Microglia-RTS: Instance Segmentation of Microglia in Hypoxic Ischemic Fetal Sheep Brain Histology

Callan Loomes[1[0009-0001-5669-2048]], Joanne Davidson[1[0000-0002-6781-4258]], Grace Tian[1[0009-0002-9293-9471]], Alistair Gunn[1[0000-0003-0656-7035]], Laura Bennet[1[0000-0002-4336-7596]], and Hamid Abbasi[1,2[0000-0003-1136-3280]]

[1] Department of Physiology, Faculty of Medical and Health Sciences, University of Auckland, Auckland, New Zealand

[2] Auckland Bioengineering Institute, University of Auckland, Auckland, New Zealand

callan.loomes@auckland.ac.nz, h.abbasi@auckland.ac.nz

Abstract - Hypoxic–ischemic encephalopathy (HIE) is a major cause of neonatal death and lifelong neuro-disability, with limited treatment options beyond therapeutic hypothermia within 6 hours of birth. In translational fetal sheep models, microglia are central to inflammatory injury cascades and are a key immunohistochemical readout for evaluating therapeutic efficacy. Still, their quantification and morphological grading in histological images remains largely manual, limiting throughput, reproducibility, and mechanistic insights. We propose Microglia-RTS, a two-stage pipeline for microglia instance segmentation in high-resolution histology for preclinical drug screening. The first detection stage adapts a modified RT-DETR with a pretrained DINOv3 ConvNeXt backbone that incorporates a novel adaptive query selection module to tune the number of object queries per image based on local cell density. The second segmentation stage fine-tunes the Segment Anything Model (SAM2) to produce cell masks from detector proposals. Patch-based inference with overlap-aware stitching enables high-resolution processing of images, preserving intricate cellular details. We evaluate Microglia-RTS on a new fully annotated dataset comprising ~7,000 cell instances, classified into ameboid, intermediate or ramified, across 45-Iba1 stained hypoxic–ischemic fetal sheep brain images from three white matter regions (IGWM1, IGWM2 and PVWM). The modified detector achieved an F1-Detection score of 0.55±0.01 during validation, considerably surpassing the baseline model, and the fine-tuned SAM2 achieved an average Dice-score of 0.86±0.01. Across 5-fold cross-validation on whole images, Microglia-RTS achieved a mean panoptic quality of 0.48±0.02, supporting scalable microglia analysis for pre-clinical HIE drug discovery by reducing observer bias and accelerating quantitative endpoint extraction.

Keywords: Instance Segmentation, Deep Learning, Brain Cells

1 Introduction

Hypoxic-ischemic encephalopathy (HIE) is a neurological disorder caused by inadequate fetal brain perfusion around the time of birth, commonly due to umbilical cord entanglement or impaired placental blood flow [1,2]. Affecting 0.15% of live births,

HIE is associated with serious outcomes including epilepsy, cerebral palsy, neurodevelopmental impairment, and in severe cases, infantile death [1,3]. The pathophysiology of HIE unfolds in phases: during the initial 6-7 hours after insult reperfusion, denoted the latent phase, cell metabolism is partially restored [2,4]. However, inflammatory and apoptotic cascades remain active, progressing to a secondary phase (roughly 6 to 72 hours post-injury) characterized by mitochondrial dysfunction and apoptosis, and often accompanied by seizures [4,5].

Microglia, the resident immune cells of the brain, are central to this inflammatory progression [6]. Under hypoxic-ischemic (HI) conditions they transition from ramified toward ameboid morphologies, reflecting activation states linked to neuroinflammation, phagocytosis and neuronal injury [7-9]. Therefore, their quantification and morphological phenotyping are important for tracking disease progression and evaluating mechanisms of candidate therapies and drugs [10,11].

Therapeutic hypothermia (TH) is currently the only clinically validated treatment for term HIE-impacted infants; yet its efficacy depends on early initiation within the latent phase, which is often difficult in practice because injury timing is uncertain [12]. Preclinical research using fetal sheep models for adjunctive therapy discovery is constrained by manual cell counting and classification, which are time-intensive and prone to inter/intra-observer biases [4,13,14]. Robust automated, reproducible approaches are needed to enable scalable histology phenotyping and richer mechanistic insight [15,16]. While biological instance segmentation methods exist [17-20] and increasingly utilize transformer-based detection (e.g., DEtection TRansformer (DETR) [21] and its variants such as Real-Time DETR (RT-DETR) [22]) and foundation segmentation models such as the Segment Anything Model (SAM2) [23], their application to microglia in HI-impacted brain histology remains limited.

This paper, for the first time, addresses this gap by introducing a novel two-stage pipeline for automated microglia analysis in fetal HI sheep tissue. This work presents two main contributions: (i) a new annotated dataset of ~7000 Iba1-stained microglia across 3 morphological classes in 3 white-matter regions in HI fetal sheep brains, and (ii) a robust instance-segmentation framework combining modified RT-DETR and SAM2 with adapted pretraining, novel architectural modifications, optimized training methodology, and patching/re-stitching for scalable high-resolution inference.

2 Methods

2.1 Ethics, Data collection, and Preparation

All procedures were approved by the Animal Ethics Committee of The University of Auckland under the New Zealand Animal Welfare Act. Animal care, handling, and surgical procedures followed previously reported protocols [4,14]. Data were collected from 28 Romney/Suffolk fetal sheep at 0.85 gestation (near-term brain-equivalent), across three experimental groups: HI-normothermia (n=8), HI-hypothermia (72-hr hypothermia; n=14), and sham-normothermia (no HI; n=6). After each experimental protocol, brains were collected post-mortem, formalin-fixed, paraffin-embedded, sectioned coronally at 10μm, and immunolabelled with Iba1. Microglia were imaged by

light microscopy at ×20 magnification in three white-matter regions: the intragyral white matter of the first and second parasagittal gyri (IGWM1 and IGWM2) and the periventricular white matter (PVWM). Across the 28 subjects, 45 images (15 per region; 2048 × 2880 px) were selected for annotation.

While generalist microglia classifications typically range from four canonical morphotypes (ramified, rod-like, activated, ameboid) to six or more categories [24-26], here we use a three-class scheme (ramified, intermediate, ameboid) based on representative examples in Fig. 1. This approach preserves biological interpretability while simplifying annotation by merging 'activated' and 'rod-like' morphologies into a single 'Intermediate' class, spanning the continuum between highly branched 'Ramified' cells and blob-like 'Ameboid' cells with retracted processes.

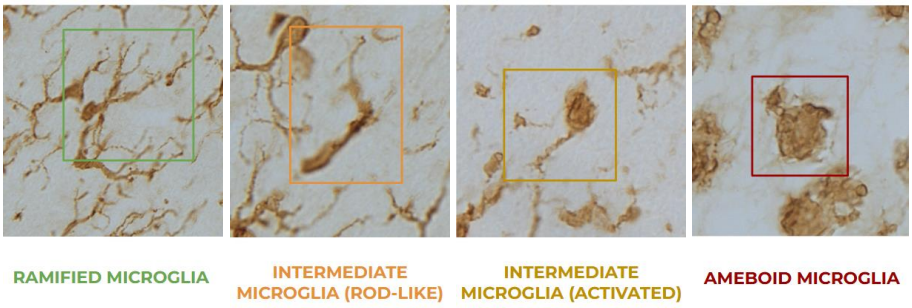

Fig. 1. Representative microglia for the three-class classification scheme.

Annotation followed a three-stage workflow: first, clinicians led by J.D, annotated microglial instances with points in the raw images using ImageJ. Second, expert annotators (C.L & G.T) segmented and classified each marked instance in Napari using a custom plugin [27]. Third, all images underwent final review by C.L before model training. Final class-wise and region-wise counts are summarized in Table 1.

Table 1. Summary of the total number of annotated microglia instances in the created dataset.

Region	AMEBOID	INTERMEDIATE	RAMIFIED	TOTAL
IGWM1	307	587	801	1695
IGWM2	505	837	803	2145
PVWM	1017	1312	773	3102
TOTAL	2256	2689	1812	**6942**

2.2 Detection and Classification with Modified RT-DETR

For microglia detection, a modified RT-DETRv2 [28] was trained with 5-fold cross validation (80:20 train to validation split, with brain regions balanced across folds). RT-DETRv2 [28] was adopted as the base detector over the more recent RT-DETRv3 [29] on grounds of reproducibility and pipeline portability. At project inception, RT-

DETRv2 was officially supported in the HuggingFace Transformers Library [30], providing well-documented APIs and streamlined deployment across compute environments, whereas RT-DETRv3 was released as a separate research repository.

Each fold was trained for 80 epochs, where visual inspection indicated a plateau in training performance across all models, and the best validation model was retained. A three-phase unfreezing strategy was applied: backbone and encoder frozen for epochs 1–5, encoder unfrozen at epoch 6, and full end-to-end training from epoch 11, with layer-wise adjusted learning rates. An AdamW optimizer was used with a base learning rate of 1×10^{-4}, a 2-epoch linear warmup, and cosine annealing decay.

Data augmentations including vertical/horizontal flips (50% each), mild affine transforms (rotation, scaling, translation), and random gamma adjustments were applied. To increase training variability and support scalable patch-based inference, we used random patch-sampling during training: five 1024×1024 px patches were sampled from each image per epoch using a greedy selection strategy, exposing the model to different image regions across epochs. Detection performance was measured using the F1-detection score at Intersection over Union (IoU) ≥ 0.5, where a true positive required both correct class prediction and overlap ≥ 0.5 with the ground-truth instance (Eq. 1).

$$F_{1,d@IoU=0.5} = \frac{2TP_{d@IoU=0.5}}{2TP_{d@IoU=0.5}+FN_{d@IoU=0.5}+FP_{d@IoU=0.5}} \quad (1)$$

Backbone Comparison: Our first architectural modification replaced the standard RT-DETRv2 ResNet-50 backbone with DINOv3 self-supervised pretrained backbones (pretrained on the LVD-1689M dataset [31]). With query count fixed at 300, we evaluated three different DINOv3-pretrained backbone archetypes: ConvNeXt (small, large), ViT (small, large), and a hybrid ConvNeXt-ViT architecture with three-scale feature fusion via concatenation (ConvNeXt-Small + ViT-Small and ConvNeXt-Large + ViT-Base variants). As they are included with the "off-the-shelf" RT-DETRv2 model, ResNet-34, ResNet-50, and ResNet-101 backbones were tested as baselines.

Deformable Query Modifications: Using the best-performing backbone (ConvNeXt-Large), we then evaluated the effect of modifying query count. In RT-DETRv2, queries are encoder features selected by "objectness" scoring, which determines how many candidate regions are refined and thus mediates recall-precision trade-off [28]. Query count was varied between 20 and 500, with performance binned by object count per image to validate the hypothesis that fewer queries benefit lower-density images. We also introduced a novel adaptive query branch that dynamically predicts query count at runtime using a lightweight multi-layer perceptron (MLP) applied to the deepest backbone feature map. This branch is supervised with an auxiliary mean squared error (MSE) loss against a target derived from ground-truth object count. A learned MLP was preferred to other density estimation methods (e.g. intensity statistics or classical blob counting) as the deepest backbone feature map encodes scale, morphology and staining-variation cues relevant to microglia density; end-to-end optimisation with the MLP aligns query-count prediction with the downstream detection objective.

2.3 Microglia-RTS: Combined Pipeline with SAM2 Segmentation

Microglia segmentation used a fine-tuned SAM2-small model trained on the same 5-fold cross-validation scheme as the detection phase. Each fold was trained for 50 epochs with the image encoder frozen, fine-tuning only the prompt encoder and mask decoder using bounding box prompts. AdamW (lr=1×10^{-4}, weight decay=0.01) with a combined binary cross-entropy (BCE)-Dice loss [32] was used, with the best checkpoint per fold selected by validation Dice [33] (Eq. 2).

$$Dice = \frac{2|P \cap G|}{(|P|+|G|\}} \quad (P = predicted\ mask, G = ground\ truth\ mask) \tag{2}$$

The final pipeline, denoted Microglia-RTS for the remainder of this study, allowed for full image inference. Microglia-RTS patchifies the input image, processes each patch with the modified RT-DETRv2 (ConvNeXt-Large backbone with adaptive query branch) followed by fine-tuned SAM2, and then merges patch-level predictions via majority voting and overlap-aware re-stitching. A schematic of the final modified RT-DETRv2 used in the Microglia-RTS pipeline is shown in Fig. 2.

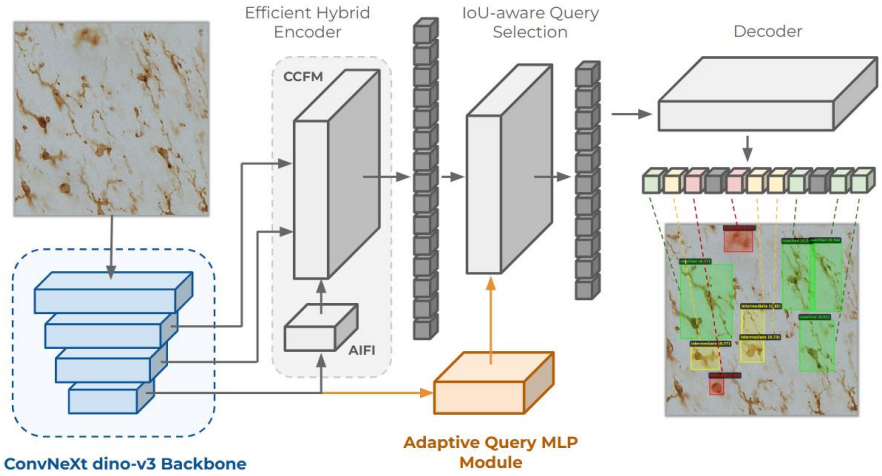

Fig. 2. Detection module of the final Microglia-RTS pipeline, including enhanced DINO-v3 ConvNeXt backbone (blue) and adaptive query MLP module (orange).

The combined pipeline was evaluated on a consistent validation set using the panoptic quality (PQ) metric [34] at IoU=0.5, defined as the product of detection quality (DQ) and segmentation quality (SQ) (Eq. 3). TP, FP and FN denote true positive, false positive and false negative instance matches respectively (at IoU=0.5), and IoU(P,G) is the intersection-over-union between matched predicted (P) and ground-truth (G) instance masks. Mean PQ (mPQ) was computed as the average of per-class PQ scores.

$$PQ = DQ \times SQ = \frac{TP}{(TP+0.5\cdot FP+0.5\cdot FN)} \times \frac{1}{|TP|}\sum_{(P,G)\in TP} IoU(P,G) \tag{3}$$

2.4 Computing Infrastructure

Microglia-RTS was implemented and trained using the PyTorch library in Python using New Zealand eScience Infrastructure (NeSI) high-performance computing facilities. The training and inference process was executed using NVIDIA A100 GPUs.

3 Results

3.1 Architectural Modifications for RT-DETRv2

Fig. 3 summarizes RT-DETRv2 detection performance across backbone choices. Among all tested configurations, the ConvNeXt-Large backbone achieved the best overall result, reaching a validation F1-detection score of 0.489±0.033. Notably, a clear size effect was observed with larger backbones generally outperforming their smaller counterparts within the same family. The hybrid backbones consistently performed better than the ResNet-50 baseline.

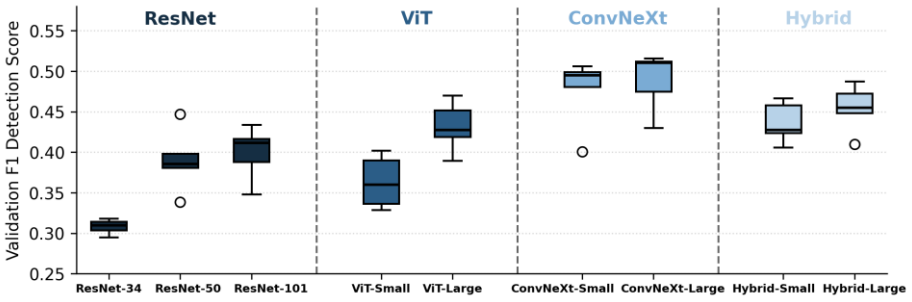

Fig. 3. Summary of RT-DETRv2 detection performance across backbone configurations.

Fig. 4 summarizes the query-count exploration; as shown in Fig. 4A, optimal query count depended on patch density, where lower query counts were more favorable for sparse patches whereas higher counts were better suited to denser patches. Consistent with this pattern, the adaptive query branch achieved the strongest overall performance across fixed-query settings (Fig. 4B), with the adaptive configuration (dynamic range: 20 to 100 queries) achieving the highest mean F1-detection score of 0.548±0.014.

3.2 Performance of Fine-tuned SAM2 and Combined pipeline

The fine-tuned SAM2 model achieved a Dice score of 0.862±0.008 across folds, indicating strong mask prediction performance on the microglia dataset. Full-pipeline evaluation on the validation set (25% patch overlap) is summarized in Table 2 and provides an ablative-style comparison of detection backbone and query strategy contributions.

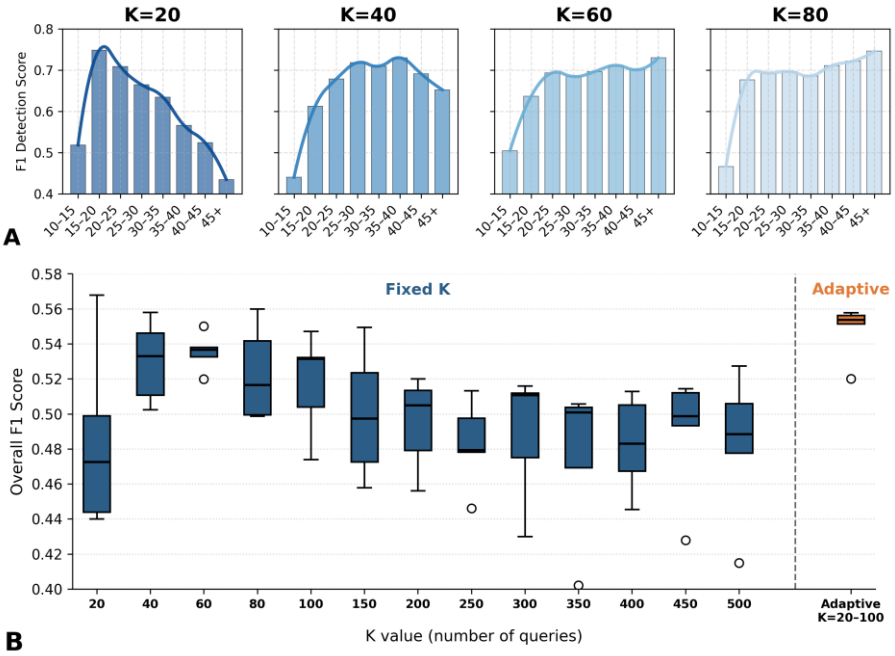

Fig. 4. Performance of different fixed query models at different instance-thresholds (**A**), and performance of different fixed query models against adaptive-query model (**B**).

Representative outputs from the final Microglia-RTS pipeline (ConvNeXt-Large + adaptive queries) are shown in Fig. 5 for relatively healthy and relatively damaged fetal sheep tissue, illustrating qualitatively strong agreement with ground truth across distinct tissue states.

Table 2. Ablation results summary of Microglia-RTS for different detection backbones. Bolded values represent highest performance across models.

Backbone	Number of Queries	mPQ	mSQ	mDQ	Class-Wise PQ		
					Ameboid	Int.	Ramified
ResNet50	K=300	0.35±0.03	0.74±0.01	0.47±0.04	0.35±0.10	0.31±0.02	0.38±0.02
ConvNeXt-Large	K=300	0.45±0.06	**0.76±0.01**	0.59±0.07	**0.47±0.08**	0.42±0.06	0.46±0.06
ConvNeXt-Large	K=60	0.46±0.05	**0.76±0.01**	0.61±0.06	0.45±0.10	0.44±0.04	0.49±0.05
ConvNeXt-Large	Adaptive (K=20-100)	**0.48±0.02**	**0.76±0.01**	**0.63±0.03**	0.46±0.06	**0.45±0.03**	**0.52±0.04**

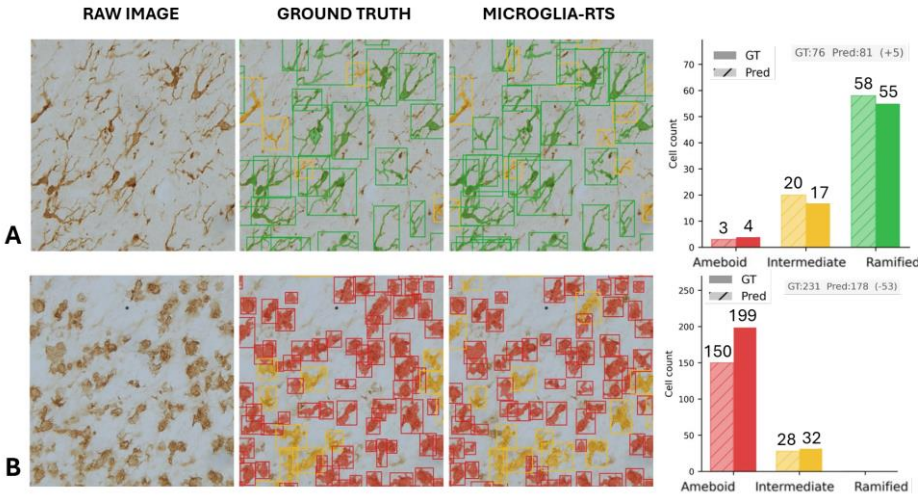

Fig. 5. Representative outputs (images are cropped and zoomed) from the full Microglia-RTS pipeline for relatively healthy (**A**) and relatively damaged (**B**) fetal sheep brain tissues.

4 Discussion and Conclusions

This study presented Microglia-RTS, a two-stage pipeline that integrated a modified RT-DETRv2 detector with a fine-tuned SAM2 model for automated instance segmentation of Iba1-stained microglia. The pipeline's primary novelty lies in the detector design: substituting the standard ResNet-50 backbone with a DINOv3-pretrained ConvNeXt-Large backbone [28,31] produced the greatest improvement in end-to-end performance, increasing mPQ from 0.35±0.03 to 0.45±0.06 (Table 2). This supports the value of the DINOv3 self-supervised pretraining framework for domain transfer to histology [31]. While inclusion of the adaptive query branch provided smaller mPQ gains, its practical advantage was improved consistency across folds (Fig. 4B). Because microglial density varies markedly with injury severity, from sparse ramified cells to dense ameboid clusters in inflamed tissue [4,9,14], adaptive query selection appeared to act as a density-aware regulator that reduces over-prediction in sparse patches while preserving capacity in dense regions. At larger training scales, this advantage may narrow if fixed higher-query models learn stronger implicit suppression through Hungarian loss matching (a critical mechanism for DETRs) during optimization [21].

Across detection module variants in Table 2, Microglia-RTS's mSQ remained stable and consistently higher than mDQ, while mPQ closely tracked gains in mDQ; this indicates that overall performance is currently detection-limited rather than segmentation-limited. Further, the relatively low detection F1 score of 0.55 reflects the dataset's intrinsic difficulty, possibly due to combining instance detection in clustered microglial environments with three-way morphological classification on a biologically continuous ramified-to-ameboid spectrum. In contrast, strong segmentation results likely stem from SAM2's pre-trained mask priors transferring well to cellular morphologies [23].

Furthermore, class-wise analysis highlighted lower PQ for intermediate and ameboid cells than for ramified microglia (Table 2), likely due to morphological overlap in transitional states and the tendency of injured microglia to form dense aggregates that complicate detection and classification (e.g. Fig. 5B).

Overall, Microglia-RTS offers a scalable, reproducible, density-robust framework for high-throughput microglia detection, phenotyping, and segmentation in fetal sheep HI histology. By addressing a major throughput bottleneck in translational HIE research, it enables faster, more quantitative evaluation of neuroinflammatory responses and pre-clinical assessment of adjunctive therapies to therapeutic hypothermia.

Future Directions and Limitations: Expanding the dataset could potentially improve detection accuracy, and also support finer-grained phenotyping beyond the current three-class scheme (e.g., separating activated and rod-like intermediate morphologies), improving biological specificity for meaningful distinctions relevant to inflammatory progression. As detection remains the primary performance bottleneck, alternative detection architectures, such as RT-DETRv3 [29], may warrant further investigation.

Data Availability: The dataset is not publicly available due to ethical restrictions; data and code may be shared upon reasonable request, subject to formal agreements.

Acknowledgments. The authors acknowledge the use of the New Zealand eScience Infrastructure (NeSI) for the results of this research. https://www.nesi.org.nz.

Disclosure of Interests. The authors have no conflicts of interest to disclose.

References

1. Sánchez-Rodríguez, E.C., López, V.J.: Hypoxic ischemic encephalopathy (HIE). Frontiers in Neurology. 15, (2024). https://doi.org/10.3389/fneur.2024.1389703.
2. Yıldız, E.P. et al.: Neonatal hypoxic ischemic encephalopathy: an update on disease pathogenesis and treatment. Expert Review of Neurotherapeutics. 17, 5, 449–459 (2016). https://doi.org/10.1080/14737175.2017.1259567.
3. Douglas-Escobar, M., Weiss, M.D.: Hypoxic-ischemic encephalopathy: a review for the clinician. JAMA Pediatrics. 169, 4, 397 (2015). https://doi.org/10.1001/jamapediatrics.2014.3269.
4. Davidson, J.O. et al.: How long is sufficient for optimal neuroprotection with cerebral cooling after ischemia in fetal sheep? Journal of Cerebral Blood Flow & Metabolism. 38, 6, 1047–1059 (2017). https://doi.org/10.1177/0271678x17707671.
5. Bennet, L. et al.: Potential biomarkers for hypoxic–ischemic encephalopathy. Seminars in Fetal and Neonatal Medicine. 15, 5, 253–260 (2010).
6. Wolf, S.A. et al.: Microglia in Physiology and Disease. Annual Review of Physiology. 79, 1, 619–643 (2016). https://doi.org/10.1146/annurev-physiol-022516-034406.
7. Lv, Y. et al.: The role of microglia mediated pyroptosis in neonatal hypoxic-ischemic brain damage. Biochemical and Biophysical Research Communications. 521, 4, 933–938 (2019). https://doi.org/10.1016/j.bbrc.2019.11.003.
8. Bernis, M.E. et al.: Temporal Characterization of Microglia-Associated Pro- and Anti-Inflammatory Genes in a Neonatal Inflammation-Sensitized Hypoxic-Ischemic Brain Injury Model. Oxidative Medicine and Cellular Longevity. 2022, 23, 1–16 (2022).
9. Liu, Y. et al.: Comparisons in microglial responses to ischemic–hypoxic brain injury in neonatal and adult mice analyzed by transcriptome integration and deconvolution

technologies. Advanced Technology in Neuroscience. 2, 4, 161–172 (2025). https://doi.org/10.4103/atn.atn-d-25-00004.
10. Green, T.R.F., Rowe, R.K.: Quantifying microglial morphology: an insight into function. Clinical and Experimental Immunology. 216, 3, 221–229 (2024).
11. Morrison, H. et al.: Quantitative microglia analyses reveal diverse morphologic responses in the rat cortex after diffuse brain injury. Scientific Reports. 7, 1, (2017). https://doi.org/10.1038/s41598-017-13581-z.
12. Wassink, G. et al.: Therapeutic Hypothermia in Neonatal Hypoxic-Ischemic Encephalopathy. Current Neurology and Neuroscience Reports. 19, 2, (2019). https://doi.org/10.1007/s11910-019-0916-0.
13. Reddaway, J. et al.: Microglial morphometric analysis: so many options, so little consistency. Frontiers in Neuroinformatics. 17, (2023). https://doi.org/10.3389/fninf.2023.1211188.
14. Wassink, G. et al.: Partial white and grey matter protection with prolonged infusion of recombinant human erythropoietin after asphyxia in preterm fetal sheep. Journal of Cerebral Blood Flow & Metabolism. 37, 3, 1080–1094 (2016).
15. Loomes, C. et al.: Automated Cell Quantification in Hypoxic-Ischemic Fetal Sheep Brain Histology: A Two-Step Segmentation and Classification Pipeline. Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference. 2025, (2025). https://doi.org/10.1109/embc58623.2025.11253411.
16. Abbasi, H. et al.: EEG sharp waves are a biomarker of striatal neuronal survival after hypoxia-ischemia in preterm fetal sheep. Scientific reports. 8, 1, (2018). https://doi.org/10.1038/s41598-018-34654-7.
17. Hörst, F. et al.: CellViT: Vision Transformers for precise cell segmentation and classification. Medical Image Analysis. 94, 103143 (2024). https://doi.org/10.1016/j.media.2024.103143.
18. Pang, M. et al.: CelloType: a unified model for segmentation and classification of tissue images. Nature Methods. 22, 2, 348–357 (2024). https://doi.org/10.1038/s41592-024-02513-1.
19. Archit, A. et al.: Segment Anything for Microscopy. Nature Methods. 22, 3, 579–591 (2025). https://doi.org/10.1038/s41592-024-02580-4.
20. Zhou, C. et al.: MSRT-DETR: A novel RT-DETR model with multi-scale feature sequence for cell detection. Biomedical Signal Processing and Control. 103, 107378 (2025). https://doi.org/10.1016/j.bspc.2024.107378.
21. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) ECCV 2020, LNCS, vol. 12346, pp. 213–229. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-58452-8_13
22. Zhao, Y., Lv, W., Xu, S., Wei, J., Wang, G., Dang, Q., Liu, Y., Chen, J.: DETRs beat YOLOs on real-time object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 16965–16974 (2024)
23. Ravi, N. et al.: SAM 2: Segment Anything in Images and Videos, (2024). https://doi.org/10.48550/arxiv.2408.00714.
24. Leyh, J. et al.: Classification of Microglial Morphological Phenotypes Using Machine Learning. Frontiers in Cellular Neuroscience. 15, (2021). https://doi.org/10.3389/fncel.2021.701673.

25. Hsu, C.-H. et al.: StainAI: quantitative mapping of stained microglia and insights into brain-wide neuroinflammation and therapeutic effects in cardiac arrest. Communications Biology. 8, 1, (2025). https://doi.org/10.1038/s42003-025-07926-y.
26. Xiong, H. et al.: μGlia-Flow, an automatic workflow for microglia segmentation and classification. Journal of neuroscience methods. 419, 110446 (2025). https://doi.org/10.1016/j.jneumeth.2025.110446.
27. Chiu, C.-L., Clack, N.: napari: a Python Multi-Dimensional Image Viewer Platform for the Research Community. Microscopy and Microanalysis. 28, S1, 1576–1577 (2022).
28. Lv, W. et al.: RT-DETRv2: Improved Baseline with Bag-of-Freebies for Real-Time Detection Transformer, (2024). https://doi.org/10.48550/arxiv.2407.17140.
29. Wang, S., Xia, C., Lv, F., Shi, Y.: RT-DETRv3: Real-time End-to-End Object Detection with Hierarchical Dense Positive Supervision. arXiv:2409.08475 (2024).
30. Wolf, T. et al.: Transformers: State-of-the-Art Natural Language Processing. In: Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pp. 38–45. Association for Computational Linguistics (2020).
31. Siméoni, O. et al.: DINOv3, (2025). https://doi.org/10.48550/arxiv.2508.10104.
32. Milletari, F., Navab, N., Ahmadi, S.-A.: V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. In: 2016 Fourth International Conference on 3D Vision (3DV), pp. 565–571. IEEE (2016). https://doi.org/10.1109/3DV.2016.79
33. Dice, L.R.: Measures of the Amount of Ecologic Association Between Species. Ecology. 26, 3, 297–302 (1945). https://doi.org/10.2307/1932409
34. Kirillov, A., He, K., Girshick, R., Rother, C., Dollár, P.: Panoptic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 9404–9413. IEEE (2019). https://doi.org/10.1109/CVPR.2019.00963