

Sensor Co-design for *smartpixels*

Danush Shekar^a, Ben Weiss^b, Morris Swartz^c, Corrinne Mills^a, Jennet Dickinson^b, Lindsey Gray^d, David Jiang^e, Mohammad Abrar Wadud^a, Daniel Abadjiev^f, Anthony Badea^f, Douglas Berry^d, Alec Cauper^b, Arghya Ranjan Das^g, Giuseppe Di Guglielmo^{d,j}, Karri Folan DiPetrillo^f, Farah Fahim^{d,j}, Rachel Kovach Fuentes^f, Abhijith Gandrakota^d, James Hirschauer^d, Eliza Howard^f, Shiqi Kuang^g, Carissa Kumar^f, Ron Lipton^d, Mia Liu^g, Petar Maksimovic^c, Nick Manganelli^h, Mark S Neubauer^e, Aidan Nicholas^f, Emily Pan^f, Benjamin Parpillon^d, Jannicke Pearkes^h, Gauri Pradhan^d, Shruti R Kulkarniⁱ, Ricardo Silvestre^h, Chinar Syal^d, Nhan Tran^{d,j}, Amit Trivedi^a, Keith Ulmer^h, Manuel Blanco Valentin^j, Dahai Wen^c, Jieun Yoo^a, Eric You^f, Aaron Youngⁱ

^aUniversity of Illinois at Chicago Chicago IL 60607 USA

^bCornell University Ithaca NY 14853 USA

^cJohns Hopkins University Baltimore MD 21218 USA

^dFermi National Accelerator Laboratory Batavia IL 60510 USA

^eUniversity of Illinois Urbana-Champaign Champaign IL 61801 USA

^fThe University of Chicago Chicago IL 60637 USA

^gPurdue University West Lafayette IN 47907 USA

^hUniversity of Colorado Boulder Boulder CO 80309 USA

ⁱOak Ridge National Laboratory Oak Ridge TN 37831 USA

^jNorthwestern University Evanston IL 60208 USA

Abstract

Pixel tracking detectors at upcoming collider experiments will see unprecedented charged-particle densities. Real-time data reduction on the detector will enable higher granularity and faster readout, possibly enabling the use of the pixel detector in the first level of the trigger for a hadron collider. This data reduction can be accomplished with a neural network (NN) in the readout chip bonded with the sensor that recognizes and rejects tracks with low transverse momentum (p_T) based on the geometrical shape of the charge deposition (“cluster”). To design a viable detector for deployment at an experiment, the dependence of the NN as a function of the sensor geometry, external magnetic field, and irradiation must be understood. In this paper, we present first studies of the efficiency and data reduction for planar pixel sensors exploring these parameters. A smaller sensor pitch in the bending direction improves the p_T discrimination, but a larger pitch can be partially compensated with detector depth. An external magnetic field parallel to the sensor plane induces Lorentz drift of the electron-hole pairs produced by the charged particle, broadening the cluster and improving the network performance. The absence of the external field diminishes the background rejection compared to the baseline by $\mathcal{O}(10\%)$. Any accumulated radiation damage also changes the cluster shape, reducing the signal efficiency compared to the baseline by $\sim 30 - 60\%$, but nearly all of the performance can be recovered through retraining of the network and updating the weights. Finally, the impact of noise was investigated, and retraining the network on noise-injected datasets was found to maintain performance within 6% of the baseline network trained and evaluated on noiseless data.

Contents

1	Introduction	2
2	Simulated datasets	2

Email address: dsheka3@uic.edu (Danush Shekar)

3	Filtering algorithm for pixel clusters	5
3.1	Performance metrics	5
4	Studies on performance of filtering algorithm	6
4.1	Sensor geometry	6
4.2	Lorentz drift	7
4.3	Irradiation	7
4.4	Noise	8
5	Conclusions and outlook	9

1. Introduction

Next-generation collider detectors will see unprecedented collision rates and particle densities, putting severe pressure on detector readout bandwidth. The track density is highest near to the collision point, where, conventionally, highly granular silicon “pixel” tracking detectors are placed. That increased density motivates ever smaller pitch in the detector design, which further exacerbates the bandwidth needs, even with zero-suppressed readout. The *smartpixels* project is co-designing machine learning (ML) algorithms and application-specific integrated circuits (ASICs) that can be deployed on a pixel detector to perform inference and data reduction on-detector, reducing the bandwidth requirements. As an example, the *smartpixels* technology introduced to the innermost sensors in the CMS experiment would help unlock new physics capabilities by facilitating the potential inclusion of inner-tracker data to the Level 1 trigger.

The *smartpixels* project established the feasibility of filtering the tracks of incident particles by transverse momentum through a neural network implementable on an ASIC [1]. The filtering algorithm is a classifier that uses the shape of the cluster of adjacent pixels that see a signal from the particle crossing the detector. An ASIC with this algorithm implemented has been produced and characterized in the lab [2]. On the algorithm side, further work explores the possibility of regressing the incident track parameters and their uncertainties using the cluster shape [3].

A necessary step to the deployment of *smartpixels* in a collider detector is the construction of a physical prototype bonded to a sensor so that its performance can be measured in beam, including after irradiation. Another necessary step is understanding how the physical design choices affect the performance of the algorithm(s) implemented. This paper therefore extends that program of co-design to include the specifications and performance of the sensor, including the pixel pitch, active sensor thickness, and radiation hardness, using the classifier network architecture developed in Ref. [1] as the benchmark.

The datasets to test and train the filtering algorithm are defined in Section 2, and Section 3 reviews the structure and training of the filter algorithm. Then, we show the filter algorithm’s performance for different pixel pitch in two dimensions and sensor depth (Section 4.1), for sensors oriented in the magnetic field so that charges have no Lorentz drift (Section 4.2), for sensors that have been irradiated to fluences typical of the High-Luminosity Large Hadron Collider (HL-LHC) (Section 4.3), and for sensors with noise-injected outputs (Section 4.4).

2. Simulated datasets

The PixelAV [4] program is used to generate the datasets required for training and validating the filtering network. PixelAV provides a detailed model of silicon pixel detectors, including drift, diffusion, and the Hall effect. It is also capable of simulating radiation effects such as charge trapping and the signals induced from trapped charges. Electric and weighting field maps of the sensor design were simulated using Silvaco TCAD [5] and sent as input to PixelAV. The version of PixelAV used to generate the datasets in this paper also provides granular information on the time evolution of signals in the sensor, and the datasets here record 20 “slices” in steps of 200 ps. PixelAV calculates the signals induced by incident charged pions in an array of 21×13 pixels based on the trajectory of the incident particles. The incident position (x, y) of the charged particle on the sensor mid-plane is uniformly distributed across the center 3×3 pixels of the 21×13

array. The collection of datasets considered for training and validation has been expanded compared to the ones generated for the previous study [1].

The baseline performance of the filtering algorithm [1] was evaluated using simulated datasets from a futuristic silicon planar sensor (pitch = $50 \times 12.5 \mu m^2$ along $x \times y$ and thickness = $100 \mu m$ along z) assumed to be in the barrel region of the CMS [6] experiment at CERN. The incident particle makes an angle with the sensor array in the $x-z$ and $y-z$ planes denoted by α and β , respectively. The trajectories of particles emanating from the interaction point at the origin curve due to the 3.8 T external magnetic field along x before encountering the sensor array placed along a cylinder at approximately 30 mm from the origin. Due to that curvature, the shape of the cluster depends on the particle's transverse momentum, which is highly correlated with β . In particular, the cluster shape along the y -direction is sensitive to the incident angle β as well as the p_T of the particle. The size of the pixel array by choice is large enough to fully encompass a cluster for these incident positions and angles (α, β).

There are typically regions in collider experiments, like the endcap pixel sensors in the CMS experiment, where sensors are oriented such that the external magnetic field is parallel to the sensor electric field. In such scenarios, there is an absence of Lorentz drift (deflection of charge carriers by the magnetic field) that otherwise improves impact position reconstruction via interpolation. Thus, the co-design program was expanded to include simulations of sensors in endcap regions using the kinematic properties taken from fitted tracks in CMS 13 TeV collision data[7]. While the orientation of the sensors in the endcaps would be different with respect to the barrel region, α and β are still defined as the angles made with the sensor length-thickness ($x-z$) and sensor width-thickness ($y-z$) planes, respectively. This definition is maintained across both detector regions by transforming the coordinate systems as illustrated in Figure 1.




Figure 1: A schematic diagram of two arbitrary charged particles traversing sensors in the barrel and endcap regions. The diagram highlights the definition of α and β angles in both these regions for arbitrary incident particles, along with the magnetic field direction. The change in orientation of the sensors in the endcap regions with respect to sensors in the barrel region requires defining a coordinate system for each regions to ensure the definition of incident angles remain the same.

The distribution of induced charges in the pixel array forms the bulk of input information sent to the classification network, and thus, the granularity of the pixel pitch and the sensor thickness influences the network's performance. To understand this impact, five sensor geometries that vary in pitch along y as listed in Table 1 were investigated. To offer comparison to an existing detector, we include a sensor matching the inner tracker sensor geometry in the CMS Phase 2 upgrade in the survey, along with a $100 \mu m$ thick variant of the same.

Datasets with different p_T distributions were also created for each geometry. One, called the “flat- p_T ” dataset, contains 2 million events with their p_T distributed uniformly in $[0.03, 5]$ GeV. This dataset is used for training in order to improve training performance across a large momentum range and avoid bias. The other, called the “physical- p_T ” dataset, follows a falling momentum distribution typical of what would be seen in the 13 TeV pp collision data at CMS. These datasets, also containing 2 million events each, were used to test the performance of the network with a realistic distribution. The distributions of p_T , α , and β

Sl. No.	Sensor name	Sensor geometry			Bias voltage [V]
		Length [μm]	Width [μm]	Thickness [μm]	
1	S1	50	10	100	100
2	S2	50	12.5	100	100
3	S3	50	15	100	100
4	S4	50	20	100	100
5	S5	50	25	100	100
6	S6	100	25	100	100
7	S7	100	25	150	175

Table 1: Specifications of the sensor geometries surveyed.

from all events in both physical- and flat- p_T datasets are shown in Figure 2.




Figure 2: Distributions of p_T (left), α (center), and β (right) across the various datasets produced. A symmetric distribution of α values in the center plot is not seen in the endcap regions because only a section of the endcap disk was simulated, ignoring the section diametrically opposite due to radial-symmetry considerations.

The change in cluster size for different sensor geometries is visible in Fig. 3, which shows the distribution of cluster sizes along the y -direction for the physical- p_T datasets of all sensor geometries studied.

Pixel detectors are typically installed closest to the interaction point, and face the brunt of radiation damage. This alters sensor characteristics, so an assessment of the evolution of the classifier’s performance with radiation damage is crucial. Therefore, datasets at three different irradiation levels 0, 370, and 1100 fb^{-1} (equivalent to 0, 3.3×10^{15} , and 1.0×10^{16} 1 MeV equivalent neutrons per cm^2 , respectively) were produced for the baseline and CMS Phase-2 sensor geometries [8]. A per-pixel charge threshold of 400 e- was applied to remove induced signals from trapped charges. The datasets from the un-irradiated sensors however had no charge threshold as noise hasn’t been introduced in these studies yet.

Table 2 summarizes all datasets produced to obtain the results presented in this paper. The bias voltages mentioned in this table ensure similar electric field values in the bulk region for the un-irradiated sensors. The bias was increased in case of radiation damaged sensors to help retain optimal electric field magnitudes.

Sl. No.	Detector region	Sensors simulated	Angle between $\mathbf{E}$ and $\mathbf{B}$	Irradiation level [fb^{-1}]
1	Barrel	S1 - S7	90°	0
2	Endcap	S1 - S7	180°	0
3	Barrel	S1, S7	90°	370, 1100

Table 2: Categorization of the simulations datasets based on their use-cases in studies presented in this paper.




Figure 3: Distribution of cluster size along the Y-direction for sensor geometries in the barrel region, using the physical- p_T dataset. The distributions in left and right correspond to sensors with varying y -pitch and thickness, respectively.

3. Filtering algorithm for pixel clusters

Three classification models were evaluated in the baseline study [1] and the model chosen for deployment on-sensor is the one used for the studies presented in this paper. The classifier is trained using the “cluster y -profile” of each event, which contains more information than the size alone. The cluster y -profile is calculated for each event as the sensor output at 4000 ps summed over pixel rows (x) to project the integrated charge of the cluster onto the y -axis. Additionally, y_0 , the azimuthal position of the particle’s incident position on the sensor array in the global detector coordinates, is passed as input. The inter-dependence of the cluster y -profile, y_0 , and p_T was examined in [1]. This network, therefore taking in 14 input values (y_0 and 13 values from cluster y -profile), contains one dense hidden layer with 128 neurons and 1920 parameters, followed by one dense output layer with 3 neurons and 387 parameters. A softmax activation is used to generate classification probabilities between 0 and 1, and each event is assigned the classification label corresponding to the highest probability.

The three output categories are positively charged and $p_T < 200$ MeV; negatively charged and $p_T < 200$ MeV; and $p_T > 200$ MeV, both positively and negatively charged. The rationale to have two background categories comes from differences in cluster profiles for oppositely charged particle tracks that curve in opposite directions in the magnetic field. The outputs are also categorized on the p_T value of 200 MeV which is a hyper-parameter that will be referred to as “ p_T boundary”. This value was chosen previously as it was found to result in a flat signal efficiency for tracks with $p_T > 2$ GeV, which is a useful range for physics studies [1].

Neural network trainings were run for 200 epochs and stopped early if the loss function showed no improvement after 20 epochs. A batch size of 1024 was used in all models. The Adam optimizer with a learning rate of 0.001 was used in conjunction with the Keras Sparse Categorical Cross entropy loss function in all models.

3.1. Performance metrics

The model performance is evaluated using three metrics:

- Signal efficiency: fraction of clusters with $p_T > 2$ GeV that are classified as high p_T .
- Background rejection: fraction of clusters with $p_T < 2$ GeV that are classified as low p_T .
- Data-reduction: fraction of events classified as low p_T irrespective of true class

Each metric is calculated as a fraction of the total number of clusters in the dataset. We calculate a statistical uncertainty based on the dataset size using the binomial uncertainty formula. All training and evaluation cycles were performed 10 times with different random seeds to obtain stochastic errors from the ML algorithm. The final uncertainty is defined as the addition of the statistical and stochastic errors in quadrature.

4. Studies on performance of filtering algorithm

The following sub-sections describe the findings of the studies done in terms of the performance metrics highlighted in Section 3.

4.1. Sensor geometry

The impact of sensor geometry on model performance was studied using datasets from seven configurations, varying in both pitch and thickness. Fig. 4 contains results of the algorithm trained on the flat- p_T and evaluated on the physical- p_T datasets. The performance is plotted as a function of p_T boundary for four sensor geometries. We observe a performance drop across all p_T boundaries for increasing y -pitch and decreasing thickness.




Figure 4: Performance results of sensors in the barrel region for four sensor geometries as a function of p_T boundary. The legend describes the pixel dimensions along the x (length), y (width), and z (thickness) directions, respectively.

The hyperparameter optimization study carried out in prior work [1] for the $50 \times 12.5 \times 100 \mu\text{m}^3$ sensor geometry found the p_T boundary value = 0.2 GeV to be a nominal one. We observe the same value produces acceptable performance across other sensor geometries as well. Results shown in Fig. 5 correspond to model performance across all sensor geometries with the p_T boundary hyperparameter set to 0.2 GeV.




Figure 5: Performance results of sensors in barrel region for sensor geometries varying in pitch and thickness. The X-axis labels indicate unit pixel dimensions along the x (length), y (width), and z (thickness) directions, respectively. The solid, dashed, and dotted lines correspond to geometries with varying y -pitch, x -pitch, and thickness, respectively.

As seen in Figure 3, the discriminating power on cluster profile features is lost with increasing sensor y -pitch. In line with this observation, we observe all three performance metrics to degrade as sensor y -pitch increases, and improve for increases in sensor thickness. The performances of the $50 \times 25 \times 100$ and $100 \times 25 \times 100 \mu\text{m}^3$ sensors are comparable within error margins, since the model input does not include any x -direction-dependent information. Notably, the increase in thickness from $100 \times 25 \times 100 \mu\text{m}^3$ to $100 \times 25 \times 150 \mu\text{m}^3$ recovers some of the performance loss from the larger pixel pitch, though it does not perform as well as the $50 \times 12.5 \times 100 \mu\text{m}^3$ baseline.

4.2. Lorentz drift

Model performance was also evaluated in regions of the detector in which there is no Lorentz drift because the electric field in the sensor is parallel to the external magnetic field. This is typical of disk-type (endcap) geometries in conventional pixel detectors when planar-sensor modules are oriented orthogonal to the magnetic field. Figure 6 shows these results as a function of the p_T boundary for four sensor geometries.




Figure 6: Performance results of sensors in the endcap region for four sensor geometries as a function of p_T boundary. The legend describes the pixel dimensions along the x (length), y (width), and z (thickness) directions, respectively.

At the 0.2 GeV boundary, sensors without Lorentz drift show reduced performance compared to those with Lorentz drift. However, as the p_T boundary increases, signal efficiency improves while data reduction and background rejection decrease across all geometries.

4.3. Irradiation

Using the datasets produced for radiation damaged sensors for the $50 \times 12.5 \times 100 \mu\text{m}^3$ and Phase 2 sensor geometry, the performance of the classification network was examined along with the importance of having reprogrammable weights in the ASIC.

As shown in Figure 7, if the model weights remain fixed while radiation damage accumulates, classification accuracy deteriorates significantly - leading to high rejection rates even for relevant events. In contrast, retraining the models at each irradiation level (and subsequently updating ASIC weights) preserves model performance, demonstrating that the performance of the classifier can be mostly maintained even at fluences typical of the HL-LHC through simply updating the weights. For example, if retrained, the baseline sensor geometry's signal efficiency and data reduction remain within 7% and 2% of their original values, respectively, across 1100 fb^{-1} of radiation damage.




Figure 7: Performance results of the baseline and CMS Phase 2 sensors in the barrel region as a function of irradiation levels. "Original" means using the weights resulting from training on simulation of un-irradiated detectors. Results in this figure use a p_T boundary = 0.2 GeV with the networks trained and evaluated on datasets with a 400 electron per-pixel threshold to address leakage currents arising from trapped charges in irradiated sensors.

4.4. Noise

To estimate the effect of detector noise on the performance of the filtering algorithm, a noise model is implemented by applying Gaussian smearing to the simulated charge collected in each pixel. The width of this Gaussian depends on the baseline charge fluctuation Q_{base} , the charge mismatch Q_{th} , and the Equivalent Noise Charge (ENC):

$$\sigma_{\text{noise}} = \sqrt{Q_{\text{base}}^2 + Q_{\text{th}}^2 + \text{ENC}^2} \quad (1)$$

Q_{base} is the measure of the fluctuation in the pre-amplifier output baseline arising from event-by-event amplitude variation (due to variation in deposited charges by minimum-ionizing particles in Silicon) and variation in time between subsequent events. Q_{th} represents the non-uniformity in the charge threshold across a pixel matrix that can arise from effects like systematic offset or random process variations in the digitization stage. ENC refers to the amount of charge that, if deposited in the sensor, would produce a signal equal to the root-mean-square value of the noise of the system.

Simulations of the architecture and ROIC process used for implementation of the p_T filter [9] were designed in the Cadence Virtuoso flow and simulated with Spectre [10] in the baseline $50 \times 12.5 \times 100 \mu\text{m}^3$ sensor geometry. The resulting Q_{base} is found to be 12 electrons, and Q_{th} lies in the 50-60 electron range. The ENC varies with sensor geometry and capacitance, temperature, and radiation damage. Simulations performed at -40°C and sensor capacitance 30 fF yield an ENC of 55 electrons before radiation damage. For a sensor leakage current of 1 nA, the same range of capacitance yields ENC of 65 electrons at -40°C . This corresponds to the effect of leakage current compensation after irradiation to approximately 500 Mrad [11].

The performance of the p_T filtering algorithm is evaluated at four different noise levels up to 100 electrons, as shown in Figure 8. For each data point, noise sampled from a Gaussian distribution of $\mu = 0$ and $\sigma = \sigma_{\text{noise}}$ was injected into the simulation datasets for further training/evaluation. We observe that there is a considerable drop in performance with increasing levels of noise. However, upon retraining the network on the noise-injected simulation datasets, the performance of the network can be retained. For instance, at $\sigma_{\text{noise}} = 100$ electrons, the retrained model's signal efficiency and data reduction lie within 3% and 6% to the original model's results at $\sigma_{\text{noise}} = 0$ electrons, respectively.




Figure 8: Performance results of the baseline $50 \times 12.5 \times 100 \mu\text{m}^3$ sensor geometry as a function of noise at p_T boundary = 0.2 GeV. “Original” corresponds to the model weights obtained by training on simulation with no Gaussian noise applied.

We also evaluated the model performance on the radiation damaged sensor dataset (1100 fb^{-1}) injected with noise for the baseline ($50 \times 12.5 \times 100 \mu\text{m}^3$) and CMS Phase 2 ($100 \times 25 \times 150 \mu\text{m}^3$) sensor geometries. The ENC values calculated from simulation for the baseline and CMS Phase 2 sensor geometries were 65 and 105 electrons, respectively, which results in σ_{noise} of 90 and 125 electrons respectively. To assess model robustness under extreme operating conditions, we evaluated performance on datasets that include both noise and radiation damage effects. Retraining the model on such noise-injected, irradiated sensor datasets was found to substantially preserve performance. For the baseline sensor geometry, the retrained model's the signal efficiency and data reduction remained within 10% and 9% of the baseline model (trained on unirradiated, noiseless data). For the CMS Phase-2 geometry, these metrics were within 6% and 4%, respectively.

5. Conclusions and outlook

On-detector inference of incident particle characteristics and on-sensor data filtering for highly granular pixel detectors has been explored in previous work [1], but only for a single benchmark sensor geometry. In this paper, we have systematically evaluated the performance of the baseline filtering network across variations in sensor geometry, Lorentz drift, radiation damage, and noise. Our results highlight that smaller sensor pitch and increased thickness improve classification performance due to better resolved cluster features. The presence of Lorentz drift enhances the discriminative power of the classification algorithm, particularly at low transverse momentum thresholds, but tuning of the p_T boundary used for training labels can help recover lost performance in those conditions. Additionally, we demonstrated that noise and radiation damage severely degrades filtering performance. While retraining can compensate for noise effects, periodic retraining is required to address radiation damage, highlighting the importance of having reprogrammable logic in the ASICs.

These findings inform design considerations for future detectors aiming to implement intelligent, on-sensor data reduction. The proof-of-concept from prior work, augmented with co-design work presented in this paper, also provides valuable direction for future R&D under the *smartpixels* project. In particular, we are extending simulations to account for scenarios one can expect to occur in the experiment, studying a variety of machine learning algorithms, and prototyping and characterizing new ASICs developed for *smartpixels*. Continued efforts have the potential to reduce data throughput while maintaining high physics performance, which is essential for the success of next-generation collider experiments.

Acknowledgements

This work was completed using computing resources at the University of Chicago’s Research Computing Center and the Fermilab Elastic Analysis Facility. We thank Burt Holzman (Fermilab) for computing support. C M, D S, and J Y are supported by NSF award PHY-2208803 and, together with A T, by DOE Office of Science award DE-SC0023715 from Funding Opportunity Announcement for Artificial Intelligence Research for High Energy Physics, DE-FOA-0002705. M L and A R D are supported by the A3D3 HDR Institute through the NSF award PHY-2117997. M S N and D J are supported through the NSF cooperative agreement OAC-2117997 the DOE Office of Science award DE-SC0023365. M S and P K are supported by NSF award PHY-2310072. K F D and E H are supported by the NSF CAREER Program through award 2443370, and K F D is additionally supported by the Neubauer Family Assistant Professor Program. E Y is supported by the University of Chicago’s Quad Undergraduate Research Scholar program, and D A is supported by the University of Chicago’s Sachs Fellowship. AB is supported by the Schmidt Sciences Foundation.

References

- [1] J. Yoo, J. Dickinson, M. Swartz, G. Di Guglielmo, A. Bean, D. Berry, M. Blanco Valentin, K. DiPetrillo, F. Fahim, L. Gray, J. Hirschauer, S. R. Kulkarni, R. Lipton, P. Maksimovic, C. Mills, M. S. Neubauer, B. Parpillon, G. Pradhan, C. Syal, N. Tran, D. Wen, A. Young, Smart pixel sensors: towards on-sensor filtering of pixel clusters with deep learning, *Machine Learning: Science and Technology* 5 (3) (2024) 035047. doi:10.1088/2632-2153/ad6a00.
URL <https://dx.doi.org/10.1088/2632-2153/ad6a00>
- [2] A. Badea, A. Bean, D. Berry, J. Dickinson, K. DiPetrillo, F. Fahim, L. Gray, G. D. Guglielmo, D. Jiang, R. Kovach-Fuentes, P. Maksimovic, C. Mills, M. S. Neubauer, B. Parpillon, D. Shekar, M. Swartz, C. Syal, N. Tran, J. Yoo, Intelligent pixel detectors: Towards a radiation hard asic with on-chip machine learning in 28 nm cmos (2024). *arXiv:2410.02945*.
URL <https://arxiv.org/abs/2410.02945>

- [3] J. Dickinson, R. Kovach-Fuentes, L. Gray, M. Swartz, G. D. Guglielmo, A. Bean, D. Berry, M. B. Valentin, K. DiPetrillo, F. Fahim, J. Hirschauer, S. R. Kulkarni, R. Lipton, P. Maksimovic, C. Mills, M. S. Neubauer, B. Parpillon, G. Pradhan, C. Syal, N. Tran, D. Wen, J. Yoo, A. Young, Smartpixels: Towards on-sensor inference of charged particle track parameters and uncertainties (2023). **arXiv:** 2312.11676.
URL <https://arxiv.org/abs/2312.11676>
- [4] M. Swartz, A Detailed Simulation of the CMS Pixel Sensor, Tech. rep., CERN, Geneva (2002).
URL <https://cds.cern.ch/record/687440>
- [5] Silvaco TCAD.
URL <https://silvaco.com/tcad/>
- [6] The cms experiment at the cern lhc, Journal of Instrumentation 3 (08) (2008) S08004. doi:10.1088/1748-0221/3/08/S08004.
URL <https://dx.doi.org/10.1088/1748-0221/3/08/S08004>
- [7] D. Shekar, M. Swartz, J. Dickinson, M. A. Wadud, Smartpixels co-design datasets (Sep. 2025). doi: 10.5281/zenodo.17180303.
URL <https://zenodo.org/records/17180303>
- [8] A. Tumasyan, et al., The Phase-2 Upgrade of the CMS Tracker (6 2017). doi:10.17181/CERN.QZ28.FLHW.
- [9] B. Parpillon, C. Syal, J. Yoo, J. Dickinson, M. Swartz, G. D. Guglielmo, A. Bean, D. Berry, M. B. Valentin, K. DiPetrillo, A. Badea, L. Gray, P. Maksimovic, C. Mills, M. S. Neubauer, G. Pradhan, N. Tran, D. Wen, F. Fahim, Smart pixels: In-pixel ai for on-sensor data filtering (2024). **arXiv:** 2406.14860.
URL <https://arxiv.org/abs/2406.14860>
- [10] "cadence custom ic / analog / rf design", https://www.cadence.com/en_US/home/tools/custom-ic-analog-rf-design.html, accessed: 2025-09-04.
- [11] M. Garcia-Sciveres, RD53A Integrated Circuit Specifications, Tech. rep., CERN, Geneva (2015).
URL <https://cds.cern.ch/record/2113263>