



Zero-Shot Segmentation of Estuary Mudflats Using the Segment Anything Model

Jaren Unzen

Computer Science, University of New Hampshire, Durham, NH 03824

Abstract

The Problem: While **mudflats** are vital for flood defense and carbon sequestration, their rapid disappearance goes unmonitored because traditional manual data annotation is too expensive and labor-intensive to scale.

The Approach: We evaluate the **Segment Anything Model 3 (SAM3)** foundation model, utilizing its **zero-shot** capabilities to automate **mudflat segmentation** without the need for additional training data or manual labeling.

Why it Matters: By bypassing traditional resource bottlenecks, this method enables the **low-cost, high-frequency monitoring** essential for protecting these critical carbon sinks, ecological habitats, and natural flood barriers.

Dataset

Table 1: Various attributes of the 200 image evaluation data set.

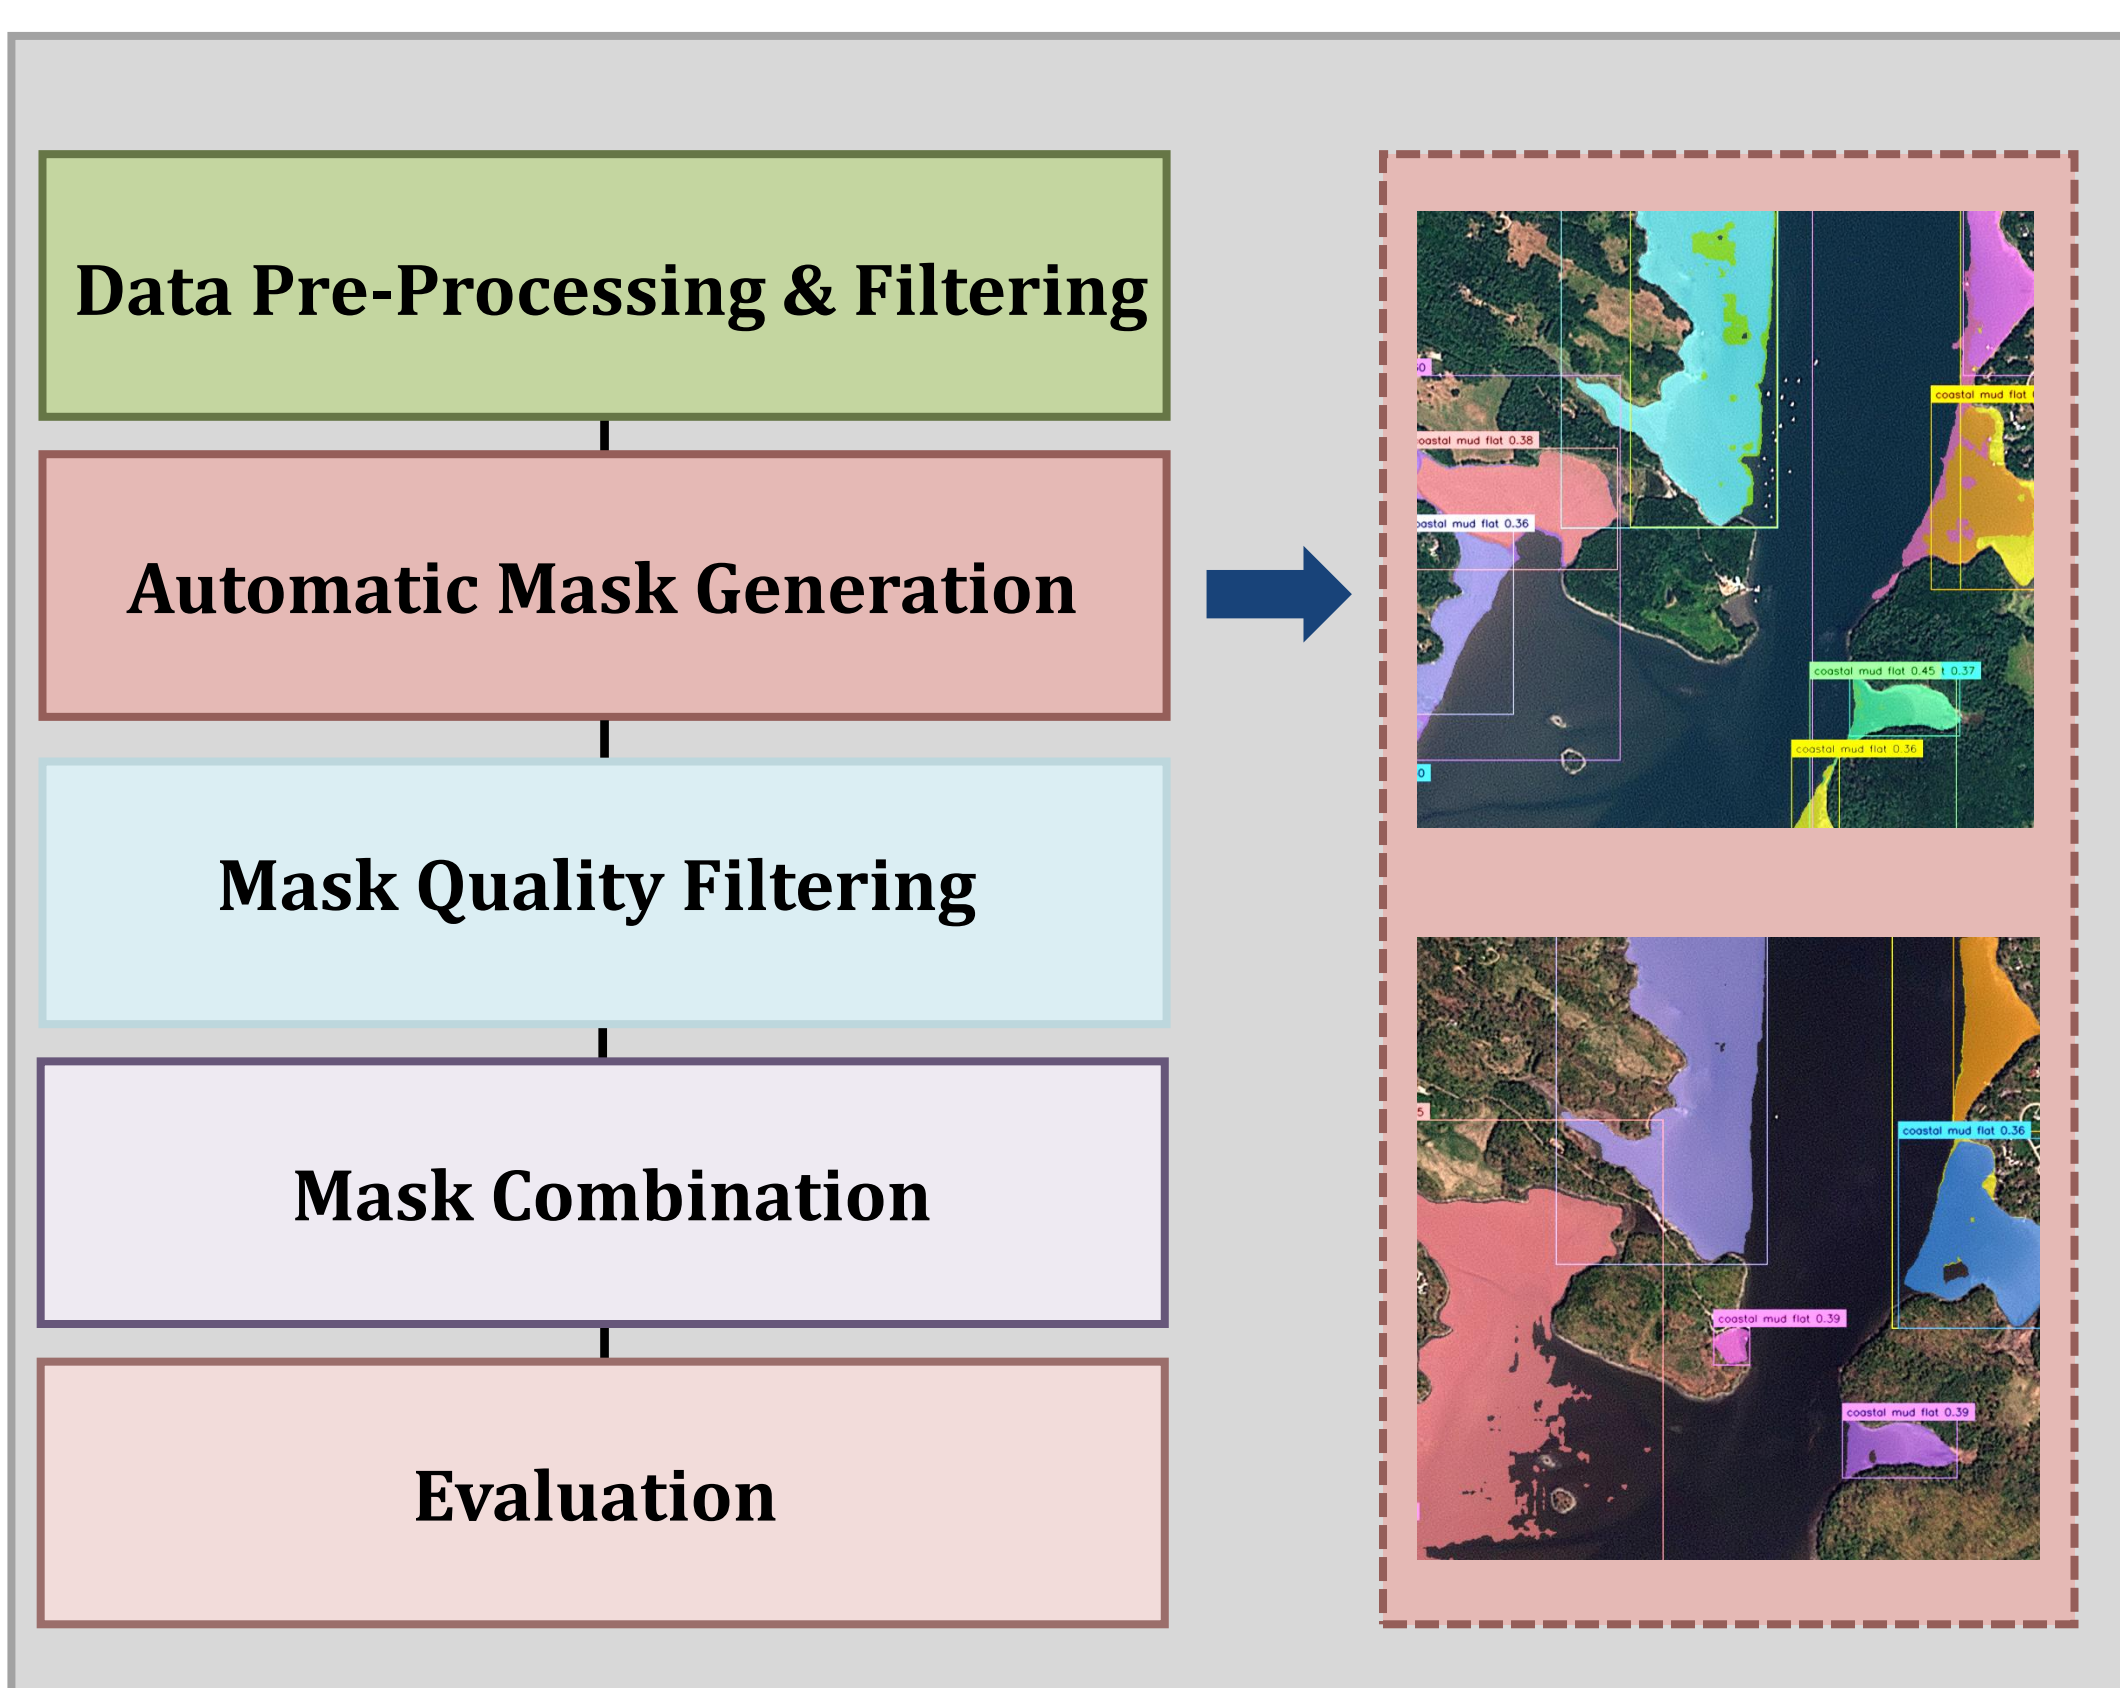
| Study Site | Filter Date | Image Bands | Resolution | Spatial Res. | Cloud Cover |
|-------------------|--------------------------|-------------|-------------|--------------|-------------|
| Great Bay Estuary | 2024-04-25 to 2025-08-13 | RGB | 1024 x 1024 | 3.7 m/px | <20% |

Pipeline

Data Source

planet


We applied various pre- and post-processing techniques to the satellite imagery to improve visual clarity and exclude low-quality masks.



SAM 3

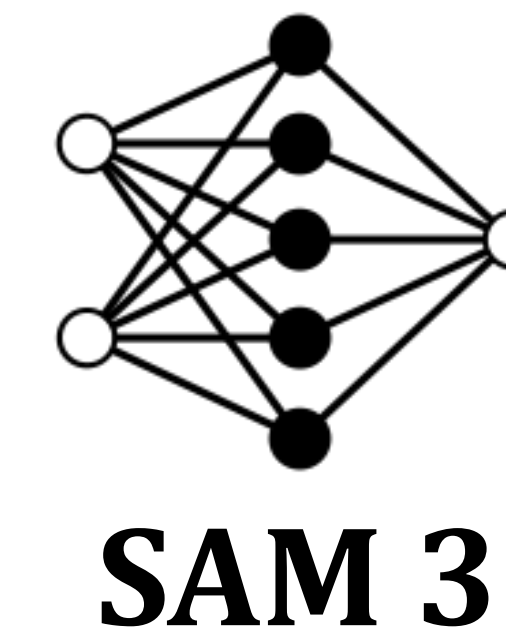
Segmentation

Input

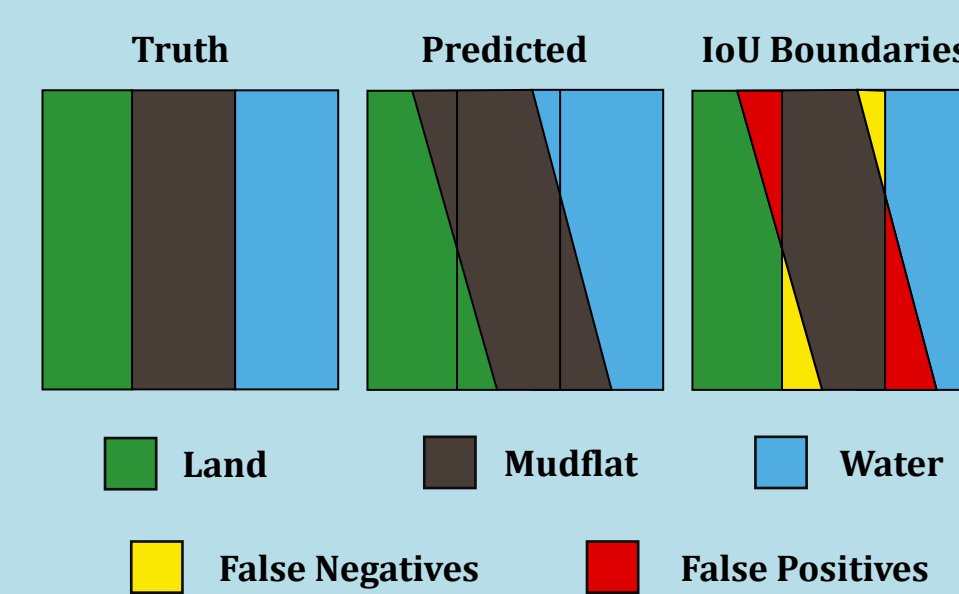


Text Prompt:


mud flat



Segmentation



MASK GENERATION

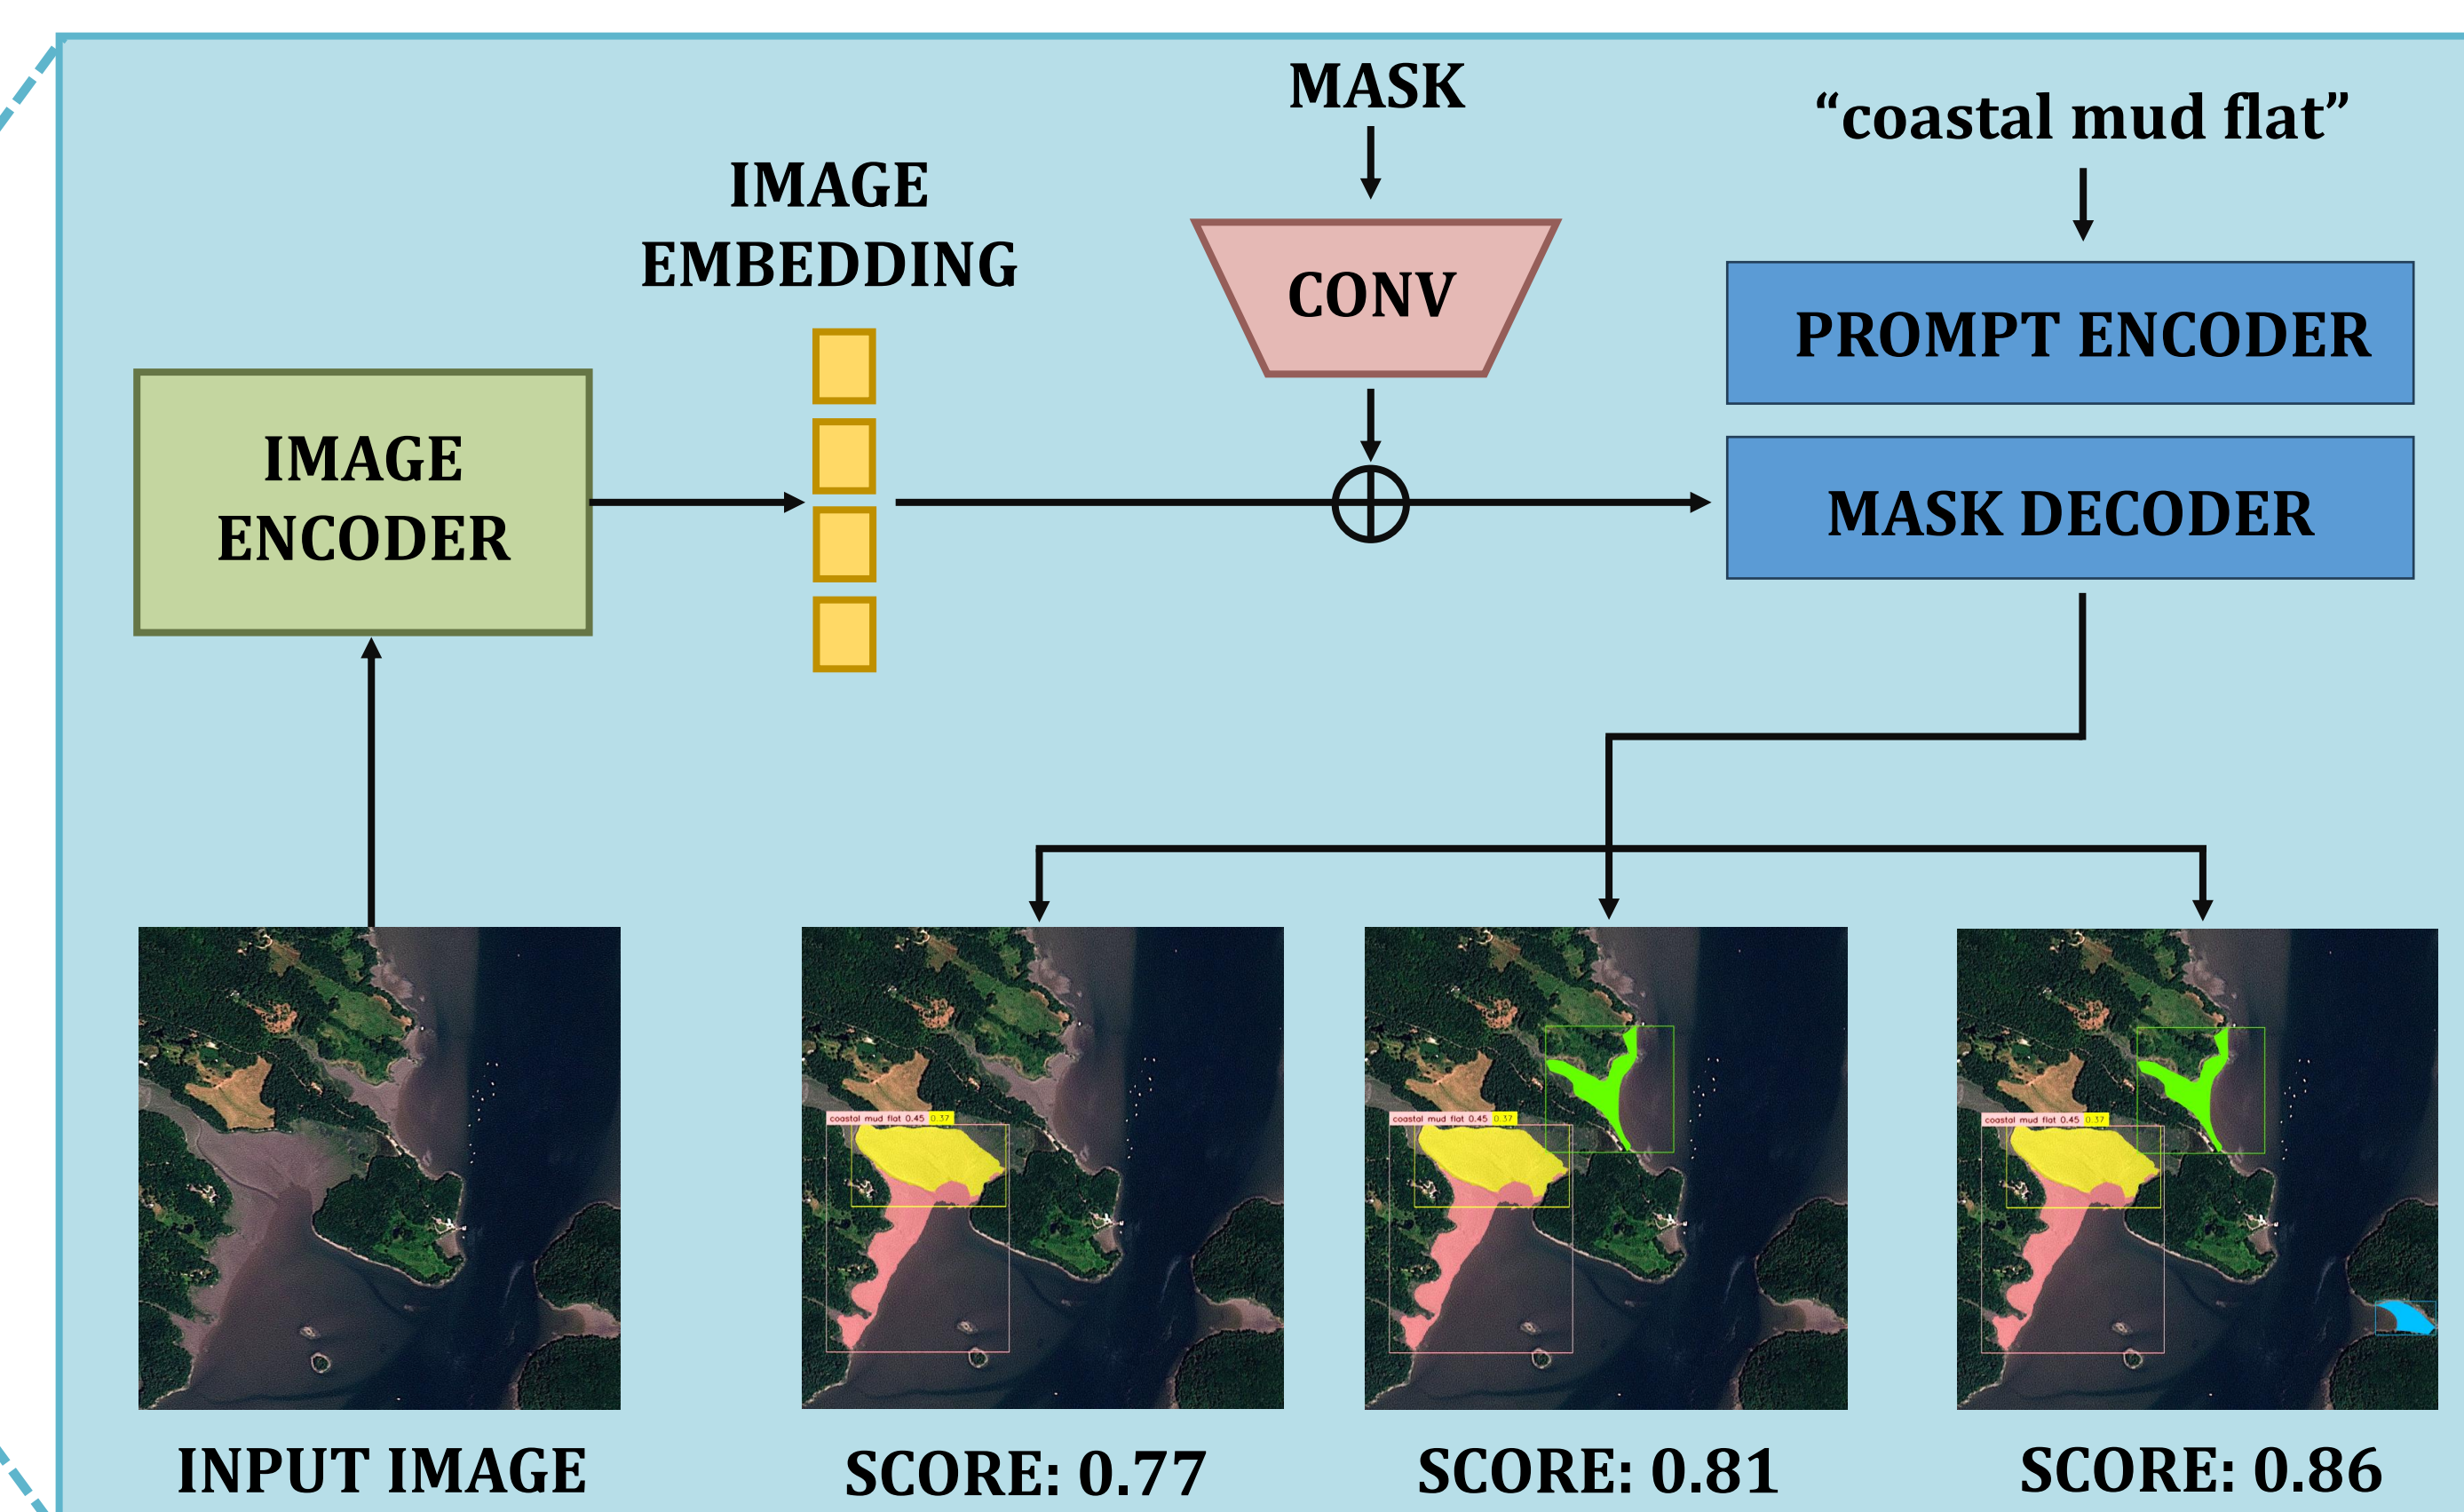


MASK FILTERING & COMBINATION

SAM3 produces multiple masks for each prompt, which we filter and combine into one high-quality result.



Architecture

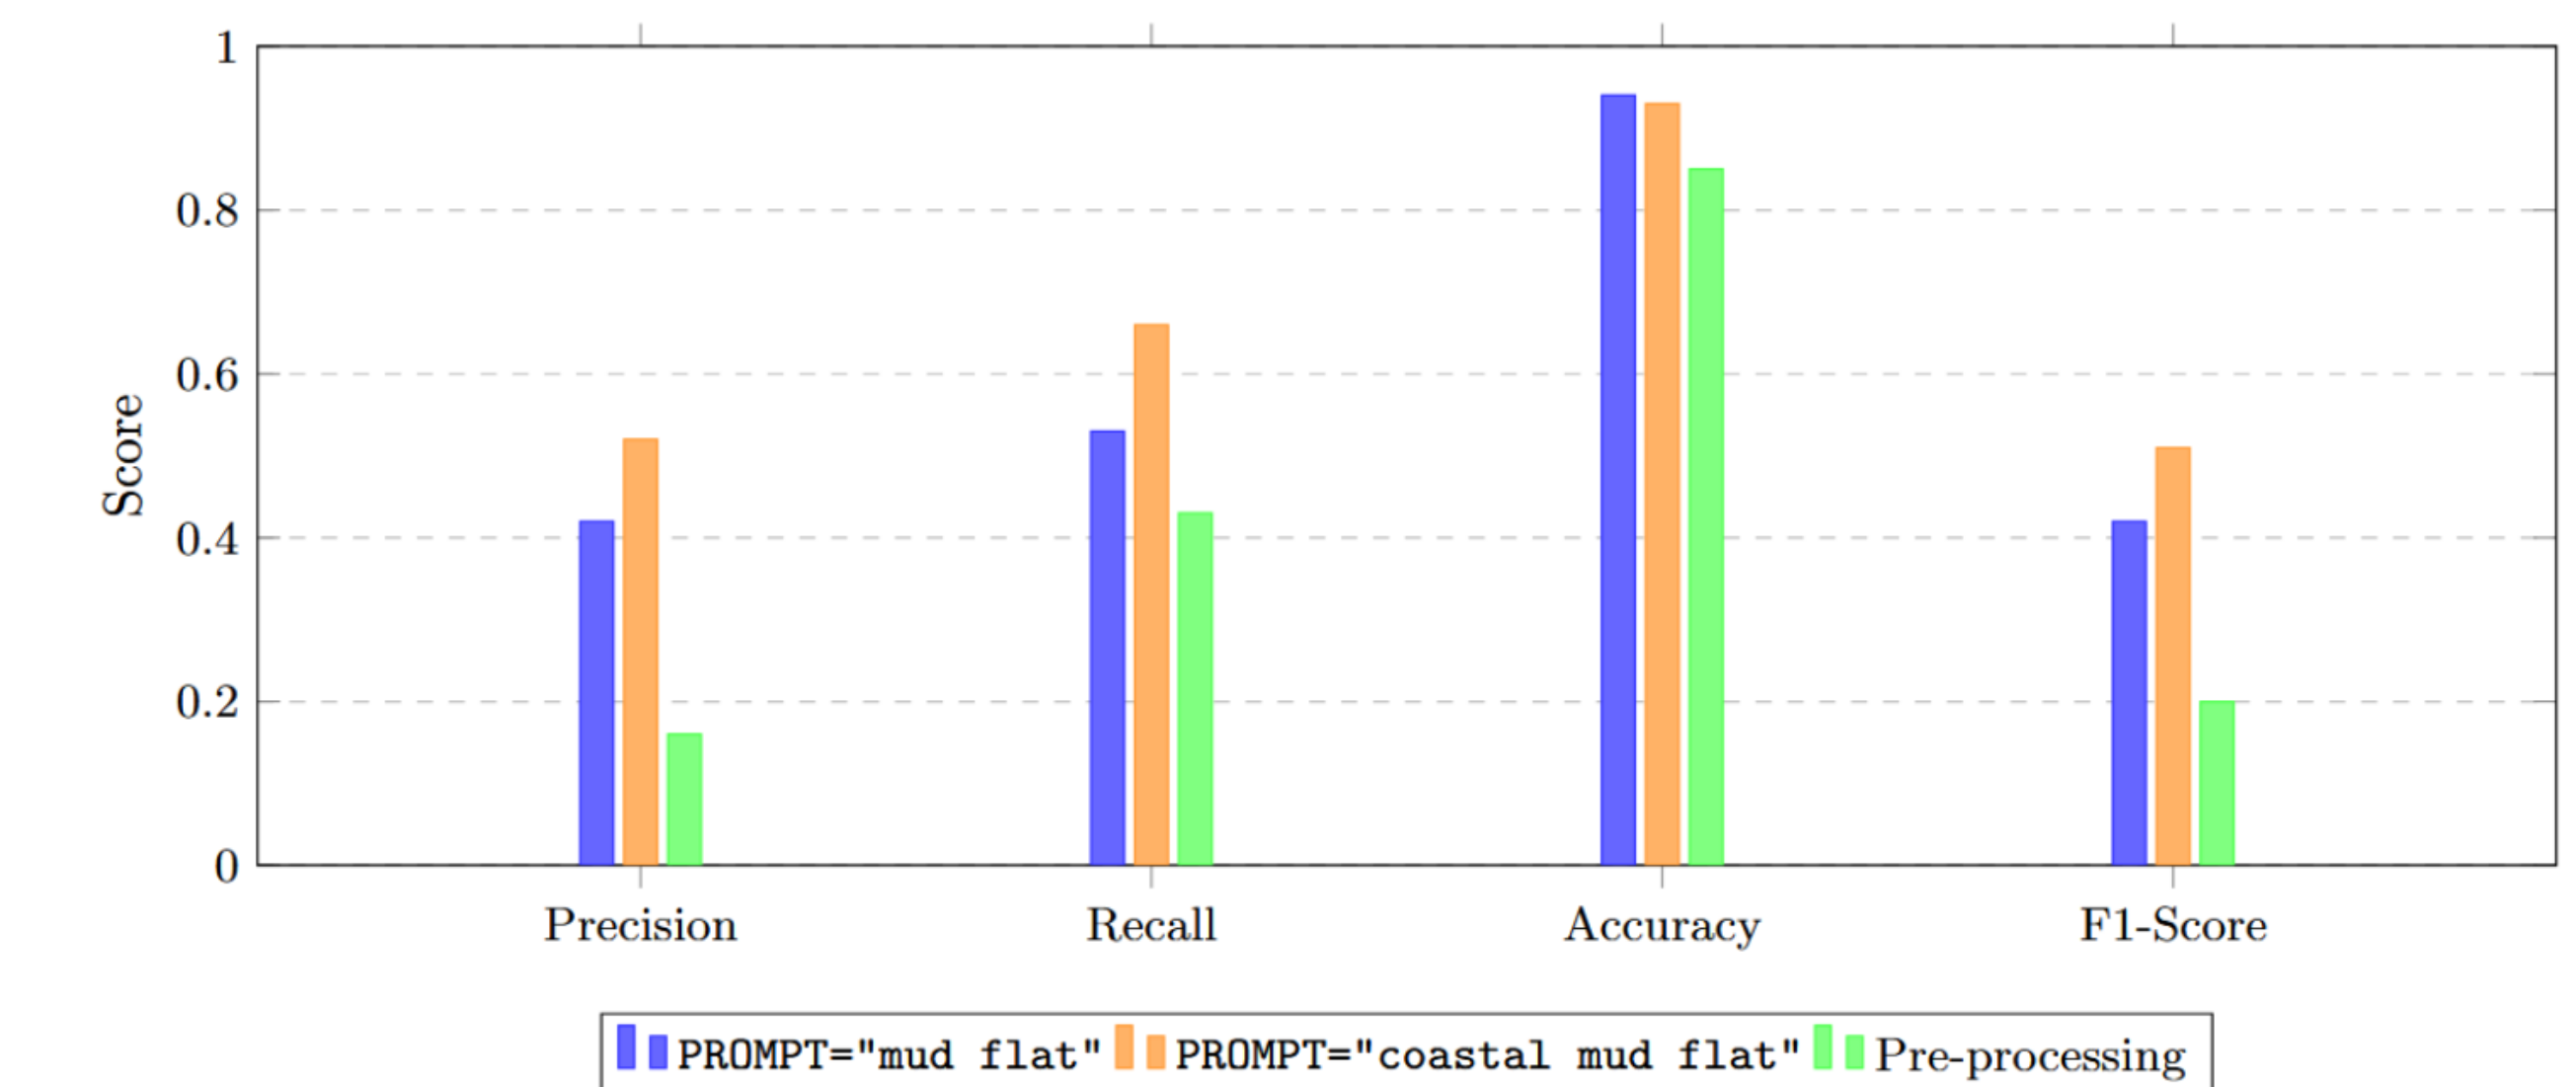


Experiment Results

Table 2: Various metrics for evaluating segmentation quality across our dataset.

| SAM3 Mud Flat Segmentation Results | | | | |
|------------------------------------|-----------|--------|----------|-------------|
| Method | Precision | Recall | Accuracy | F1-Score |
| PROMPT="coastal mud flat" | 0.52 | 0.66 | 0.93 | 0.51 |
| PROMPT="mud flat" | 0.42 | 0.53 | 0.94 | 0.42 |
| Pre-processing | 0.17 | 0.43 | 0.85 | 0.2 |

Figure 1: Grouped bar chart comparing segmentation metrics across methods.



HIGH QUALITY MASK



POOR QUALITY MASK



IMAGE TRUTH MASK

Conclusion

This study demonstrates that **Segment Anything 3** provides a viable, low-resource alternative to traditional supervised learning for mud flat monitoring. By leveraging a zero-shot framework, we successfully bypassed the need for expensive data annotation while maintaining a baseline **F1-score of 0.51**. While similarities between mud and sediment remain a challenge, this study shows that **foundational models can be effectively adapted for niche remote sensing tasks**. Future work will apply these strategies to other land cover types and explore fine-tuning strategies for improved performance.