AI.FISH

Computer Vision Results

Review of Approach

AI.FISH

Goal: Predict species apportionment in an image

This involves:

- Localizing all fish in an image
  - 100s or 1000s of targets
- Classifying the species of all localized fish
- Computing proportions from these counts

Review of Approach

AI.FISH

Existing computer vision algorithms inadequate.

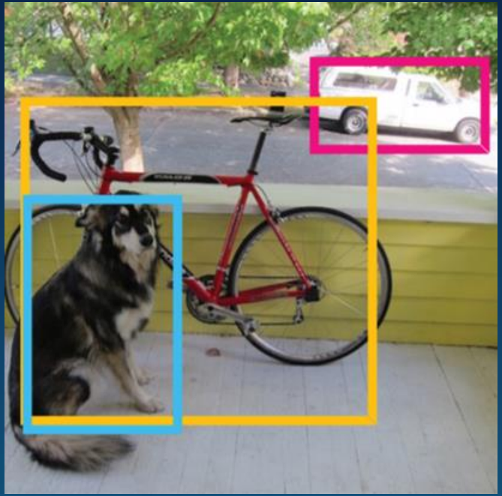

Object detection

- Localization (boxes) + classification
- But **cannot handle dense scenes**

Crowd localization

- Localization in dense scenes
- But **cannot classify**

Review of Approach

AI.FISH

Goal: Dense localization and classification.

Dataset

AI.FISH

- 9 vessels, 39 trips, 471 tows of raw data
- Three step annotation procedure:
    1. Localizations (“dots”)
    2. Expert species classifications
    3. Expert review
- Collected 477,889 annotations in 3,362 images

Algorithm Details

Starting point: Crowd Localization Transformer (CLTR)

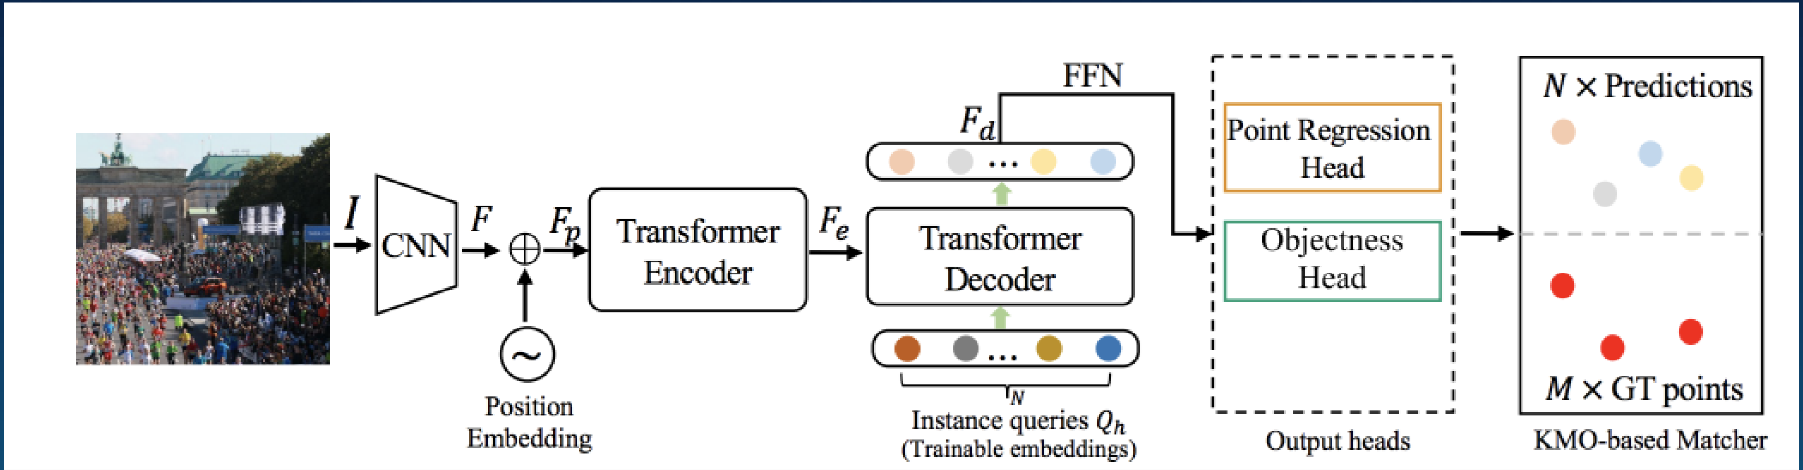

- State of the art for crowd localization
- But **cannot perform classification**

Liang, D., Xu, W. and Bai, X., 2022, October. An end-to-end transformer model for crowd localization. In European Conference on Computer Vision (pp. 38-54). Cham: Springer Nature Switzerland

Algorithm Details

Add a classification branch

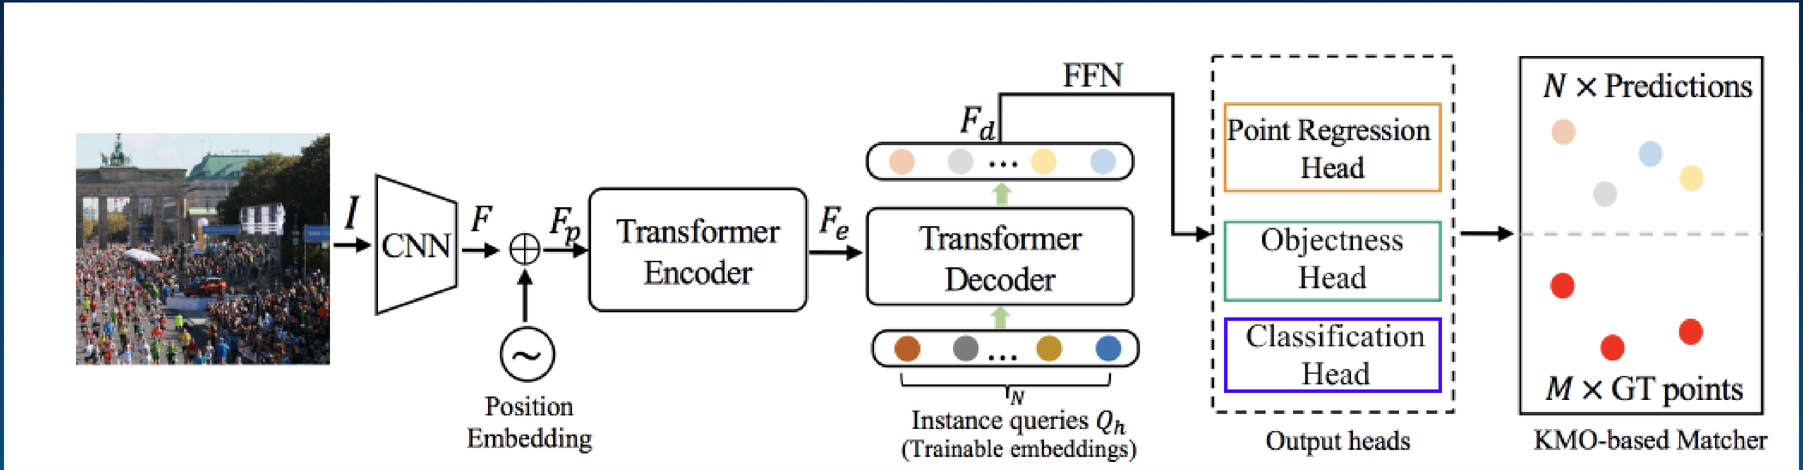

- Additional neural network layers
- Additional classification loss during training (softmax + cross-entropy)
- Now **each point has a species classification**

Metrics

AI.FISH

Two metrics were developed to evaluate our model in the context of the apportionment task:

1. **Dominant Species Accuracy**
2. **Weighted Classification Error**

We evaluated the model on its own (using our test dataset) as well as in comparison to trained reviewers.

Metrics

AI.FISH

Dominant Species Accuracy

Accuracy of predicting the most common class.

Typically an image is dominated by one species; predicting this species correctly will have the largest effect on overall accuracy.

This metric provides a simple “at a glance” measure of how well we do at identifying the majority class.

Metrics

Weighted Classification Error

A more complex metric that takes into account all species present as well as the apportionment goal.

Metrics

Weighted Classification Error

Ground Truth

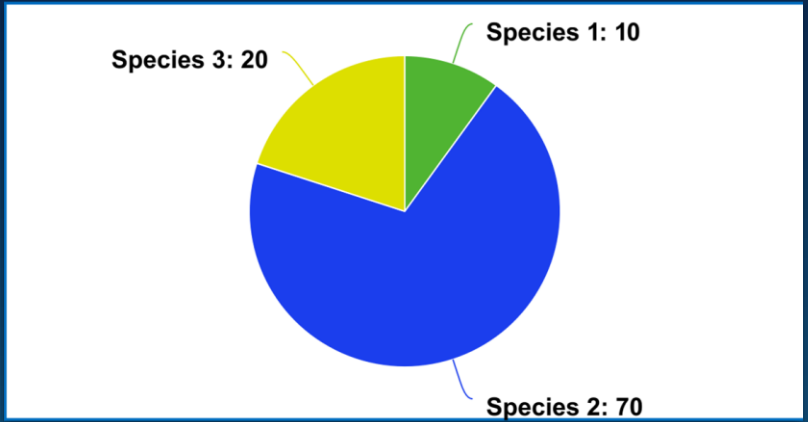

Predicted

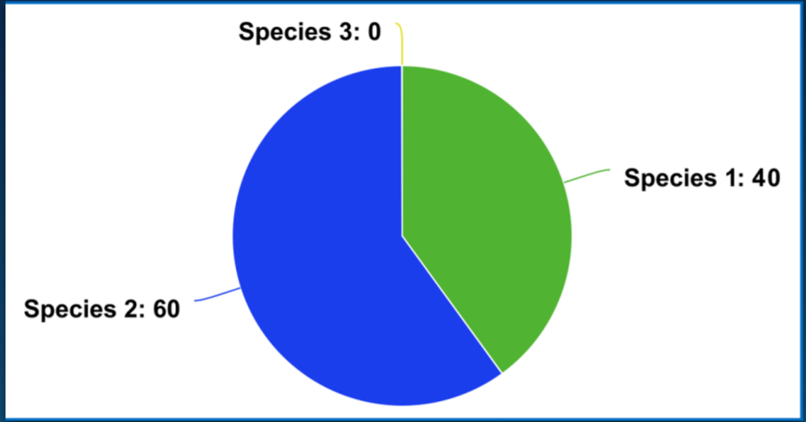

Metrics

Weighted Classification Error

Ground Truth

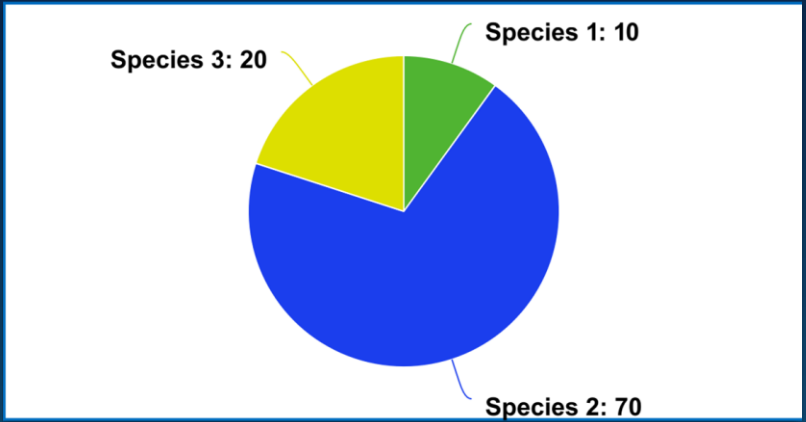

Predicted

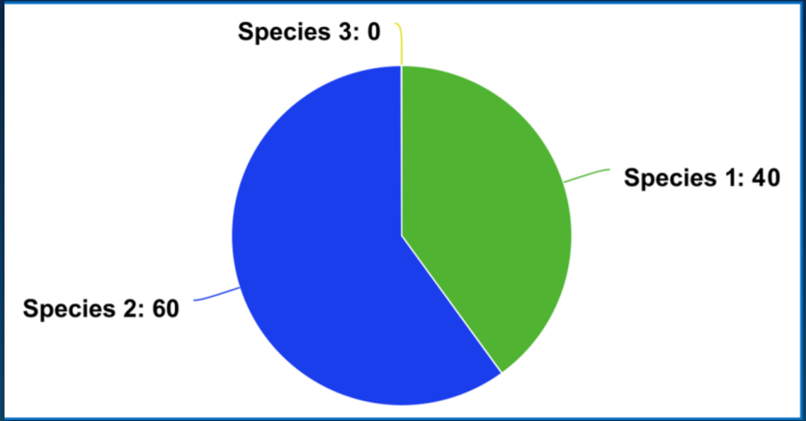

- **Mean absolute error:** (|40 - 10| + |60 - 70| + |0 - 20|) ÷ 3 = 20%

Metrics

Weighted Classification Error

Ground Truth

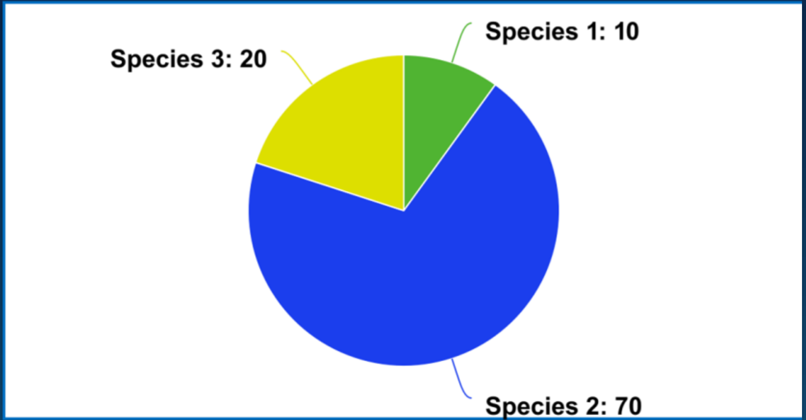

Predicted

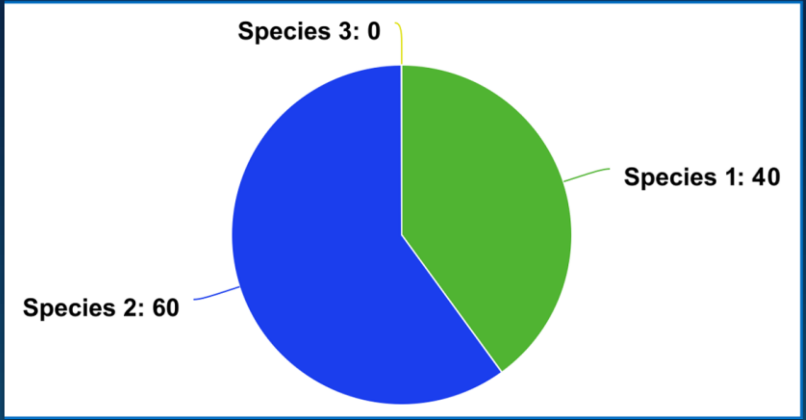

- **Mean absolute error:** (|40 - 10| + |60 - 70| + |0 - 20|) ÷ 3 = 20%
- But this gives equal weight to all classes, which might not be appropriate.
    - E.g. there are 5 classes total, but only 3 present; now divisor is 5, so error is artificially reduced

Metrics

AI.FISH

Weighted Classification Error

Ground Truth

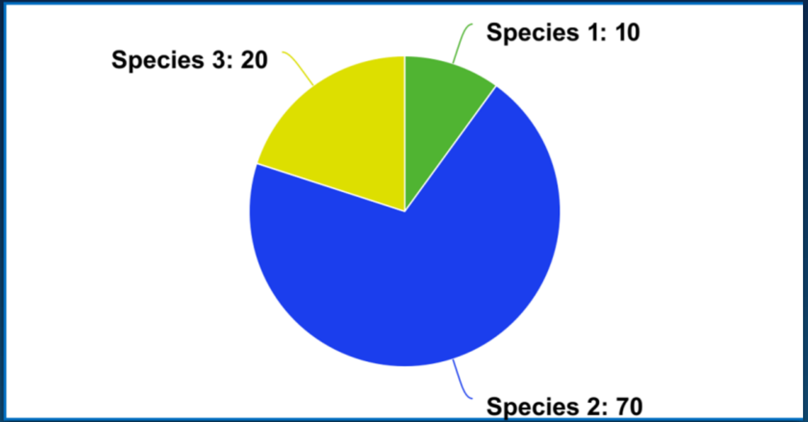

Predicted

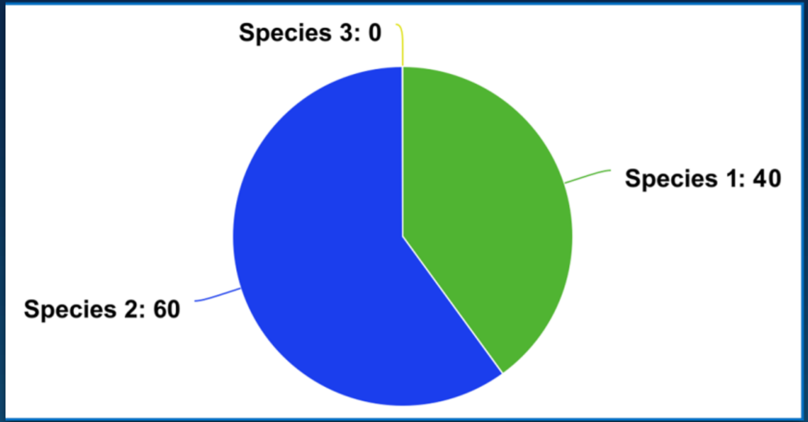

- **Weight by ground truth:** 0.1*|40 - 10| + 0.7*|60 - 70| + 0.2*|0 - 20| = 14%

Metrics

Weighted Classification Error

Ground Truth

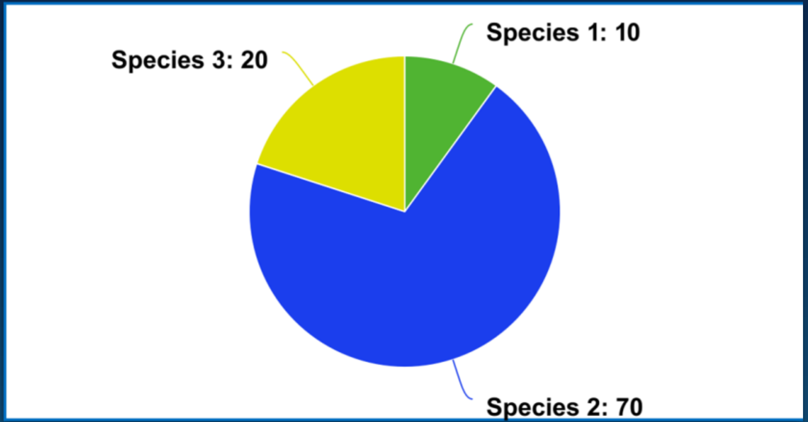

Predicted

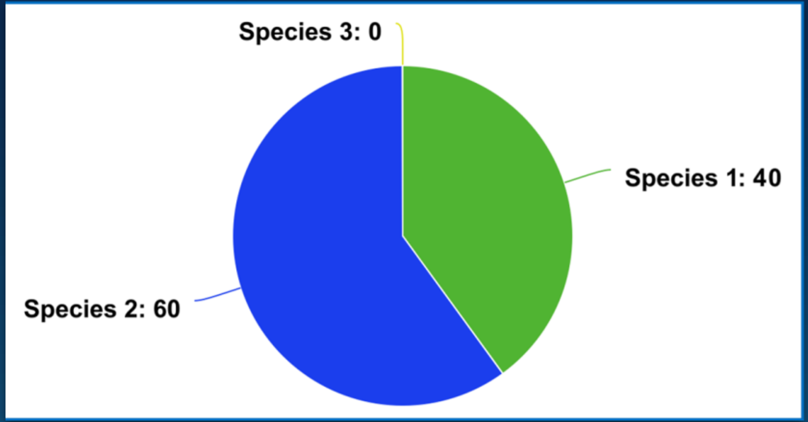

- **Weight by ground truth:** 0.1*|40 - 10| + 0.7*|60 - 70| + 0.2*|0 - 20| = 14%
- But what if a ground truth class is not present, but you predict it? Error for that class would be 0.

Metrics

Weighted Classification Error

Ground Truth

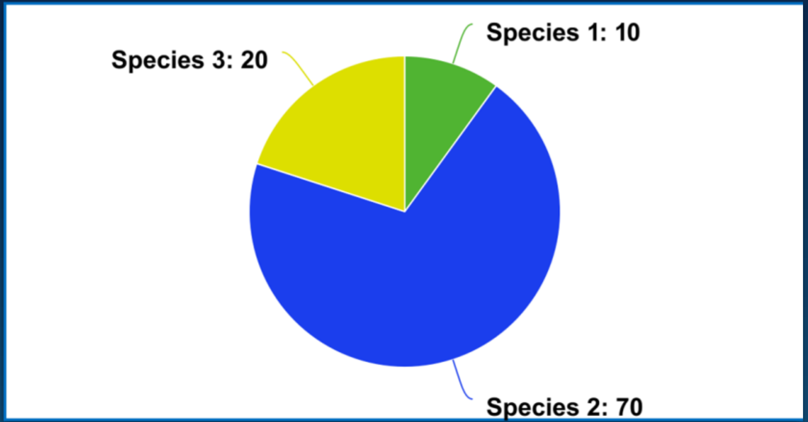

Predicted

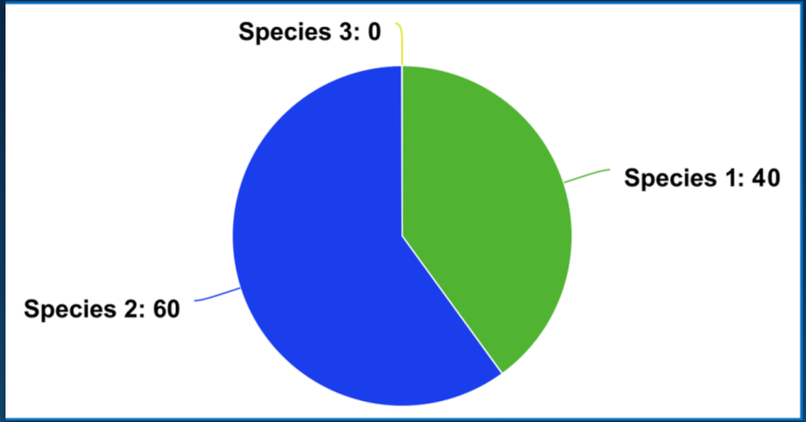

- **Weight by avg of GT + predicted:** 0.25*|40 - 10| + 0.65*|60 - 70| + 0.1*|0 - 20| = 19.2%

Metrics

Weighted Classification Error

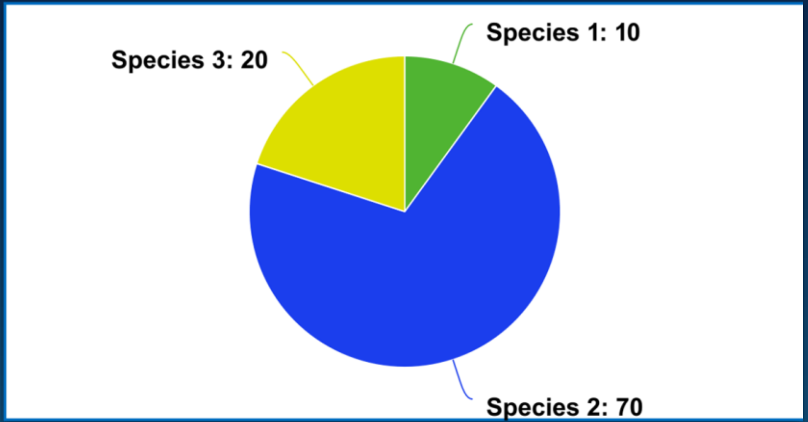

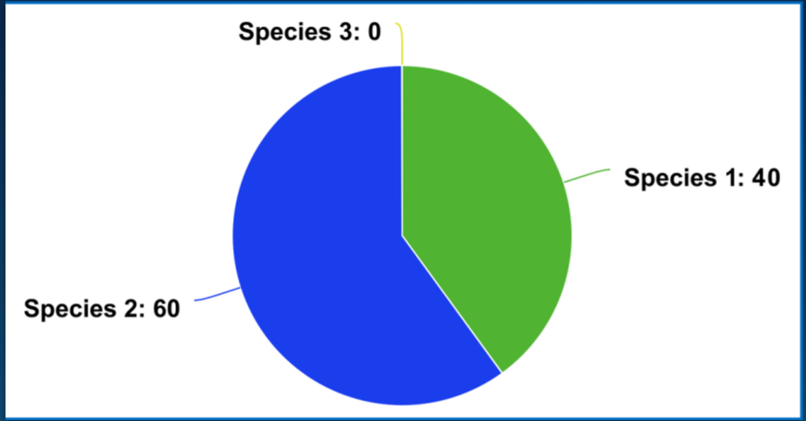

- **Weight by avg of GT + predicted:** 0.25*|40 - 10| + 0.65*|60 - 70| + 0.1*|0 - 20| = 19.2%
- Bonus: Now the measure is symmetric, so neither needs to be considered the "ground truth": we can compare the discrepancy of human reviewers, for example.

Results

AI.FISH

Evaluate:

- Model performance
- Human expert performance

Using:

- Test set of 100 held-out examples
    - Sampled from *tows not present in training data*

Results

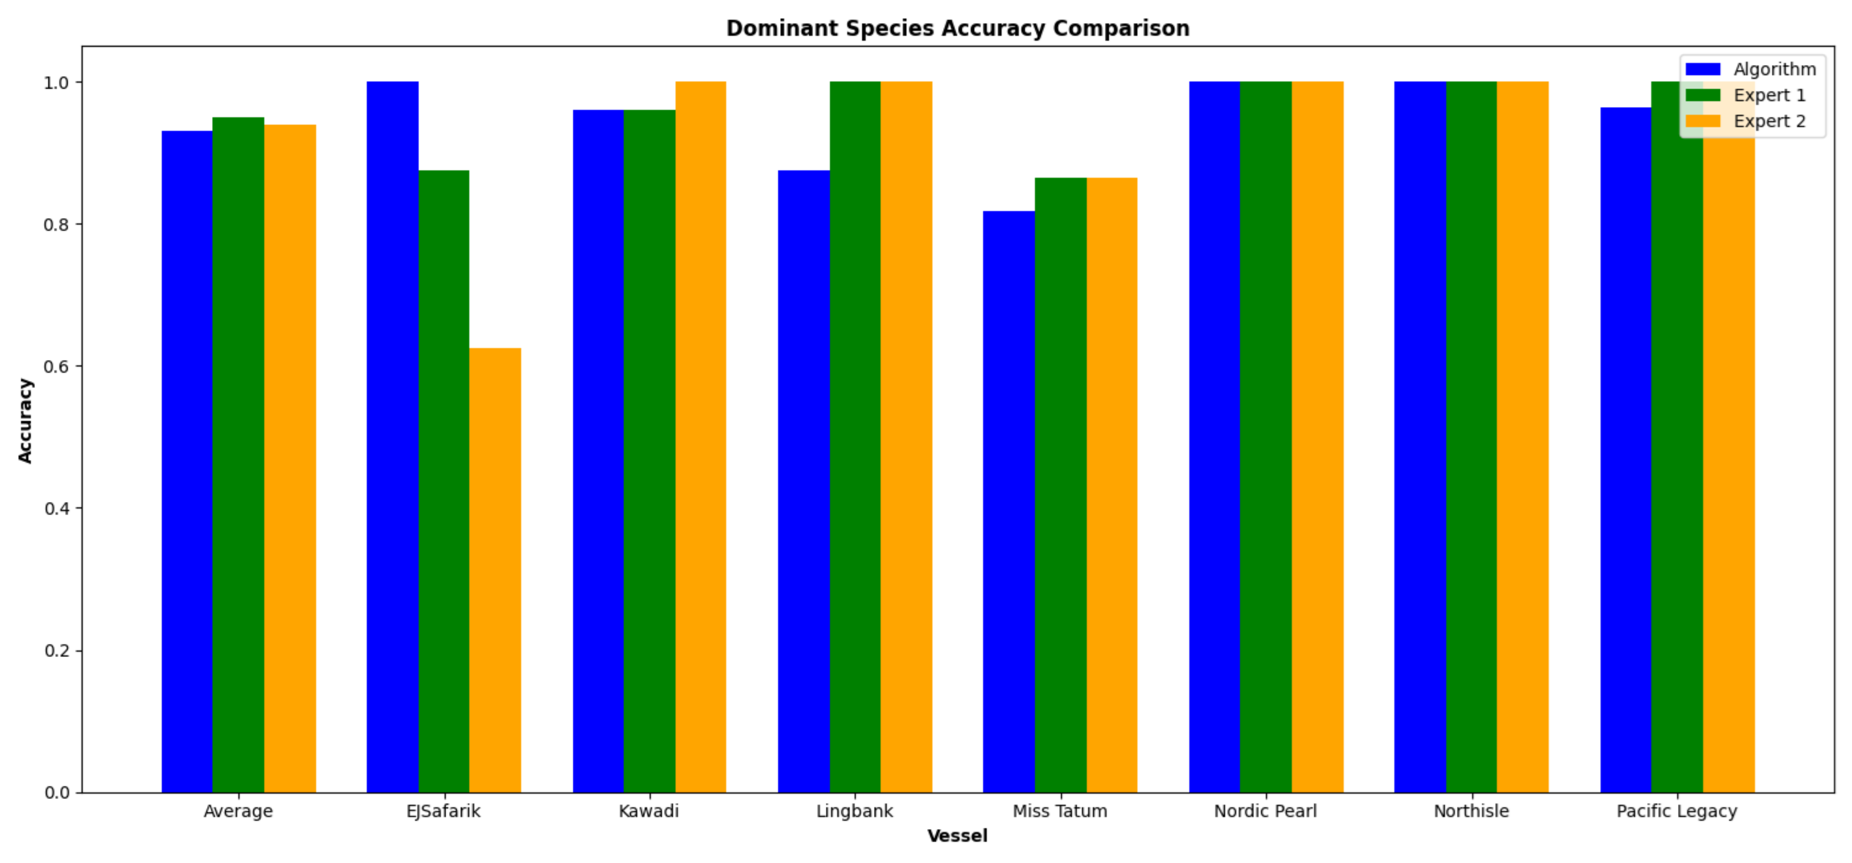
Dominant Species Accuracy Comparison
Algorithm
Expert 1
Expert 2
Accuracy
1.0
0.8
0.6
0.4
0.2
0.0
Average
EJSafarik
Kawadi
Lingbank
Miss Tatum
Nordic Pearl
Northisle
Pacific Legacy
Vessel

Results

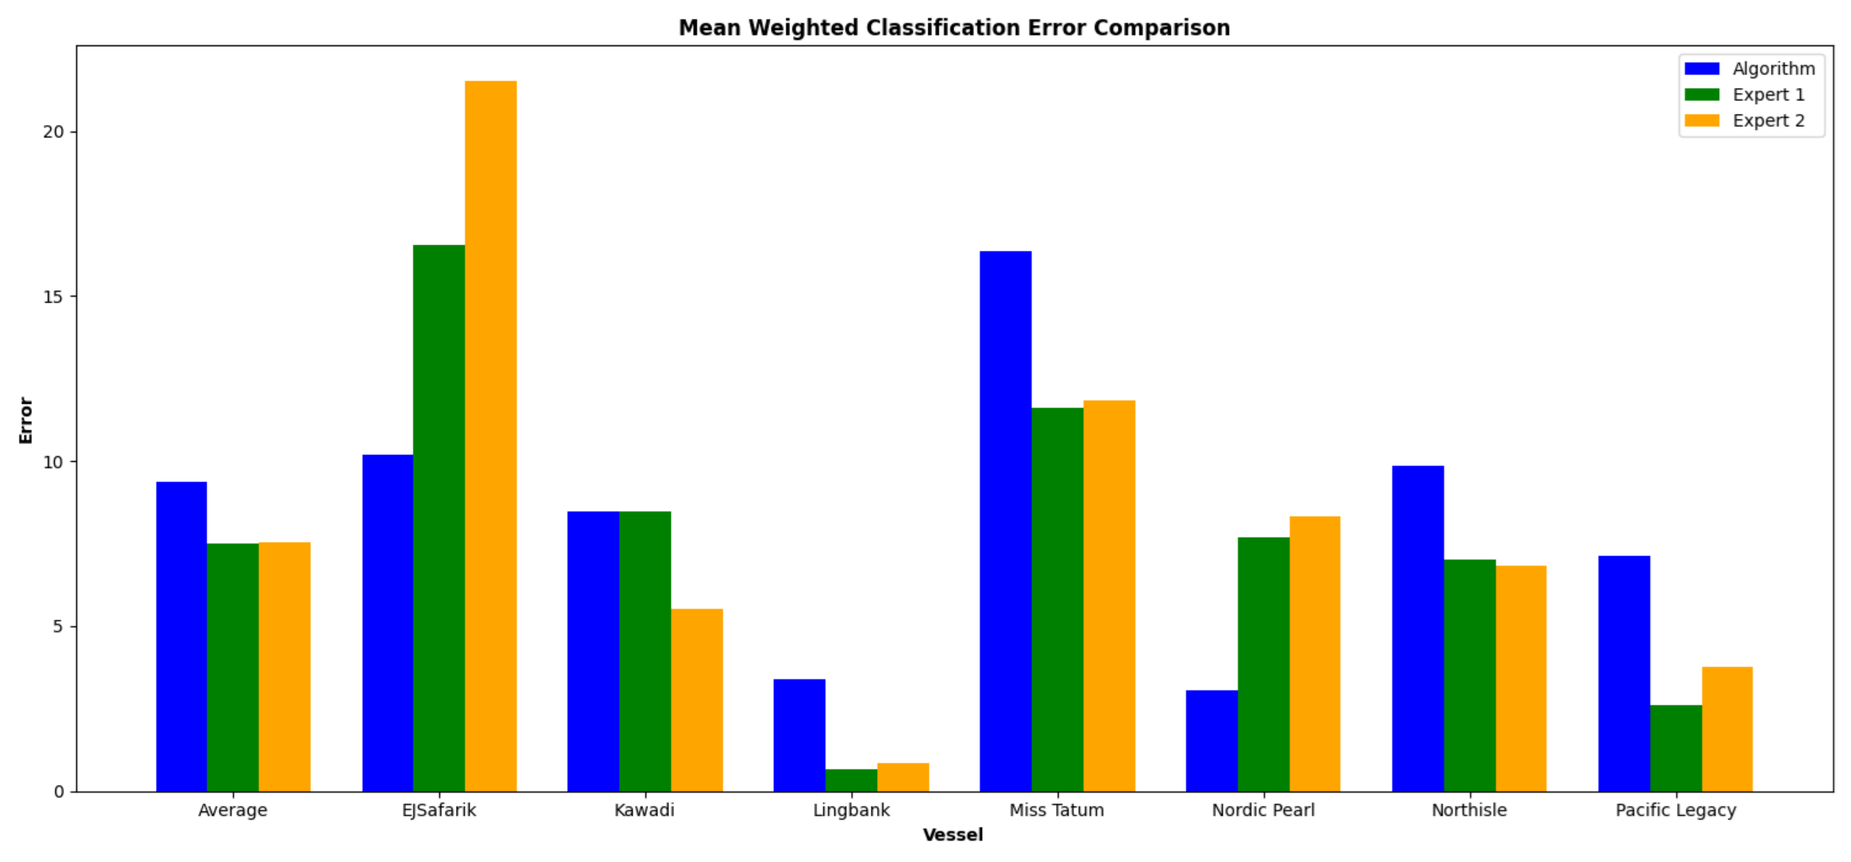
Mean Weighted Classification Error Comparison
Algorithm
Expert 1
Expert 2
Error
0
5
10
15
20
Average
EJSafarik
Kawadi
Lingbank
Miss Tatum
Nordic Pearl
Northisle
Pacific Legacy
Vessel

Results

AI.FISH

Metric	Algorithm	Expert 1	Expert 2
Dominant Species Accuracy	94%	95%	94%
Mean Weighted Classification Error	9.4%	7.5%	7.5%

- Algorithm achieves human expert-level performance on dominant species classification
- Algorithm is within 2% of human expert performance when considering mean weighted classification error
- Demonstrates the **feasibility of our approach** for producing **accurate automated apportionment estimates**