Using Computer Vision to Annotate and Count Dense 0-2mm Oyster Seeds

Fall 2025 DeVision Oyster Team: Maxwell Shroyer, Emma Kate Conner, Caleb Garza

Department of Mathematics, Louisiana State University,
Baton Rouge Advisors: Dr. Nadejda Drenska, Dr. Peter Wolenski

Continuation of Past Work

Over the past two years, the DeVision Oyster Team has been working on a project to automate the counting of oyster seeds for Louisiana Sea Grant Research Lab. We have developed models to count seeds from 2-4mm and 4-6mm with great success using a StarDist-based model. However, our previous attempt at a 0-2mm model only achieved 75% accuracy, which drastically lagged behind the 90-95% accuracy we achieved with the 2-4mm and 4-6mm models. Therefore, the focus for this semester was to rework the 0-2mm model to achieve a more satisfactory accuracy.

Model	2–4mm Acc% (MAE)	4–6mm Acc% (MAE)	4mm Acc% (MAE)	0–2mm Acc%
4-6mm_12_epochs_300	87.31% (10.12)	**98.48% (1.0)**	**93.34% (2.0)**	-
4-6mm_27_epochs_500	88.22% (9.62)	93.94% (4.0)	91.67% (2.4)	-
oyster-model-2-4mm-500	**93.24% (5.12)**	87.88% (8.0)	82.45% (5.8)	-
oyster-model-0-2mm (previous)	-	-	-	75%

MAE = Mean Absolute Error (average absolute deviation in shell counts)

StarDist

StarDist is a deep learning method for instance segmentation of star-convex shaped objects. It uses a CNN (convolutional neural network) based on U-Net architecture, trained on images and their corresponding instance masks.

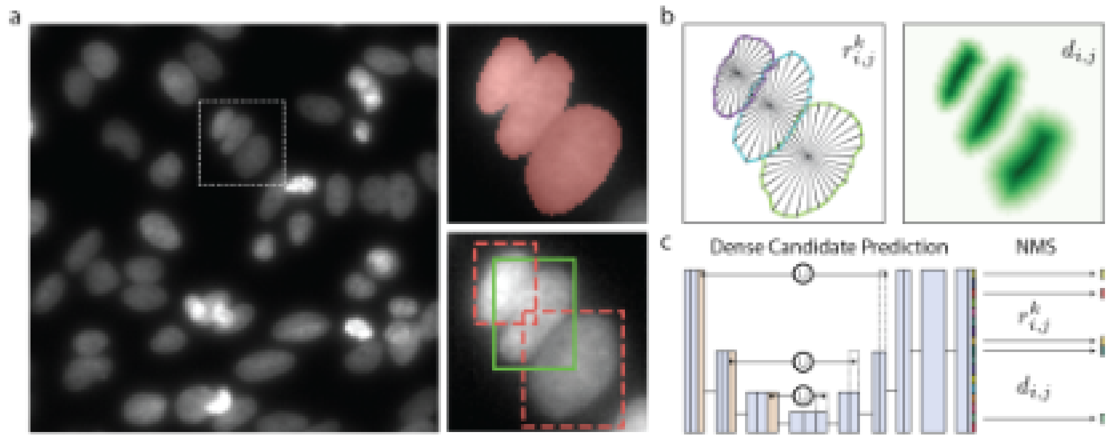

Figure 1. Stardist Architecture High Level Diagram

YOLOv10s

YOLO (You Only Look Once) is a family of real-time object detection models. We selected YOLOv10s for its small model size, making it suitable for edge deployment scenarios such as the Raspberry Pi. To enable efficient inference on resource-constrained hardware, we apply knowledge distillation techniques to further compress the model while maintaining detection accuracy for counting oyster seeds.

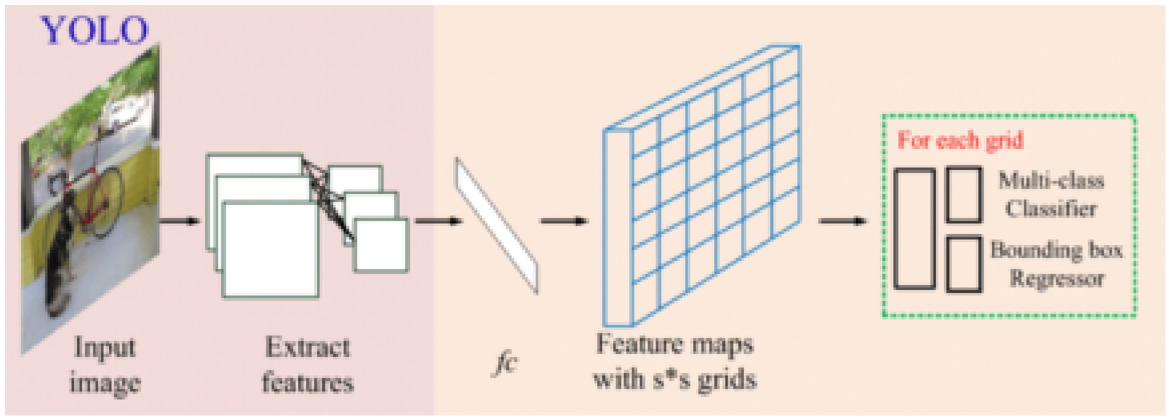

Figure 2. YOLO Architecture High Level Diagram

Dataset

One of the main issues with our first attempt at the 0-2mm model was the lack of a sizable and accurate dataset. To address this, we manually annotated a new set of 50 high-quality images, then used data augmentation techniques to generate synthetic images, resulting in a total dataset of 200 images. This expanded dataset provided significantly more training examples, allowing the model to learn more robust features and improving its overall performance.

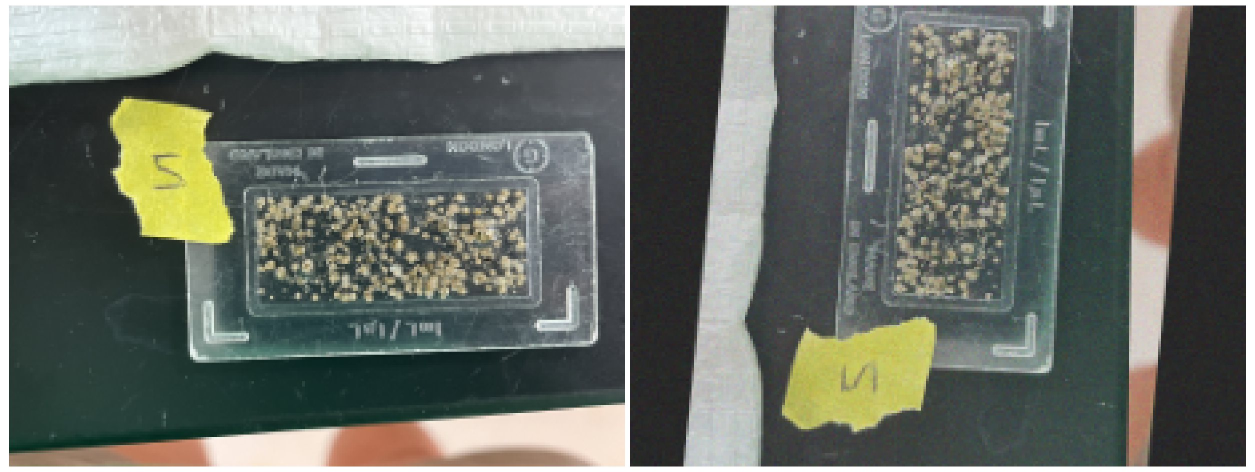

(a) Original annotated image (b) Augmented Image

Figure 3. Example of data augmentation applied to 0-2mm oyster seed images

Model Comparison Results

We evaluated both StarDist and YOLOv10s models on a test set of image. From visual inspection The 9-16-25 images were captured at close range with slight motion blur, while the 5-13-24 images were captured from a greater distance with sharper focus. StarDist was trained on a oyster seed dataset with minimal augmentation, as augmentations were found to decrease accuracy for this size class so this may have influenced its performance.

Table 1. Comparison of StarDist and YOLOv10s predictions on test images

Image	Actual Count	StarDist Pred.	YOLOv10s Pred.	StarDist Error	YOLOv10s Error
Sample19-16-25	118	62	117	-56	-1
Sample2_9-16-25	128	99	136	-29	+8
Sample3_9-16-25	161	128	184	-33	+23
Sample4_9-16-25	138	116	151	-22	+13
Sample5_9-16-25	186	153	214	-33	+28
Sample1_5-13-24	254	112	70	-142	-184
Sample2_5-13-24	275	252	179	-23	-96
Sample3_5-13-24	221	201	145	-20	-76
Sample4_5-13-24	249	235	187	-14	-62
MAE				41.3	54.6
Accuracy				77.4%	76.4%

Both models achieved comparable overall accuracy (StarDist: 77.4%, YOLOv10s: 76.4%), but with complementary strengths. **YOLOv10s** excelled on close-range images with motion blur, achieving near-perfect counts (mean error ±14.6 seeds). **StarDist** performed better on distant images with smaller, densely packed seeds (errors as low as -14), as its star-convex predictions handle overlapping circular objects well. These complementary patterns highlight the need for future research into a unified approaches.

Model Comparison Results Cont.

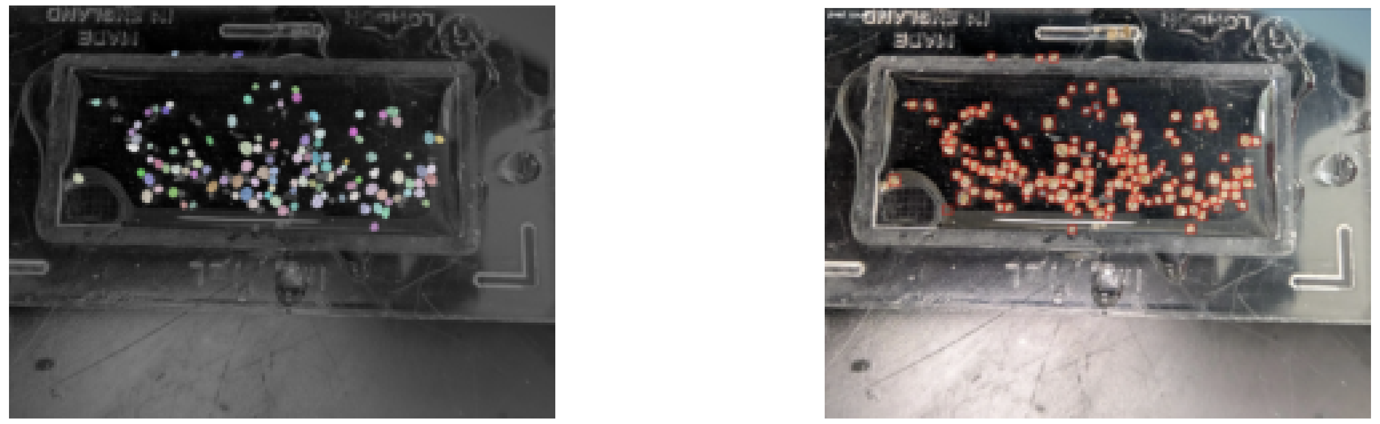

(a) StarDist (predicted count: 116) (b) YOLOv10 (predicted count: 151)

Figure 4. Model outputs on Sample4_9-16-25

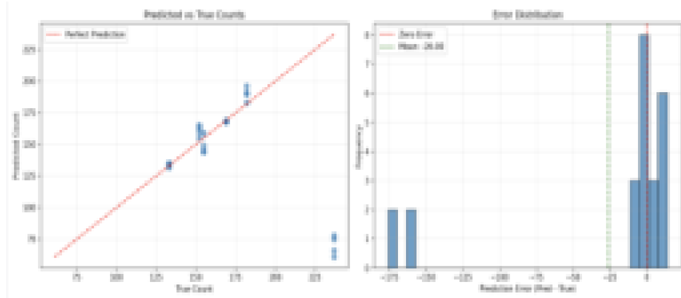

Figure 5. YOLOv10s detection results across test images

Camera distance significantly impacts accuracy: YOLOv10s achieved near-perfect counts on close-range images (Sample1: 117 predicted vs 118 actual) but severely undercounted on distant 5-13-24 images where seeds occupy fewer pixels we suspect this is because of downsampling to 640px input resolution.

Future Work

- **Improve model accuracy**: Continue training and refining both StarDist and YOLOv10s models to reduce systematic undercounting bias and achieve accuracy levels comparable to the 2-4mm and 4-6mm models (95+%)
- **Address domain shift**: Develop strategies to improve model generalization between iPhone training images and Raspberry Pi camera inference, potentially through additional augmentation techniques or collecting more training data from the deployment camera

Acknowledgments

We would like to thank the LSU Department of Mathematics, Dr. Peter Wolenski, Dr. Nadejda Drenska, Gowri Priya Sunkara, Maganizo Kapita, and the Louisiana Sea Grant

References

[1] M. Weigert, U. Schmidt, R. Haase, K. Sugawara, and G. Myers, "Star-convex polyhedra for 3d object detection and segmentation in microscopy," *IEEE (WACV)*, 2020.

[2] A. Wang, H. Chen, L. Liu, K. Chen, Z. Lin, J. Han, and G. Ding, "Yolov10: Real-time end-to-end object detection," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. arXiv:2405.14458.

[3] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2015*, vol. 9351 of *Lecture Notes in Computer Science*, pp. 234–241, 2015.

[4] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," *arXiv preprint arXiv:1503.02531*, 2015. NIPS 2014 Deep Learning Workshop.