0.0123	0.0181	0.0064	0.0091	0.0053	0.0077
		0.0118		0.0042		0.0031
ΔM = 25 GeV	0.0394	0.0576	0.0230	0.0339	0.0192	0.0272
		0.0282		0.0156		0.0119

Table 12: stop production cross section limits in pb, obtained from using MC background and stop samples scaled to 300 fb^{-1}. Values correspond to the mean and the upper/lower 68% CL. (seen in Fig. 19)

2018 data + c1n2 @ 59.8 fb^{-1}						
sample	600 GeV		1000 GeV		1400 GeV	
ΔM = 12 GeV	0.5071	0.7096	0.2266	0.3539	0.1981	0.2946
		0.3627		0.1678		0.1391
ΔM = 15 GeV	0.1755	0.2678	0.1000	0.1502	0.0699	0.1037
		0.1283		0.0693		0.0472
ΔM = 20 GeV	0.2425	0.3720	0.1389	0.2130	0.0908	0.1420
		0.1815		0.1007		0.0624
ΔM = 25 GeV	0.7606	1.1747	0.5365	0.7937	0.4470	0.6647
		0.5411		0.3251		0.2634

Table 13: c1n2 production cross section limits in pb, obtained from using 2018 data and c1n2 samples scaled to 59.8 fb^{-1}. Values correspond to the mean and the upper/lower 68% CL. (seen in Fig. 20)

ZjetsWjets + c1n2 @ 300 fb^{-1}						
sample	600 GeV		1000 GeV		1400 GeV	
ΔM = 12 GeV	0.1778	0.2567	0.0846	0.1289	0.0774	0.1109
		0.1264		0.0609		0.0523
ΔM = 15 GeV	0.0640	0.0897	0.0389	0.0549	0.0273	0.0400
		0.0458		0.0271		0.0182
ΔM = 20 GeV	0.0939	0.134	0.0529	0.0751	0.0441	0.0576
		0.0532		0.0375		0.0309
ΔM = 25 GeV	0.2532	0.3713	0.1807	0.2634	0.1480	0.2109
		0.1771		0.1311		0.1031

Table 14: c1n2 production cross section limits in pb, obtained from using MC background and c1n2 samples scaled to 300 fb^{-1}. Values correspond to the mean and the upper/lower 68% CL. (seen in Fig. 21)