

Human-in-the-Loop Segmentation of Multi-species Coral Imagery

Scarlett Raine^{1,2}, Ross Marchant³, Brano Kusy², Frederic Maire¹,
Niko Sünderhauf¹ and Tobias Fischer¹

¹QUT Centre for Robotics, Australia {sg.raine, f.maire, niko.suenderhauf, tobias.fischer}@qut.edu.au

²CSIRO Data61, Australia {scarlett.raine, brano.kusy}@csiro.au

³Image Analytics, Australia ross.g.marchant@gmail.com

Abstract

Broad-scale marine surveys performed by underwater vehicles significantly increase the availability of coral reef imagery, however it is costly and time-consuming for domain experts to label images. Point label propagation is an approach used to leverage existing image data labeled with sparse point labels. The resulting augmented ground truth generated is then used to train a semantic segmentation model. Here, we first demonstrate that recent advances in foundation models enable generation of multi-species coral augmented ground truth masks using denoised DINOv2 features and K-Nearest Neighbors (KNN), without the need for any pre-training or custom-designed algorithms. For extremely sparsely labeled images, we propose a labeling regime based on human-in-the-loop principles, resulting in significant improvement in annotation efficiency: If only 5 point labels per image are available, our proposed human-in-the-loop approach improves on the state-of-the-art by 17.3% for pixel accuracy and 22.6% for mIoU; and by 10.6% and 19.1% when 10 point labels per image are available. Even if the human-in-the-loop labeling regime is not used, the denoised DINOv2 features with a KNN outperforms the prior state-of-the-art by 3.5% for pixel accuracy and 5.7% for mIoU (5 grid points). We also provide a detailed analysis of how point labeling style and the quantity of points per image affects the point label propagation quality and provide general recommendations on maximizing point label efficiency.

1. Introduction

Effective and informed management of marine ecosystems requires data at a range of spatio-temporal scales [8]. Marine surveys are being increasingly performed using autonomous underwater and surface vehicles [9, 16]. However, these approaches generate large quantities of images

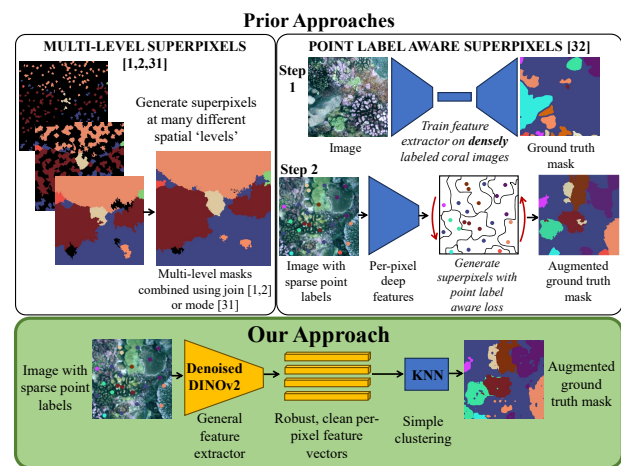


Figure 1. Our point label propagation approach leverages the DINOv2 foundation model without any fine-tuning to generate augmented ground truth masks for complex underwater imagery. **Top:** Prior approaches relied on layering superpixels containing point labels (left), or pre-training a feature extractor on labeled coral imagery, and then clustering pixels using a custom-designed superpixel algorithm (right). **Bottom:** Our method uses KNN to cluster deep features from the denoised DINOv2 foundation model.

of the seafloor which must first be analyzed to obtain usable outputs such as coverage estimations of different substrates and coral species [33, 37]. Coral images are often highly complex, with indistinct boundaries, high variation in color and texture among coral species and poor clarity [20, 36]. The intricate image characteristics and the difficulty in accurately identifying coral species requires domain experts to annotate underwater imagery, preventing the use of common computer vision tools such as the crowd sourcing annotation platform Amazon Turk.

Traditionally, marine scientists annotate underwater images using a method called Coral Point Count [24], where randomly or grid-spaced sparse pixels are labeled. These

pixels are called *point labels* [32]. Although there is a large quantity of historic data available in both grid and random formats [4, 13], the optimal style of point placement for the purpose of training deep learning models to perform semantic segmentation has not been explored.

In recent years, superpixels based on color information [1, 2, 31] and deep features [32] have been used for propagating point labels into dense, pixel-wise augmented ground truth masks used to train deep neural networks to perform semantic segmentation of unseen coral images. Most recently, Raine *et al.* [32] introduced a novel point label aware approach to superpixels, which clustered pixels based on deep features. While this method improved on the state-of-the-art, it required training on coral imagery to provide the deep features for the superpixel method, and suffered from performance degradation when small quantities of point labels are available.

In this work, we tackle the regime in which extremely few labels are provided. This setting is critical as marine survey projects often have limited budgets for labeling data [8]. Furthermore, a common use case for survey data processing involves quickly iterating and retraining models during field trips as new species or environmental conditions are encountered [8]. We propose using the general foundation model DINOv2 [28, 38] to provide the per-pixel deep features (Fig. 1). We use the simple K-Nearest Neighbor algorithm to generate the augmented ground truth, and outperform the state-of-the-art for small numbers of point labels. We also demonstrate further performance improvements by using a human-in-the-loop point selection regime in which the knowledge of the human expert is leveraged to reduce uncertainty in the KNN’s feature space.

This paper demonstrates the relevance of general foundation models for multi-species segmentation of domain-specific underwater imagery, while improving label efficiency when few points are available (Fig. 1). Our contributions are summarized as follows:

1. We propose to leverage a general purpose foundation model to generate per-pixel deep features for domain-specific coral images and establish that the features are effective without any training or fine-tuning on coral imagery. Combining these features with the simple K-Nearest Neighbors algorithm is sufficient for generating accurate augmented ground truth masks from sparse point labels and removes the need for complex superpixel algorithms.
2. For extremely sparse point labels, *i.e.* 5-25 points per image, we propose a human-in-the-loop labeling regime, which combines human knowledge with the model’s introspective uncertainty to select informative point label locations. We outperform the previous state-of-the-art by 17.3% pixel accuracy and 22.6% mIoU when there are 5 point labels available per image, and by 10.6% and

19.1% if 10 point labels are available.

3. Even without the human-in-the-loop labeling regime, using DINOv2 denoised features with a KNN improves on the label propagation task for small numbers of point labels per image. On the UCSD Mosaics dataset we see improvements of 3.5% for pixel accuracy and 5.7% for mIoU when 5 points are labeled; and 3.3% in pixel accuracy and 10.2% in mIoU for 10 points.
4. We perform thorough experiments to determine the effect of the number of point labels and the point labeling style on the point propagation task, and provide meaningful recommendations for efficient annotation.

We make our code¹ and video² available to foster future research.

2. Related Work

Broad-scale marine survey technologies such as autonomous underwater and surface vehicles enable the collection of large quantities of imagery [9, 16, 25, 27]. Automating the analysis of this imagery requires solutions which combine computer vision, deep learning and domain-specific expertise in marine biology [14, 37]. This section discusses approaches for semantic segmentation of underwater imagery and weakly supervised methods for point label propagation, recent advances in foundation models, and human-in-the-loop principles.

2.1. Segmentation of Underwater Imagery

Semantic segmentation of underwater imagery is complicated by a range of factors including the visual traits of coral species, which can appear similar between different species, and which are often intricate and highly textured [4, 15]. The image quality and clarity can also be affected by turbidity, scattering and attenuation of sunlight, blur and changes in coloration due to depth [20, 36]. These image characteristics, combined with the lack of semantic “objectness” of coral instances, makes semantic segmentation of underwater imagery a unique and challenging problem in computer vision.

There have been numerous approaches for fully supervised segmentation of corals in underwater imagery [11, 18, 35, 35, 41, 42, 44, 45], where the model is trained on pairs of images and densely labeled, pixel-wise ground truth masks. The TagLab annotation tool [30] makes dense pixel-wise annotation of large orthoimages faster, but relies on a model trained on 15,000 densely labeled coral images.

There are fewer approaches for weakly supervised segmentation of corals [1, 2, 31, 32, 39, 40]. These approaches are based on custom-designed superpixel methods which

¹<https://github.com/sgraine/HIL-coral-segmentation>

²<https://youtu.be/YBTUCECu30M>

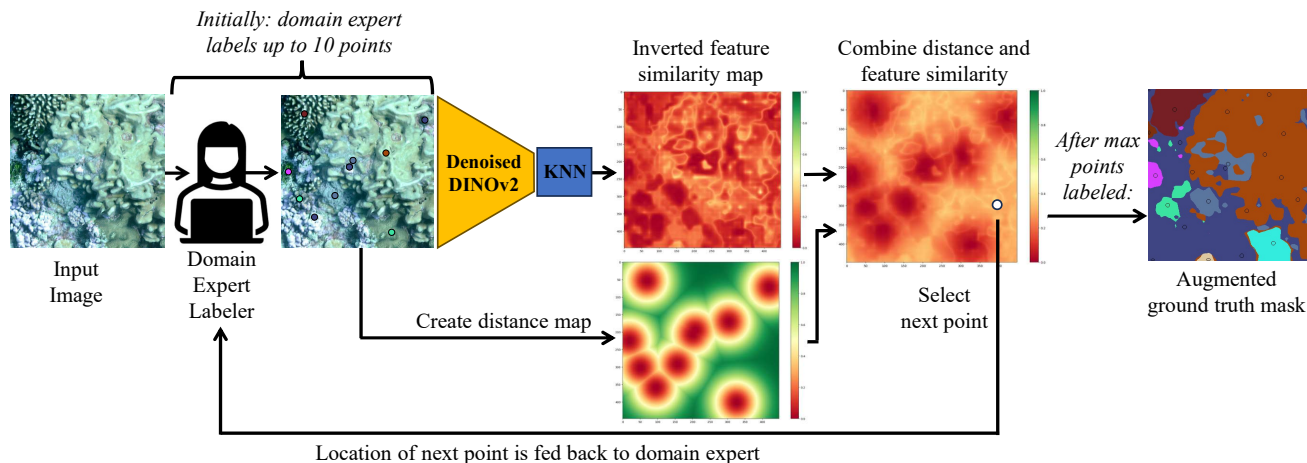


Figure 2. Proposed Algorithm Schematic. Our smart labeling scheme combines domain expert knowledge and our model’s internal uncertainty to optimize point label selection in a human-in-the-loop framework. Our method starts by taking a coral image as input and requesting the domain expert to label up to 10 points centrally in the largest instances. Then a feature similarity map is generated by calculating the cosine similarities between the labeled points and every other pixel. We encourage exploration by incorporating a distance map and then combine both maps to obtain an overall probability mask for pixel selection. The selected pixel is fed back to the domain expert for labeling, and then the KNN is updated. Once the maximum points have been labeled, the augmented ground truth mask is generated and can be later used for training a model to perform semantic segmentation.

generate dense ground truth masks from sparse point labels. The multi-level superpixel method [1, 2, 31] generates superpixels from color features at many spatial scales, and then propagates point labels within each segment before joining together the different “levels” of superpixels. The most recent approach, Point Label Aware Superpixels [32] describes a novel superpixel algorithm which uses the point labels directly in generating the superpixel segments.

These prior superpixel approaches rely on having sufficient points available, and experience degraded performance in the very sparse setting. There is an opportunity to generalize and simplify these approaches by leveraging the recent advances in general foundation models.

2.2. Foundation Models

Recent works developed foundation models for learning robust pre-trained feature representations that are task-agnostic [23, 28, 46]. Foundation models are trained on large-scale datasets and are designed to learn highly generalized representations which allow the model to transfer to tasks and data outside of the training distribution [23].

Some works aim to personalize foundation models for specific visual concepts, e.g. the user’s pet, as in [43], or adapt the model by training a task-specific decoder or adapter [5]. For leaf counting, instance segmentation, and disease classification for plant phenotyping, the adapted general foundation models did not outperform the task-specific methods [5]. In medical image analysis, [3] demonstrate the cross-task generalizability of DINOv2 and re-

port competitive performance when the features are used with KNN for disease classification. Other works have performed self-supervised object localization without labels [34], however they do not perform segmentation of the entire image.

Although some research has investigated the application of the DINOv2 foundation model for specialized problems [3, 5, 17], the performance of DINOv2 for underwater coral segmentation has not been studied. As described in Section 2.1, underwater images have unique visual characteristics, including abstract textures, fractal-like boundaries, and overlapping instances [4, 32]. It is unknown whether foundation models trained on general images are able to produce meaningful feature embeddings for coral imagery.

2.3. Human-in-the-Loop

Human-in-the-Loop describes machine learning which involves interaction between humans and the algorithm [26]. Specifically, the Human-in-the-Loop sub-field of Interactive Machine Learning describes a framework in which control is shared between the human and the model: the human supplies information to the model in a focused, frequent and interactive way [19, 26].

While using models to predict labels on new data has been performed in the ecology domain previously [6, 22], the use of foundation models as part of an interactive labeling framework has not been implemented for coral point label propagation.

To our knowledge, an approach for multi-species coral

segmentation which combines the general knowledge of foundation models with sparse domain-specific labeling has not been proposed in the literature. This is an opportunity to decrease the associated time and cost of manually labeling domain-specific imagery, while improving the accuracy of propagated ground truth masks when few labels are available.

3. Method

3.1. Method Overview

Our proposed point label propagation approach leverages the denoised DINOv2 foundation model [38], based on [28]. We generated the augmented ground truth mask by clustering pixels in the deep feature space with K-Nearest Neighbors.

Our approach takes a photo-quadrat coral image and a set of sparse point labels as input, and outputs a dense pixel-wise augmented ground truth mask. The sparse point labels are either randomly distributed in the image, spaced evenly as a grid³, or selected using our proposed Human-in-the-Loop smart labeling regime (Section 3.2).

We take our image as input and obtain a set of feature vectors of length 768 from the denoised DINOv2 feature extractor [38]. This feature extractor is based on a transformer architecture [28], which outputs a deep feature for every 14x14 pixel patch in the input image. The feature extractor also outputs a ‘CLS’ token for the whole image, which we do not use. We spatially upsample the feature vectors with bilinear interpolation, such that we obtain one deep feature vector for each pixel in the input image. The per-pixel feature vectors are then L2 normalized.

We take our set of L sparse labeled pixels and store the normalized feature embeddings $\{\mathbf{v}_1, \dots, \mathbf{v}_l, \dots, \mathbf{v}_L\}$ for these pixels. We also store $X = \{(x_1, y_1), \dots, (x_L, y_L)\}$ where (x_l, y_l) are the pixel coordinates of \mathbf{v}_l . We calculate the cosine similarity between the feature embedding \mathbf{v}_l for $l \in \{1, \dots, L\}$ and the feature embedding \mathbf{v}_p for every other pixel p in the image:

$$\text{sim}(\mathbf{v}_p, \mathbf{v}_l) = \mathbf{v}_p \cdot \mathbf{v}_l. \quad (1)$$

We determine the augmented ground truth mask by performing K-Nearest Neighbors with $k = 1$. Note that we trialed different values for k and did not see any improvement for $k > 1$, as seen in Fig. 5 and discussed further in the Supplementary Material.

3.2. Human-in-the-Loop Pixel Selection

To further improve our performance in the extremely sparse case, we propose a novel labeling regime (Fig. 2). In con-

³In the case that grid-spaced points are used and the number of points cannot be equally distributed into rows and columns, the nearest quantity is used e.g. for 5 point labels, a grid of 2x2 points is used.

trast to prior approaches, which have leveraged randomly distributed or grid-spaced sparse point labels, we consider the point labeling problem as a human-in-the-loop task. To this end, we assume that we have a domain expert available to collaboratively label a certain number of points, which are then used in our KNN and DINOv2 point label propagation approach.

To select informative points that we want the human to label, we consider the cosine similarity between labeled and unlabeled pixel features in the DINOv2 deep feature space. We start by asking the domain expert to label up to 10 pixels in the middle of the largest instances they can see in the image. For more than 10 point labels, the smart pixel regime iteratively proposes one point at a time for labeling based on which parts of the image have the most uncertainty. This uncertainty is modeled as the cosine similarity to the closest labeled pixel.

To do this, we first follow the method described in Section 3.1 to obtain, upsample and normalize the per-pixel feature embeddings and we find a cosine similarity map (Eq. 1) between the starting labeled pixels and every other pixel in the image. We invert this map such that pixel locations which have a low cosine similarity to the closest labeled pixel have a higher probability of selection:

$$C(x, y) = 1 - \max_{l \in \{1, \dots, L\}} \text{sim}(\mathbf{v}_q, \mathbf{v}_l), \quad (2)$$

where \mathbf{v}_q is the feature vector at location (x, y) .

We encourage exploration across the whole image by creating a probabilistic distance map for the labeled pixels. We first compute the Euclidean distance transform on a binary mask denoting the location of our set of labeled pixels, where initially $L = 10$:

$$D(x, y) = \min_{(x', y') \in X} \sqrt{(x - x')^2 + (y - y')^2}. \quad (3)$$

We then perform Gaussian smoothing over the distance transform and tune the smoothing parameter σ in the ablation study in Section 5.2:

$$D_{\text{smooth}}(x, y) = 1 - \exp\left(-\frac{D(x, y)^2}{2\sigma^2}\right). \quad (4)$$

We combine the probabilistic cosine similarity map with the distance map, and weight the two terms with $\lambda = 2.2$ (see hyperparameter tuning in Section 5.2):

$$M(x, y) = \frac{D_{\text{smooth}}(x, y) + \lambda \text{sim}(\mathbf{v}_x, \bar{\mathbf{v}}_l)}{\lambda + 1}. \quad (5)$$

After combining the distance map, we select the next pixel for labeling by taking the location $(\hat{x}, \hat{y}) = \arg \max_{(x, y)} M(x, y)$ corresponding to the highest selection probability in M .




Figure 3. Point Label Propagation Pixel Accuracy and Mean IoU: our KNN and DINOv2 approach is shown in orange for random and grid labeling, and the red line depicts performance of the KNN and DINOv2 approach with the Human-in-the-Loop collaborative labeling scheme. Our approach significantly outperforms prior works when there are very sparse point labels available, *i.e.* 5-25 points. When a larger quantity of points are used (300 points), the performance of the different approaches converges.

Table 1. Performance of Point Label Propagation Approaches (Refer to Section 4.3 for Metric Definitions), for 5 / 10 / 25 / 300 Point Labels. ‘F-MSS’ is Fast MSS [31], ‘PLAS’ is Point Label Aware Superpixels [32], and ‘D+NN’ is KNN with Denoised DINOv2 [38] (Ours).

Method	Label Style	PA				mPA				mIoU				Time per Image (s)			
		5	10	25	300	5	10	25	300	5	10	25	300	5	10	25	300
F-MSS	Rand.	7.29	13.49	30.09	86.81	6.60	12.34	29.26	82.70	6.55	12.11	28.53	80.12	2.14	2.19	2.21	2.76
F-MSS	Grid	6.36	13.94	39.18	89.98	5.95	13.01	36.72	88.17	5.93	12.83	35.51	86.44	2.43	2.45	2.36	2.96
PLAS - Single	Rand.	48.45	55.26	65.16	86.68	32.03	41.44	57.65	81.74	23.86	32.22	47.76	77.56	1.71	2.00	2.17	1.93
PLAS - Single	Grid	51.88	58.11	72.96	89.28	38.24	46.30	64.91	86.16	28.93	36.94	58.00	82.73	1.55	1.80	2.06	1.81
PLAS - Ens.	Rand.	52.73	62.00	71.11	92.47	36.48	49.04	63.21	89.93	25.91	35.6	50.46	85.45	4.27	4.55	5.02	5.35
PLAS - Ens.	Grid	54.23	65.76	76.31	94.60	40.08	53.20	69.13	92.49	30.00	40.34	59.82	89.38	4.06	4.25	5.15	5.28
D+NN (Ours)	Rand.	55.72	64.51	75.07	88.77	39.94	50.91	65.80	83.84	32.09	42.79	58.04	81.75	4.88	4.55	4.74	4.90
D+NN (Ours)	Grid	57.73	69.07	78.74	89.86	44.40	58.08	70.05	87.41	35.74	50.58	64.40	85.77	4.78	4.70	4.79	4.69
D+NN (Ours)	Smart	71.56	76.38	81.27	89.61	61.46	69.87	75.91	86.45	52.60	59.48	67.97	85.00	4.74	4.98	20.0	273.08

4. Experimental Setup

This section describes the implementation details (Section 4.1), evaluation datasets (Section 4.2), and evaluation metrics (Section 4.3).

4.1. Implementation

All experiments are conducted with a Quadro RTX 6000, and inference times are with respect to this GPU. Our approach is implemented using Python and PyTorch [29]. We use the Faiss module to enable fast K-Nearest Neighbors on GPU [21]. We use the denoised DINOv2 model and implementation from [38].

4.2. Datasets

The UCSD Mosaics dataset is a multi-species coral dataset labeled with dense ground truth masks [2, 10]. We use the

version of the dataset provided by [2], but we notice some ground truth masks are corrupted so we exclude these (we remove 219 from the training split, yielding 3,974 images and 32 from the test split, resulting in 696 images; further details can be found in the Supplementary Material). Each image is 512 by 512 pixels and the dataset contains 33 types of corals and an ‘unknown’ or ‘unlabeled’ class. For consistency with [2, 32], we ignore this class during evaluation. To simulate the domain expert in our human-in-the-loop smart labeling regime, we take the point label from the ground truth mask at the location proposed, or in the case of the first 10 pixels, we select the middle pixel of the largest instances in the mask.

4.3. Evaluation Metrics

We use three frequently used metrics [2, 12, 32] to establish and compare the performance of our approach with prior

Table 2. Effect of Denoising on Point Propagation (Refer to Section 4.3 for Metric Definitions). ‘Raw’ refers to the original DINOv2 [28] and ‘Denoise’ refers to the Denoising ViT implementation [38].

Labels	PA	mPA	mIoU
	Raw / Denoise	Raw / Denoise	Raw / Denoise
5	68.58 / 71.57	60.23 / 61.46	50.28 / 52.60
10	73.32 / 76.38	68.04 / 69.87	55.76 / 59.48
25	76.94 / 81.27	70.97 / 75.91	61.61 / 67.97
50	79.71 / 82.80	76.68 / 77.52	68.43 / 72.96
100	82.87 / 85.60	79.45 / 81.82	74.92 / 78.77
200	86.14 / 88.06	84.20 / 84.83	81.44 / 82.75
300	88.10 / 89.61	85.58 / 86.45	83.79 / 85.00

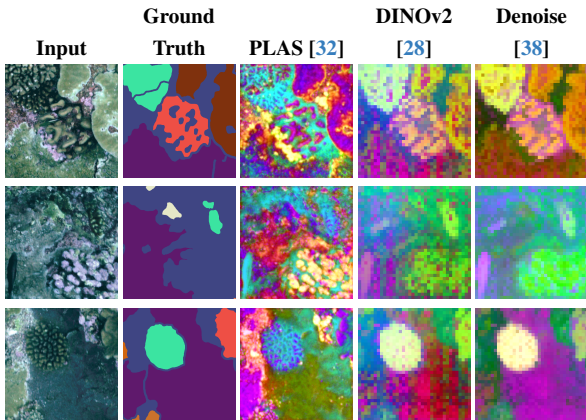


Figure 4. Comparison of Point Label Aware Superpixels (PLAS) [32] features, DINOv2 features [28], and denoised DINOv2 features [38]. For the transformer approaches, features for every 14x14 patch in the original image have been upsampled with bilinear interpolation. All features are reduced to RGB for visualisation with Principal Components Analysis (PCA). Pixels with similar RGB colors are similar in the deep feature space. The CNN features used by PLAS do not effectively group pixels into meaningful segments. The denoising model clearly reduces the artefacts, resulting in smoother, cleaner features and improved clustering.

methods: 1) Pixel Accuracy (PA) is the sum of correctly classified pixels divided by the predicted pixels; 2) The mean pixel accuracy (mPA) is the pixel accuracy averaged over the classes; and 3) The mean intersection over union (mIoU) denotes the average of the per-class IoU scores. A higher score indicates better performance for all metrics.

5. Results

We first compare our method to the state-of-the-art for point label propagation in Section 5.1 and then provide ablation studies in Section 5.2.

5.1. Comparison to State-of-the-art Methods

We compare the performance of our novel method to state-of-the-art approaches, namely Fast Multi-level Superpixel

Segmentation (*Fast MSS*) [31], a faster implementation of CoralSeg [2], and Point Label Aware Superpixels, for which we compare against both the single method (*Single*) and ensemble of three Point Label Aware algorithms (*Ensemble*), as described in [32].

As shown in Table 1, leveraging K-Nearest Neighbors with features extracted by the denoised DINOv2 foundation model [38] for point label propagation outperforms prior approaches for small numbers of point labels (5, 10 and 25 per image). The absolute increase in mIoU is 46.1% and 22.6% when compared to Fast MSS and Point Label Aware Superpixels respectively for five point labels and our human-in-the-loop labeling regime (Fig. 3); and we improve by 64.3% and 17.3% for pixel accuracy (Fig. 3). If the human-in-the-loop labeling regime is not used, we still outperform the state-of-the-art by 3.5% for pixel accuracy and 5.7% for mIoU (for 5 grid points). Even in the setting that we do not target in this paper, *i.e.* if there are as many as 300 point labels available, our approach exhibits comparable performance to the single classifier methods but is outperformed by the ensemble of three Point Label Aware Superpixel classifiers (Table 1).

Our approach, which leverages DINOv2 and our smart labeling regime, exhibits similar computation times per image as the ensemble for the Point Label Aware Superpixel method for small quantities of point labels (Table 1). However, when smart labeling is used for large quantities of points *i.e.* 300 points, the computation time is prohibitively high as clustering in the deep feature space occurs every iteration to generate the feature similarity map (Eq. 2). For large quantities of points, the grid-spaced version should be used instead. It is important to note that the intended use case for smart labeling is to improve performance in the extremely sparse (5-25 points) setting.

Fig. 6 presents qualitative results demonstrating that our method generates a dense augmented ground truth mask which closely matches the ground truth provided, even for very sparsely labeled images.

5.2. Ablation Study

5.2.1 Denoising DINOv2 Features

We use the denoised feature extractor for DINOv2 described in [38] and demonstrate the utility of this method by results as seen in Table 2. Fig. 4 compares the original [28] and denoised [38] DINOv2 deep feature embeddings. We provide an additional ablation in the Supplementary Material which also evaluates the performance when DINOv2 is trained with registers [7], both with and without denoising [38], although we did not see an improvement with this approach. Fig. 4 compares the deep features to the ground truth for each of the images, as well as the deep CNN features used in the Point Label Aware Superpixel method [32]. We show that the raw DINOv2 features




Figure 5. Pixel accuracy of label propagation for 25 point labels. Our human-in-the-loop smart pixel selection regime is robust to changes in hyperparameter values. (a) If the feature similarity map is weighted more highly, there is a small improvement (we choose $\lambda = 2.2$). (b) The pixel accuracy is highest when the distance map Gaussian smoothness is set to $\sigma = 50$. (c) Clustering with KNN results in higher performance when $k = 1$.

exhibit artefacts caused by the position embeddings used during training. These artefacts are unhelpful during clustering because coral instances of the same class can appear in different areas of the same image. The denoising model effectively removes these artefacts, resulting in cleaner features and improved point propagation (Table 2).

5.2.2 Weighting the Probability Maps (λ)

We evaluate the impact of weighting term λ which balances the importance of the cosine similarity map (Eq. 2) and the distance map (Eq. 3). As seen in Fig. 5, we find that a value of $\lambda = 2.2$ results in the best pixel accuracy, although our approach is not sensitive to the exact value of λ .

5.2.3 Exclusion Distance

Our human-in-the-loop labeling regime considers the distance between each pixel selected for labeling by incorporating a Gaussian-smoothed distance mask between all pixels and the previously labeled pixels. The Gaussian smoothing introduces the σ hyperparameter which controls how closely pixels can be selected to previously labeled pixels. We demonstrate the impact of this hyperparameter on the point label propagation task through the ablation study results in Fig. 5, and find that our approach is robust to different values.

5.3. Effect of Point Label Quantity

Greater quantities of point labels resulted in improved performance for the point label propagation task (Fig. 3). Having a sufficient number of points is especially critical for the Fast MSS [31] approach. For grid-spaced point labels, the Fast MSS approach improves from 5.9% to 86.4% mIoU when increasing from 5 to 300 labels, a difference of 80.5%,

whereas the equivalent difference in mIoU is 59.4% and 50.0% for the Point Label Aware Superpixels [32] and our DINOv2 and KNN approach, respectively (Table 1). We include the results for all values of point labels (5, 10, 25, 50, 100, 200, 300) in the Supplementary Material.

Although the point label propagation improves for all methods as the quantity is increased, there is a decrease in the rate of improvement as the number of labels is increased from 100 to 300 points. When the quantity of grid labels is increased from 100 to 300 points per image, the Fast MSS [31] approach improves by 11.7% for mIoU, as compared to an improvement of 68.9% for mIoU when increasing from 5 to 100 points. For the Point Label Aware Superpixels, the mIoU improves by 49.9% when increasing from 5 to 100 grid points, and by 9.5% when increasing from 100 to 300 points. For our denoised DINOv2 and KNN approach, the mIoU improves by 26.2% when increasing from 5 to 100 smart points, and by 6.2% when increasing from 100 to 300 smart points.

5.4. Effect of Point Label Placement Style

All of the approaches evaluated benefit from grid placement of point labels over randomly placed point labels (Fig. 3). The effect is particularly pronounced for the multi-level superpixels (Fast MSS) [31], which exhibits absolute improvements for mIoU when using grid-spaced pixels of 13.3%, 9.9%, 6.8% and 6.3% for label quantities of 50, 100, 200 and 300 points respectively. The Point Label Superpixels also benefit from grid-spaced points, with improvements of 8.4%, 4.8%, 4.4% and 3.9% also for 50, 100, 200 and 300 points respectively. Grid-spaced labels ensure consistent coverage of the whole image and make effective use of every point label. As seen in Fig. 6, randomly placed point labels can often be placed very close together, reducing the information gained.

Fig. 6 demonstrates that for very small numbers of point labels, *e.g.* 5 points per image, a significant benefit is gained from leveraging the knowledge of the domain expert in selecting points in the center of instances for up to 10 points, and then iteratively selecting further pixels with the point propagation model, as described in Section 3. The augmented ground truth masks obtained by our proposed approach (top two rows of Fig. 6) are significantly closer to the ground truth than the prior approaches. We note that our smart label regime could be applied to the other techniques as well, and we will investigate this in future work.

When few labels are available, the multi-level superpixel methods [2, 31] suffer as the method relies on layering labeled regions from different scales. The point label superpixel method suffers in the sparse label case as the boundaries of the superpixels are not forced to conform to the instance boundaries by conflicting point labels [32]. Our method performs well in the sparse label cases because pix-

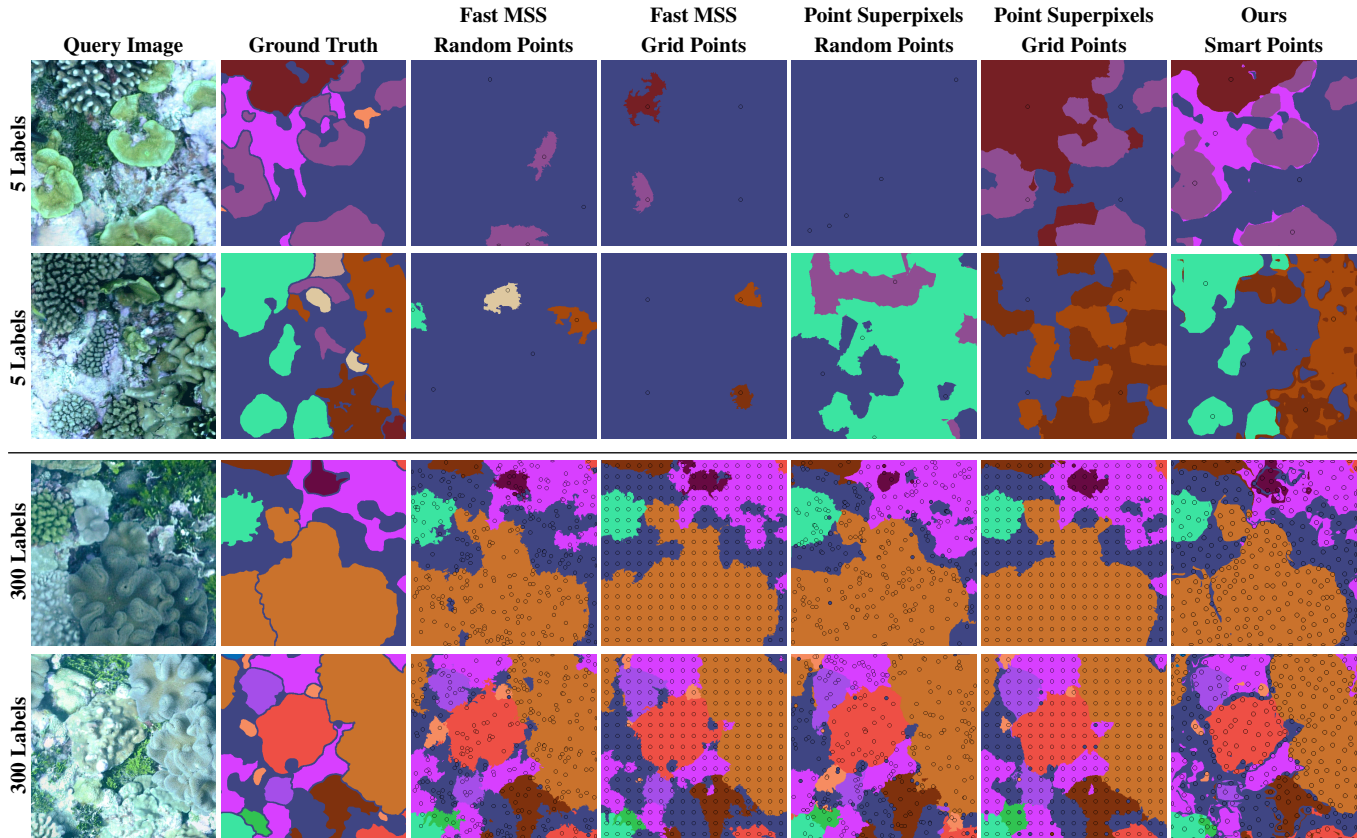


Figure 6. Qualitative comparison between the Fast MSS approach [31], the Point Label Aware Superpixel approach [32] and our approach, based on denoised DINOv2 features [38], K-Nearest Neighbors and our Human-in-the-Loop labeling regime. The top two rows show point propagation for 5 labels, and the bottom two rows demonstrate point propagation when there are 300 labels available. The pixels used in the point label propagation are shown as black circles within the output augmented ground truth masks. Additional qualitative results are in the Supplementary Material.

els can be assigned as the correct class even if spatially far away from labeled points, because the clustering occurs only in the deep feature space.

6. Conclusion

This work has demonstrated that the DINOv2 foundation model can be used without any fine-tuning to perform point label propagation in underwater imagery. Our approach uses denoised DINOv2 features and the simple KNN algorithm to generate augmented ground truth masks. When used with our proposed human-in-the-loop labeling regime, we improve by 22.6%, 19.1% and 8.2% mIoU for 5, 10 and 25 point labels as compared to the prior state-of-the-art, significantly improving label efficiency. Even if the DINOv2 features and KNN is used for grid-spaced point labels, we outperform prior approaches on the UCSD dataset by 5.7%, 10.2% and 4.6% mIoU for 5, 10 and 25 point labels respectively.

We also perform comprehensive studies on the effect of the number of point labels and the point label placement

style on the accuracy of the point propagation task, and recommend that grid labels improve point label propagation performance over random labels for all approaches. For greater than 100 points per image, the improvement in point label propagation is marginal if using the Point Aware Superpixels [32] method or the DINOv2 and KNN method presented in this paper. For very few points per image, it is beneficial to incorporate domain expert knowledge in selecting which pixels to label. This work has demonstrated the relevance of general foundation models for complex domain specific tasks, and has significantly improved performance and annotation efficiency of point label propagation in the extremely sparse label setting.

Acknowledgments

This work was done in collaboration between QUT and CSIRO Data61. S.R., F.M., N.S., and T.F. acknowledge continued support from the Queensland University of Technology (QUT) through the Centre for Robotics. T.F. acknowledges funding from Intel Labs via grant RV3.290.Fischer and an ARC Laureate Fellowship FL210100156 to Prof. Michael Milford.

References

- [1] Iñigo Alonso and Ana C Murillo. Semantic segmentation from sparse labeling using multi-level superpixels. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5785–5792, 2018. 2, 3
- [2] Iñigo Alonso, Matan Yuval, Gal Eyal, Tali Treibitz, and Ana C Murillo. CoralSeg: Learning coral segmentation from sparse annotations. *Journal of Field Robotics*, 36(8):1456–1477, 2019. 2, 3, 5, 6, 7
- [3] Mohammed Baharoon, Waseem Qureshi, Jiahong Ouyang, Yanwu Xu, Kilian Phol, Abdulrhman Aljouie, and Wei Peng. Towards general purpose vision foundation models for medical image analysis: An experimental study of DINOv2 on radiology benchmarks. *arXiv preprint arXiv:2312.02366*, 2023. 3
- [4] Oscar Beijbom, Peter J Edmunds, David I Kline, B Greg Mitchell, and David Kriegman. Automated annotation of coral reef survey images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1170–1177, 2012. 2, 3
- [5] Feng Chen, Mario Valerio Giuffrida, and Sotirios A Tsafaris. Adapting vision foundation models for plant phenotyping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 604–613, 2023. 3
- [6] Qimin Chen, Oscar Beijbom, Stephen Chan, Jessica Bouwmeester, and David Kriegman. A new deep learning engine for CoralNet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3693–3702, 2021. 3
- [7] Timothée Darcet, Maxime Oquab, Julien Mairal, and Piotr Bojanowski. Vision transformers need registers. *arXiv preprint arXiv:2309.16588*, 2023. 6
- [8] Ellen M Ditria, Christina A Buelow, Manuel Gonzalez-Rivero, and Rod M Connolly. Artificial intelligence and automated monitoring for assisting conservation of marine ecosystems: A perspective. *Frontiers in Marine Science*, 9: 918104, 2022. 1, 2
- [9] Matthew Dunbabin, Justin Manley, and Peter L Harrison. Uncrewed maritime systems for coral reef conservation. In *Global Oceans*, 2020. 1, 2
- [10] Clinton B Edwards et al. Large-area imaging reveals biologically driven non-random spatial patterns of corals at a remote reef. *Coral Reefs*, 36(4):1291–1305, 2017. 5
- [11] Daniel P Furtado, Edson A Vieira, Wildna Fernandes Nascimento, Kelly Y Inagaki, Jessica Bleuel, Marco Antonio Zanata Alves, Guilherme O Longo, and Luiz S Oliveira. #DeOlhoNosCorais: A polygonal annotated dataset to optimize coral monitoring. *PeerJ*, 11:e16219, 2023. 2
- [12] Alberto Garcia-Garcia, Sergio Orts-Escolano, Sergiu Oprea, Victor Villena-Martinez, Pablo Martinez-Gonzalez, and Jose Garcia-Rodriguez. A survey on deep learning techniques for image and video semantic segmentation. *Applied Soft Computing*, 70:41–65, 2018. 5
- [13] Manuel González-Rivero, Pim Bongaerts, Oscar Beijbom, Oscar Pizarro, Ariell Friedman, Alberto Rodriguez-Ramirez, Ben Upcroft, Dan Laffoley, David Kline, Christophe Bailhache, et al. The Catlin seaview survey—kilometre-scale seascape assessment, and monitoring of coral reef ecosystems. *Aquatic Conservation: Marine and Freshwater Ecosystems*, 24(S2):184–198, 2014. 2
- [14] Manuel González-Rivero, Oscar Beijbom, Alberto Rodriguez-Ramirez, Tazio Holtrup, Yeray González-Marrero, Anjani Ganase, Chris Roelfsema, Stuart Phinn, and Ove Hoegh-Guldberg. Scaling up ecological measurements of coral reefs using semi-automated field image collection and analysis. *Remote Sensing*, 8(1):30, 2016. 2
- [15] Manuel González-Rivero et al. Monitoring of coral reefs using artificial intelligence: A feasible and cost-effective approach. *Remote Sensing*, 12(3):489, 2020. 2
- [16] Daniel Gregorek, Abraham Tibebe, Eduardo Caudet, Carlos Barrera, and Ralf Bachmayer. Long-endurance optical seafloor imaging using underwater gliders: Concept, development and initial trials. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 6162–6168, 2023. 1, 2
- [17] Joana Palés Huix, Adithya Raju Ganeshan, Johan Fredin Haslum, Magnus Söderberg, Christos Matsoukas, and Kevin Smith. Are natural domain foundation models useful for medical image classification? In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 7634–7643, 2024. 3
- [18] Md Jahidul Islam, Chelsey Edge, Yuyang Xiao, Peigen Luo, Muntaqim Mehtaz, Christopher Morse, Sadman Sakib Enan, and Junaed Sattar. Semantic segmentation of underwater imagery: Dataset and benchmark. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2020. 2
- [19] Liu Jiang, Shixia Liu, and Changjian Chen. Recent research advances on interactive machine learning. *Journal of Visualization*, 22:401–417, 2019. 3
- [20] Leilei Jin and Hong Liang. Deep learning for underwater image recognition in small sample size situations. In *MTS/IEEE OCEANS*, 2017. 1, 2
- [21] Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3):535–547, 2019. 5
- [22] Benjamin Kellenberger, Devis Tuia, and Dan Morris. AIDE: Accelerating image-based ecological surveys with interactive machine learning. *Methods in Ecology and Evolution*, 11(12):1716–1727, 2020. 3
- [23] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. 3
- [24] Kevin E Kohler and Shaun M Gill. Coral Point Count with Excel extensions (CPCe): A Visual Basic program for the determination of coral and substrate coverage using random point count methodology. *Computers & Geosciences*, 32(9): 1259–1269, 2006. 1
- [25] Yang Li, Jiajun Liu, Brano Kusy, Ross Marchant, Brendan Do, Torsten Merz, Joey Crosswell, Andy Steven, Lachlan Tychsen-Smith, David Ahmedt-Aristizabal, et al. A real-time edge-AI system for reef surveys. In *Proceedings of the*

- Annual International Conference on Mobile Computing And Networking*, pages 903–906, 2022. 2
- [26] Eduardo Mosqueira-Rey, Elena Hernández-Pereira, David Alonso-Ríos, José Bobes-Bascarán, and Ángel Fernández-Leal. Human-in-the-loop machine learning: A state of the art. *Artificial Intelligence Review*, 56(4):3005–3054, 2023. 3
- [27] Serena Mou, Dorian Tsai, and Matthew Dunbabin. Reconfigurable robots for scaling reef restoration. *arXiv preprint arXiv:2205.04612*, 2022. 2
- [28] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. DINOv2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3, 4, 6
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Proceedings of the Advances in Neural Information Processing Systems*, 32, 2019. 5
- [30] Gaia Pavoni, Massimiliano Corsini, Federico Ponchio, Alessandro Muntoni, Clinton Edwards, Nicole Pedersen, Stuart Sandin, and Paolo Cignoni. TagLab: AI-assisted annotation for the fast and accurate semantic segmentation of coral reef orthoimages. *Journal of Field Robotics*, 39(3):246–262, 2022. 2
- [31] Jordan P Pierce, Yuri Rzhano, Kim Lowell, and Jennifer A Dijkstra. Reducing annotation times: Semantic segmentation of coral reef survey images. In *Global Oceans*, pages 1–9, 2020. 2, 3, 5, 6, 7, 8
- [32] Scarlett Raine, Ross Marchant, Brano Kusy, Frederic Maire, and Tobias Fischer. Point label aware superpixels for multi-species segmentation of underwater imagery. *IEEE Robotics and Automation Letters*, 7(3):8291–8298, 2022. 2, 3, 5, 6, 7, 8
- [33] Hugh Runyan, Vid Petrovic, Clinton B Edwards, Nicole Pedersen, Esmeralda Alcantar, Falko Kuester, and Stuart A Sandin. Automated 2D, 2.5D, and 3D segmentation of coral reef pointclouds and orthoprojections. *Frontiers in Robotics and AI*, 9:884317, 2022. 1
- [34] Oriane Siméoni, Gilles Puy, Huy V Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing objects with self-supervised transformers and no labels. *arXiv preprint arXiv:2109.14279*, 2021. 3
- [35] Yee Wang Sui, Koh Xian Ming, Malika Meghjani, Nagarajan Raghavan, Cyrille Jegourel, and Keegan Kang. An automated data processing pipeline for coral reef monitoring. In *MTS/IEEE OCEANS*, 2022. 2
- [36] Xin Sun, Junyu Shi, Lipeng Liu, Junyu Dong, Claudia Plant, Xinhua Wang, and Huiyu Zhou. Transferring deep knowledge for object recognition in low-quality underwater videos. *Neurocomputing*, 275:897–908, 2018. 1, 2
- [37] Lian Xu, Mohammed Bennamoun, Senjian An, Ferdous Sohel, and Farid Boussaid. Deep learning for marine species recognition. In *Handbook of Deep Learning Applications*, pages 129–145. 2019. 1, 2
- [38] Jiawei Yang, Katie Z Luo, Jiefeng Li, Kilian Q Weinberger, Yonglong Tian, and Yue Wang. Denoising vision transformers. *arXiv preprint arXiv:2401.02957*, 2024. 2, 4, 5, 6, 8
- [39] Xi Yu, Ying Ma, Stephanie Farrington, John Reed, Bing Ouyang, and Jose C Principe. Fast segmentation for large and sparsely labeled coral images. In *IEEE International Joint Conference on Neural Networks*, 2019. 2
- [40] Xi Yu, Bing Ouyang, Jose C Principe, Stephanie Farrington, John Reed, and Yanjun Li. Weakly supervised learning of point-level annotation for coral image segmentation. In *MTS/IEEE Oceans*, pages 1–7, 2019. 2
- [41] Hanqi Zhang, Armin Grün, and Ming Li. Deep learning for semantic segmentation of coral images in underwater photogrammetry. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2:343–350, 2022. 2
- [42] Hanqi Zhang, Ming Li, Jiageng Zhong, and Jianguying Qin. CNet: A novel seabed coral reef image segmentation approach based on deep learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 767–775, 2024. 2
- [43] Renrui Zhang, Zhengkai Jiang, Ziyu Guo, Shilin Yan, Junting Pan, Hao Dong, Peng Gao, and Hongsheng Li. Personalize segment anything model with one shot. *arXiv preprint arXiv:2305.03048*, 2023. 3
- [44] Jiageng Zhong, Ming Li, Hanqi Zhang, and Jianguying Qin. Combining photogrammetric computer vision and semantic segmentation for fine-grained understanding of coral reef growth under climate change. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 186–195, 2023. 2
- [45] Zheng Ziqiang, Xie Yaofeng, Liang Haixin, Yu Zhibin, and Sai-Kit Yeung. CoralVOS: Dataset and benchmark for coral video segmentation. *arXiv preprint arXiv:2310.01946*, 2023. 2
- [46] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *Advances in Neural Information Processing Systems*, 36, 2024. 3