Department of Physics and Astronomy
Heidelberg University

Bachelor Thesis in Physics
submitted by

Tobias Rudolf

born in Bietigheim-Bissingen (Germany)

2025

Material Budget Imaging of Thin Samples in view of the ALICE ITS3 Upgrade

This Bachelor Thesis has been carried out by Tobias Rudolf at the
Physikalisches Institut, Heidelberg University (Germany)
under the supervision of
Prof. Dr. Silvia Masciocchi

Abstract:

“Material-budget imaging“ (MBI) based on multiple Coulomb scattering is an established method for quantifying material budget. Prior work has steadily improved its sensitivity to small material differences, a key requirement for ultra-thin trackers such as ALICE ITS3. Building on this foundation, this thesis addressed the challenge of measuring thin samples with high precision and well-controlled systematics. Using aluminium, nickel, and silicon targets in the material budget regime of 0.1 ‰ to 1 ‰, i.e. below the canonical validity range of the Highland relation. It is demonstrated that the expected scaling $\theta_0 \propto \sqrt{x/X_0}$ continued to hold. A data-driven streamline of treating systematics, combining reference-based deconvolution with transfer factors and momentum corrections, stabilised absolute comparisons between measurement and expectation. At the thinnest budgets, an additive resolution floor was identified where absolute accuracy begins to degrade even though relative sensitivity persists. Together, these results extended MBI into the ultra-thin domain and provided a practical pathway for characterising the minute material contributions relevant to ITS3 and future low-mass detector technologies.

Zusammenfassung:

„Material-Budget Imaging“ (MBI) auf Basis der mehrfachen Coulomb-Streuung ist eine etablierte Methode zur Quantifizierung von Material Budget. Frühere Arbeiten haben die Empfindlichkeit gegenüber kleinen Materialunterschieden stetig verbessert, eine zentrale Voraussetzung für ultradünne Tracker wie ALICE ITS3. Aufbauend auf dieser Grundlage befasste sich diese Arbeit mit der Herausforderung, dünne Proben mit hoher Präzision und gut kontrollierten Systematiken zu vermessen. Hierzu wurden Aluminium-, Nickel- und Silizium-Proben im Material Budget Regime von 0.1 ‰ bis 1 ‰ eingesetzt, d.,h. unterhalb des kanonischen Gültigkeitsbereichs der Highland-Relation. Es wurde gezeigt, dass die erwartete Skalierung $\theta_0 \propto \sqrt{x/X_0}$ weiterhin galt. Ein datengetriebener Arbeitsablauf zur Behandlung der Systematiken, der eine referenzbasierte Dekonvolution mit Transferfaktoren und Impulskorrekturen kombinierte, stabilisierte die absoluten Vergleiche zwischen Messung und Erwartung. Bei den dünnsten Budgets wurde eine additive Auflösungsuntergrenze identifiziert, ab der die absolute Genauigkeit nachließ, obwohl die relative Sensitivität erhalten blieb. Zusammengenommen erweiterten diese Ergebnisse MBI in den ultradünnen Bereich und lieferten einen praktikablen Weg zur Charakterisierung der minimalen Materialbeiträge, die für ITS3 und zukünftige Detektortechnologien mit geringem Materialbudget relevant sind.

Contents

1 ALICE — A Large Ion Collider Experiment 1
- 1.1 Quark-Gluon Plasma . . . 3
- 1.2 Coordinates, Layout, and Acceptance . . . 4
- 1.3 The Inner Tracking System 2 . . . 5
- 1.4 Other Major Detector Systems . . . 8
- 1.5 The Inner Tracking System 3 . . . 8

2 Particle Detection 12
- 2.1 Ionisation and excitation (Bethe–Bloch) . . . 12
- 2.2 Bremsstrahlung and the radiation length . . . 14
- 2.3 Multiple Coulomb scattering: Molière core and Highland form . . . 15
- 2.4 From angle histograms to material budget . . . 17

3 Acquisition & Reconstruction 18
- 3.1 ALPIDE Sensor . . . 18
- 3.2 Telescope Setup . . . 18
- 3.3 Measurement Samples . . . 19
- 3.4 Test Beam Campaign . . . 22
- 3.5 Alignment & Tracking . . . 23
- 3.6 Pixelisation . . . 25

4 Material Budget 28
- 4.1 Scattering Angle Distributions . . . 28
- 4.2 Momentum Gradient Correction . . . 31
- 4.3 Fluctuations in Uniform Material Budget . . . 34
- 4.4 Material Budget Determination at Small x/X0 . . . 43

5 Summary, conclusion and outlook 48

A Additional Information 51
- A.1 Gaussian Fit Values . . . 51
- A.2 Central-interval robustness . . . 53

Nomenclature 54

Bibliography 54

Chapter 1

ALICE — A Large Ion Collider Experiment

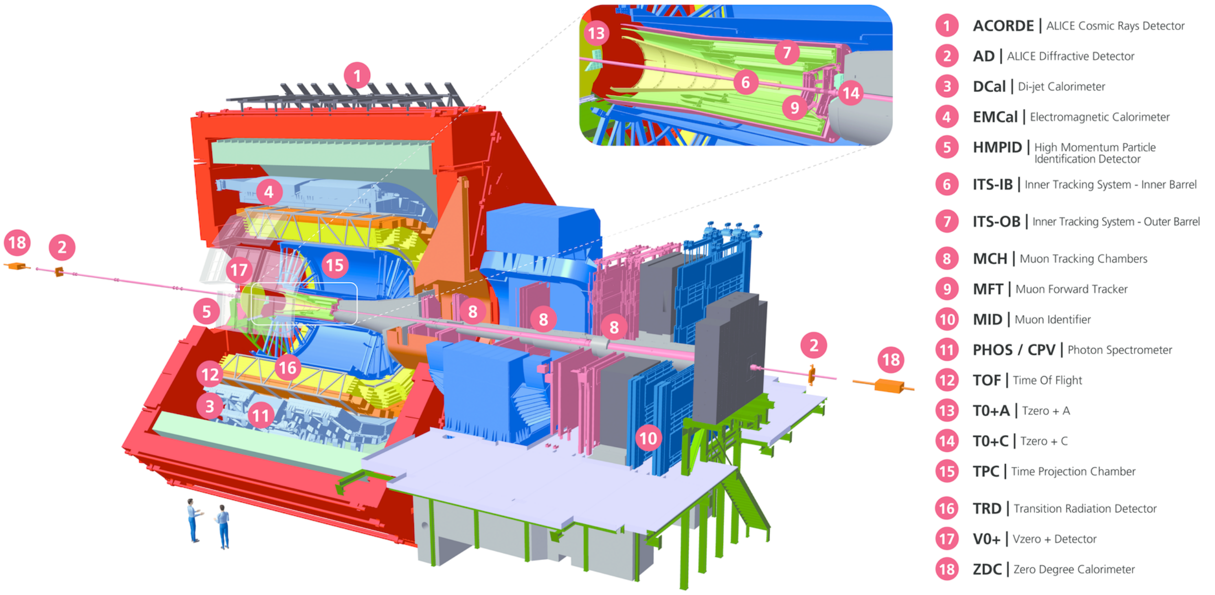

Figure 1.1 – Overview of the ALICE detector: the central barrel inside the L3 solenoid is shown with principal subsystems indicated. Taken from [1].

ALICE is conceived and built to answer a simple but profound question: *how do quarks and gluons behave when matter is so hot and dense that protons and neutrons have not yet formed?* [2, 3] At the Large Hadron Collider (LHC), heavy nuclei collide at ultrarelativistic energies, creating a droplet of deconfined quark–gluon matter. The role of ALICE is to record those collisions with the precision needed to reconstruct the ensuing story told by hadrons, photons, electrons, and muons. Doing so demands excellent tracking, particle identification (PID) in the busiest collision environment at the LHC and very low material in the innermost tracking layers so that multiple scattering does not blur the picture at low transverse momentum [4].

CERN [5] is the European laboratory for particle physics on the French–Swiss border near Geneva. It operates a network of accelerators and experiments used by a global collaboration of institutes and universities. The flagship machine is the LHC, a 26.7 km ring in which counter-rotating proton or hadron beams are brought into collision at several dedicated interaction points (IPs). Four large experiments are located on the ring, ALICE occupies IP2 and is optimised for heavy-ion physics. In everyday language, “heavy-ion” means nucleus–nucleus collisions (e.g. Pb–Pb), in contrast to proton–proton (pp) collisions. ALICE records pp and heavy-ion collisions. pp data provide the baseline for interpreting Pb–Pb, while the detector design prioritises performance in the high-multiplicity ion environment through a very low material budget at small radii and precise inner tracking. The LHC sits at the end of an injector chain that prepares and accelerates beams step by step before injection into the main ring. The schematic of the accelerator complex (Fig. 1.2) provides a visual map from source to collider and marks the locations of the major experiments, including IP2 for ALICE. In routine operation the LHC alternates between pp periods and dedicated heavy-ion periods with detector readout and reconstruction being configured accordingly.

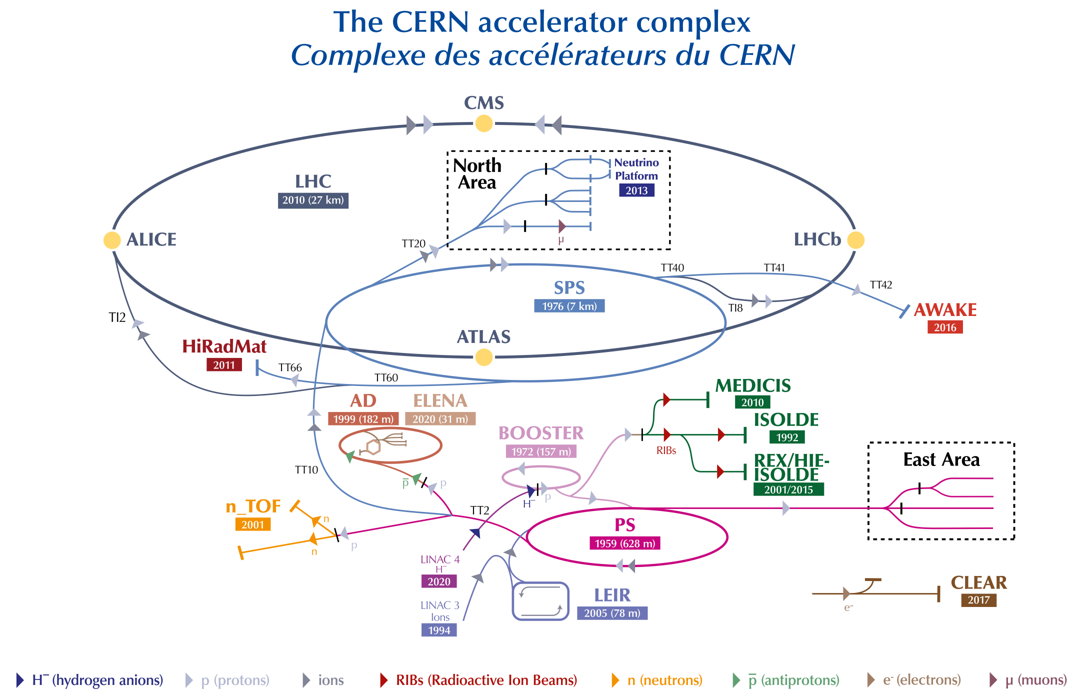

[redacted]

Figure 1.2 – CERN accelerator complex and injector chain leading to the LHC. The sequence of machines prepares beams for collision and indicates the positions of the LHC experiments. ALICE is located at IP2. Taken from [6].

Functionally, ALICE acts as a coherent mid-rapidity spectrometer inside a 0.5 T solenoidal magnetic field. A thin, all-pixel Inner Tracking System 2 (ITS2, see Fig. 1.1) sits close to the beam pipe to pin down the primary vertex and provide the earliest space points. The large Time Projection Chamber (TPC) supplies the long lever arm and dE/dx information, fast timing and dedicated electron identification complete particle identification (PID). A forward muon spectrometer extends the reach to quarkonia and heavy flavour. These subsystems operate in tandem with continuous data taking in the current data taking period (Run 3).

For the aims of this thesis, material at the smallest radii is the decisive constraint. Every additional fraction of a radiation length in front of the first precise space points increases multiple Coulomb scattering and, at low p_{T}, degrades the impact-parameter and momentum resolutions. The ITS2 design addresses exactly this by placing monolithic active pixel sensors as close as practicable to the interaction point and by minimising supports, cooling and services. In practice, these choices keep soft tracks and displaced vertices measurable even in the most crowded heavy-ion events.

1.1 Quark-Gluon Plasma

In a head-on Pb–Pb collision at the LHC, a tiny, short-lived droplet of deconfined quark–gluon matter forms and expands before it hadronises [2, 3]. Figure 1.3 summarises this evolution at a glance. The widths and positions of the stages are illustrative.

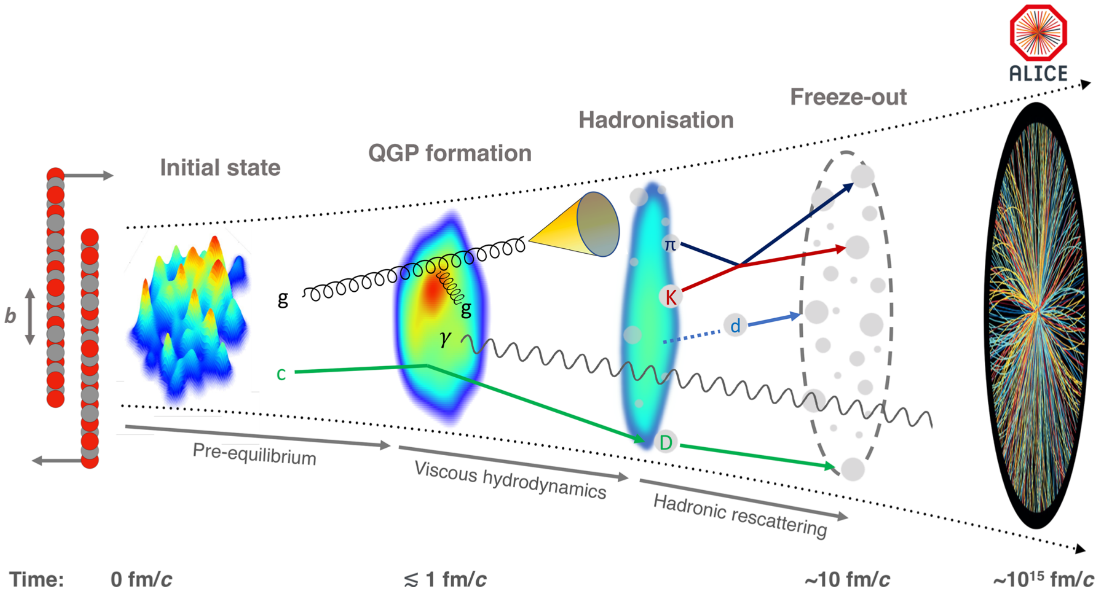

Figure 1.3 – Schematic space–time journey of a central heavy-ion collision: initial state and pre-equilibrium → Quark-Gluon plasma (QGP) formation and viscous hydrodynamic expansion → hadronisation and hadronic rescattering → kinetic freeze-out. Example probes (gluons/jets, direct photons γ, heavy flavour $c \to D$, identified hadrons π, K) indicate where key observables originate. Taken from [7].

At the moment of overlap (time ~ 0 fm/c), the two Lorentz-contracted nuclei pass through each other and deposit energy in an almond-shaped region whose size and anisotropy are set by the impact parameter b (left of Fig. 1.3). A short *pre-equilibrium* stage follows, and within $\lesssim 1$ fm/c the system approaches local equilibrium so that a hydrodynamic description becomes applicable. The droplet then undergoes *viscous hydrodynamic* expansion, converting the initial spatial anisotropy into momentum-space anisotropy observed as collective flow coefficients v_n [2].

As the temperature falls toward the quantum chromodynamics (QCD) crossover, partons convert into hadrons (*hadronisation*). The tracks in Fig. 1.3 highlight typical probes: high-p_{T} partons that lose energy while traversing the medium (jet quenching), direct photons γ that escape largely unmodified, and heavy quarks (c) whose diffusion and hadronisation into D mesons are sensitive to transport properties. After hadronisation, the system enters a *hadronic rescattering* phase in which elastic interactions continue until they cease at *kinetic freeze-out*. Particle yields are effectively fixed earlier at *chemical freeze-out* [8, 3]. Characteristic timescales are order-of-magnitude: QGP formation within $\mathcal{O}(1\,\mathrm{fm/c})$ and kinetic freeze-out around $\mathcal{O}(10\,\mathrm{fm/c})$, with details depending on centrality and beam energy [2].

Taken together, these signatures motivate the ALICE detector requirements. Collective flow and identified-hadron spectra live at low to intermediate p_{T}, heavy flavour and quarkonia need excellent vertexing and jet quenching benefits from large acceptance and robust tracking. All three depend critically on precise, low-material inner tracking and reliable PID at mid-rapidity, hence the central role of the ITS–TPC–TOF ensemble introduced in this chapter [4].

1.2 Coordinates, Layout, and Acceptance

The QGP observables just introduced translate into concrete geometric and kinematic requirements. This section fixes the coordinate conventions and gives a compact map of the central barrel acceptance and radial layout that matter for tracking, PID, and, most directly for this work, the inner material budget.

ALICE uses a right-handed system with the z-axis along the LHC beam, x pointing toward the LHC ring centre, and y upward. Azimuth is denoted by ϕ, and the polar angle θ is measured from the $+z$ axis. Pseudorapidity is defined as $\eta = -\ln[\tan(\theta/2)]$, thus small polar angles (forward) correspond to large $|\eta|$, while mid-rapidity sits near $|\eta| \approx 0$ [8].

The left panel of Fig. 1.4 summarises the η coverage of principal subsystems. Around mid-rapidity, tracking and PID are provided by the ITS, TPC, and TOF with near-2π azimuthal acceptance. The TRD overlaps this region to enhance e/π separation. Electromagnetic calorimetry (EMCal/DCal) covers a central band for jets and photons, while the photon spectrometer (PHOS) provides high-resolution photons in a narrower window. Forward devices extend timing and multiplicity measurements outside the barrel [4].

The right panel of Fig. 1.4 shows a transverse (x–y) cut through the barrel. From the beam pipe outward one sees: the concentric ITS layers very close to the interaction point,

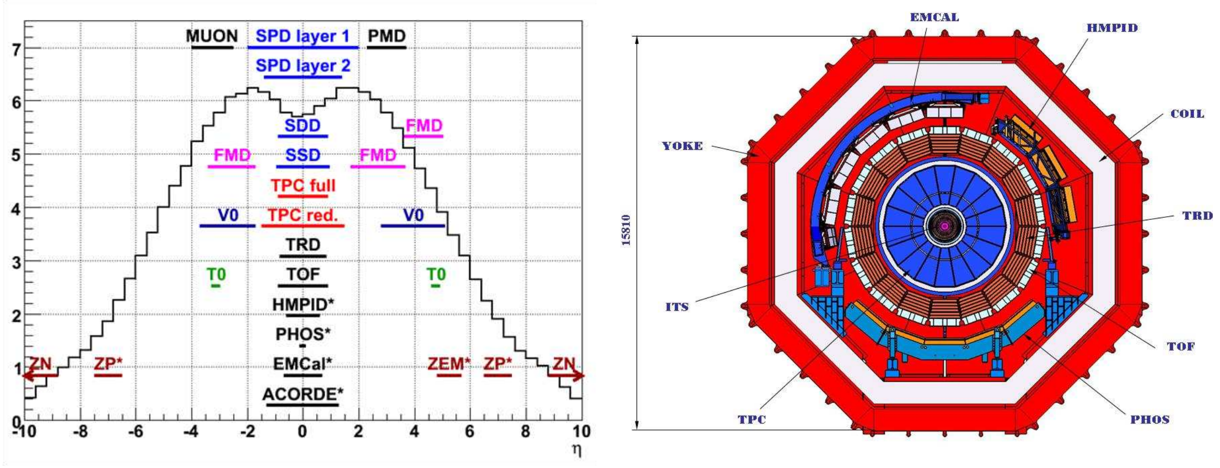

(a) Pseudorapidity coverage of principal ALICE sub-detectors. Taken from [9]

(b) Transverse (x–y) cross-section of the ALICE central barrel. Taken from [10]

Figure 1.4 – ALICE η-acceptance and transverse barrel layout.

the large TPC drift volume forming the blue annulus, time-of-flight counters on the outer radius and, in selected sectors, the TRD and electromagnetic calorimeters. The coil and iron yoke of the L3 magnet surround the acceptance and provide the $\sim 0.5\,\mathrm{T}$ solenoidal field used for momentum measurement.

Tracking in the solenoidal field. ALICE operates in a $0.5\,\mathrm{T}$ solenoidal magnetic field that bends charged particles. The momentum and charge sign follow from the curvature of the helical track in the magnetic field [8]. The attainable precision is governed by three ingredients that recur throughout this thesis: the geometric lever arm of the measurement, the precision of individual space points and the amount of material traversed. A large radial span with widely separated measurements increases bending information, here the TPC, with many three-dimensional points over a long drift, dominates [11]. The innermost ITS2 layers provide the earliest and most precise points and constrain the primary vertex [12, 13, 14]. Finally, multiple Coulomb scattering grows with the accumulated material budget x/X_0 and is most detrimental at low p_T [8]. In the continuous readout of Run 3, pattern recognition seeds in one system and are prolonged into the other, after which a global fit combines all hits [15, 16].

1.3 The Inner Tracking System 2

Within the central barrel, the ITS2 provides the earliest and most precise space points and thereby anchors primary-vertex and impact-parameter measurements that many heavy-ion observables rely on. It sits exactly where performance is won or lost: at the smallest radii, where low material suppresses multiple scattering and fine granularity fixes the first precise points of a trajectory. High track densities at mid-rapidity in heavy-ion events make this role even more stringent. Efficiency for soft tracks, clean separation of primaries and secondaries, and stable pattern recognition all depend on a thin, precise inner tracker that can operate continuously in Run 3 conditions [12, 14]. Sensor and readout specifics are described in Section 3.1.

Seven concentric pixel layers wrap the beam pipe as shown in Fig. 1.5. The three innermost layers form the Inner Barrel and are placed as close as practicable to the interaction region to capture the first measurements with minimal extrapolation. The remaining four layers extend the radial lever arm and stabilise the pattern recognition as tracks transition from the silicon into the TPC. The staves are arranged with overlaps (see Fig. 1.7) so that services and mechanical boundaries do not open gaps relevant for vertexing. The reduced beam-pipe radius, visible in the figure, was an enabling step, allowing the innermost layer to move inward while the overall support, cooling and services were redesigned to remove unnecessary mass [12].

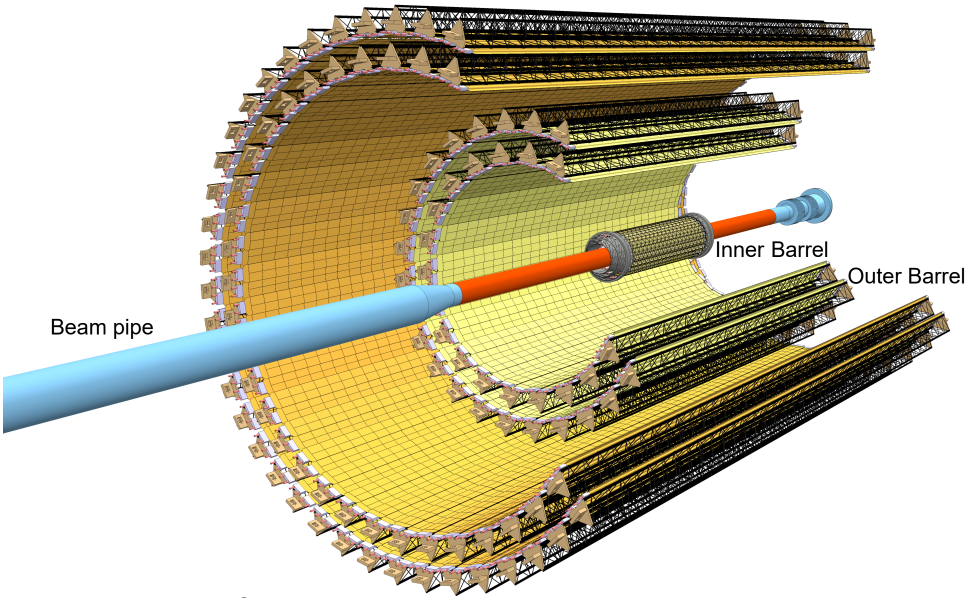

Figure 1.5 – ITS2 layout and segmentation: seven layers of monolithic pixel sensors wrap the beam pipe. The three innermost layers form the Inner Barrel, the outer four extend the lever arm on lightweight staves. The reduced beam-pipe radius allows the first precise space points to be placed very close to the interaction point. Taken from [17].

Placing thin, precise pixels at small radii translates directly into better impact-parameter resolution, most noticeably at low p_{T} where multiple scattering otherwise dominates. The data comparison in Fig. 1.6a shows this effect in the $r\phi$ projection: the Run 3 curve (with ITS2) lies well below the Run 2 curve across the full momentum range, with the largest relative gain below a few GeV/c. This improvement strengthens charm and beauty measurements, improves the association of soft tracks to the primary vertex in high multiplicity, and generally lowers the systematic floor for analyses that depend on small displacements [14, 12].

The material story behind this performance is captured by the stave-level "bill of materials" in Fig. 1.7. Expressed as a fraction x/X_0, the total along a track results from the sum of

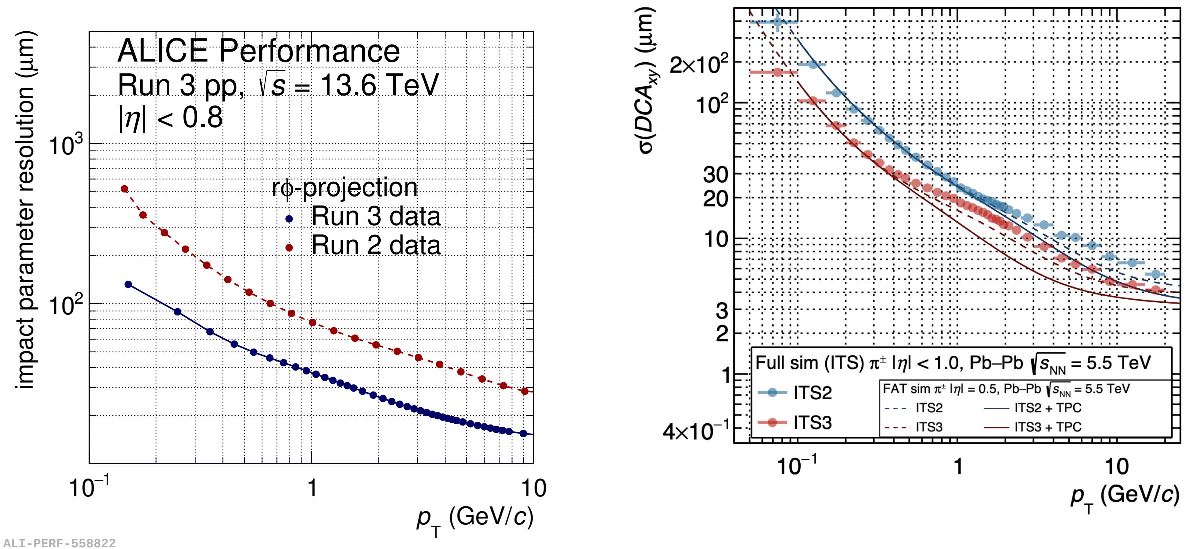

(a) Run 2 vs Run 3 $r\phi$ impact-parameter resolution as a function of p_{T}. Taken from [1].

(b) ITS2 vs ITS3 $r\phi$ impact-parameter resolution as a function of p_{T}. Only tracks with a hit on each ITS layer were considered. Taken from [18].

Figure 1.6 – Impact-parameter (d_0) resolution in the $r\phi$ projection vs. p_{T}.

several contributions: the silicon sensor itself, thin adhesive layers, carbon-fibre and foam structures, low-mass cooling, and the power/data services. Because these contributions add, shaving even a few tenths of a percent from any element at the smallest radii lowers the multiple-scattering floor for every low-p_{T} trajectory, as illustrated in Fig. 1.6b [12, 13].

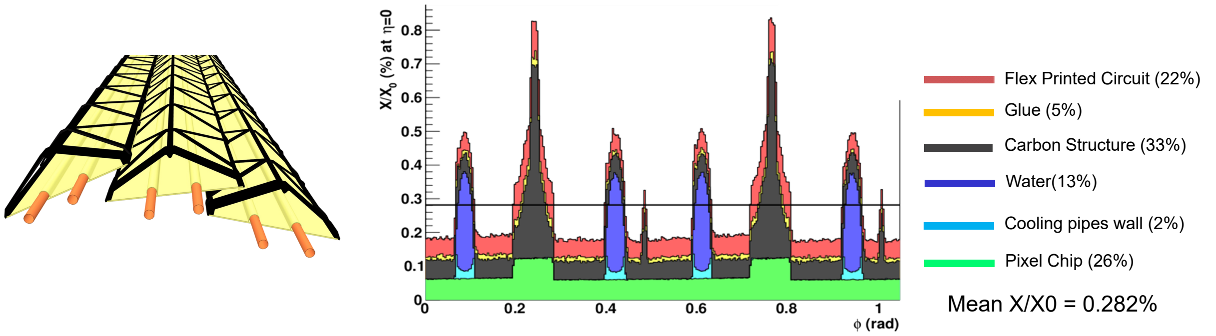

Figure 1.7 – Stave-level material budget for ITS2, expressed as a fraction x/X_0. The principal components (sensor, adhesives, carbon structures, cooling, services) are shown to indicate where reductions most effectively lower multiple scattering at low p_{T}. Taken from [12]

Operationally, the ITS2 works as a pair with the TPC in continuous Run 3 running. Pattern-recognition seeds formed in one detector are prolonged into the other, and a global fit combines all space points to yield momenta and impact parameters with the correct covariances. Alignment and calibration are maintained within the analysis framework so that the per-hit precision delivered by the silicon actually appears as stable track-parameter resolutions over long time frames [15, 16]. In the context of this thesis, this means two things: the own low material of the ITS2 sharpens the sensitivity of MBI to added or misplaced components, and its precise first space points improve the reconstruction of the projected angles that our estimators use.

1.4 Other Detector Systems and a Look Ahead

Beyond the inner silicon, the remaining central-barrel and the forward instrumentation complete the measurement chain required by the QGP programme. This section summarises the contribution of the major system to tracking and PID at mid-rapidity and its integration with continuous Run 3 readout, emphasising the aspects relevant for the analyses that follow.

TPC — main tracker and dE/dx PID The Time Projection Chamber surrounds the silicon with a large gas volume instrumented at both endplates. Charged particles ionise the gas and electrons drift in a uniform electric field to segmented pads where position and drift time form three-dimensional space points. In Run 3, multi-GEM stacks provide stable amplification with low ion backflow, enabling continuous, dead-time–free operation [15]. The TPC supplies the dominant lever arm for momentum measurement and precise dE/dx for PID, while its many points per track stabilise the global fit that combines ITS2 and TPC information [11].

TOF — timing that closes the mid-p_T PID gap The Time-of-Flight barrel measures arrival times with $\mathcal{O}(1 \times 10^2\,\mathrm{ps})$ resolution [19]. Together with the momentum from the fit, this yields mass-squared and therefore species separation in the intermediate momentum window where TPC dE/dx alone is less discriminating. TOF also provides clean timing anchors in continuous readout.

TRD — targeted e/π separation The Transition Radiation Detector wraps selected sectors outside the TPC. Radiators generate X-ray transition radiation for highly relativistic particles. Multi-wire proportional chambers detect it alongside ionisation signals. Likelihoods built from these responses enhance electron identification where it is most challenging and can be folded into reconstruction [20].

1.5 The Inner Tracking System 3

The physics programme outlined above puts a premium on minimal inner material: multiple scattering at the first precise space points sets the low-p_T floor for impact-parameter and momentum resolutions. ITS3 targets a further reduction of x/X_0 at the smallest radii beyond ITS2 by replacing stave-based mechanics with wafer-scale, bent monolithic sensors on an ultra-light cylindrical support. The qualitative implication is direct: a thinner

inner barrel lowers multiple-scattering, tightening vertexing and improving the efficiency and purity of soft-track reconstruction in high-multiplicity events.

Concept and mechanics. ITS3 is the first collider detector to employ *wafer-scale, bent* monolithic active pixel sensors (MAPS), with individual devices up to $\sim 26\,\mathrm{cm} \times 10\,\mathrm{cm}$. Sensors are cylindrically bent and held by minimal carbon-foam structures: two longitudinal *longerons* per sensor ensure uniform bending without warping, and *half-rings* at the endcaps define the radii. Services (power, cooling, data) are routed from the accessible side and air provides low-mass cooling. The innermost half-layers are positioned at radii of approximately 19 mm, 25.2 mm, and 31.5 mm, enabled by a smaller beam pipe. With carbon foam only at the sensor edges (longerons and half-rings) and no traditional stave stack-up in front of the active area, the *average* material budget of a full half-layer is at the $\sim 0.09\%\,X_0$ level, with the central sensor region setting the baseline and edge structures providing the only localized [18]. These choices are expressly aimed at suppressing multiple scattering for soft tracks and tightening the primary/secondary vertex resolutions that heavy-flavour and quarkonia analyses rely on.

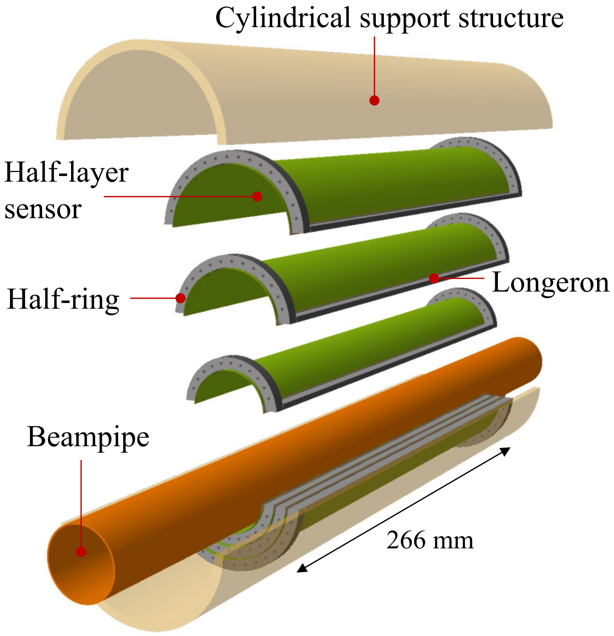

Figure 1.8 – ITS3 inner-barrel concept: wafer-scale, bent MAPS mounted on a minimal cylindrical support. Edge carbon-foam longerons and half-rings define shape and radii while keeping material out of the active acceptance. Taken from [18].

Material-budget regions and edge structures. Three regions govern x/X_0 in a half-layer: (i) the *sensor-only* central band, (ii) the *longerons* at top/bottom in ϕ, and the endcap (iii) *half-rings* on both endings. In addition to the silicon, carbon foams are bonded via controlled adhesive layers. A carbon fleece acts as a barrier to limit glue wicking into the foam, and a thin polyimide film is used at the longeron interface. The baseline x/X_0 is thus set by the bent sensor, with localized increases where glue, foam, and films appear at the edges.

Simulated material-budget distributions. Figure 1.9 collects the canonical simulations used to characterize x/X_0 versus (η, ϕ) for tracks originating at $z_{\text{vtx}} = 0$. The central sensor band is azimuthally uniform and longerons produce narrow enhancements at $\phi \approx 0, \pi$. The average within $|\,\eta\,| < 1$ and $0 < \phi < \pi$ yields $\langle x/X_0 \rangle \approx 0.086\%$, while extending to $|\,\eta\,| < 2$ gives $\langle x/X_0 \rangle \approx 0.149\%$. A profile at $\eta = 0$, $\phi = \pi/2$ averaged over $|z_{\text{vtx}}| \leq 14\,\text{cm}$ gives $\langle x/X_0 \rangle \approx 0.082\%$. In all cases, the sensor dominates at higher $|\,\eta\,|$ where path length increases with incidence angle. These maps explicitly locate where edge structures and adhesives contribute, informing both reconstruction and metrology requirements.[1]

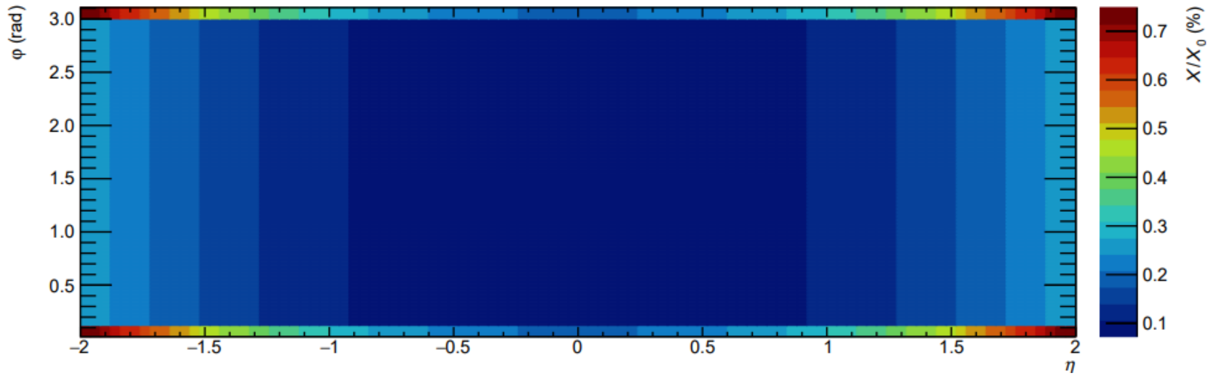

(a) x/X_0 for $z_{\text{vtx}} = 0$ as a function of ϕ (rad) and η.

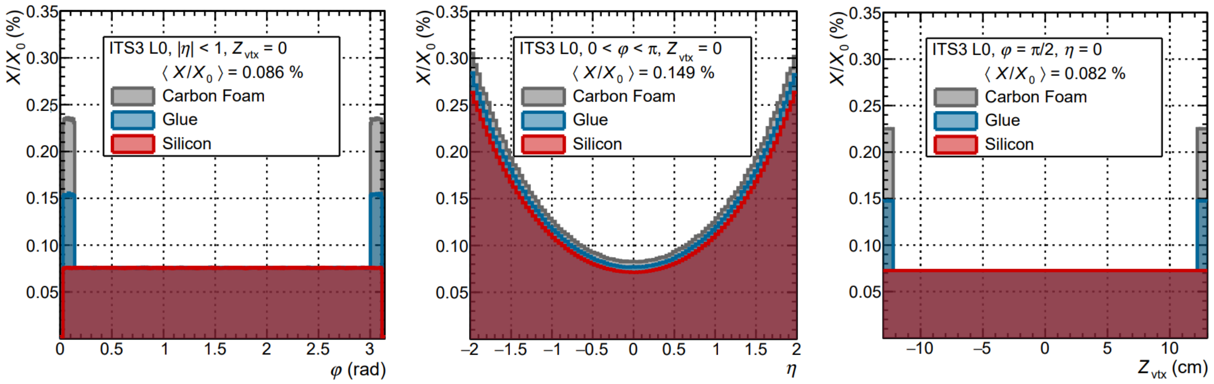

(b) $\langle x/X_0 \rangle$ vs. ϕ for $|\,\eta\,| < 1$ and $0 < \phi < \pi$.

(c) $\langle x/X_0 \rangle$ vs. η for $|\,\eta\,| < 2$ and $0 < \phi < \pi$.

(d) $\langle x/X_0 \rangle$ vs. z_{vtx} for $\eta = 0$ and $\phi = \pi/2$.

Figure 1.9 – Simulated average material budget for the innermost ITS3 half-layer. Taken from [18].

Implications for this work. For material–budget imaging, ITS3 provides both the *motivation* and the *targets*: The central sensor band defines the absolute baseline to be

---

[1] As discussed in the TDR, early simulations may omit some thin layers (e.g. carbon fleece or polyimide at the longeron). This motivates an experimental validation of the effective x/X_0 of each stack-up element.

validated at the $\mathcal{O}(10^{-3})$ level, edge regions (longerons, half-rings) and adhesive/fleece/-film interfaces introduce localized, mechanically necessary inhomogeneities that must be quantified and any omission or mis-modelling of ultra-thin layers in simulation translates into biases in the effective x/X_0 that MBI can detect. The methodology developed in this thesis is therefore tuned to isolate small differences in x/X_0 over $\mathcal{O}(\mathrm{mm})$ scales of material thickness and to provide a data-driven cross-check of the ITS3 material model used in reconstruction and performance projections.

Chapter 2

Particle Detection

Charged particles traversing matter lose energy and undergo angular deflections due to repeated electromagnetic interactions with atomic electrons and nuclei. This chapter collects the ingredients needed later for material–budget imaging (MBI): from the mean energy loss and its fluctuations to the definition of the radiation length and the multiple–scattering laws used to connect measured angular widths to the underlying material budget x/X_0 [21, 22].

The focus is on heavy charged particles (e.g. pions, muons, protons) for which collisional energy loss dominates up to moderate $\beta\gamma$. Electrons follow the same principles but develop substantial radiative losses already at GeV energies. Three quantitative pillars are required: the mean energy loss governed by the Bethe–Bloch formula, the statistical fluctuations of the energy deposit (Landau/Vavilov straggling), and multiple Coulomb scattering (Molière theory and its practical Highland form).

2.1 Ionisation and excitation (Bethe–Bloch)

The mean collisional stopping power of a heavy charged particle of charge $z \cdot e$ and velocity $v = \beta \cdot c$ in a material with atomic number Z and atomic mass A is

$$-\left\langle \frac{dE}{dx} \right\rangle = K\, z^2 \, \frac{Z}{A} \, \frac{1}{\beta^2} \left[\frac{1}{2} \ln\!\left(\frac{2\, m_e c^2\, \beta^2 \gamma^2\, T_{\max}}{I^2} \right) - \beta^2 - \frac{\delta(\beta\gamma)}{2} - \frac{C}{Z} \right], \tag{2.1}$$

with the symbols and constants defined in Table 2.1. The form follows the modern convention while the historical derivations are due to Bethe and Bloch [23, 24, 25]. The maximum transferable kinetic energy in a single collision with an atomic electron is

$$T_{\max} = \frac{2\, m_e c^2\, \beta^2 \gamma^2}{1 + 2\gamma\, \frac{m_e}{M} + \left(\frac{m_e}{M}\right)^2}, \tag{2.2}$$

in the standard notation used across literature [21]. For $M \gg m_e$ this reduces to $T_{\max} \simeq 2\, m_e c^2\, \beta^2 \gamma^2$. At intermediate $\beta\gamma$ the stopping power exhibits a shallow minimum (minimum–ionising region) around $\beta\gamma \approx$ 3–4. Toward low β shell corrections increase the

stopping power, while at high $\beta\gamma$ the density effect reduces the rise by polarisation of the medium. The characteristic material dependence and the MIP region are illustrated in Fig. 2.1. Gases (e.g. He) sit higher because the density effect sets in later, while heavy solids cluster tightly.
For $M \gg m_e$ this reduces to $T_{\max} \simeq 2\,m_e c^2\,\beta^2\gamma^2$. At intermediate $\beta\gamma$ the stopping power exhibits a shallow minimum (minimum–ionising region) around $\beta\gamma \approx 3$–4. Toward low β shell corrections increase the stopping power, while at high $\beta\gamma$ the density effect reduces the rise by polarisation of the medium [26, 27].

Table 2.1 – Variables to calculate the mean energy loss with the Bethe–Bloch formula.

K	$4\pi N_A r_e^2 m_e c^2$	z	Charge of projectile
N_A	Avogadro constant	$\delta(\beta\gamma)$	Density–effect correction
r_e	Electron radius	Z	Atomic number of absorber
m_e	Electron mass	I	Mean excitation energy
A	Atomic mass of absorber	C/Z	Shell correction

Detector mapping. In tracking detectors, the energy lost by ionisation appears as a stochastic energy deposit in the active volume. In thin layers this deposit exhibits pronounced fluctuations and is typically summarised by robust estimators (e.g. truncated mean) rather than the arithmetic mean [28, 29].

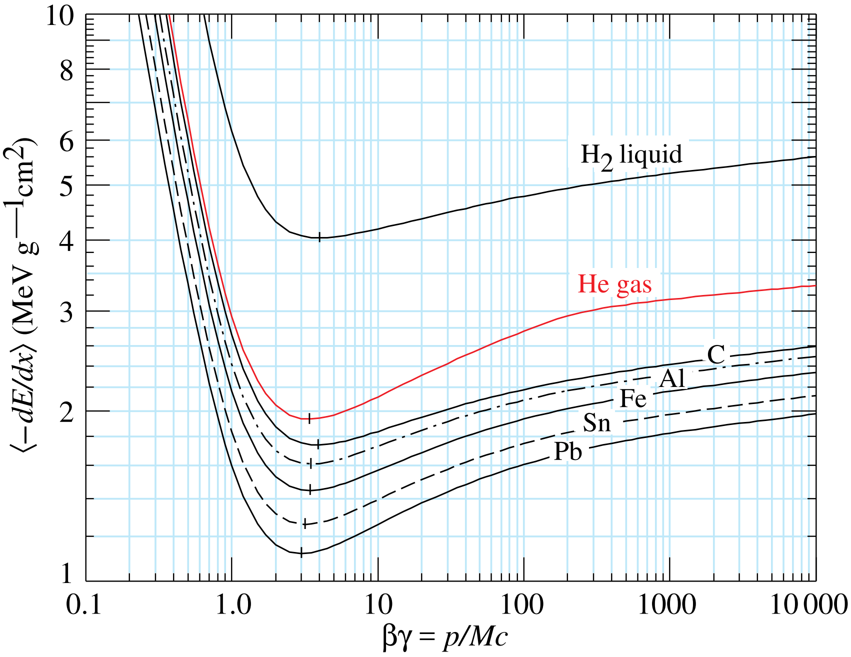

Figure 2.1 – Mean mass stopping power $\langle -(1/\rho)\,dE/dx\rangle$ as a function of $\beta\gamma$ for representative materials, the MIP minimum is annotated. Taken from [30].

Straggling and δ electrons Single–track energy deposits fluctuate strongly in thin absorbers due to rare hard collisions that liberate energetic δ electrons. The resulting distribution is asymmetric with a long high–energy tail and is well described by Landau–Vavilov theory [31]. Figure 2.2 shows simulated spectra for 500 MeV pions in Si at several thicknesses: the most–probable loss scales approximately with thickness, while the width and the non-Gaussian tail are increasingly prominent for thinner layers. For sufficiently thick absorbers many collisions average out and the distribution tends toward a Gaussian [32].

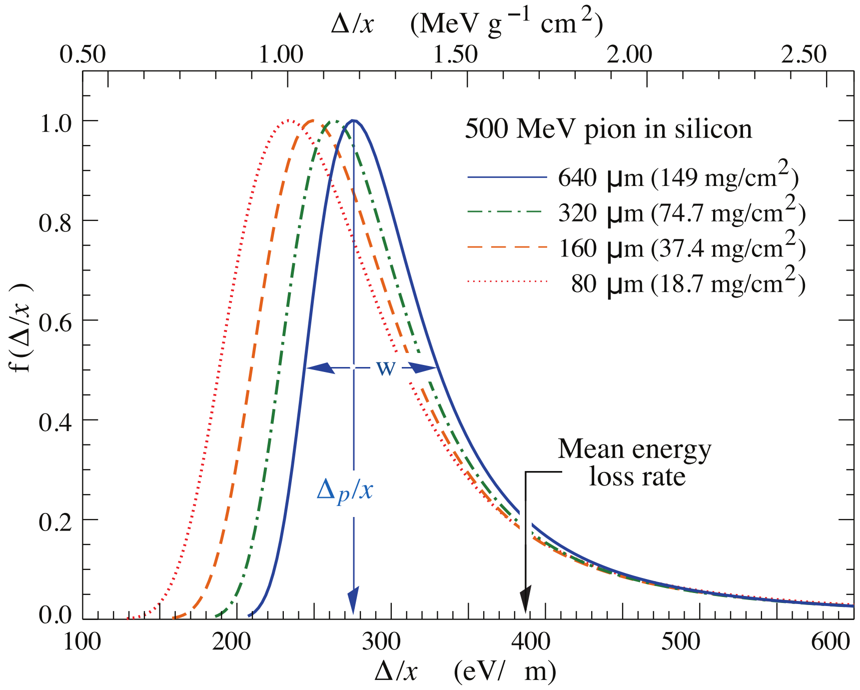

Figure 2.2 – Simulated energy–deposit distribution in 285 µm Si for high–energy pions. Taken from [30].

2.2 Bremsstrahlung and the radiation length

For electrons and photons, radiative processes are quantified by the radiation length X_0. The mean energy of an ultrarelativistic electron decreases exponentially with traversed thickness:

$$E(x) = E_0 \exp\left(-\frac{x}{X_0}\right), \tag{2.3}$$

which defines X_0 in the high-energy limit, so that X_0 is the mean distance over which the energy is reduced by a factor e [33, 34]. For elements, a widely used empirical approximation in units of $\mathrm{g\,cm^{-2}}$ is

$$X_0 \simeq \frac{716.405\,A}{Z\,(Z+1)\,\ln\!\left(287/\sqrt{Z}\right)}\ \mathrm{g\,cm^{-2}}, \tag{2.4}$$

and the *length* radiation length follows from $X_0[\mathrm{cm}] = (X_0[\mathrm{g\,cm^{-2}}])/\rho$ with the mass density ρ [30].

Radiation length X_0 characterises both bremsstrahlung energy loss of electrons and the scale of multiple scattering for all charged particles. It depends primarily on Z and A and is quoted either as a mass thickness ($\mathrm{g\,cm^{-2}}$) or as a length (mm or cm) via division by the density. High–Z materials have short X_0, low–Z materials have long X_0.

Mixtures, compounds, and stacks For a mixture or compound, the radiation length follows the standard mass–fraction rule

$$\frac{1}{X_0} = \sum_i \frac{w_i}{X_{0,i}}, \tag{2.5}$$

with w_i the mass fraction and $X_{0,i}$ the elemental radiation length of component i [21]. The *material budget* is the traversed thickness expressed in units of radiation length,

$$\frac{x}{X_0} \qquad \text{(dimensionless material budget)}, \tag{2.6}$$

and for a stack of layers it adds linearly:

$$\left(\frac{x}{X_0}\right)_{\text{stack}} = \sum_i \frac{x_i}{X_{0,i}}. \tag{2.7}$$

For uncorrelated input uncertainties, a compact Gaussian error–propagation for the stack is

$$\sigma^2_{x/X_0} = \sum_i \left(\frac{\sigma_{x_i}}{X_{0,i}}\right)^2 + \sum_i \left(\frac{x_i}{X^2_{0,i}}\,\sigma_{X_{0,i}}\right)^2, \tag{2.8}$$

which is used later when quoting nominal budgets and their uncertainties.

Table 2.2 – Radiation lengths used in this thesis (values consistent with PDG material tables [35]).

Material	X_0 (mm)	Density
Aluminium (Al)	89.0 ± 0.4	$\rho \simeq 2.70\,\mathrm{g\,cm^{-3}}$
Silicon (Si)	93.7	$\rho \simeq 2.33\,\mathrm{g\,cm^{-3}}$
Nickel (Ni)	14.240 ± 0.014	$\rho \simeq 8.90\,\mathrm{g\,cm^{-3}}$

2.3 Multiple Coulomb scattering: Molière core and Highland form

Single-scatter tails (Rutherford) Elastic Coulomb scattering of a charge $z \cdot e$ off a nucleus of charge $Z \cdot e$ is described at leading order by the Rutherford differential cross section,

$$\frac{d\sigma}{d\Omega} = \left(\frac{z\,Z\,\alpha\,\hbar c}{2\,p\,\beta\,\sin^2\frac{\phi}{2}} \right)^2 , \tag{2.9}$$

which at small angles behaves as $d\sigma/d\Omega \propto 1/\phi^4$ [36]. The strong $1/\phi^4$ tail implies that even when the *core* of the multiple–scattering distribution is near–Gaussian (by the central–limit theorem), the *tails* remain non–Gaussian due to occasional large single scatters.

Multiple Coulomb scattering (MCS) arises from the accumulation of many small deflections and is well described by Molière theory [37, 38, 39]. For practical use, the projected RMS scattering angle is given, with excellent accuracy for thin scatterers, by the Highland relation including the Lynch–Dahl logarithmic refinement [40]:

$$\boxed{\theta_0 = \frac{13.6\,\mathrm{MeV}}{\beta\,c\,p}\; z\,\sqrt{\frac{x}{X_0}}\,\left[1 + 0.038\,\ln\!\left(\frac{x}{X_0}\cdot\frac{z^2}{\beta^2}\right)\right]} \tag{2.10}$$

where p is the particle momentum and the angles on the left are *projected (planar)* widths θ_x or θ_y as used throughout the analysis [41, 42]. As an approximation to the Gaussian core of the Molière distribution, Eq. (2.10) is accurate for thin, approximately homogeneous scatterers with small deflections and negligible energy loss, it does not capture the non-Gaussian Rutherford tails or large-angle scatters. At very small x/X_0 the weak logarithmic term, together with instrumental broadening or momentum uncertainty, makes inversion of Eq. (2.10) numerically delicate, whereas for thicker layers or low momenta energy loss and straggling reduce its accuracy. The 3D space–angle RMS is larger by a factor $\sqrt{2}$ for an azimuthally symmetric distribution. The associated RMS lateral displacement over a path length x follows to leading order as

$$\sigma_y \simeq \frac{x}{\sqrt{3}}\,\theta_0. \tag{2.11}$$

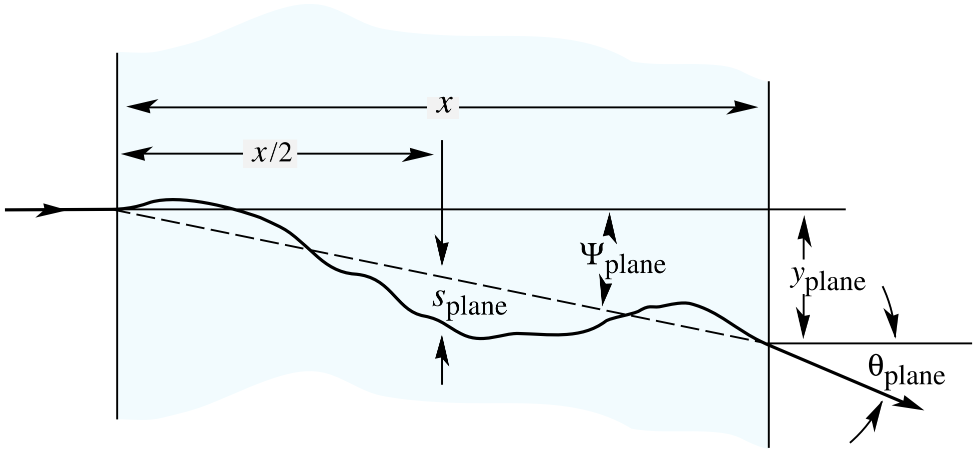

Figure 2.3 – Geometry for projected angles and lateral displacement in a thin scatterer. Taken from [30].

For traversals of several layers it is advantageous to form the total material budget first and apply Eq. (2.10) once to the sum, rather than adding individual θ_0 in quadrature. This avoids overcounting of the logarithmic term and reflects that the scattering kernel depends on the *total* x/X_0 [42].

2.4 From angle histograms to material budget

Experimentally one obtains, per spatial region, a distribution of projected deflection angles. Because Rutherford tails are present, the width used as MCS proxy is extracted from the central bulk of the distribution to reduce tail bias. In this work a Gaussian is fitted to a central quantile (e.g. 98 %) of the distribution, the resulting standard deviation defines θ_0 for that region. This procedure yields a stable estimator with minimal sensitivity to rare large scatters. The leading scaling $\theta_0 \propto \sqrt{(x/X_0)}/p$ provides immediate intuition: at fixed momentum, thicker regions widen the core as $\theta_0 \propto \sqrt{x/X_0}$, at fixed thickness, higher momentum reduces the width as $\theta_0 \propto 1/p$ [41, 43].
Given a measured projected width θ_0 and a known momentum p, the material budget is obtained by inverting Eq. (2.10). Because of the weak logarithmic term, the inversion is done iteratively: starting from $x^{(0)}/X_0 = (\theta_0 p/13.6\,\mathrm{MeV})^2$, update $x^{(i+1)}$ from $x^{(i)}$ by evaluating the logarithmic term at $x^{(i)}$. Iteration stops when $\left|x^{(i+1)} - x^{(i)}\right|/x^{(i)} < \varepsilon = 1 \times 10^{-12}$, or after $N_{\max} = 50$ steps. The same scheme applies pixel by pixel to form spatially resolved x/X_0 maps once a reliable momentum scale has been established by calibration targets.

Applicability to electrons and x/X_0 budgeting For electrons the collisional picture is supplemented by sizeable bremsstrahlung. Nevertheless, Eqs. (2.10)–(2.11) remain valid for the angular core as they depend on x/X_0 and kinematics, not on how energy is lost. When nominal budgets of foils or grids are stated, Eq. (2.7) is used with the X_0 values of Table 2.2, and uncertainties are propagated with Eq. (2.8). The inversion described above is the backbone of the material–budget maps reconstructed later.

Chapter 3

Acquisition & Reconstruction

To determine the material budget of next generation ultra-thin detectors, such as the ITS3, the assembly has to be decomposed into its constituent components. For each component, the thickness x and the radiation length X_0 are required. A similar procedure is done in the measurements here. Instead of measuring whole structures, material layers in the form of individual foils are measured by calculating the scattering angle of a traversing particle through the foil with a precise sensor telescope setup. The sensors used for this therefore had to have a very precise spatial resolution. The measurements were then performed at two beam energies, 2.4 GeV and 1.0 GeV, to provide a cross-check at a second momentum setting.

3.1 ALPIDE Sensor

Testing the limits of the Highland parametrisation requires both suitable material samples and a measurement setup with sufficient angular resolution to resolve small multiple–scattering deflections, hence spatial and angular resolutions are both relevant. The sensor used is the ALICE Pixel Detector (ALPIDE). It measures $15\,\mathrm{mm} \times 30\,\mathrm{mm}$ and is segmented into 512×1024 digital pixels of $26.88\,\mu\mathrm{m} \times 29.24\,\mu\mathrm{m}$ [44]. The ALPIDE is a monolithic active pixel sensor (MAPS): the silicon diodes and readout circuitry are integrated within the pixels; implemented in a 180 nm CMOS process [45] by Tower Semiconductor Ltd. [46]. It provides individual pixel readout and achieves a minimum ionising particle detection efficiency above 99 % [17]. The sensors provide a binary-hit mode with globally tuned and equalised thresholds to ensure a uniform response across the matrix (see [44]), thereby stabilising the single-point resolution entering the track-angle estimate.

3.2 Telescope Setup

The telescope geometry was optimised for angular resolution. Previous material-budget studies showed that good angular resolution is obtained when one measurement plane is placed as close to the target as possible and two additional planes are positioned further away and very close to each other [47]. This layout was adopted here symmetrically about the target and used six ALPIDE planes in total. Three of those being upstream and three being downstream. Here, upstream and downstream denote the sides before and after the target along the nominal beam direction.

Two constraints limit the best resolution: the minimum distance between the innermost ALPIDE and the target, and the minimum spacing between the two outer planes forming the lever arm. The latter is limited to 20 mm by the carrier-card geometry, while the target–sensor distance can be reduced to 15 mm. The remaining free parameter is the lever-arm distance between the inner plane and its neighbouring outer plane.

The lever arm was fixed at 140 mm because it maximised the angular resolution at the nominal beam momentum $p = 2.4\,\mathrm{GeV/c}$ (see Fig. 3.1), the regime of the majority of the data taking, and is consistent with EUDET-type telescope performance [48]. The resulting optimal telescope layout is shown in Fig. 3.2.

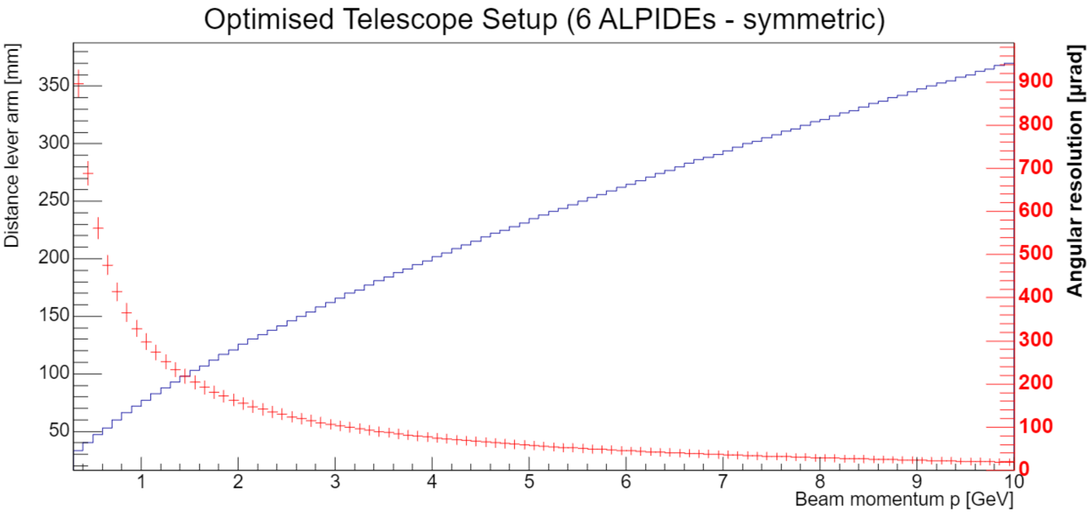

Figure 3.1 – Optimal telescope geometry and resolution[1]

3.3 Measurement Samples

The choice of samples fixes the accessible range in x/X_0. As the objective is to test the Highland relation at small material budgets, the samples were selected to populate that regime. Eight targets spanning three materials were prepared and are listed in Table 3.1. Six are uniform foils of size 30 mm × 13 mm covering most of the sensitive area of the ALPIDE. Aluminium, nickel and silicon were used. The silicon sample is a blank wafer and is treated as a uniform foil. In addition, two composite samples with 3 × 3 grid structure, illustrated in Figure 3.3, were prepared. These grids comprise nine regions

---

[1] Taken from GSI PI MAPS collaboration wiki, *Material Budget*, GitLab CERN (restricted access), accessed August 2025.

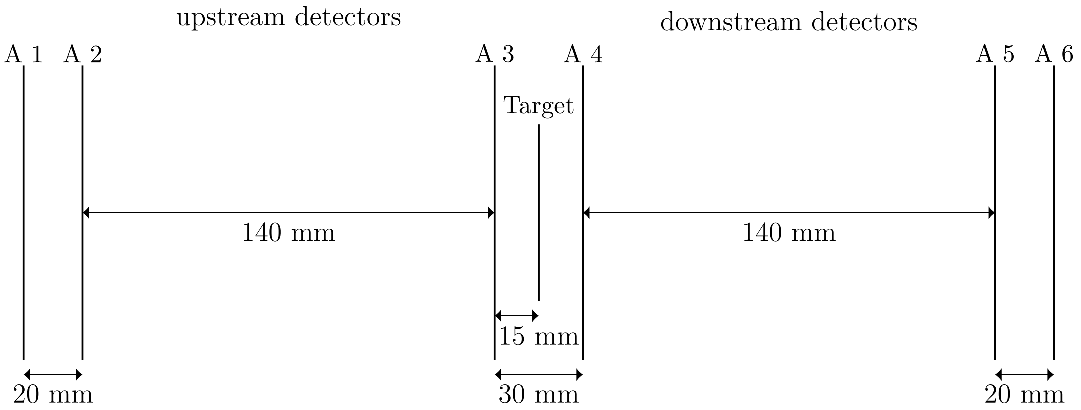

Figure 3.2 – Geometric arrangement of the six-plane ALPIDE telescope with the target at the centre.

with stepwise thickness differences created by stacking foils such that overlaps define the regional path lengths. One cell is empty and eight cells have increasing thickness. This is helpful, since it enables to cancel uncertainties due to the setup within the same measurement. Moreover, for some structured samples, there will be multiple regions with the same material thickness, giving a cross check for the systematics involved.

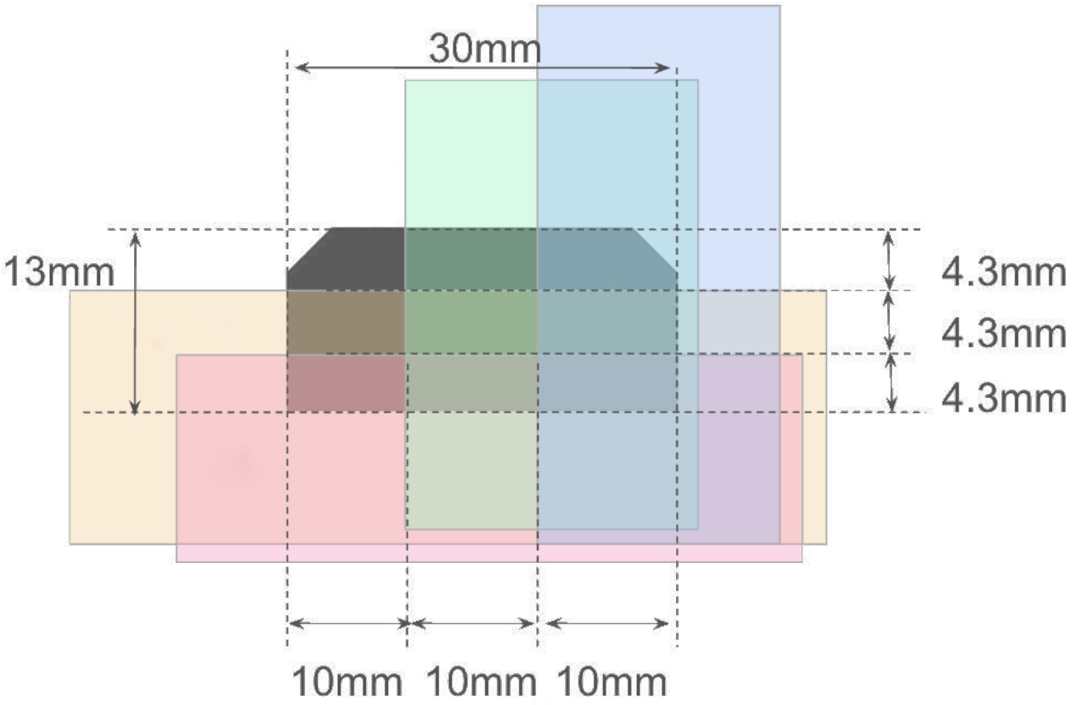

Figure 3.3 – Structure of a complex target composed of four stacked foils forming a 3×3 grid of regions with different thicknesses.[2] The top-left cell is empty (no foil), the bottom-right cell has the maximum thickness.

The nominal thicknesses were chosen to sample the small-x/X_0 region evenly with a limited number of foils. All key specifications (materials, nominal foil thicknesses x,

---

[2] Taken from GSI PI MAPS collaboration wiki, *Material Budget*, GitLab CERN (restricted access), accessed August 2025.

Radiation length X_0, and x/X_0 with uncertainties) are compiled in Tab. 3.1. To increase granularity in total effective thickness without increasing the number of samples, multiple uniform foils were combined and measured at once. The uniform-foil thicknesses were set to (10.0 ± 1.0) µm, (30.0 ± 1.0) µm, and (100.0 ± 1.0) µm for aluminium, and (6.0 ± 1.0) µm and (12.5 ± 1.0) µm for nickel, yielding in total 13 individual samples and allowing effective thicknesses between 6 µm and 280 µm.

Table 3.1 – Nominal thicknesses, radiation lengths, and corresponding material budgets for all measured targets. Radiation length values from the Particle Data Group [30].

Material	**Thickness** x (µm)	**Radiation length** X_0 (mm)	**Material Budget** $\frac{x}{X_0}$ (‰)
Aluminium	10.0±1.0	89.0 ± 0.4	0.112±0.011
	30.0±1.0		0.337±0.011
	100.0±1.0		1.124±0.012
Aluminium Grid	4.5±1.0		0.051±0.011
	9.0±1.4		0.101±0.016
	10.0±1.0		0.112±0.011
	14.5±1.4		0.163±0.016
	19.0±1.7		0.213±0.019
	20.0±1.4		0.225±0.016
	24.5±1.7		0.275±0.019
	29.0±2.0		0.326±0.023
Silicon	40.0±1.0	93.70	0.427±0.011
Nickel	6.0±1.0	14.240 ± 0.014	0.421±0.070
	12.5±1.0		0.878±0.070
Nickel Grid	3.0±1.0		0.211±0.070
	6.0±1.4		0.421±0.098
	9.0±1.7		0.63±0.12
	12.0±2.0		0.84±0.14

Uncertainties in the expected material budgets reflect both thickness tolerances and material purity. Each foil carries a thickness tolerance of $\sigma_{t,\text{foil}} = 1$ µm.[3] Assuming uncorrelated tolerances, a stack of N layers (including grid stacks) has

$$\sigma_{t,\text{stack}} = \sqrt{\sum_{i=1}^{N} \sigma_{t,i}^2} = \sqrt{N}\,\sigma_{t,\text{foil}}\,. \tag{3.1}$$

For example, with $\sigma_{t,\text{foil}} = 1$ µm and $N = 4$, one obtains $\sigma_{t,\text{stack}} = 2$ µm. Material purities are specified by the providers as 99.5 % for aluminium and 99.9 % for nickel, which enter as the uncertainty on X_0. The wafer-grade silicon impurity is negligible for the present purpose.

The distribution of all effective samples (including subtractions) is shown in Fig. 3.4: each

---

[3] Personal communication with B. Blidaru (Aug. 2025) based on discussions with the sample manufacturers.

point represents a possible measurement with a material combination. The x-axis value is the corresponding x/X_0 and the y-axis values represents the scattering angle distribution width, with horizontal error bars from the x/X_0 uncertainty and vertical bars from its propagation to θ_0. Together they span the small-x/X_0 regime up to $\sim 3.5 \times 10^{-3}$ and provide dense coverage down close to 10^{-4}, which would correspond to about 10 µm of aluminium.

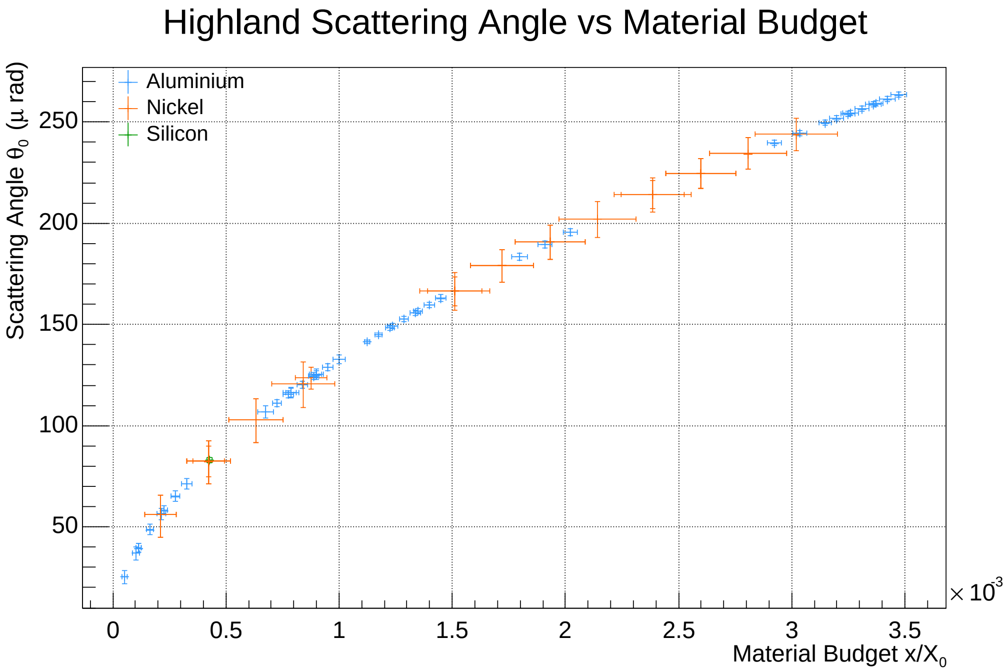

Figure 3.4 – Material budgets x/X_0 of all prepared targets, including uniform and grid structures.

3.4 Test Beam Campaign

Data were recorded at the test-beam facility of DESY [49, 50]. Seven days of beam time were allocated, with the beamline of area 22 shared with another setup, that did not affect the measurements reported here. Including setup, occasional access for the downstream experiment, and target exchanges, the net data-taking time amounted to 156 h and an average trigger rate of 7.5 kHz.

The ALPIDE telescope was installed in a light-tight enclosure, where the front panel could be opened manually for access. Two scintillators were placed upstream and downstream of the telescope to provide a coincidence trigger, ensuring that only electrons traversing all six planes were recorded (see Fig. 3.5). The electron beam entered from the right in the photograph.

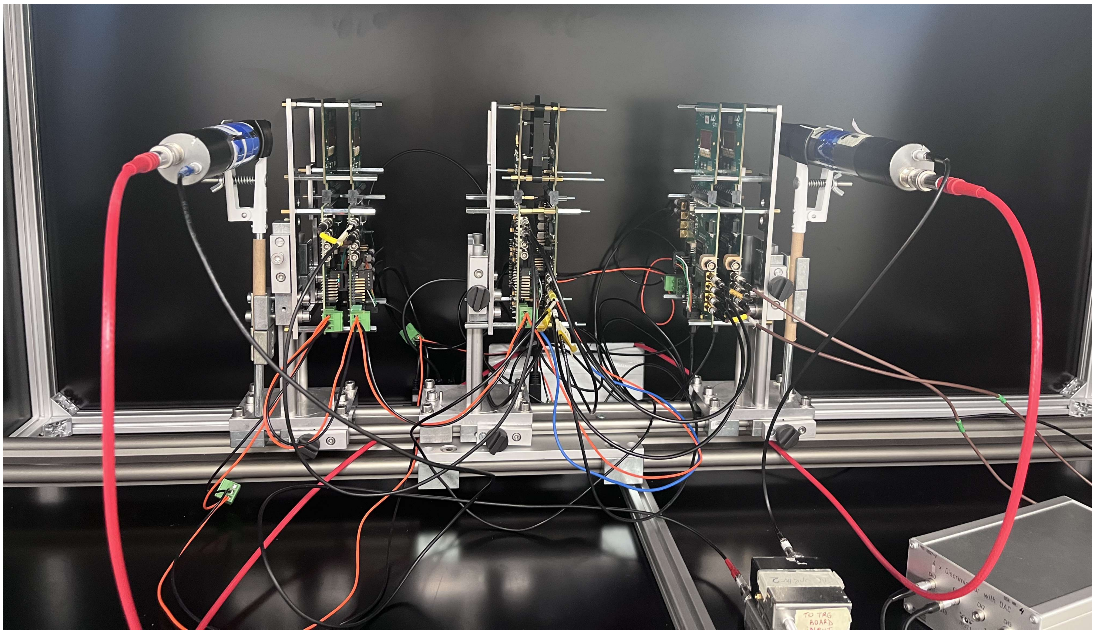

Figure 3.5 – Photograph of the ALPIDE telescope setup inside the light-shielding enclosure at the DESY test beam facility.

3.5 Alignment & Tracking

An initial mechanical alignment was performed using the laser guidance of the facility to centre the beam on the sensitive area. A mechanically stable reference was required for the track-based alignment and had to be decoupled from the target holder to avoid shifts during target exchanges. Therefore the second upstream plane was chosen as the reference, and all plane coordinate systems were defined relative to it. During beam operation, a precise alignment of all six planes was then obtained with the `Corryvreckan` [51] module *AlignmentTrackChi2*. The module varied sensor positions in x and y and the three rotational angles, while keeping z positions and inter-plane distances fixed, since small changes in beam direction had a negligible impact on the alignment for the chosen geometry.

The alignment proceeded as a standard global χ^2 minimisation: with the current geometry, straight tracks were fitted, residuals $r_i \equiv M_i - F_i$ between measured and fitted hit positions were computed for all planes, a linearised least-squares system was solved for the alignment corrections, and the geometry was updated. These steps were iterated to convergence. In compact form,

$$\chi^2 = \sum_{i \in \text{planes}} \frac{(M_i - F_i)^2}{F_i}, \tag{3.2}$$

where M_i and F_i denote measured and fitted positions, respectively. Eight iterations were sufficient for the telescope used here.

Because the target holder and central planes can move slightly during exchanges, a fresh alignment was performed for each measurement, enabling a comparison of alignment quality across runs (Fig. 3.6). The alignment remained stable over the week with an average $\chi^2/n_{\mathrm{dof}} = 1.46$. Runs at 1.0 GeV exhibit larger χ^2/n_{dof} due to increased large-angle scattering at lower momentum. The power-regulator target seen as the third data point in Figure 3.6 shows a poorer fit, since it has an inhomogeneous structure. Nevertheless, this target is not important for the following work done here.

To test for slow drifts within a run, the aluminium 280 µm data set was realigned from scratch every 10 min (17 alignments). No significant trend was observed, supporting the strategy of using one alignment per measurement.

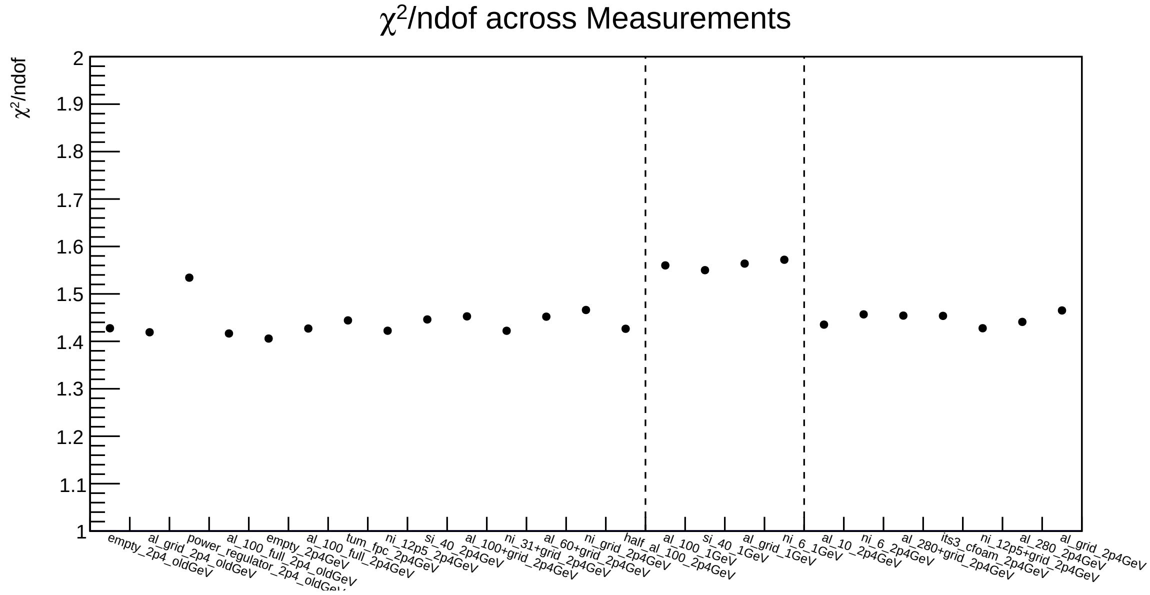

Figure 3.6 – Alignment quality, quantified by χ^2/n_{dof}, for each measurement in the test beam campaign. Vertical dashed lines indicate the change in beam momentum from 2.4 GeV to 1.0 GeV and back.

After alignment, tracks were reconstructed. Straight-line fits were applied separately to the upstream and downstream triplets. Candidate pairs were associated when their extrapolated impact points at the target agreed within 100 µm × 100 µm. This configuration in the `Corryvreckan` module *TrackingMultiplet* yields the scattering angle between the two lines at the target position. A General Broken Line (GBL) [52] approach was considered but the two-segment model with a single scattering point was found to avoid overfitting minor effects while matching the actual conditions.

As seen in Figure 3.7, a first measure of the material under observation is found in the scattering angle θ. The measured scattering angle is three-dimensional. The projected scattering angle is obtained from the transverse components as

$$\theta_{\mathrm{net}} = \sqrt{\theta_x^2 + \theta_y^2}\,. \tag{3.3}$$

Because multiple Coulomb scattering is stochastic, repeated traversals do not have the same angle. Collecting many events therefore yields a distribution of scattering angles. For the subsequent analysis, the relevant observable derived from this distribution is its

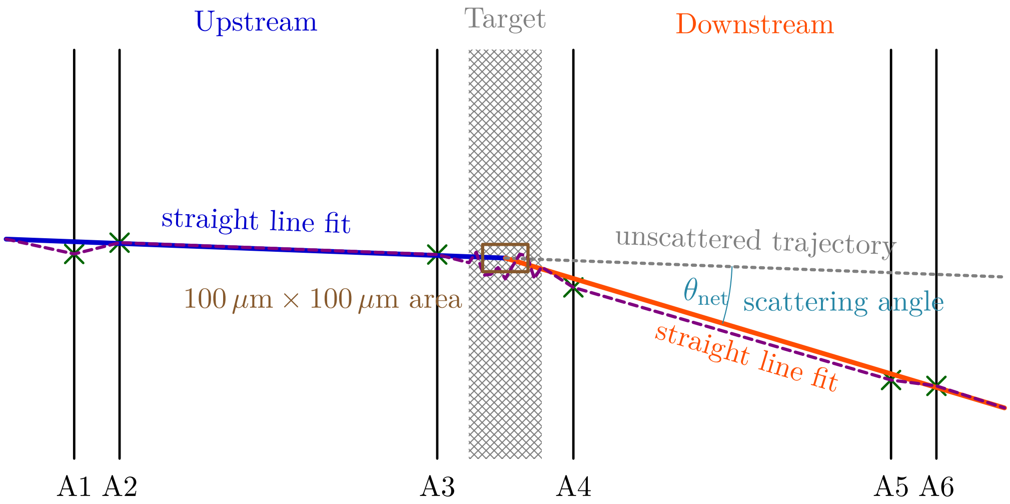

Figure 3.7 – Schematic of the tracking telescope with an example event showing multiple scattering in the target. The scattering angle θ is the angle between the downstream track and the unscattered reference line at the target position. Distances in beam direction are to scale, the target thickness is exaggerated for visibility.

width. Rather than a single global width, spatially resolved widths across the sensor are required, which motivates the *Pixelisation* step below.

3.6 Pixelisation

Pixelisation divides the sensitive sensor area into equally sized cells (the so-called *pixels*). In each cell an independent scattering-angle distribution is accumulated and fitted. The number of entries per cell decreases with the cell size, while spatial resolution improves. An appropriate balance between statistical precision and granularity is therefore required.

Again, a single measurement is taken to examine the relation between the resolution and the statistical precision of the scattering distribution. To quantify statistical robustness locally, the fluctuation of Gaussian widths between neighbouring cells is used (concept shown in Fig. 3.8). Denote by $\sigma_{\text{RMS},c}$ the Gaussian width in the central (orange) cell c and by $\mathcal{N}_4(c)$ its four von–Neumann neighbours (yellow). The pairwise fluctuation is

$$\Delta\sigma_{c\to n} = \left|\sigma_{\text{RMS},c} - \sigma_{\text{RMS},n}\right| . \tag{3.4}$$

The local mean fluctuation is

$$\overline{\Delta\sigma}_c = \frac{1}{|\mathcal{N}_4(c)|} \sum_{n\in\mathcal{N}_4(c)} \Delta\sigma_{c\to n} , \tag{3.5}$$

and the global indicator is the average over all valid cells $\mathcal{C}$ of the sensor area,

$$\langle \Delta\sigma \rangle = \frac{1}{|\mathcal{C}|} \sum_{c \in \mathcal{C}} \overline{\Delta\sigma}_c \,. \tag{3.6}$$

Edge cells with fewer than four neighbours are normalised by $|\mathcal{N}_4(c)|$. Averaging this value over all cells in the sensor area, a quantified indicator for the data statistics is formed.

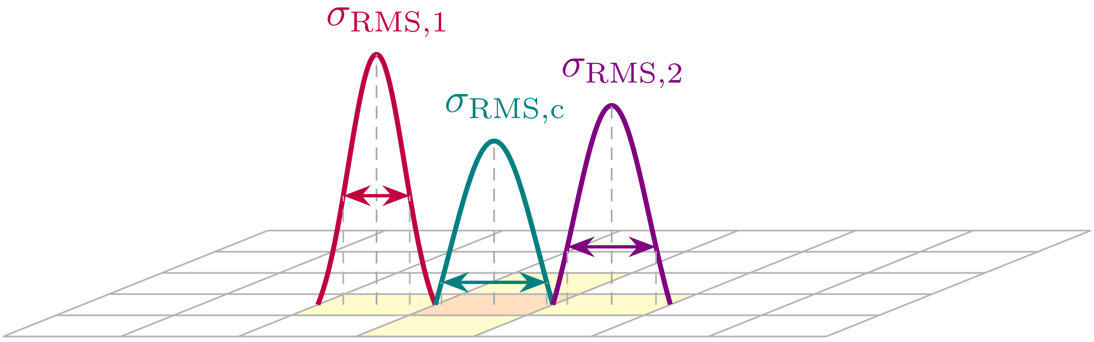

Figure 3.8 – Illustration of the *local* mean fluctuation $\overline{\Delta\sigma_c}$: for the orange cell c, take the absolute differences in RMS width σ_{RMS} to its four von–Neumann neighbours (yellow) and average them.

Figure 3.9 shows the fluctuations versus cell size from 1 mm to 0.06 mm. Smaller cells exhibit larger fluctuations due to reduced statistics, while larger cells improve stability at the cost of spatial resolution. For the purposes here, sub-100 µm resolution was not required, but precise widths were. A cell size of 250 µm × 250 µm was therefore adopted, yielding more than 5000 cells per measurement area.

The choice of fit function also matters. Alternatives such as a Gaussian convolved with Student-T tails introduce additional degrees of freedom compared with a pure Gaussian (mean, standard deviation only). In low-statistics cells or with varied fit intervals this flexibility can destabilise parameter estimates and bias widths, compared to the gaussian, if no pronounced non-Gaussian features are present [53]. Given the observed residuals, the Gaussian model provides a stable and reproducible width estimate with minimal risk of overfitting and is adopted as default.

Because a Gaussian does not accurately describe the long tails of the scattering angle distributions, fits were restricted to a central interval [54]. Quantile cuts of 98 % and 95 % were compared (motivated by [53]). A 98 % cut retains the innermost 98 % of the distribution and excludes the outer 1 % on each side. Repeating the fluctuation study of Fig. 3.9 for both central-interval cuts yielded curves that are distinguishable but differ only marginally. Retaining the additional 3 % under the 98 % cut was judged more informative and was adopted for the main analysis. The 95 % analogue is shown in App. A.2, Fig. A.1. Consequently, the 98 % quantile was adopted, consistent with the estimator used by Lynch and Dahl [42]. An example final pixelised map for a 6 µm thick nickel sample is shown in Fig. 3.10.

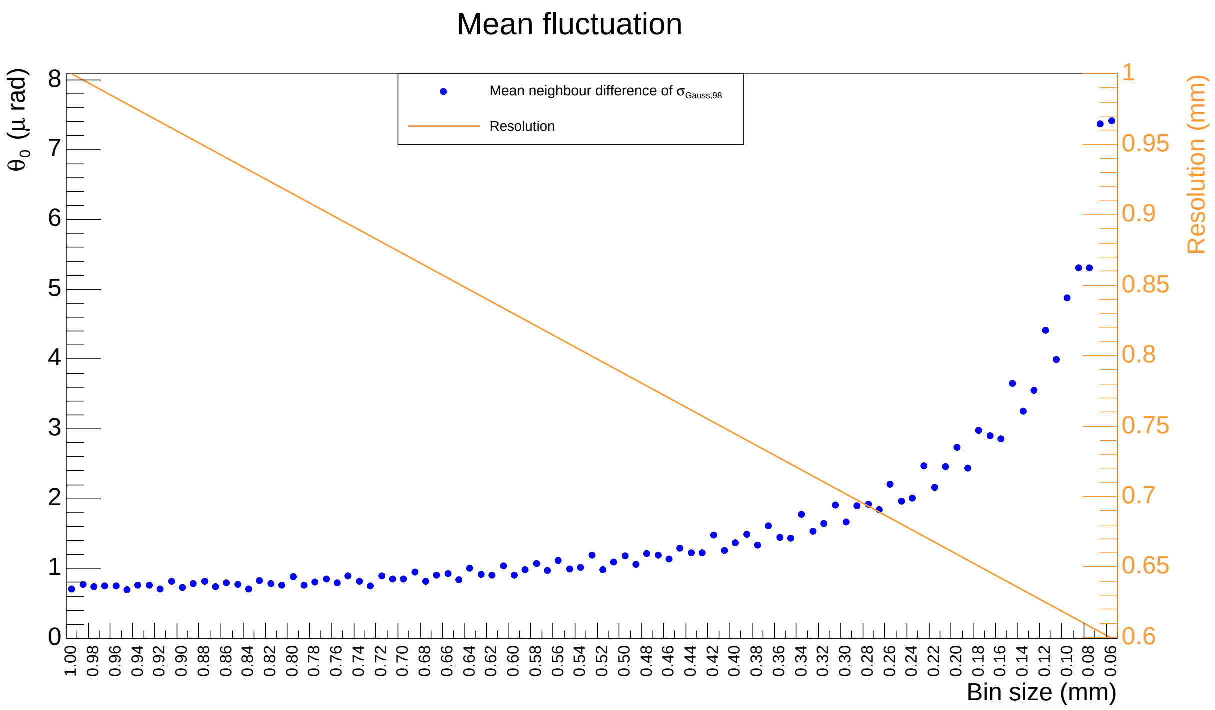

Figure 3.9 – Mean fluctuation $\overline{\Delta\sigma}$ of Gaussian widths between von Neumann neighbours as a function of cell size using a 98 % central interval [54]. App. A.2, Fig. A.1 shows the 95 % variant. The curves differ slightly but the effect is negligible for the choice of pixelisation resolution.

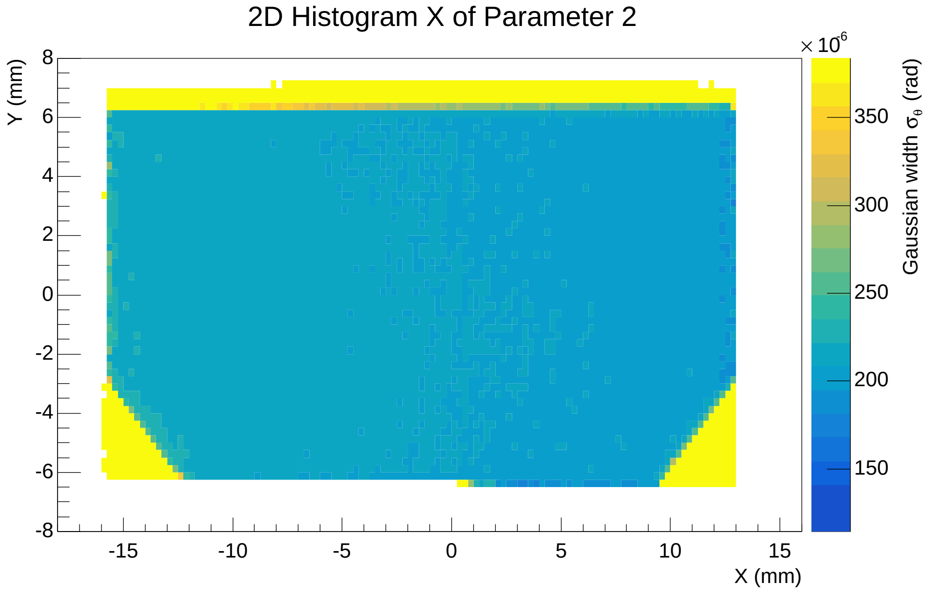

Figure 3.10 – Example of a final pixelised 2D histogram of Gaussian widths after applying the 98 % quantile cut, shown for Nickel 6 µm with a cell size of 250 µm × 250 µm.

Chapter 4

Material Budget

After completing the pixelisation step described in the previous chapter, the analysis proceeds to the core objective of this study: converting the measured scattering angles into quantitative material budget values. This transformation is not straightforward, as the raw scattering angles are influenced by a range of additional effects beyond the interaction with the target material itself.

To obtain accurate and physically meaningful two-dimensional material budget distributions, these influences must first be corrected for. The process therefore involves a sequence of intermediate steps: analysing and refining the scattering angle distributions, applying corrections for systematic offsets and beam momentum gradients, quantifying statistical fluctuations, and calibrating the results against well-characterised reference targets. Each of these stages plays a crucial role in ensuring that the final material budget values reflect only the true properties of the investigated structures. The methodology and results of these steps are presented in the following sections.

4.1 Scattering Angle Distributions

The first stage after completing the pixelisation step is a thorough investigation of the measured scattering angle distributions for each target configuration. This step is critical because it forms the foundation for all subsequent corrections and calibrations in the material budget reconstruction.

For structured targets such as the Nickel grid, the spatial variation of the scattering angles can be visually identified in the 2D distributions of the fitted Gaussian widths. The per-cell Gaussian width are denoted by σ_θ. After subtraction and scaling, $\sigma_\theta^{\mathrm{corr}}$ is used. These distributions reveal distinct patterns that directly reflect the underlying target geometry: areas with an increasing thickness appear with higher scattering angles, reflecting the corresponding increase in material thickness (see Fig. 4.1). This relation is a direct consequence of the material-dependent nature of multiple Coulomb scattering, as described by the Highland formula [41, 42].

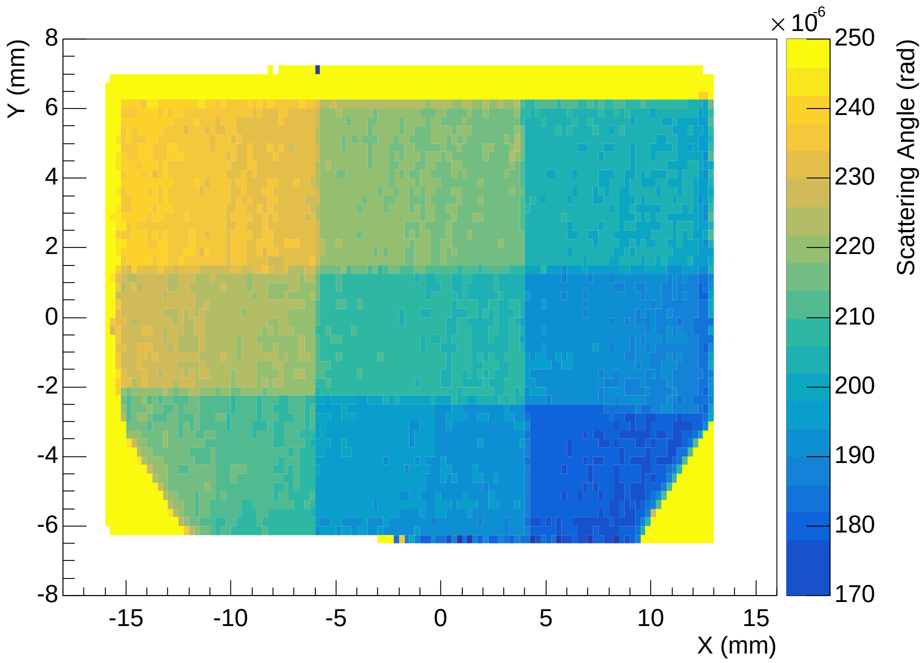

Figure 4.1 – Two-dimensional map of the fitted Gaussian widths of the scattering angle distributions for the Nickel grid target, showing spatial variations that correspond to the underlying grid structure. Regions with more foils exhibit larger scattering angles due to the increased material thickness, in agreement with the material-dependent multiple Coulomb scattering as parametrised by Highland/Lynch-Dahl [41, 42].

In addition to the intended geometric pattern, the two-dimensional distributions also contain features unrelated to the target design. For instance, in the lower region of the Nickel grid map, one can observe two sharp, straight-edged areas with elevated scattering angles. These structure originates from mounting the pixel sensor chip onto its carrier card.

The straight-edged bands at the top and bottom are consistent with the adhesive and the carrier card, which locally increase the effective material budget and thus the measured scattering. By contrast, the left-hand feature is most likely a residual misalignment (e.g. a small relative shift or rotation) that inflates the apparent scattering in that region.

It is important to emphasise that the measured scattering angle in a given pixel is not purely determined by the interaction of the beam particles with the target material. Multiple additional contributions affect the measured widths, including:

- Instrumental resolution of the tracking telescope, which inherently broadens the reconstructed scattering angle distributions and imposes a lower bound on the measurable widths.
- Residual multiple scattering from non-target materials in the beam line, such as air, upstream and downstream sensors.
- Systematic effects in the reconstruction process, for example from imperfect alignment or resolution variations across the sensor planes, which can introduce biases in the fitted widths.

Previous work systematically studied these contributions and demonstrated that the commonly used quadratic subtraction of a reference measurement (e.g., an *empty* target holder) is insufficient to fully remove them [47, 53]. The limitation of the quadratic subtraction arises from the fact that not all contributions scale quadratically with material thickness. Some effects, instrumental resolution in particular, are constant offsets that remain after subtraction, while others scale differently and cannot be captured by a simple quadratic calculation.

To overcome this limitation, a multiplicative scaling factor is applied in the present analysis. After subtracting the reference measurement, the resulting scattering angle distributions are scaled by this factor, which is determined from calibration measurements. This procedure ensures that the combined background and resolution effects are properly normalised across all measured thicknesses.

In this work, the scaling correction is applied consistently to all target measurements, both uniform and structured. As a result, the corrected two-dimensional scattering angle distributions represent only the scattering induced by the target material itself, free from contributions due to the telescope resolution, residual background scattering, or unrelated upstream/downstream materials. These refined distributions form the input for the subsequent steps in the material budget extraction.

4.2 Momentum Gradient Correction

The multiple–Coulomb–scattering width scales inversely with the particle momentum. In the Highland parametrisation (see Eq. 2.10), a bias in the momentum p translates directly into a bias in the inferred material budget x/X_0. To first order (neglecting the logarithmic term), $x/X_0 \propto (p\,\theta_0)^2$, so an overestimate of p by $\Delta p/p$ biases x/X_0 high by approximately $2\,\Delta p/p$. If, for instance, the analysis assumes the nominal beam setting $p_{\mathrm{nom}} = 2.4\,\mathrm{GeV/c}$ while the true momentum at the device under test is lower, the measured widths θ_0 are larger than expected and the inversion of Eq. (2.10) overestimates x/X_0.

At the DESY test beam, a momentum shortfall, due to energy loss, of order $0.2\,\mathrm{GeV/c}$ between the nominal beamline setting and the effective momentum within the ALPIDE telescope is expected. In this data set, the effective momentum inside the telescope was found to be $(2.20 \pm 0.10)\,\mathrm{GeV/c}$, i.e. about 9.2 % lower than the nominal $2.4\,\mathrm{GeV/c}$. Without correcting for this, the reconstructed material budget would be biased high by over 18 % (again ignoring the weak logarithmic dependence), which is unacceptably large for the precision targeted here.

The correction proceeded in three connected stages: (i) determine the effective momentum scale p_{true} from uniform calibration targets, (ii) cancel a residual transverse momentum gradient (common to all targets) at the map level and (iii) apply the resulting momentum description consistently in the inversion from θ_0 to x/X_0. The subsections below follow this workflow.

Principle of the correction

The strategy is to *measure* the effective momentum p_{true} at the telescope position from the data and the nominal setting using calibration targets of known thickness [55, 56], and then to use p_{true} consistently in all subsequent conversions between θ_0 and x/X_0.

This step used uniform calibration targets with well-known thickness x and radiation length X_0, scattering-width distributions already corrected for instrumental/background effects using the subtraction and global scaling described in the previous section and regions of interest on each calibration target where the material is spatially uniform that the statistical width is well above the residual resolution floor.

Operationally, this principle was implemented first on a per-target basis to obtain apparent momenta $\hat{p}_i$ by inverting Eq. (2.10), as detailed next.

Inversion per target

For each calibration target i, a corrected width $\hat{\sigma}_{\theta,i}$ (ROI mean of σ_θ) was extracted as the mean of the per-cell Gaussian widths of the scattering angle distribution over the chosen region. With $\beta \simeq 1$ for $2.4\,\mathrm{GeV/c}$ electrons, Eq. (2.10) was inverted to obtain the apparent momentum

$$\hat{p}_i = \frac{13.6\,\mathrm{MeV/c}}{\hat{\sigma}_{\theta,i}}\sqrt{\frac{x_i}{X_{0,i}}}\left[1 + 0.038\ln\left(\frac{x_i}{X_{0,i}}\right)\right] . \tag{4.1}$$

Very thin targets were not used for this step, because for small x_i the width approaches the resolution floor and the fractional uncertainty on $\hat{p}_i$ becomes large. Moreover, this is also where the Highland approximation becomes less accurate (Sec. 2.3). This regime was analysed separately and was not used for calibration.

Uncertainties on $\hat{p}_i$ were propagated from (a) the statistical error on $\hat{\theta}_i$, estimated from the dispersion of per-pixel widths divided by $\sqrt{N_{\mathrm{pix}}}$, (b) thickness tolerances on x_i and (c) material constants $X_{0,i}$. Linear propagation using Eq. (4.1) gives, to good approximation,

$$\left(\frac{\sigma_{\hat{p}_i}}{\hat{p}_i}\right)^2 \simeq \underbrace{\left(\frac{\sigma_{\hat{\theta}_i}}{\hat{\theta}_i}\right)^2}_{\text{(a) statistical error}} + \left(\frac{1}{2} + \frac{0.038}{1 + 0.038\ln\left(\frac{x_i}{X_{0,i}}\right)}\right)^2 \left[\underbrace{\left(\frac{\sigma_{x_i}}{x_i}\right)^2}_{\text{(b) thickness}} + \underbrace{\left(\frac{\sigma_{X_{0,i}}}{X_{0,i}}\right)^2}_{\text{(c) material constant}}\right] . \tag{4.2}$$

The global scaling factor applied to the widths was treated as a common systematic. Its contribution was accounted for when assigning the final systematic on p_{true} (it cancels in relative comparisons among the calibration targets).

The set of $\hat{p}_i$ then provided the inputs for a consolidated estimate of the momentum field and its average, obtained through an inverse-variance combination and examined versus transverse position X.

Global momentum estimate

After computing the per-target momenta $\hat{p}_i$ for all calibration targets, the results were compared as functions of the transverse position X inside the ALPIDE telescope (Fig. 4.2). Each curve in the figure corresponds to one calibration target, with the data points representing the average momentum extracted in bins of X. The near-parallel behaviour of the individual targets reflects the consistency of the measurement across different materials and thicknesses. A band indicated the average momentum p_{true} obtained via the inverse-variance weighted combination (Eq. 4.3), which follows the observed gradient in X. Most targets lie within this common band, while a few configurations exhibited systematic offsets and are treated as outliers in Fig. 4.2. Two very thin foils (Ni 6 µm, Al 10 µm) trended to lower apparent momenta because their small material budgets drove the projected widths toward the instrumental floor, degrading the precision of the extracted angles (cf. Sec. 2.3). These were therefore not used for calibration. By contrast, one configuration (Al 100 µm) showed a higher apparent momentum attributable to a residual misalignment during data taking. That dataset was discarded and remeasured. The second recorded dataset fell within the common band, essentially on top of the average, and was included in the final calibration. Accordingly, the two very thin foils and the misaligned original dataset were excluded from the global average, whereas the remeasured Al 100 µm dataset was included.

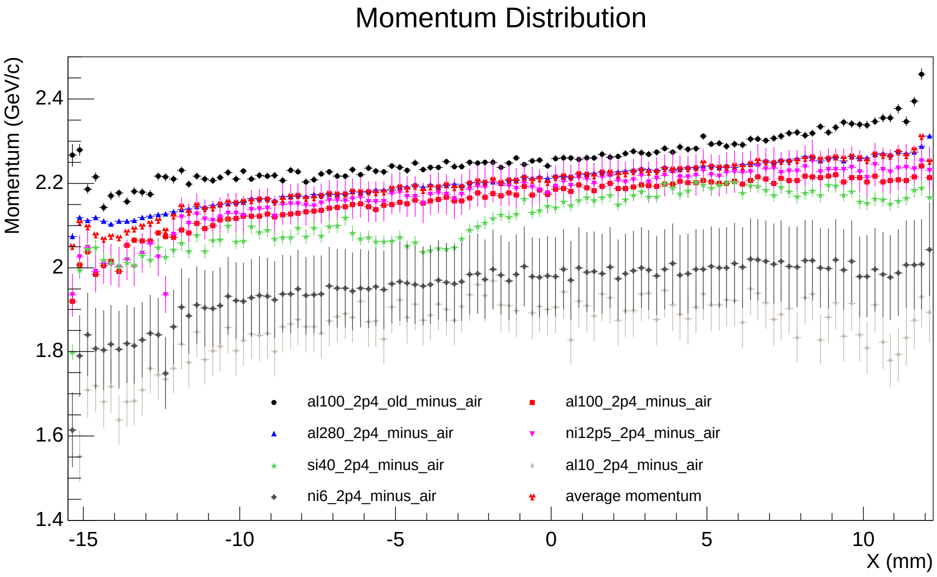

Figure 4.2 – Measured momentum distributions as a function of transverse position X for the calibration targets at 2.4 GeV/c nominal beam energy. Each set of markers corresponds to a different target configuration, the red series with three-triangle markers indicates the inverse-variance weighted average momentum p_{true} inside the telescope.

The per-target momenta $\hat{p}_i$ were combined into a single estimate using an inverse-variance weighted average,

$$p_{\text{true}} = \frac{\sum_i \hat{p}_i \, \sigma_{p,i}^{-2}}{\sum_i \sigma_{p,i}^{-2}}, \qquad \sigma^2(p_{\text{true}}) = \frac{1}{\sum_i \sigma_{p,i}^{-2}}. \tag{4.3}$$

Here the variance of p_{true} is given by the inverse of the total weight, with each calibration target contributing with weight $1/\sigma_{p,i}^2$ according to its uncertainty. Consequently, a measurement with twice the standard uncertainty carries one quarter of the weight of the more precise one.

Applying the procedure to the selected set of calibration targets yielded

$$p_{\text{true}} = (2.20 \pm 0.10)\,\text{GeV/c} \tag{4.4}$$

i.e. a shortfall of $\Delta p = p_{\text{nom}} - p_{\text{true}} = 0.202\,\text{GeV/c}$ relative to the nominal setting. The corresponding momentum scale factor is

$$c_p \equiv \frac{p_{\text{true}}}{p_{\text{nom}}} = 0.916\,. \tag{4.5}$$

The comparison across targets also revealed a mild, common slope versus X, consistent with a position-dependent momentum. The common transverse gradient was corrected in the next section by applying a position-dependent momentum rescaling based on the global gradient determined here.

Local momentum gradient

A mild transverse gradient was observed in the per-target widths across the telescope acceptance (Fig. 4.2). Since $\theta_0 \propto 1/p$, this behaviour is consistent with a position-dependent momentum that varies with the horizontal coordinate X. A one-dimensional momentum profile $\bar{p}(X)$ was extracted from the calibration data (Fig. 4.2). To remove the common gradient from any two-dimensional width map $\theta(X, Y)$, a column-wise rescaling was applied using this profile:

$$\theta^{\text{flat}}(X, Y) \;\equiv\; \theta(X, Y)\, \bar{p}(X)\,, \tag{4.6}$$

which cancels the X-dependent $1/p$ variation to first order. After this rescaling, the residual X-dependence was consistent with zero within uncertainties. No local modelling beyond the measured $\bar{p}(X)$ and no local inversion of Eq. (2.10) were applied at this stage.

Uncertainties. Uncertainties were propagated through the multiplicative correction. Assuming independent θ and $\bar{p}(X)$,

$$\sigma^2\big(\theta^{\text{flat}}\big) \;=\; \bar{p}(X)^2\, \sigma_\theta^2 \;+\; \theta(X, Y)^2\, \sigma_{\bar{p}(X)}^2\,. \tag{4.7}$$

4.3 Fluctuations in Uniform Material Budget

The next step of the analysis is to quantify pixel-to-pixel fluctuations in the reconstructed material budget and to assess whether discrete thicknesses of different target types, uniform foils, grid structures, and combined uniform+grid structures, can be separated reliably. The inputs to this study are the per-pixel corrected scattering widths, denoted θ_0^{corr}, which already incorporate both the momentum calibration and the background subtraction described in the previous sections. These widths are converted to material-budget values x/X_0 through inversion of the Highland parametrisation (Eq. 2.10).
Neglecting the weak logarithmic term in Eq. (2.10), the Highland expression reduces to

$$\theta_0^{\text{corr}} \propto \sqrt{\frac{x}{X_0}}\,, \tag{4.8}$$

so that the material budget can be written in terms of the measured width as

$$\frac{x}{X_0} \propto \left(\theta_0^{\text{corr}}\right)^2. \tag{4.9}$$

The relative uncertainty on x/X_0 therefore follows from Gaussian error propagation:

$$\frac{\sigma_{x/X_0}}{x/X_0} = 2\,\frac{\sigma_{\theta_0^{\mathrm{corr}}}}{\theta_0^{\mathrm{corr}}}\,. \tag{4.10}$$

Because $x/X_0 \propto \left(\theta_0^{\mathrm{corr}}\right)^2$, small relative fluctuations of the per-cell width carry linearly (with a factor of two) into the relative fluctuations of x/X_0. The sampling uncertainty of a width estimated from N independent angles in a pixelised cell decreases generically as $\sqrt{N}$. Consequently, the per-cell relative spread of x/X_0 falls as $\sqrt{N}$ and is controlled primarily by the local occupancy. With the adopted cell size of $0.25\,\mathrm{mm} \times 0.25\,\mathrm{mm}$ and the requirement $N \geq 500$ angles per cell, this places a statistical floor in the few-percent range. Cells with higher occupancy fall correspondingly below this level. Each uniform-target map comprised ~ 5300 cells, while individual grid regions contributed ~ 600 cells. These counts stabilise the Gaussian fit of the *across-cell* distribution, whereas the per-cell precision remains governed by N as above.

This factor-of-two amplification and the $\sqrt{N}$ scaling explain why the observed spread of x/X_0 across pixels is dominated by statistical precision rather than by systematic effects, and they motivate the choice of sufficiently large per-cell statistics to keep $\sigma_{x/X_0}/(x/X_0)$ in the few-percent regime wherever N comfortably exceeds the selection threshold.

Method

For each target, the corrected per-pixel scattering widths θ_0^{corr} are converted to material budget values x/X_0 using Eq. (2.10). Each pixelisation cell contributes one entry to a one-dimensional histogram, such that the histogram represents the distribution of material budget values across all cells of the target.

To extract quantitative fluctuation metrics, the histograms are fitted with a single Gaussian distribution,

$$f(x) = \mathcal{N}(\mu, \sigma)\ , \tag{4.11}$$

with μ representing the mean material budget of the target and σ quantifying the fluctuations around this mean. The Gaussian fit is chosen over a simple root-mean-square (RMS) evaluation because it naturally downweights outliers in the histogram tails, which typically arise from edge effects or statistical noise in low-occupancy pixels. To further suppress any noise, only cells with at least 500 entries were included in the histograms, cells with fewer entries were excluded. Since the distribution is expected to approach a Gaussian in the limit of large data, this fit-based approach yields a robust and unbiased estimate of σ, which is used consistently as the fluctuation metric throughout this analysis.

Figure 4.3 shows the material budget histogram for the uniform aluminium target together with its Gaussian fit. The fitted Gaussian width σ_{uni} quantifies the fluctuations in a homogeneous sample and serves as a reference for the following analyses.

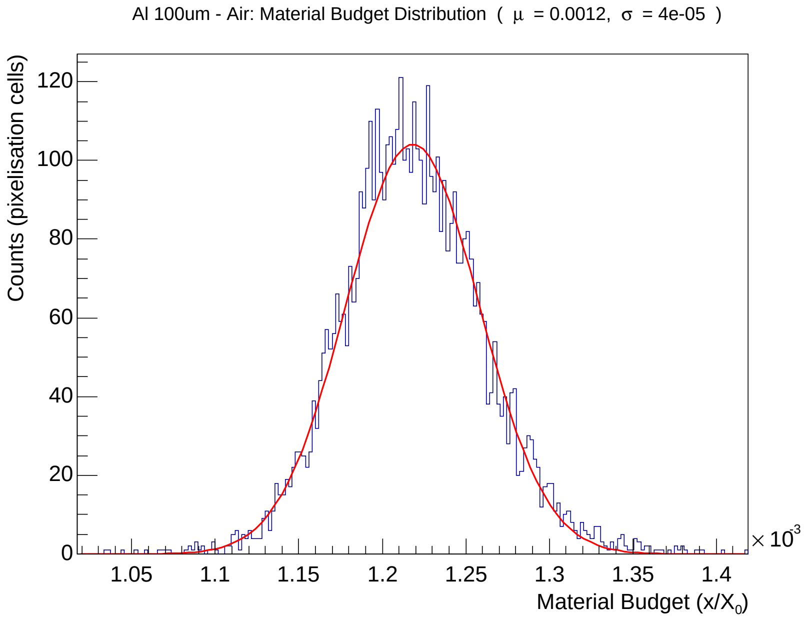

Figure 4.3 – Material budget histogram for the uniform 100 µm aluminium foil. The Gaussian fit (red) captures the spread of per-cell values and provides the reference fluctuation width σ_{uni} used in the analysis.

Uniform targets

Uniform foils produced a single, relatively narrow peak centred at μ in the per-pixel material-budget distribution. The fitted width σ_{uni} set the baseline fluctuation scale for a given run and thickness, and the ratio $\sigma_{\mathrm{uni}}/\mu$ quantified the relative precision of the total estimate. In this dataset, typical orders of magnitude were $\mu \sim 10^{-4}$–10^{-3} and $\sigma_{\mathrm{uni}} \sim 10^{-5}$, corresponding to $\sigma_{\mathrm{uni}}/\mu$ in the few-percent range. These values are consistent with the expectation that fluctuations in x/X_0 are dominated by statistical precision. Representative Gaussian fit parameters for a selection of uniform foils are summarised in Table 4.1, while the list of the most important fitted targets is provided in Appendix A.1.

It is further noteworthy that the relative fluctuation widths are independent of the material type. Expressing all results in units of radiation length x/X_0 absorbs differences in density and atomic composition into X_0, such that the ratio σ/μ reflects purely statistical counting effects. This explains why aluminium, nickel, and silicon targets alike exhibit fluctuations in the few-percent regime once normalised to their radiation lengths.

The observed σ_{uni} was directly linked to the chosen pixelisation in Sec. 3.6. With a cell size of 0.25 mm × 0.25 mm, each cell accumulated a finite number of scattering-angle measurements. Smaller cells would reduce this number and increase σ (while improving spatial granularity), whereas larger cells would lower σ at the cost of spatial resolution. The adopted cell size therefore represented a balanced choice that kept $\sigma_{\mathrm{uni}}/\mu$ at the few-percent level while preserving the ability to resolve discrete thickness classes in structured targets.

Table 4.1 – Uniform aluminium foils: Gaussian fit parameters for a representative selection of targets. The second column lists the *expected* material budget from nominal thickness using Al $X_0 = 8.897\,\text{cm}$. Relative widths σ_{uni}/μ remain at the few-percent level across most thicknesses, consistent with statistical precision–dominated fluctuations. A more detailed list of fitted values is given in Appendix A.1.

Nominal thickness (µm)	$(x/X_0)_{\text{exp}}$ $(\times 10^{-4})$	μ_k $(x/X_0 \times 10^{-4})$	σ_k $(x/X_0 \times 10^{-5})$	σ_k/μ_k (%)
10	1.12	1.8	3	17.22
100	11.24	12.0	4	3.33
280	31.47	32.0	5	1.59

Structured targets

Nickel grid

The grid geometry implied discrete path lengths and thus discrete material budgets. Accordingly, the per-pixel x/X_0 histogram for the nickel grid, seen in Figure 4.4, showed multiple, well-separated peaks at distinct μ_k values corresponding to the grid thickness classes. A residual *empty* peak (from subtracting the empty reference in the empty grid region) was ignored for the fluctuation analysis. For all populated classes, the per-peak widths were comparable to the uniform baseline, $\sigma_k \simeq \sigma_{\text{uni}}$, indicating statistical-precision–dominated spreads. Relative widths clustered at a few percent, while the peak separations $\Delta\mu$ satisfied $\Delta\mu \gg \sigma_k$, ensuring clean separability across the full range. The corresponding Gaussian fit parameters are listed in Table 4.2.

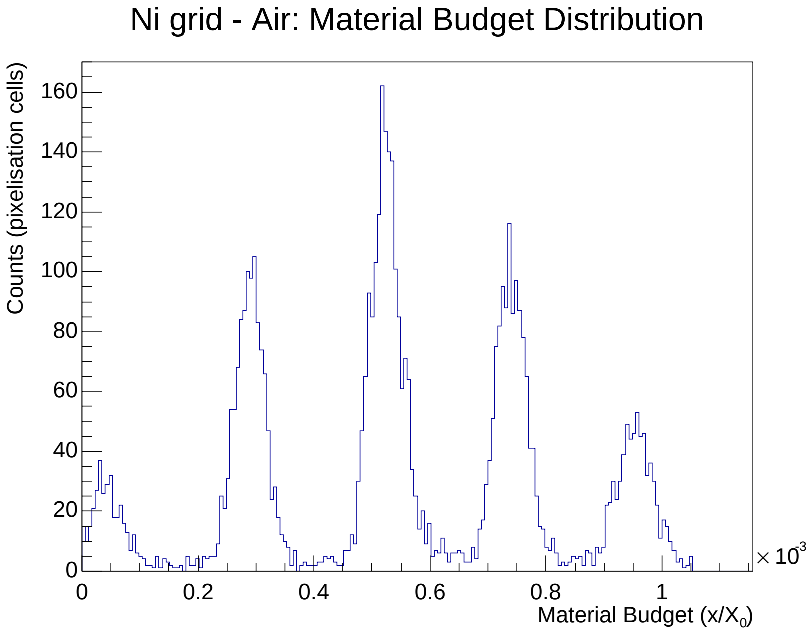

Figure 4.4 – Per-pixel material-budget distribution x/X_0 for the nickel grid target. Four relevant peaks are visible, corresponding to the discrete grid thickness levels.

Table 4.2 – Nickel grid: Gaussian fit parameters for the four populated thickness classes (l. to r.). The third column lists the *expected* material budget from the nominal thickness using Ni $X_0 = 1.48\,\text{cm}$. The relative widths σ_k/μ_k remain at the few-percent level, consistent with the uniform reference and ensuring clear separability of the classes.

Class	Nominal thickness (µm)	$(x/X_0)_{\text{exp}}$ $(\times 10^{-4})$	μ_k $(x/X_0 \times 10^{-4})$	σ_k $(x/X_0 \times 10^{-5})$	σ_k/μ_k (%)
1	3	2.03	2.90	2.6	8.79
2	6	4.05	5.23	3.0	5.74
3	9	6.08	7.45	2.7	3.62
4	12	8.11	9.60	3.5	3.65

Aluminium grid

For the aluminium grid, the histogram shows multiple bands of increasing material budget, reflecting the grouping of nominal path lengths in the structure as can be seen in Figure 4.5. Within the higher-budget bands, sub-peaks could be assigned to individual thickness classes, whereas in the lowest-budget band the sub-peaks merged into a single broader structure. This behaviour is consistent with the smaller absolute x/X_0 values and the finer step sizes in means of material budget between adjacent classes in aluminium compared to nickel, which reduce $\Delta\mu$ relative to the (statistical precision-dominated) per-peak σ and lead to partial overlap.

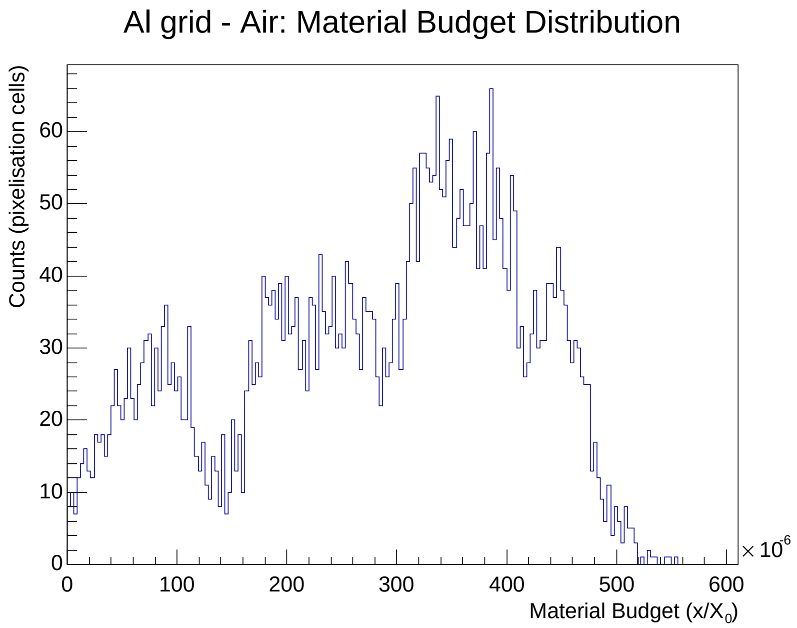

Figure 4.5 – Distribution of per-pixel material budget values x/X_0 for the Aluminium grid target. The spectrum is clearly multimodal, organising into different clusters. Sub-structure is visible in the higher-budget clusters, while the lowest-budget cluster appears as a single broad peak.

Convolution: uniform & grid

Conceptually, a complex target consisting of a uniform foil over (or under) a periodic grid produces a material–budget spectrum that is well described as the convolution of the uniform-foil response with the discrete Gaussian levels of the grid. Writing the uniform response as a single Gaussian $G(x;\mu_{\mathrm{uni}},\sigma_{\mathrm{uni}})$ and the grid as $\sum_k w_k\, G(x;\mu_k,\sigma_{\mathrm{grid},k})$, the expected distribution is

$$P_{\mathrm{conv}}(x) = G(x;\mu_{\mathrm{uni}},\sigma_{\mathrm{uni}}) \,*\, \sum_k w_k\, G(x;\mu_k,\sigma_{\mathrm{grid},k}) = \sum_k w_k\, G\big(x;\, \mu_{\mathrm{uni}}+\mu_k,\ \sigma_{\mathrm{conv},k}\big),$$

with $\sigma^2_{\mathrm{conv},k} = \sigma^2_{\mathrm{uni}} + \sigma^2_{\mathrm{grid},k}$ (independent Gaussian broadening). Thus the peaks remain spaced by the grid steps but are broadened by the quadrature sum of the uniform and grid widths. It is convenient to parameterise the uniform layer by its relative width $\varepsilon \equiv \sigma_{\mathrm{uni}}/\mu_{\mathrm{uni}}$ and to define the control ratio

$$r \;\equiv\; \frac{t_{\mathrm{uni}}}{\Delta t_{\mathrm{grid}}}\,,$$

which sets the separation-to-width ratio $\Delta\mu/\sigma_{\mathrm{conv},k}$ and therefore the degree of peak resolvability. Increasing r broadens and eventually merges the initially discrete sub-peaks into bands and, in the thick-baseline limit, into a single broad mode. This progression is illustrated by toy-model convolutions for nickel (constant step $\Delta t_{\mathrm{grid}} = 3\,\mu\mathrm{m}$) in Fig. 4.6 and for aluminium (mixed steps 4.5 µm and 1 µm) in Fig. 4.7.
A complication is that the grid step size is constant for the nickel grids (3 µm), whereas the aluminium grids contain both 4.5 µm and 1 µm increments. Consequently, peak separability varies across Al regions: larger steps remain discernible longer under convolution than 1 µm steps, so the 1 µm sub-series merges earlier, producing bands sooner (cf. Fig. 4.7).

Ni grid + uniform Ni 31 µm

Adding a thin uniform nickel layer left the discrete grid regions cleanly resolvable (Fig. 4.8). Each region was shifted by the uniform offset while its width increased only mildly, consistent with $\sigma^2_{\mathrm{conv},k} = \sigma^2_{\mathrm{uni}} + \sigma^2_{\mathrm{grid},k}$. Because the uniform contribution was small compared with the grid step sizes (here $r \approx 10.33$ for $\Delta t_{\mathrm{grid}} = 3\,\mu\mathrm{m}$), the separations $\Delta\mu$ between neighbouring regions remained much larger than the convolved widths, $\Delta\mu \gg \sigma_{\mathrm{conv},k}$, and the spectrum preserved a clearly multi-peaked structure.

Al grid + uniform Al 60 µm and 100 µm

For aluminium, a moderately thick uniform layer broadened the response sufficiently that the characteristic bands of the grid remained visible, while most intra-band sub-peaks became only weakly indicated (Fig. 4.9 and Fig. 4.10). This followed from the reduced step separations within bands—particularly those arising from 1 µm increments, so that

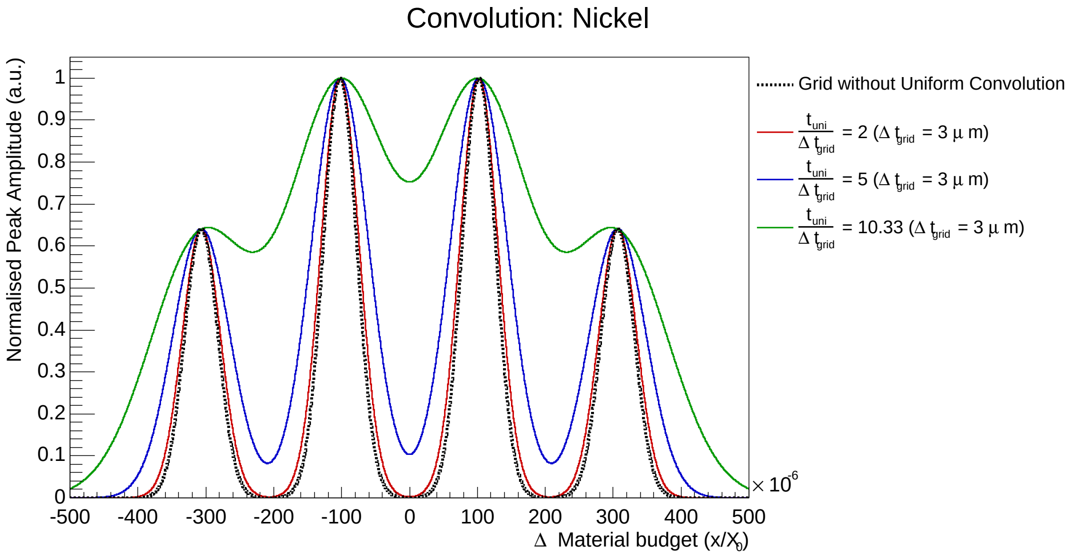

Figure 4.6 – Nickel grid: toy-model convolution showing the transition from resolved sub-peaks to bands as the ratio $t_{\text{uni}}/\Delta t_{\text{grid}}$ increases. Axes: $\Delta(x/X_0)$ relative to the central peak and normalised amplitude.

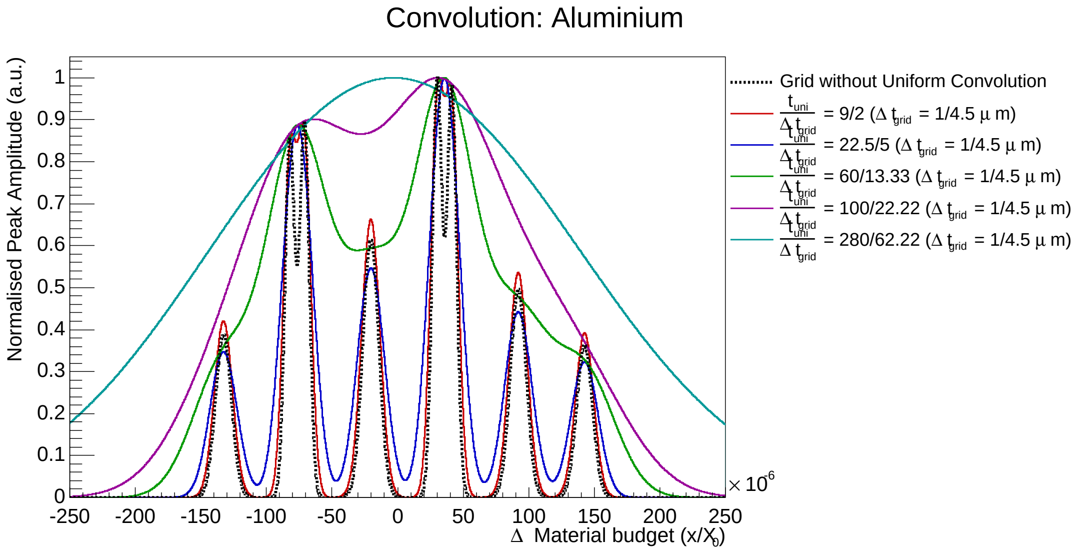

Figure 4.7 – Aluminium grid: toy-model convolution for several $t_{\text{uni}}/\Delta t_{\text{grid}}$ values mirroring measured configurations. Mixed steps (4.5 µm, 1 µm) cause earlier band merging for the 1 µm series. Axes: $\Delta(x/X_0)$ and normalised amplitude.

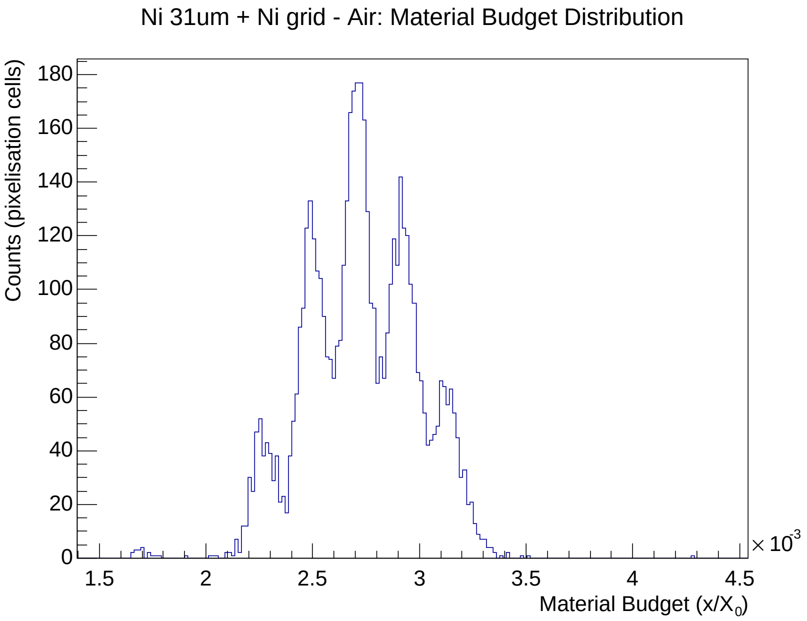

Figure 4.8 – Nickel grid with an additional 31 µm uniform Ni layer: per-pixel x/X_0 histogram with all grid-induced peaks still distinctly separated.

$\Delta\mu$ approached a few times the convolved width. Larger 4.5 µm steps still left hints of sub-structure in the higher-budget bands, but the overall appearance was that of smeared bands rather than fully resolved peaks.

Al grid + uniform Al 280 µm

With a thick aluminium baseline the convolution was dominated by the uniform component (Fig. 4.11). The increased σ_{uni} reduced the separation-to-width ratio within bands ($\Delta\mu/\sigma_{\text{conv},k} \lesssim 3$), so that even the band structure smeared into a single broad mode and discrete grid levels were no longer discernible. Quantitatively, the relevant control parameter was $r = t_{\text{uni}}/\Delta t_{\text{grid}}$. For the aluminium grid,

$$\frac{t_{\text{uni}}}{\Delta t_{\text{grid}}} \simeq \frac{280\,\mu\text{m}}{1\,\mu\text{m}} \approx 280 \quad (\text{fine steps})\,, \qquad \frac{280\,\mu\text{m}}{4.5\,\mu\text{m}} \approx 62 \quad (\text{coarser steps})\,. \tag{4.12}$$

With a baseline relative width $\varepsilon \equiv \sigma_{\text{uni}}/\mu_{\text{uni}}$ in the few-percent range, the modal separability scales as

$$\frac{\Delta\mu}{\sigma_{\text{conv}}} \approx \frac{\Delta t_{\text{grid}}/t_{\text{uni}}}{\varepsilon}\,, \tag{4.13}$$

yielding $\Delta\mu/\sigma_{\text{conv}} \sim 0.1$–$0.5$ for the values above, which is well below the $\sim$ 3–4 range required for practical resolvability of distinct modes in this work, as discussed below. The loss of visible sub-structure was therefore expected: in this regime only the overall mean material budget remained informative, while fine structure was lost.

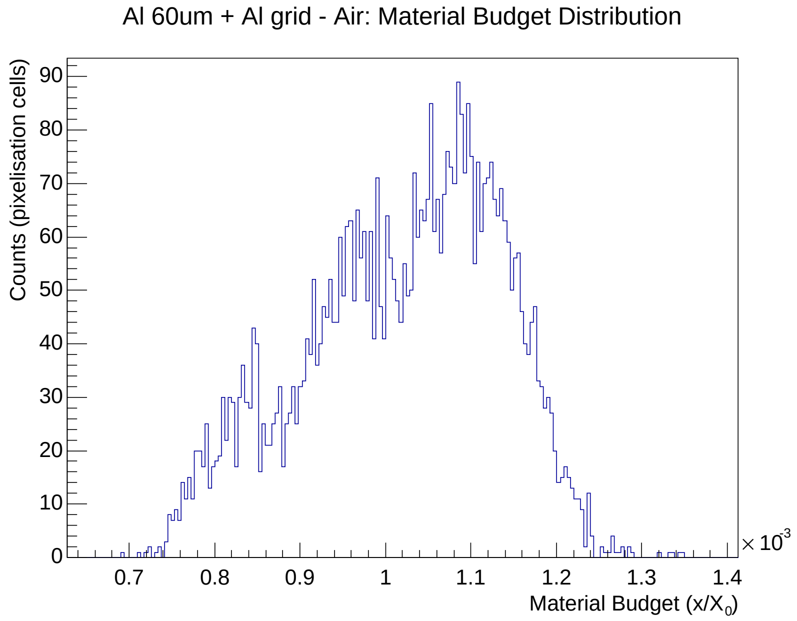

Figure 4.9 – Aluminium grid with an additional 60 µm uniform Al layer: band structure persists, but intra-band sub-peaks largely merge under convolution.

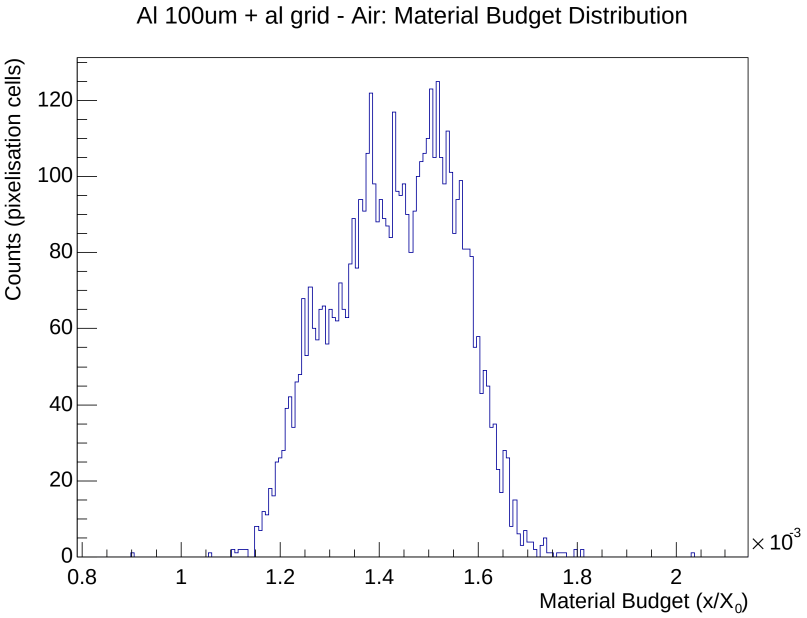

Figure 4.10 – Aluminium grid with a 100 µm uniform Al layer: bands smear further and individual grid levels are no longer resolvable.

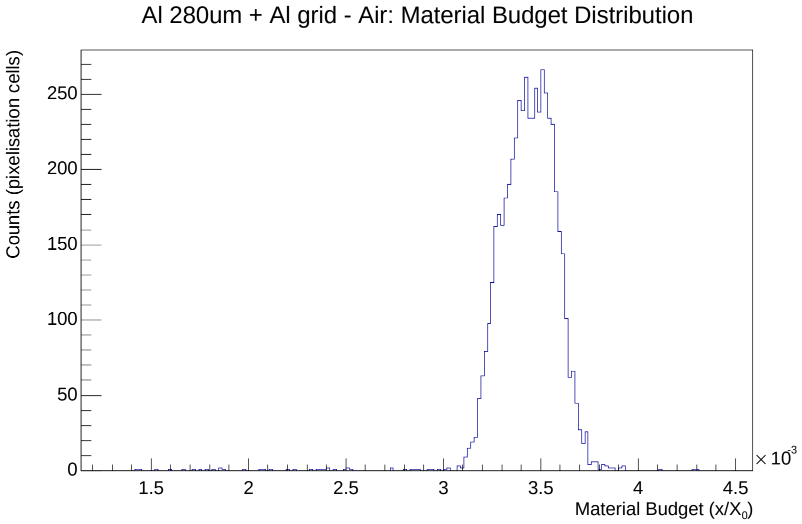

Figure 4.11 – Aluminium grid with a 280 µm uniform Al layer (thick-baseline regime): spectrum collapses into a single broad peak.

Resolution Criterion and Empirical Estimate

The resolvability of neighbouring modes is governed by the ratio of their mean separation to the common convolved width:

$$\text{peaks resolvable if} \qquad \frac{\Delta\mu}{\sigma_{\text{conv}}} \gtrsim 3\text{–}4\,, \qquad \text{with } \sigma_{\text{conv}} \simeq \sigma_{\text{uni}}\,. \tag{4.14}$$

To first order $\mu \propto x$, and with $\sigma_{\text{uni}} = \varepsilon\,\mu_{\text{uni}}$ (baseline relative width ε at the few-percent level), this leads to the thickness-based requirement

$$\frac{t_{\text{uni}}}{\Delta t_{\text{grid}}} \lesssim \frac{1}{\epsilon} \times \frac{1}{3\text{–}4} \approx 10\text{–}20\,, \tag{4.15}$$

that is, clear separation persists while the uniform thickness does not exceed $\mathcal{O}(10\text{–}20)$ times a single grid step. This bound accounts for the observations: configurations with small absolute material budgets and step sizes of similar order (Ni grid and Al grid with thin uniform layers) satisfy it and exhibit resolved peaks, whereas small differences superimposed on thick baselines (e.g. Al grid on $\gtrsim$ 100 µm Al) violate it, leading to merging into bands and, ultimately, a single broad mode.

4.4 Material Budget Determination at Small x/X_0

In this section the absolute scale of the reconstructed material budget in the low–thickness regime was established and validated. For each calibration target i, two dimensionless quantities were defined: a *theoretical* expectation x_i^{th}/X_0 computed from the nominal thickness x_i and the tabulated radiation length $X_{0,i}$, and a *measured* value x_i^{meas}/X_0 obtained by inverting the Highland relation using the data–driven momentum calibration from Sec. 4.2. Throughout, x^{th}/X_0 used manufacturer specifications for x^{th} supplemented by quoted tolerances, and X_0 values from standard tables (PDG). Very thin effective budgets were realised both with single ultra–thin foils and by subtracting distributions of materials with closely matched thickness, after the global subtraction and scaling correction established earlier to control instrumental and background contributions.

Uncertainty model Unless stated otherwise, each point carried (i) a horizontal uncertainty given by propagating the foil thickness tolerance $\sigma_{x^{\mathrm{th}}}$ into $\sigma_{x^{\mathrm{th}}/X_0} = \sigma_{x^{\mathrm{th}}}/X_0$, (ii) a vertical statistical uncertainty from the scatter of the per–pixel Gaussian widths θ_0 within the selected ROI propagated through the inversion of Highland and (iii) a small model component from the propagation of the weak $\ln(x/X_0)$ dependence. The latter enters in the same way as in the momentum–inversion error budget of Sec. 4.2 (Eq. 4.1–4.2). Thin nickel foils carried the largest *relative* $\sigma_{x^{\mathrm{th}}/X_0}$ and therefore show visibly larger horizontal error bars than aluminium and silicon.

Expectation If the reconstruction is unbiased, one expects $x^{\mathrm{meas}}/X_0 = x^{\mathrm{th}}/X_0$, i.e. the points should lie on the line of unity in a $y = x$ comparison.

Physics cross–check on $\boldsymbol{\theta_0(x/X_0)}$ Before confronting measured and theoretical budgets directly, the multiple–scattering shape was verified by fitting a two–parameter Highland form to the dependence of the corrected widths on material budget,

$$\theta_0 \;=\; S_2 \, \frac{1}{p_{\mathrm{true}}} \sqrt{\frac{x}{X_0}} \, \left[1 + \varepsilon \, \ln(x/X_0)\right] , \tag{4.16}$$

with S_2 and ε free and p_{true} fixed to the data–driven value above. The resulting curve described aluminium, nickel, and silicon with a single universal trend once expressed in x/X_0, as expected for multiple Coulomb scattering. Residual material–to–material spreads were consistent with thickness tolerances. The lowest budgets tended to sit mildly above the nominal curve, whereas toward larger x/X_0 the fit converged to the canonical Highland prediction. This behaviour is characteristic of a small additive resolution floor (discussed below) rather than a breakdown of the scattering model. An explicit overlay of θ_0 vs. x/X_0 in Figure 4.12 shows the Highland prediction and the best fit.

Using p_{true} tracks the data better than a curve computed at the nominal 2.4 GeV/c, thereby validating the momentum used in the inversion to x/X_0. In contrast, a fit restricted to the very thinnest budgets returned $S_2 \simeq 10.3\,\mathrm{MeV}$ and $\varepsilon \simeq -0.003$, these numbers are not robust and reflect parameter degeneracy over a limited lever arm rather than a physical deviation from the canonical Highland parameters. Taken together with Sec. 4.2, where using the nominal 2.4 GeV/c would have biased x/X_0 high by nearly 19 %, this cross–check favours the data–driven momentum scale adopted throughout.

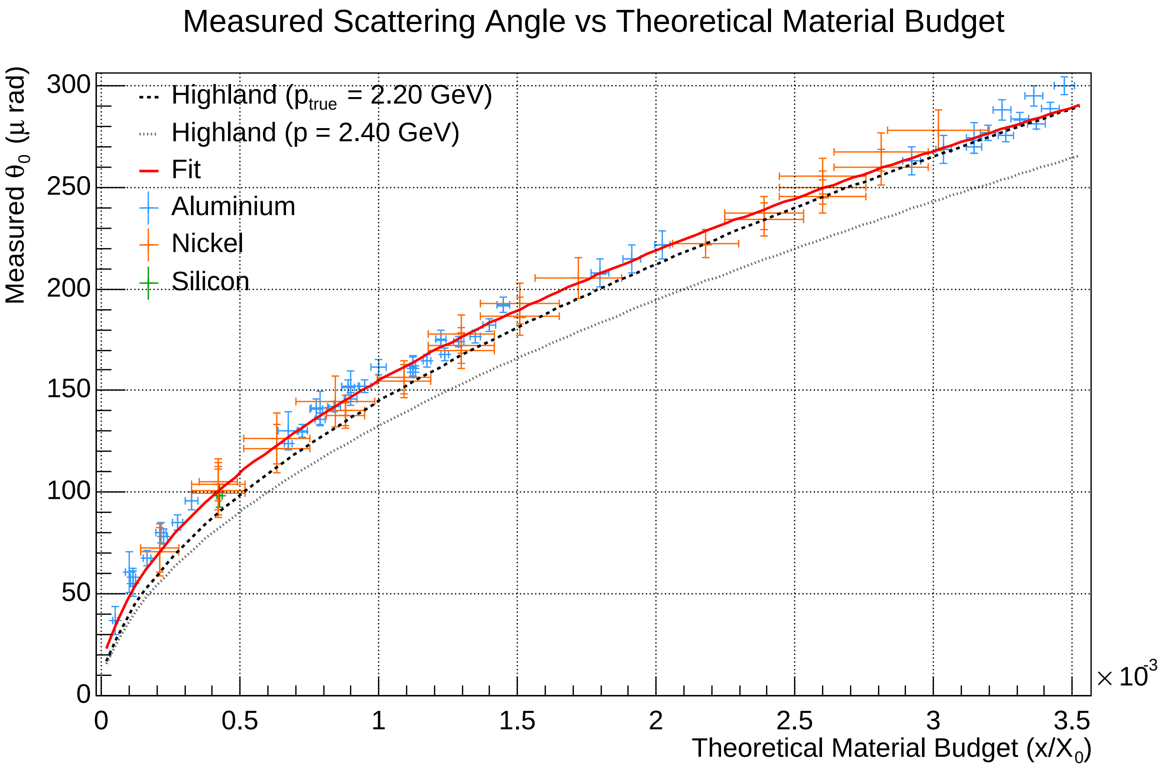

Figure 4.12 – Highland validation: corrected scattering angle θ_0 vs. material budget x/X_0. The dashed black curve shows the Highland prediction at the data–driven momentum $p_{\text{true}} = 2.20\,\text{GeV/c}$, the dotted grey curve shows the prediction at $2.40\,\text{GeV/c}$. A data–driven two–parameter Highland fit (free S_2, ε) evaluated at p_{true} is overlaid to guide the eye. Materials align on a common curve in x/X_0. The data favour the data–driven momentum over the nominal $2.4\,\text{GeV/c}$.

Comparison with theory with $\chi^2/\mathrm{ndof} = 123/73 = 1.68$, i.e. acceptable but mildly tensioned. The points in Figure 4.13 lie slightly above the unity line at the smallest budgets, while the agreement tightens with increasing x/X_0. The leverage is dominated by the thinnest Ni points, which also carry the largest relative horizontal uncertainties.

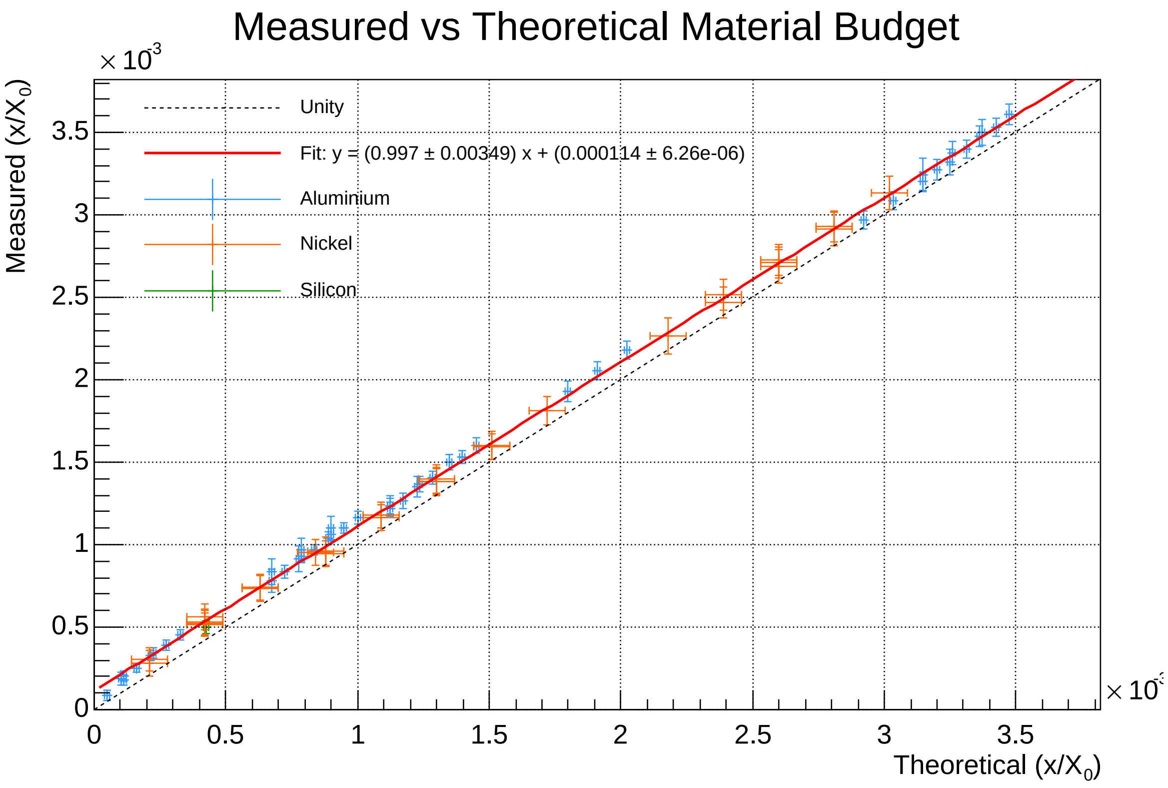

Figure 4.13 – Measured vs. theoretical material budget for all calibration targets. The unity line ($y = x$) and the best–fit straight line $y = c_1 x + c_0$ are overlaid. Fit: $c_1 = 0.997 \pm 0.003$, $c_0 = (1.14 \pm 0.06) \times 10^{-4}$, $\chi^2/\mathrm{ndf} = 123/73 = 1.68$. Points sit slightly above unity, but agreement improves with x/X_0.

These parameters have a compact interpretation:

A positive intercept $c_0 > 0$ indicates a small additive residual in x/X_0, as expected from an un-subtracted constant term in θ_0^2 after the reference subtraction/scaling. If a small constant width remains, $\theta^2_{\mathrm{meas}} = \theta^2_{\mathrm{true}} + \delta\theta^2$, then $x^{\mathrm{meas}}/X_0 = x^{\mathrm{true}}/X_0 + K$, $\delta\theta^2 \equiv x^{\mathrm{true}}/X_0 + c_0$, hence a constant offset most visible at low budgets.

A slope $c_1 < 1$ at the sub–percent level points to a minute global scale effect, compatible with a per–mille residual in momentum/normalisation or from using an effective ε over a finite x/X_0 range. This is within the systematic envelope of the method and consistent with Sec. 4.2.

For intuition, the intercept $c_0 = (1.14 \pm 0.06) \times 10^{-4}$ ($\simeq$ $\theta_0 = (43.5 \pm 1.3)\,\mu\mathrm{rad}$ for the scattering angle) corresponds to about $(10.1 \pm 0.5)\,\mu\mathrm{m}$ of aluminium or approximately $(1.62 \pm 0.09)\,\mu\mathrm{m}$ of nickel in thickness, and its *relative* impact scales as $c_0/(x/X_0)$. It is $\sim 11.4\,\%$ at $x/X_0 \simeq 1 \times 10^{-3}$ but already only $\sim 3.8\,\%$ at 3×10^{-3} — precisely the trend seen in Fig. 4.13.

Sensitivity threshold at low budgets The additive floor c_0 defines a regime where the reconstruction remains *resolving* (differences are still visible) but ceases to be *absolutely accurate*. A convenient criterion follows from requiring that the relative bias $c_0/(x/X_0)$ stay below a tolerance δ:

$$x_{\min}(\delta) \;\approx\; \frac{c_0}{\delta}\,. \tag{4.17}$$

With $c_0 = (1.14 \pm 0.06) \times 10^{-4}$, a 5 % criterion yields $x_{\min} \simeq 2.3 \times 10^{-3}$, and, accordingly, absolute material budgets were reported only for regions with $x/X_0 \gtrsim 2.3 \times 10^{-3}$, while for thinner regions ($x/X_0 < 2.3 \times 10^{-3}$) only relative contrasts (e.g. $\Delta\mu/\mu$) were quoted. Below this threshold the offset dominates the fractional error, whereas above it the agreement becomes acceptable, consistent with the direct θ_0–vs–x/X_0 validation (see Fig. 4.12) which showed the expected low–x curvature. The general resolvability arguments derived from the convolution study (Sec. 4.3) likewise emphasised that separating fine structures on top of thick baselines becomes progressively harder once the baseline fluctuations (few–percent) approach the step size, with a practical separability condition of $\Delta\mu/\sigma_{\mathrm{conv}} \gtrsim 3$ to 4.

Taken together, these observations place the small–x/X_0 behaviour into a consistent picture. The two–parameter Highland fit to θ_0 as a function of x/X_0 reproduced the expected shape of multiple scattering and clearly preferred the data–driven momentum scale. This momentum was therefore adopted for all subsequent inversions to x/X_0. In the direct comparison of x^{meas}/X_0 with x^{th}/X_0, the fit returned $c_1 = 0.997 \pm 0.003$ and $c_0 = (1.14 \pm 0.06) \times 10^{-4}$ with a reduced χ^2 of 1.68. The positive intercept points to a small additive floor that remains after subtraction and scaling, while the sub–percent slope indicates only a very mild global scale effect. Overall, the Highland description reproduces the data within about 5 % down to $x/X_0 \simeq 2.3 \times 10^{-3}$. At smaller budgets the reconstruction remains resolving, but absolute accuracy is limited by the additive term c_0 rather than by any intrinsic breakdown of the scattering model.

This behaviour follows naturally from the corrections established in Sec. 4.2. These steps were essential to reduce systematic distortions to the point where only the small additive floor remains visible.

Chapter 5

Summary, conclusion and outlook

This thesis established and validated a data-driven workflow for material–budget imaging (MBI) in the regime relevant to modern inner trackers. Using a six-plane ALPIDE telescope operated at the DESY electron test beam, multiple Coulomb scattering was exploited as a quantitative proxy for x/X_0, and measured projected-angle widths were converted to material budget through an explicitly inverted Highland/Lynch–Dahl relation (see Eq. 2.10). The analysis chain was tailored to thin, quasi-uniform scatterers and small local inhomogeneities, and was designed to be robust against the core–tail interplay characteristic of multiple scattering at the ultra-thin limit.

At the detector and beamline level, the telescope geometry and operational choices were optimised for angular resolution and stability: the closest approach of $\sim 15\,\mathrm{mm}$ to the target and a 20 mm outer-pair spacing defined a lever arm of 140 mm, which was found to have the best angular resolution at 2.4 GeV nominal momentum (see Fig. 3.2). Data were recorded over a week-long campaign with uniform aluminium, nickel, and silicon foils and two composite 3×3 grids providing discrete thickness classes (see Tab. 3.1), predominantly at the nominal 2.4 GeV setting and, in addition, at 1 GeV for a momentum-scaling cross-check (the latter not analysed in detail here). Straight-line fits upstream and downstream of the target yielded per-event scattering angles at the interaction plane, and these angles were accumulated in a pixelised map for spatially resolved width estimation.

Two ingredients were key to the stability of the per-pixel estimator. First, pixelisation at $250\,\mu\mathrm{m}\times250\,\mu\mathrm{m}$ balanced statistical precision against granularity: it produced maps with $\mathcal{O}(5\times10^3)$ cells while controlling inter-cell fluctuations at fixed statistics (see Fig. 3.9). Second, widths were extracted from Gaussian fits restricted to a central quantile. A 98 % two-sided cut was adopted after verifying that 95 % and 98 % gave indistinguishable behaviour in the pixel-size optimisation (see Fig. A.1). This restricted-interval strategy intentionally de-emphasised tails, trading a small loss of efficiency for greatly improved robustness and reproducibility of the core-width estimator.

Scattering-width maps reproduced the expected geometric patterns at the mm scale. For the nickel grid, the stepwise thickness pattern was mirrored by areas in the θ-width map.

Ancillary structures (adhesive seams, carrier-card material) also appeared where present, confirming the sensitivity of the method to local services (see Fig. 4.1). Converting widths to x/X_0 and compressing by material revealed the core statistical structure of the measurement. Uniform foils yielded single, relatively narrow peaks in the per-pixel x/X_0 distributions with means μ in the range 10^{-4}–10^{-3} and Gaussian widths $\sigma_{\text{uni}} \sim 10^{-5}$, corresponding to few-percent relative fluctuations σ_{uni}/μ. Expressed in radiation-length units, this behaviour was material-agnostic (Al/Ni/Si), underscoring that residual spread at fixed statistics is dominated by counting precision rather than by material-specific effects (see Fig. 4.3).

A central lesson of the work concerns the treatment of non-target broadening. Instrumental resolution, residual air, telescope material upstream/downstream of the target, and small reconstruction biases contribute to the measured core widths and cannot be removed by a naive quadratic subtraction alone. Following prior evidence, a multiplicative scaling was applied after reference subtraction, with the scale determined from calibration measurements. This brought different thickness settings onto a common normalisation and prevented thickness-dependent biases in the thin limit.

Momentum control was the other decisive ingredient for quantitative x/X_0. A data-driven momentum calibration using uniform foils yielded an effective momentum of $p_{\text{true}} = (2.20 \pm 0.10)\,\text{GeV/c}$ at the telescope for the nominal 2.4 GeV setting (see Fig. 4.2,4.4), i.e. a global scale factor $c_p = 0.916$ (see Eq. 4.5). Without this correction, the inversion $\theta_0 \to x/X_0$ would overestimate x/X_0 by nearly 20 % at first order. In addition, a mild transverse gradient common to all targets was cancelled scaling with a determined momentum calibration, removing position-dependent momentum effects. Together, the global p_{true} calibration and the gradient cancellation stabilised the conversion to x/X_0 across the field of view. The auxiliary 1 GeV sample, not pursued further in this work, offers a direct $\theta_0 \propto 1/p$ scaling cross-check.

With the momentum description and background normalisation in place, the inversion remained stable for relative contrasts. Thin-to-thick ladders in Al/Ni/Si aligned on a universal $\theta_0(\sqrt{x/X_0})$ trend once expressed in radiation-length units, validating the material scaling expected for multiple scattering. Toward the smallest budgets, points sat mildly above the nominal curve, consistent with a residual additive floor rather than a failure of the scattering model (see Eq. 4.16,4.12). In the absolute comparison x_{meas}/X_0 versus x_{th}/X_0 constructed from nominal areal thicknesses and PDG X_0 values, agreement was acceptable overall, with leverage dominated by the thinnestl points. In this regime, the log term in Highland, the non-Gaussian single-scatter tails, and the instrumental floor collectively limit the precision of an absolute inversion.

The resulting performance picture is therefore twofold. On the one hand, the chain clearly *resolved* permille-level material budgets and local differences at the 10^{-4}–10^{-3} level, sufficient to map uniform foils and to separate discrete classes in grid targets. A simple resolvability criterion, $\Delta\mu/\sigma_{\text{conv}} \gtrsim 3$–$4$ with $\sigma_{\text{conv}} \approx \sigma_{\text{uni}}$ (see Eq. 4.14-4.15), captured when multiple modes separated cleanly and when thin contrasts riding on thick baselines merged into a single band, matching the empirical behaviour across configurations (see Fig. 4.11,4.10). On the other hand, at the ultra-thin end the absolute scale became model-limited: even large statistics could not fully overcome the combined impact of momentum-scale uncertainty, residual non-target scattering after subtraction, and small model mis-specification at the core–tail transition. Statistical errors were subleading in the main operating region due to the chosen cell size and event counts.

These findings corroborate earlier MBI studies that highlighted the decisive role of proper resolution/background handling in the small-x/X_0 regime and, in a practical sense, integrate sample-tolerance knowledge into the uncertainty model. They also quantify how much momentum control is needed to claim absolute scales below 10^{-3} with confidence, providing a concrete target for future campaigns.

The path forward follows directly from the limitations identified. First, further reduce momentum uncertainty and position dependence, including an explicit re-analysis of the 1 GeV dataset as a scaling cross-check, with central-interval choices re-tuned if needed. Second, suppress non-target scattering by migrating to in-vacuum or reduced-air paths, minimising upstream/downstream material, and stabilising the post-subtraction scaling with per-run controls and empty/thin-silicon references. Third, make the inversion more model-aware in the ultra-thin regime: benchmark the adopted central-interval Gaussian against robust M-estimators and explicit Molière-core plus Rutherford-tail templates, and explore iterative inversions that include a core–tail prior. Fourth, revisit reconstruction under controlled conditions: comparing the present two-segment model to General Broken Lines with identical selections and synthetic injections would quantify any achievable gain in per-pixel θ resolution without overfitting. Fifth, tighten the absolute scale by cross-metrology where feasible, profilometry or targeted X-ray/CT on regions that light up in $\sigma(\theta)$, and by densifying the calibration ladder with certified ultra-thin Al/Ni and wafer-scale Si in the 10^{-4}–10^{-3} band.

Finally, the technique is ready to be applied to assembled ITS3 prototypes. The present pipeline already provides the relative contrast and spatial reach needed to map seams, overlaps, and services on bent-MAPS cylinders. The recommended upgrades concentrate on raising absolute precision to turn those relative maps into quantitative x/X_0 statements in the ultra-thin domain. With automated quality control, versioned configurations, and open analysis notes encompassing calibration targets, air runs, and uncertainty budgets, future campaigns will be directly comparable across sites and iterations. In summary, the work demonstrated precise sensitivity with quantified uncertainties and identified the specific, achievable steps required to lift the model-limited floor at the thinnest budgets—closing the loop from *resolving* to *determining* material budgets where it matters most for low-p_T tracking and impact-parameter resolution.

Appendix A

Additional Information

A.1 Gaussian Fit Values

For completeness, the full set of Gaussian fit parameters is compiled here. Table A.1 lists all uniform foils, while Table A.2 contains the results for all grid regions (Ni and Al). These extend the representative values shown in the main text.

Table A.1 – Complete set of Gaussian fit parameters for all uniform foils. The third column lists the *expected* material budget from nominal thickness using material-specific radiation lengths (Al: $X_0 = 8.897\,\mathrm{cm}$, Ni: $X_0 = 1.48\,\mathrm{cm}$, Si: $X_0 = 9.36\,\mathrm{cm}$). Relative widths σ/μ remain consistently in the few-percent range, confirming that fluctuations are statistics-dominated.

Material	Nominal thickness (µm)	$(x/X_0)_{\mathrm{exp}}$ ($\times 10^{-4}$)	μ_k ($x_k/X_{0,k} \times 10^{-4}$)	σ_k ($x_k/X_{0,k} \times 10^{-5}$)	σ_k/μ_k (%)
Ni	6	4.05	5.6	3	5.4
Al	10	1.12	1.8	3	16.7
Ni	12.5	8.45	9.4	4	4.3
Si	40	4.27	5.0	4	8.0
Al	60	6.74	8.3	8	9.6
Al	70	7.87	9.7	7	7.2
Al	80	8.99	11.0	7	6.4
Al	100	11.24	12.0	4	3.3
Al	160	17.98	19.0	6	3.2
Al	170	19.11	21.0	6	2.9
Al	180	20.23	22.0	6	2.7
Al	260	29.22	30.0	5	1.7
Al	270	30.35	31.0	5	1.6
Al	280	31.47	32.0	5	1.6

Table A.2 – Complete set of Gaussian fit parameters for all shown grid targets (Ni and Al). The third column lists the *expected* material budget from the nominal total thickness using material-specific radiation lengths (Al: $X_0 = 8.897\,\mathrm{cm}$, Ni: $X_0 = 1.48\,\mathrm{cm}$). Relative widths σ_k/μ_k remain at the few-percent level, comparable to the uniform baseline and ensuring clear separability of thickness classes where applicable.

Sample	Nominal thickness (µm)	$(x/X_0)_{\mathrm{exp}}$ $(\times 10^{-4})$	μ_k $(x_k/X_{0,k} \times 10^{-4})$	σ_k $(x_k/X_{0,k} \times 10^{-5})$	σ_k/μ_k (%)
Ni grid	3.0	2.03	2.9	2.6	9.0
Ni grid	6.0	4.05	5.2	3.0	5.8
Ni grid	9.0	6.08	7.5	2.7	3.6
Ni grid	12.0	8.11	9.6	3.5	3.7
Al grid	4.5	0.51	0.9	3.0	33.3
Al grid	9.0	1.01	1.9	3.7	19.5
Al grid	10.0	1.12	2.0	3.2	16.0
Al grid	14.5	1.63	2.5	2.1	8.4
Al grid	19.0	2.14	3.3	2.9	8.8
Al grid	20.0	2.25	3.4	2.8	8.2
Al grid	24.5	2.75	3.9	2.5	6.4
Al grid	29.0	3.26	4.5	2.6	5.8
Ni 31 µm + grid	34.0	22.97	25	6.7	2.7
Ni 31 µm + grid	37.0	25	27	6.3	2.3
Ni 31 µm + grid	40.0	27.03	29	6.9	2.4
Ni 31 µm + grid	43.0	29.05	31	7.7	2.5
Al 60 µm + grid	64.5	7.25	8.3	4.0	4.8
Al 60 µm + grid	69.0	7.76	9.1	8.1	8.9
Al 60 µm + grid	70.0	7.87	9.3	3.9	4.2
Al 60 µm + grid	74.5	8.37	9.7	3.2	3.3
Al 60 µm + grid	79.0	8.88	10	3.6	3.6
Al 60 µm + grid	80.0	8.99	11	3.5	3.2
Al 60 µm + grid	84.5	9.50	11	2.9	2.6
Al 60 µm + grid	89.0	10.00	12	3.6	3.0
Al 100 µm + grid	104.5	11.75	13	4.3	3.3
Al 100 µm + grid	109.0	12.25	13	5.9	4.5
Al 100 µm + grid	110.0	12.36	14	4.3	3.1
Al 100 µm + grid	114.5	12.87	14	4.1	2.9
Al 100 µm + grid	119.0	13.38	15	5.0	3.3
Al 100 µm + grid	120.0	13.49	15	4.7	3.1
Al 100 µm + grid	124.5	13.99	15	3.7	2.5
Al 100 µm + grid	129.0	14.50	16	4.7	2.9
Al 280 µm + grid	284.5	31.98	33	6.5	2.0
Al 280 µm + grid	289.0	32.48	33	7.4	2.2
Al 280 µm + grid	290.0	32.60	34	7.3	2.2
Al 280 µm + grid	294.5	33.10	34	5.7	1.7
Al 280 µm + grid	299.0	33.61	35	6.0	1.7
Al 280 µm + grid	300.0	33.72	35	7.5	2.1
Al 280 µm + grid	304.5	34.23	35	5.3	1.5
Al 280 µm + grid	309.0	34.73	36	6.4	1.8

A.2 Central-interval robustness

For robustness, the mean–fluctuation analysis of Sec. 3.6 was repeated with a two-sided 95 % central interval, keeping the full pipeline identical to the main study (same bin-size scan as in Fig. 3.9, von–Neumann neighbours, and edge-cell normalisation by $|\mathcal{N}_4(c)|$). The resulting curve, shown in Fig. A.1, lies minimally below the 98 % result, but was neglected due to less data being included with this method. Therefore, the 98 % cut was retained in the main analysis for slightly higher statistical precision.

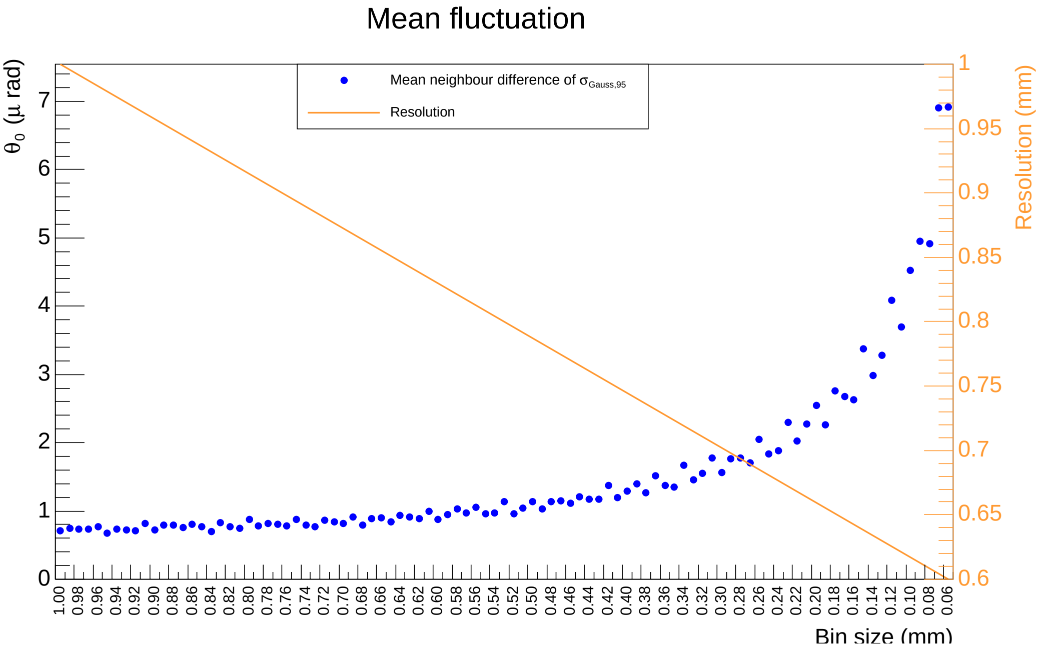

Figure A.1 – Mean fluctuation $\overline{\Delta\sigma}$ between von–Neumann neighbours vs. bin size using a 95 % central interval. The curve overlaps the 98 % result of Fig. 3.9 within plotting resolution.

Nomenclature

ALICE	A Large Ion Collider Experiment
ALPIDE	Alice Pixel Detector
CERN	Conseil européen pour la recherche nucléaire
CMOS	Complementary Metal-Oxide-Semiconductor
DCal	Di-jet Calorimeter
DESY	Deutsches Elektronen-Synchroton
EMCal	Electro Magnetic Calorimeter
FIT	Fast Interaction Trigger
GBL	General Broken Line
HMPID	High Momentum Particle Identification Detector
IP	Interaction Point
ITS2	Inner Tracking System 2
ITS3	Inner Tracking System 3
LHC	Large Hadron Collider
MAPS	Monolithic Active Pixel Sensors
MBI	Material Budget Imaging
MCS	Multiple Coulomb Scattering
MIP	Minimum Ionising Particle
ndof	Number of Degrees Of Freedom
PHOS	Photon Spectrometer
PID	Particle Identification
QCD	Quantum Chromodynamics
QGP	Quark-Gluon Plasma
RMS	Root Mean Square
ROI	Region Of Interest
TOF	Time-Of-Flight
TPC	Time Projection Chamber
TRD	Transition Radiation Detector

Bibliography

[1] ALICE Collaboration. Alice figure repository, 2025. Official repository of ALICE figures. (cit. on pp. 1 and 7)

[2] Wit Busza, Krishna Rajagopal, and Wilke Van Der Schee. Heavy ion collisions: The big picture and the big questions. *Annual Review of Nuclear and Particle Science*, 68:339–376, 2018. (cit. on pp. 1, 3, and 4)

[3] Peter Braun-Munzinger and Johanna Stachel. The quest for the quark–gluon plasma. *Nature*, 448(7151):302–309, 2007. (cit. on pp. 1, 3, and 4)

[4] ALICE Collaboration. The alice experiment at the cern lhc. *JINST*, 3:S08002, 2008. (cit. on pp. 1 and 4)

[5] Conseil Européen pour la Recherche Nucléaire. Official website. (cit. on p. 1)

[6] Ewa Łopienska. The cern accelerator complex, layout in 2022. `https://cds.cern.ch/record/2800984`, 2022. CERN Document Server — General Photo. (cit. on p. 2)

[7] The ALICE Collaboration. The ALICE experiment: A journey through QCD. *The European Physical Journal C*, 84(8):813, 2024. (cit. on p. 3)

[8] R. L. Workman et al. (Particle Data Group). Review of particle physics. *Progress of Theoretical and Experimental Physics*, 2022(8):083C01, 2022. (cit. on pp. 4 and 5)

[9] P.G. Kuijer. Commissioning and prospects for early physics with alice. *Nuclear Physics A*, 830(1–4):81c–88c, November 2009. (cit. on p. 5)

[10] The ALICE Collaboration. The alice experiment at the cern lhc. *Journal of Instrumentation*, 3(08):S08002, aug 2008. (cit. on p. 5)

[11] J. Alme et al. The alice tpc, a large 3-dimensional tracking device with fast readout for ultra-high multiplicity events. *Nucl. Instrum. Methods A*, 622(1):316–367, 2010. (cit. on pp. 5 and 8)

[12] B. Abelev et al. (ALICE Collaboration). Technical design report for the upgrade of the alice inner tracking system. *J. Phys. G: Nucl. Part. Phys.*, 41:087002, 2014. (cit. on pp. 5, 6, and 7)

[13] M. Mager. Alpide, the monolithic active pixel sensor for the alice its upgrade. *Nucl. Instrum. Methods A*, 824:434–438, 2016. (cit. on pp. 5 and 7)

[14] F. Reidt. Upgrade of the alice its detector. *Nucl. Instrum. Methods A*, 1032:166632, 2022. (cit. on pp. 5 and 6)

[15] ALICE Collaboration. The upgrade of the alice tpc with gems and continuous readout. *JINST*, 16(03):P03022, 2021. (cit. on pp. 5, 7, and 8)

[16] ALICE Collaboration. Real-time data processing in the alice high level trigger at the lhc. *Computer Physics Communications*, 242:25–48, 2019. (cit. on pp. 5 and 7)

[17] The ALICE Collaboration. Technical design report for the upgrade of the alice inner tracking system. *Journal of Physics G: Nuclear and Particle Physics*, 41(8):087002, jul 2014. (cit. on pp. 6 and 18)

[18] The ALICE Collaboration. Technical design report for the alice inner tracking system 3 - its3 ; a bent wafer-scale monolithic pixel detector. Technical report, CERN, Geneva, 2024. Co-project Manager: Magnus Mager, magnus.mager@cern.ch. (cit. on pp. 7, 9, and 10)

[19] LHC Experiments Committee. Alice time-of-flight system (tof): Technical design report. Technical report, CERN, Geneva, 2000. (cit. on p. 8)

[20] ALICE Collaboration. The alice transition radiation detector: Construction, operation, and performance. *Nucl. Instrum. Methods A*, 881:88–127, 2018. (cit. on p. 8)

[21] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments: A How-To Approach.* Springer, Berlin, Heidelberg, 2 edition, 1994. (cit. on pp. 12 and 15)

[22] C. Grupen and B. Shwartz. *Particle Detectors.* Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology. Cambridge University Press, Cambridge, 2 edition, 2008. (cit. on p. 12)

[23] H. Bethe. Zur theorie des durchgangs schneller korpuskularstrahlen durch materie. *Annalen der Physik*, 397(3):325–400, 1930. (cit. on p. 12)

[24] F. Bloch. Zur bremsung rasch bewegter teilchen beim durchgang durch materie. *Annalen der Physik*, 408:285–320, 1933. (cit. on p. 12)

[25] F. Bloch. Bremsvermögen von atomen mit mehreren elektronen. *Zeitschrift für Physik*, 81:363–376, 1933. (cit. on p. 12)

[26] R. M. Sternheimer and R. F. Peierls. General expression for the density effect for the ionization loss of charged particles. *Phys. Rev. B*, 3:3681–3692, 1971. (cit. on p. 13)

[27] R. M. Sternheimer, M. J. Berger, and S. M. Seltzer. Density effect for the ionization loss of charged particles in various substances. *Atomic Data and Nuclear Data Tables*, 30(2):261–271, 1984. (cit. on p. 13)

[28] H. Spieler. *Semiconductor Detector Systems.* Oxford University Press, Oxford, 2005. (cit. on p. 13)

[29] G. F. Knoll. *Radiation Detection and Measurement.* John Wiley & Sons, Hoboken, NJ, 4 edition, 2010. (cit. on p. 13)

[30] S. Navas et al. Review of particle physics. *Phys. Rev. D*, 110(3):030001, 2024. (cit. on pp. 13, 14, 16, and 21)

[31] P. V. Vavilov. Ionization losses of high-energy heavy particles. *Sov. Phys. JETP*, 5:749–751, 1957. (cit. on p. 14)

[32] H. Bichsel. Straggling in thin silicon detectors. *Rev. Mod. Phys.*, 60:663–699, 1988. (cit. on p. 14)

[33] Y.-S. Tsai. Pair production and bremsstrahlung of charged leptons. *Rev. Mod. Phys.*, 46:815–851, 1974. (cit. on p. 14)

[34] H. W. Koch and J. W. Motz. Bremsstrahlung cross-section formulas and related data. *Rev. Mod. Phys.*, 31:920–955, 1959. (cit. on p. 14)

[35] Particle Data Group. Atomic and nuclear properties of materials for particle physics (2024 update). `https://pdg.lbl.gov/2024/AtomicNuclearProperties/`, 2024. Accessed 2025-09-01. (cit. on pp. 15 and 60)

[36] E. Rutherford. The scattering of α and β particles by matter and the structure of the atom. *Philosophical Magazine*, 21:669–688, 1911. (cit. on p. 16)

[37] G. Molière. Theorie der streuung schneller geladener teilchen i. einzelstreuung am abgeschirmten coulomb-feld. *Z. Naturforsch. A*, 2:133–145, 1947. (cit. on p. 16)

[38] G. Molière. Theorie der streuung schneller geladener teilchen ii. mehrfach- und vielfachstreuung. *Z. Naturforsch. A*, 3:78–97, 1948. (cit. on p. 16)

[39] H. A. Bethe. Molière's theory of multiple scattering. *Phys. Rev.*, 89:1256–1266, 1953. (cit. on p. 16)

[40] Particle Data Group and Workman. Review of particle physics. *Progress of Theoretical and Experimental Physics*, 2022(8):083C01, 08 2022. (cit. on p. 16)

[41] V. L. Highland. Some practical remarks on multiple scattering. *Nucl. Instrum. Methods*, 129(2):497–499, 1975. (cit. on pp. 16, 17, 28, and 29)

[42] G. R. Lynch and O. I. Dahl. Approximations to multiple coulomb scattering. *Nucl. Instrum. Methods Phys. Res. B*, 58(1):6–10, 1991. (cit. on pp. 16, 17, 26, 28, and 29)

[43] D. E. Groom, N. V. Mokhov, and S. I. Striganov. Muon stopping power and range tables 10 mev–100 tev. *Atomic Data and Nuclear Data Tables*, 78(2):183–356, 2001. (cit. on p. 17)

[44] ALICE ITS ALPIDE development team. *ALPIDE Operations Manual*, 2016. (cit. on p. 18)

[45] M Šuljić. Alpide: the monolithic active pixel sensor for the alice its upgrade. *Nuovo Cimento C*, 41(1-2):91, 2018. (cit. on p. 18)

[46] Tower Semiconductor Ltd. Official website. (cit. on p. 18)

[47] Simon Groß-Bölting. Material budget imaging for mockup samples of the alice its3 with an electron beam at energies of 2.4 gev and 4.8 gev. Master's thesis, Heidelberg University, 2024. (cit. on pp. 19 and 30)

[48] H. Jansen, S. Spannagel, J. Behr, et al. Performance of the eudet-type beam telescopes. *EPJ Techniques and Instrumentation*, 3(7), 2016. (cit. on p. 19)

[49] Roland Diener, Hendrik Jansen, et al. The desy ii test beam facility. *Nucl. Instrum. Methods Phys. Res. A*, 922:265–286, 2019. (cit. on p. 22)

[50] Deutsches Elektron-Synochroton. Official website. (cit. on p. 22)

[51] D. Dannheim et al. Corryvreckan: a modular 4d track reconstruction and analysis software for test beam data. *Journal of Instrumentation*, 16(03):P03008, mar 2021. (cit. on p. 23)

[52] Claus Kleinwort. General broken lines as advanced track fitting method. *Nucl. Instrum. Methods Phys. Res. A*, 673:107–110, 2012. (cit. on p. 24)

[53] Mihail-Bogdan Blidaru. *Bent monolithic active pixel sensors and material budget imaging studies towards the ALICE ITS3 upgrade.* PhD thesis, Heidelberg University, 2024. (cit. on pp. 26 and 30)

[54] G. Cowan. Statistics. *Review of Particle Physics*, 2024. Particle Data Group, Phys. Rev. D 110, 030001 (2024); chapter 40. (cit. on pp. 26 and 27)

[55] N. Agafonova, A. Aleksandrov, O. Altinok, and others (OPERA Collaboration). Momentum measurement by the multiple coulomb scattering method in the opera lead-emulsion target. *New Journal of Physics*, 14:013026, 2012. (cit. on p. 31)

[56] K. Kodama, N. Saoulidou, G. Tzanakos, et al. Momentum measurement of secondary particle by multiple coulomb scattering with emulsion cloud chamber in donut experiment. *Nucl. Instrum. Methods Phys. Res. A*, 574(1):192–198, 2007. (cit. on p. 31)

List of Figures

1.1 ALICE detector overview . . . 1

1.2 CERN accelerator complex . . . 2

1.3 Space–time journey of a heavy-ion collision . . . 3

1.4 ALICE η-acceptance and transverse barrel layout. . . . 5

1.5 ITS2 layout and segmentation . . . 6

1.6 Impact-parameter resolution vs. p_{T} . . . 7

1.7 ITS2 stave material budget . . . 7

1.8 ITS3 inner barrel concept . . . 9

1.9 Simulated average ITS3 material budget . . . 10

2.1 Mean mass stopping power vs $\beta\gamma$ . . . 13

2.2 Simulated energy–deposit distribution . . . 14

2.3 MCS geometry . . . 16

3.1 Optimal telescope geometry and resolution . . . 19

3.2 Adopted telescope geometry . . . 20

3.3 Illustration of a complex target . . . 20

3.4 Material budget of prepared targets . . . 22

3.5 ALPIDE telescope setup in the test beam . . . 23

3.6 Alignment comparison between measurements . . . 24

3.7 Schematic passage of a particle through the telescope . . . 25

3.8 Local mean fluctuation . . . 26

3.9 Cell size dependence of mean fluctuation . . . 27

3.10 Pixelised 2D histogram for Ni 6 µm . . . 27

4.1 2D map of scattering angle widths for the Nickel grid target . . . 29

4.2 Momentum distributions at 2.4 GeV/c nominal beam energy . . . 33

4.3 Material-budget histogram: 100 µm uniform Al layer . . . 36

4.4 Material-budget histogram: nickel grid . . . 37

4.5 Material-budget histogram: aluminium grid . . . 38

4.6 Nickel grid: convolution toy model . . . 40

4.7 Aluminium grid: convolution toy model . . . 40

4.8 Material-budget histogram: nickel grid + 31 µm Ni 41
4.9 Material-budget histogram: aluminium grid + 60 µm Al 42
4.10 Material-budget histogram: aluminium grid + 100 µm Al 42
4.11 Material-budget histogram: aluminium grid + 280 µm Al 43
4.12 θ_0 vs. x/X_0 Highland validation 45
4.13 Measured vs. theoretical material budget 46

A.1 95 % central interval check 53

List of Tables

2.1 Variables to calculate the mean energy loss with the Bethe–Bloch formula. 13
2.2 Radiation lengths used in this thesis (values consistent with PDG material tables [35]). 15

3.1 Overview of target properties 21

4.1 Uniform aluminium foils: Gaussian fit parameters 37
4.2 Nickel grid: Gaussian fit parameters 38

A.1 Complete Gaussian fit parameters for uniform foils 51
A.2 Complete Gaussian fit parameters for mentioned complex targets 52

Declaration:

I hereby declare that I have written this thesis independently and have used no sources or aids other than those indicated.

Erklärung:

Ich versichere, dass ich diese Arbeit selbstständig verfasst habe und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, 17.09.2025