KLP: Kelp Location Profiling Using Segmentation of Aerial Drone Imagery for Assessing Percent Kelp-Canopy Coverage

Tim Sananikone

Mentors: Danelle Cline, Duane Edgington, Steve Haddock

Summer 2025

Keywords: Segment Anything Model (SAM), UAV Aerial Drone, Kelp-Canopy Coverage

ABSTRACT

Kelp forests are an integral part of the marine ecosystem, providing habitat and food for more than 1,000 species. In recent years, due to environmental changes such as marine heat-waves and explosions in urchin populations, there has been a substantial decline in kelp-forest health around the world. Effective monitoring is integral to ensuring conservation of these habitats, and uncrewed aerial vehicles (UAVs) offer a cost-effective monitoring tool. Previous work done to calculate accurate kelp coverage relied upon multi-spectral cameras, specifically the NIR color band, to generate their results. This study introduces an automated deep-learning pipeline that calculates kelp canopy coverage using only standard RGB imagery. Our workflow applies the Segment Anything Model (SAM) to generate segmentation over subsetted slices of the original 42-megapixel images. SAM was seeded by using color thresholds to identify regions of water, which were usually more uniform and contiguous than the patches of kelp. Results have been qualitatively investigated and compared to manually annotated, ground-truth datasets representing the minimum and maximum potential coverage in a photo.

1 INTRODUCTION

1.1 Importance of Kelp

Kelp is a foundational species, with kelp forests supporting more than 1,000 different animals [U.S. National Park Service, 2019]. These forests are central to the life cycles of many fish, such as the Pacific Herring (*Clupea pallasii*). Each year, herring schools migrate to lay their eggs on kelp fronds, fueling both the local ecosystem with a nutrient-rich food source and the broader oceanic ecosystem through a new generation of herring. The loss of kelp forests would rob these herring of a habitat critical to their reproductive cycle, impacting a key source of food for many other predators. The collapse of a kelp forest would not result in the loss of a single habitat, but instead have a cascading effect on the entire marine ecosystem.

1.2 Declining Kelp

In recent years, primarily due to changing ecological factors, the overall health of kelp forests is at risk. Events such as marine heatwaves and booming urchin populations have accelerated this process. It is estimated that Northern California has lost as much as 96% of its Bull Kelp in less than a decade [Zuckerman, 2023]. Thus, developing papers that are both accessible and coss-effective is critical.

1.3 Monitoring Problems

Methods for monitoring kelp coverage often have to balance between scale and resolution. Satellites offer a large-scale solution with large fields of view while sacrificing the fine nuances of the kelp or small-scale changes within the data [Cavanaugh et al., 2010]. On the other side of the spectrum, high-fidelity methods, such as underwater surveys conducted by organizations such as Reef Check, provide a much more detailed assessment of the health of the environment. This process, however, is time intensive requiring divers to be actively present in these areas of interest [Piñeiro Corbeira et al., 2023]. UAVs represent an effective middle ground, offering both high spatial resolution and the ability to survey entire sites cost-effectively. Previous work generating kelp coverage calculations from UAV relied upon access to multi-spectral color bands, particularly the NIR band [Cavanaugh et al., 2021], while this study demonstrates that using standard RGB imagery can also be used to effectively monitor kelp coverage.

2 MATERIALS AND METHODS

2.1 Study Area

Images used for this study were captured from 6 different sites around the Monterey Bay Area from 2024-2025, chosen due to their biological diversity. Flight times were chosen to minimize environmental interference such as sun glint.

Figure 1: *Example UAV image taken from Seymour survey site*

2.2 Camera and Drone details

Images were captured using a Sony Cyber-shot DSC-RX1R II 42-megapixel camera recording standard RGB color channels. High-resolution 8k images (7952x5304 pixels) were captured at a flight altitude of 60m, providing a resolution of approximately 1cm per-pixel. The images were taken using a Trinity F90+ VTOL UAV. Flight paths were designed to provide comprehensive coverage of sites whilst still maintaining line of sight with the drone operator.

Figure 2: *Example UAV image taken from Seymour survey site*

2.3 Segment Anything Model (SAM)

The core of the pipeline utilizes the Segment Anything Model (SAM). A foundational deep learning model, SAM was designed to generate high-fidelity segmentation masks of objects from images [Kirillov et al., 2023]. Unlike traditional segmentation models, SAM generates its segmentation through input prompts, such as seed points or bounding boxes, removing the need for a large manually annotated training dataset to use the model within a new domain. The SAM model has no semantic understanding of what each object it is segmenting is, only that there is a distinct object being segmented. This characteristic makes it highly useful for the problem of kelp coverage segmentation because of the tendency of kelp to appear in highly variable shapes and sizes.

2.4 Image Pre-processing

Images were loaded directly from a Tator database through a CSV metadata file. Once loaded, the images were converted to LAB color space using OpenCV. This transformation separates lightness from color information, allowing for a more distinct representation of the pixels helping the model more effectively differentiate between water and kelp pixels. To generate a representative water color, which is required for generating the seed points needed for the SAM model, for each individual site, 50,000 pixels were randomly sampled from each image within that site. This sample size was selected as it provided a consistent color value whilst remaining computationally efficient. The median L, A, and B values from the sampled pixels are then calculated to generate a baseline water color used as a reference for segmentation. Given that within a majority of images water is the most dominant object, by sampling from each of the images, an accurate representation of the water within a certain site can be computed without looking over all pixels. This site-specific approach helps account for local variations in water color between sites and times of day.

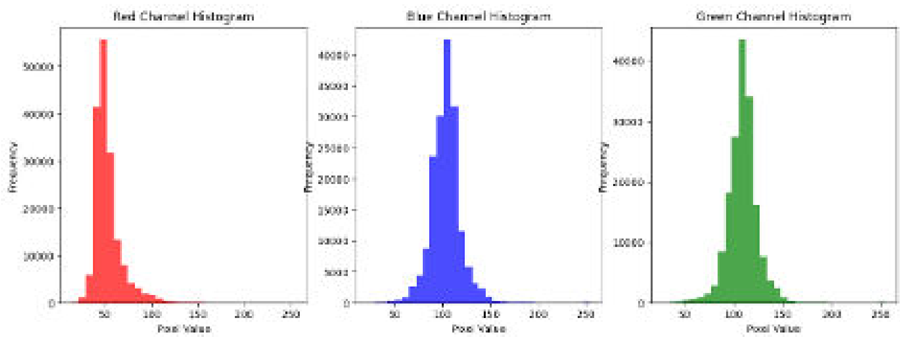

(a) *Histogram showing sampled points.*

(b) *9 of the most common pixel values.*

Figure 3: *Visualizations of the site-specific water color generation process.*

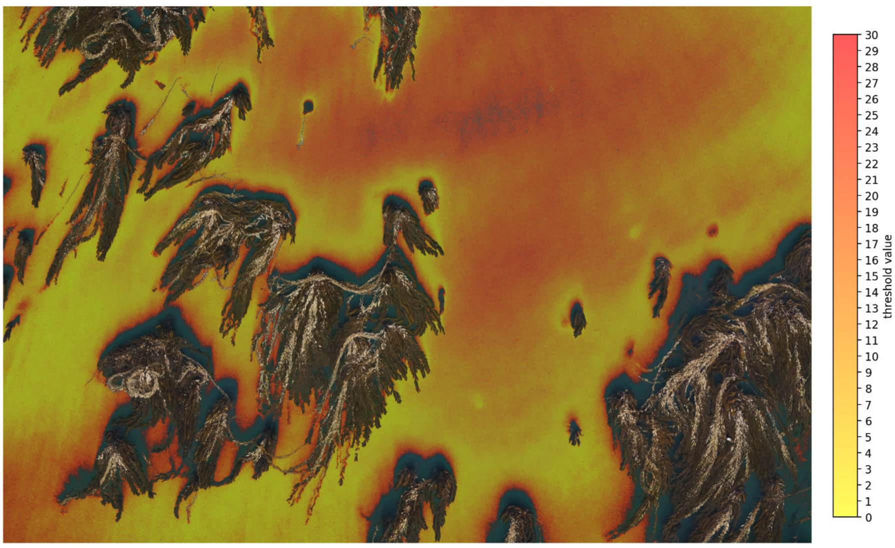

Figure 4: *Visualization of euclidean distance from pixel value to median water pixel value.*

2.5 SAHI-SAM Pipeline

The pipeline integrates SAM with the Slicing Aided Hyper Inference (SAHI) framework to effectively and quickly process large, high resolution UAV imagery [Akyon et al., 2022]. In order to reduce inference time, the mobileSAM created by Ultralytics was utilized [Zhang et al., 2023].=

2.5.1 Problem Inversion

The key step in generating segmentations was inverting the problem: instead of directly segmenting all kelp within the image, the pipeline creates a segmentation of the water and inverses the mask. In aerial images of kelp forests, the majority of each image consists of a single continuous body of water. In contrast, kelp can appear fragmented and in smaller individual patches or large patches of kelp. Thus it is significantly easier for SAM to accurately segment out one cohesive 'water' object than to directly identify and segment out many dispersed 'kelp objects'. This strategy ensures a more robust segmentation of the primary background from which the kelp canopy coverage can be calculated.

2.5.2 Image Slicing

The SAM model was trained on images of size 1024x1024. Thus in order to efficiently process high-resolution UAV images whilst also utilizing all the features learned by the SAM model through training, we employ the slicing aided hyper inference (SAHI) framework. SAHI is used to slice the large 8k image into smaller overlapping patches (0.2) of size 1024x1024. each slice is then processed independently by SAM using generated seed points with the resulting segmentation being stitched back together.

2.5.3 Automated Seed Point Generation

The quality of the segmentation thus is highly dependent on the input seed points. If seed points are clustered too closely, they may cause the SAM model to segment out small localized sections of water caused by ripples or other surface level effects. This in turn means that the larger connected water bodies are not segmented properly. Additionally, if pockets of water are contained within the image, clustering the seed points too close may lead to the pockets not being accurately segmented out. Conversely, if seed points are placed too close to Kelp, the SAM model may misidentify the object it is being asked to segment out. In order to create these seed points, multiple heuristics are applied. First in order to reduce the search space, each slice is divided into a grid of 64x64 pixel cells. Each grid cell is then analyzed for its validity based on two criteria. The thresholds for these checks were determined through an iterative process based on a qualitative assessment of the segmentation performance.

- **Color Check:** A cell was considered valid if the mean Euclidean distance between its pixels and the site's reference water color was below a threshold.

- **Uniformity Check:** To avoid noisy regions indicative of kelp edges or kelp itself, a cell was only considered valid if the standard deviation of its color values was below a threshold.

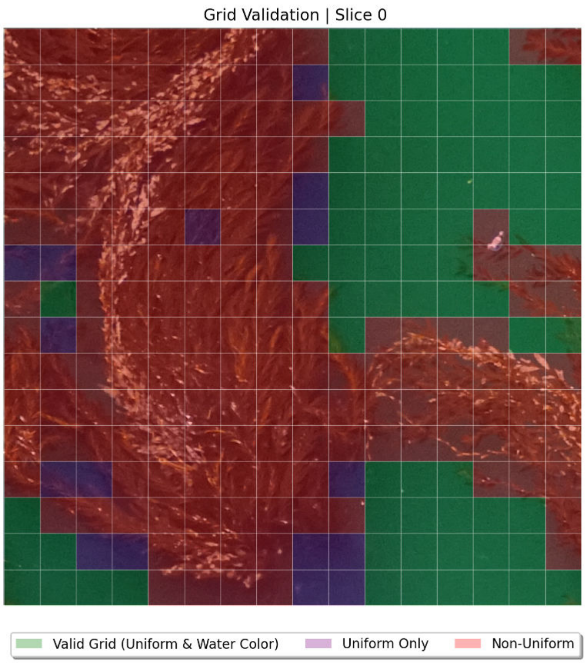

Figure 5: *Visualization of the grid-based color and uniformity check.*

If a majority of grids (98%) within a slice passed or failed the uniformity check and the color check, the slice was assumed to be entirely water or kelp and SAM inference was skipped in order to reduce time spent on inference. From the valid grid cells passing both the color and uniformity check, an initial point is sampled from within each grid., In order to ensure that the seed points are spatially disturbed and capturing the entire water region, a Poisson disk sampling strategy is used to select the final set of seed points. This method generates points that are randomly selected whilst also ensuring a minimum distance. This helps to prevent clustering of the points and increases the likelihood that disconnected water pockets within the image are included in the generated segmentation.

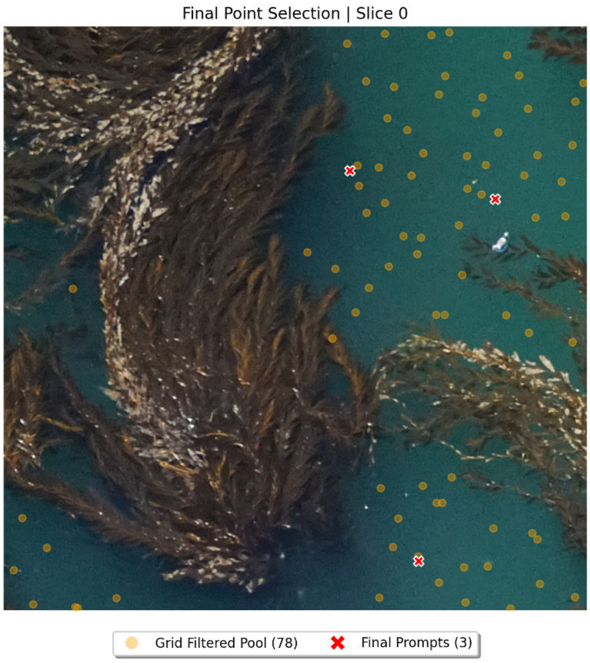

Figure 6: *Seed points capturing two distinct water pockets within a slice.*

2.5.4 Hierarchical Segmentation

While splitting the images into smaller slices of 1024x1024 helps to generate fine segmentations of the kelp, it can lead to patches including large sections of kelp to be unsegmented, leading to large gaps within the final segmentation. In order to mediate this, the coarse structure of the kelp is captured through the use of a coarse model.

1. **Coarse Pass:** An initial segmentation pass is performed on slices of size 4096x4096. This pass focuses on capturing the broad outlines of major kelp beds and large cohesive water bodies.

2. **Fine Pass:**A second pass is performed on slices of 1024x1024. This fine pass focuses on refining the segmentation, particularity at the edges of the kelp and capturing small pieces of kelp that may have been missed by the coarse pass.

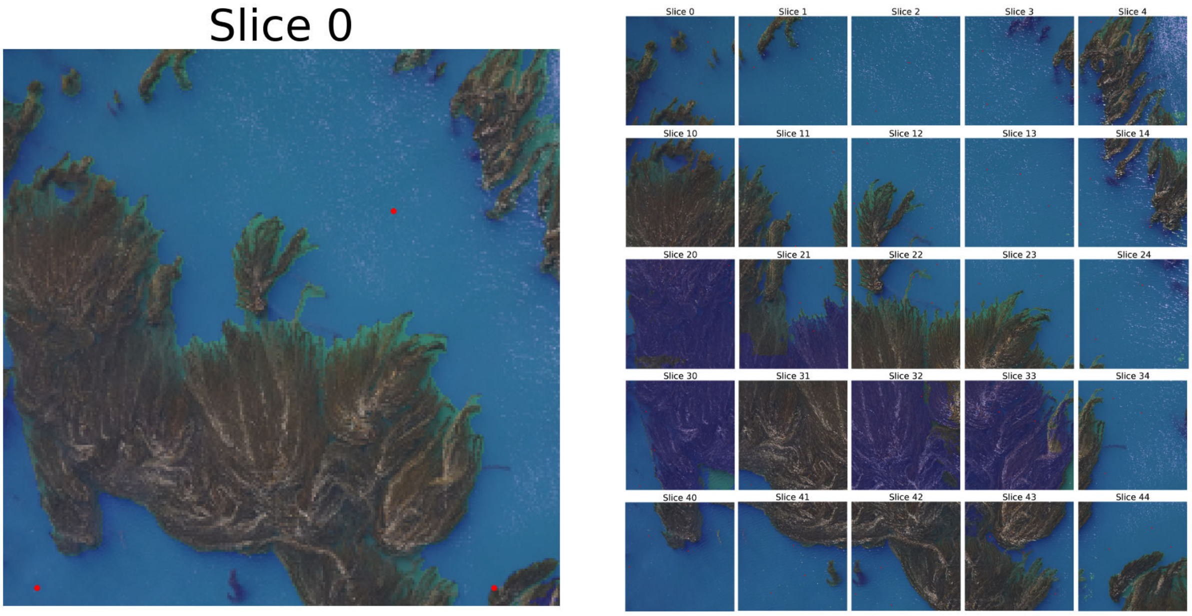

(a) *Coarse pass (4096x4096 slice).* **(b)** *Fine pass (1024x1024 slice).*

Figure 7: *A comparison of the coarse vs. fine mask over the same area.*

2.5.5 Mask Reconstruction

The binary water masks from both passes are then merged to create a high-fidelity kelp mask. These masks are reconstructed from the SAHI generated slices using a logical OR. A pixel was classified as water if any slice within the image identified the pixel as being water. Then both water masks are inverted to produce the binary kelp masks. An erosion with a 51x51 kernel was applied to the coarse mask in order to remove a segmentation artifacts introduced when the model down samples the coarse slices of size 4096x4096 to its native 1024x1024 input size. This kernel size was chosen through a qualitative analysis as it provided the best balance between removing excess points and preserving the overall shape. The final kelp mask was generated by resolving disagreements between the two masks. Pixels classified as kelp in both the coarse and fine mask were retained as kelp within the final mask. Pixels who have classifications that disagree between masks have their pixel values compared to the site-specific water color value. If the euclidean distance is below a color threshold, the pixel is classified as water. Otherwise the pixel is classified as kelp. This final check helps to ensure that ambiguous regions generated from merging the two kelp masks together are resolved.

2.6 Coverage Calculation

Kelp coverage calculation was then calculated by running the pipeline over an entire survey site. Coverage values were saved in site specific json files.

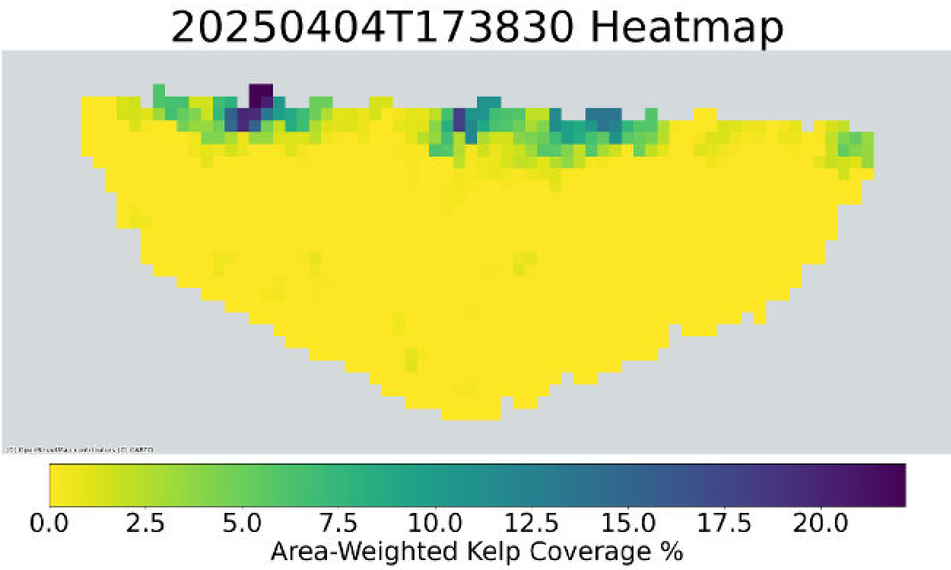

Figure 8: *Heatmap generated over an entire survey site representing kelp coverage*

3 RESULTS

TO BE CHANGED AFTER RUNNING OVER MIN / MAX THRESHOLDING

The pipeline was able to process high-resolution UAV imagery from all six sites across a number of different months generating binary segmentation masks that were then used to calculate a metric for the kelp coverage within an area. However without having access to a ground truth dataset, the viability could only be assessed qualitatively. A rigorous quantitative metric for segmentation accuracy could not be calculated.

4 DISCUSSION

TO BE CHANGED AFTER RUNNING OVER MIN / MAX THRESHOLDING

This pipeline reinforces the potential of using high resolution UAV imagery with SAM to generate kelp coverage calculations without the need for specialized multi-spectral cameras. The strength of the pipeline lies in its ability to be applied without any retraining. However, while qualitatively promising, without having access to a ground truth segmentation mask the true efficacy of this approach cannot be calculated. The evaluation of the segmentation masks was purely qualitative, based on visual comparisons with the original UAV imagery. To address this, future work should prioritize the creation of a ground truth dataset. By generating a segmentation mask by hand for UAV images and having it reviewed by a domain expert, a benchmark could be generated. The pipeline's output could then be rigorously tested using standard segmentation metrics such as the Dice coefficient or Intersection over Union. Furthermore, due to the nature of the imagery being high resolution, other color information may be utilized in order to generate more accurate readings of the kelp. Looking at RGB values within the image may allow researchers to extrapolate how healthy the kelp is based on the color.

5 CONCLUSIONS & RECOMMENDATIONS

TO BE CHANGED AFTER RUNNING OVER MIN / MAX THRESHOLDING

In conclusion this study provides a preliminary analysis on the efficacy of a SAM based method over high resolution UAV imagery. By creating segmentations of the water within images and subsequently inverting the mask, this pipeline offers an approach to estimating kelp coverage without the need for specialized sensors or model retraining. Qualitatively, the pipeline is capable of producing consistent segmentation masks across multiple sites and conditions, however, quantitative results are still required.

For future work, several different avenues can be explored. The highest priority is the generation of a ground truth dataset. Basic thresholding and manual annotations can be utilized to generate an estimation of the true coverage of kelp within an image, with further validation done by domain experts. A ground truth dataset would also allow for a more rigorous approach to hyperparameter tuning as segmentation accuracy could be utilized as a metric for optimizing the performance of the pipeline. Additionally, applying this pipeline to the available UAV imagery would allow for a temporal analysis of the sites already surveyed. Seeing the change in kelp coverage could allow for analysis of seasonal dynamics, kelp growth, and species specific behavior. Beyond generation of a ground truth, access to this high-resolution UAV imagery may allow for the generation of species-specific kelp coverage maps, distinguishing between *Macrocystis* (giant kelp) from *Nereocystis* (bull kelp). Creating a distinction between these species is challenging, as they often appear structurally similar from above with NIR bands commonly used to differentiate them. Future work may explore including the additional spectral bands on the UAV as to increase the information available and allow for species discrimination if an RGB only based approach proves to be insufficient. Species-specific mapping would allow for a more accurate understanding of kelp health within regions.

ACKNOWLEDGEMENTS

The MBARI Summer Internship Program is generously supported through a gift from the Dean and Helen Witter Family Fund and the Rentschler Family Fund in memory of former MBARI board member Frank Roberts (1920-2019) and by the David and Lucile Packard Foundation. Additional funding is provided by the Maxwell/Hanrahan Foundation.

References

[Akyon et al., 2022] Akyon, F. C., Altinuc, S. O., and Temizel, A. (2022). Slicing aided hyper inference and fine-tuning for small object detection. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 966–970.

[Cavanaugh et al., 2021] Cavanaugh, K. C., Bell, T. W., and Hockridge, E. G. (2021). An automated method for mapping giant kelp canopy dynamics from uav. *Frontiers in Environmental Science*, 8.

[Cavanaugh et al., 2010] Cavanaugh, K. C., Siegel, D. A., Kinlan, B. P., and Reed, D. C. (2010). Scaling giant kelp field measurements to regional scales using satellite observations. *Marine Ecology Progress Series*, 403:13–27.

[Kirillov et al., 2023] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y., Dollár, P., and Girshick, R. (2023). Segment anything. *arXiv preprint arXiv:2304.02643*.

[Piñeiro Corbeira et al., 2023] Piñeiro Corbeira, C., Barrientos, S., Provera, I., García, M. E., Díaz-Tapia, P., Peña, V., Bárbara, I., and Barreiro, R. (2023). Kelp forests collapse reduces understorey seaweed β-diversity. *Annals of Botany*.

[U.S. National Park Service, 2019] U.S. National Park Service (2019). Kelp forests - channel islands national park. `https://www.nps.gov/chis/learn/nature/kelp-forests.htm`.

[Zhang et al., 2023] Zhang, C., Han, D., Qiao, Y., Kim, J. U., Bae, S.-H., Lee, S., and Hong, C. S. (2023). Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*.

[Zuckerman, 2023] Zuckerman, C. (2023). The vanishing kelp forest. *Nature Conservancy Magazine*.