

Figure 1: Example of an OrthoSAM segmentation of two SediNet images. We provide a few examples and their code and parameters for selected SediNet images on our GitHub repository.

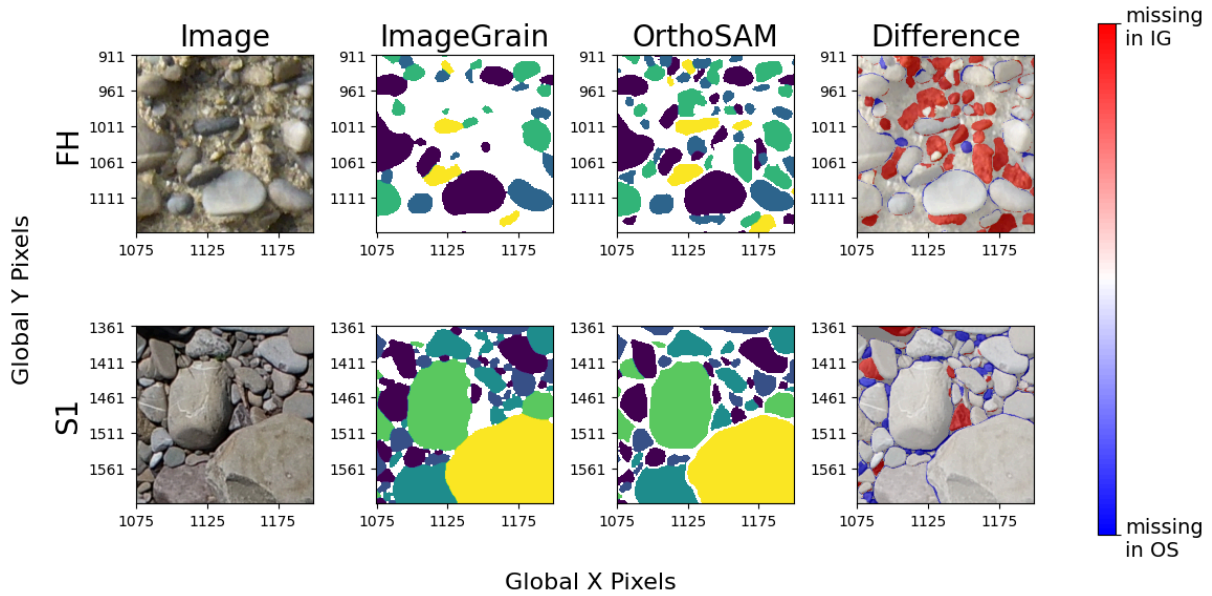


Figure 2: Comparison of OrthoSAM (OS) and ImageGrains (IG) predictions for center crops of image FH and S1. We note that the number of detected objects varies between the two models for each dataset. For example, for image S1, the OrthoSAM prediction identified 4316 objects, while ImageGrains identified 4660 objects. This ratio of 0.9 ($4316/4660$) is different from image FH ($2819/1742=1.6$). This does not allow a precision and accuracy assessment with the same training data. The difference plot visualizes the agreement and disagreement between two predictions. Red regions highlight areas where OrthoSAM identifies an object, and ImageGrains does not. While blue regions highlight areas where ImageGrains identifies an object, and OrthoSAM does not. Both examples demonstrated OrthoSAM’s capability in fine object segmentation. However, due to the lack of a classification component, OrthoSAM has the inherent limitation that irrelevant objects may remain in the segmentation results. In particular, we see patches of sand that were falsely segmented in FH. Here, we see that lower resolution or blurriness in the image can exaggerate the issue, resulting in more false positives.