

# Energy-Aware NECO for Single-Pass Pixel-wise Out-of-Distribution Detection in Semantic Segmentation

Boyuan Zhang<sup>1</sup>, Huanshan Huang<sup>1</sup>, Yifei Cao<sup>2,3\*</sup>

**Abstract**—Reliable semantic segmentation for mobile robots requires both accurate dense prediction and robust uncertainty estimation under distribution shift. In practice, however, strong uncertainty baselines such as Monte Carlo (MC) Dropout often require repeated stochastic forward passes and are difficult to deploy on resource-constrained edge platforms. We propose Energy-Aware NEural Collapse based Out-of-distribution detection (NECO), a single-pass pixel-wise out-of-distribution (OOD) score for semantic segmentation. The method is motivated by two observations. First, preliminary experiments on the Multiple Uncertainties for Autonomous Driving (MUAD) dataset showed that uncertainty-based methods are useful for dense OOD detection, but repeated inference is expensive. Second, our earlier study on Neural Collapse (NC) and NECO suggested that purely geometric projection scores can miss simple anomalies whose features remain embedded in the dominant in-distribution subspace. We therefore combine a centered NECO-style geometric ratio computed from decoder features with a logit-based Energy score. Both components are standardized using statistics fitted on a pure in-distribution validation split and then fused by a convex combination.

We evaluate the method on the `version="small"` miniMUAD subset using true pixel-level OOD labels. The proposed hybrid score achieves an Area Under the Receiver Operating Characteristic curve (AUROC) of 0.8539, compared with 0.8280 for NECO-only, 0.8171 for Energy-only, and 0.8124 for an ensemble predictive-entropy baseline. Additional qualitative, distributional, condition-wise, and operating-point analyses show a consistent picture: the hybrid detector is stronger on the main ranking metric and across a broad operating range, while the ensemble baseline remains preferable in the extreme high-recall tail. This clarifies both the practical value and the current trade-off of the proposed single-pass design. Our code is publicly available at [https://github.com/boyuan-zhangx/Energy-Aware\\_NECO.git](https://github.com/boyuan-zhangx/Energy-Aware_NECO.git).

## I. INTRODUCTION

Semantic segmentation is a central dense perception task in computer vision, autonomous driving, and mobile robotics because it provides pixel-level scene understanding for downstream reasoning and control [1]–[4]. Modern segmentation systems achieve strong in-distribution performance, but they are still trained under a closed-set assumption and therefore remain vulnerable to previously unseen objects, textures, scene layouts, and environmental conditions at deployment time [5]–[7]. This gap makes out-of-distribution detection

particularly important in safety-critical applications such as autonomous driving and mobile robotics [5], [8], [9].

A common way to address this issue is to estimate predictive uncertainty. Bayesian approximations such as Monte Carlo Dropout by Gal and Ghahramani [10] provide useful uncertainty estimates through repeated stochastic forward passes. Deep Ensembles [12] and post-hoc calibration methods such as Temperature Scaling [13] and Local Temperature Scaling [14] further improve predictive reliability and confidence alignment. However, repeated inference or multiple independently trained models are expensive, which makes these approaches less attractive for resource-constrained edge deployment [10], [12].

The computational cost of repeated inference motivates single-pass post-hoc OOD scoring. Early confidence-based baselines such as Maximum Softmax Probability [15] are simple to compute. Stronger methods include ODIN by Liang *et al.* [16], Mahalanobis scoring by Lee *et al.* [17], Energy-based scoring by Liu *et al.* [18], and ViM by Wang *et al.* [19]. In parallel, Neural Collapse [20] provides a geometric interpretation of learned representations near the terminal phase of training, and NECO by Ben Ammar *et al.* [21] explicitly exploits this structure for OOD detection through principal-component geometry. These developments suggest that feature geometry and logit evidence are both meaningful signals for OOD detection.

Dense OOD detection in semantic segmentation remains challenging because the task is spatially structured and sensitive to object boundaries, context, and scale [6], [8], [9]. In addition, it is important to distinguish semantic anomalies from broader environmental covariate shifts, as the two may produce different uncertainty behaviors and influence OOD evaluation differently. Existing methods therefore span several directions, including entropy-based scoring, hybrid anomaly detection, synthetic negative data, nearest-neighbor methods, and object-centric anomaly segmentation [8], [9], [22]–[24]. This diversity indicates that no single score family is sufficient across all anomaly types.

Our design is motivated by two empirical observations. First, dense uncertainty maps are informative, but repeated inference is expensive. Second, purely geometric NECO-style scoring can miss simple anomalies that remain close to dominant in-distribution directions. We therefore propose Energy-Aware NECO, a single-pass pixel-wise OOD score for semantic segmentation. The method combines a centered NECO-style geometric ratio computed from decoder features with a logit-based Energy score. We evaluate it on miniMUAD, a lightweight subset of MUAD introduced by

<sup>1</sup>Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France. [firstname.lastname@polytechnique.edu](mailto:firstname.lastname@polytechnique.edu)

<sup>2</sup>CIAD, UTBM, Université Marie et Louis Pasteur, France. [firstname.lastname@utbm.fr](mailto:firstname.lastname@utbm.fr)

<sup>3</sup>U2IS, ENSTA, Institut Polytechnique de Paris, France. [firstname.lastname@ensta.fr](mailto:firstname.lastname@ensta.fr)

\*Corresponding Author. Yifei Cao was funded by the French National Research Agency (ANR) under grant agreement No ANR-23-CE10-0016.

Franchi *et al.* [25], and show that the hybrid score improves over both NECO-only and Energy-only under the same pixel-wise AUROC protocol.

Our contribution is threefold:

- 1) We adapt a NECO-inspired geometric OOD score to dense decoder features for semantic segmentation.
- 2) We propose a practical hybrid score that fuses centered feature-space geometry and logit-space Energy using only one forward pass.
- 3) We evaluate the method on miniMUAD with true pixel-level OOD masks and show that the hybrid score improves over both individual components and a predictive-entropy ensemble baseline under the same protocol.

## II. RELATED WORK

### A. Semantic Segmentation for Dense Perception

Semantic segmentation has evolved from fully convolutional dense prediction to encoder–decoder and transformer-based architectures [1]–[4]. These models provide strong in-distribution scene understanding, but they do not solve the open-world perception problem by themselves [5], [6]. This limitation is especially important in autonomous driving and mobile robotics, where deployment conditions often differ from the training distribution [7], [25].

### B. Uncertainty Estimation and Calibration

Uncertainty estimation is a standard route to robust prediction under distribution shift. MC Dropout interprets dropout at inference time as approximate Bayesian inference and estimates predictive uncertainty through repeated stochastic forward passes [10]. In computer vision, both epistemic and aleatoric uncertainty have been studied as useful signals for reliability-aware prediction [11]. Deep Ensembles and post-hoc calibration methods such as Temperature Scaling improve confidence estimation and predictive reliability without changing the main task definition [12], [13]. For dense prediction, local calibration methods have also been proposed to better account for spatially varying confidence [14]. Despite their effectiveness, repeated inference or multiple models remain costly for edge deployment [10], [12].

### C. Dense OOD and Anomaly Segmentation

OOD and anomaly segmentation have become established research topics for safety-critical road-scene perception [5], [6]. Fishyscapes introduced a benchmark for safe semantic segmentation in autonomous driving, while SegmentMeIfYouCan emphasized both pixel-wise and component-wise evaluation [5], [6]. On the method side, entropy-based meta classification, hybrid anomaly detection, synthetic negative data, nearest-neighbor scoring, and mask-level reasoning have all been explored for dense OOD detection [8], [9], [22]–[24]. This line of work shows that dense OOD detection often benefits from combining complementary cues rather than relying on a single scalar confidence measure [9], [22].

### D. Single-Pass OOD Scoring and Neural Collapse

A large family of OOD detectors operates in a post-hoc and often single-pass manner [15]–[19], [21]. Maximum Softmax Probability is a simple confidence baseline, while ODIN improves softmax-based OOD detection through temperature scaling and input perturbation [15], [16]. Energy-based scoring later showed that logits can provide a stronger OOD signal than softmax confidence, and ViM further combined feature statistics with virtual-logit matching [18], [19]. Mahalanobis-based methods also demonstrated that class-conditional feature geometry can be useful for post-hoc OOD detection [17]. On the representation side, Neural Collapse describes a family of geometric regularities emerging near the terminal phase of training, and NECO uses this geometry to construct a post-hoc OOD score [20], [21]. Our method follows this general direction, but adapts the idea to dense decoder features and explicitly complements the geometric score with Energy.

## III. BACKGROUND AND MOTIVATION

The motivation above is developed here in more detail. First, on MUAD, uncertainty-based dense prediction is meaningful. In our preliminary exploration, Deep Ensemble improved Expected Calibration Error (ECE) relative to a single model, and predictive-entropy maps from MC Dropout produced sensible spatial patterns, with uncertainty concentrated around semantic boundaries for in-distribution scenes and more globally elevated on OOD scenes. This observation is consistent with the general use of uncertainty estimation for reliability-aware prediction in computer vision and dense perception [10]–[12], [14]. At the same time, it also highlights the practical drawback of repeated inference for edge deployment [10], [12].

Second, in our earlier OOD and Neural Collapse exploration on image classification, purely geometric NECO-style scores worked better on high-energy, directionally distinct anomalies than on simple anomalies that remained close to the learned in-distribution subspace. This observation is consistent with the broader intuition that different OOD types may require different scoring cues and that hybrid feature-logit methods can be more robust than purely geometric or purely confidence-based methods alone [19], [21]–[23]. It therefore suggests that geometry alone is not always sufficient and that a complementary score based on logit magnitude or evidence may recover part of the missing signal [18], [19].

## IV. ENERGY-AWARE NECO FORMULATION

### A. Problem Setting

An input image is denoted by  $x \in \mathbb{R}^{H \times W \times 3}$ , where  $H$  and  $W$  are the image height and width. The segmentation network produces pixel-wise logits  $z(x) \in \mathbb{R}^{H \times W \times C}$  and decoder features  $h(x) \in \mathbb{R}^{H \times W \times d}$ , where  $C$  is the number of in-distribution (ID) semantic classes and  $d$  is the decoder feature dimension. Our goal is to assign to each valid pixel an OOD score such that larger values indicate stronger OOD evidence.

### B. Dense NECO Adaptation

The original NECO score measures how strongly a feature projects onto the principal ID subspace [21]. In our dense adaptation, we collect decoder pixel features from a pure ID split and compute their empirical mean  $\mu$  together with a principal-component basis  $P$ . For a pixel feature  $h(x)$ , the centered geometric ratio is defined as

$$\text{NECO}(x) = \frac{\|P(h(x) - \mu)\|_2}{\|h(x) - \mu\|_2 + \epsilon}, \quad (1)$$

where  $\epsilon > 0$  is a numerical-stability constant.

This formulation is motivated by the observation that deep representations tend to form structured class-wise geometries near the end of training, consistent with the Neural Collapse phenomenon [20]. This formulation is not an exact reproduction of the original NECO score. The modification is intentional: the original formulation was designed for classification representations, whereas our setting uses dense pixel-level decoder features. The centering step and projection pipeline therefore provide a stable and tractable adaptation to dense segmentation features while preserving the original geometric intuition.

### C. Energy Term

For the same pixel, we compute the logit-based Energy score [18]

$$\text{Energy}(x) = T \cdot \log \sum_{c=1}^C \exp\left(\frac{\text{logit}_c(x)}{T}\right), \quad (2)$$

where  $c$  indexes the semantic classes,  $\text{logit}_c(x)$  denotes the logit of class  $c$  at the considered pixel, and  $T > 0$  is the temperature parameter. The Energy term captures the magnitude and evidence structure of the logits. It complements the geometric ratio because some anomalous pixels may remain aligned with dominant ID directions while still producing atypical logit evidence.

### D. Validation-Fitted Standardization

The raw NECO and Energy values operate on different scales. We therefore standardize them using a pure ID validation split. Let  $\mu_{\text{NECO}}$ ,  $\sigma_{\text{NECO}}$ ,  $\mu_{\text{Energy}}$ , and  $\sigma_{\text{Energy}}$  denote the empirical mean and standard deviation of the two raw scores on that split. We define

$$Z_{\text{NECO}} = \frac{\text{NECO}(x) - \mu_{\text{NECO}}}{\sigma_{\text{NECO}}} \quad (3)$$

and

$$Z_{\text{Energy}} = \frac{\text{Energy}(x) - \mu_{\text{Energy}}}{\sigma_{\text{Energy}}}. \quad (4)$$

Only ID data are used for score fitting, which matches a realistic OOD protocol in which unknown categories are unavailable at fitting time [15], [18].

### E. Hybrid OOD Score

The fused ID-oriented score is defined as

$$S_{\text{ID}}(x) = \alpha \cdot Z_{\text{NECO}} + (1 - \alpha) \cdot Z_{\text{Energy}}, \quad (5)$$

where  $\alpha \in [0, 1]$  controls the trade-off between the geometric and logit-based terms. In our implementation,  $\alpha = 0.6$ , which gives slightly more weight to the geometric component, as it empirically provides a more stable reference under dense spatial representations.

$$S_{\text{OOD}}(x) = -S_{\text{ID}}(x), \quad (6)$$

so that larger values correspond to stronger OOD evidence.

### F. Complementarity of Geometry and Energy

The proposed fusion is motivated by complementary failure modes in post-hoc OOD detection. If an anomalous pixel deviates directionally from the learned ID principal subspace, the geometric term can detect it, which matches the Neural Collapse intuition behind NECO [20], [21]. If an anomalous pixel remains close to the dominant ID subspace but produces atypical logit magnitude or evidence, the Energy term can still provide useful discrimination [18]. By combining both signals, the proposed hybrid score can handle a wider range of dense anomalies under a single forward pass.

## V. EXPERIMENTAL SETUP

### A. Dataset

We evaluate on MUAD [25]. In the present implementation, we use the `version="small"` miniMUAD subset. This controlled subset is intentionally used to focus the evaluation on semantic anomaly detection while reducing the influence of severe environmental covariate shifts such as extreme lighting and weather conditions. We therefore focus our analysis on the scenes and proxy conditions that are available within this setup.

In our protocol, the OOD mask is a true pixel-level label mask, not a proxy. The segmentation model is trained with  $C = 19$  ID classes. Pixels with labels greater than or equal to the class count are treated as OOD labels in the evaluation split, while ignored pixels are removed from evaluation.

### B. Preprocessing and Score Fitting Split

All images and masks are resized to (256, 512). During training, we apply random horizontal flipping, convert images to `float32`, and normalize channels with ImageNet statistics. Labels with index  $\geq 19$  are remapped to the ignore value 255, so rare or anomalous classes do not contribute to the segmentation loss.

All score statistics are fitted on a pure ID validation split. Concretely, the validation loader is used to estimate the PCA basis  $P$ , the feature mean  $\mu$ , and the mean and standard deviation of the raw NECO and Energy scores. No OOD labels or OOD score statistics are used during score fitting.

TABLE I

COMPARISON OF OOD SCORES ON MINIMUAD UNDER THE SAME PIXEL-WISE AUROC PROTOCOL. HIGHER AUROC IS BETTER.

Method	AUROC	ID mean	OOD mean
Ensemble (predictive entropy)	0.8124	0.3021	0.8271
NECO-only	0.8280	0.3916	2.1789
Energy-only	0.8171	-0.0601	1.2303
Hybrid (NECO + Energy)	<b>0.8539</b>	0.2110	1.7995

### C. Baselines

We compare our method with the following baselines:

- Ensemble predictive entropy,
- NECO-only,
- Energy-only,
- Hybrid (NECO + Energy).

The ensemble baseline uses predictive entropy as the OOD score. It is therefore a reasonable uncertainty reference, but it is not the same scoring family as the proposed hybrid.

### D. Metrics

We report pixel-wise AUROC as the primary quantitative metric, which is standard in OOD and anomaly segmentation benchmarks [5], [6]. We also report the mean score on ID pixels and the mean score on OOD pixels to illustrate score separation. In addition, we analyze the false positive rate at 95% true positive rate (FPR95), the false positive rate at 98% true positive rate (FPR98), and the shape of the ROC curve, because operating-point behavior matters for safety-oriented dense OOD detection.

## VI. RESULTS

Table I summarizes the main quantitative comparison.

The hybrid score achieves the best AUROC, improving over both NECO-only and Energy-only. This supports the intended complementarity between feature-space geometry and logit-space evidence. The separation between ID and OOD mean scores is also clear, which suggests cleaner thresholding than with either individual component.

The comparison with the ensemble baseline should be interpreted as a comparison of deployment-oriented detection effectiveness rather than as a claim that both methods quantify the same notion of uncertainty. The ensemble baseline is a multi-model predictive-uncertainty estimate, whereas the proposed method is a single-pass fused OOD score. The main result is therefore not that the two methods are theoretically equivalent, but that the proposed single-pass detector remains competitive with a stronger but more expensive uncertainty baseline.

## VII. QUALITATIVE AND DISTRIBUTIONAL ANALYSIS

The main quantitative result in Table I is complemented by four figures that explain why the hybrid detector works, where the gain comes from, and why a geometric term alone is not sufficient.

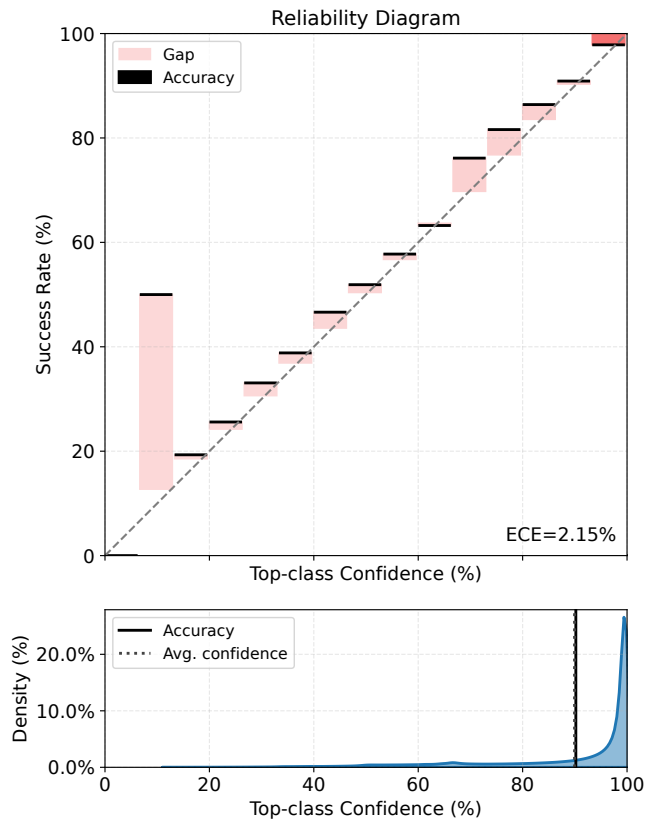


Fig. 1. Calibration evidence for the baseline segmentation model. The reliability diagram and confidence histogram show that the backbone already exhibits non-trivial confidence structure, which motivates uncertainty-aware dense OOD analysis on top of this predictor.

### A. Calibration Evidence from the Baseline Segmenter

Figure 1 shows the calibration behavior of the segmentation backbone. The reliability diagram indicates that the baseline predictor is not arbitrarily overconfident, and the confidence histogram remains concentrated near the high-confidence region with a limited calibration gap. This figure establishes that uncertainty-like outputs from the backbone already contain meaningful structure, which makes it reasonable to compare entropy-based and post-hoc OOD scores on top of the same predictor.

### B. Multi-Scene Qualitative OOD Maps

Figure 2 replaces the previous single-scene qualitative comparison with three representative miniMUAD scenes. Across these scenes, the hybrid detector produces more spatially continuous responses over the annotated OOD objects, with fewer internal holes and less fragmentation than the ensemble entropy maps. The ensemble baseline remains strong around generic scene edges, but its high-score regions are more diffuse and more tightly coupled to boundary structure. The figure therefore supports a more precise statement than the original draft: the advantage of Hybrid is not only numerical, but also structural at the object level.

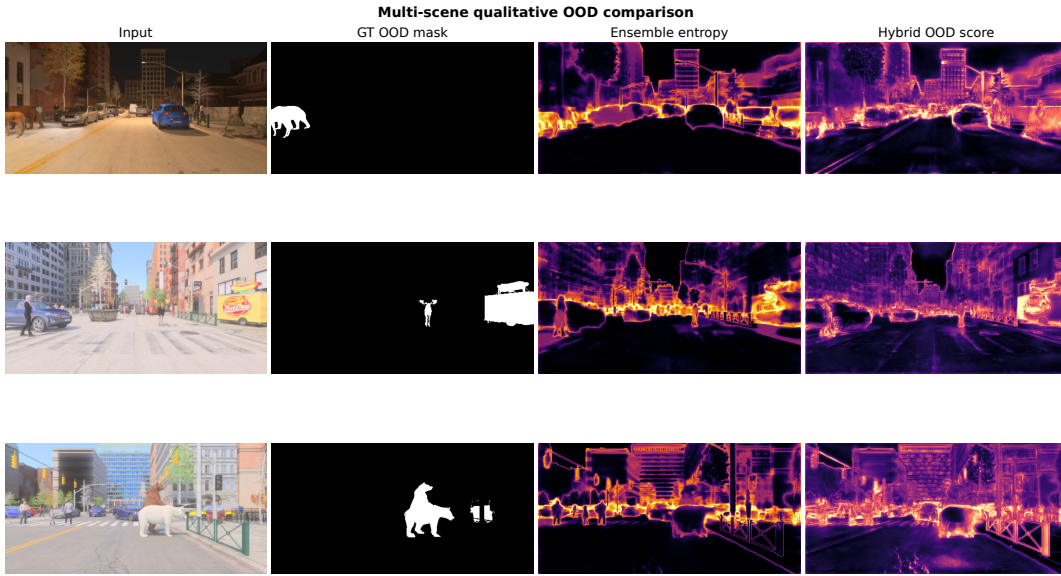


Fig. 2. Multi-scene qualitative OOD comparison on three representative miniMUAD scenes. Each row shows the input image, the ground-truth OOD mask, the ensemble predictive-entropy map, and the proposed hybrid OOD score. Across the three scenes, the hybrid detector produces more spatially continuous responses over the annotated OOD objects, while the ensemble baseline remains more edge-dominated.




Fig. 3. Distribution of ID and OOD scores for the compared methods. The hybrid score shows the clearest reduction of ID/OOD overlap, which is consistent with its AUROC advantage in Table I.

### C. Score Distribution Analysis

Figure 3 explains the AUROC improvement in Table I. The hybrid detector shows the clearest reduction of overlap between ID and OOD score distributions. Relative to the ensemble baseline, the OOD mass is shifted farther toward larger scores while the ID mass remains more concentrated at lower values. This is exactly the distributional pattern that supports better global ranking and more stable threshold selection.

### D. Feature-Geometry Failure Case

Figure 4 clarifies why the Energy term is necessary. It illustrates a representative failure case of a purely geometric detector: some simple anomalies remain embedded in dominant in-distribution directions and are therefore not cleanly separated by subspace projection alone. The figure justifies the design decision in Section IV: if geometry alone can miss a simple anomaly mode, then a complementary logit-based cue becomes well motivated.

## VIII. DISCUSSION AND FUTURE WORK

The main conclusion from Table I is not merely that the hybrid score attains the largest AUROC. The more important point is that the hybrid score improves over both NECO-only and Energy-only, which supports the intended complementarity between feature-space geometry and logit-space evidence. This interpretation is consistent with prior work showing that different OOD detectors capture different aspects of distribution shift and that hybrid dense OOD methods can benefit from combining multiple cues [18], [19], [22], [23].

This observation is also consistent with the design logic of our method. The centered NECO term captures subspace alignment with the learned ID representation. The Energy term captures abnormal activation magnitude and logit evidence. Some anomalies are well described by geometry, while others are better described by Energy. The hybrid combines both under a single-pass constraint.

The result is practically meaningful for edge deployment. Uncertainty-based baselines such as predictive entropy from

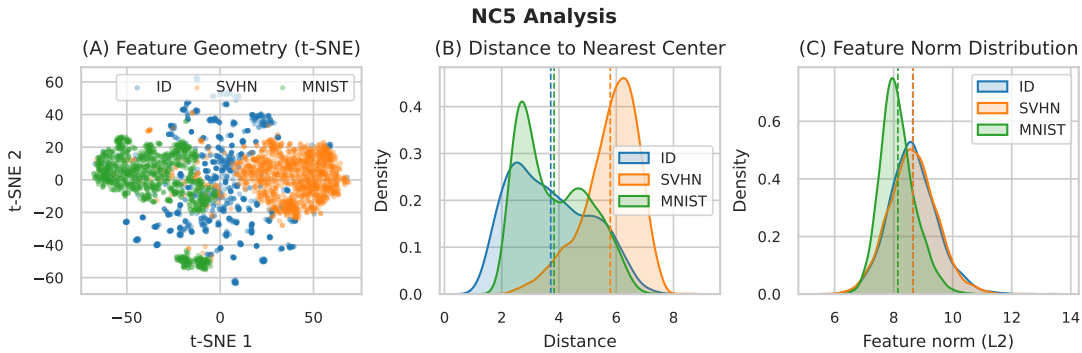


Fig. 4. Feature-geometry failure case from the earlier OOD and Neural Collapse exploration. The figure illustrates why a purely geometric detector may miss simple anomalies that remain embedded in dominant ID directions, which motivates the addition of the Energy term.




Fig. 5. ROC and operating-point behavior. Top: full ROC curve. Bottom: zoom on the high-TPR tail. Hybrid dominates over a broad operating range, but the ensemble baseline becomes preferable near the extreme recall regime.

ensembles or MC Dropout are useful, but they are more expensive at inference time [10], [12]. By contrast, the proposed score reuses features and logits that are already available in a standard forward pass.

#### A. Operating-Point Trade-Off in the High-TPR Tail

#### B. Condition-Wise Robustness Under Proxy Conditions

Figure 5 makes the trade-off explicit. In the full ROC curve, the hybrid detector dominates the ensemble baseline over a broad range of false-positive rates, which is consistent with the AUROC result in Table I. However, the tail zoom reveals a crossover in the high-TPR region. Around the very high recall regime, the ensemble baseline becomes preferable: at the marked  $\text{TPR} = 0.98$  operating point,

the ensemble reaches the target recall with a visibly lower false-positive rate than the hybrid detector. In contrast, at  $\text{TPR} = 0.95$ , the two methods are closer and Hybrid remains competitive.

This figure clarifies the main trade-off of the paper. The proposed hybrid score is stronger on the global ranking metric and on much of the operating range, but it is less well controlled in the extreme high-recall tail. A plausible explanation is that the Energy component produces a heavier high-score tail on hard ID pixels. When the threshold is lowered to capture the final few percent of OOD pixels, these difficult ID pixels cross the decision boundary in large numbers. By contrast, the ensemble baseline benefits from variance reduction through multi-model averaging, which yields better false-positive control in the most conservative regime.

Figure 6 analyzes four proxy conditions available in min-iMUD: `low_light`, `high_light`, `low_contrast`, and `high_contrast`. Three observations matter. First, Hybrid remains above Ensemble in AUROC for all four subsets. Second, FPR95 remains broadly comparable, with only modest condition-dependent differences. Third, the condition-wise FPR98 bars remain close as well. This means that the operating-point weakness of the hybrid detector does not appear as a uniform collapse across all proxy conditions. Instead, the more severe weakness emerges in the extreme high-recall tail of the global ROC curve.

The current study also points to several natural future directions. First, a local spatial smoothing stage or lightweight structured post-processing, such as average pooling, Gaussian smoothing, or a compact Conditional Random Field (CRF), could suppress isolated false positives on hard ID pixels. Second, outlier exposure or virtual outlier synthesis could explicitly tighten the ID energy tail during training. Third, logit normalization or a distance-based alternative to raw Energy, such as Mahalanobis or nearest-neighbor scoring in decoder space, could improve behavior near extreme operating points. Fourth, a head-only lightweight ensemble, for example dropout restricted to the prediction head, could recover part of the variance reduction of deep ensembles without paying their full computational cost.




Fig. 6. Condition-wise robustness under four proxy conditions available in miniMUAD. Hybrid improves AUROC across all four subsets. FPR95 remains broadly comparable. The condition-wise FPR98 bars stay close as well, suggesting that the most severe weakness of Hybrid is not a uniform per-condition degradation, but a high-recall tail effect revealed by the global ROC analysis.

## IX. CONCLUSION

We presented Energy-Aware NECO, a single-pass pixel-wise OOD score for semantic segmentation. The method fuses a centered NECO-style geometric ratio with a logit-based Energy score, using score statistics fitted on a pure ID validation split. On miniMUAD with true pixel-level OOD masks, the proposed hybrid achieves the best AUROC among the compared methods.

The broader conclusion is nuanced. The hybrid detector offers a strong efficiency–performance compromise for dense OOD detection, with cleaner global ranking, stronger qualitative object-level responses, and robust behavior across several proxy-condition splits. At the same time, the ROC-tail analysis shows that repeated-inference uncertainty estimation still retains an advantage in the most conservative high-recall regime. This makes the proposed method a promising deployment-oriented detector rather than a universal replacement for ensemble-based uncertainty estimation.

## REFERENCES

- [1] J. Long, E. Shelhamer, and T. Darrell, “Fully Convolutional Networks for Semantic Segmentation,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 3431–3440.
- [2] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in *Medical Image Computing and Computer-Assisted Intervention*, 2015, pp. 234–241.
- [3] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, “Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation,” in *Proceedings of the European Conference on Computer Vision*, 2018, pp. 801–818.
- [4] E. Xie et al., “SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers,” in *Advances in Neural Information Processing Systems*, 2021.
- [5] H. Blum, P.-E. Sarlin, J. Nieto, R. Siegwart, and C. Cadena, “Fishyscapes: A Benchmark for Safe Semantic Segmentation in Autonomous Driving,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2019.
- [6] R. Chan, K. Lis, S. Uhlemeyer, H. Blum, S. Honari, R. Siegwart, P. Fua, M. Salzmann, and M. Rottmann, “SegmentMelfYouCan: A Benchmark for Anomaly Segmentation,” in *NeurIPS Datasets and Benchmarks*, 2021.
- [7] D. Bogdoll et al., “AnoVox: A Benchmark for Multimodal Anomaly Detection in Autonomous Driving,” arXiv preprint, 2024.
- [8] R. Chan, M. Rottmann, and H. Gottschalk, “Entropy Maximization and Meta Classification for Out-of-Distribution Detection in Semantic Segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 5128–5137.
- [9] S. N. Rai, F. Cermelli, D. Fontanel, C. Masone, and B. Caputo, “Unmasking Anomalies in Road-Scene Segmentation,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 4037–4048.
- [10] Y. Gal and Z. Ghahramani, “Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning,” in *Proceedings of the 33rd International Conference on Machine Learning*, 2016, pp. 1050–1059.
- [11] A. Kendall and Y. Gal, “What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision?” in *Advances in Neural Information Processing Systems*, 2017.
- [12] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles,” in *Advances in Neural Information Processing Systems*, 2017.
- [13] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, “On Calibration of Modern Neural Networks,” in *Proceedings of the 34th International Conference on Machine Learning*, 2017, pp. 1321–1330.
- [14] Y. Ding et al., “Local Temperature Scaling for Probability Calibration,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [15] D. Hendrycks and K. Gimpel, “A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks,” in *International Conference on Learning Representations*, 2017.
- [16] S. Liang, Y. Li, and R. Srikant, “Enhancing the Reliability of Out-of-distribution Image Detection in Neural Networks,” in *International Conference on Learning Representations*, 2018.
- [17] K. Lee, K. Lee, H. Lee, and J. Shin, “A Simple Unified Framework for Detecting Out-of-Distribution Samples and Adversarial Attacks,” in *Advances in Neural Information Processing Systems*, 2018.
- [18] W. Liu, X. Wang, J. D. Owens, and Y. Li, “Energy-based Out-of-distribution Detection,” in *Advances in Neural Information Processing Systems*, 2020.
- [19] H. Wang, Z. Li, L. Feng, and W. Zhang, “ViM: Out-of-Distribution with Virtual-logit Matching,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2022, pp. 4921–4930.
- [20] V. Pappayan, X. Y. Han, and D. L. Donoho, “Prevalence of Neural Collapse During the Terminal Phase of Deep Learning Training,” *Proceedings of the National Academy of Sciences*, vol. 117, no. 40, pp. 24652–24663, 2020.
- [21] M. Ben Ammar, N. Belkhir, S. Popescu, A. Manzanera, and G. Franchi, “NECO: Neural Collapse Based Out-of-distribution Detection,” in *International Conference on Learning Representations*, 2024.
- [22] M. Grcić, P. Bevandić, and S. Šegvić, “DenseHybrid: Hybrid Anomaly Detection for Dense Open-set Recognition,” arXiv preprint, 2022.
- [23] M. Grcić and S. Šegvić, “Hybrid Open-set Segmentation with Synthetic Negative Data,” arXiv preprint, 2023.
- [24] S. Galesso, M. Argus, and T. Brox, “Far Away in the Deep Space: Dense Nearest-Neighbor-Based Out-of-Distribution Detection,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2023.
- [25] G. Franchi, X. Yu, A. Bursuc, A. Tena, R. Kazmierczak, S. Dubuisson, E. Aldea, and D. Filliat, “MUAD: Multiple Uncertainties for Autonomous Driving, a Benchmark for Multiple Uncertainty Types and Tasks,” in *Proceedings of the British Machine Vision Conference*, 2022.