EEVEE AND GATE: FINDING THE RIGHT BENCH MARKS AND HOW TO RUN THEM SEAMLESSLY

Anonymous authors

004

010

Paper under double-blind review

ABSTRACT

011 Model evaluation is a cornerstone of machine learning, guiding model design and 012 progress measurement. Designing generalizable evaluation processes remains a 013 challenge, however, partly due to the vast number of possible domain, task and modality combinations and lack of knowledge of how informative they are. In 014 this paper, we propose *EEVEE* (Efficient Evaluation process Evolution Engine)¹, a 015 method that frames evaluation process design as a learning problem. By analyzing 016 a large number of evaluation metrics from diverse benchmarks and models, EEVEE 017 identifies a smaller subset of tasks with high predictive power over the full set of 018 evaluation metrics, reducing evaluation time. To find the optimal subset maximiz-019 ing signal while minimizing GPU hours, EEVEE evaluates pre-trained models of various architectures, pretraining schemes, and modalities on diverse downstream 021 tasks and datasets including image classification, segmentation, relational reasoning, zero-shot image-to-text tasks, medical classification and segmentation, video 023 classification, and regression. Our results identify three subsets of benchmarks, with 8, 15 and 21 tasks, providing high quality signal for model generalization. Key benchmarks selected include iWildCam, CLEVR-Math, ACDC, WinoGround, 025 CIFAR100, Fungi, and ADE20K. We structure the subsets into three tiers for 026 12, 24, and 36 GPU-hour budgets and package them into a unified, efficient, and 027 user-friendly Python framework that we built with the researcher in mind – which 028 we refer to as the GATE engine. Our experiments reveal ConvNextV2, SigLIP 029 and CLIP as top-performing model encoders, with EfficientNetV2 and ResNext50 excelling in medical tasks and challenging image classification, in particular in 031 Happy Whale Individual classification, ConvNet based models seem to outperform 032 transformer models by a factor of 2.5x, which is surprising. The top performing encoder being ConvNextV2 followed by CLIP seems to agree with other recent large 034 scale evaluations. We also demonstrate the framework's versatility in fine-tuning models from text and audio modalities, paving the way for future cross-modal evaluations.

037 038

1 INTRODUCTION

Increasing Complexities of Benchmarking: As we create benchmarks for expanding model capability evaluation, the growing number and complexity of these benchmarks inadvertently complicates evaluation, requiring more resources like engineering, computation, and research time. Consequently, prioritizing which benchmarks to use becomes challenging. The high costs and longer wait times of newer, complex benchmarks often deter their adoption, leading researchers to rely on older, simpler benchmarks. This risks missing valuable insights from innovative ideas that may underperform on simpler benchmarks but have broader applicability, while promoting incremental improvements that overfit to simpler benchmarks but underperform in comprehensive evaluations.

To illustrate the mounting increase in available benchmarks, we can look at the historical benchmarks in deep learning. Few benchmarks have had as much impact as ImageNet (30), which remains a rich resource for model training and evaluation, particularly in visuo-linguistic models. As key capabilities for deep neural networks were discovered, more benchmarks were generated to measure and stimulate progress in those areas. In natural language processing, the GLUE benchmark (69),

¹Pronounced as /'i:vi:/ EE-vee

SQuAD (48), and CoNLL-2003 (51) have been instrumental. In audio processing, LibriSpeech (42), TIMIT (16), and VCTK (72) are widely used. For machine translation, WMT (3), IWSLT (23), and Europarl (26) have driven advancements. Relational reasoning has been advanced by benchmarks such as CLEVR (24), bAbI (70), and RAVEN (75). In segmentation, PASCAL VOC (14), Cityscapes (8), and COCO (36) remain crucial. Large language models are often evaluated using benchmarks like SuperGLUE (68), LAMBADA (43), and MMLU (20). Vision-language models are typically evaluated using benchmarks such as VQA (1), Visual7W (81), and Flickr30k (45).

061 As a result, a researcher has to choose from all these options, and even more, and then find a 062 way to unify and experiment with their models across all of them. The lack of unification, and 063 the lack of guarantees for their generalization signal, quickly becomes a kind of "evaluation hell", 064 where researchers waste a lot of time just doing redundant things like fixing the same bugs to download datasets, preprocess them etc, while at the same time not having any real signal as to which 065 benchmarks are more informative, other than just knowing what has been used the most – which is 066 usually a function of popularity, and not real informativeness. To elaborate, the adoption of complex 067 evaluation processes that could enhance research efficiency and impact is often hindered by the 068 engineering effort required to evaluate machine learning models. Researchers must create involved 069 pipelines across multiple datasets demanding high data engineering efforts, develop task-specific 070 adapters, and derive nuanced training recipes, which is time-consuming. As a result, researchers 071 often revert to simpler evaluation strategies instead of comprehensive assessments. 072

A good benchmark should alleviate these burdens by automating dataset handling, integrating task 073 adapters, optimizers, schedulers, and logging mechanisms seamlessly. It should provide broad and 074 meaningful signals with minimal GPU time, accommodating various computational budgets, ensuring 075 inclusivity. Furthermore, an increasingly important factor for a robust modern benchmark engine 076 is its support for multi-modal learning and early fusion techniques. AI systems must seamlessly 077 integrate and reason across multiple modalities, such as text, images, audio, and more. Multi-modal learning enhances self-supervised learning opportunities and provides inherent supervision through 079 natural alignments, like audio-visual synchronization in videos. Early fusion, where data from 080 different modalities is combined at the initial stages of processing, allows models to leverage shared 081 representations, improving generalization and reasoning capabilities across varied tasks and domains. These key desiderata are what motivates the production of this work.

With the desiderata in mind, we next introduce EEVEE, a methodology developed for building high-signal low-cost evaluation routines, and GATE, the resulting benchmark that is designed to be extensible, readable, flexible, modular and robust, supported by a new efficient, easy to use framework.

EEVEE, Learning Optimal Benchmarks: The ability to find which benchmarks offer the most signal with respect to a given goal, such that we can optimize our compute time, research iteration speed, and engineering time is increasingly crucial. In this work, rather than just manually designing a new set of benchmarks, we propose a methodology, called *EEVEE (Empirical Evaluation process Evolution Engine)* that frames evaluation design as a learning problem and then leverages machine learning to automate the discovery and refinement of evaluation processes.

More specifically, EEVEE operates by taking in a large set of performance metrics from diverse models applied across various benchmarks and identifies a smaller subset of benchmarks with high predictive power over the entire set. EEVEE achieves this through two main components: (a) an evolutionary algorithm to optimize the selection of benchmark combinations based on a computed score, and (b) a meta-model trained to predict a model's performance on the full set of benchmarks using performance metrics from a chosen subset. We parameterize the meta-model as as a small neural network.

The meta-model receives input performance metrics from a subset of benchmarks and predicts performance on the full set of performance metrics. Through careful *k*-fold cross-validation and leveraging a diverse set of models and benchmarks, EEVEE iteratively evolves benchmark combinations that offer high information content with respect to the entire spectrum of benchmarks, ensuring robust, efficient and comprehensive evaluation that can be targeted to computational budgets ranging from more "GPU Poor" users to high-budget organizations.

Taking the desiderata explained above and the resulting understanding of what a good evaluation engine should look like, we demonstrate the effectiveness of EEVEE by tasking it with the discovery 108 of benchmark combinations that offer good signal-to-GPU-time ratio, for the evaluation of model 109 encoders – also referred to as backbones, on their ability to adapt to new tasks, domains, and 110 modalities. For this purpose, we choose a pool of 20 models, varying in their pretraining schemes 111 (e.g CLIP, DINO, ImageNet Classification), architectures (e.g. ResNets, ViTs, ConvNext) and even 112 their source modalities (e.g. Whisper, BERT), which we adapt on 31 benchmarks ranging from image classification, segmentation, relational reasoning, zero-shot image-to-text tasks, medical classification 113 and segmentation, video classification, and regression, using robust fine tuning recipes, and training 114 for 10K iterations, ensuring that the signal we get is about models that are adaptable, generalizable 115 and efficient in their adaptation. 116

117 By applying 20 models on 31 benchmarks and employing EEVEE on their resulting metrics, we 118 identify three subsets of benchmarks, each targeted to a specific computational budget range. Some of the key benchmarks that have been selected include iWildCam, CLEVR-Math, ACDC, WinoGround, 119 mini-ImageNet, Fungi, ADE20K, and dtextures. We refer to the discovered subsets as Tiers, and 120 assign to them identifiers for their sizes, specifically, small (n=8, 12 GPU hours), base (n=15, 24 GPU 121 hours) and big (n=21, 36 GPU hours). We package these tiers into our comprehensive benchmarking 122 suite and software framework (called GATE) designed for domain, task and modality transferability 123 evaluation, which facilitates the transfer of neural network encoders to different modalities, domains, 124 and tasks. GATE's architecture caters to the research community, enabling straightforward replace-125 ment of these transferable encoders with minimal effort. With these innovations, GATE seeks to 126 evolve the landscape of model encoder evaluation, championing a deeper understanding of transfer 127 learning and model adaptability.

128 Contributions: 1. We introduce EEVEE, a machine learning approach for selecting subsets of 129 benchmarks optimized to offer maximal predictive power over a larger benchmark set. 2. We conduct 130 a comprehensive investigation of diverse benchmarks within the space of image, image+text and 131 video modalities, pinpointing those with the highest predictive value for a model's performance 132 in downstream tasks. We apply EEVEE to model encoder evaluation by training 20 models on 31 133 benchmarks, identifying subsets of 8, 15 and 21 benchmarks that offer high signal-to-GPU-hour 134 ratios. 3. We pack the EEVEE-discovered subsets (of 8, 15 and 21 benchmarks out of 31 benchmarks) 135 into targeted benchmark packs, referred to as tiers, designed for specific compute budgets (of 12, 24 and 36 GPU hours) and project phases, and establish standard experimental settings for these 136 tiers. We call these collectively as the GATE Benchmarks. 4. We develop the GATE engine, a 137 unified benchmark suite and software framework that automates dataset downloading, preprocessing, 138 and pipelining for fine tuning and evaluation. GATE facilitates the incorporation of new model 139 encoders, adapts input modalities, fine-tunes with robust recipes, and logs critical information such 140 as training and evaluation metrics, power, energy, computational usage, task visualizations, and 141 model gradients per layer. 5. Through our extensive investigation, we identify foundation models 142 demonstrating superior transferability across diverse tasks. 6. We run a range of modality-shifting 143 transfer experiments in the standard evaluation process for ML researchers, so that future work can 144 potentially further probe into how pretraining on one set of modalities transfers to downstream (and 145 potentially different) modalities.

146 147

148

2 RELATED WORK

149 **On the Diversity of Benchmarks:** There is a vast array of benchmark suites in machine learning. 150 To the best of our knowledge, the benchmark suites relating strongly to GATE are ImageNet (9), 151 VTAB (74), VLMBench (78) and WILDS (27). ImageNet has been of tremendous importance and 152 interest to the transfer learning community. Nevertheless, there has been skepticism about overfitting 153 to such datasets resulting from implicitly qualifying models using the test set performance over 154 the years (49; 6) or the test set not being challenging enough to gauge model generalization power 155 (50). Although ImageNet pre-training helps transfer performance to the many-shot classification 156 setting (13), it provides minimal to no gains on more challenging datasets such as fine-grained 157 classification (28). Similarly, with a larger distribution shift, ImageNet pre-trained models was 158 found to offer limited benefits for medical imaging tasks due to large distribution shifts induced by 159 fundamental differences in data sizes, features, and task specifications; that is, lightweight models perform comparably to standard architectures (47). To make matters worse, ImageNet performance 160 is less correlated with and less predictive of downstream performance on diverse tasks beyond 161 classification such as object detection, few-shot classification, and segmentation (13). On top of it all,

162	Desiderata \downarrow Benchmark \rightarrow	ImageNet	VTAB	VLMBench	WILDS	GATE (Ours)
163	Diversity of Tasks Diversity of Domains	<i></i>	<i></i>			
	Diversity of Modalities	¥	~~	~~	~~~	<i>````</i>
164	Automatic Dataset Download/Preparation	\checkmark	\checkmark	\checkmark	\checkmark	$\checkmark\checkmark\checkmark$
165	Code allows for easy switch of encoders	\checkmark	\checkmark	\checkmark	\checkmark	$\checkmark\checkmark\checkmark$
105	Optimized for fast and effective research iteration	×	\checkmark	\checkmark	×	
166	Run Time	~~		~~~	\checkmark	
	Includes Medical Domains	×	1	×.	××	
167	Includes Environmental domains	\checkmark	\checkmark	×		
168	Tiered compute budgets GPU poor optimized	~~	~~	~~	~~	

172

Table 1: Our Desiderata (first column) VS Benchmarks (first row). More ticks means better, from one/red/lacking, two/gray/OK, three/green/good

when ImageNet is extended with a perturbed temporal dimension, models performance significantlyworsen (55).

175 On the Usability of Benchmarks: Beyond ImageNet, VTAB introduced a benchmark with a wider 176 diversity of tasks and domains (74). Nevertheless, it does not offer task and domain shifts offered 177 in GATE, such as medical segmentation and video classification and regression that are known to 178 be ill-measured and gauged by ImageNet alone (47; 55). That said, VTAB offers satellite imaging 179 and 3D tasks which GATE does not. Nevertheless, GATE as a software framework was optimized to minimise usage friction, to take no more than 12 GPU hours on our smallest tier, and, to only require 180 approximately 1 hour of adding the new encoder and wrapping it into GATE wrappers for GATE to be 181 able to go away and take care of everything, including dataset downloading, task adapter integration 182 and full train/val and test cycles with logging of various key metrics. VTAB, in our experience, 183 requires a lot more manual work in getting the datasets, and integrating new models to be adapted. 184 Similarly, VLMBench (78) and WILDS (27) offer more diverse datasets beyond previous work but 185 neither offer a tiered approach that enables iterative development of models during pre-training, nor produce extensible and flexible benchmarks that can be easily glued into researchers experimentation 187 code without friction. 188

On the Systematic Selection of Benchmarks: Previous work investigated the properties inherit 189 in multi-task benchmarks that trade-off diversity and sensitivity where the latter is how robust a 190 benchmark ranking is to the inclusion of irrelevant models or minute changes in the tasks themselves 191 (76). It was found that multi-task benchmark are unstable to irrelevant changes in tasks design. 192 Nevertheless, this is related to how the benchmark ranks models; whether it compares how model often 193 ranks higher than another in cardinal benchmarks or if the performance across tasks is averaged to 194 produce a single rank in cardinal ones. Meanwhile, our benchmark produces fine-grained information 195 to model performances across diverse tasks rather than producing specific ranking which is delegated 196 to the user analysis. Another complementary thread of work investigates dynamic benchmarks where model training and data collection is interleaved to continually challenge model knowledge (56). To 197 the best of our knowledge, this is the first work that studies the selection of multi-task, multi-domain 198 benchmarks that satisfy limited compute budgets while maximizing research signal. 199

In summary, Table 1 shows the desiderata that we believe a good evaluation suite and framework
 should have such that they can both offer the community useful signal, and also balance that with
 being practical so that people can adopt it.

203 204

205

3 EEVEE METHODOLOGY

206 EEVEE is our proposed method for automating the selection of Pareto-optimal benchmark subsets. 207 By analyzing benchmark performance metrics, EEVEE identifies a small, highly informative subset 208 that maximizes information relative to the entire benchmark pool. This ensures that, as machine 209 learning benchmark breadth and depth increases, we will always be able to identify and select few that 210 offer high information about the whole. We strike a balance between providing rich evaluation signals 211 and maintaining simplicity, reducing computational costs and human efforts required for adopting 212 new benchmarks. EEVEE enables the production of a tiered evaluation engine accommodating 213 various computational budgets, fostering an inclusive and accessible research environment, and improving the quality of insights derived from machine learning research while addressing reluctance 214 towards resource-intensive evaluation processes. This balance between efficiency, simplicity, and 215 signal richness presents EEVEE's value proposition for advancing machine learning research.

Working Principle of EEVEE: EEVEE works by building a *meta-model* over the performance metrics of models sufficient both in number and diversity, on the full benchmark pool from which we want to choose our subset. With the term *benchmark* in this paper we refer to a dataset + task pairs. The meta-model is parameterized as a 2-layer MLP network with 128 hidden neurons and leaky relu activation.

Formally, given a large benchmark pool $B = \{b_0, b_1, \dots, b_K\}$, where *B* is the full set of benchmarks, and b_i are individual benchmarks therein, we have a sufficiently large and diverse pool of model performance metrics $M = \{m_0^0, m_1^0, \dots, m_K^N\}$. Here, m_i^j is the performance metric of model *j* on benchmark b_i . We aim to discover a subset of *B* of size *k*. This means *k* total benchmarks make up the subset. If we build a meta-model $g(M_{selected}, \theta)$ to predict all of *M* given only the selected subset $M_{selected}$, it should minimize the following loss:

227 228

229 230

$$L_{EEVEE} = MSE(M, g(M_{selected}, \theta)) \tag{1}$$

In this equation, MSE is the mean squared error. M represents the full set of performance metrics of all our models on the full benchmark pool B. The term $g(M_{selected}, \theta)$ represents the predictions of the meta-model g with parameters θ when it is given the performance metrics of all models from the selected subset of benchmarks $B_{selected}$, referred to as $M_{selected}$.

However, our main focus lies in the selected combination of performance metrics $M_{selected}$ that can generalize well on previously unseen models. To that end, we must split M into train, validation and test sets, each consisting of performance metrics acquired from different models (e.g. train \rightarrow ResNet50, ViT-Base, CLIP, and val \rightarrow ResNet50, DINO, DeIT), and explicitly optimize the inner loop test loss rather than the training loss, while we use the validation loss to select the best meta-model for test. Hence the loss we wish to minimize is:

244

 $L_{EEVEE}^{test} = MSE(M^{test}, g(M_{selected}^{test}, \theta))$ ⁽²⁾

245 We need a non-differentiable method for choosing the k benchmarks in $M_{selected}$, since brute 246 force becomes intractable very quickly, so we employ evolutionary methods to learn the k selected 247 benchmarks.

248 This results in a bi-level optimization, with an evolutionary method on the outer loop $e(B_{selected})$, 249 where e is the evolutionary method, and $B_{selected}$ are the benchmarks being selected – or indeed, the 250 genes being optimized, and a small meta-model parameterized as a neural network $q(\theta)$ that receives a train/val split from $B_{selected}$ and trains itself to do the task described in Equation 1, after which 251 process it is scored using the val set using the loss in Equation 2. Then, once a given candidate of 252 benchmarks $B_{selected}$ is scored, in this way, the outer loop performs a tournament selection where 253 only the top 50 candidates are preserved and mutated by removing one benchmark at random, and 254 adding another at random. Each winning candidate mutates into 10 children, and the parent is 255 also preserved in the gene pool, producing a gene pool with 550 candidates for every cycle. At 256 initialization, we sample 1000 random combinations. We have found that 1000 is a good starting 257 population that is both tractable to score and facilitates the necessary diversity that enables limited 258 variation in results across several runs, showcasing convergent behaviour. We include full pseudocode 259 showcasing all the details related to how we performed EEVEE for our experiments in Algorithm 1, 260 2 and 3 in Figure 1.

261 Balancing the different metrics: In a given set of tasks, domains and datasets there can be an 262 imbalance in terms of how many metrics each one has and what types of metrics. Some metrics 263 are higher-is-better, while others are lower-is-better. We follow a simple way to balance this out, 264 which is, for a given metric that is higher-is-better we simply apply standard normalisation, and for 265 those lower is better, we first reverse their polarity and apply standard normalization. Then, for a 266 given dataset with meny metrics, we take the mean of those metrics, and, for a task containing many datasets, we take the mean of the mean of the per-dataset metrics. Therefore optimizing for what can 267 be considered a per-task equally weighted reward. There are many other ways to do this, and those 268 can depend on the context and what one is trying to achieve, but we chose this general one, since our 269 context was such.

295

296

Require: Performance metrics M , Input metrics M_{selected} , Epochs	Algorithm 3 Evolution
$E = 20, \text{Hidden dimension } d_{\text{hidden}} = 100, \text{Learning rate } \alpha = 0.01, \text{Weight decay } \lambda = 0.01, \text{Optimizer type } \omega = "AdamW"$ Ensure: Evaluation score mean(scores) 1: Convert data to tensors $x = M_{\text{selected}}$ and $y = M$ 2: Normalize x and y 3: Initialize ShuffleSplit cross-validation kf 4: Initialize shuffleSplit cross-validation kf 4: Initialize empty list scores 5: for each train, val split in kf do 6: Divide x into x_{train} and x_{val} ; y into y_{train} and y_{val} 7: Build meta-model $g(\theta)$ with hidden dimension d_{hidden} 8: Train $g(\theta)$ on x_{train} and y_{train} for E epochs with learning rate α , weight decay λ , and optimizer ω 9: Predict $y_{\text{pred}} = g(x_{\text{val}}, \theta)$ 10: Compute mean squared error score = MSE($y_{\text{pred}}, y_{\text{val}}$) 11: Append score to scores 12: end for 13: return mean(scores)	Require: Performance metrics $M = \{m_1^1, m_1^2, \dots, m_K^N\}$, Benchmark set B , Combination size k , Number of winners W , Number of children per winner C , Number of generations G , Initial combinations size I , Training epochs E , Hilden dimension $d_{\text{hidden}} = 100$, Learning rate $\alpha = 0.01$, Weight decay $\lambda = 0.01$, Optimizer type $\omega = "\text{AdamW"}$ Ensure: Evolved benchmark combinations B_{winners} 1: Initialize initial combinations B_{initial} with I random samples from B of size k 2: Evaluate performance of B_{initial} using SCOR-ING $(M, B_{\text{initial}}, E, d_{\text{hidden}}, \alpha, \lambda, \omega)$ and store scores in S 3: Select top W combinations from S as B_{winners} 4: for generation $g = 1$ to G do 5: Initialize a new set of combinations B_{new} 6: for each combination $B_{\text{selected}} \in B_{\text{winners}}$ do 7: Add B_{selected} to B_{new} 8: for each child $c = 1$ to C do 9: Mutate B_{selected} using MUTATION(B_{selected}, B) to
Algorithm 2 Mutation	10: create a new combination $B'_{selected}$ Add $B'_{selected}$ to B_{new}
Require: $B_{\text{selected}} \subset B, B = \{b_1, b_2, \dots, b_K\}$ Ensure: New B'_{selected} 1: Select $b_{\text{remove}} \in B_{\text{selected}}$ 2: Select $b_{\text{add}} \in B$ 3: while $b_{\text{add}} \in B_{\text{selected}}$ do 4: Select another $b_{\text{add}} \in B$ 5: end while 6: Create B'_{selected} by replacing b_{remove} with b_{add} 7: return B_{selected}	 end for end for end for Evaluate performance of B_{new} using SCOR- ING(M, B_{new}, E, d_{hidden}, α, λ, ω) and store scores in S Select top W combinations from S as B_{winners} end for return B_{winners}

Figure 1: (a) EEVEE Scoring algorithm, Mutation algorithm, and (b) Evolution algorithm.

Architecture of Meta-Model: We attempted deeper and shallower MLPs and transformers with
 various activations and hidden sizes but the chosen network balances speed of training with general ization. We have run experiments using 1-5 layer MLPs and transformers, with varying activation
 functions and hidden sizes ranging from 8-256. We found that a 2 layer MLP with 128 hidden size
 and leaky relu activation function offered the best generalization performance as a 2 layer transformer,
 but was much cheaper to train. Therefore, we used the 2-layer MLP throughout.

303 Applying EEVEE on Model Encoder Generalization

304 Why Model Encoder Evaluation? A common practice across machine learning applications involves 305 augmenting general model encoders with task-oriented heads (13). The adaption of this paradigm 306 can be attributed to the computational efficiency associated with training model encoders, over 307 more expensive setups. Much of computer vision, as well as vision to text search and retrieval 308 happen using model encoders (15; 77). Similarly, various applications requiring translation from one 309 domain/modality/task to another require an encoder of some sort (34). Even the "decoder-only" LLM models that have demonstrated incredible capabilities in the last few years, internally can be seen as a 310 series of representation encoders, a series of refinement before they reach the decoding stage (63). 311

Multi-modal early fusion is another topic closely related with model encoders – as research in early fusion can be done most efficiently when trying to learn data encoders rather than a full encoderdecoder, or decoder-only models (35). World model research in multi-modal dimensions can also take place most efficiently within a model-encoder context. Recent works like I/VJEPA (2) for example have paved the way for self-supervised learning which functions using model encoders, and has been demonstrated to be more efficient and more generalizable than full pixel decoding variants.

The goal of focusing on Model Encoder Evaluation: By applying EEVEE to search for a paretooptimal set of benchmarks, *and* packaging it up in a unified framework that is built for the researcher in mind from the ground up, one which offers out of the box automated downloading, pipeline building, task adapters, and a very mature training and eval loop. Within this framework, we facilitate, all relevant logging information, including key training and eval metrics, rich gradient information, power and computational information, as well as visualizations where relevant. Finally, we support easy switching of model encoders, no matter what source modality they come from – our framework




Figure 2: GATE Framework Pipeline

dubbed *GATE* is a one stop shop for ones model representation research needs, both during research, debugging, as well as at the evaluation phase.

GATE comes in three tiers *small*, *base* and *big*-GATE. Each having 8, 15 and 21 benchmarks within it, and targetted towards 12/24 and 36 GPU hours on a A100 40GB. We hope that by making it very easy for the end user and offering such rich signal for machine learning research, many researchers will choose to use GATE, to enhance their research signal, whilst keeping the compute budgets relatively feasible.

Preparations: Choosing Models, Benchmarks and Adaptation Processes: EEVEE will vield 348 better results if the space of models, benchmarks and adaptation processes we use is diverse, but also 349 thorough in numbers. A. Adaptation Process We wanted GATE to cover multiple domains, tasks 350 and modalities when shifting from the source to the target setting. For that reason we decided that if 351 a model encoder has an input layer that does not fit the target modality, we simply remove that input 352 layer and replace it with a relevant ViT-like patchification (12) followed by a linear combination for 353 each patch. For tasks where we have text, we would tokenize the text using BPE (54), and for tasks 354 where we have video we would use the model encoder on each image, to acquire an image-level 355 vector representation, and then follow that up with a simple 4 layer transformer that receives a 356 sequence of image-vector tokens, to produce a video-level embedding, on top of which we apply the 357 task-specific head at hand. The task-adapters we used leaned on established methods, and where 358 possible we just used a transformer head, which includes segmentation, relational reasoning and video classification, with everything just using a linear head, full details available at H. After these 359 modifications, described in Figure 2, we use a fine tuning scheme – this decision was informed by 360 preliminary experiments on both full fine tuning and linear probe with a frozen backbone, in which 361 we found that there was a clear superiority of fine tuning over linear probing for the benchmarks we 362 chose in our pool. Full details of these preliminary experiments can be found in Appendix C.1. In our 363 preliminary experiments we were able to identify three recipes, one for ConvNet-style architectures, 364 one for ViT-style architectures and one for Hybrid architectures such as ConvNext and ResNext that worked well for all tasks, details in C.1. 366

B. Model Pool We wanted the space of models used to cover many important pretraining schemes, 367 architectures, and source modalities. The details of these choices are provided next: 1. Pre-368 training Task and Dataset Variation: With a consistent architecture, models were subjected 369 to various pretraining tasks and datasets. Model instances representing this category include 370 CLIPViT (46), ConvNextV2 (38), Siglip, FlexViT (7), LaionViT, ImageNet1K ViT 371 (11) with Random Augment, SAM-ViT, DiNoViT, EfficientFormerV2 (33) and DeiT3 372 (62). Further to these, we include models initialized from scratch, specifically, ViT, ResNet50 373 (19), FlexViT, EfficientNetV2 (60), and then fine-tuned on the GATE tasks. 2. Archi-374 tectural Variation: We explored models having the same pretraining dataset (ImageNet), but 375 differing in their architecture. This group encompassed a mix of standard CNN models such as EffNetV2, ResNet50, ResNext50 (71), ConvNextV2_Base (38) and transformer-based 376 models like EfficientFormer (33) and FlexViT (7). 3. Modality and Dataset Variation: 377 This axis comprised models trained on modalities other than vision such as Whisper, coming from




Figure 3: The EEVEE MSE Loss (k=1) shows "predictiveness over the whole," with lower values being better. Benchmarks like iWildcam, HappyWhale, and WinoGround test unique capabilities and may not predict all tasks, yet EEVEE often includes at least two of these in its top combinations along with a "natural-image representative" such as CIFAR100, ADE20K or Flickr30K.

an audio to text task and Bert (10), Bart (32) and Mpnet (58) coming from various text-based tasks. These models had their original input processing systems replaced by a Vision Transformer style embedding and were subsequently fine-tuned on the GATE tasks.

393 C. Benchmark Pool The benchmark pool, detailed in the Appendix, includes Image Classification 394 (ImageNet1k (9), CIFAR100 (29), Places365 (79), Food101 (39), HappyWhale (18)), Few Shot 395 Image Classification (Aircraft (40), Fungi (53), MiniImageNet (66), CUB200 (67), Describable 396 Features (73)), Zero Shot Text-Image Classification (Flickr30K (44), New Yorker Caption Context 397 (21), Winoground (61)), Visual Relational Reasoning (CLEVR (24), CLEVRMath (37)), Image 398 Semantic Segmentation (ADE20K (80), COCO10K (36), COCO164K (36), NYU-Depth-v2 (57), 399 PascalContext (41), Cityscapes (8)), Medical Image Classification (Chexpert (22), Diabetic Retinopa-400 thy (17), HAM10000 (64)), Medical Segmentation (ACDC (5)), Video Classification (HMDB51 (31), 401 UCF-101 (59), Kinetics400 (25)) and Video Regression (iWildcam (4)).

402
 403
 404
 404
 405
 405
 406
 407
 408
 408
 409
 409
 409
 409
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400
 400

D. Experimental Approach We wanted our research environment to reflect the end user, so we 406 can properly understand their needs, and to offer a *pragmatic* experimental setup of in-the-wild 407 researchers with little time to hyperparameter optimize, and which have to make decisions on small 408 amounts of preliminary experiments – someone choosing a model encoder off the shelf and adapting 409 it to downstream setting. For that reason, we kept any hyperparameter tuning, or human attention 410 when it came to specific models to a minimum. Instead, we relied on existing good recipes, and 411 did some preliminary experiments as explained in detail in C.1. Briefly, we discovered specific 412 adjustments for each architecture type: for Convolutional Architectures, we used AdamW with a 413 learning rate of 1e-3, and 6e-4 for segmentation tasks; for Vision Transformer Architectures, AdamW with a learning rate of 1e-5; and for Convolutional + Transformer Hybrid Architectures, AdamW 414 with a learning rate of 2e-5. A plateau learning rate scheduler was configured with parameters like 415 mode "min", factor 0.5, patience 1000, and threshold 1e-4, allowing models to effectively choose 416 their own schedules based on their learning progress. This adaptive scheduling facilitated "good 417 enough" learning rates and enhanced performance across different architectures. 418

4 Results


419 420 421

390

391

392


421 Single Benchmark Predictiveness: As demonstrated in Figure 3, using EEVEE we quantified the 422 predictive power of each benchmark **on its own**, when not in a combination with others. We have 423 found that ADE20K, Flickr30K, and the New York Caption Competition lead in their predictive 424 power, with few-shot tasks, and relational reasoning, being very close to the best in predictive power. 425 ImageNet1K sits squarely in the middle of the competition. Furthermore, some of the most "novel" benchmarks like iwildcam, happy whale, ACDC, NYU and Winoground are the least predictive tasks, 426 427 Winoground being magnitudes less predictive. We argue that this is mainly due to the tasks being "harder", and our models being less designed for those. The results in WinoGround were bearly better 428 than chance for example. However, when once we move to combinations of benchmarks, these 'less' 429 predictive benchmarks become key contributors to better predictive power, as they represent edge 430 cases, as can be seen in Figures 6g 7c, 7i in Appendix J, where these have the highest importance 431 when removed from a given set.



443 444

tier

tier



445 Figure 4: Degradation of predictive power when a 446 trained from scratch, for different GATE tiers.

447 Predictiveness of Discovered Combinations In 448 Figure 5, we can see how the top-50 performing 449 candidate combinations perform as we vary the number of benchmarks per combination from 450 1 to 26. We can see that there is a point of di-451 minishing returns around the k = 8 point, after 452 which there appears to be some "overfitting" oc-453 curing. We verified that the overfitting was a 454 result of having a small sample number of 20 455 models, to train, val and test our meta-models 456 with. We tried our level best to find the best 457 architecture and regularization schemes for our




Figure 5: Performance of Models build with Kbest datasets: We do a search over the space of all k for EEVEE and box plot the population summary statistics of the top 50 combination candidates.

458 meta-model, and this was the best we could do given available compute and (human) time. We 459 chose 8, 15, and 21 as the combination-threshold to make our packs out of as they satisfied the 460 computational budgets we set for ourselves, and they have very diverse and predictive tasks, as can be seen in Figures 6g 7c, 7i. For full details on all the discovered top-k combinations please 461 look at Appendix Section J.1. Best Models based on GATE: As can be seen in Table 2, or the 462 Appendix extended Table 3, the best overall models are CONVNextV2, SigLIP and CLIP in that 463 order, with SigLIP and CLIP often exchanging ranks between themselves. However, it is worth noting 464 that EfficientNetV2 demonstrated exceptional performance/compute across all tasks, and even 465 outperformed all models in many medical tasks. Finally, ConvNet based models, and particularly 466 ResNext50 seem to have done exceptionally well in the edge-case scenarios of ACDC, Happy Whale 467 Individual identification, and general medical tasks, which indicates perhaps some sort of learning 468 efficiency advantages related to their inductive biases. 469

Limitations: We empirically evaluate EEVEE on a relatively large pool of models and benchmarks, 470 however, with more models, and benchmarks it could yield much more general results. Especially 471 with benchmarks targetting the text and audio modalities, as well as potentially offline RL. 472

473 5 CONCLUSION 474

In this paper, we propose EEVEE, an evolutionary-method-based search algorithm that can discover 475 out of a large collection of benchmarks, the ones that can offer the most predictive value on the 476 original collection, for a given set of models. We apply EEVEE on the task of model-encoder 477 evaluation in the context of images, image-text, videos, and medical domains. As a result, we obtain 478 the GATE Benchmark, which consists of 3 tiers, each targeted to a particular GPU budget, from 12, 479 24 and 36 GPU hours, per model evaluation. We then introduce the GATE engine, which takes these 480 benchmarks, and offers a researcher-designed environment in which one can easily port their own 481 model encoder, and run the full GATE tiers, and automatically produce a variety of performance, 482 energy/power, hardware utilization metrics and task visualizations. We evaluated 20 representative models ranging from image, image-text, text and audio pretrained models, on the GATE tiers, and we 483 discovered that ConvNextV2 and SigLIP seem to lead the pack overall, with EfficientNetV2 being an 484 exceptional, efficient alternative for the medical domain and for *unique scenario* tasks, such as Happy 485 Whale, ACDC and WinoGround. Finally, ConvNet based models, and ResNext50 in particular, seem

486	$\textbf{Metric} \downarrow \textbf{Model} \rightarrow$	cvnxtv2	siglip	clip	flex	deit	laion	vit	dino	smvit	rnx50	effv2	r50a1	effrmr	seffv2	sflex	svit	whspr s	sr50a1	bert ba	art mpnet
487	Img Class CIFAR-100 Acc@1	84.2	74.6	76.9	75.1	66 7	75.1	66.6	55.7	50.3	69.3	67.3	34.3	15.6	37.6	10.3	78	11.0	15.9	14.5 9	.0 1.0
488	Food-101 Acc@1	92.9	91.6	93.3	89.1	87.3	91.4	86.5	84.8	75.7	86.1	86.4	69.4	61.6	36.5	24.5	25.8	17.0	16.3	18.7 11	l.6 <mark>8.5</mark>
489	HWhale Individual Acc@1 HWhale Species Acc@1	$75.6 \\ 99.8$					21.0 99.7			$3.6 \\ 95.4$	78.7 99.7	77.1 99.7	5.2 92.1	4.4 92.8	$33.2 \\ 96.5$	2.8 76.5		2.2 64.3	2.1 65.8	2.3 1 71.2 5	
490	ImageNet-1K Acc@1 ImageNet-1K Acc@5	85.3 96.8					$74.1 \\ 93.1$			75.5 90.8	$77.6 \\ 93.3$	73.5 91.4	72.5 90.5	$44.6 \\ 72.5$	$16.9 \\ 37.3$	3.2 10.1		2.2 7.7	$\frac{1.3}{4.7}$	1.5 0 5.2 3	
491	Places365 Acc@1	54.7	53.5	54.1	52.1	49.0	53.7	47.5	47.3	27.1	51.8	51.5	40.9	25.2	26.6		8.6	7.5	5.0	5.2 3	.0 2.2
492	Task Mean Few-Shot Img Class	84.2	75.6	75.0	11.5	/1.8	72.6	09.3	00.7	59.8	79.5	78.1	57.8	45.2	40.7	19.5	18.0	16.0	15.9	17.0 12	2.6 11.1
493	Aircraft Acc@1 CUBirds Acc@1	96.7 98.0								92.9 93.4	91.6 92.8	90.6 92.1	$\frac{86.2}{89.4}$	$78.2 \\ 86.3$	$59.2 \\ 52.5$	54.9 50.0		$55.1 \\ 44.4$		61.2 60 48.4 50	
	DTextures Acc@1	85.0	85.2	88.6	78.9	81.9	86.1	80.8	79.4	81.9	77.7	60.3	77.2	68.5	46.6	50.2	50.5	50.0	33.1	44.6 49	9.8 38.3
494	Fungi Acc@1 Mini-Imagenet Acc@1	85.8 97.0					85.2 90.8			77.7 92.9	74.1 94.1	$73.7 \\ 63.2$	$67.1 \\ 93.2$	$59.2 \\ 90.9$	$27.6 \\ 36.7$	$\frac{38.0}{45.9}$		$33.9 \\ 44.8$	34.2	32.9 33 39.7 37	7.3 36.8
495	Omniglot Acc@1 VGG Flowers Acc@1	98.6 99.7								$98.6 \\ 93.4$	$98.5 \\ 87.9$	98.7 91.3	$95.5 \\ 89.3$	$95.8 \\ 90.6$	$98.2 \\ 59.6$	$93.4 \\ 69.4$		82.9 63.0	$ 80.5 \\ 53.4 $	90.2 84 59.1 59	
496	Task Mean	94.4					93.1				88.1	81.4	85.4	81.4	54.3	57.4		53.4	45.6	53.7 53	
497	Img Seg ADE20K mIoU	46.8	47.1	44.0	43.7	37.8	43.4	33.2	33.3	25.9	18.2	14.2	11.7	9.8	1.5	0.5	0.4	0.6	0.4	0.4 0	.5 0.4
498	Cityscapes mIoU COCO-10K mIoU	62.3 26.9								$59.5 \\ 28.6$	$40.8 \\ 18.4$	$64.2 \\ 10.2$	40.2 5.7	$2.5 \\ 14.0$	$46.7 \\ 1.1$	22.8 0.9	23.5 0.8	17.1 0.4	$18.6 \\ 1.6$	2.7 2 0.1 1	.0 2.7 .3 0.1
499	COCO-164K mIoU	32.7	36.7	33.8	33.0	30.5	32.4	27.0	28.9	25.7	16.8	9.7	4.7	13.7	1.0	0.7	0.7	0.5	0.7	0.1 1	.1 0.1
500	NYU mIoU Pascal mIoU	$7.5 \\ 32.8$	$7.7 \\ 34.8$							$\frac{11.0}{24.0}$	$5.9 \\ 16.6$	$\frac{8.3}{11.7}$	$6.4 \\ 6.8$	$10.5 \\ 14.0$	$6.8 \\ 1.7$		$3.7 \\ 1.1$	2.9 1.4	$7.2 \\ 2.3$	5.4 5 1.0 1	
501	Task Mean Img Relational	34.8	39.3	37.4	36.2	34.8	35.2	31.3	32.8	29.1	19.5	19.7	12.6	10.8	9.8	4.9	5.0	3.8	5.1	1.6 1	.9 1.6
502	CLEVR Acc@1	52.5								51.6	50.1		49.3	45.2	39.3	46.1		46.4			2.5 41.2
503	CLEVR Colour CLEVR Count	$35.4 \\ 45.8$								$34.2 \\ 45.6$	$26.8 \\ 45.3$	$15.7 \\ 39.0$	$24.7 \\ 45.1$	$14.7 \\ 44.8$	$12.5 \\ 37.9$	25.7 45.1		$28.8 \\ 44.8$		13.2 13 44.7 44	3.0 13.2 4.7 43.0
504	CLEVR Material CLEVR Shape	$60.5 \\ 52.1$	60.6	60.5	60.0	60.5	60.6	61.4	61.3		$\frac{58.6}{50.2}$	$52.1 \\ 34.3$	$57.5 \\ 50.2$	$53.7 \\ 44.8$	$49.8 \\ 33.3$	$53.7 \\ 35.8$		$54.0 \\ 36.1$		49.8 50 34.6 33	
505	CLEVR Size	61.0	61.1	61.3	60.7	61.1	60.8	62.0	62.3	60.9	59.6	53.5	58.3	55.7	50.6	56.2	55.2	55.2	54.6	54.2 54	4.1 50.1
506	CLEVR Yes/No CLEVR-Math Acc@1	60.7 79.3									$59.8 \\ 55.6$	$53.3 \\ 44.0$	$59.9 \\ 56.0$	$59.6 \\ 56.6$	$51.4 \\ 30.2$	60.1 46.9		$59.5 \\ 46.2$		59.5 59.44.8 42	
507	Task Mean Medical Class	55.9	54.4	54.9	53.1	55.2	53.9	53.9	53.6	52.6	50.8	41.6	50.1	46.9	38.1	46.2	45.9	46.4	45.0	42.9 42	2.5 40.7
508	Chexpert APS Macro	61.6								61.5	59.8	60.2		55.2	48.0	33.9				36.9 33	
	Chexpert AUC Macro Chexpert BS Macro	$82.5 \\ 84.3$					82.5 84.6			82.8 87.0	81.1 86.3	81.9 84.8	$79.1 \\ 86.1$	$79.9 \\ 86.4$	74.7 84.6	64.7 82.9		$65.5 \\ 83.0$	67.0 83.1	67.6 65 83.1 <mark>82</mark>	
509	Diabetic APS Macro Diabetic AUC Macro	56.9 87.5					56.4 85.3			$45.2 \\ 81.2$	$55.6 \\ 85.6$	58.7 86.1	$35.5 \\ 76.0$	$36.6 \\ 79.0$	$20.6 \\ 53.4$	$21.6 \\ 55.7$		22.5 57.8		$22.4 \ 21 \\ 59.4 \ 55$	
510	Diabetic BS Macro	94.5	94.0	93.9	93.9	93.8	93.6	93.7	93.6	93.0	93.9	94.2	92.3	92.6	91.6	91.3	91.4	91.4	91.5	91.8 91	1.6 91.6
511	HAM10K APS Macro HAM10K AUC Macro	94.5 99.1								83.4 97.8	$87.9 \\ 97.9$	87.1 97.5	$43.7 \\ 89.3$	$46.9 \\ 90.1$	$\frac{38.8}{85.6}$	38.0 86.1		32.2 82.8	$48.5 \\ 91.0$	50.6 37 91.1 85	
512	HAM10K BS Macro Task Mean	98.4 84.4								97.2 81.0	$97.6 \\ 82.9$	97.2 83.1	$95.2 \\ 72.4$	$95.5 \\ 73.6$	$94.6 \\ 65.8$	$94.5 \\ 63.2$					4.4 94.2 3.1 62.0
513	Medical Seg																				
514	ACDC Dice Score Task Mean	63.1 63.1					48.0 48.0			$44.6 \\ 44.6$	$44.2 \\ 44.2$	$61.0 \\ 61.0$	$40.2 \\ 40.2$	$18.7 \\ 18.7$	46.0 46.0	$16.5 \\ 16.5$		32.2 32.2	28.7 28.7	23.2 20 23.2 20	
515	Img to Txt ZS Flickr30K Img2Txt	6.3	6.3	7.0	5.9	5.6	6.8	5.0	5.2	4.5	4.1	3.7	4.7	4.2	1.6	1.8	2.0	1.9	2.0	1.9 1	.8 1.6
516	Flickr30K Txt2Img	5.7	5.9	6.0	5.3	5.1	6.5	6.0	5.1	5.0	3.8	4.0	4.2	3.9	1.7	1.8	2.0	2.2	2.3	1.9 1	.7 1.6
517	NYCC Img2Txt NYCC Txt2Img	6.9 6.1	$6.6 \\ 5.9$	$\frac{6.9}{6.4}$	$5.8 \\ 5.5$	$6.5 \\ 6.0$	6.9 6.2	6.4 6.4	$\frac{6.0}{5.8}$	4.7 4.8	$\frac{4.9}{4.3}$	4.1 4.1	$\frac{4.6}{3.9}$	4.2 3.7	1.6 1.6	2.1 2.0		1.9 2.0	$2.1 \\ 2.4$	2.0 1 1.9 1	
518	Winoground Img2Txt Winoground Txt2Img	$51.0 \\ 50.0$					$50.3 \\ 55.5$			$\frac{53.8}{48.6}$	61.9 54.8	$50.0 \\ 50.0$	$48.9 \\ 49.6$	$47.3 \\ 52.4$	$43.9 \\ 52.8$	50.0 50.0		$50.0 \\ 51.8$		49.6 50 51.8 48	
519	Task Mean	21.0								20.2	22.3	19.3	49.0 19.3	19.3	17.2	18.0		18.3		18.2 17	
520	Video Class HMDB-51 Acc@1	52.5	40.7	40.6	32.2	39.3	24.9	27.4	32.8	33.1	5.6	11.5	1.8	2.1	3.8	8.3	7.9	6.1	5.4	6.4 7	.5 4.0
521	Kinetics Acc@1 UCF-101 Acc@1	48.8 84.4	44.2	51.4	43.7	40.3		33.2	36.4	25.8	2.7 19.7	1.0 11.1	0.2 2.8	0.3 0.8	0.4 2.1	2.0 15.2	1.6	1.0 6.6	0.5 8.7	0.3 0 6.5 7	.3 0.3
522	Task Mean	61.9								40.7 35.9	9.4	7.8	1.6	1.1	2.1	8.5		4.6	4.9	4.4 4	
523	Video Reg IWildCam MAE Score	55.2												13.9	29.6	41.3	39.3	36.3	40.3	27.5 38	8.7 29.2
524	Task Mean GATE	55.2											37.3			41.3					8.7 29.2
525	Full GATE Mean	69.0 76.6								$58.5 \\ 66.8$		54.4		42.8		37.5					1.9 <u>31.8</u>
526	Big GATE Mean Base GATE Mean	76.6 68.3	65.6	65.7	62.6	63.7	60.7	60.2	60.7	58.6	$\frac{66.7}{55.1}$	53.5	$\frac{58.5}{48.2}$	$53.1 \\ 42.8$	38.0	$43.8 \\ 36.5$	36.3	35.4	36.6	$34.8 \ 34$	0.8 37.1 4.8 30.4
520	Small GATE Mean Full GATE Rank	77.7 1.0	$74.9 \\ 3.0$				71.2 6.0			65.3 9.0	$65.7 \\ 10.0$	61.7 11.0		$49.3 \\ 13.0$		35.7 15.0					4.4 30.4).0 21.0
	Big GATE Rank	1.0	2.0	3.0	4.0	5.0	6.0	7.0	8.0	9.0	10.0	11.0	12.0	13.0	14.0	$15.0 \\ 16.0$	16.0	17.0	18.0	19.0 20).0 21.0
528	Base GATE Rank Small GATE Rank	1.0 1.0	$\frac{3.0}{2.0}$		$5.0 \\ 4.0$		$7.0 \\ 6.0$		$6.0 \\ 7.0$	$9.0 \\ 10.0$		$\begin{array}{c} 11.0\\ 11.0 \end{array}$	12.0 12.0	$13.0 \\ 13.0$	$14.0 \\ 14.0$						$\begin{array}{ccc} 0.0 & 21.0 \\ 0.0 & 21.0 \end{array}$
529																					

 Table 2: Summary of experiments: Black/Bold best model, Green second best, Blue third best, and red the worst performing model. Models prefixed with 's' refer to 'from scratch' trained models, rather than pretrained. For the full table look at Appendix Table 3

to have a lot more *learning* efficiency, as they are the best adapted models on very novel domains, such as Happy Whale individual prediction challenge, ACDC and medical tasks.

540 REFERENCES

542

543

544

546

547

548

549

550

551

552 553

554

556

558

559

560 561

562

563

565

566

567 568

569

570

571 572

573

574

575 576

577

578 579

580

581

582

583 584

585

586

587 588

589

590

591

- [1] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings* of the IEEE International Conference on Computer Vision (ICCV), pages 2425–2433, 2015.
- [2] Adrien Bardes, Quentin Garrido, Jean Ponce, Xinlei Chen, Michael Rabbat, Yann LeCun, Mahmoud Assran, and Nicolas Ballas. Revisiting feature prediction for learning visual representations from video, 2024.
- [3] Loic Barrault, Ondrej Bojar, Marta R Costa-jussa, Christian Federmann, Mark Fishel, Yvette Graham, Barry Haddow, Matthias Huck, Philipp Koehn, Shervin Malmasi, Christof Monz, et al. Findings of the 2019 conference on machine translation (wmt19). In *Proceedings of the Fourth Conference on Machine Translation (Volume 2: Shared Task Papers, Day 1)*, pages 1–61, 2019.
- [4] Sara Beery, Grant Van Horn, and Pietro Perona. The iwildcam 2018 challenge dataset. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, pages 54–60, 2018.
- [5] Olivier Bernard, Alain Lalande, Caio Zotti, Florence Cervenansky, Xin Yang, Pheng-Ann Heng, Ismail Cetin, Karim Lekadir, Oscar Camara, Miguel A Gonzalez Ballester, et al. Deep learning techniques for automatic mri cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging*, 37(11):2514–2525, 2018.
 - [6] Lucas Beyer, Olivier J. Hénaff, Alexander Kolesnikov, Xiaohua Zhai, and Aäron van den Oord. Are we done with imagenet?, 2020.
 - [7] Lucas Beyer, Pavel Izmailov, Alexander Kolesnikov, Mathilde Caron, Simon Kornblith, Xiaohua Zhai, Matthias Minderer, Michael Tschannen, Ibrahim M. Alabdulmohsin, and Filip Pavetic. Flexivit: One model for all patch sizes. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 14496–14506, 2022.
- [8] Marius Cordts, Mohamed Omran, Sebastian Ramos, Tobias Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3213–3223, 2016.
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Li Kai, and Li Fei-Fei. Imagenet: A largescale hierarchical image database. In 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009.
- [10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805, 2018.
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929, 2020.
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations (ICLR)*, 2021.
- [13] Linus Ericsson, Henry Gouk, and Timothy M. Hospedales. How well do self-supervised models transfer? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5414–5423, June 2021.
- [14] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.

600

601

602

603

604

606

607

608

609

619

620

621

622

623

624

625 626

627

628

629

630

631 632

633

634

635 636

637

638

639

640

641 642

643

644

645

- 594 [15] Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, 595 Meng Wang, and Haofen Wang. Retrieval-augmented generation for large language models: A 596 survey, 2024. 597
 - [16] John S Garofolo, Lori F Lamel, William M Fisher, Jonathan G Fiscus, David S Pallett, and Nancy L Dahlgren. Timit acoustic-phonetic continuous speech corpus ldc93s1, 1993.
 - [17] Varun Gulshan, Lily Peng, Marc Coram, Michael C Stumpe, Derek Wu, Arunachalam Narayanaswamy, Subhashini Venugopalan, Kasumi Widner, Travis Madams, Jorge Cuadros, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. JAMA, 316(22):2402-2410, 2016.
- [18] Happywhale. Happywhale whale and dolphin identification challenge, 2022. 605
 - [19] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 770–778. IEEE, 2016.
- [20] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Mantas He, Dawn 610 Song, and Jacob Steinhardt. Measuring massive multitask language understanding. arXiv 611 preprint arXiv:2009.03300, 2020. 612
- 613 [21] Jack Hessel. New yorker caption contest corpus, 2023. 614
- [22] Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, 615 Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large 616 chest radiograph dataset with uncertainty labels and expert comparison. Proceedings of the 617 AAAI Conference on Artificial Intelligence, 33:590–597, 2019. 618
 - [23] Niehues Jan et al. Iwslt 2017: Proceedings of the 14th international workshop on spoken language translation. 2017.
 - [24] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In Proceedings of the IEEE Conference on Computer Vision and Pattern *Recognition*, pages 2901–2910, 2017.
 - [25] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. In arXiv preprint arXiv:1705.06950, 2017.
 - [26] Philipp Koehn. Europarl: A parallel corpus for statistical machine translation. MT summit, 5:79-86, 2005.
 - [27] Pang Wei Koh, Shiori Sagawa, Henrik Marklund, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Phillips, Irena Gao, et al. Wilds: A benchmark of in-the-wild distribution shifts. In International Conference on Machine Learning, 2021.
 - [28] Simon Kornblith, Jonathon Shlens, and Quoc V. Le. Do better imagenet models transfer better? In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), June 2019.
 - [29] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical Report UTML TR 2009, University of Toronto, Toronto, Ontario, Canada, 2009.
 - [30] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25:1097-1105, 2012.
- [31] Hildegard Kuehne, Hueihan Jhuang, Estibaliz Garrote, Tomaso Poggio, and Thomas Serre. 646 Hmdb: A large video database for human motion recognition. In 2011 International Conference on Computer Vision (ICCV), pages 2556-2563. IEEE, 2011.

653

654

655 656

657

658 659

660

661 662

663

664

665 666

667

668

669

670 671

672

673 674

675

676 677

678

679

680

681

682

683

684 685

686

687

688

689

690

691

692

693 694

695

696

697

- [32] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2020.
 - [33] Xiuyu Li, Yutong Yuan, Shu Chen, Martin Danelljan, Radu Timofte, and Luc Van Gool. Efficientformer: Vision transformers at mobilenet speed. *arXiv preprint arXiv:2206.01191*, 2022.
 - [34] Long Lian, Boyi Li, Adam Yala, and Trevor Darrell. Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models, 2024.
 - [35] Paul Pu Liang, Amir Zadeh, and Louis-Philippe Morency. Foundations & trends in multimodal machine learning: Principles, challenges, and open questions. *ACM Computing Surveys*, 56(10):1–42, 2024.
 - [36] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision*, pages 740–755, 2014.
 - [37] Adam Dahlgren Lindström and Savitha Sam Abraham. Clevr-math: A dataset for compositional language, visual and mathematical reasoning. *ArXiv*, abs/2208.05358, 2022.
 - [38] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *arXiv preprint arXiv:2201.03545*, 2022.
 - [39] Matthieu Guillaumin Lukas Bossard and Luc Van Gool. Food-101 mining discriminative components with random forests. In *European Conference on Computer Vision*, 2014.
 - [40] Subhransu Maji, Esa Rahtu, Juho Kannala, Matthew Blaschko, and Andrea Vedaldi. Finegrained visual classification of aircraft. In 2013 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 554–561. IEEE, 2013.
 - [41] Roozbeh Mottaghi, Xiaobai Chen, Xiaofeng Liu, Nam-Gyu Cho, Seong-Whan Lee, Sanja Fidler, Raquel Urtasun, and Alan Yuille. The role of context for object detection and semantic segmentation in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 891–898, 2014.
 - [42] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. Librispeech: An asr corpus based on public domain audio books. In 2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 5206–5210, 2015.
 - [43] Denis Paperno, Germán Kruszewski, Angeliki Lazaridou, Ngoc Quan Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernández. The lambada dataset: Word prediction requiring a broad discourse context. In *Proceedings of the 54th Annual Meeting* of the Association for Computational Linguistics (Volume 1: Long Papers), pages 1525–1534, 2016.
 - [44] Bryan Plummer, Liwei Wang, Chris Cervantes, Juan Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for grounded image descriptions. arXiv preprint arXiv:1505.04870, 2015.
 - [45] Bryan A Plummer, Liwei Wang, Christopher M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2641–2649, 2015.
- [46] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual
 models from natural language supervision. In *International conference on machine learning*,
 pages 8748–8763. PMLR, 2021.

707

708

709

716

717

718

722

723

724 725

726

727 728

729

730

736

737

738 739

740 741

742

743

744

745

746 747

748

749

750

751

752

- [47] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
 - [48] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, 2016.
- [49] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do cifar-10 classifiers generalize to cifar-10?, 2018.
- [50] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do ImageNet classifiers generalize to ImageNet? In *Proceedings of the 36th International Conference on Machine Learning*, Proceedings of Machine Learning Research. PMLR, 2019.
 - [51] Erik F Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Languageindependent named entity recognition. In *Proceedings of the seventh conference on Natural language learning at HLT-NAACL 2003*, pages 142–147, 2003.
- [52] Adam Santoro, David Raposo, David G.T. Barrett, Mateusz Malinowski, Razvan Pascanu, Peter
 Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. In
 Advances in Neural Information Processing Systems (NeurIPS), pages 4967–4976, 2017.
 - [53] Dirk Schroeder, Yin Cui, Yang Chai, Daniel Kristensen, Evangelos Kalogerakis, and Serge Belongie. The fgvcx fungi classification challenge. In CVPR Workshops, 2018.
 - [54] Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, 2016.
 - [55] Vaishaal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. Do image classifiers generalize across time? In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9661–9669, October 2021.
- [56] Ali Shirali, Rediet Abebe, and Moritz Hardt. A theory of dynamic benchmarks, 2023.
- [57] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgbd images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 746–760, 2012.
 - [58] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. Mpnet: Masked and permuted pre-training for language understanding. arXiv preprint arXiv:2004.09297, 2020.
 - [59] Khurram Soomro, Amir Roshan Zamir, and Mubarak Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. In *arXiv preprint arXiv:1212.0402*, 2012.
 - [60] Mingxing Tan and Quoc V Le. Efficientnetv2: Smaller models and faster training. *arXiv* preprint arXiv:2104.00298, 2021.
 - [61] Tristan Thrush, Hongyu Jiang, Goutham Prasad, and Jacob Andreas. Winoground: Probing vision and language models for visio-linguistic compositionality. *arXiv preprint arXiv:2204.03162*, 2022.
 - [62] Hugo Touvron, Piotr Bojanowski, Mathilde Caron, Matthieu Cord, Alaaeldin El-Nouby, Edouard Grave, Armand Joulin, Gabriel Synnaeve, and Jakob Verbeek. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022.
 - [63] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. arXiv preprint arXiv:2307.09288, 2023.
- [64] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The ham10000 dataset, a large collection of multi-source dermatoscopic images of common pigmented skin lesions. *Scientific Data*, 5:180161, 2018.

756 [65] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, 757 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In Advances in Neural Informa-758 tion Processing Systems (NeurIPS), pages 5998-6008, 2017. 759 [66] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Koray Kavukcuoglu, and Daan Wierstra. 760 Matching networks for one shot learning. In Advances in Neural Information Processing 761 Systems (NeurIPS), pages 3630–3638, 2016. 762 763 [67] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge Belongie. The caltech-764 ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of 765 Technology, 2011. 766 767 [68] Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. Superglue: A stickier benchmark for general-purpose 768 language understanding systems. In Proceedings of the 33rd International Conference on 769 Neural Information Processing Systems, pages 3266–3280, 2019. 770 771 [69] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 772 Glue: A multi-task benchmark and analysis platform for natural language understanding. In 773 Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural 774 Networks for NLP, pages 353-355, 2018. 775 [70] Jason Weston, Antoine Bordes, Sumit Chopra, Alexander M Rush, Bart van Merrienboer, 776 Armand Joulin, and Tomas Mikolov. Towards ai-complete question answering: A set of 777 prerequisite toy tasks. arXiv preprint arXiv:1502.05698, 2015. 778 779 [71] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual 780 transformations for deep neural networks. In Proceedings of the IEEE Conference on Computer 781 Vision and Pattern Recognition (CVPR), pages 1492–1500. IEEE, 2017. 782 783 [72] Junichi Yamagishi, Christophe Veaux, and Kirsten MacDonald. Cstr vctk corpus: English 784 multi-speaker corpus for cstr voice cloning toolkit (version 0.92). University of Edinburgh. The Centre for Speech Technology Research (CSTR), 2019. 785 786 [73] Fisher Yu, Yinda Zhang, Shuran Song, Ari Seff, and Jianxiong Xiao. Fine-grained visual 787 comparisons with local learning. In 2014 IEEE Conference on Computer Vision and Pattern 788 Recognition (CVPR), pages 192–199. IEEE, 2014. 789 790 [74] Xiaohua Zhai, Joan Puigcerver, Alexander Kolesnikov, Pierre Ruyssen, Carlos Riquelme, Mario 791 Lucic, Josip Djolonga, Andre Susano Pinto, Maxim Neumann, Alexey Dosovitskiy, et al. A large-scale study of representation learning with the visual task adaptation benchmark. In 792 793 International Conference on Learning Representations, 2020. 794 [75] Chi Zhang, Feng Gao, Baoxiong Jia, Song-Chun Zhu, and Yixin Zhu. Raven: A dataset for 795 relational and analogical visual reasoning. In Proceedings of the IEEE Conference on Computer 796 Vision and Pattern Recognition, pages 5317–5327, 2019. 797 798 [76] Guanhua Zhang and Moritz Hardt. Inherent trade-offs between diversity and stability in 799 multi-task benchmarks, 2024. 800 [77] Ruochen Zhao, Hailin Chen, Weishi Wang, Fangkai Jiao, Xuan Long Do, Chengwei Qin, 801 Bosheng Ding, Xiaobao Guo, Minzhi Li, Xingxuan Li, and Shafiq Joty. Retrieving multimodal 802 information for augmented generation: A survey, 2023. 803 804 [78] Kaizhi Zheng, Xiaotong Chen, Odest Chadwicke Jenkins, and Xin Eric Wang. VLMbench: A 805 compositional benchmark for vision-and-language manipulation. In Thirty-sixth Conference on 806 Neural Information Processing Systems (NeurIPS) Datasets and Benchmarks, 2022. 807 [79] Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 808 million image database for scene recognition. In IEEE Transactions on Pattern Analysis and 809 Machine Intelligence, volume 40, pages 1452–1464. IEEE, 2017.

[80] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 5122–5130, 2017. [81] Yuke Zhu, Olaf Groth, Michael S Bernstein, and Li Fei-Fei. Visual7w: Grounded question answering in images. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 4995-5004, 2016.

864 A END-USER GUIDELINES

866 867	For an e	nd-user to use GATE, they need to:
868	1.	Install the GATE framework python package, as described in the Github repo's readme page.
869	2	Choose a path for implementing the new foundation model encoder they wish to evaluate.
870	2.	This is either cloning the full GATE repo and modifying existing components directly,
871		or, importing the GATEncoder and GATEModel classes from GATE, and wrapping up
872		their model within it. Doing so requires the researcher to implement a relevant forward
873		function that can take in the modalities their model needs to process, as well as defining a
874		configuration that tells GATE what modalities a model can receive and output features on,
875 876		as well as any transforms needed for a batch to be ready for their model.
877	3.	The user chooses a GATE tier to use (from smallGATE, baseGATE and bigGATE).
878		Based on the configuration defined by the user in step 2.
879	4.	GATE generates a list of commands, each representing an experiment that needs to be run,
880		and can then run these commands on your local GPU box, parallelizing the tasks, one on
881		each available GPU, or, can provide a list of commands or json file that one can use to run
882		these commands on a GPU cluster, or other hardware.
883	5.	GATE emits a wandb project, with metrics, visualizations and other measures, allowing easy
884		tracking of experiments, and sharing thereof, as well as huggingface model weights for each
885		model being trained – which is also used to achieve a <i>stateless</i> execution.
886	6.	Once the experiments are completed, one can invoke the produce-analysis.py file
887		within GATE to get tables and figures that analyse the data, similar to what appears in
888		this paper. Those results can then be used to report results in a paper, or, be used to make
889		decisions for production models.
890	T 1.	
891		becess ensures the GATE framework is aware of what a model's supported modalities are, as now to produce modality-specific features, given the model. Once this is completed, the user,
892		ngle line of code, can select a GATE tier, and launch all jobs needed to produce results for that
893		portantly, GATE is made to facilitate and encourage foundation models that are diverse in their
894		ties, and allow the researchers to focus on what matters – that is, designing and training their
895		on model – rather than spending the majority of their time building and optimizing evaluation
896		ate. Furthermore, the diversity of signal that GATE provides allows better understanding of a
897		odel's strengths and weaknesses, which as a result makes the research, review and iteration
898	process	of the field as a whole more efficient. This is because there is a consistent boilerplate that

902 A.1 PRINCIPAL USE CASES

the overly optimistic, or overly pessimistic side of things.

1. Model Development and Iteration: GATE serves as a valuable tool during the model research and development phase. By integrating the model into GATE and running either the smallGATE or baseGATE tiers, developers can obtain a comprehensive and robust performance evaluation of their model across diverse domains, tasks, and modalities. Worth noting that GATE allows easy inclusion of foundation models pretrained on images, video, audio, text, etc, to be fine-tuned on pixel-based tasks. It achieves this by replacing a model's root layer / embedding layer, with one appropriate for a given task's modality, and adding on top a relevant task adapter head.

runs all models, with broad signal that reduces probability of making erroneous conclusions - both in

911
92. Model Evaluation for Machine Learning Research: GATE enhances the communication of research findings and their potential applications, a vital aspect of scientific collaboration. By using GATE as a benchmark, even at the most cost-efficient GPU hour level of smallGATE, the clarity and depth of future ML papers can be significantly improved. GATE's explicit evaluation of modality, domain, and task shifts in a given foundation model provides a nuanced and informative perspective on a model's true capabilities, offering a more detailed understanding of a model's strengths and weaknesses than optimizing a single metric, such as ImageNet validation error.

918 B RESULT EXTRAS

The results were logged in WandB, and then further processed after all experiments were completed to generate the tables and figures in this paper. Much of the logged information outside of testing metrics were not used for any of the figures and tables in this paper. The full set of experiments and all the logged results can be found at our wandb gate project repo².

B.1 RESULT PROCESSING

927 Once all experiments were completed, we queried our wandb project repository and returned test
928 results from all our experiments, if an experiment name was duplicated, we used the latest entries,
929 and, for each experiment type there existed three independent runs. We averaged the results of any
930 metrics across such independent runs to acquire a better approximation to the true performance of
931 those models.

932 933

924 925

926

C PRELIMINARY EXPERIMENTS DETAILS

934 935

936

C.1 PRELIMINARY EXPERIMENTS

First, we trained models on ImageNet1k, CIFAR100, CLEVR, ADE20K, CityScapes, and, ACDC
for 5K iterations, using cosine annealing learning schedule or plateau annealing, with AdamW,
weight decays varying from 0.1 - 0.0001, and applied models from each major architecture category –
specifically, the CLIPViT, ImageNet pretrained ViT, ResNext, ResNet and ConvNextV2. The results
from these experiments pointed to the fact that there exists one general and good recipe for each
architecture style. The recipes that we discovered were as follows:

943 C.1.1 ACROSS ARCHITECTURE SETTINGS

⁹⁴⁵ Unless otherwise stated, the settings here are applied universally in all experiments.

946
 947
 948
 948
 949
 949
 940
 941
 941
 942
 942
 943
 944
 944
 944
 944
 945
 946
 946
 947
 948
 948
 948
 948
 948
 948
 949
 948
 944
 944
 944
 944
 945
 946
 947
 948
 947
 948
 948
 948
 948
 948
 949
 949
 949
 949
 949
 940
 940
 941
 941
 942
 942
 944
 944
 944
 944
 945
 945
 946
 947
 946
 947
 947
 948
 948
 948
 949
 949
 949
 949
 949
 949
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940
 940

Training Details: Training iterations: 10K, validate every 500 iterations.

950 Test Details: Top-3 validation models (across all validated checkpoints) are ensembled by prediction averaging.
 952

953 C.1.2 ARCHITECTURE SPECIFIC SETTINGS 954

Convolutional Architectures: Optimizer: AdamW, learning rate 1e-3, and for segmentation tasks only, we used learning rate 6e-4

957 Vision Transformer Architectures: Optimizer: AdamW, learning rate 1e-5

958959959Convolutional + Transformer Hybrid Architectures Optimizer: AdamW, learning rate 2e-5

- 960 The above recipes were what we used throughout all our experiments unless otherwise stated.
- 961 962

963 964

965 966

967

968

969

970

971

D GATE GUIDING PRINCIPLES

The fundamental values driving the design decisions behind GATE are the following:

1. Maximizing Generalization Signal: GATE is designed to provide a high signal-to-noise ratio concerning a model's ability to generalize in diverse downstream contexts, that vary in domain, task and modality. This allows for a more robust assessment of a model's capacity for adaptation and versatility. By noise here we refer to how clear a given signal response is. For example, an image classification test accuracy signal on ImageNet, would provide clear

²omitted until double blind is over

signal with respect to the natural domain and the classification task, but would be blurry for more compositional, object disentanglement and relational tasks, such as segmentation, or, visual question answering.

- 2. Time Efficiency: Acknowledging the importance of computational resources and time, GATE operates within set benchmarks of 12, 24, and 36 GPU hours (established on A100 @ 40GB). These set timeframes ensure GATE's assessments are both thorough and expedient.
 - 3. Minimizing Usage Friction: The framework supporting GATE is designed to be user-friendly, enabling easy integration of new backbones and facilitating smooth experimentation. This low-friction approach ensures a streamlined experience when using GATE, making the process of evaluation more efficient.

We argue that a good balance of the above can generate a pragmatic, yet thorough foundation model evaluation suite, that will, importantly, be of real use to most researchers in the field.

984 985 986

987

972

973

974

975

976

977

978

979

980

981

982 983

E DEFINING THE GATE BENCHMARK

GATE is a comprehensive evaluation engine designed to advance the development of more general
 machine learning models. It improves on existing benchmarks by enabling the evaluation of models across diverse modalities, domains, and tasks.

GATE is composed of three key components. The first is a benchmark *pool*, a broad collection of 992 datasets, tasks, and processes that measure a model's performance across various domains, tasks, 993 and modalities. The second component is a set of benchmark *tiers*, which are meticulously curated 994 subsets from the GATE benchmark pool, tailored to specific compute budgets and project phases. 995 The final, and is a software framework, designed to seamlessly integrate new foundation models and 996 execute the GATE tiers, thereby enabling efficient performance evaluation across a diverse range of 997 downstream modalities, domains, and tasks. Practically, GATE is directed towards machine learning 998 researchers and developers as a means to efficiently, and with little friction, get broad signal about 999 how their model performs after transfer in diverse contexts, specifically selected for their empirically 1000 evaluated high signal-to-noise ratio with respect to predictive power in how a model performs in previously unseen contexts. 1001

Building GATE was a careful balancing act. We needed to respect specific time budgets while also aiming for a wide variety of evaluation scenarios. Our approach was as follows:

- 1. Select a diverse set of learning contexts, spanning multiple domains, tasks and modalities. We refer this as the *Benchmark Pool*.
- 2. Select a broad set of key foundation models, varying in their architecture, pretraining scheme and source modality. We refer to this as the *Model Pool*.
 - 3. Fine tune each of the models in the model pool, on each of the contexts in the benchmark pool. Evaluate trained models on each context's test sets.
- 4. Use the test set results acquired to quantify the predictive power each benchmark holds with respect to previously unseen benchmarks, both at the individual level and the collection level. We call this measure, the *downstream generalization predictability measure* (**DGPM**).
- 5. Use the DGPM values of the various combinations of benchmarks to build the three GATE tiers, selecting combinations of benchmarks that can provide the most information within a target time budget.
- We elaborate on each of the above steps in the following subsections.

1021 F BENCHMARK POOL SELECTION DETAILS

1022

1020

1005

1008

1009

1010

1011

1012

1013

1014

1015

1016

1017

Medical Image Classification: Medical data are known to present a substantial shift in both domain and even modality depending on their format. We have selected datasets that not only pose significant challenges for foundation models but also align with the broader imperative to deliver real-world benefits downstream.

1026
 1027
 1028
 1028
 1028
 1029
 1029
 1029
 1029
 1020
 1020
 1021
 1021
 1022
 1022
 1023
 1023
 1024
 1025
 1026
 1027
 1028
 1028
 1029
 1029
 1029
 1020
 1021
 1021
 1022
 1022
 1023
 1024
 1025
 1025
 1026
 1026
 1027
 1028
 1028
 1029
 1029
 1020
 1020
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 1021
 <li

Diabetic Retinopathy Classification: Early detection of diabetic retinopathy from retinal images
 is a public health priority; models fine-tuned on this dataset can have immediate implications for
 preventing vision loss on a global scale. This dataset requires models to decipher fine-grained,
 progressive changes indicative of the disease, reflecting the precision necessary for medical AI
 applications.

HAM10000 (Human Against Machine with 10000 dermatoscopic images): The dataset provides
 a diverse spectrum of skin lesion images vital for differentiating between benign and malignant
 conditions. Incorporating this dataset not only challenges the pattern recognition prowess of AI but
 also contributes to the advancement of dermatology through machine learning technologies.

Metrics: We collect Average Precision Score (APS), Area Under the Receiver Operating Characteristics Curve (AUC), and Brier Score (BS) both overall (i.e. macro) as well as for individual pathologies/classes.

Medical Segmentation: This category evaluates foundational models' ability to generalize from natural to medical image modalities and to perform domain-specific tasks that require precision and complex spatial understanding:

ACDC (Automated Cardiac Diagnosis Challenge): This dataset is aimed at assessing models' generalization to the medical domain, particularly the transferability of representations for segmenting anatomical structures in cardiac MRI images. By focusing on the heart's intricate anatomy, ACDC tests the models' ability to adapt to clinically relevant shapes and patterns—a shift from common visual recognition tasks to precise medical delineation. Metrics: We collect dice loss, mIoU, mean accuracy and overall accuracy.

1051

1052 1053 G BENCHMARK POOL DETAILS

1053

Having a set of diverse benchmarks ranging in challenge factor, as well as modality, task and domain
shift was key. We explain in more detail why why consider these factors important in Appendix in
more detail. We refer to this as our *benchmark pool*, and it consists of the following:

Image Classification: We employ ImageNet1k (9), CIFAR100 (29), Places365 (79), and Food101 (39) to cover diverse natural image domains. Additionally, we include HappyWhale (18) for a more challenging domain shift, aiding in wildlife research and providing an interesting test case for model evaluation.

Few Shot Image Classification: We use the MetaDataset task recipe on the Aircraft (40), Fungi (53), MiniImageNet (66), CUB200 (67), and Describable Features (73) datasets to evaluate task and domain shift robustness for an evaluation model.

Zero Shot Text-Image Classification: Another key setting is that of zero-shot text-image classification, on which many current key models were trained and evaluated (46). We utilize Flickr30K, New Yorker Caption Context (a challenging humor task), and Winoground–a task requiring the model to match two texts with their corresponding images, focusing on compositional differences.

Visual Relational Reasoning: A context where earlier models, such as ResNet50 (19) had low performance without layers with associative inductive biases (e.g., relational neural networks or transformers (52; 65)). This ensures we are aware of any trade-offs in relational compositional abilities in our models. We use CLEVR (24) and CLEVRMath (37).

Image Semantic Segmentation: Essential for various real-world applications, serving as an indicator of a model's ability to retain spatial information and identify objects at a per-pixel level. ADE20K (80), COCO10K (36), COCO164K (36), NYU-Depth-v2 (57), PascalContext (41), and Cityscapes (8).

Medical Image Classification: Medical data exhibit substantial domain and modality shifts, posing significant challenges for machine learning models while aligning with the imperative to deliver real-world benefits.*Chexpert* (22) (chest X-rays annotated for thoracic disease diagnosis), *Dia*-

betic Retinopathy Classification (17) (retinal images for early detection of diabetic retinopathy),
 HAM10000 (64) (dermatoscopic images for differentiating skin lesions).

Medical Segmentation \rightarrow ACDC (Automated Cardiac Diagnosis Challenge) (5): This dataset as-
sesses models' generalization to the medical domain, particularly the transferability of representations
for segmenting anatomical structures in cardiac MRI images. By focusing on the heart's intricate
anatomy, ACDC tests the models' ability to adapt to clinically relevant shapes and patterns.

Video Classification: Video classification tasks test models on their temporal generalization abilities and require an understanding of not only individual frame content but also the transition and context between frames. *HMDB51 (Human Motion Database)* (31), *UCF-101 (University of Central Florida - 101 action categories)* (59), *Kinetics400* (25).

Video Regression: Where classification tasks gauge categorical distinctions, video regression tasks assess models' ability to make continuous numerical predictions from temporal data, serving as an indicator of a model's capability to process and quantify dynamic content. *iWildcam (International Wildlife Camera Trap Challenge)* (4): This dataset targets estimating animal species abundance from videos and is a direct test of modality and task shift, and showcases a models' potential impact on ecological monitoring and species conservation efforts.

- 1. **Modality shifting** contexts: Contexts where the foundation model is asked to learn to do well at a task that requires understanding of a previously unseen modality. More specifically, assuming a foundation model has been trained on natural images, this would be transferring to medical imaging, video, audio and test contexts. This would shed light on the performance of a model's middle layers.
- 2. **Task shifting** contexts: Contexts where a model is tasked with performing a previously unseen task, for example, transferring from classification to segmentation or relational reasoning.
- 3. Domain shifting contexts: Contexts where a model is required to perform a task on a domain that is different from the one it was trained on. For example moving from natural images on ImageNet at 224x224 resolution to black and white Omniglot characters at 28x28 resolution, or, moving from ImageNet to images of fungi. More extreme domain shifts would be going from natural images to medical images for example.
- 1112 1113

1098

1099

1100

1101

1102 1103

1104

1105

1106

1114 H TASK ADAPTER DETAILS

1115 1116

Classification: For classification after a given encoder's output features we apply a linear layer as is standard.

Segmentation: We extract features after every stage within an encoder, i.e. before each pooling layer
 in conv-net architectures, and, after a transformer block in ViT-derivative encoders. We then upscale
 those to 64x64 before we concatenate and feed to a transformer decoder.

Relational Reasoning: We process images with our designated encoder, and text with a CLIP text
encoder. We then concatenate the features, and feed into a transformer that considers receives as
tokens each feature map of the image encoder and each token in the CLIP text encoder output,
therefore allowing relational associations between these to be learned. After the transformer, we take
the mean of the output tokens and apply a linear layer.

Few Shot Img Classification: We use the encoders as they are and employ a prototypical network as our method of achieving few shot learning.

- Image to Text Zero Shot: We use the standard CLIP cosine distance-based matching, and we employ BERT embeddings for text and for images we apply our chosen encoder.
- 1132 Video Classification and Regression: We process each frame with our chosen encoder which
 produce one vector per frame, and then use a 4-layer transformer to process the temporal axis before we apply a linear layer mapping to our classes or our single value output.

1134 I EXPERIMENTAL DETAILS

Experimental Environment Details: GPUs: 4 x A6000 Ada @ 48GB, CPUs: 128 Core AMD EPYC 7713 64-Core Processor, RAM: 1 TB, HD: 15TB NVME. All experiments were done with BF16 precision.

- J ADDITIONAL RESULTS
- 1143 J.1 Full details on discovered combinations

1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177 1178
1179
1180
1181 1182
1182
1184
1185



1239

Figure 6: Degradation of predictive power when a given benchmark is removed and the meta-model

trained from scratch, for different best combinations in varying k.

1242																						
1243	$\frac{\text{Metric} \downarrow \text{Model} \rightarrow}{\text{I}}$	cvnxtv2	siglip	clip	flex	deit	laion	vit	dino	smvit	rnx50	effv2	r50a1	effrmr	seffv2	sflex	svit	whspr	sr50a1	bert	bart	mpnet
1244	Img Class CIFAR-100 Acc@1	84.2	74.6	76.9	75.1	66.7	75.1	66.6	55.7	50.3	69.3	67.3	34.3	15.6		10.3	7.8	11.0			9.0	1.0
1245	CIFAR-100 Acc@5 CIFAR-100 Loss	97.4 0.6	$93.8 \\ 0.9$	$95.1 \\ 0.8$	$94.4 \\ 0.9$	$90.9 \\ 1.2$	$93.9 \\ 0.9$	$89.7 \\ 1.2$	$83.6 \\ 1.6$	$ 80.1 \\ 1.9 $	$91.9 \\ 1.2$	$90.7 \\ 1.3$	$\frac{65.9}{2.5}$	$\frac{42.3}{3.5}$	$67.6 \\ 2.4$	$\frac{30.6}{3.9}$	$25.5 \\ 4.1$	$31.6 \\ 3.9$	$\frac{40.2}{3.6}$	$\frac{38.1}{3.7}$	$29.2 \\ 4.0$	$5.0 \\ 4.6$
1246	Food-101 Acc@1	92.9	91.6	93.3	89.1	87.3	91.4	86.5	84.8	75.7	86.1	86.4	69.4	61.6	36.5	24.5	25.8	17.0	16.3	18.7	11.6	8.5
1247	Food-101 Acc@5 Food-101 Loss	$\begin{array}{c} 99.0 \\ 0.3 \end{array}$	98.7 0.3	99.1 0.2	$\frac{98.1}{0.4}$	$97.8 \\ 0.4$	$\frac{98.7}{0.3}$	$97.4 \\ 0.5$	$97.0 \\ 0.5$	$93.5 \\ 1.0$	$97.2 \\ 0.6$	$97.1 \\ 0.6$	$91.0 \\ 1.1$	$\frac{86.6}{1.5}$	$\frac{66.1}{2.6}$	$\frac{51.0}{3.2}$	$52.8 \\ 3.1$	$\frac{41.1}{3.6}$	$\frac{38.9}{3.6}$	$\frac{43.0}{3.5}$	$32.2 \\ 3.9$	$26.1 \\ 4.1$
1248	HWhale Individual Acc@1 HWhale Individual Acc@5	$75.6 \\ 84.6$	$31.7 \\ 49.5$	$35.2 \\ 53.9$	$ 48.4 \\ 64.5 $	$23.7 \\ 40.9$	$21.0 \\ 37.9$	$27.5 \\ 46.0$	$9.1 \\ 22.0$	$3.6 \\ 11.0$	78.7 86.7	$77.1 \\ 83.6$	$5.2 \\ 14.8$	$4.4 \\ 11.9$	$33.2 \\ 52.5$	$2.8 \\ 9.2$	$2.5 \\ 8.1$	$2.2 \\ 6.9$	$2.1 \\ 6.8$	$2.3 \\ 7.6$	$1.7 \\ 5.7$	$1.5 \\ 5.4$
1249	HWhale Individual Loss	1.6	4.6	4.3	3.6	4.9	5.1	4.7	5.9	6.7	1.3	1.5	6.4	6.6	3.9	7.0	7.1	7.3	7.3	7.2	7.5	7.4
1250	HWhale Species Acc@1 HWhale Species Acc@5	$\frac{99.8}{100.0}$	99.8 100.0	99.7 100.0	$99.8 \\ 100.0$	$99.5 \\ 100.0$	$99.7 \\ 100.0$	$99.7 \\ 100.0$	$99.2 \\ 99.9$	$95.4 \\ 99.6$	$99.7 \\ 100.0$	99.7 100.0	$92.1 \\ 98.9$	$92.8 \\ 99.1$	$96.5 \\ 99.8$	$76.5 \\ 96.1$	$74.5 \\ 95.8$	$64.3 \\ 92.0$	$65.8 \\ 92.6$		59.3 89.8	$62.9 \\ 91.1$
	HWhale Species Loss	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.2	0.0	0.0	0.3	0.2	0.1	0.8	$0.8 \\ 2.4$	$1.2 \\ 2.2$	1.1	0.9	1.4	1.2
1251	ImageNet-1K Acc@1 ImageNet-1K Acc@5	85.3 96.8	$81.9 \\ 95.8$	$76.0 \\ 93.7$	$82.3 \\ 95.5$	$\frac{82.1}{94.7}$	$74.1 \\ 93.1$		$77.9 \\ 93.0$	$75.5 \\ 90.8$	$77.6 \\ 93.3$	$73.5 \\ 91.4$	$72.5 \\ 90.5$	$44.6 \\ 72.5$		$3.2 \\ 10.1$	$\frac{2.4}{8.2}$	7.7	$1.3 \\ 4.7$	$1.5 \\ 5.2$	$0.8 \\ 3.2$	$\begin{array}{c} 0.2 \\ 1.2 \end{array}$
1252	ImageNet-1K Loss Places365 Acc@1	0.6 54.7	$\frac{0.8}{53.5}$	$\begin{array}{c} 1.0 \\ 54.1 \end{array}$	$\frac{0.8}{52.1}$	$0.8 \\ 49.0$	$1.1 \\ 53.7$	$\frac{1.3}{47.5}$	$\frac{1.0}{47.3}$	$2.3 \\ 27.1$	$1.0 \\ 51.8$	$\frac{1.2}{51.5}$	$1.1 \\ 40.9$	$\frac{2.8}{25.2}$	$\frac{4.3}{26.6}$	$6.0 \\ 9.0$	$6.1 \\ 8.6$	$6.1 \\ 7.5$	$6.5 \\ 5.0$	$6.4 \\ 5.2$	$6.6 \\ 3.0$	$\frac{6.8}{2.2}$
1253	Places365 Acc@5	85.3	84.1	84.7	83.3	80.8	84.3	79.9	79.5	59.9	82.9	82.6	73.5	55.2	55.5	26.3	25.0	22.4	16.4	16.4	11.0	9.0
1254	Places365 Loss Task Mean	1.7 88.0	$\frac{1.7}{79.6}$	1.7 80.1	$1.8 \\ 81.9$	$\frac{1.9}{76.1}$	$1.7 \\ 76.9$	$2.0 \\ 74.8$	$2.0 \\ 70.8$	$3.1 \\ 63.5$	$\frac{1.8}{84.6}$	$1.8 \\ 83.4$	$2.3 \\ 62.4$	$3.3 \\ 51.0$	$\frac{3.2}{52.2}$	$\frac{4.5}{29.1}$	$\frac{4.6}{28.1}$	$\frac{4.6}{25.5}$	$\frac{5.0}{25.5}$	$\frac{5.0}{26.5}$	$5.3 \\ 21.4$	$5.3 \\ 17.8$
1255	Few-Shot Img Class Aircraft Acc@1	96.7	96.6	97.4	95.9	95.3	96.7	96.3	94.4	92.9	91.6	90.6	86.2	78.2	59.2	54.9	50.4	55.1	58.2	61.2	60.8	57.2
1256	Aircraft Loss	0.2	0.2	0.2	0.2	0.3	0.2	0.2	0.3	0.3	0.4	1.2	0.4	311.5	44.1	2.1	2.1	1.6	2.3	2.5	1.2	1.6
1257	CUBirds Acc@1 CUBirds Loss	98.0 0.2	97.9 0.2	$97.2 \\ 0.2$	$96.4 \\ 0.3$	$96.2 \\ 0.2$	$96.6 \\ 0.2$	$95.9 \\ 0.2$	$94.4 \\ 0.3$	$93.4 \\ 0.3$	$92.8 \\ 0.4$	$92.1 \\ 0.5$	$89.4 \\ 0.4$	86.3 33.7	$\frac{52.5}{2.5}$	$\frac{50.0}{3.6}$	$\frac{45.2}{3.5}$	$\frac{44.4}{2.3}$	$\frac{31.9}{8.8}$	$\frac{48.4}{3.2}$	$\frac{50.3}{2.0}$	$\frac{48.5}{1.6}$
1258	DTextures Acc@1	85.0	85.2	88.6	78.9	81.9	86.1	80.8	79.4 1.2	81.9	77.7	60.3	77.2	68.5	46.6	50.2	50.5	50.0	33.1	44.6	49.8	38.3
1259	DTextures Loss Fungi Acc@1	0.9 85.8	$\begin{array}{c} 0.7\\ 85.6\end{array}$	0.5 85.7	$1.1 \\ 83.7$	$0.9 \\ 80.6$	$0.7 \\ 85.2$	$1.1 \\ 81.3$	1.2 77.4	$0.9 \\ 77.7$	$0.7 \\ 74.1$	$\frac{14.3}{73.7}$	$\frac{0.6}{67.1}$	$\frac{3.6}{59.2}$	$\frac{1.8}{27.6}$	$2.5 \\ 38.0$	$2.7 \\ 37.0$	$2.4 \\ 33.9$	$5.0 \\ 28.2$	$2.0 \\ 32.9$	$1.9 \\ 33.8$	$1.4 \\ 7.6$
1260	Fungi Loss Mini-Imagenet Acc@1	<mark>0.6</mark> 97.0	$0.6 \\ 96.2$	0.6 93.1	0.7 99.1	$\begin{array}{c} 0.8\\ 98.8 \end{array}$	$\frac{0.6}{90.8}$	$0.8 \\ 89.9$	$\begin{array}{c} 0.9 \\ 98.7 \end{array}$	$0.8 \\ 92.9$	$1.1 \\ 94.1$	$5.8 \\ 63.2$	$1.1 \\ 93.2$	$\begin{array}{c} \textbf{1031.2}\\ \textbf{90.9} \end{array}$	$2.6 \\ 36.7$	$2.2 \\ 45.9$	$2.2 \\ 47.2$	$2.2 \\ 44.8$	$2.4 \\ 34.2$	$2.4 \\ 39.7$	$\frac{2.3}{37.3}$	$2.9 \\ 36.8$
1261	Mini-Imagenet Loss	0.1	0.1	0.3	0.0	0.0	0.3	0.4	0.1	0.2	0.3	23.7	0.3	0.6	2.4	1.6	1.6	1.6	2.1	1.8	1.9	1.9
1262	Omniglot Acc@1 Omniglot Loss	$98.6 \\ 0.1$	$98.9 \\ 0.1$	99.0 0.1	$\frac{98.9}{0.1}$	$98.7 \\ 0.1$	$98.9 \\ 0.1$	$98.8 \\ 0.1$	$98.6 \\ 0.1$	$98.6 \\ 0.1$	$98.5 \\ 0.1$	$98.7 \\ 0.1$	$95.5 \\ 0.2$	$95.8 \\ 0.2$	$98.2 \\ 0.1$	$93.4 \\ 0.3$	$93.6 \\ 0.2$	$82.9 \\ 0.6$	$\begin{array}{c} 80.5\\ 0.7\end{array}$	$ 90.2 \\ 0.4 $	0.6	$90.7 \\ 0.3$
1263	VGG Flowers Acc@1 VGG Flowers Loss	99.7 0.1	98.9 0.1	$98.6 \\ 0.1$	$96.7 \\ 0.2$	$96.2 \\ 0.2$	$97.0 \\ 0.1$	$95.9 \\ 0.2$	$95.5 \\ 0.2$	$93.4 \\ 0.2$	$87.9 \\ 0.4$	$91.3 \\ 0.5$		$90.6 \\ 0.3$	$59.6 \\ 1.6$	$69.4 \\ 1.8$	$69.4 \\ 1.6$	$63.0 \\ 1.4$	$53.4 \\ 4.2$	$\frac{59.1}{2.5}$	$59.4 \\ 1.6$	$ 60.8 \\ 1.5 $
1264	Task Mean	94.4	94.2	94.2	92.8	92.5	93.1	91.3	91.2	90.1	88.1	81.4	85.4	81.4	54.3	57.4		53.4	45.6	53.7		48.6
1265	Img Seg ADE20K CE Loss	1.1	1.0	1.1	1.1	1.3	1.0	1.3	1.4	1.7	2.0	2.2	2.8	2.8	3.3	3.8	3.8	3.7	3.7	3.7	3.7	3.8
1266	ADE20K Focal Loss	0.2	0.2	$0.2 \\ 57.5$	$\frac{0.2}{56.0}$	0.2	$\frac{0.2}{57.3}$	0.3	$0.3 \\ 45.1$	0.3	$0.4 \\ 26.8$	$0.5 \\ 20.4$	0.6	0.6	0.8	0.9	0.9	0.9	0.9	0.9	0.9	0.9
1267	ADE20K Mean Acc@ ADE20K Overall Acc@	$\frac{59.8}{71.8}$	60.8 74.4	72.6	71.4	$49.1 \\ 66.9$	$\frac{57.5}{72.4}$	$44.2 \\ 64.2$	63.5	$36.3 \\ 57.5$	49.6	$\frac{20.4}{43.9}$	$17.9 \\ 34.6$	$15.2 \\ 39.7$	$\frac{3.6}{21.3}$	$1.6 \\ 11.7$	1.6 11.9	$1.8 \\ 13.1$	$1.8 \\ 14.1$	$1.8 \\ 14.0$	$1.8 \\ 14.4$	$1.8 \\ 14.2$
1268	ADE20K mIoU Cityscapes CE Loss	$\begin{array}{c} 46.8 \\ 0.2 \end{array}$	47.1 0.2	$\begin{array}{c} 44.0 \\ 0.2 \end{array}$	$43.7 \\ 0.2$	$37.8 \\ 0.2$	$43.4 \\ 0.2$	$33.2 \\ 0.2$	$33.3 \\ 0.2$	$25.9 \\ 0.2$	$ \begin{array}{c} 18.2 \\ 0.4 \end{array} $	$14.2 \\ 0.2$	$11.7 \\ 0.4$	$9.8 \\ 4.1$	$1.5 \\ 0.3$	$0.5 \\ 0.7$	$0.4 \\ 0.7$	$0.6 \\ 0.9$	0.4 0.9	$0.4 \\ 3.9$	$0.5 \\ 4.0$	$0.4 \\ 3.8$
1269	Cityscapes Focal Loss	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.1	0.0	0.1	1.0	0.0	0.1	0.1	0.1	0.1	0.9	0.9	0.9
1209	Cityscapes Overall Acc@ Cityscapes mIoU	$92.5 \\ 62.3$	94.2 69.8	$93.9 \\ 67.6$	$93.6 \\ 67.5$	$93.1 \\ 63.9$	$93.7 \\ 67.7$	$93.4 \\ 63.9$	$93.1 \\ 61.4$	$92.8 \\ 59.5$	$\frac{88.5}{40.8}$	$93.2 \\ 64.2$	$\frac{87.4}{40.2}$	$\frac{41.5}{2.5}$	$90.4 \\ 46.7$	$78.1 \\ 22.8$	$78.6 \\ 23.5$	$72.2 \\ 17.1$	$75.4 \\ 18.6$	$47.4 \\ 2.7$	$37.7 \\ 2.0$	$\frac{47.3}{2.7}$
	COCO-10K CE Loss COCO-10K Focal Loss	$3.0 \\ 0.7$	1.3 0.3	$1.5 \\ 0.3$	$1.4 \\ 0.3$	$1.5 \\ 0.3$	$1.4 \\ 0.3$	$1.5 \\ 0.3$	$1.6 \\ 0.3$	$1.6 \\ 0.3$	$2.1 \\ 0.4$	$2.6 \\ 0.6$	$3.3 \\ 0.8$	$3.5 \\ 0.8$	$3.6 \\ 0.8$	$\frac{4.5}{1.1}$	$3.8 \\ 0.9$	$4.0 \\ 0.9$	$3.6 \\ 0.8$	$4.1 \\ 1.0$	$3.7 \\ 0.9$	$4.1 \\ 1.0$
1271	COCO-10K Mean Acc@	38.8	50.6	47.2	46.0	43.4	44.9	41.2	43.5	40.7	27.0	15.8	8.2	20.9	2.2	1.7	1.9	1.3	2.9	0.6	2.5	0.6
1272	COCO-10K Overall Acc@ COCO-10K mIoU	$57.9 \\ 26.9$	69.8 39.5	$\begin{array}{c} 66.4\\ 35.6\end{array}$	$\frac{66.0}{35.1}$	$64.4 \\ 32.8$	$65.9 \\ 33.6$	$62.8 \\ 29.8$	$63.1 \\ 31.0$	$61.2 \\ 28.6$	$\frac{51.3}{18.4}$	$40.1 \\ 10.2$	$23.3 \\ 5.7$	$45.2 \\ 14.0$	$20.9 \\ 1.1$	$15.2 \\ 0.9$	$20.5 \\ 0.8$	$14.7 \\ 0.4$	$24.6 \\ 1.6$	$9.4 \\ 0.1$	$22.5 \\ 1.3$	$9.3 \\ 0.1$
1273	COCO-164K CE Loss COCO-164K Focal Loss	$1.9 \\ 0.4$	1.4 0.3	$1.5 \\ 0.3$	$1.5 \\ 0.3$	$1.6 \\ 0.3$	$1.5 \\ 0.3$	$1.6 \\ 0.3$	$1.7 \\ 0.4$	$1.8 \\ 0.4$	$2.2 \\ 0.5$	$2.7 \\ 0.6$	$3.5 \\ 0.8$	$7.0 \\ 1.7$	$3.7 \\ 0.9$	$\frac{4.3}{1.0}$	$3.9 \\ 0.9$	$4.0 \\ 0.9$	$4.0 \\ 0.9$	$\frac{4.2}{1.0}$	$3.7 \\ 0.9$	$4.2 \\ 1.0$
1274	COCO-164K Mean Acc@	45.9	50.1	46.9	45.3	42.6	44.5	38.6	43.0	38.7	25.4	14.7	7.0	21.3	2.0	1.5	1.9	1.5	1.8	0.6	2.5	0.7
1275	COCO-164K Overall Acc@ COCO-164K mIoU	$60.9 \\ 32.7$	65.8 36.7	$63.5 \\ 33.8$	$63.0 \\ 33.0$		$\frac{63.2}{32.4}$	$59.5 \\ 27.0$	$\frac{59.1}{28.9}$	$\frac{55.6}{25.7}$	$47.9 \\ 16.8$	$39.3 \\ 9.7$	$20.3 \\ 4.7$	$39.3 \\ 13.7$	$19.2 \\ 1.0$	$13.6 \\ 0.7$	$19.4 \\ 0.7$	$15.6 \\ 0.5$	$18.3 \\ 0.7$	$9.5 \\ 0.1$	$21.7 \\ 1.1$	$9.6 \\ 0.1$
1276	NYU CE Loss NYU Dice Score	2.5	1.5	$2.0 \\ 0.8$	$2.3 \\ 0.8$	$1.5 \\ 0.8$	$\frac{2.5}{0.8}$	2.3 0.8	$1.5 \\ 0.8$	1.6	$1.6 \\ 0.8$	$1.8 \\ 0.8$	$1.6 \\ 0.8$	1.4 0.8	1.6	1.6	1.6	$1.7 \\ 0.7$	$1.5 \\ 0.8$	$1.5 \\ 0.8$	$1.5 \\ 0.8$	$\frac{1.5}{0.8}$
1277	NYU Focal Loss	0.8 0.5	$0.8 \\ 0.2$	0.4	0.5	0.3	0.5	0.5	0.3	$0.8 \\ 0.3$	0.3	0.3	0.3	0.2	$0.8 \\ 0.3$	$0.8 \\ 0.3$	$0.8 \\ 0.3$	0.3	0.2	0.2	0.2	0.2
1278	NYU Mean Acc@ NYU Overall Acc@	$19.7 \\ 19.0$	$21.5 \\ 37.2$	$13.0 \\ 30.8$	$19.6 \\ 30.0$	22.7 42.8	$19.4 \\ 25.2$	$19.7 \\ 27.3$	23.0 34.7	$\frac{22.9}{31.2}$	$18.5 \\ 33.4$	$18.3 \\ 30.7$	$12.7 \\ 33.4$	$ 18.9 \\ 39.1 $		$\frac{10.0}{34.6}$			$13.0 \\ 36.3$	$11.9 \\ 37.2$		$12.0 \\ 37.4$
1279	NYU mIoU	7.5	7.7	7.8	$6.9 \\ 0.6$	12.2 0.9	5.7	6.1	$12.1 \\ 0.8$	11.0	5.9	8.3	$6.4 \\ 2.2$	$10.5 \\ 3.1$	6.8	$3.5 \\ 2.4$	$3.7 \\ 2.4$	$\frac{2.9}{2.4}$	$7.2 \\ 2.4$	$5.4 \\ 2.6$	$5.0 \\ 2.5$	$5.4 \\ 2.6$
1280	Pascal CE Loss Pascal Dice Loss	1.0 0.8	0.5 0.6	$\begin{array}{c} 0.5 \\ 0.4 \end{array}$	0.5	0.5	$\frac{0.5}{0.4}$	$0.8 \\ 0.5$	0.5	$0.9 \\ 0.4$	$1.4 \\ 0.5$	$1.5 \\ 0.4$	0.5	0.5	2.3 0.2	0.4	0.4	0.4	0.4	0.5	0.5	0.4
1281	Pascal Focal Loss Pascal Loss	$0.2 \\ 1.4$	0.1 0.5	0.1 0.1	$0.1 \\ 0.3$	0.2 0.3	$\frac{0.1}{0.4}$	$0.2 \\ 0.3$	$0.1 \\ 0.6$	$0.2 \\ 0.6$	$0.2 \\ 0.5$	$0.3 \\ 0.4$	$0.4 \\ 1.4$	$\begin{array}{c} 0.7 \\ 4.2 \end{array}$	$0.5 \\ 1.6$	$0.5 \\ 1.6$	$0.5 \\ 1.6$	$0.5 \\ 1.6$	$0.5 \\ 1.6$	$0.5 \\ 3.4$	$0.5 \\ 1.7$	$0.5 \\ 3.5$
1282	Pascal Mean Acc@	42.2	43.5	44.2	39.6	38.8	37.4	34.7	40.3	29.1	20.7	16.2	10.6	18.0	3.5	3.1	2.8	3.3	4.5	2.6	3.3	2.5
1283	Pascal Overall Acc@ Pascal mIoU	$75.1 \\ 32.8$	87.6 34.8	87.2 35.7	<mark>86.6</mark> 30.6	$77.5 \\ 31.4$	$\frac{86.6}{28.3}$	$78.9 \\ 27.5$	$79.5 \\ 29.8$	$76.7 \\ 24.0$			$49.7 \\ 6.8$	$\begin{array}{c} 66.6 \\ 14.0 \end{array}$	1.7	$34.2 \\ 1.3$	1.1	1.4	$\frac{39.6}{2.3}$		1.4	$32.3 \\ 0.9$
1284	Task Mean Img Relational	44.1	49.6	47.1	46.4	45.1	45.7	41.8	43.6	39.9	31.9	28.5	21.2	24.0	17.0	13.1	13.9	12.7	14.7	10.0	11.2	9.9
1285	CLEVR Acc@1	52.5	52.7	52.7	52.1	52.6	52.6	52.8	52.8	51.6	50.1	40.6	49.3	45.2	39.3	46.1		46.4	44.9	42.6		41.2
1286	CLEVR Colour Acc@1 CLEVR Colour Loss	35.4 1.5	$\frac{36.1}{1.5}$	36.4 1.5	$35.0 \\ 1.5$	$35.5 \\ 1.5$	$35.6 \\ 1.5$	$35.3 \\ 1.5$	$\frac{36.1}{1.5}$	$34.2 \\ 1.6$	$26.8 \\ 1.9$	$\frac{15.7}{2.1}$	$24.7 \\ 2.0$	$\frac{14.7}{2.1}$	$\frac{12.5}{2.1}$	$25.7 \\ 2.0$	1.9	$\frac{28.8}{1.9}$	$22.8 \\ 2.0$	$13.2 \\ 2.1$	2.1	13.2 2.1
1287	CLEVR Count Acc@1 CLEVR Count Loss	45.8 1.1	$45.8 \\ 1.2$	$ 45.8 \\ 1.1 $	45.9 1.2	$\frac{45.8}{1.2}$	$45.7 \\ 1.2$	$\frac{45.7}{1.2}$	$45.6 \\ 1.2$	$45.6 \\ 1.2$	$\frac{45.3}{1.2}$	$39.0 \\ 1.3$	$\frac{45.1}{1.2}$	$44.8 \\ 1.2$	$37.9 \\ 1.4$	$45.1 \\ 1.2$	$44.7 \\ 1.2$	$44.8 \\ 1.2$	$\frac{44.9}{1.2}$	$44.7 \\ 1.2$	$ \begin{array}{r} 44.7 \\ 1.2 \end{array} $	$43.0 \\ 1.2$
1288	CLEVR Material Acc@1	60.5	60.6	60.5	60.0	60.5	60.6	61.4	61.3	60.2	58.6	52.1	57.5	53.7	49.8	53.7	51.7	54.0	53.0	49.8	50.5	49.9
1289	CLEVR Material Loss CLEVR Shape Acc@1	$\frac{0.7}{52.1}$	$0.7 \\ 52.4$	$\begin{array}{c} 0.7 \\ 52.5 \end{array}$	$0.7 \\ 51.1$	$0.7 \\ 52.2$	$\begin{array}{c} 0.7 \\ 52.4 \end{array}$	0.6 52.9	0.6 51.2	$0.7 \\ 49.9$	$0.7 \\ 50.2$	$0.7 \\ 34.3$	$0.7 \\ 50.2$	$0.7 \\ 44.8$	$\begin{array}{c} 0.7\\ 33.3\end{array}$	$0.7 \\ 35.8$	$0.7 \\ 34.9$	$0.7 \\ 36.1$	$0.7 \\ 34.6$	$0.7 \\ 34.6$	$0.7 \\ 33.7$	0.7 33.4
1290	CLEVR Shape Loss CLEVR Size Acc@1	0.9 61.0	$0.9 \\ 61.1$	0.9 61.3	$1.0 \\ 60.7$	$0.9 \\ 61.1$	$0.9 \\ 60.8$	$\begin{array}{c} 0.9 \\ 62.0 \end{array}$	1.0 62.3	$1.0 \\ 60.9$	$1.0 \\ 59.6$	$1.1 \\ 53.5$	$1.0 \\ 58.3$	$1.1 \\ 55.7$	$1.1 \\ 50.6$	$1.1 \\ 56.2$	1.1	$1.1 \\ 55.2$	$1.1 \\ 54.6$		1.1	$1.1 \\ 50.1$
1291	CLEVR Size Loss	0.6	0.6	0.6	0.6	0.6	0.6	0.6	0.6	0.6	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7
1292	CLEVR Yes/No Acc@1 CLEVR Yes/No Loss	$\begin{array}{c} 60.7 \\ 0.6 \end{array}$	$60.5 \\ 0.6$	60.8 0.6	60.6 0.6		$\frac{60.7}{0.6}$				$59.8 \\ 0.6$	$53.3 \\ 0.7$	$59.9 \\ 0.6$	$59.6 \\ 0.6$	$51.4 \\ 0.7$		$59.2 \\ 0.6$	$59.5 \\ 0.6$	$59.8 \\ 0.6$	$59.5 \\ 0.6$	$59.3 \\ 0.6$	$58.6 \\ 0.6$
1292	CLEVR-Math Acc@1	79.3	65.9	68.8	59.9	73.7	62.9	60.5	59.3	58.3	55.6	44.0	56.0	56.6	30.2	46.9	46.5	46.2	45.7	44.8	42.1	36.4
1293	CLEVR-Math Acc@5 CLEVR-Math Loss	99.8 0.5	$99.5 \\ 0.8$	$\begin{array}{c} 99.6 \\ 0.7 \end{array}$	$98.9 \\ 0.9$	$\begin{array}{c} 99.7 \\ 0.6 \end{array}$	$99.3 \\ 0.8$	$ 99.2 \\ 0.9 $	$ 98.9 \\ 0.9 $	$98.9 \\ 1.0$	$98.8 \\ 1.0$	$97.7 \\ 1.3$	$98.8 \\ 1.0$	$98.8 \\ 1.0$	$\frac{86.1}{1.7}$	1.2	$98.1 \\ 1.2$	1.2	$97.7 \\ 1.2$		1.3	$92.8 \\ 1.5$
1294	Task Mean Medical Class	60.8	59.4	59.8	58.2	60.2	59.0	58.9	58.7	57.8	56.1	47.8	55.5	52.7	43.5	52.0	51.7	52.1	50.9	49.0	48.5	46.5
1233	Chexpert 0 APS	75.7	76.5	76.6	76.8	76.8	74.7	76.0	75.8	76.3	75.1	75.2	69.1	70.3	65.3	20.6	22.3	21.9	29.4	31.6	25.2	23.2

1296																				
1297	Chexpert 0 AUC	91.3	$92.1 \\ 7.4$	92.5	$92.3 \\ 7.4$	92.6	91.4	92.2	$92.3 \\ 7.3$	92.6 6.9	$91.0 \\ 7.9$	91.6	89.9	$90.5 \\ 7.7$	88.5	61.5 64.0		71.3	72.3 66	
1298	Chexpert 0 BS Chexpert 1 APS	$7.8 \\ 55.3$	55.2	$\begin{array}{c} 7.3 \\ 55.5 \end{array}$	55.8	$\frac{7.0}{54.2}$	$7.5 \\ 54.4$	$7.3 \\ 54.2$	52.1	55.9	53.1	$7.3 \\ 53.3$	$7.9 \\ 44.2$	43.0	$\frac{8.4}{33.5}$	12.6 12.8 28.9 30.3	31.0	$11.9 \\ 28.9$	$12.1 \ 12 \\ 30.1 \ 29$.9 28.5
1299	Chexpert 1 AUC Chexpert 1 BS	$75.7 \\ 18.8$	$\frac{76.0}{18.5}$	$75.3 \\ 19.4$	77.0 20.2	$75.4 \\ 18.7$	$75.3 \\ 20.6$	$\frac{75.3}{20.6}$	$73.8 \\ 18.6$	76.1 16.3	$74.9 \\ 17.6$	$75.2 \\ 20.3$	$\begin{array}{c} 69.4 \\ 17.2 \end{array}$	$\begin{array}{c} 69.8 \\ 17.2 \end{array}$	$64.1 \\ 18.3$	56.3 56.9 18.7 18.0		$57.1 \\ 18.6$	57.6 57 18.5 18	
1300	Chexpert 2 APS Chexpert 2 AUC	$43.8 \\ 71.8$	$43.8 \\ 71.2$	$43.5 \\ 71.8$	45.1 72.4	45.5 72.1	$44.8 \\ 72.0$	$43.9 \\ 71.7$	$\frac{42.3}{71.3}$	$43.6 \\ 71.7$	$\frac{43.5}{70.5}$	$44.4 \\ 71.1$	$41.8 \\ 69.9$	$\frac{42.8}{70.9}$	$35.3 \\ 63.1$	30.1 31.3 58.6 59.0		$32.3 \\ 60.7$	$32.6 31 \\ 60.5 60$	
1301	Chexpert 2 BS	18.5	17.8	21.0	21.1	18.6	21.2	20.5	19.1	17.0	16.0	20.4	16.2	16.2	17.4	18.4 18.5	2 18.1	17.9	$17.9 \ 17$.8 17.9
1302	Chexpert 3 APS Chexpert 3 AUC	$ 80.7 \\ 86.8 $	$\frac{80.9}{86.8}$	$ 80.8 \\ 86.5 $	82.1 87.9	$ 81.7 \\ 87.2 $	$ 80.5 \\ 86.6 $	$79.7 \\ 85.9$	$78.6 \\ 84.6$	$79.1 \\ 85.8$	$79.2 \\ 84.9$	$\begin{array}{c} 80.6 \\ 87.0 \end{array}$	$73.5 \\ 82.3$	$75.3 \\ 83.5$	$\frac{58.6}{73.0}$	51.7 $50.365.6$ 65.3	65.2	$53.2 \\ 65.6$	54.0 48 67.2 64	.2 64.2
1303	Chexpert 3 BS Chexpert 4 APS	$17.4 \\ 53.4$	$16.4 \\ 49.5$	$16.4 \\ 50.1$	$15.6 \\ 53.4$	15.3 54.5	$16.2 \\ 50.9$	$16.9 \\ 52.6$	$17.2 \\ 50.8$	$\frac{15.9}{52.3}$	$17.6 \\ 49.9$	$16.3 \\ 50.7$	$18.1 \\ 41.7$	$17.1 \\ 44.9$	$23.5 \\ 47.3$	26.1 26.0 38.4 36.7		$25.2 \\ 39.2$	$24.8 \ 26 \ 37.9 \ 35$	
1304	Chexpert 4 AUC Chexpert 4 BS	$87.5 \\ 10.4$	$\frac{86.7}{10.0}$	$87.0 \\ 10.9$	88.1 10.2	$\frac{88.0}{9.1}$	$87.0 \\ 10.9$	$87.3 \\ 10.2$	$\frac{86.8}{9.9}$	87.7 8.8	$\frac{86.0}{9.4}$	$\frac{86.4}{11.6}$	$84.1 \\ 10.1$	$85.1 \\ 9.6$	84.8 9.4	81.7 80.3 9.7 10.0		$81.5 \\ 9.9$	81.3 79 10.7 10	
1305	Chexpert APS Macro	61.6	61.0	61.2	62.6	62.3	60.9	61.2	59.9	61.5	59.8	60.2	54.1	55.2	48.0	33.9 34.	34.3	35.7	36.9 33	.7 33.0
1306	Chexpert AUC Macro Chexpert BS Macro	$\frac{82.5}{15.7}$	$82.5 \\ 15.6$	$\frac{82.3}{15.5}$	83.2 14.9	$\frac{82.9}{13.8}$	$82.5 \\ 15.4$	$\frac{82.4}{15.1}$	$\begin{array}{c} 81.8\\14.4\end{array}$	82.8 13.0	$\begin{array}{c} 81.1 \\ 13.7 \end{array}$	$\frac{81.9}{15.2}$	$79.1 \\ 13.9$	$\begin{array}{c} 79.9 \\ 13.6 \end{array}$	$74.7 \\ 15.4$	64.7 65.1 17.1 17.1	17.0	$67.0 \\ 16.9$	$\begin{array}{ccc} 67.6 & 65 \\ 16.9 & 17 \end{array}$.2 17.2
1307	Chexpert Loss Diabetic 0 APS	0.3 93.0	$0.4 \\ 91.8$	$\begin{array}{c} 0.5\\ 91.5\end{array}$	$0.3 \\ 91.3$	0.3 90.9	$\frac{0.3}{91.3}$	$0.4 \\ 90.6$	$0.4 \\ 90.4$	$\begin{array}{c} 0.3\\ 88.3 \end{array}$	$0.3 \\ 90.8$	$0.4 \\ 91.5$	$0.3 \\ 85.4$	$0.4 \\ 87.2$	$\begin{array}{c} 0.4 \\ 75.5 \end{array}$	0.5 0.5 76.3 75.0		$0.5 \\ 79.8$	0.4 0.79.4 76	
1308	Diabetic 0 AUC Diabetic 0 BS	86.3 10.7	$84.6 \\ 11.9$	$84.0 \\ 12.3$	$83.9 \\ 12.4$	$83.0 \\ 12.6$	$83.6 \\ 12.6$	$81.7 \\ 13.0$	$80.9 \\ 13.0$	$77.2 \\ 14.6$	$83.9 \\ 12.1$	$84.3 \\ 11.7$	$72.2 \\ 16.5$	$75.1 \\ 15.7$	$\frac{52.4}{19.0}$	54.3 53.0 19.5 19.4			58.6 54 18.6 19	
1309	Diabetic 1 APS Diabetic 1 AUC	14.0 69.6	$13.6 \\ 67.2$	$14.0 \\ 67.4$	$13.0 \\ 66.0$	$13.0 \\ 65.3$	12.9 66.1	14.5 66.5	$10.8 \\ 65.3$	9.0 59.7	$12.6 \\ 66.5$	$13.5 \\ 66.4$	$8.4 \\ 54.4$	9.0 59.5	$7.2 \\ 51.4$	8.4 8.8 54.9 56.9	8.9	8.4 53.9	7.7 7 54.9 52	4 7.3
1310	Diabetic 1 BS	6.1	6.4	6.5	6.1	6.0	6.8	6.4	5.8	5.8	6.0	6.4	6.9	5.3	6.5	6.7 - 6.5	6.9	6.4	6.3 6	4 6.3
1311	Diabetic 2 APS Diabetic 2 AUC	65.5 88.5			$61.4 \\ 86.0$	$\frac{58.4}{84.7}$	$57.1 \\ 85.3$	$54.2 \\ 84.3$	$\frac{51.1}{82.5}$	$\frac{44.3}{79.6}$	$59.7 \\ 85.5$		$28.9 \\ 71.6$	$32.2 \\ 73.8$	$14.6 \\ 50.9$	$17.0 \ 16.' \\ 53.4 \ 52.'$		$20.2 \\ 61.2$	$17.8 17 \\ 57.7 54$	
1312	Diabetic 2 BS Diabetic 3 APS	8.0 41.6	$\frac{8.5}{49.7}$	$9.0 \\ 47.6$	$9.3 \\ 48.4$	$9.7 \\ 45.3$	9.0 53.1	$9.5 \\ 46.5$	$9.9 \\ 38.8$	$10.7 \\ 37.1$	$9.2 \\ 47.2$	$\frac{8.3}{50.7}$	$\frac{11.7}{22.4}$	$11.7 \\ 32.0$	$\frac{12.1}{2.8}$	12.7 12.8 3.1 3.1		$\frac{12.7}{4.6}$	11.9 12 4.0 3	
1313	Diabetic 3 AUC	94.8 1.9	96.5 1.6	$95.7 \\ 1.6$	$95.6 \\ 1.6$	$93.9 \\ 1.7$	$95.1 \\ 1.9$	$95.0 \\ 1.7$	94.1 1.8	$93.5 \\ 1.8$	$95.1 \\ 1.7$	96.2 1.5	87.2 2.0	92.3 2.1	56.0 2.4	56.1 57.2 2.4 2.5	2 59.1	$64.2 \\ 2.2$	64.9 58 2.3 2	.4 52.3
1314	Diabetic 3 BS Diabetic 4 APS	73.9	74.3	73.0	75.3	67.5	68.7	70.2	72.3	47.5	67.5	74.6	32.4	23.7	2.9	3.1 3.0	4.4	3.9	3.7 <mark>2</mark>	5 2.6
1315	Diabetic 4 AUC Diabetic 4 BS	98.7 1.0	$98.2 \\ 1.1$	$97.7 \\ 1.0$	$\begin{array}{c} 98.7 \\ 0.9 \end{array}$	$98.0 \\ 1.1$	$97.4 \\ 1.1$	$ 98.4 \\ 0.9 $	$98.3 \\ 0.9$	$\frac{96.9}{1.3}$	$97.2 \\ 1.1$	97.9 0.8	$94.7 \\ 1.4$	$94.3 \\ 1.8$	$\frac{56.4}{1.9}$	60.1 58.0 1.9 1.8			64.3 56 1.8 1	
1316	Diabetic APS Macro Diabetic AUC Macro	56.9 87.5	$57.2 \\ 86.7$	$\frac{56.4}{86.0}$	$\frac{56.3}{85.7}$	$54.2 \\ 85.0$	$\frac{56.4}{85.3}$	$54.4 \\ 84.7$	$51.9 \\ 83.8$	$45.2 \\ 81.2$	$55.6 \\ 85.6$	58.7 86.1	$35.5 \\ 76.0$	$36.6 \\ 79.0$	$20.6 \\ 53.4$	21.6 21.8 55.7 55.7		$23.3 \\ 61.3$	$22.4 \ 21$ 59.4 55	
1317	Diabetic BS Macro Diabetic Loss	5.5 0.2	$6.0 \\ 0.1$	$6.1 \\ 0.2$	$6.1 \\ 0.1$	6.2 0.1	$6.4 \\ 0.1$	$6.3 \\ 0.2$	$6.4 \\ 0.2$	$7.0 \\ 0.2$	$6.1 \\ 0.2$	$5.8 \\ 0.2$	$7.7 \\ 0.2$	$7.4 \\ 0.3$	$\frac{8.4}{0.2}$	8.7 8.6 0.3 0.3	8.6	$8.5 \\ 0.3$		4 8.4
1318	HAM10K 0 APS	94.3	90.0	90.3	88.7	89.2	90.9	89.8	89.0	83.3	88.2	84.1	47.4	58.0	30.4	32.8 25.8	3 25.0	41.2	$46.2 \ 34$.4 33.8
1319	HAM10K 0 AUC HAM10K 0 BS	99.1 2.1	98.2 2.9	$\frac{98.3}{3.5}$	$97.6 \\ 2.8$	$97.8 \\ 3.1$	$\frac{98.2}{3.1}$	$97.7 \\ 3.1$	$\frac{98.0}{3.4}$	$\frac{96.7}{3.8}$	$97.6 \\ 3.4$	$97.0 \\ 4.0$	$\frac{89.0}{7.0}$	$91.7 \\ 6.3$	$\frac{80.6}{8.2}$	81.2 78. 8.1 8.5	8.6	$\frac{85.2}{7.6}$	86.9 82 7.3 8	1 8.3
1320	HAM10K 1 APS HAM10K 1 AUC	99.2 98.9	$99.2 \\ 98.7$	$99.1 \\ 98.4$	$99.2 \\ 98.5$	$99.2 \\ 98.4$	$99.1 \\ 98.4$	$99.1 \\ 98.4$	$99.2 \\ 98.4$	$98.7 \\ 97.3$	$98.9 \\ 98.1$	$98.1 \\ 97.1$	$96.2 \\ 92.7$	$96.5 \\ 93.5$	$94.2 \\ 89.7$	93.9 93.' 88.7 88.'		$95.5 \\ 91.0$	96.0 94 91.9 88	
1321	HAM10K 1 BS HAM10K 2 APS	3.1 95.5	3.7 98.6	$4.5 \\ 89.0$	$\frac{4.2}{94.4}$	$4.5 \\ 88.7$	$4.4 \\ 92.1$	$4.6 \\ 92.4$	$4.4 \\ 95.3$	$6.2 \\ 69.7$	$5.0 \\ 81.6$	$6.3 \\ 89.0$	$10.0 \\ 11.3$	$9.4 \\ 5.0$	$\frac{11.7}{5.7}$	$12.5 12.8 \\ 5.2 8.1$		$11.3 \\ 19.5$	$10.7 \ 13$ $12.2 \ 7$	
1322	HAM10K 2 AUC	99.9	100.0	99.7	99.9	99.7	99.8	99.8	99.9	99.3	98.4	99.8	81.1	75.6	79.4	79.6 73.0	68.2	90.8	87.2 81	.2 78.3
1323	HAM10K 2 BS HAM10K 3 APS	0.3 88.0	$\begin{array}{c} 0.3 \\ 85.5 \end{array}$	$0.4 \\ 83.9$	0.3 85.2	$\begin{array}{c} 0.5 \\ 86.2 \end{array}$	$\begin{array}{c} 0.3 \\ 83.0 \end{array}$	$0.3 \\ 84.0$	$0.3 \\ 82.5$	$0.8 \\ 74.2$	$\begin{array}{c} 0.5 \\ 80.8 \end{array}$	$\begin{array}{c} 0.4 \\ 74.3 \end{array}$	$1.3 \\ 41.9$	$1.3 \\ 46.7$	$1.3 \\ 34.7$	1.3 1.3 35.1 33.1	2 31.5	$1.2 \\ 42.5$	$1.3 \ 1$ $48.4 \ 42$.4 35.2
1324	HAM10K 3 AUC HAM10K 3 BS	96.7 3.5	$95.5 \\ 3.7$	$95.6 \\ 4.2$	$95.9 \\ 3.9$	$\frac{96.1}{3.5}$	$95.3 \\ 4.1$	$95.9 \\ 4.2$	$\frac{96.1}{4.4}$	$94.4 \\ 5.0$	$95.4 \\ 4.7$	$92.5 \\ 5.1$	$83.8 \\ 7.9$	$\frac{84.9}{7.6}$	$81.7 \\ 8.4$	80.0 80.3 8.4 8.4	8.5	$\frac{85.9}{7.7}$	88.1 84 7.2 8	0 8.2
1325 1326	HAM10K 4 APS HAM10K 4 AUC	$99.5 \\ 100.0$	100.0 100.0	99.7 100.0	$98.2 \\ 100.0$	$100.0 \\ 100.0$	$98.5 \\ 100.0$	100.0 100.0	$98.5 \\ 100.0$	$98.7 \\ 100.0$	$96.4 \\ 99.9$	$96.9 \\ 100.0$	$26.8 \\ 92.3$	$21.9 \\ 94.5$	$33.6 \\ 84.2$	32.3 24.0 89.0 89.5		$52.8 \\ 97.7$	73.8 34 97.6 92	
1320	HAM10K 4 BS HAM10K 5 APS	0.0 95.6	$\begin{array}{c} 0.0 \\ 94.8 \end{array}$	$0.1 \\ 94.5$	$0.1 \\ 91.5$	<mark>0.0</mark> 90.3	$0.1 \\ 93.7$	0.0 91.6	$0.0 \\ 90.8$	$0.2 \\ 83.2$	$0.1 \\ 88.0$	$0.2 \\ 91.4$	$1.1 \\ 54.1$	$1.2 \\ 67.0$	$1.2 \\ 41.8$	1.0 1.1 36.4 36.8	1.2 3 22.3	$0.9 \\ 48.1$	$\begin{array}{ccc} 0.5 & 1 \\ 41.3 & 36 \end{array}$	
1328	HAM10K 5 AUC	99.7	99.7	99.6	99.5	99.2	99.5	99.4	99.0	98.8	98.0	99.4	94.7	96.6	92.9	92.0 91.3	5 87.5	94.0	$92.5 \ 90$.3 88.2
1329	HAM10K 5 BS HAM10K 6 APS	1.1 89.2	1.1 85.2	$1.1 \\ 83.9$	$1.1 \\ 86.3$	$\begin{array}{c} 1.3 \\ 88.0 \end{array}$	1.0 87.6	$1.2 \\ 84.7$	$1.3 \\ 83.3$	$\frac{1.7}{75.8}$	$1.5 \\ 81.3$	$1.4 \\ 83.2$	$3.4 \\ 28.4$	$2.8 \\ 33.4$	$3.8 \\ 31.4$	3.9 3.9 3.9 30.5 29.8	524.5	$3.6 \\ 39.6$	$3.9 4 \\ 36.6 25$.8 23.5
1330	HAM10K 6 AUC HAM10K 6 BS	99.3 1.0	$98.3 \\ 1.5$	$99.1 \\ 1.5$	$98.6 \\ 1.1$	$\begin{array}{c} 98.9 \\ 1.1 \end{array}$	$\frac{99.1}{1.3}$	99.3 1.4	$98.6 \\ 1.3$	$98.0 \\ 1.7$	$98.0 \\ 1.4$	$98.6 \\ 1.6$	$91.6 \\ 3.0$	$93.6 \\ 2.9$	$91.2 \\ 3.0$	91.9 91.4 3.0 3.0		$92.3 \\ 2.8$	93.4 90 2.9 3	
1331	HAM10K APS Macro HAM10K AUC Macro	94.5 99.1	$\frac{93.3}{98.6}$	$91.4 \\ 98.7$	$\frac{92.2}{98.5}$	$91.3 \\ 98.6$	$92.1 \\ 98.6$	$91.6 \\ 98.7$	$90.8 \\ 98.5$	$83.4 \\ 97.8$	$87.9 \\ 97.9$	$87.1 \\ 97.5$	$43.7 \\ 89.3$	$46.9 \\ 90.1$	$38.8 \\ 85.6$	38.0 35.9 86.1 84.0		$48.5 \\ 91.0$	$50.6 \ 37$ 91.1 85	
1332	HAM10K BS Macro	1.6 0.3	1.9 0.2	$2.2 \\ 0.3$	$\frac{1.9}{0.3}$	$2.0 \\ 0.3$	$2.1 \\ 0.2$	$2.1 \\ 0.2$	$2.1 \\ 0.2$	$2.8 \\ 0.2$	$2.4 \\ 0.2$	2.8 0.6	$4.8 \\ 0.2$	$4.5 \\ 0.2$	$5.4 \\ 0.2$	5.5 $5.60.2$ 0.2	5.7	$5.0 \\ 0.2$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	6 5.8
1333	HAM10K Loss Task Mean	57.0	56.7	56.5	$\frac{0.3}{56.7}$	56.2	$\frac{0.2}{56.4}$	56.2	55.5	53.6	55.3	0.0 56.0	$\frac{0.2}{45.0}$	46.0	39.4	37.4 37.0		40.6	40.8 37	
1334	Medical Seg ACDC Dice Score	0.6	0.5	0.5	0.5	0.4	0.5	0.5	0.5	0.4	0.4	0.6	0.4	0.2	0.5	0.2 0.2	0.3	0.3	0.2 0	
1335	ACDC Mean Acc@ ACDC Overall Acc@	$\frac{86.3}{86.5}$	$\frac{85.8}{86.2}$	$83.4 \\ 83.2$	$78.5 \\ 78.7$	75.5 75.1	$78.0 \\ 78.3$	$76.9 \\ 77.0$	$79.4 \\ 79.0$	$74.0 \\ 73.5$	$93.4 \\ 93.5$	94.1 94.2	$71.7 \\ 71.5$	$67.6 \\ 67.5$	$76.0 \\ 76.0$	46.7 53.7 47.2 53.4			56.1 50 55.5 51	
1336	ACDC mIoU Task Mean	$57.9 \\ 57.8$	$57.0 \\ 57.3$	$57.4 \\ 56.1$	$53.1 \\ 52.7$	$50.2 \\ 50.3$	$53.0 \\ 52.4$	$47.7 \\ 50.5$	$54.3 \\ 53.3$	$50.1 \\ 49.5$	$66.9 \\ 63.6$	67.2 64.0	$47.5 \\ 47.8$	$47.9 \\ 45.8$		27.6 30.4	1 35.6	$35.1 \\ 39.0$	32.1 24 36.0 31	.3 26.9
1337	Img to Txt ZS																			
1338	Flickr30K Img2Txt Acc@1 Flickr30K Img2Txt Acc@5	$6.3 \\ 20.9$	$\begin{array}{c} 6.3 \\ 21.3 \end{array}$	7.0 21.0	$5.9 \\ 20.0$	$5.6 \\ 19.3$	6.8 22.1	$5.9 \\ 20.4$	$5.2 \\ 18.8$	$4.5 \\ 18.0$	$4.1 \\ 16.0$	$3.7 \\ 16.1$	$\frac{4.7}{16.9}$	$\frac{4.2}{15.5}$	$1.6 \\ 8.1$	$ \begin{array}{ccc} 1.8 & 2.0 \\ 8.6 & 8.4 \end{array} $	8.9	$2.0 \\ 9.1$	$ \begin{array}{ccc} 1.9 & 1 \\ 9.1 & 8 \end{array} $	5 8.4
1339	Flickr30K Img2Txt Loss Flickr30K Txt2Img Acc@1	$3.8 \\ 5.7$	$\frac{3.8}{5.9}$	3.8 6.0	$3.8 \\ 5.3$	$3.9 \\ 5.1$	3.7 6.5	$3.8 \\ 6.0$	$3.9 \\ 5.1$	$3.9 \\ 5.0$	$3.9 \\ 3.8$	$3.9 \\ 4.0$	$\frac{4.0}{4.2}$	$\frac{4.0}{3.9}$	$\frac{4.2}{1.7}$	$\begin{array}{ccc} 4.1 & 4.1 \\ 1.8 & 2.0 \end{array}$		$\frac{4.1}{2.3}$	$ \begin{array}{ccc} 4.1 & 4 \\ 1.9 & 1 \end{array} $	
1340	Flickr30K Txt2Img Acc@5 Flickr30K Txt2Img Loss	20.9 3.8	22.1 3.8	$21.6 \\ 3.8$	20.8 3.9	20.0 3.9	23.0 3.8	$21.0 \\ 3.8$	19.8 3.9	$18.9 \\ 3.9$	$16.5 \\ 3.9$	$17.3 \\ 4.0$	$17.1 \\ 4.0$	$15.5 \\ 4.0$	$\frac{7.8}{4.2}$	8.9 8.4 4.2 4.2	9.2	9.4 4.1	9.5 8 4.1 4	8 8.3
1341	NYCC Img2Txt Acc@5	21.4	21.4	22.0	20.0	21.2	22.1	21.4	20.0	17.8	17.1	17.0	15.9	15.8	7.9	8.7 8.9	8.7	9.5	8.9 8	5 7.9
1342	NYCC Img2Txt Loss NYCC Img2Txt	$\frac{3.8}{6.9}$	$\frac{3.8}{6.6}$	3.8 6.9	$\frac{3.8}{5.8}$	$3.8 \\ 6.5$	3.8 6.9	$3.8 \\ 6.4$	$3.8 \\ 6.0$	$3.9 \\ 4.7$	$3.9 \\ 4.9$	$3.9 \\ 4.1$	$4.0 \\ 4.6$	$4.0 \\ 4.2$	$\frac{4.2}{1.6}$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$		$\frac{4.1}{2.1}$	$ \begin{array}{ccc} 4.1 & 4 \\ 2.0 & 1 \end{array} $	
1343	NYCC Loss NYCC Txt2Img Acc@5	$\frac{3.8}{21.9}$	$3.8 \\ 21.6$	3.8 22.5	$\frac{3.8}{20.2}$	$\frac{3.8}{21.9}$	$3.8 \\ 21.9$	3.8 22.7	$\frac{3.8}{20.7}$	$3.9 \\ 18.4$	$3.9 \\ 17.3$	$3.9 \\ 17.4$	$4.0 \\ 16.0$	$\frac{4.0}{15.3}$	$\frac{4.9}{7.9}$	4.1 4.1 9.4 8.3		$4.1 \\ 9.9$	4.1 4 8.9 8	
1344	NYCC Txt2Img Loss NYCC Txt2Img	3.8 6.1	3.8 5.9	3.8 6.4	3.8 5.5	3.8 6.0	3.8 6.2	3.8 6.4	3.9 5.8	3.9 4.8	3.9 4.3	3.9 4.1	4.0 3.9	4.0 3.7	$\frac{5.5}{1.6}$	$4.1 4.1 \\ 2.0 1.7$	4.1	$\frac{4.1}{2.4}$	4.1 4	2 4.2
1345	Winoground Img2Txt Loss	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7	0.7 0.7	0.7	0.7	0.7 0.	7 0.7
1346	Winoground Img2Txt Winoground Txt2Img Loss	$51.0 \\ 0.7$	$53.4 \\ 0.7$	$\begin{array}{c} 59.5 \\ 0.7 \end{array}$	49.7 0.7	$\begin{array}{c} 50.0 \\ 0.7 \end{array}$	$\frac{50.3}{0.7}$	$49.5 \\ 0.7$	$43.5 \\ 0.7$	$\frac{53.8}{0.7}$	61.9 0.7	$\begin{array}{c} 50.0 \\ 0.7 \end{array}$	48.9 0.7	47.3 0.7	$\begin{array}{c} 43.9 \\ 0.7 \end{array}$	50.0 41. 0.7 0.7	0.7	$53.2 \\ 0.7$	$\begin{array}{ccc} 49.6 & 50 \\ 0.7 & 0 \end{array}$	7 0.7
1347	Winoground Txt2Img Task Mean	$50.0 \\ 21.1$	$55.2 \\ 22.0$	56.2 22.9	$53.1 \\ 20.6$	$\frac{50.0}{20.6}$	$55.5 \\ 22.1$	$\frac{48.3}{20.8}$	$54.2 \\ 19.9$	$\frac{48.6}{19.4}$	$54.8 \\ 20.0$	$\frac{50.0}{18.4}$	$49.6 \\ 18.2$	$52.4 \\ 17.8$	$52.8 \\ 13.5$	50.0 54.2 14.3 13.2		$52.2 \\ 15.2$	$51.8 \ 48$ 14.6 14	
1348	Video Class HMDB-51 Acc@1	52.5	40.7	40.6	32.2	39.3	24.9	27.4	32.8	33.1	5.6	11.5	1.8	2.1	3.8	8.3 7.9		5.4	6.4 7	
1349	HMDB-51 Acc@5	81.4	70.0	70.5	60.9	68.6	54.2	58.5	59.8	63.8	23.0 4.7	28.8	10.4	$\frac{10.2}{4.1}$	13.6	$26.4 \ 25.3$	3 17.8	23.6	24.4 24	.9 15.6
	HMDB-51 Loss	2.1	2.8	3.1	3.4	2.7	3.8	3.3	3.1	3.0	4.1	4.4	4.7	4.1	3.9	4.2 4.3	4.4	3.7	3.8 3	1 3.9

1050																						
1350	Kinetics Acc@1	48.8	44.2	51.4	43.7	40.3	44.6	33.2	36.4	25.8	2.7	1.0	0.2	0.3	0.4	2.0	1.6	1.0	0.5	0.3	0.3	0.3
1351	Kinetics Acc@5 Kinetics Loss	$\frac{75.5}{2.4}$	$\frac{70.9}{2.6}$	77.9 2.1	$\frac{70.7}{2.5}$	$67.6 \\ 2.7$	$\frac{71.7}{2.5}$	$\frac{59.9}{3.2}$	$63.0 \\ 3.0$	$\frac{51.8}{3.5}$	$9.7 \\ 5.5$	$4.3 \\ 6.1$	$1.3 \\ 6.1$	$1.4 \\ 6.1$	$1.7 \\ 6.1$	$7.0 \\ 5.7$	$\frac{6.5}{5.8}$	$3.5 \\ 6.0$	$2.2 \\ 6.1$	$1.3 \\ 6.1$	$1.3 \\ 6.1$	$\frac{1.3}{6.1}$
1352	UCF-101 Acc@1 UCF-101 Acc@5	84.4 95.4	$75.1 \\ 92.5$	$69.9 \\ 89.1$	$63.2 \\ 82.3$	$75.0 \\ 91.6$	$63.4 \\ 86.2$	$\frac{58.8}{81.7}$	$66.6 \\ 86.3$	$\frac{48.7}{75.3}$	$19.7 \\ 42.2$	$\frac{11.1}{28.9}$	$2.8 \\ 8.5$	$\begin{array}{c} 0.8 \\ 5.0 \end{array}$	$2.1 \\ 8.2$	$\frac{15.2}{35.5}$	$13.3 \\ 33.8$	$6.6 \\ 17.9$	$\frac{8.7}{25.2}$	$6.5 \\ 23.1$	$7.0 \\ 20.2$	$2.7 \\ 11.2$
1353	UCF-101 Loss Task Mean	0.6 73.0	$1.0 \\ 65.6$	$1.3 \\ 66.6$	$1.7 \\ 58.8$	$\begin{array}{c} 1.0\\ 63.7\end{array}$	$1.5 \\ 57.5$	$1.7 \\ 53.3$	$1.4 \\ 57.5$	$2.3 \\ 49.8$	$4.3 \\ 17.2$	5.0 14.3	$4.8 \\ 4.2$	$\frac{4.7}{3.3}$	4.6	3.7	$3.8 \\ 14.7$	$4.5 \\ 8.8$	4.0	4.2 10.3	4.2	4.5
1354	Video Reg														5.0				10.9			5.8
1355	IWildCam MAE Score IWildCam MSE Loss	1.3 3.7	$1.4 \\ 4.4$	$\frac{1.3}{4.0}$	1.4 4.0	$1.4 \\ 4.1$	$\frac{1.6}{5.4}$	$\frac{1.4}{4.3}$	$1.5 \\ 5.0$	$\frac{1.6}{5.9}$	$2.0 \\ 7.1$	$1.9 \\ 6.5$	$1.9 \\ 6.2$	2.6 12.5	$2.1 \\ 8.5$	$1.8 \\ 5.1$	$1.8 \\ 6.3$	$1.9 \\ 6.0$	$1.8 \\ 6.2$	$\frac{2.2}{8.6}$	$1.8 \\ 6.4$	$\frac{2.1}{8.4}$
1356	Task Mean GATE	1.3	1.4	1.3	1.4	1.4	1.6	1.4	1.5	1.6	2.0	1.9	1.9	2.6	2.1	1.8	1.8	1.9	1.8	2.2	1.8	2.1
1357	Full GATE Mean	69.0	66.8	66.8	64.6	64.3	63.4	62.1	62.2	58.5	56.3	54.4	48.4	42.8			37.2		36.9	35.0		31.8
1358 1359	Big GATE Mean Base GATE Mean	76.6 68.3	$74.5 \\ 65.6$	$\begin{array}{c} 74.4 \\ 65.7 \end{array}$	$72.8 \\ 62.6$	$72.0 \\ 63.7$	$71.9 \\ 60.7$	$70.6 \\ 60.2$	$\begin{array}{c} 70.0 \\ 60.7 \end{array}$	$\frac{66.8}{58.6}$	$\frac{66.7}{55.1}$	$64.8 \\ 53.5$	$\frac{58.5}{48.2}$	$53.1 \\ 42.8$	$ 46.8 \\ 38.0 $		$\frac{43.4}{36.3}$	$\frac{41.9}{35.4}$	$\frac{41.5}{36.6}$	$\frac{40.9}{34.8}$		$37.1 \\ 30.4$
1360	Small GATE Mean Full GATE Rank	77.7 1.0	$\frac{74.9}{3.0}$	$\begin{array}{c} 74.6 \\ 2.0 \end{array}$	$73.3 \\ 4.0$	$72.4 \\ 5.0$	$71.2 \\ 6.0$	$\frac{68.9}{8.0}$	$69.1 \\ 7.0$	$65.3 \\ 9.0$	$65.7 \\ 10.0$	$61.7 \\ 11.0$	$58.5 \\ 12.0$	$49.3 \\ 13.0$			$35.4 \\ 16.0$		$35.3 \\ 17.0$	$34.1 \\ 19.0$		$30.4 \\ 21.0$
1361	Big GATE Rank Base GATE Rank	1.0 1.0	$2.0 \\ 3.0$	$3.0 \\ 2.0$	$4.0 \\ 5.0$	$5.0 \\ 4.0$	$6.0 \\ 7.0$	$7.0 \\ 8.0$	$8.0 \\ 6.0$	$9.0 \\ 9.0$	$\begin{array}{c} 10.0 \\ 10.0 \end{array}$	11.0	$12.0 \\ 12.0$	$13.0 \\ 13.0$	14.0	15.0	$\begin{array}{c} 16.0 \\ 17.0 \end{array}$	17.0	$18.0 \\ 15.0$	$19.0 \\ 20.0$		21.0 21.0
1362	Small GATE Rank	1.0	2.0	3.0	4.0	5.0	6.0	8.0	7.0	10.0	9.0	11.0		13.0			17.0		18.0			21.0
1363	Tabl	e 3: Fu	ll exp	erime	ents ta	able:	Black	c/Bol	d best	t mod	el. Gi	reen s	secon	d best	. Blu	e thi	rd be	est. a	nd re	d		
1364		worst p																				
1365		pretrai																				
1366																						
1367																						
1368																						
1369																						
1370																						
1371																						
1372																						
1373																						
1374																						
1375																						
1376																						
1377																						
1378																						
1379																						
1380																						
1381																						
1382																						
1383																						
1384																						
1385																						
1386																						
1387 1388																						
1389																						
1390																						
1390																						
1392																						
1393																						
1394																						
1395																						
1396																						
1397																						
1398																						
1399																						
1400																						
1401																						
1402																						
1403																						







Figure 7: Degradation of predictive power when a given benchmark is removed and the meta-model trained from scratch, for different best combinations in varying k.



0.03



1511






Figure 11: Ranking Heatmap for smallGATE: We show how the various models on the y-axis rank on the metrics on the x-axis, where brighter is higher/better rank. From left to right we apply a spearman correlation sorting to capture tasks more similar to imagenet1k more towards the leftmost side, and, dissimilar ones towards the rightmost side. From top to bottom we rank models based on average rank.



