Finding Berries: Segmentation and Counting of Cranberries using Point Supervision and Shape Priors Supplementary Material

1. Huber Loss

The loss function used for convexity, circularity, and count are modified versions of the Huber loss [1]. This loss has the general form of

$$H(R) = \begin{cases} \frac{1}{2}(R)^2, & \text{for } |R| \leq k \\ k|R| - \frac{1}{2}k, & \text{for } |R| > k. \end{cases} \tag{1}$$

For small residuals R, this loss acts as L_2 loss, and for large residuals it acts as L_1 loss. We can find the gradient of this loss by differentiating with respect to the residual

$$\frac{dH}{dR} = \begin{cases} R, & \text{for } |R| \leq k \\ k * sgn(R), & \text{for } |R| > k \end{cases} = \begin{cases} -k, & \text{for } R < -k \\ R, & \text{for } -k \leq R \leq k \\ k, & \text{for } k < R. \end{cases} \tag{2}$$

2. Additional Qualitative Results

Here we present additional visual comparisons between our method and LC-FCN [2] on CRAID.

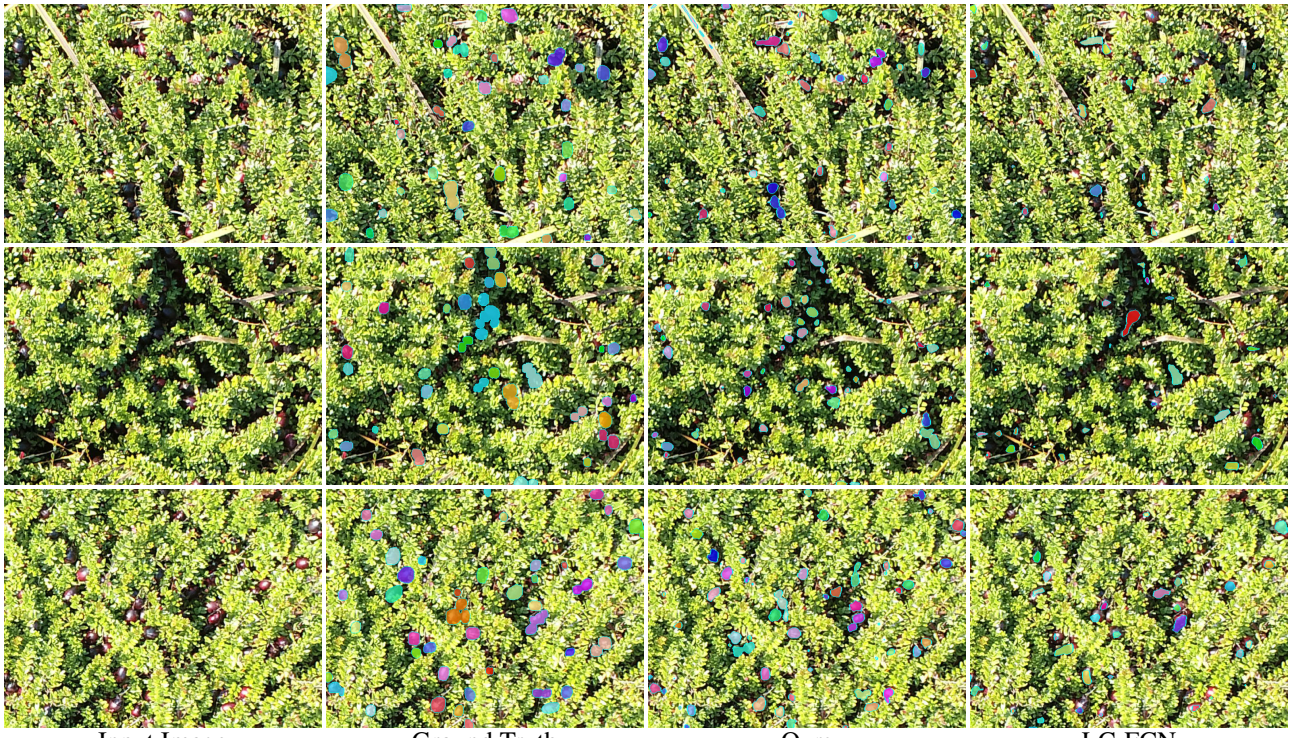

Input Image | Ground Truth | Ours | LC-FCN

Figure 1: Qualitative comparison with SOTA methods on CRAID. Our method ($\mathcal{L}_{Seg}$ + $\mathcal{L}_{Split}$ + $\mathcal{L}_{Convex}$) shows that using shape priors and better boundary and region selection allows robust segmentation and counting performance. Colors in prediction masks are random and are used to represent instances. Colors may repeat. Best viewed in color and zoomed.

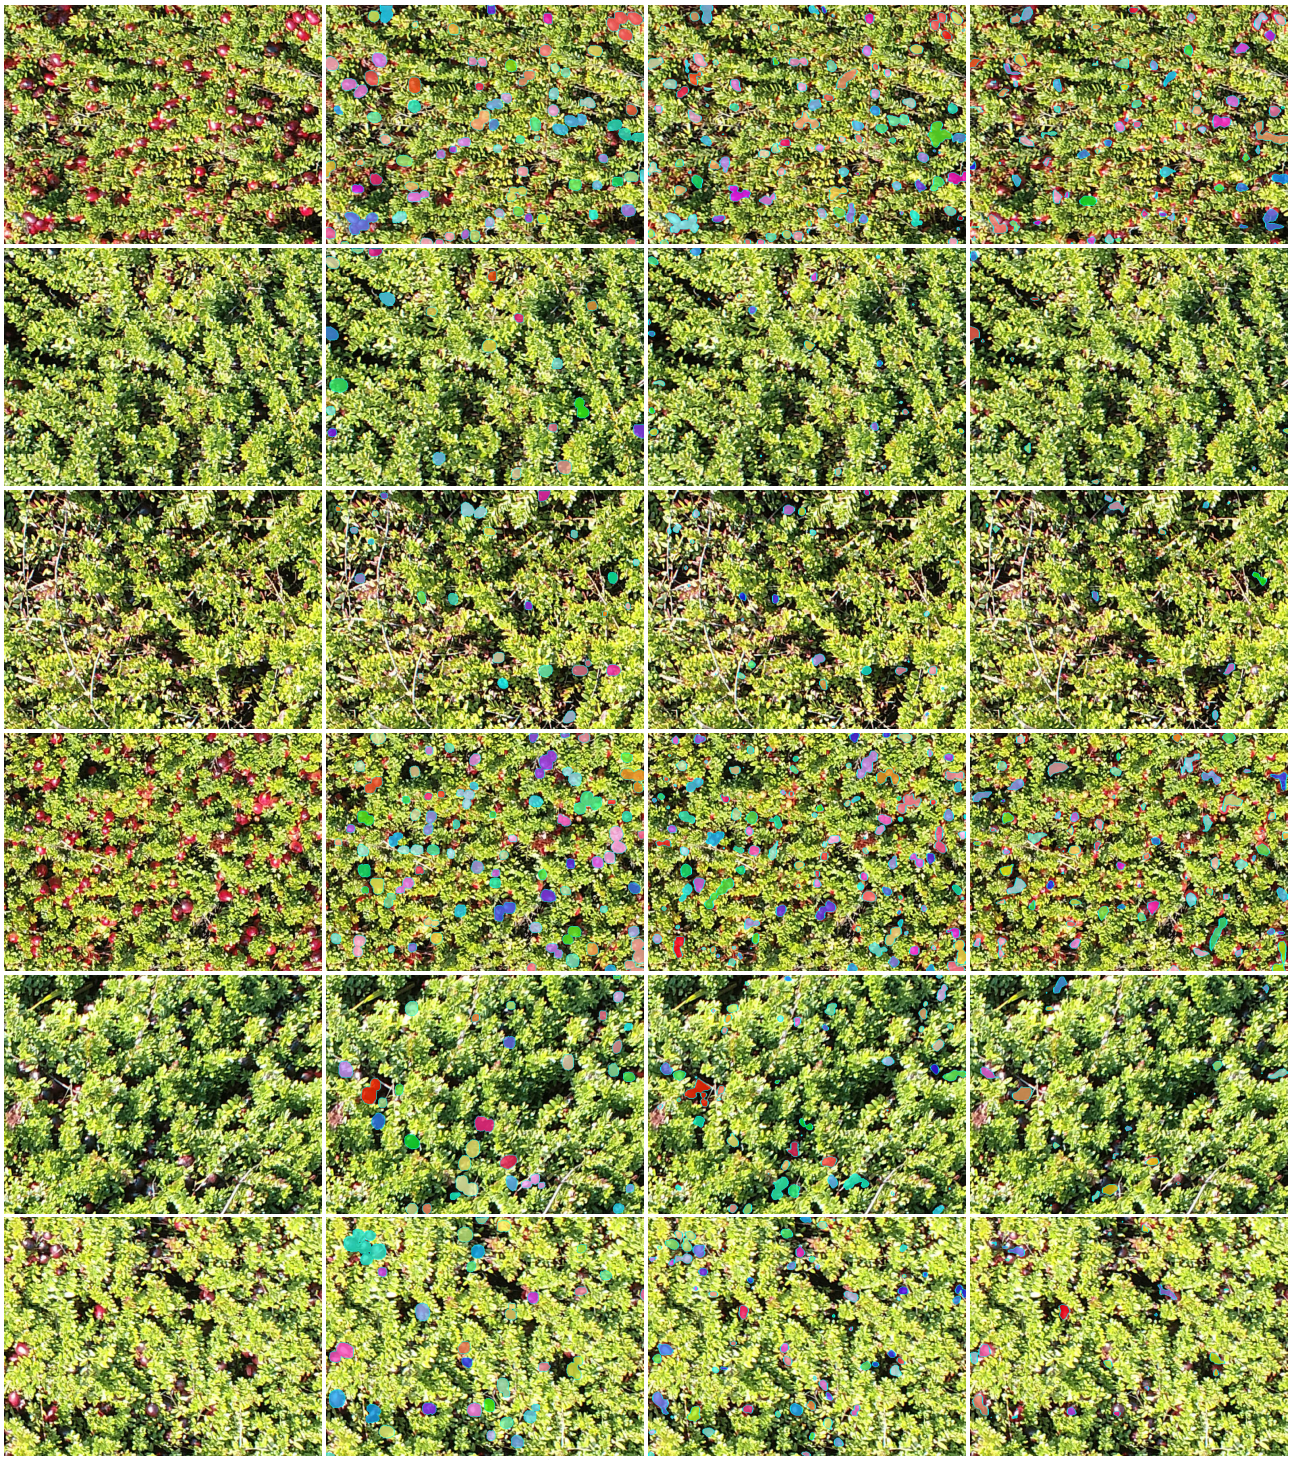

Figure 2: Qualitative comparison with SOTA methods on CRAID. Our method ($\mathcal{L}_{Seg}$ + $\mathcal{L}_{Split}$ + $\mathcal{L}_{Convex}$) shows that using shape priors and better boundary and region selection allows robust segmentation and counting performance. Colors in prediction masks are random and are used to represent instances. Colors may repeat. Best viewed in color and zoomed.

References

[1] Peter J Huber. Robust estimation of a location parameter. In *Breakthroughs in statistics*, pages 492–518. Springer, 1992. 1

[2] Issam H Laradji, Negar Rostamzadeh, Pedro O Pinheiro, David Vazquez, and Mark Schmidt. Where are the blobs: Counting by localization with point supervision. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 547–562, 2018. 1