

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DATA³: DATASET SUBSET SELECTION FOR SPECIALIZATION

Anonymous authors
Paper under double-blind review




Figure 1: Foundation model training aims for broad generalization, by using all data available, usually from massive internet-scale datasets. In practice, we find these models are often suboptimal for specific deployments, which may exhibit different distributions over categories or data characteristics from the general training data pool. Dataset subset selection for specialization seeks to identify model training subsets closely aligned with the target deployment, achieving superior performance under the given distribution and attribute shifts.

ABSTRACT

In many real-world machine learning applications (e.g. detecting broken bones in x-rays or species in camera traps), models need to perform well on specific deployments (e.g. a specific hospital or national park) rather than the domain broadly. However, deployments often have imbalanced, unique data distributions. Discrepancies between training and deployment distributions lead to suboptimal performance, highlighting the need to curate training data for *specialized models for specific deployment needs*. We formalize **dataset subset selection for specialization (DS3)**: given a training set drawn from a general distribution and a (potentially unlabeled) query set drawn from a deployment-specific distribution, the goal is to select a subset of the training data that optimizes deployment performance.

We introduce **DATA³**, the first dataset and benchmark designed specifically for the DS3 problem. DATA³ encompasses five *real-world* application domains, each with a set of distinct deployments to specialize in. We conduct a comprehensive study evaluating different state-of-the-art data curation algorithms and find that methods trained on general distributions consistently fail to perform optimally on deployment tasks. Additionally, we demonstrate the existence of expert-curated (deployment-specific) subsets that outperform training on all available data by up to 51.3%. Our benchmark highlights the critical role of tailored dataset curation in enhancing performance and training efficiency on deployment-specific distributions, which we posit will only become more important as global, public datasets become available across domains and ML models are deployed in the real world.

1 BACKGROUND AND MOTIVATION

Machine learning models are typically trained on large datasets with the assumption that the training distribution closely matches the distribution of the deployment where the model will be applied. However, in real-world applications, deployment data distributions often diverge from general and/or

054 global training set distributions (Shen et al., 2024; Taori et al., 2020). Selecting relevant data subsets
 055 aligned with specific deployments is crucial to maximize field performance. The problem of *data*
 056 *subset selection for specialization* (DS3) is thus critical: given all available training data for a domain
 057 and a small (usually unlabeled) query set that represents the desired deployment, the goal is to identify
 058 a subset of the training data, such that training the ML model on this subset maximizes performance
 059 on the deployment distribution.

060 **Real world example.** Consider a wildlife ecologist who aims to build a classifier to detect the presence
 061 of invasive species in camera trap images collected at the Channel Islands. Existing labeled training
 062 data in this context is limited, thus training a classifier from scratch is likely to be unsuccessful.
 063 A common approach is to finetune a general pre-trained model (such as ViT or CLIP) on all *relevant*
 064 camera trap data. But *what does "relevant data" mean?* Would using similar species data from
 065 other camera trap locations (perhaps on the mainland) improve performance, or introduce noise?
 066 What about including data from non-similar species at that location? While adding data to a training
 067 set can sometimes improve performance, it can also decrease individual subgroup performance in a
 068 biased way (Compton et al., 2023) and introduce spurious correlations that can enable models
 069 to learn potentially dangerous "shortcuts," resulting in biased predictions, shown across various
 070 domains (Geirhos et al., 2020; Badgeley et al., 2018; Wang et al., 2021; Beery et al., 2022a).

071 **Our contributions.** Our key contributions are the following:

- 072 (i) We are the first to identify and formalize the challenge of sub-selecting training data to
 073 specialize models to new deployments (dataset subset selection for specialization).
- 074 (ii) We propose **DATAS³**: A novel benchmark that enables the AI community to investigate and
 075 make progress on DS3. **DATAS³** reformulates, adapts, and adds to five diverse datasets, each
 076 from a different application domain. We worked directly with domain experts throughout
 077 the curation and reformulation process to ensure that **DATAS³** accurately reflects (1) real-
 078 world dataset distribution challenges that require model specialization (i.e., covariate shifts,
 079 subpopulation shifts, and long-tailed distributions), and (2) evaluation settings (test splits)
 080 representative of real-world deployment scenarios in each domain.
- 081 (iii) We show that a well-curated subset can consistently outperform models trained on the entire
 082 dataset for each deployment.
- 083 (iv) We also conduct an extensive experimental study comparing current SOTA subset selection
 084 methods on **DATAS³**. After training a suite of baselines, our results clearly show that current
 085 subset selection methods fail on DS3, highlighting the need future research to solve the
 086 DS3 problem on **DATAS³**.
- 087 (v) We release a codebase, python package, and public leaderboard for submission to the
 088 benchmark, available at [datas3-benchmark.github.io](https://github.com/datas3-benchmark/datas3-benchmark)

091 2 PROBLEM STATEMENT

093 **DS3 problem formulation.** Let X be a pool of data points, $T \subset X$ be a given *training set*
 094 drawn from a training (pool) distribution P_T over X , and let $Q \subset X$ be a *query set* drawn from the
 095 desired **deployment-specific distribution** P_Q over X . Given a model θ , the objective of **dataset**
 096 **subset selection for specialization (DS3)**, is to design an algorithm **SubsetSelection-ALG**,
 097 which takes T (the training set) and Q (the deployment representative query set) as input, and
 098 outputs a subset $S^* \subset T$ that minimizes the expected loss of θ trained on S^* over the desired
 099 deployment-specific distribution P_Q . More formally:

$$100 \quad 101 \quad 102 \quad S^* = \arg \min_{S \subset T} \mathbb{E}_{q \sim P_Q} [\mathcal{L}(\theta(S), q)], \quad (1)$$

103 where $\theta(S)$ denotes the model trained on the subset $S \subset T$, and $\mathcal{L}(\theta(S), q)$ is the loss function
 104 evaluated on a single point q sampled from P_Q and the trained model $\theta(S)$. The term $\mathbb{E}_{q \sim P_Q}$ denotes
 105 the expected value over the distribution P_Q . Hence, the algorithm **SubsetSelection-ALG**
 106 outputs S^* , the subset of T that minimizes the expected loss of the entire desired deployment
 107 distribution P_Q . Notably, **SubsetSelection-ALG** can only access the desired deployment-specific
 108 distribution via the query set Q . Unlike complementary lines of work such as active domain

adaptation (ADA), which assumes real-time compute and focuses on actively/iteratively selecting and then collecting labels for data within the deployment during the specialization process, DS3 selects data in a single-shot approach prior to specialization on an already available pool of data (potentially for use in resource-constrained applications).

Is the query set labeled? This formalization can be divided into two cases. In the first, the query set Q is annotated with a set of labels: Q is a set of $m > 0$ pairs $Q = \{(q_1, y_1), \dots, (q_m, y_m)\}$, where for every $i \in [m]$, q_i is the i th feature vector describing the i th input, and y_i is its corresponding label/annotation. In this case the algorithm `SubsetSelection-ALG` has access to the set of labels $\{y_1, \dots, y_n\}$. In the second scenario, no labels are provided for Q , meaning that the `SubsetSelection-ALG` does not have access to the set $\{y_1, \dots, y_n\}$ and consequently $Q = \{q_1, \dots, q_m\}$. In a real-world example, Q can be thought of as the data collected from a deployment thus far, enabling additional selection from a larger database (the training pool). Annotating Q for any specific deployment is quite expensive, requiring time, money, and expertise, so progress on methods without query labels would be helpful for real-world applications.

Is `SubsetSelection-ALG` model agnostic? Similarly, this formalization can be approached in two different ways: one where the computation of S^* depends on a given specific model θ , i.e., `SubsetSelection-ALG` is model dependent, and has access to the model θ we wish to train on. Ideally, a well-performing, robust method should work well for multiple models, and will be more generalizable than a model-dependent algorithm. We test several different models on our benchmark for various `SubsetSelection-ALG` baselines to test this.


Should `SubsetSelection-ALG` be sample efficient? The goal of our benchmark is to specialize on a desired deployment distribution. Unlike standard subset selection, where subset size is often a primary concern, our focus is on selecting subsets based on relevance based on a particular deployment that yield highest performance evaluated on that deployment. Smaller subsets offer many advantages, such as training efficiency, lower memory/storage, etc; we analyze these tradeoffs in Appendix C.

3 RELATED WORK

It has become increasingly clear that data work is equally important to architecture design for increased model performance (Compton et al., 2023). Data curation for better quality training pools has been identified as an important line of research within this field. Many methods have been proposed for data curation and subset selection – we provide a comprehensive overview of these methods in Appendix B. Current benchmarks for data curation include Gadre et al. (2024), Mazumder et al. (2023) and Feuer et al. (2024). However, these benchmarks focus on data curation for a single higher quality training pool meant for better performance across many different downstream tasks, in contrast to specialization for a particular deployment. Additionally, data selection methods (Killamsetty et al., 2021b; Tukan et al., 2023) are often evaluated on standard CIFAR10/100 (Krizhevsky et al., 2009) or ImageNet (Deng et al., 2009b) datasets, where test and validation sets have similar distribution to their training sets. No existing benchmarks focus on the DS3 challenge. `DATAS`³ is the first benchmark specifically designed to evaluate subset selection methods for *deployment-specific specialization*, rather than generalization, where the training and testing data exhibit distributional shifts representative of real-world deployment challenges (Figure 1).

4 THE DATAS³ BENCHMARK

Datasets. We describe each `DATAS`³ dataset. Our benchmark includes five datasets, each capturing a unique and diverse application of ML: AutoArborist for tree classification (Beery et al., 2022b), iWildCam for camera trap species identification (Beery et al., 2021), GeoDE for diverse object classification (Ramaswamy et al., 2023), NuScenes for autonomous driving footage steering regression (Caesar et al., 2020), and FishDetection for underwater video fish detection (Dawkins et al., 2017). To make the `DATAS`³ datasets usable, we have made considerable changes to them to better highlight deployment challenges, augment with additional data, or preprocess the data for use with standard ML pipelines. For each dataset, we provide a proof-of-concept "oracle" / knowledge-driven subset that demonstrates the usefulness of subset selection, with improvements over using training on all data. These subsets were created using information that benchmark users are not provided (e.g. metadata, GPS location, region, etc). Additional details about each dataset and can be found in Apdx. E.



175 4.1 iWILDCAM 176

177 **Motivation:** Animal populations have declined by 68% on average since 1970 (Staub, 2020). To mon-
178 itor this biodiversity loss, ecologists deploy camera traps—motion-activated cameras placed in the
179 wild (Wearn & Glover-Kapfer, 2017)—and process the data with machine learning models (Norouz-
180 zadeh et al., 2019; Beery et al., 2019). However, variations in illumination, camera angle, background,
181 vegetation, color, and animal frequencies across different locations cause these models to generalize
182 poorly to new deployments. To specialize models for specific locations, selecting appropriate data
183 subsets for deployment-specific (in this case location) specialization becomes essential.

184 **Problem Setting & Data:** To study this problem, we use the iWildCam 2020 dataset, comprising
185 of 203,029 images from 323 different camera traps spread across multiple countries in different
186 parts of the world. The task is multi-class species classification from 182 different animal species.
187 Performance is measured by overall classification accuracy for species identification. The original
188 camera trap data comes from the Wildlife Conservation Society (link).

189 **Deployments:** Our deployments were defined to be split across camera trap locations to simulate
190 the common scenario of researchers setting up new cameras within a region, with poor model
191 generalization on the new cameras (Wearn & Glover-Kapfer, 2017). Our train/test split was done
192 randomly across the 200 locations, with the five downstream test tasks created by clustering by the
193 latitude and longitude of camera GPS location in 4 deployments: (1) Central America, (2) Eastern
194 Africa, (3) Southern Africa, and (4) Southeast Asia. Similar to most other camera trap datasets,
195 iWildCam has significant long-tailed label distributions, with variation in species and backgrounds
196 between locations, as can be seen in Figure 3.


197 **Knowledge-driven Subset:** These subsets were created by only choosing training data from camera
198 locations that are within 100km of the camera locations in the deployments (the relevant geographical
199 area) and eliminating irrelevant classes that are not present in the deployment.

200 201 4.2 GODE 202


203 **Motivation:** Object classification datasets are often constructed by scraping images from the web
204 but contain geographical biases (Shankar et al., 2017). Instead of scraping images from the web,
205 GeoDE (Ramaswamy et al., 2023) crowdsources a dataset that is roughly balanced across 40 different
206 objects and six world regions, showing that common objects (stoves, bicycles, etc), vary in appearance
207 across the world. Accordingly, specializing models to different regions becomes useful when the
208 objects have strong covariate shift.

209 **Problem setting & Data:** GeoDE is a diverse dataset of 61,490 images comprising 40 different
210 objects collected from 6 world regions (Africa, Americas, East Asia, Europe, Southeast Asia, West
211 Asia). The associated task is multiclass classification, where the goal is to predict the object depicted
212 in each image.



213 **Deployments:** We propose 4 different deployments: (1) objects in Indonesia, (2) objects in Nigeria,
214 (3) indoor objects, and (4) outdoor objects, as shown in Figure 3. Nigeria and Indonesia were selected
215 as the two countries with the poorest performance, and the indoor/outdoor deployment tasks were

216
217
218
219

249 **GeoDE**
 250 (object classification)



273 **FishDetection**
 274 (Underwater video fish detection)



303 Figure 3: The five datasets in our benchmark: iWildCam, GeoDE, AutoArborist, FishDetection, and
 304 NuScenes each have real-world applications in deployment. In iWildCam, GeoDE, and AutoArborist,
 305 we show the class distributions of each deployment; in FishDetection, the number of detections per
 306 image is shown, and in NuScenes environment features. These diagrams show that each dataset
 307 has unique challenges in the deployments that lead to a need for model specialization, including
 308 long-tailedness (AutoArborist, iWildCam), covariate shift (all), subpopulation shifts (GeoDE, FishDe-
 309 tection), and more. These axes of variation are described in depth in Section 4 and further in Apdx E.

270 selected for enabling model specialization. The training dataset includes images from all countries,
 271 and the test data contains only images from Nigeria and Indonesia.
 272

273 **Knowledge-driven Subset:** These subsets were generated by selecting data from the relevant
 274 countries/categories in the training data, ie. only selecting African subcontinent data for the Nigeria
 275 deployment, Asian subcontinent data for the Indonesia deployment, and indoor/outdoor objects within
 276 the training pool for these deployments.

277 4.3 AUTOARBORIST 278

279 **Motivation:** Ecological imagery for environmental monitoring, such as automated tree classification,
 280 provides policymakers with critical, data-driven insights to support climate adaptation, urban planning,
 281 and more (Brandt et al., 2016). This task is associated with fundamental challenges such as noisy
 282 labels, non-iid data, fine-grained and long-tailed class distribution, and geospatial distribution shift.
 283 These challenges lead to a need for specialization of models where general-purpose models fail.

284 **Problem Setting & Data:** The AutoArborist dataset is a multi-view, fine-grained visual tree categori-
 285 zation dataset containing street-level images of over 1 million public zone trees from 300 genus-level
 286 categories across 23 major cities in the US and Canada.

287 **Deployments:** Deployments in AutoArborist correspond to the development models for use by
 288 individual cities. The deployment cities of (1) Surrey with 66 distinct tree genus classes, (2) Calgary
 289 with 30 classes, (3) Los Angeles with 175 classes, and (4) Washington DC with 67 classes were
 290 chosen due to their diverse climates, species distributions, and urban structures, as seen in Figure 3.
 291 Surrey and Calgary were treated as our in-distribution (ID) deployments, with some of these cities
 292 data in the training pool. Washington DC and LA were the out-of-distribution deployments, with no
 293 city data in the training pool.

294 **Knowledge-driven Subset:** We used the relevant data from Surrey and Calgary in the training pool
 295 for these ID deployments. Accordingly, we used data from San Francisco and San Jose for Los
 296 Angeles and Charlottesville, Pittsburgh, and New York for Washington DC. Label distribution shift
 297 and covariate shift are visualized in Figure 10 and 9, respectively.

298 4.4 FISHDETECTION 299

300 **Motivation:** Climate change, pollution, and overfishing continue to threaten marine biodiversity
 301 across the globe (United Nations, 2023; Di Lorenzo et al., 2022). Marine imagery is an increasingly
 302 common resource to monitor fish stocks and biodiversity. However, ML methods are difficult to apply
 303 across various environmental settings due to differences in lighting, turbidity, species, vegetation,
 304 camera sensors, etc. (Borremans et al., 2024; Jerlov, 1976; Akkaynak & Treibitz, 2019), creating a
 305 need for specialized models.

306 **Problem Setting & Data:** We use the public VIAME FishTrack23 dataset (Dawkins et al., 2017)
 307 consisting of 854,078 images across various environmental settings, ranging from freshwater rivers
 308 to deeper benthos. Specifically, the task is to predict bounding box localizations around every fish
 309 present in each image. Performance is measured by mAP across various IoU thresholds. Most of
 310 the images across all datasets are taken from video streams, and can be grouped as such, that were
 311 deployed primarily on camera traps, both baited and unbaited.

312 **Deployments:** Deployments are split according to the subsets of the VIAME dataset, which roughly
 313 correspond to geographic regions. Train, test and subset splits are either taken as provided or randomly
 314 sampled frames from each dataset, roughly corresponding to: (1) freshwater Pacific Northwest lakes;
 315 (2) Pacific Ocean; (3) East Atlantic Ocean; and (4) Gulf of Mexico.

316 **Knowledge-driven Subset:** For each deployment, we use the subset in the relevant geographical
 317 area (e.g., images from Gulf of Mexico for the Gulf of Mexico deployment).

318 4.5 NUSCENES 319

320 **Motivation:** End-to-end autonomous driving systems streamline vehicle control by directly mapping
 321 sensory inputs, such as images, to control outputs like steering angles (Wang et al., 2024). Adapting
 322 these systems to specialize in particular streets or environments is made easier as a single model

324 encompasses the full system. Thus, training this model to specialize in a specific environment brings
 325 advantages, capturing detailed local road layouts, traffic patterns, area-specific obstacles, and more.
 326

327 **Problem Setting & Data:** We explore vision-based control for self-driving across diverse environments (e.g., different city areas) and driving scenarios (e.g., pedestrians crossing, construction zones),
 328 formulated as a regression task. This dataset includes 88,461 images from the NuScenes dataset,
 329 subsampled from the image sweeps at a rate of 2. The images were captured from a video stream
 330 recorded while driving a car. Each image is paired with a steering angle control from the CAN bus,
 331 synchronized with the sensor timestamps of both the camera and CAN bus data. The model’s goal is
 332 to predict a single scalar value representing the car’s steering angle. Performance is evaluated in an
 333 open-loop manner using metrics like mean squared error.

334 **Deployments:** Deployments are organized by the geographic locations where the data was collected,
 335 including (1) Boston Seaport, (2) Singapore Holland Village, (3) Singapore One-North, and (4) Singa-
 336 pore Queenstown. While all tasks are based on expert demonstrations of driving and general driving
 337 behaviors, each location presents varying environmental features—such as vegetation, road types,
 338 roadside infrastructure, and weather—as well as differences in driving style and road regulations.
 339 Train/test splits are randomly sampled within each deployment.

340 **Knowledge-driven subset:** Since this training pool is a combination of the four deployment
 341 locations, we simply use the relevant location’s data as the training subset. For example, we use the
 342 subset of the training pool with Boston Seaport data for the Boston Seaport deployment.
 343

344 4.6 BENCHMARK PIPELINE

345 To compete on our benchmark, models must select relevant data from the training pool and then
 346 finetune models on that relevant data. Explicitly, **(i)** given a small query set representing the
 347 deployment data (we consider both labeled and unlabeled query sets), curate a subset of data from
 348 the training for a specific deployment, **(ii)** finetune/train a fixed model on the chosen subset from the
 349 training pool and **(iii)** evaluate on the deployment (test) set (Figure 2).

350 For each dataset, we fix the training procedure for all subsets, fixing model architecture, optimizers,
 351 and loss functions. We match the label distribution of the query set to the deployment/test set as
 352 closely as possible using stratified sampling, but from each class of the training pool, we sample
 353 uniformly at random. We run a small hyperparameter sweep for each training subset across batch
 354 sizes {32, 64, 128} and learning rates {0.01, 0.001, 0.0001} for each deployment. For all classifica-
 355 tion/regression datasets, we use ResNet50 for full-finetuning (He et al., 2015) (table 1) and a ViT for
 356 LoRA finetuning (Apdx Table 3), as well as a ViT (Dosovitskiy et al., 2020) for linear probes (Apdx
 357 Table 2). For the detection dataset, we use a YOLOv8n model, using default parameters, though we
 358 subsample images to 640p. Full details are in Apdx D.

361 4.7 METRICS

362 Participants are evaluated across 12 deployments from five datasets, as outlined in Section 4. For
 363 the classification task datasets of GeoDE, AutoArborist, and iWildCam, we report accuracy for each
 364 deployment, for the regression task dataset NuScenes, we report mean squared error, and for the
 365 detection task FishDetection, we report mAP50. For each deployment, we evaluate participants of
 366 the benchmark on overall accuracy of training subset; we also report subset size – while the less data
 367 used the better, we mainly focus on optimal performance, in line with the DS3 formulation.
 368

370 5 BASELINES

371 We compare performance of dataset subset selection algorithms across our benchmark, across different
 372 scenarios: (a) access to an unlabeled query set, and (b) access to a labeled query set. We also curate a
 373 third category, (c), which leverages domain expertise to generate expert-selected subsets, in order to
 374 demonstrate the existence of better-than-all subsets for these deployments.

375 **Non-subset comparisons:**

376 No filtering: Performance of a model trained on the entire training pool, without any filtering.

378 *Query Sets:* As a comparison, we also include performance of a model trained directly on the labeled
 379 query set for each deployment. Note that this would require access to query labels, which are not
 380 always available. When labels are available, performance of models trained on the small query sets
 381 are often poor, hence the value of learning from larger-scale general-pool data. As a logistical point,
 382 none of the baselines we show in our results train on query set data.

383 **Expert-Driven Subsets:** We contribute curated, "expert knowledge" subsets using domain knowledge
 384 and/or metadata. We find these knowledge-guided subsets often outperform using all samples in the
 385 training pool (no filtering). The creation of these subsets is described per-dataset in section 4.
 386

387 **Unlabeled-query baselines:**

388 *Image-alignment (Image-Align):* We take the cosine similarity between the training and query embedding
 389 space, using examples that exceed a threshold for at least x samples, where x is a hyperparameter
 390 chosen from $\{1, 10, 100\}$.

391 *Nearest neighbors features (Near-Nbors):* To better align our method with the downstream deployment,
 392 we explore using examples whose embedding space overlaps with the query set of data. To
 393 do so, we cluster image embeddings extracted by an OpenAI ViT model for each image into 1000
 394 clusters using Faiss (Johnson et al., 2019). Then, we find the nearest neighbor clusters for every query
 395 set example and keep the training cluster closest to each query set cluster. This method was inspired
 396 by the similar DataComp baseline (Gadre et al., 2024).

397 **Labeled-query baselines:**

398 *CLIP score filtering (CLIP-score):* We also experiment with CLIP score filtering, using examples
 399 that exceed a threshold for cosine similarity between CLIP image and text similarity. Text for each
 400 image was created with manual captioning (e.g. for iWildCam, *"This is a camera trap image of a*
 401 *lion taken at time 10-2-2016 at 04:26:13 in Nigeria"*). We select the subset that exceeds a threshold
 402 of CLIP-score similarity, with the threshold calculated for subsets that make up 25%, 50%, 75%, and
 403 90% of the dataset.

404 *Matching relative frequency (Match-Dist):* We explore having access to the relative frequency of
 405 each label in the downstream deployment. For example, a domain expert at a national park might
 406 know the relative frequency of species (deployment-specific domain knowledge). We create subsets
 407 by sampling 25%, 50%, 75%, and 90% of the training pool to match the label distribution of the
 408 deployment.

409 *Matching labels (Match-Label):* Similarly, a domain expert may know the classes present in the
 410 downstream deployment. For example, a domain expert at a national park might know the species
 411 present (deployment-specific knowledge) that we can utilize for subset selection. For these subsets,
 412 we simply remove the classes present in the training pool that are not present in the testing pool.

414
 415 **6 RESULTS**

416
 417 **Well chosen subsets outperform training on all data.** The knowledge-driven subsets in Table 1
 418 show that deployment-specific well-chosen subsets of the data can significantly outperform models
 419 trained on all the data, with improvements in deployment accuracy up to 3.6% for GeoDE, 11.9%
 420 for iWildCam, 51.3% for AutoArborist, a 0.03 reduction in MSE for NuScenes, and 0.13 increase
 421 in mAP50 for FishDetection. Even when the knowledge-driven subsets underperform all training
 422 data, as in NuScenes Deployment 2, there exist subsets from other baselines that outperform using all
 423 the data. In Appendix E, we provide a additional breakdown of the key factors that contributed to
 424 performance gain on these knowledge-driven subsets.

425 **Training on more data has diminishing returns.** For all deployments, we see that we can achieve
 426 near-optimal performance with subsets of the data. The knowledge-driven subsets are significantly
 427 smaller than the total training data size, with the average percentage of the total training pool used
 428 being 4% for GeoDE, 11% for iWildCam, 8% for AutoArborist, 10% for NuScenes, and 20% for
 429 FishDetection. Appendix C shows that even 25% of the data can perform near-optimally in some
 430 cases, with little performance loss with 50% of the data on the algorithmic baselines. Overall, these
 431 results demonstrate that greater efficiency for training specialized ML models is possible, potentially
 reducing computational and data storage burdens in deployable settings.

432 433	Dataset	Metric	Deploy #	Non subset		Knowledge-driven	Unlabeled query set		Labeled query set		
				Query-set	All-data		Image-Align	Near-Nbors	CLIP-score	Match-Label	Match-Dist
434 435	GeoDE	Acc (#)	Deploy 1	0.87 (500)	0.89 (53k)	0.92 (2.9k)	0.88 (26k)	0.88 (48k)	0.89 (40k)	0.88 (53k)	0.89 (48k)
			Deploy 2	0.45 (500)	0.89 (53k)	0.91 (2.6k)	0.90 (26k)	0.89 (48k)	0.90 (40k)	0.90 (53k)	0.88 (27k)
			Deploy 3	0.95 (500)	0.82 (53k)	0.85 (1.4k)	0.85 (24k)	0.76 (48k)	0.84 (40k)	0.83 (1.4k)	0.88 (48k)
			Deploy 4	0.83 (500)	0.83 (53k)	0.85 (2.6k)	0.79 (24k)	0.78 (48k)	0.83 (40k)	0.84 (2.6k)	0.83 (13k)
436 437	iWildCam	Acc (#)	Deploy 1	0.70 (301)	0.66 (130k)	0.65 (8.5k)	0.56 (36k)	0.50 (117k)	0.50 (97k)	0.74 (8.1k)	0.74 (117k)
			Deploy 2	0.78 (302)	0.34 (130k)	0.35 (9.2k)	0.44 (45k)	0.47 (98k)	0.46 (97k)	0.35 (55k)	0.49 (65k)
			Deploy 3	0.44 (301)	0.72 (130k)	0.75 (19k)	0.54 (24k)	0.45 (98k)	0.42 (97k)	0.72 (60k)	0.75 (117k)
			Deploy 4	0.46 (309)	0.66 (130k)	0.67 (21k)	0.60 (22k)	0.60 (33k)	0.29 (97k)	0.69 (57k)	0.74 (33k)
438 439	AutoArborist	Acc (#)	Deploy 1	0.16 (1.5k)	0.35 (781k)	0.86 (70k)	0.38 (44k)	0.39 (391k)	0.38 (47k)	0.67 (368k)	0.74 (703k)
			Deploy 2	0.20 (1.5k)	0.48 (781k)	0.86 (123k)	0.11 (49k)	0.14 (703k)	0.14 (47k)	0.65 (532k)	0.56 (391k)
			Deploy 3	0.12 (1.5k)	0.16 (781k)	0.38 (35k)	0.16 (46k)	0.10 (703k)	0.17 (47k)	0.16 (534k)	0.23 (703k)
			Deploy 4	0.12 (1.5k)	0.14 (781k)	0.39 (26k)	0.10 (48k)	0.11 (391k)	0.11 (47k)	0.10 (527k)	0.23 (195k)
440 441	NuScene	MSE (#)	Deploy 1	0.063 (6.0k)	0.050 (100k)	0.029 (20k)	0.040 (35k)	0.040 (90k)	0.073 (31k)	-	-
			Deploy 2	0.070 (1.0k)	0.021 (100k)	0.049 (4.6k)	0.15 (17k)	0.042 (90k)	0.032 (31k)	-	-
			Deploy 3	0.089 (2.7k)	0.068 (100k)	0.038 (10k)	0.049 (28k)	0.13 (90k)	0.071 (31k)	-	-
			Deploy 4	0.12 (1.9k)	0.048 (100k)	0.039 (7.0k)	0.086 (26k)	0.39 (90k)	0.050 (31k)	-	-
442 443	FishDetection	mAP50 (#)	Deploy 1	0.22 (500)	0.68 (841k)	0.69 (179k)	0.50 (630k)	0.60 (103k)	-	-	-
			Deploy 2	0.26 (600)	0.32 (841k)	0.45 (152k)	0.31 (630k)	0.40 (120k)	-	-	-
			Deploy 3	0.13 (541)	0.32 (841k)	0.39 (6.0k)	0.28 (630k)	0.23 (204k)	-	-	-
			Deploy 4	0.079 (519)	0.59 (841k)	0.60 (320k)	0.54 (630k)	0.39 (45k)	-	-	-

Table 1: Best-performing subsets across hyperparameters for baseline methods across all datasets and deployments (abbreviated as Deploy) for YOLOv8 full-finetuning for FishDetection and ResNet50 full-finetuning for the rest. Accuracy is reported for the classification tasks of GeoDE, iWildCam, and Auto Arborist, mAP50 for FishDetection (greater is better), and MSE for NuScenes (smaller is better). We include subset size in parentheses. We include results for **ViT LoRA finetuning** and **ViT linear probes** in Appendix C in Table 3 and Table 2, which display similar trends. Match-Dist and Match-Label are not applicable for NuScenes, as it is a regression task and does not have clear classes/labels for these methods. FishDetection only uses the unlabeled query set, as the ground truth is bounding boxes, rather than labels themselves. Baselines are distinguished from one another by their access to information, with each baseline having access to expert knowledge, or a labeled/unlabeled query set. We do not report the random baseline in this table, but demonstrate results in Appendix C as it mainly refers to subset size. For each deployment, there exists a subset that outperforms training on all data, indicated in bold.

Methods without access to supervision perform poorly. While the knowledge-driven subsets in Table 1 demonstrate that a well-chosen subset *does exist* for all deployments, finding this subset without extra knowledge is still an open problem. Some of our baselines require access to query labels, this requirement can in many cases be unrealistic in the deployable ML setting (labels can be expensive or difficult to collect). The two unsupervised baselines, the nearest neighbors and image alignment methods, do not perform optimally on the deployments, often underperforming using all the training data. Our benchmark opens up the line of research for potential unsupervised methods for this data subset selection process.

7 DISCUSSION AND CONCLUSIONS

DATAS³ is the first benchmark to promote the development of dataset subset selection methods capable of specialization to diverse real-world deployments. The benchmark is both open-source and easy to use, lowering the barrier to entry for this new important problem.

DATAS³ highlights open challenges for the research community. In our experimental study, we show that there is no winning baseline that performs well across multiple domains/datasets. Additionally, while some methods perform well when given access to labeled query sets, no methods perform well in the unsupervised setting. Finally, some datasets are more challenging than others—methods may need to specifically target different types of distribution shifts.

DATAS³ has value beyond subset selection. In addition to the DS3 problem, DATAS³ can be used as a testbed for various other complementary lines of work, such as domain adaptation, active learning, coresnet selection, and more. We highlight these relevant methods in Appendix B.

Extensions to other domains. Model specialization for deployments isn’t limited to the domains we include. We are open to expanding this benchmark to capture more scientific domains, and welcome further dataset contributions from the broader ML and scientific research community.

486 REFERENCES
487

488 Derya Akkaynak and Tali Treibitz. Sea-thru: A method for removing water from underwater images.
489 In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp.
490 1682–1691, 2019.

491 Anonymous. When less is more: Investigating data pruning for pretraining LLMs at scale. In *NeurIPS*
492 *Workshop on Attributing Model Behavior at Scale*, 2023. URL <https://openreview.net/forum?id=XUIYn3jo5T>.
493

494 Koichiro Asano, Akira Hebisawa, Takashi Ishiguro, Noboru Takayanagi, Yasuhiko Nakamura, Junko
495 Suzuki, Naoki Okada, Jun Tanaka, Yuma Fukutomi, Shigeharu Ueki, Koichi Fukunaga, Satoshi
496 Konno, Hiroto Matsuse, Katsuhiko Kamei, Masami Taniguchi, Terufumi Shimoda, and Tsuyoshi
497 Oguma. New clinical diagnostic criteria for allergic bronchopulmonary aspergillosis/mycosis
498 and its validation. *The Journal of allergy and clinical immunology*, 2020. URL <https://api.semanticscholar.org/CorpusID:221673224>.
499

500 Marcus A. Badgeley, John R. Zech, Luke Oakden-Rayner, Benjamin Scott Glicksberg, Manway Liu,
501 William Gale, Michael V. McConnell, Bethany Percha, Thomas M. Snyder, and Joel T. Dudley.
502 Deep learning predicts hip fracture using confounding patient and healthcare variables. *NPJ*
503 *Digital Medicine*, 2, 2018. URL <https://api.semanticscholar.org/CorpusID:53250171>.
504

505 Fred Bane, Celia Soler Uguet, Wiktor Stribiżew, and Anna Zaretskaya. A comparison of data filtering
506 methods for neural machine translation. In *Proceedings of the 15th Biennial Conference of the*
507 *Association for Machine Translation in the Americas (Volume 2: Users and Providers Track*
508 *and Government Track)*, pp. 313–325, Orlando, USA, September 2022. Association for Machine
509 Translation in the Americas. URL <https://aclanthology.org/2022.amta-upg.22>.
510

511 Cenk Baykal, Lucas Liebenwein, Igor Gilitschenski, Dan Feldman, and Daniela Rus. Sensitivity-
512 informed provable pruning of neural networks. *SIAM J. Math. Data Sci.*, 4(1):26–45, 2022.
513

514 Sara Beery, Dan Morris, and Siyu Yang. Efficient pipeline for camera trap image review, 2019. URL
515 <https://arxiv.org/abs/1907.06772>.
516

517 Sara Beery, Arushi Agarwal, Elijah Cole, and Vighnesh Birodkar. The iwildcam 2021 competition
518 dataset, 2021.
519

520 Sara Beery, Guanhang Wu, Trevor Edwards, Filip Pavetic, Bohdan S. Majewski, Shreyasee Mukherjee,
521 Stanley Chan, John Morgan, Vivek Rathod, and Jonathan Huang. The auto arborist dataset: A
522 large-scale benchmark for multiview urban forest monitoring under domain shift. 2022 *IEEE/CVF*
523 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21262–21275, 2022a. URL
524 <https://api.semanticscholar.org/CorpusID:250476238>.
525

526 Sara Meghan Beery, Guanhang Wu, Trevor Edwards, Filip Pavetić, Bo Majewski, Shreyasee Mukher-
527 jee, Stan Chan, John Morgan, Vivek Mansing Rathod, and Jonathan Chung-kuan Huang. The
528 auto-arborist dataset: A large-scale benchmark for generalizable, multimodal urban forest monitor-
529 ing. 2022b.
530

531 Stella Biderman, Hailey Schoelkopf, Quentin Anthony, Herbie Bradley, Kyle O’Brien, Eric Hallahan,
532 Mohammad Aflah Khan, Shivanshu Purohit, USVSN Sai Prashanth, Edward Raff, Aviya Skowron,
533 Lintang Sutawika, and Oskar van der Wal. Pythia: A suite for analyzing large language models
534 across training and scaling, 2023.
535

536 Sidney Black, Stella Biderman, Eric Hallahan, Quentin Anthony, Leo Gao, Laurence Golding, Horace
537 He, Connor Leahy, Kyle McDonell, Jason Phang, Michael Pieler, Usvsn Sai Prashanth, Shivanshu
538 Purohit, Laria Reynolds, Jonathan Tow, Ben Wang, and Samuel Weinbach. GPT-NeoX-20B: An
539 open-source autoregressive language model. In *Proceedings of BigScience Episode #5 – Workshop*
540 *on Challenges & Perspectives in Creating Large Language Models*, pp. 95–136, virtual+Dublin,
541 May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.bigsience-1.9.
542 URL <https://aclanthology.org/2022.bigsience-1.9>.
543

540 Catherine Borremans, Jennifer Durden, Timm Schoening, Emma Curtis, Luther Adams, Alexandra Branzan Albu, Aurélien Arnaubec, Sakina-Dorothée Ayata, Reshma Baburaj, Corinne Bassin, et al. Report on the marine imaging workshop 2022. *Research Ideas and Outcomes*, 10:e119782, 2024.

544 Leslie A. Brandt, Abigail Derby Lewis, Robert T. Fahey, Lydia Scott, Lindsay E. Darling, and Christopher W. Swanston. A framework for adapting urban forests to climate change. *Environmental Science & Policy*, 66:393–402, 2016. URL <https://api.semanticscholar.org/CorpusID:156304225>.

548 Vladimir Braverman, Dan Feldman, and Harry Lang. New frameworks for offline and streaming coresets constructions. *arXiv preprint arXiv:1612.00889*, 2016.

551 Vladimir Braverman, Petros Drineas, Cameron Musco, Christopher Musco, Jalaj Upadhyay, David P. Woodruff, and Samson Zhou. Near optimal linear algebra in the online and sliding window models. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS*, pp. 517–528, 2020.

554 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeff Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Ma teusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. *ArXiv*, abs/2005.14165, 2020. URL <https://api.semanticscholar.org/CorpusID:218971783>.

561 Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 11621–11631, 2020.

566 Ke Chen. On coresets for k-median and k-means clustering in metric and euclidean spaces and their applications. *SIAM J. Comput.*, 39(3):923–947, 2009.

568 Rachit Chhaya, Anirban Dasgupta, and Supratim Shit. On coresets for regularized regression. In *Proceedings of the 37th International Conference on Machine Learning, ICML*, 2020.

571 Kenneth L. Clarkson. Coresets, sparse greedy approximation, and the frank-wolfe algorithm. *ACM Trans. Algorithms*, 6(4):63:1–63:30, 2010.

573 Michael B. Cohen, Cameron Musco, and Christopher Musco. Input sparsity time low-rank approximation via ridge leverage score sampling. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pp. 1758–1777, 2017.

576 Vincent Cohen-Addad, Kasper Green Larsen, David Saulpic, Chris Schwiegelshohn, and Omar Ali Sheikh-Omar. Improved coresets for euclidean k -means. *Advances in Neural Information Processing Systems*, 35:2679–2694, 2022.

580 Cody Coleman, Christopher Yeh, Stephen Mussmann, Baharan Mirzasoleiman, Peter Bailis, Percy Liang, Jure Leskovec, and Matei Zaharia. Selection via proxy: Efficient data selection for deep learning. In *International Conference on Learning Representations*, 2019.

583 Rhys Compton, Lily H. Zhang, Aahlad Manas Puli, and Rajesh Ranganath. When more is less: Incorporating additional datasets can hurt performance by introducing spurious correlations. *ArXiv*, abs/2308.04431, 2023. URL <https://api.semanticscholar.org/CorpusID:265819574>.

587 Anirban Dasgupta, Petros Drineas, Boulos Harb, Ravi Kumar, and Michael W. Mahoney. Sampling algorithms and coresets for l_p regression. In *Proceedings of the Nineteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pp. 932–941, 2008.

591 Matthew Dawkins, Linus Sherrill, Keith Fieldhouse, Anthony Hoogs, Benjamin Richards, David Zhang, Lakshman Prasad, Kresimir Williams, Nathan Lauffenburger, and Gaoang Wang. An open-source platform for underwater image and video analytics. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 898–906, 2017. doi: 10.1109/WACV.2017.105.

594 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
 595 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,
 596 pp. 248–255, 2009a. doi: 10.1109/CVPR.2009.5206848.

597 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
 598 archical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*,
 599 pp. 248–255, 2009b. doi: 10.1109/CVPR.2009.5206848.

600 Emanuele Di Lorenzo, Christian Lønborg, Jesper H Andersen, Elva G Escobar-Briones, Michelle Jil-
 601 lian Devlin, Angel Borja, Marius Nils Müller, Carol Robinson, Alex Ford, Anna Milena Zivian,
 602 et al. *Sustainable Development Goal 14-Life Below Water: Towards a Sustainable Ocean*. Frontiers
 603 Media SA, 2022.

604 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 605 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit,
 606 and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale.
 607 *ArXiv*, abs/2010.11929, 2020. URL <https://api.semanticscholar.org/CorpusID:225039882>.

608 Liat Ein-Dor, Alon Halfon, Ariel Gera, Eyal Shnarch, Lena Dankin, Leshem Choshen, Marina
 609 Danilevsky, Ranit Aharonov, Yoav Katz, and Noam Slonim. Active Learning for BERT: An Em-
 610 pirical Study. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language
 Processing (EMNLP)*, pp. 7949–7962, Online, November 2020. Association for Computational
 611 Linguistics. doi: 10.18653/v1/2020.emnlp-main.638. URL <https://aclanthology.org/2020.emnlp-main.638>.

612 Abolfazl Farahani, Sahar Voghieri, Khaled Rasheed, and Hamid R. Arabnia. A brief review of domain
 613 adaptation, 2020. URL <https://arxiv.org/abs/2010.03978>.

614 Benjamin Feuer, Jiawei Xu, Niv Cohen, Patrick Yubeaton, Govind Mittal, and Chinmay Hegde.
 615 Select: A large-scale benchmark of data curation strategies for image classification. *arXiv preprint
 arXiv:2410.05057*, 2024.

616 Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smirnis, Thao Nguyen,
 617 Ryan Marten, Mitchell Wortsman, Dhruba Ghosh, Jieyu Zhang, et al. Datacomp: In search of the
 618 next generation of multimodal datasets. *Advances in Neural Information Processing Systems*, 36,
 619 2024.

620 Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard S. Zemel, Wieland Brendel,
 621 Matthias Bethge, and Felix Wichmann. Shortcut learning in deep neural networks. *Nature
 Machine Intelligence*, 2:665 – 673, 2020. URL <https://api.semanticscholar.org/CorpusID:215786368>.

622 Sariel Har-Peled and Soham Mazumdar. On coresets for k-means and k-median clustering. In
 623 *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pp. 291–300, 2004.

624 Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition.
 625 *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2015.
 626 URL <https://api.semanticscholar.org/CorpusID:206594692>.

627 Duojun Huang, Jichang Li, Weikai Chen, Jun Steed Huang, Zhenhua Chai, and Guanbin Li. Divide
 628 and adapt: Active domain adaptation via customized learning. *2023 IEEE/CVF Conference
 629 on Computer Vision and Pattern Recognition (CVPR)*, pp. 7651–7660, 2023. URL <https://api.semanticscholar.org/CorpusID:260067836>.

629 Lingxiao Huang and Nisheeth K. Vishnoi. Coresets for clustering in euclidean spaces: importance
 630 sampling is nearly optimal. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on
 631 Theory of Computing, STOC*, pp. 1416–1429, 2020.

632 Nils Gunnar Jerlov. *Marine optics*, volume 14. Elsevier, 1976.

633 Glenn Jocher, Jing Qiu, and Ayush Chaurasia. Ultralytics YOLO, January 2023. URL <https://github.com/ultralytics/ultralytics>.

648 Jeff Johnson, Matthijs Douze, and Hervé Jégou. Billion-scale similarity search with GPUs. *IEEE*
 649 *Transactions on Big Data*, 7(3):535–547, 2019.
 650

651 Ibrahim Jubran, Murad Tukan, Alaa Maalouf, and Dan Feldman. Sets clustering. In *International*
 652 *Conference on Machine Learning*, pp. 4994–5005. PMLR, 2020.

653 KrishnaTeja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, Abir De, and Rishabh K.
 654 Iyer. GRAD-MATCH: gradient matching based data subset selection for efficient deep model
 655 training. In *Proceedings of the 38th International Conference on Machine Learning, ICML*, pp.
 656 5464–5474, 2021a.

657 KrishnaTeja Killamsetty, Durga Sivasubramanian, Ganesh Ramakrishnan, and Rishabh K. Iyer.
 658 GLISTER: generalization based data subset selection for efficient and robust learning. In *Thirty-*
 659 *Fifth AAAI Conference on Artificial Intelligence, AAAI*, 2021b.

660 Krishnateja Killamsetty, Xujiang Zhao, Feng Chen, and Rishabh Iyer. Retrieve: Coreset selection
 661 for efficient and robust semi-supervised learning. In *Advances in Neural Information Processing*
 662 *Systems (NeurIPS)*, 2021c.

663 Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsub-
 664 ramani, Weihua Hu, Michihiro Yasunaga, Richard Lanas Phillips, Irena Gao, Tony Lee, Etiene
 665 David, Ian Stavness, Wei Guo, Berton A. Earnshaw, Imran S. Haque, Sara Meghan Beery, Jure
 666 Leskovec, Anshul Kundaje, Emma Pierson, Sergey Levine, Chelsea Finn, and Percy Liang. Wilds:
 667 A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*,
 668 2020. URL <https://api.semanticscholar.org/CorpusID:229156320>.

669 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images, 2009.

670 Alina Kuznetsova, Hassan Rom, Neil Gordon Alldrin, Jasper R. R. Uijlings, Ivan Krasin, Jordi Pont-
 671 Tuset, Shahab Kamali, Stefan Popov, Matteo Malloci, Alexander Kolesnikov, Tom Duerig, and
 672 Vittorio Ferrari. The open images dataset v4. *International Journal of Computer Vision*, 128:1956
 673 – 1981, 2018. URL <https://api.semanticscholar.org/CorpusID:53296866>.

674 Hugo Laurençon, Lucile Saulnier, Thomas Wang, Christopher Akiki, Albert Villanova del Moral,
 675 Teven Le Scao, Leandro Von Werra, Chenghao Mou, Eduardo González Ponferrada, Huu Nguyen,
 676 Jörg Frohberg, Mario Šaško, Quentin Lhoest, Angelina McMillan-Major, Gerard Dupont, Stella
 677 Biderman, Anna Rogers, Loubna Ben allal, Francesco De Toni, Giada Pistilli, Olivier Nguyen,
 678 Somaieh Nikpoor, Maraim Masoud, Pierre Colombo, Javier de la Rosa, Paulo Villegas, Tristan
 679 Thrush, Shayne Longpre, Sebastian Nagel, Leon Weber, Manuel Muñoz, Jian Zhu, Daniel Van
 680 Strien, Zaid Alyafeai, Khalid Almubarak, Minh Chien Vu, Itziar Gonzalez-Dios, Aitor Soroa,
 681 Kyle Lo, Manan Dey, Pedro Ortiz Suarez, Aaron Gokaslan, Shamik Bose, David Adelani, Long
 682 Phan, Hieu Tran, Ian Yu, Suhas Pai, Jenny Chim, Violette Lepercq, Suzana Ilic, Margaret Mitchell,
 683 Sasha Alexandra Luccioni, and Yacine Jernite. The bigscience roots corpus: A 1.6tb composite
 684 multilingual dataset, 2023.

685 Lucas Liebenwein, Cenk Baykal, Harry Lang, Dan Feldman, and Daniela Rus. Provable filter pruning
 686 for efficient neural networks. In *International Conference on Learning Representations*, 2019.

687 Alaa Maalouf, Ibrahim Jubran, and Dan Feldman. Fast and accurate least-mean-squares solvers. In
 688 *Proceedings of the 33rd International Conference on Neural Information Processing Systems*, pp.
 689 8307–8318, 2019.

690 Alaa Maalouf, Ibrahim Jubran, Murad Tukan, and Dan Feldman. Coresets for the average case error
 691 for finite query sets. *Sensors*, 21(19):6689, 2021.

692 Alaa Maalouf, Gilad Eini, Ben Mussay, Dan Feldman, and Margarita Osadchy. A unified approach to
 693 coreset learning. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.

694 Mark Mazumder, Colby Banbury, Xiaozhe Yao, Bojan Karlaš, William Gaviria Rojas, Sudnya
 695 Diamos, Greg Diamos, Lynn He, Alicia Parrish, Hannah Rose Kirk, Jessica Quaye, Charvi Rastogi,
 696 Douwe Kiela, David Jurado, David Kanter, Rafael Mosquera, Juan Ciro, Lora Aroyo, Bilge Acun,
 697 Lingjiao Chen, Mehul Smriti Raje, Max Bartolo, Sabri Eyuboglu, Amirata Ghorbani, Emmett
 698 Goodman, Oana Inel, Tariq Kane, Christine R. Kirkpatrick, Tzu-Sheng Kuo, Jonas Mueller, Tristan

702 Thrush, Joaquin Vanschoren, Margaret Warren, Adina Williams, Serena Yeung, Newsha Ardalani,
 703 Praveen Paritosh, Lilith Bat-Leah, Ce Zhang, James Zou, Carole-Jean Wu, Cody Coleman, Andrew
 704 Ng, Peter Mattson, and Vijay Janapa Reddi. Dataperf: Benchmarks for data-centric ai development,
 705 2023. URL <https://arxiv.org/abs/2207.10062>.

706 Raphael A. Meyer, Cameron Musco, Christopher Musco, David P. Woodruff, and Samson Zhou. Fast
 707 regression for structured inputs. In *The Tenth International Conference on Learning Representa-
 708 tions, ICLR*, 2022.

709 Sören Mindermann, Jan Brauner, Muhammed Razzak, Mrinank Sharma, Andreas Kirsch, Winnie
 710 Xu, Benedikt Höltgen, Aidan N. Gomez, Adrien Morisot, Sebastian Farquhar, and Yarin Gal.
 711 Prioritized training on points that are learnable, worth learning, and not yet learnt, 2022.

712 Baharan Mirzasoleiman, Jeff A. Bilmes, and Jure Leskovec. Coresets for data-efficient training
 713 of machine learning models. In *Proceedings of the 37th International Conference on Machine
 714 Learning, ICML*, pp. 6950–6960, 2020a.

715 Baharan Mirzasoleiman, Kaidi Cao, and Jure Leskovec. Coresets for robust training of deep neural
 716 networks against noisy labels. In *Advances in Neural Information Processing Systems 33: Annual
 717 Conference on Neural Information Processing Systems, NeurIPS*, 2020b.

718 Niklas Muennighoff, Alexander M. Rush, Boaz Barak, Teven Le Scao, Aleksandra Piktus, Nouamane
 719 Tazi, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. Scaling data-constrained language models,
 720 2023.

721 Mohammad Sadegh Norouzzadeh, Dan Morris, Sara Beery, Neel Joshi, Nebojsa Jojic, and Jeff Clune.
 722 A deep active learning system for species identification and counting in camera trap images, 2019.
 723 URL <https://arxiv.org/abs/1910.09716>.

724 Mansheej Paul, Surya Ganguli, and Gintare Karolina Dziugaite. Deep learning on a data diet: Finding
 725 important examples early in training. In *Association for the Advancement of Artificial Intelligence
 726 (AAAI)*, 2021.

727 Guilherme Penedo, Quentin Malartic, Daniel Hesslow, Ruxandra Cojocaru, Alessandro Cappelli,
 728 Hamza Alobeidli, Baptiste Pannier, Ebtesam Almazrouei, and Julien Launay. The refinedweb
 729 dataset for falcon llm: Outperforming curated corpora with web data, and web data only, 2023.

730 Viraj Prabhu, Arjun Chandrasekaran, Kate Saenko, and Judy Hoffman. Active domain adaptation
 731 via clustering uncertainty-weighted embeddings. *2021 IEEE/CVF International Conference on
 732 Computer Vision (ICCV)*, pp. 8485–8494, 2020. URL <https://api.semanticscholar.org/CorpusID:224714171>.

733 Jack W. Rae, Sebastian Borgeaud, Trevor Cai, Katie Millican, Jordan Hoffmann, Francis Song, John
 734 Aslanides, Sarah Henderson, Roman Ring, Susannah Young, Eliza Rutherford, Tom Hennigan,
 735 Jacob Menick, Albin Cassirer, Richard Powell, George van den Driessche, Lisa Anne Hendricks,
 736 Maribeth Rauh, Po-Sen Huang, Amelia Glaese, Johannes Welbl, Sumanth Dathathri, Saffron
 737 Huang, Jonathan Uesato, John Mellor, Irina Higgins, Antonia Creswell, Nat McAleese, Amy Wu,
 738 Erich Elsen, Siddhant Jayakumar, Elena Buchatskaya, David Budden, Esme Sutherland, Karen
 739 Simonyan, Michela Paganini, Laurent Sifre, Lena Martens, Xiang Lorraine Li, Adhiguna Kuncoro,
 740 Aida Nematzadeh, Elena Gribovskaya, Domenic Donato, Angeliki Lazaridou, Arthur Mensch,
 741 Jean-Baptiste Lespiau, Maria Tsimpoukelli, Nikolai Grigorev, Doug Fritz, Thibault Sottiaux,
 742 Mantas Pajarskas, Toby Pohlen, Zhitao Gong, Daniel Toyama, Cyprien de Masson d’Autume,
 743 Yujia Li, Tayfun Terzi, Vladimir Mikulik, Igor Babuschkin, Aidan Clark, Diego de Las Casas,
 744 Aurelia Guy, Chris Jones, James Bradbury, Matthew Johnson, Blake Hechtman, Laura Weidinger,
 745 Iason Gabriel, William Isaac, Ed Lockhart, Simon Osindero, Laura Rimell, Chris Dyer, Oriol
 746 Vinyals, Kareem Ayoub, Jeff Stanway, Lorrayne Bennett, Demis Hassabis, Koray Kavukcuoglu,
 747 and Geoffrey Irving. Scaling language models: Methods, analysis and insights from training
 748 gopher, 2022.

749 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi
 750 Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text
 751 transformer, 2020.

756 Vikram V. Ramaswamy, Sing Yu Lin, Dora Zhao, Aaron B. Adcock, Laurens van der Maaten, Deepti
 757 Ghadiyaram, and Olga Russakovsky. Geode: a geographically diverse evaluation dataset for object
 758 recognition, 2023. URL <https://arxiv.org/abs/2301.02560>.

759

760 Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng
 761 Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Li Fei-Fei.
 762 Imagenet large scale visual recognition challenge. *International Journal of Computer Vision*, 115:
 763 211 – 252, 2014. URL <https://api.semanticscholar.org/CorpusID:2930547>.

764

765 Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set
 766 approach. In *International Conference on Learning Representations (ICLR)*, 2018.

767

768 Shreya Shankar, Yoni Halpern, Eric Breck, James Atwood, Jimbo Wilson, and D. Sculley. No
 769 classification without representation: Assessing geodiversity issues in open data sets for the
 770 developing world. *arXiv: Machine Learning*, 2017. URL <https://api.semanticscholar.org/CorpusID:26262581>.

771

772 Judy Hanwen Shen, Inioluwa Deborah Raji, and Irene Y. Chen. The data addition dilemma, 2024.
 773 URL <https://arxiv.org/abs/2408.04154>.

774

775 Shoaib Ahmed Siddiqui, Nitarshan Rajkumar, Tegan Maharaj, David Krueger, and Sara Hooker.
 776 Metadata archaeology: Unearthing data subsets by leveraging training dynamics, 2022.

777

778 Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari S. Morcos. Beyond neural
 779 scaling laws: beating power law scaling via data pruning. *arXiv*, 2022.

780

781 Ben Sorscher, Robert Geirhos, Shashank Shekhar, Surya Ganguli, and Ari S. Morcos. Beyond neural
 782 scaling laws: beating power law scaling via data pruning, 2023.

783

784 Francis Staub. Living planet report 2020: Bending the curve of
 785 biodiversity loss, Sep 2020. URL <https://icriforum.org/living-planet-report-2020-bending-the-curve-of-biodiversity-loss/>.

786

787 Alex Tamkin, Dat Nguyen, Salil Deshpande, Jesse Mu, and Noah Goodman. Active learning helps
 788 pretrained models learn the intended task, 2022.

789

790 Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt.
 791 Measuring robustness to natural distribution shifts in image classification. *ArXiv*, abs/2007.00644,
 792 2020. URL <https://api.semanticscholar.org/CorpusID:220280805>.

793

794 Elad Tolochinsky, Ibrahim Jubran, and Dan Feldman. Generic coresnet for scalable learning of
 795 monotonic kernels: Logistic regression, sigmoid and more. In *International Conference on
 796 Machine Learning, ICML*, 2022.

797

798 Murad Tukan, Cenk Baykal, Dan Feldman, and Daniela Rus. On coresets for support vector machines.
 799 *Theor. Comput. Sci.*, 890:171–191, 2021.

800

801 Murad Tukan, Loay Mualem, and Alaa Maalouf. Pruning neural networks via coresets and convex
 802 geometry: Towards no assumptions. In *Proceedings of the 36th International Conference on
 803 Neural Information Processing Systems*, 2022.

804

805 Murad Tukan, Samson Zhou, Alaa Maalouf, Daniela Rus, Vladimir Braverman, and Dan Feldman.
 806 Provable data subset selection for efficient neural networks training. In *International Conference
 807 on Machine Learning*, pp. 34533–34555. PMLR, 2023.

808

809 United Nations. Agreement under the United Nations Convention on the Law of the Sea on the
 810 Conservation and Sustainable Use of Marine Biological Diversity of Areas Beyond National
 811 Jurisdiction. United Nations Treaty Collection, 2023. Opened for signature on 20 September 2023.

812

813 Jindong Wang, Cuiling Lan, Chang Liu, Yidong Ouyang, and Tao Qin. Generalizing to unseen
 814 domains: A survey on domain generalization. *IEEE Transactions on Knowledge and Data Engi-
 815 neering*, 35:8052–8072, 2021. URL <https://api.semanticscholar.org/CorpusID:232110832>.

810 Tsun-Hsuan Wang, Alaa Maalouf, Wei Xiao, Yutong Ban, Alexander Amini, Guy Rosman, Sertac
 811 Karaman, and Daniela Rus. Drive anywhere: Generalizable end-to-end autonomous driving
 812 with multi-modal foundation models. In *2024 IEEE International Conference on Robotics and*
 813 *Automation (ICRA)*, pp. 6687–6694. IEEE, 2024.

814 Xinyi Wang, Hieu Pham, Paul Michel, Antonios Anastasopoulos, Jaime Carbonell, and Graham
 815 Neubig. Optimizing data usage via differentiable rewards. In *International Conference on Machine*
 816 *Learning (ICML)*, 2020.

817 Xudong Wang, Long Lian, and Stella X Yu. Unsupervised selective labeling for more effective
 818 semi-supervised learning. In *European Conference on Computer Vision*, pp. 427–445. Springer,
 820 2022.

821 Oliver Wearn and Paul Glover-Kapfer. Camera-trapping for conservation: a guide to best-practices,
 822 10 2017.

823 Guillaume Wenzek, Marie-Anne Lachaux, Alexis Conneau, Vishrav Chaudhary, Francisco Guzmán,
 824 Armand Joulin, and Edouard Grave. CCNet: Extracting high quality monolingual datasets from
 825 web crawl data. In *Proceedings of the Twelfth Language Resources and Evaluation Conference*,
 826 pp. 4003–4012, Marseille, France, May 2020. European Language Resources Association. ISBN
 827 979-10-95546-34-4. URL <https://aclanthology.org/2020.lrec-1.494>.

828 Michelle Yuan, Hsuan-Tien Lin, and Jordan Boyd-Graber. Cold-start active learning through
 829 self-supervised language modeling. In *Proceedings of the 2020 Conference on Empirical*
 830 *Methods in Natural Language Processing (EMNLP)*, pp. 7935–7948, Online, November 2020.
 831 Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.637. URL
 832 <https://aclanthology.org/2020.emnlp-main.637>.

833 Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuhui Chen, Christopher
 834 Dewan, Mona Diab, Xian Li, Xi Victoria Lin, Todor Mihaylov, Myle Ott, Sam Shleifer, Kurt
 835 Shuster, Daniel Simig, Punit Singh Koura, Anjali Sridhar, Tianlu Wang, and Luke Zettlemoyer.
 836 Opt: Open pre-trained transformer language models, 2022.

837 Chunting Zhou, Pengfei Liu, Puxin Xu, Srinivasan Iyer, Jiao Sun, Yuning Mao, Xuezhe Ma, Avia
 838 Efrat, Ping Yu, Lili Yu, et al. Lima: Less is more for alignment. *Advances in Neural Information*
 839 *Processing Systems*, 36, 2024.

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863