



OPEN

## FathomNet: A global image database for enabling artificial intelligence in the ocean

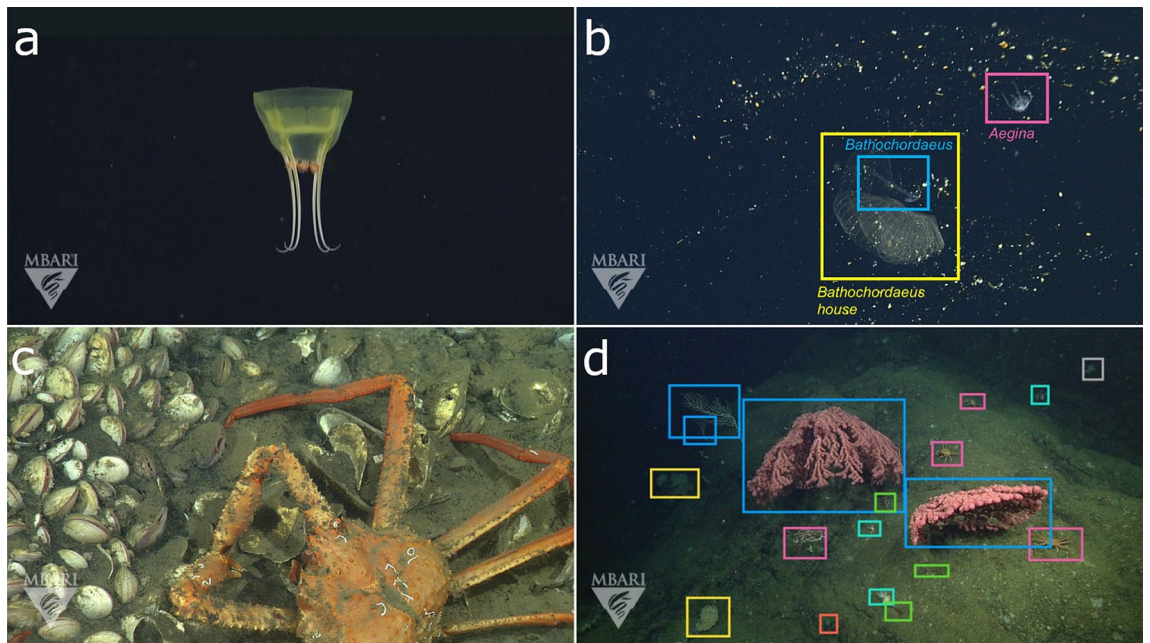
Kakani Katija<sup>1,2,3</sup>✉, Eric Orenstein<sup>1</sup>, Brian Schlining<sup>1</sup>, Lonny Lundsten<sup>1</sup>, Kevin Barnard<sup>1</sup>, Giovanna Sainz<sup>1</sup>, Oceane Boulais<sup>4</sup>, Megan Cromwell<sup>5</sup>, Erin Butler<sup>6</sup>, Benjamin Woodward<sup>6</sup> & Katherine L. C. Bell<sup>7</sup>

The ocean is experiencing unprecedented rapid change, and visually monitoring marine biota at the spatiotemporal scales needed for responsible stewardship is a formidable task. As baselines are sought by the research community, the volume and rate of this required data collection rapidly outpaces our abilities to process and analyze them. Recent advances in machine learning enables fast, sophisticated analysis of visual data, but have had limited success in the ocean due to lack of data standardization, insufficient formatting, and demand for large, labeled datasets. To address this need, we built FathomNet, an open-source image database that standardizes and aggregates expertly curated labeled data. FathomNet has been seeded with existing iconic and non-iconic imagery of marine animals, underwater equipment, debris, and other concepts, and allows for future contributions from distributed data sources. We demonstrate how FathomNet data can be used to train and deploy models on other institutional video to reduce annotation effort, and enable automated tracking of underwater concepts when integrated with robotic vehicles. As FathomNet continues to grow and incorporate more labeled data from the community, we can accelerate the processing of visual data to achieve a healthy and sustainable global ocean.

Monitoring a space as vast as the ocean<sup>1</sup> that is filled with life that we have yet to describe<sup>2</sup>, using traditional, resource-intensive (e.g., time, person-hours, cost) sampling methodologies are limited in their ability to scale in spatiotemporal resolution and engage diverse communities<sup>3</sup>. However, with the advent of modern robotics<sup>4</sup>, low-cost observation platforms, and distributed sensing<sup>5</sup>, we are beginning to see a paradigm shift in ocean exploration and discovery. This shift is evidenced in oceanographic monitoring via satellite remote sensing of near-surface ocean conditions and the global ARGO float array, where distributed platforms and open data structures are propelling the chemical and remote sensing communities to new scales of observation<sup>6,7</sup>. Due to a variety of constraints, large-scale sampling of biological communities or processes below the surface waters of the ocean has largely lagged behind.

There are three common modalities for observing biology and biological processes in the ocean—acoustics, “-omics”, and imaging—each with their strengths and weaknesses. Acoustics allow for observations of population- and group-scale dynamics, however individual-scale observations, especially determination of animals down to lower taxonomic groups like species, are challenging tasks<sup>8</sup>. The promising field of eDNA allows for identification of biological communities based on their shed DNA in the water column. While eDNA studies provide broad-scale views of biological communities with only a few discrete samples, determining the spatial source of the DNA, relating the measurements to population sizes, and the presence of confounding non-marine biological markers in samples are active areas of research that still need to be addressed<sup>9</sup>. Ultimately, -omics and acoustics approaches rely on visual observations for verification. Imaging, a non-extractive method for ocean observation, enables the identification of many animals to the species level, elucidates community structure and spatial relationships in a variety of habitats, and reveals fine-scale behavior of animal groups<sup>10</sup>. However, processing visual data, particularly data with complex scenes and organisms that require expert classifications,

<sup>1</sup>Monterey Bay Aquarium Research Institute, Research and Development, Moss Landing 95039, USA. <sup>2</sup>California Institute of Technology, Graduate Aerospace Laboratories, Pasadena 91125, USA. <sup>3</sup>Smithsonian Institution, National Museum of Natural History, Washington, DC 37012, USA. <sup>4</sup>NOAA, Southeast Fisheries Science Center, Key Biscayne 33149, USA. <sup>5</sup>NOAA, National Centers for Environmental Information, Stennis Space Center, St. Louis 39529, USA. <sup>6</sup>CVision AI Inc., Research and Development, Medford 02155, USA. <sup>7</sup>Ocean Discovery League, Saunderstown 02874, USA. ✉email: kakani@mbari.org



**Figure 1.** Labeled image data for deep learning algorithms require an annotation and localization for a single concept. Images showing the disparity between (a, c) traditionally annotated visual data in MBARI's VARS database and (b, d) the level of annotations and localizations (or bounding boxes) that are required for the data science community. Top (a, b) and bottom (c, d) rows show images for midwater and benthic environments, respectively. (a, b) and (c, d) are example images from VARS that have been assigned a single-concept annotation for the genera *Aegina* (b, pink) and *Chionoecetes* (d, pink), respectively. Although images in the left column (a, c) are largely iconic views of a single-concept, images in the right column (b, d) show non-iconic views of multiple concepts that additionally include (b) *Bathochordaeus* (blue), *Bathochordaeus house* (yellow), (d) *Paragorgia* (blue), *Asteroidea* (pink), *Psolus* (cyan), *Pandalopsis* (green), *Heterochone* (yellow), and *Tunicata* (red). Label text has been removed for clarity in (d).

is a resource-intensive process that cannot be scaled without significant investment, capacity building, and advances in automation<sup>11,12</sup>.

Imaging is an increasingly common modality for sampling biological communities in a variety of environments due to the ease by which the technology can be deployed, and the number of remotely controlled and autonomous platforms that can be used<sup>13</sup>. Imaging has also been used for real-time, underwater vehicle navigation and control while performing difficult tasks in complex environments<sup>14</sup>. Moreover, imaging is an effective engagement tool to share marine life information and the issues facing the ocean with broader communities<sup>15,16</sup>. In short, visual data is an invaluable tool for better understanding the ocean and conveying that information broadly.

Given all the applications of marine imaging, a number of annotation tools have been developed to manage and analyze visual data. These efforts have resulted in many capable software solutions that can either be deployed locally on a computer, during field expeditions, or broadly on the worldwide web<sup>17</sup>. However, with the limited availability of experts and the prohibitive costs to annotate and store footage, novel methods for automated annotation of marine visual data are desperately needed. This need is motivating the development and deployment of artificial intelligence and data science tools for ocean ecology.

Artificial intelligence (AI) is a broad term that encompasses many different approaches<sup>18</sup>, some of which have already been used to study marine systems. Statistical learning methods like random forests have been used in the plankton imaging community, achieving automated classification of microscale plants and animals with accuracies greater than 90%<sup>11</sup>. Unsupervised learning can be deployed with minimal data and a priori knowledge of marine environments, however these algorithms have limited application for automating the detection and classification of objects in marine imagery with sufficient granularity and detail to be used for annotation<sup>19</sup>. Deep learning algorithms trained on visual data where all objects have been identified (e.g., Fig. 1) have improved performance of automated annotation and classification tasks to finer taxonomic levels<sup>20,21</sup>, however this approach requires publicly available, large-scale labeled image datasets for training<sup>22–24</sup>.

Image repositories for terrestrial applications have been available to the computer vision (CV) community for many years. ImageNet was the first labeled set based on a hierarchy, or the number of classes (or “things”) in WordNet, with the long-term goal to collect 500 to 1 k full-resolution images for 80 k concepts, or ~ 50 M images<sup>22</sup>. In order to reach this scale, ImageNet used images scraped from Flickr, resulting in a collection of largely iconic images (e.g., centered objects in relatively uncluttered environments). Like ImageNet, Microsoft's COCO<sup>23</sup> used workers with Amazon's Mechanical Turk to generate labels (5 k labeled instances) for images with 90 classes, resulting in 2.5 M labeled instances in more than 328 k images. More recently, iNat2017 is a biologically

	Data entry categories		
	Required	Recommended	Suggested
Collections	Owner's institution	Bibliographic citation	Collection code
	Rights holder (use owner's institution if not specified)	Access rights (for more generous use)	Collection ID
	Contributor's email	Basis of record (manual or machine)	Dataset generalizations
	Record type (images)	Dataset language	Dataset name
	Modified field (upload date)		Dynamic properties
	UUID		Information withheld
	URL		Institution code
	Data format (CSV+)		Institution ID
			References
Images	Image URL	Latitude/Longitude	Group of
	Bounding box coordinates (x, y, width, height)	Depth	Occluded
	Concept name	Timestamp (ISO8601)	Truncated
		Imaging type	AltConcept
		Observer	
		Altitude	

**Table 1.** Data entry categories for Collections (corresponding to a single upload) and Images that make up the collections, divided into required, recommended, and suggested fields. The data entry categories for Collections are from the Darwin CORE Archive data standard.

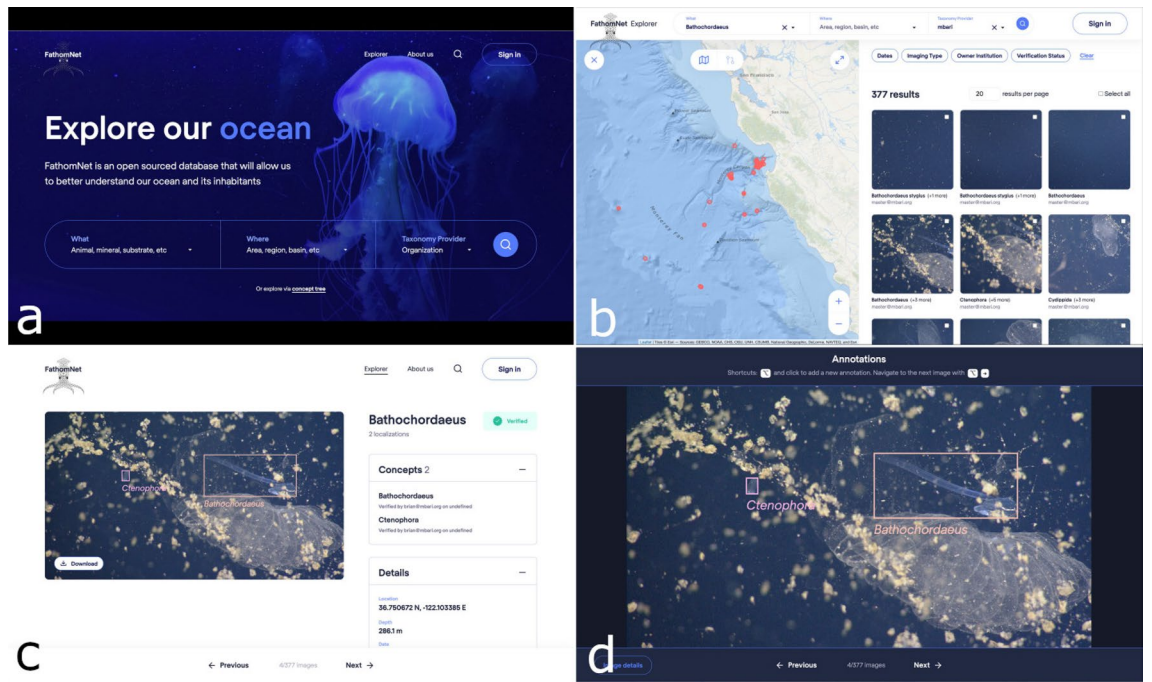
focused dataset built from 675 k images of 5 k species of animals that have been collected and verified by users of iNaturalist<sup>25</sup>. Unlike ImageNet and iNat2017, COCO was specifically built to include non-iconic views of “things” that provides imagery with contextual relationships between objects, which is especially relevant to marine environments.

Large, publicly available labeled image datasets for marine concepts primarily represent planktonic communities. Imaging at these spatial scales (~10–1000 s of microns) requires controlled lighting and imaging conditions that utilize darkfield, brightfield, or holographic illumination<sup>26</sup>, and these datasets include WHOI-Plankton<sup>27</sup> (3M images representing 109 concepts collected by 1 imaging system) and EcoTaxa<sup>11</sup> (150 M images collected by 10 imaging systems) among others. While the plankton imaging community has made significant progress in creating image data portals, these large datasets are primarily built for classification tasks (i.e., contains regions of interest only) and necessarily exclude larger animals and other marine animals found in midwater and benthic environments. CoralNet<sup>28</sup>, a portal and platform for working with coral imagery, is of a similar order of magnitude to these plankton datasets but is likewise restricted to a particular type of organism. There are no equivalent datasets for macroscopic ocean organisms; image sets of higher trophic-level animals<sup>29</sup> are typically distributed in individual repositories that can be difficult to find for non-subject matter experts in the CV and AI communities. Thus, there is a clear need for a labeled image dataset that is representative of diverse biological communities across spatial scales in the ocean that can be readily accessed in a single, publicly available online repository (Fig. 1).

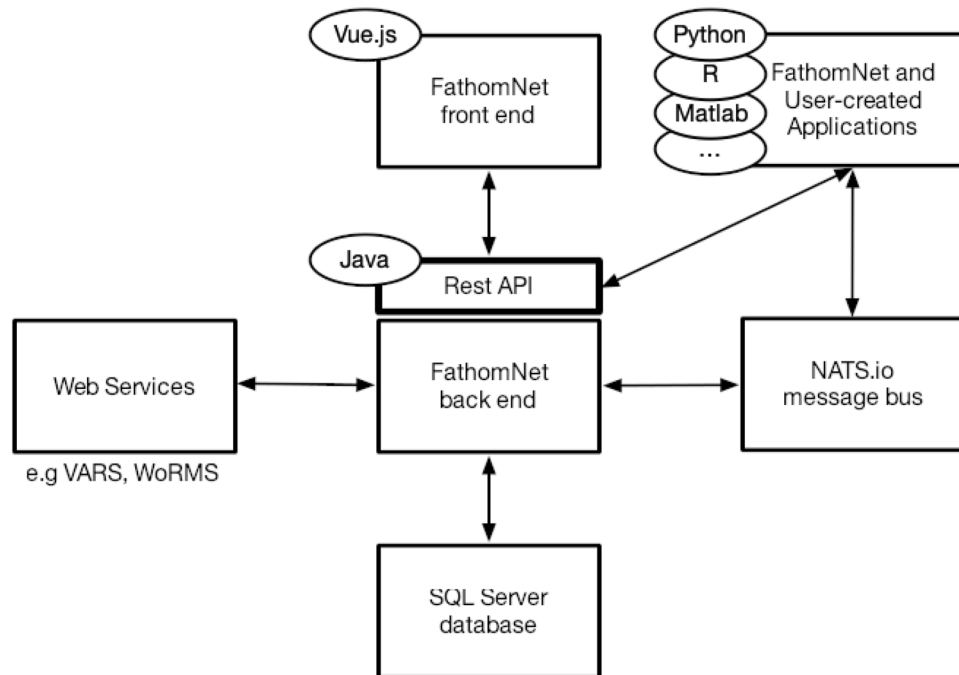
## Results

In order to process imagery and video collected anywhere in the world's ocean, we need a comprehensive labeled image dataset that can scale to the global ocean (Fig. 1). If we want to enlist individuals that are not subject-matter experts in marine imaging but rather computer scientists that are familiar with developing state-of-the-art algorithms for automated image analysis, we need an accessible central repository for this labeled data<sup>22,23</sup>. To address this need, we developed FathomNet, a distributed, publicly available labeled image dataset built on FAIR data principles<sup>30</sup>, that uses community-recognized Darwin CORE archive data formats (Table 1)<sup>31</sup>. FathomNet has a Rest API and website (www.fathomnet.org; Fig. 2) backed by a relational database (Fig. S1) that integrates widely used community web services (e.g., World Register of Marine Species or WoRMS<sup>32</sup>; Fig. 3).

**FathomNet database architecture.** FathomNet is a SQL server database that utilizes various web services (e.g., VARS, WoRMS, MarineRegions<sup>33</sup>), and NATS.io message bus (Fig. 3). Data can either be accessed through the FathomNet website<sup>34</sup> (Fig. 2) or via the FathomNet Rest API. FathomNet currently utilizes either WoRMS<sup>35</sup> or MBARI's knowledgebase (Fig. S1), and can accommodate other taxonomies that have existing APIs. Images, image URLs, and associated metadata (Table 1) can be uploaded on the FathomNet database via either the website using a CSV file and appropriate Darwin CORE archive fields (Table 1), or the Rest API. Fields such as *imaging type* and *alt concept* allow users to provide additional details about their data, including whether collected by particular platform or imaging system (e.g., ROV, Baited Remote Underwater Video<sup>36</sup>, or Underwater Video Profiler<sup>37</sup>), or identifying subfeatures (e.g., head or tail) for a particular concept, respectively. Submission of data is open to anyone granted write access to the database, and data will undergo a quality control process of verification by members of the FathomNet community.

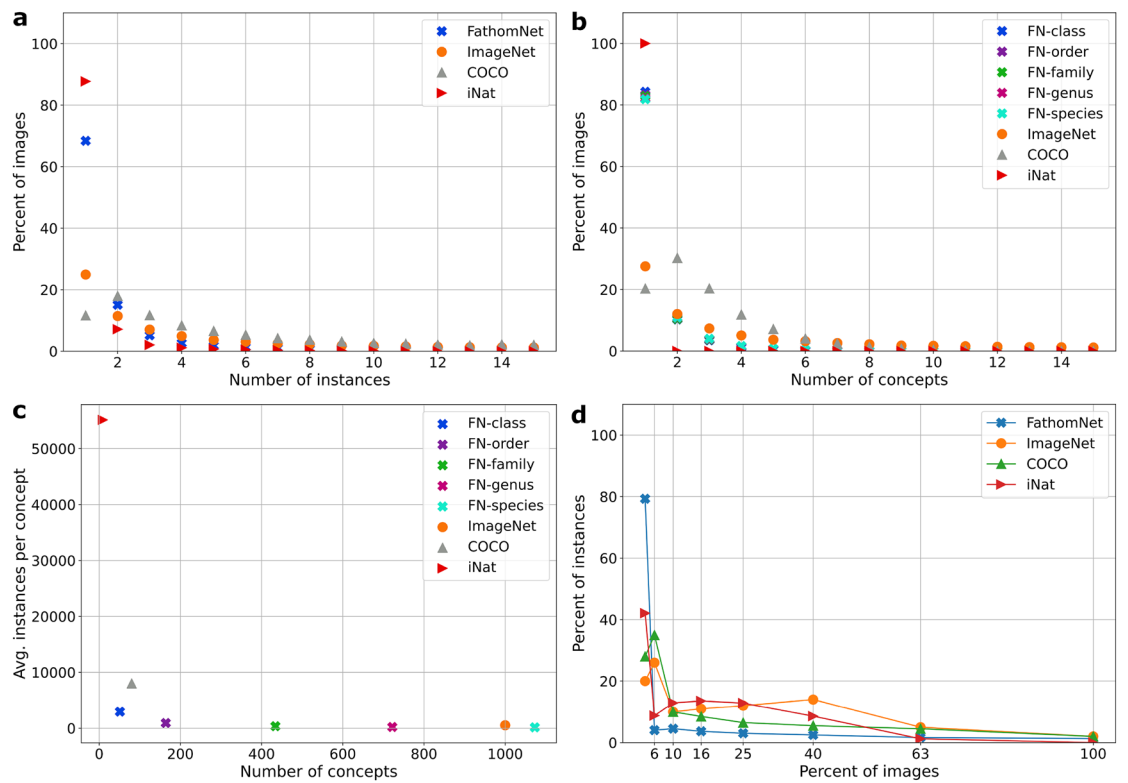


**Figure 2.** FathomNet data portal is accessed at [www.fathomnet.org](http://www.fathomnet.org). The website contains features that include (a) a simple search bar for terms in the concept tree, (b) filtered searches where images can be displayed based on geographic location or terms within the concept tree (among others), (c) image display pages where concepts, details, and contributors' information is shown, and a (d) basic annotation and localization tool to allow users to augment or correct uploaded data in the database.



**Figure 3.** FathomNet database software architecture. The FathomNet software architecture includes a Java Rest API, SQL Server database, a NATS.io messaging bus, and external web services (e.g., WoRMS, VARS, MarineRegions).

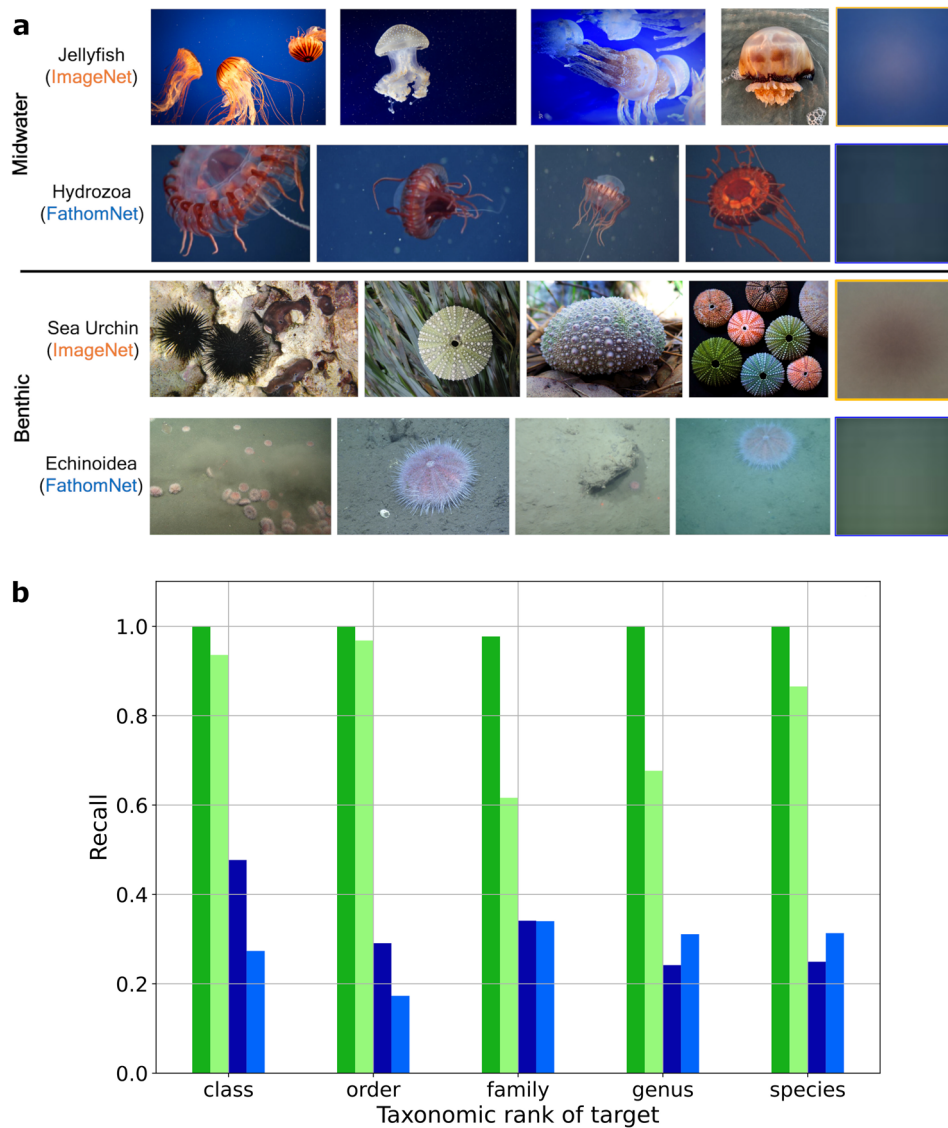




**Figure 4.** Dataset statistics contextualizing FathomNet relative to COCO, ImageNet, and iNat2017. Statistics for COCO were computed from the 2014 training split, ImageNet numbers were generated from the 2015 localization data split, and iNat2017 metrics were derived from the 2017 training data. **(a)** Percent of images containing a number of instances, or localizations. **(b)** Percent of images displaying a given number of concepts. FathomNet annotations are displayed broken down at different levels of taxonomic rank. **(c)** The average number of instances versus the average number of concepts over the entirety of each dataset. **(d)** The distribution of relative instance size. Most objects in FathomNet and iNat are small compared to ImageNet and COCO.

The FathomNet database seeks to aggregate marine labeled image data for more than 200k currently accepted species of *Animalia* in the WoRMS database<sup>32</sup> using community-based taxonomic standards<sup>38</sup>. Our goal is to obtain 1000 independent observations for each species in diverse poses and imaging conditions, resulting in more than 200M observations that will continue to grow as the number of described species increase. While we use species in *Animalia* as an initial target, the FathomNet concept tree (Fig. S1) is currently based on MBARI's Video Annotation and Reference System (VARS) knowledgebase<sup>39</sup>, and can be expanded beyond biota to include underwater instances of equipment, geological features, marine debris, etc. Additional taxonomic hierarchies<sup>40,41</sup> can be wired into FathomNet via respective web service API calls. By utilizing existing annotation tools and providing functionality that will enable reviewing, editing, augmenting, and verifying submitted data (Fig. 2), we can aggregate existing underwater image datasets that can be added to the database.

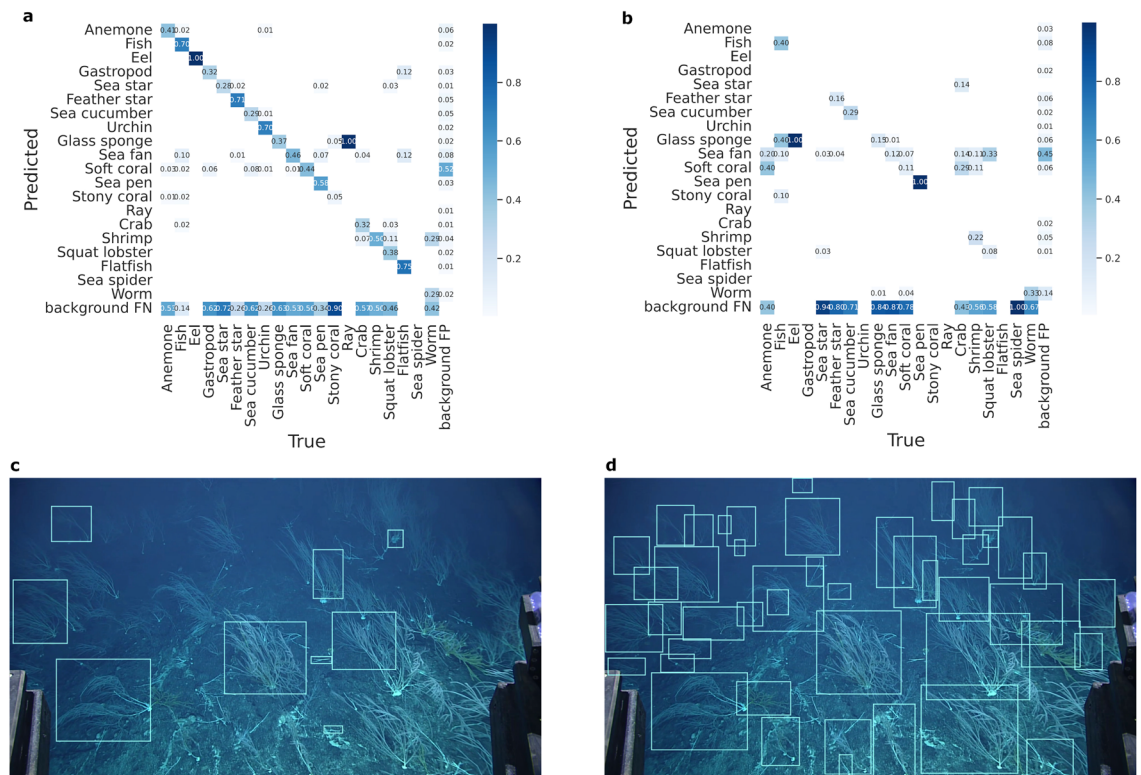
**FathomNet dataset statistics.** FathomNet is seeded with data from three sources, collectively representing 30+ years of ROV, AUV, drop camera, and stereo imaging deployments: (1) MBARI's VARS database, (2) National Geographic Society's (NGS) Deep Sea Camera System<sup>42</sup>, and (3) the National Oceanic and Atmospheric Administration's (NOAA) ROV *Deep Discoverer*<sup>43</sup>. As of July 2022, FathomNet has 84,454 images, 175,873 localizations from 204 separate collections for 2244 concepts, with additional contributions ongoing. This snapshot of the database is between ImageNet<sup>22</sup> and COCO<sup>23</sup> in terms of the numbers of categories and instances per image (Fig. 4a–c). The 13 super categories in the iNat2017<sup>25</sup> benchmark dataset (collapsed from all 5,039 fine-grained concepts of terrestrial organisms from iNaturalist) contains many more instances per concept relative to FathomNet at the class level (Fig. 4c), however FathomNet has more diversity in terms of the number of concepts and instances per image (Fig. 4a, b). Relative to ImageNet, COCO, and iNat2017, the localizations in FathomNet typically occupy less of the frame and occur in diverse poses due to inclusion of iconic and non-iconic imagery (Figs. 4d, 5a). Coverage, a characteristic that defines whether all objects within an image are annotated, varies depending on the environment the image was captured (e.g., benthic or midwater) and the taxonomic rank (Fig. 5b). Images of midwater organisms are typically more completely annotated than those taken on the seafloor on a per frame basis, and is largely a function of the population density of benthic organisms such as deep sea corals and urchins (Fig. 1).



**Figure 5.** Dataset visualizations contextualizing FathomNet relative to ImageNet. **(a)** The average pixel values of images in ImageNet and FathomNet. Images from a midwater and benthic concept in ImageNet (hosted on Flickr) were approximately matched to concepts available in FathomNet, and an equal subset of images from both repositories were resized and averaged pixel-wise (far-right column). The uniform image patches from FathomNet suggest that there is more diversity in poses and object sizes relative to labeled data from ImageNet. FathomNet image credits: MBARI. **(b)** Concept coverage at different taxonomic ranks within FathomNet. Fifty FathomNet images were selected randomly at each level of the taxonomic tree and checked for annotation completeness by a human expert. Green bars correspond to the midwater concept *Bathocordaeus mcnutti* and blue bars correspond to the benthic concept *Gersemia juliepackardae*. Darker colors are the target concept and lighter colors are any other concepts present in the frame. Benthic images are typically missing more annotations than those from midwater habitats.

**Representative FathomNet use cases.** To demonstrate the utility of FathomNet, we present three forward-looking use cases: object detection of animals in underwater images, activity recognition (presence of objects) in underwater video, and machine learning algorithms integrated in underwater vehicle controllers<sup>14</sup> for automated animal tracking.

*Benthic animal detection using NOAA footage and MBARI's training data.* One of the benefits of FathomNet will be to use the training data and models to supplement the annotation of video and imagery collected in other regions of the ocean by other researchers and imaging systems. To evaluate this functionality, we fine-tuned a RetinaNet object detection model with a ResNet backbone on 20 high-level benthic supercategories (urchin, fish, sea cucumber, anemone, sea fan, sea star, worm, sea pen, crab, glass sponge, shrimp, ray, flatfish, squat lobster, gastropod, eel, soft coral, feather star, sea spider, and stony coral) generated by combining approximately



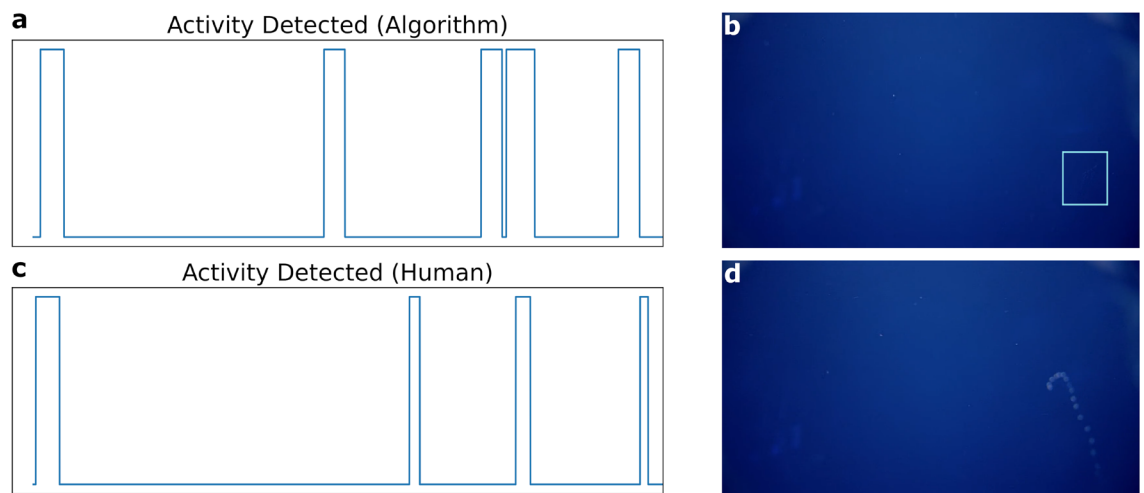
**Figure 6.** Results from the NOAA Benthic use case. **(a)** Validation confusion matrix from the RetinaNet object detector trained on supercategory data from Monterey Bay. The row of background false negatives is mostly due to incomplete original annotation data. **(b)** Confusion matrix from the model applied to NOAA data from the Musician Seamount. Note that the vast majority of the targets were sea fans. **(c)** Model output from a single frame of NOAA video data. **(d)** Ground truth annotations for the same frame. Note the density, variability in coral morphology, and overlapping individuals.

1k taxonomic classes found in the MBARI taxonomic knowledgebase (Fig. S1)<sup>44</sup>. We hypothesized that (1) the high-level taxonomic model would perform more robustly than one including all fine-grained classes and (2) would be extensible to other regions despite training data coming mostly from a single region (the Northeast Pacific). The benthic model<sup>44</sup> was then used to generate algorithm proposals on frame grabs from the NOAA CAPSTONE project<sup>45</sup> during an expedition to the Musician Seamounts. Experts in MBARI's Video Lab subsequently annotated and localized a subset of these frame grabs, and the data were compared with the model-generated proposals.

The model performed reasonably well when tested on images from Monterey Bay for the concept coverage analysis (Fig. 6a). The row of background false negatives is an artifact of the original training data coverage. The model recorded many background false positives and negatives in the target NOAA data, sometimes overwhelming the number of correct annotations of a given class (Fig. 6b). While there is substantial overlap between the data collected at the Musician Seamount and in Monterey Bay—images collected by ROVs of benthic communities—many subtle differences confounded the automated detector: slight changes in the angle of the camera relative to the sea-floor, new taxonomic classes that were not represented in the training data, and higher densities of different organisms (Fig. 6c, d). The behavior of the object detector is consistent with the effects of distribution shift, where the statistics of the target data change relative to the training data<sup>46</sup>. Indeed, the density and morphology of the corals represented by the sea fan supercategory along the Musician Seamount presented challenges to the detector (Fig. 6c, d). The output of the MBARI model on the NOAA data is not reliable on its own, but substantially eases the task of human annotators looking to begin annotation efforts on raw data.

*Midwater transect activity detection using NOAA footage and MBARI's training data.* An activity recognition (AR) routine was deployed using a midwater object detector<sup>47</sup>: a RetinaNet model with a ResNet50<sup>48</sup> backbone pre-trained on ImageNet<sup>22</sup>. Fine-tuning of parameters was done with in situ, underwater imagery from FathomNet and VARS of 17 target midwater animal classes. For the purpose of activity recognition, the model<sup>47</sup> was used as a binary object detector, and results were smoothed to identify segments of video for a human operator to review and annotate. This is a relevant use case for reviewing footage after it has been collected (e.g., post-processing of ROV expedition video), as it can be tedious and costly for an individual to watch many hours of video looking for only a few events. Presenting annotators with segments of video that only contain animals that need annotation can decrease the amount of time it takes to thoroughly annotate video footage<sup>49</sup>.

The AR routine was applied to video data collected by NOAA's ROV *Deep Discoverer*, during an expedition to the Musician Seamounts as part of the NOAA CAPSTONE project<sup>45</sup>. An MBARI researcher manually annotated



**Figure 7.** Comparison of activity recognition output from object detection (a) and a human expert (c) for one video collected by NOAA's ROV *Deep Discoverer* during the CAPSTONE expedition<sup>45</sup>. In one instance (b), the algorithm (indicated by a bounding box) correctly identifies an animal with low contrast (image has been contrast enhanced). In another (d), the algorithm fails to identify an animal that is blurry due to unfocused optics.

one hour of video with segments of interest, and these segments were compared against the AR model-generated segments (Fig. 7, left). Activity detection signals were filtered by convolving with a 10 second window to fill any gaps between object detections. Intersection over union (IOU), which takes the temporal overlap of AR routine outputs and divides by total activity duration that occurs in either method, was used to quantitatively compare approaches. The IOU over all annotated videos was 43% and event-level recall (accounting for the imperfect temporal registration between activities) was 63%. A review of these results found that the object detector did not perform well when the object was blurred (either due to motion or unfocused optics; Fig. 7, bottom-right) or when the animal was small relative to total field of view (animal  $\sim 50$  pixels in length); in some cases (e.g., animal lacked sufficient contrast to the background) the algorithm correctly identified activities that the human annotator missed (Fig. 7, top-right). Although the AR routine presented here requires further refinement (Video S1), 19% of the footage contained animals or objects of interest, and presenting annotators with video clips of interest would similarly reduce annotator effort.

*Machine learning-integrated tracking (ML-Tracking) of animals for vehicle navigation and control.* In order to observe animal behavior in the ocean, observational platforms are required to non-invasively execute targeted sampling and maintain a persistent presence to track animals over a period of time. Vision-based underwater vehicle tracking has seen renewed interest due to developments in modern computer vision and machine learning<sup>50–53</sup>, and methods like tracking-by-detection and integrating neural networks have shown promise for robust tracking<sup>54</sup>. For applications where longer duration, 24+ hr-long deployments are required (e.g., studying critical behaviors of animals particularly in the ocean's midwaters<sup>55</sup>), these modern techniques are needed.

Katija et al.<sup>14</sup> developed a Machine Learning-integrated Tracking (ML-Tracking) algorithm that incorporates a FathomNet-trained multi-class *RetinaNet*<sup>56</sup> detection model<sup>47</sup>, 3D stereo tracker subroutines, and a supervisor module that sends commands to a vehicle controller. ML-Tracking algorithms were demonstrated in midwater using ROV *MiniROV* in the Monterey Bay National Marine Sanctuary. Nearly 50 hours of stereo and high-definition footage were collected by the authors to evaluate the ML-Tracking algorithm<sup>14</sup>. The longest continuous tracking trial involved a gelatinous animal—a siphonophore *Lychnagalma* sp.—for a duration of 18,987 s or 5.27 h. Applying ML-Tracking to the long-duration observation, the percentage of time the vehicle controller received 3D position information from the ML-Tracking algorithm (binned at 1 s intervals) was 100%<sup>14</sup>.

## Discussion

FathomNet is built to aggregate data from disparate sources and accelerate discovery by giving the community access to more data and expertise. Fundamentally, the system relies on fine-grained, taxonomically correct annotations, which is a difficult standard to meet via crowd-sourcing<sup>57</sup>. The initial release of FathomNet is thus built on marine domain experts' annotation efforts; in contrast, ImageNet and COCO are built from images scraped from publicly available internet repositories and annotated by crowd-workers. Based on our estimates (see Methods), the COCO and ImageNet annotation efforts cost  $\sim$  \$295 k and  $\sim$  \$852 k, respectively, excluding the cost of image generation and other infrastructure. In contrast, MBARI's seed contribution to FathomNet cost  $\sim$  \$165 k to annotate based on estimates from MBARI's Video Lab, where experienced annotators label midwater and benthic images at  $\sim$  \$1 and  $\sim$  \$3 per image, respectively. At that same per-image rate, ImageNet would cost  $\sim$  \$6.15 M to annotate. Importantly, FathomNet currently draws largely from MBARI's VARS database, which is comprised of 6190 ROV dives that represents  $\sim$  \$144 M worth of ship time and excludes the instrument and



platform costs used to obtain that data. Including these additional costs underscores the current and future value of FathomNet, especially to groups in the ocean community that are early in their data collection process.

For FathomNet to reach its intended goals, significant community engagement, high-quality contributions across a wide range of groups and individuals, and broad utilization of the database will be needed. FathomNet can be leveraged to develop fully and semi-automated workflows that assist, but not replace, the annotation efforts of expert taxonomists and ecologists. This vision is very much in accord with decades-old aspirations of the marine biological community<sup>58</sup>, and recent findings that even state-of-the-art machine learning systems suffer from biases<sup>24</sup> that can be mitigated with human intervention<sup>59</sup>. In addition to the database itself, an ecosystem of services with FathomNet tutorials<sup>60,61</sup>, code<sup>62</sup>, and trained models have been created to sustain and grow the user community. The FathomNet website<sup>34</sup> contains community features that allow for select members to review submissions and verify data contributions. A publicly accessible FathomNet Model Zoo (FMZ)<sup>63</sup> contains FathomNet-trained machine learning models contributed by community members to be freely used on other visual data collected in various marine regions with a number of platforms containing many concepts. As FathomNet grows to include additional concepts and imagery, we envision intellectual activities around the dataset similar to ImageNet and Kaggle-style competitions, where baseline datasets and annual challenges could be leveraged to develop state-of-the-art algorithms for future deployment<sup>64,65</sup>.

FathomNet has the potential to increase the rate of image and video data analysis by human observers, contribute to generation of labeled data, and create smarter algorithms deployed on robotic vehicles to enable targeted sampling and persistent observations of marine animals. Benthic and midwater machine learning models trained on FathomNet data derived solely from MBARI<sup>44,47</sup> were used to create bounding box proposals and activity-detected video segments (Figs. 6, 7) from visual data collected by other underwater platforms and institutions (e.g., NOAA's ROV *Deep Discoverer*). This resulted in a reduction of human annotator effort by 81% in one use case, and demonstrated the potential applicability of high-level taxonomic models on other institutional data. Additionally, demonstrations using FathomNet-trained algorithms integrated with underwater vehicle controllers<sup>14</sup> provide optimism for future AI-enabled missions to fully automate targeted sampling of marine objects, persistent observations of animals, and result in less-invasive sampling of valuable resources in the ocean.

For the foreseeable future, there is not going to be a single general-use classifier or detector for all ocean image data. There simply is not enough expert annotated data and too many undiscovered organisms to train such a system. Moreover, there is substantial evidence that contemporary models trained on benchmark datasets are not robust to natural distribution shifts at either the pixel or population level<sup>46,66</sup>, an observation consistent with our experiments on out-of-distribution target datasets. While the MBARI and NOAA datasets are similar, differences in the imaging systems, deployment mode, and environment were enough to challenge the model. While such a model may not be deployable out-of-the-box, it could substantially alleviate the human cost of working with a new dataset, or be used to assess performance and suggest avenues for strategic retraining. We hope that FMZ facilitates such machine learning experiments, and enables more efficient workflows for all manner of organizations and groups. We believe this approach to open science for ocean image data has tremendous potential to accelerate the field as has been demonstrated in other areas<sup>24,67</sup>.

While growth of the FathomNet database has inherent value for the research community, the potential for public engagement and its impact on education and conservation is equally important. The imagery contained within the database can be integrated into digital video data repositories and workflows using the Rest API, and could enable a global ocean life guide that can support research and education initiatives. Social media campaigns for community engagement, similar to iNaturalist, and eBird<sup>67,68</sup> could have similar, far-reaching outcomes for FathomNet, providing a mechanism for aggregating and leveraging taxonomic knowledge shared by the community. By making FathomNet publicly accessible, we can also invite ocean enthusiasts to solve challenges related to ocean visual data at temporal and spatial scales that are impossible without widespread engagement coupled with semi-automation. Incorporating human-AI interaction research with video game design and award structures could enable widespread participation similar to other gaming platforms<sup>69</sup>, enabling direct science contributions by game players<sup>70</sup>. Through FathomNet and its community of users, we can create an ecosystem around marine visual data that can realize a more inclusive, equitable, and diverse vision of ocean exploration and discovery.

## Methods

**FathomNet seed data sources and augmentation tools.** FathomNet has been built to accommodate data contributions from a wide range of sources. The database has been initially seeded with a subset of curated imagery and metadata from the Monterey Bay Aquarium Research Institute (MBARI), National Geographic Society (NGS), and the National Oceanic and Atmospheric Administration (NOAA). Together, these data repositories represent more than 30 years of underwater visual data collected by a variety of imaging technologies and platforms around the world. To be sure, the data currently contained within FathomNet does not include the entirety of these databases, and future efforts will involve further augmenting image data from these and other resources.

**MBARI's video annotation and reference system.** Beginning in 1988, MBARI has collected and curated underwater imagery and video footage through their Video Annotation and Reference System (VARS<sup>39</sup>). This video library contains detailed footage of the biological, geological, and physical environment of the Monterey Bay submarine canyon and other areas including the Pacific Northwest, Northern California, Hawaii, the Canadian Arctic, Taiwan, and the Gulf of California. Using eight different imaging systems (mostly color imagery and video, with more recent additions that include monochrome computer vision cameras<sup>14</sup>) deployed from four different remotely operated vehicles (ROVs *MiniROV*, *Ventana*, *Tiburón* and *Doc Ricketts*), VARS contains approximately 27,400 h of video from 6190 dives, and 536,000 frame grabs. These dives are split nearly evenly between

observations in benthic (from the seafloor to 50 m above the seafloor) and midwater (from the upper surface of the benthic environment to the lower surface of the lighted shallower waters or  $\sim 200$  m) habitats. Image resolution has improved over the years from standard definition (SD;  $640 \times 480$  pixels) to high-definition (HD;  $1920 \times 1080$  pixels), with 4 K resolutions ( $3840 \times 2160$  pixels) starting in 2021. Additional imaging systems managed within VARS, which include a low-light camera<sup>1</sup>, the I2MAP autonomous underwater vehicle imaging payload, and DeepPIV<sup>71</sup>, are currently excluded from data exported into FathomNet. In addition to imagery and video data, VARS synchronizes ancillary vehicle data (e.g., latitude, longitude, depth, temperature, oxygen concentration, salinity, transmittance, and vehicle altitude), and is included as image metadata for export to FathomNet.

Of the 27,400 hours of video footage, more than 88% has been annotated by video taxonomic experts in MBARI's Video Lab. Annotations within VARS are created and constrained using concepts that have been entered into the knowledge database (or knowledgebase; see Fig. S1) that is approved and maintained by a knowledge administrator using community taxonomic standards (i.e., WoRMS<sup>35</sup>) and input from expert taxonomists outside of MBARI. To date, there are more than 7.5 M annotations across 4300 concepts within the VARS database. By leveraging these annotations and existing frame grabs, VARS data were augmented with localizations (bounding boxes) using an array of publicly available<sup>72,73</sup> and in-house<sup>74–76</sup> localization and verification tools by either supervised, unsupervised, and/or manual workflows<sup>77</sup>. More than 170,000 localizations across 1185 concepts are contained in the VARS database and, due to MBARI's embargoed concepts and dives, FathomNet contains approximately 75% of this data at the time of publication.

**NGS's benthic lander platforms and tools.** The National Geographic Society's Exploration Technology Lab has been deploying versions of its autonomous benthic lander platform (the Deep Sea Camera System, DSCS) since 2010, collecting video data from locations in all ocean basins<sup>42</sup>. Between 2010 and 2020, the DSCS has been deployed 594 times, collecting 1039 h of video at depths ranging from 28 to 10,641 m in a variety of marine habitats (e.g., trench, abyssal plain, oceanic island, seamount, arctic, shelf, strait, coastal, and fjords). Videos from deployments have subsequently been ingested into CVision AI's cloud-based collaborative analysis platform Tator<sup>73</sup>, where they are annotated by subject-matter experts at University of Hawaii and OceansTurn. Annotations are made using a Darwin Core-compliant protocol with standardized taxonomic nomenclature according to WoRMS<sup>78</sup>, and adheres to the Ocean Biodiversity Information System (OBIS<sup>79</sup>) data standard formats for image-based marine biology<sup>42</sup>. At the time of publication, 49.4% of the video collected using the DSCS has been annotated. In addition to this analysis protocol, animals have also been localized using a mix of bounding box and point annotations. Due to these differences in annotation styles, 2,963 images and 3,256 annotations using bounding boxes from DSCS has been added to the FathomNet database.

**NOAA's Office of Exploration and Research video data.** The National Oceanic and Atmospheric Administration (NOAA) Office of Ocean Exploration and Research (OER) began collecting video data aboard the RV *Okeanos Explorer* (EX) in 2010, but only retained select clips due to the volume of the video data until 2016, when deck-to-deck recording began. As NOAA's first dedicated exploration vessel, all video data collected are archived and made publicly accessible from the NOAA National Centers for Environmental Information (NCEI)<sup>80</sup>. This specialized access is dependent upon standardized ISO 19115-2 metadata records that incorporate annotations. The dual remotely operated vehicle system, ROVs *Deep Discoverer* and *Seirios*<sup>45</sup> contains 15 cameras: 6 HD and 9 SD. Two camera streams, typically the main HD camera on each ROV, are recorded per cruise. The current video library includes over 271 TB of data collected over 519 dives since 2016, including 39 dives with midwater transects. The data were collected during 3938.5 h of ROV time, 2610 h of bottom time, and 44 h of midwater transects. These data cover broad spatial areas (from the Western Pacific to the Mid-Atlantic) and depth ranges (from 86 to 5999.8 m). Ancillary vehicle data (e.g. location, depth, pressure, temperature, salinity, sound velocity, oxygen, turbidity, oxidation reduction potential, altitude, heading, main camera angle, and main camera pan angle) are included as metadata.

NOAA-OER originally crowd-sourced annotations through volunteer participating scientists, and began supporting expert taxonomists in 2015 to more thoroughly annotate collected video. In 2015, NOAA-OER and partners began the Campaign to Address Pacific Monument Science, Technology, and Ocean Needs (CAPSTONE), which was a 3 year campaign to explore US marine protected areas in the Pacific. Expert annotations generated by the Hawaii Undersea Research Laboratory<sup>45</sup> for this single campaign generated more than 90,000 individual annotations covering 187 dives (or 36% of the EX video collection) using VARS<sup>39</sup>. At the University of Dallas, Dr. Deanna Soper's undergraduate student group localized these expertly generated annotations for two cruises consisting of 37 dives (or 7% of the EX collection) from CAPSTONE, producing 8165 annotations and 2866 images using the Tator Annotation tool<sup>73</sup>. These data have formed the initial contribution of NOAA's data to FathomNet.

**Computation of FathomNet database statistics.** Drawing several metrics from the popular ImageNet and COCO image databases<sup>22,23</sup>, and additional comparisons with iNat2017<sup>25</sup>, we can generate summary statistics and characterize the FathomNet dataset. These measures serve to benchmark FathomNet against these resources, underscore how it is different, and reveal unique challenges related to working with underwater image data.

**Aggregate statistics.** Aggregate FathomNet statistics were computed from the entire database accessed via the Rest API in October 2021 (Figs. 4, 5). To visualize the amount of contextual information present in an image, we estimated the number of concepts and instances as a function of the percent of the full frame they occupy (Fig. 4a, b), with FathomNet data split taxonomically (denoted by  $x$ ) to visualize how data breaks down into

biologically relevant groupings. The taxonomic labels at each level of a given organism's phylogeny were back-propagated from the human annotator's label based on designations in the knowledgebase (Fig. S1). If an object was not annotated down to the relevant level of the taxonomic tree (e.g., species), the next closest rank name up the tree was used (e.g., genus). The average number of instances and concepts are likewise split at taxonomic rank (Fig. 4c). The percent of instances of a particular concept and how they are distributed across all images is shown in Fig. 4d.

**Concept coverage.** Coverage—an indication of the completeness of an image's annotations—is an important consideration for FathomNet. Coverage is quantified as average recall, and is demonstrated over 50 randomly selected images at each level of the taxonomic tree (between order and species; Fig. S1) for a benthic and mid-water organism, *Gersemia juliepackardae* and *Bathochordaeus mcnutti*, respectively (Fig. 5a). This is akin to examining the precision of annotations as a function of synset depth in ImageNet<sup>22</sup>. FathomNet images with expert-generated annotations at each level of the tree, including all descendent concepts, were randomly sampled and presented to a domain expert. They then evaluated the existing annotations and added missing ones until every biological object in the image was localized. The recall was then computed for the target concept and all other objects in the frame. The false detection rate of existing annotations was negligible, and was much less than 0.1% for each concept.

**Pose variability: iconic versus non-iconic imagery.** The data in FathomNet represents the natural variability in pose of marine animals, which includes both iconic and non-iconic views of the concept. A subject's position relative to the camera, relationship with other objects in the frame, the amount it is occluded, and the imaging background are all liable to change between frames. By computing the average image across each concept, an image class with high variability in pose (or non-iconic) will result in a blurrier, more uniformly gray image than a group of images with little pose diversity (or iconic)<sup>22</sup>. We computed the average image from an equivalent number of randomly sampled images across two FathomNet concepts (medusae and echinoidiae) and the closest associated synsets in ImageNet (jellyfish and starfish), which is shown in Fig. 5b.

**FathomNet data usage and ecosystem.** The FathomNet data use policy balances the need for distributed metadata sharing while simultaneously providing protection for data contributors wanting to maintain copyright of their valuable underwater image assets. All submitted annotation data are licensed under a Creative Commons Attribution—No Derivatives 4.0 International License. For image data that are submitted to FathomNet via a list of URLs, the original owner of those images maintains their copyright. The use of images that are directly or indirectly hosted by FathomNet are then licensed under a Creative Commons Attribution—Non Commercial—No Derivatives 4.0 International license, and all of the images may be used for training and development of machine learning algorithms for commercial, academic, non-profit, and government purposes. For all other uses of the images, users are to contact the original copyright holder, which can be tracked within the FathomNet database.

To grow the FathomNet community, we have created other resources that enable contributions from data scientists to marine scientists and ocean enthusiasts. Along with the FathomNet database, machine learning models that are trained on the image data can be posted, shared, and subsequently downloaded from the FathomNet Model Zoo (FMZ;<sup>63</sup>). Community members can not only contribute labeled image data, but also provide subject-matter expertise to validate submissions and augment existing labels via the web portal<sup>34</sup>. This is especially helpful when images do not have full coverage annotations. Finally, additional resources include code<sup>62</sup>, blogs<sup>60</sup>, and YouTube channel<sup>61</sup>, that contain helpful information about engaging with FathomNet.

**Estimating and contextualizing FathomNet's value.** The two most commonly used image databases in the computer vision community, ImageNet and COCO, are built from images scraped from publicly available internet repositories. Both ImageNet and COCO were built with crowd-sourced annotation via Amazon's Mechanical Turk (AMT) service, where workers are paid per label or image. The managers of these data repositories have not published the collection and annotation costs of their respective databases, however we can estimate these costs by comparing the published number of worker hours with compensation suggestions from AMT optimization studies. The recommended dollar values a study generating computer vision training data<sup>81</sup> and scientific annotations<sup>82</sup> are in keeping with several meta-analyses of AMT pay scales, suggesting that 90% of HIT rewards are less than \$0.10 a task and that average hourly wages are between \$3 and \$3.50 per hour<sup>83,84</sup>.

The original COCO release contains several different types of annotations: category labels for an entire image, instance spotting for individual objects, and pixel level instance segmentation. Each of these tasks entails different amounts of attention from annotators. Lin et al.<sup>23</sup> estimated that the initial release of COCO required over 70,000 Turker hours. If the reward was set to \$0.06 per task, category labels cost \$98,000, instance spotting was \$46,500, and segmentation cost \$150,000 for a total of about \$295,000. ImageNet currently contains 14.2M annotated images, each one observed by an average of 5 independent Turkers. At the same category label per hour rate as COCO, the dataset required ~76,850 Turker hours. Assuming a HIT reward of \$0.06, ImageNet cost \$852,000. These estimates do not include the cost for image generation, intellectual labor on the part of the managers, hosting fees, or compute costs for web scraping.

Fine-grained, taxonomically correct annotation is difficult to crowd-source on AMT<sup>57</sup>. The initial release of FathomNet annotations thus rely on domain expert annotations from the institutions generating the images. The annotation cost for MBARI's Video Lab for one technician is \$80 per hour. Expert annotators require approximately 6 months of training before achieving expert status in a new habitat, and the annotator will continue to learn taxonomies and animal morphology on the job. The bounding boxes for FathomNet require different

amounts of time in different marine environments; midwater images typically have fewer targets, and benthic images can be very dense. Based on the Video Lab's initial annotation efforts, an experienced annotator can label ~ 80 midwater images per hour for a \$1 per image cost. The same domain experts were able to label ~ 20 benthic images per hour or about \$3 per image. The 66,039 images in the initial upload to FathomNet from MBARI are approximately evenly split between the two habitats, costing ~ \$165,100 to generate the annotations. At this hourly rate, ImageNet would cost ~ \$6.15 M to annotate. We believe these costs are in-line with other annotated ocean image datasets. True domain expertise is expensive and reflects the value of an individual's training and contribution to a project. In addition to the intellectual costs of generating FathomNet, ocean data collection often requires extensive instrument development and many days of expensive ship time. To date, FathomNet largely draws from MBARI's VARS database, which is comprised of 6190 ROV dives and represents ~ \$143.7 M worth of ship time. Including these additional costs underscores the value of FathomNet, especially to groups in the ocean community that are early in their data collection process.

## Data availability

**Accession codes** All code and data used for this manuscript can be found on the FathomNet Code Repository<sup>62</sup> at [www.github.com/fathomnet](https://www.github.com/fathomnet) and the FathomNet database<sup>34</sup> at [www.fathomnet.org](https://www.fathomnet.org). The referenced machine learning models for the benthic and midwater use cases can be found either listed in the FathomNet Model Zoo<sup>63</sup> at [www.github.com/fathomnet/models](https://www.github.com/fathomnet/models), or at<sup>44</sup> and<sup>47</sup>, respectively. Supplementary information is available for this paper, which includes a figure, table, and video.

Received: 10 May 2022; Accepted: 6 September 2022

Published online: 23 September 2022

## References

- Haddock, S. H. D. *et al.* Insights into the biodiversity, behavior, and bioluminescence of deep-sea organisms using molecular and maritime technology. *Oceanography* **30**, 38–47. <https://doi.org/10.5670/OCEANOGRAPHY.2017.422> (2017).
- Appeltans, W. *et al.* The magnitude of global marine species diversity. *Curr. Biol.* **22**, 2189–2202 (2012).
- Lehman, J. From ships to robots: The social relations of sensing the world ocean. *Soc. Stud. Sci.* **48**, 57–79 (2018).
- Zhang, F., Marani, G., Smith, R. N. & Choi, H. T. Future trends in marine robotics [tc spotlight]. *IEEE Robot. Autom. Mag.* **22**, 14–122 (2015).
- Wang, Z. A. *et al.* Advancing observation of ocean biogeochemistry, biology, and ecosystems with cost-effective in situ sensing technologies. *Front. Mar. Sci.* **6**, 519 (2019).
- Claustre, H., Johnson, K. S. & Takeshita, Y. Observing the global ocean with biogeochemical-argo. *Ann. Rev. Mar. Sci.* **12**, 23–48 (2020).
- McKinna, L. I. Three decades of ocean-color remote-sensing trichodesmium spp. in the world's oceans: a review. *Progress Oceanogr.* **131**, 177–199 (2015).
- Benoit-Bird, K. J. & Lawson, G. L. Ecological insights from pelagic habitats acquired using active acoustic techniques. *Ann. Rev. Mar. Sci.* **8**, 463–490 (2016).
- Chavez, F. P. *et al.* Observing life in the sea using environmental DNA. *Oceanography* **34**, 102–119 (2021).
- Durden, J. M. *et al.* Perspectives in visual imaging for marine biology and ecology: From acquisition to understanding. *Oceanogr. Mar. Biol. Annu. Rev.* **54**, 1–72 (2016).
- Irison, J.-O., Ayata, S.-D., Lindsay, D. J., Karp-Boss, L. & Stemmann, L. Machine learning for the study of plankton and marine snow from images. *Annu. Rev. Mar. Sci.* **14**, 5 (2021).
- Zurowietz, M. & Nattkemper, T. W. Current trends and future directions of large scale image and video annotation: Observations from four years of biigle 2.0. *Front. Mar. Sci.* **2021**, 5 (2021).
- Greer, A. T. *et al.* High-resolution sampling of a broad marine life size spectrum reveals differing size- and composition-based associations with physical oceanographic structure. *Front. Mar. Sci.* **7**, 542701 (2020).
- Katija, K. *et al.* Visual tracking of deepwater animals using machine learning-controlled robotic underwater vehicles. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision* 860–869 (2021).
- Bell, A. T., Fundis, K. L. C., Fundis. Inspiring, engaging, and educating the next generation of STEM learners. *J. Ocean Technol.* **9**, 73–78 (2014).
- Fauville, G., Queiroz, A. C. M. & Bailenson, J. N. Virtual reality as a promising tool to promote climate change awareness. *Technol. Health* **2020**, 91–108 (2020).
- Gomes-Pereira, J. N. *et al.* Current and future trends in marine image annotation software. *Prog. Oceanogr.* **149**, 106–120 (2016).
- LeCun, Y., Bengio, Y. & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
- Schmarje, L., Santarossa, M., Schröder, S.-M. & Koch, R. A survey on semi-, self- and unsupervised learning for image classification. *IEEE Access* **9**, 82146–82168 (2021).
- Luo, J., Han, Y. & Fan, L. Underwater acoustic target tracking: A review. *Sens. (Switzerland)* **18**, 1–38. <https://doi.org/10.3390/s18010112> (2018).
- Branson, S., Van Horn, G., Perona, P. & Belongie, S. Improved bird species recognition using pose normalized deep convolutional nets. In *Proceedings of the British Machine Vision Conference* <https://doi.org/10.5244/C.28.87> (2014).
- Deng, J. *et al.* ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255. <https://doi.org/10.1109/CVPR.2009.5206848> (2009).
- Lin, T.-Y. *et al.* Microsoft coco: Common objects in context. In *European Conference on Computer Vision* 740–755 (Springer, 2014).
- Tsipras, D., Santurkar, S., Engstrom, L., Ilyas, A. & Madry, A. From ImageNet to image classification: Contextualizing progress on benchmarks. In *International Conference on Machine Learning* 9625–9635 (PMLR, 2020).
- Van Horn, G. *et al.* The inaturalist species classification and detection dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 8769–8778 (2018).
- Lombard, F. *et al.* Globally consistent quantitative observations of planktonic ecosystems. *Front. Mar. Sci.* **6**, 196 (2019).
- Sosik, H. M., Peacock, E. E. & Brownlee, E. F. In *WHOI-Plankton: Annotated Plankton Images—Data Set for Developing and Evaluating Classification Methods*. <https://doi.org/10.1575/1912/7341> (2014).
- Chen, Q., Beijbom, O., Chan, S., Bouwmeester, J. & Kriegman, D. A new deep learning engine for coralnet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* 3693–3702 (2021).
- Zhuang, P., Wang, Y. & Qiao, Y. WildFish: A large benchmark for fish recognition in the wild. In *Proceedings of the 26th ACM International Conference on Multimedia* 1301–1309 (2018).
- Wilkinson, M. D. *et al.* The fair guiding principles for scientific data management and stewardship. *Sci. Data* **3**, 1–9 (2016).



31. Wieczorek, J. *et al.* Darwin core: An evolving community-developed biodiversity data standard. *PLoS One* **7**, e29715 (2012).
32. Costello, M. J. *et al.* Global coordination and standardisation in marine biodiversity through the world register of marine species (worms) and related databases. *PLoS One* **8**, e51629 (2013).
33. Marine regions: Towards a standard for georeferenced marine names. <https://marineregions.org/> (2022).
34. Fathomnet website. <http://fathomnet.org/fathomnet/> (2021).
35. Horton, T. *et al.* World register of marine species (worms), accessed 30 Jan 2021. <http://www.marinespecies.org> (2021).
36. Lowry, M., Folpp, H., Gregson, M. & Suthers, I. Comparison of baited remote underwater video (bruv) and underwater visual census (uvc) for assessment of artificial reefs in estuaries. *J. Exp. Mar. Biol. Ecol.* **416**, 243–253 (2012).
37. Picheral, M. *et al.* The underwater vision profiler 5: An advanced instrument for high spatial resolution studies of particle size spectra and zooplankton. *Limnol. Oceanogr. Methods* **8**, 462–473 (2010).
38. Horton, T. *et al.* Recommendations for the standardisation of open taxonomic nomenclature for image-based identifications. *Front. Mar. Sci.* **8**, 62 (2021).
39. Schlining, B. & Stout, N. J. Mbari's video annotation and reference system. In *OCEANS 2006 1–5* (IEEE, 2006).
40. Althaus, F. *et al.* A standardised vocabulary for identifying benthic biota and substrata from underwater imagery: The catami classification scheme. *PLoS One* **10**, e0141039 (2015).
41. Howell, K. L. *et al.* A framework for the development of a global standardised marine taxon reference image database (smarter-id) to support image-based analyses. *PLoS One* **14**, e0218904 (2019).
42. Giddens, J., Turchik, A., Goodell, W., Rodriguez, M. & Delaney, D. The national geographic society deep-sea camera system: A low-cost remote video survey instrument to advance biodiversity observation in the deep ocean. *Front. Mar. Sci.* **7**, 601411 (2020).
43. Hammond, S., McDonough, J. & Russell, C. The noaa ship okeanos explorer: New ways for exploring the ocean. *Oceanography* **23**, 88–89 (2010).
44. Woodward, B., Lundsten, L. & Orenstein, E. *Mbari Benthic Supercategory Object Detector*. <https://doi.org/10.5281/zenodo.5571043> (2021).
45. Kennedy, B. R. *et al.* The unknown and the unexplored: Insights into the pacific deep-sea following noaa capstone expeditions. *Front. Mar. Sci.* **6**, 480 (2019).
46. Taori, R. *et al.* Measuring Robustness to Natural Distribution Shifts in Image Classification. *Adv. Neural Inf. Process. Syst.* **33**, 18583–18599 (2020).
47. Woodward, B. G. *et al.* *Mbari Midwater Object Detector*. <https://doi.org/10.5281/zenodo.5942597> (2022).
48. He, K., Zhang, X., Ren, S. & Sun, J. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* 770–778 (2016).
49. Nayak, R., Pati, U. C. & Das, S. K. A comprehensive review on deep learning-based methods for video anomaly detection. *Image Vis. Comput.* **106**, 104078. <https://doi.org/10.1016/j.imavis.2020.104078> (2021).
50. Lindsay, D. *et al.* The untethered remotely operated vehicle PICASSO-1 and its deployment from chartered dive vessels for deep sea surveys off Okinawa, Japan, and Osprey Reef, Coral Sea, Australia. *Mar. Technol. Soc. J.* **46**, 20–32 (2012).
51. Yoerger, D. R. *et al.* Mesobot: An autonomous underwater vehicle for tracking and sampling midwater targets. *IEEE AUV* <https://doi.org/10.1109/AUV.2018.8729822> (2018).
52. Koreitem, K. *et al.* Synthetically trained 3D visual tracker of underwater vehicles. In *OCEANS 2018 MTS/IEEE Charleston*. <https://doi.org/10.1109/OCEANS.2018.8604597> (2019).
53. Manderson, T., Higuera, J. C. G., Cheng, R. & Dudek, G. Vision-based autonomous underwater swimming in dense coral for combined collision avoidance and target selection. In *IEEE International Conference on Intelligent Robots and Systems* 1885–1891. <https://doi.org/10.1109/IROS.2018.8594410> (2018).
54. Andriluka, M., Roth, S. & Schiele, B. People-tracking-by-detection and people-detection-by-tracking. In *IEEE Conference on Computer Vision and Pattern Recognition* 1–8 (2008).
55. Yoerger, D. R. *et al.* A hybrid underwater robot for multidisciplinary investigation of the ocean twilight zone. *Sci. Robot.* **2021**, 5 (2021).
56. Lin, T.-Y., Goyal, P., Girshick, R., He, K. & Dollár, P. Focal loss for dense object detection. In *Proceedings of the IEEE International Conference on Computer Vision* 2980–2988 (2017).
57. Van Horn, G. *et al.* Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 595–604 (2015).
58. MacLeod, N., Benfield, M. & Culverhouse, P. Time to automate identification. *Nature* **467**, 154–155 (2010).
59. Orenstein, E. C. *et al.* Semi- and fully supervised quantification techniques to improve population estimates from machine classifiers. *Limnol. Oceanogr. Methods* **18**, 739–753 (2020).
60. Fathomnet medium blog. <https://medium.com/fathomnet> (2021).
61. Fathomnet youtube channel. [https://www.youtube.com/channel/UCTz\\_IVO8Q\\_FSjC5yE6sXAGg](https://www.youtube.com/channel/UCTz_IVO8Q_FSjC5yE6sXAGg) (2021).
62. Fathomnet code repository. <https://github.com/fathomnet> (2021).
63. Fathomnet model zoo. <https://github.com/fathomnet/models> (2021).
64. Russakovsky, O. *et al.* Imagenet large scale visual recognition challenge. *Int. J. Comput. Vis.* **115**, 211–252 (2015).
65. Yang, X. *et al.* Deep learning for practical image recognition: Case study on kaggle competitions. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining* 923–931 (2018).
66. Koh, P. W. *et al.* Wilds: A benchmark of in-the-wild distribution shifts. In *Proceedings of the 38th International Conference on Machine Learning*, vol. 139 of *Proceedings of Machine Learning Research* (eds Meila, M. & Zhang, T.) 5637–5664 (PMLR, 2021).
67. Unger, S., Rollins, M., Tietz, A. & Dumais, H. iNaturalist as an engaging tool for identifying organisms in outdoor activities. *J. Biol. Educ.* **2020**, 1–11 (2020).
68. Sullivan, B. L. *et al.* The ebird enterprise: An integrated approach to development and application of citizen science. *Biol. Cons.* **169**, 31–40 (2014).
69. Qian, M. & Clark, K. R. Game-based learning and 21st century skills: A review of recent research. *Comput. Hum. Behav.* **63**, 50–58 (2016).
70. Waldispuhl, J., Szantner, A., Knight, R., Caisse, S. & Pitchford, R. Leveling up citizen science. *Nat. Biotechnol.* **38**, 1124–1126 (2020).
71. Katija, K., Sherlock, R. E., Sherman, A. D. & Robison, B. H. New technology reveals the role of giant larvaceans in oceanic carbon cycling. *Sci. Adv.* **3**, e1602374 (2017).
72. Kawamura, R. Rectlabel. <https://rectlabel.com/> (2020).
73. CVision AI, Inc. Tator. <https://github.com/cvisionai/tator> (2019).
74. Barnard, K. VARS-Localize. <https://github.com/mbari-org/vars-localize> (2020).
75. Roberts, P. L. D. GridView. <https://bitbucket.org/mbari/gridview/> (2020).
76. Barnard, K. & Roberts, P. L. VARS-GridView. <https://github.com/mbari-org/vars-gridview> (2021).
77. Boulais, O. *et al.* Fathomnet: An underwater image training database for ocean exploration and discovery. [arXiv:2007.00114](https://arxiv.org/abs/2007.00114) (2020).
78. Boxshall, G. *et al.* World register of marine species (worms). In *WoRMS Editorial Board* (2014).
79. De Pooter, D. *et al.* Toward a new data standard for combined marine biological and environmental datasets-expanding obis beyond species occurrences. *Biodivers. Data J.* **2017**, 5 (2017).
80. Casey, K. S. Big data partnerships at noaa's national centers for environmental information. In *AGU Fall Meeting Abstracts*, vol. 2015, IN12A–02 (2015).

81. Sorokin, A. & Forsyth, D. Utility data annotation with amazon mechanical turk. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops* 1–8 (IEEE, 2008).
82. Hughes, A. J. *et al.* Quanti us: A tool for rapid, flexible, crowd-based annotation of images. *Nature Methods* **15**, 587–590 (2018).
83. Ipeirotis, P. G. Analyzing the amazon mechanical turk marketplace, XRDS: Crossroads. *ACM Mag. Students* **17**, 16–21 (2010).
84. Hara, K. *et al.* A data-driven analysis of workers' earnings on amazon mechanical turk. In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems* 1–14 (2018).

## Acknowledgements

Seed funding for FathomNet was provided by National Geographic Society (518018 to KK), National Oceanic and Atmospheric Administration (NA18OAR4170105 to KCB), and the Monterey Bay Aquarium Research Institute through generous support from the David and Lucile Packard Foundation (to KK). Additional funding support has been provided by National Geographic Society (NGS-86951T-21 to KCB), the National Science Foundation (OTIC 1812535 and Convergence Accelerator 2137977; to KK), and the Monterey Bay Aquarium Research Institute (to KK). Additional individuals whose contributions enriched FathomNet include members of MBARI's Video Lab (Nancy Jacobsen Stout, Kyra Schlining, Susan von Thun, Kristine Walz, Larissa Lemon), Bioinspiration Lab (Joost Daniels, Paul Roberts, Krish Mehta), and Alexandra Lapidés. National Geographic Society contributions were facilitated by Denley Delaney and Alan Turchik.

## Author contributions

K.K., K.C.B., and B.W. conceived FathomNet. E.O. generated database statistics with early contributions from O.B.; E.O., B.W., and K.K. worked on representative use cases. B.S. designed and built the database, VARS-to-FathomNet data pipelines, API, and website back-end; K.B. developed the Python API and worked on data ingestion from VARS to FathomNet. E.B. contributed to the refinement of the FathomNet website front-end. M.C. facilitated data contributions and ran dataset statistics on NOAA-OER's contributions. B.W. developed Tator-to-FathomNet data pipelines, which included NOAA and NGS data. L.L. and G.S. generated most of the labeled data from VARS that are contained in FathomNet, and L.L. conducted labeling experiments for the included use cases. K.K. wrote the manuscript with significant contributions from E.O. All authors reviewed the manuscript.

## Competing interests

The authors declare no competing interests.

## Additional information

**Supplementary Information** The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-19939-2>.

**Correspondence** and requests for materials should be addressed to K.K.

**Reprints and permissions information** is available at [www.nature.com/reprints](http://www.nature.com/reprints).

**Publisher's note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022