

352 **A Proofs**

353 *Proof of Proposition 1.* Let  $G$  denote the distribution of the score  $S = s(X, Y)$  for a randomly  
 354 sampled example  $(X, Y) \sim F$ . For any  $m \in \{1, \dots, M\}$ , let  $G^{(m)}$  denote the distribution of the score  
 355 conditioned on  $Y$  being in cluster  $m$ . Consider a randomly sampled test example  $(X_{\text{test}}, Y_{\text{test}})$  with a  
 356 label in cluster  $m$ , so its corresponding score  $s_{\text{test}} = s(X_{\text{test}}, Y_{\text{test}})$  follows distribution  $G^{(m)}$ . Now  
 357 consider  $\{s_i\}_{i \in \mathcal{I}_2(m)}$ , the scores for examples in the proper calibration dataset with labels in cluster  
 358  $m$ . Since each element of  $\{s_i\}_{i \in \mathcal{I}_2(m)}$  follows distribution  $G^{(m)}$  and is chosen to be included in the  
 359 proper calibration dataset independently from the other elements and  $s_{\text{test}}$ , the element of  $\{s_i\}_{i \in \mathcal{I}_2(m)}$   
 360 and  $s_{\text{test}}$  are independent and identically distributed, which means that they are exchangeable. Thus,  
 361 by the standard proof of coverage for conformal prediction (see, e.g., [1]), the desired result follows.

362 *Proof of Proposition 2.* This is a direct result of exchangeability and Proposition 1.

363 *Proof of Proposition 3.* This proof follows the structure of the proof of Proposition 4 in [2]. Let  
 364  $S = s(X, Y)$  for  $(X, Y) \sim F$  be a random variable representing the score of a randomly sampled  
 365 example. Let  $\mathcal{S}$  denote the set of values that  $S$  can take. Let  $G^y(s) = \mathbb{P}(S \leq s \mid Y = y)$  denote the  
 366 cdf of  $S$  when the label is  $y$ . Define  $\mathcal{Y}^{(m)} = \{y \in \mathcal{Y} : \tilde{h}(y) = m\}$  as the set of classes in cluster  $m$   
 367 and let  $G^{(m)}(s) = \mathbb{P}(S \leq s \mid Y \in \mathcal{Y}^{(m)})$  denote the cdf of  $S$  when the label  $Y$  is in cluster  $m$ . Let  
 368  $S^{(m)}$  be a random variable with cdf  $G^{(m)}$  and for an arbitrary  $y \in \mathcal{Y}^{(m)}$ , let  $S^y$  be a random variable  
 369 with cdf  $G^y$ .

370 Since we assume that the TV distance between the score distribution for every pair of classes in  
 371 cluster  $m$  is bounded by  $\epsilon$ , and  $G^{(m)}$  is a mixture of these distributions, it follows that

$$\text{TV}(S^y, S^{(m)}) \leq \epsilon.$$

372 By definition of TV distance, this is equivalent to

$$\sup_{A \in \mathcal{S}} |\mathbb{P}(S \in A \mid Y = y) - \mathbb{P}(S \in A \mid Y \in \mathcal{Y}^{(m)})| \leq \epsilon,$$

373 which we can rewrite as

$$\sup_{f \in \mathcal{F}_1} |\mathbb{E}[f(S) \mid Y = y] - \mathbb{E}[f(S) \in A \mid Y \in \mathcal{Y}^{(m)}]| \leq \epsilon,$$

374 where  $\mathcal{F}_1 = f : \mathcal{S} \rightarrow [0, 1]$ . Define  $g(s) = \mathbb{1}\{s \geq \hat{q}(m)\}$ . Since  $g \in \mathcal{F}_1$ , we have

$$\mathbb{E}[\mathbb{1}\{S \geq \hat{q}(m)\} \mid Y = y] - \mathbb{E}[\mathbb{1}\{S \geq \hat{q}(m)\} \mid Y \in \mathcal{Y}^{(m)}] \leq \epsilon,$$

375 which can be expressed as

$$\mathbb{P}(S \geq \hat{q}(m) \mid Y = y) - \mathbb{P}(S \geq \hat{q}(m) \mid Y \in \mathcal{Y}^{(m)}) \leq \epsilon.$$

376 Since the CLUSTERED procedure will exclude the true label  $Y$  from the prediction set  $C$  exactly  
 when  $S \geq \hat{q}(m)$ , the probabilities can be re-expressed in terms of mis-coverage:

$$\mathbb{P}(Y \notin C(X) \mid Y = y) - \mathbb{P}(Y \notin C(X) \mid Y \in \mathcal{Y}^{(m)}) \leq \epsilon.$$

377 By Proposition 1, we know  $\mathbb{P}(Y \notin C(X) \mid Y \in \mathcal{Y}^{(m)}) \leq \alpha$ , so

$$\mathbb{P}(Y \notin C(X) \mid Y = y) \leq \alpha + \epsilon.$$

378 Taking the complement yields

$$\mathbb{P}(Y \in C(X) \mid Y = y) \geq 1 - \alpha - \epsilon.$$

379 This is true for all  $y \in \mathcal{Y}^{(m)}$  and for every cluster  $m = 1, \dots, M$ .

380 **B Experiment details**

381 **B.1 Score functions**

382 We perform experiments using three score functions:

383 • softmax: The conformal score of an input  $x$  and a label  $y$  is one minus the softmax score:

$$s_{\text{softmax}}(x, y) = 1 - f_y(x)$$

384 where  $f_y(x)$  is entry  $y$  of the softmax vector of input  $x$ .

385 • APS: *Adaptive Prediction Sets* are designed to achieve approximate  $X$ -conditional coverage  
 386 [17]. The conformal score of input  $x$  and label  $y$  is computed as follows: For  $y = 1, \dots, |\mathcal{Y}|$ ,  
 387 let  $\hat{p}_y(x)$  be an estimate of  $\mathbb{P}(Y = y \mid X = x)$ . We use the softmax score as our  $\hat{p}$ , so  
 388  $\hat{p}_y(x) = f_y(x)$ . Let  $\hat{p}_{(i)}(x)$  be the  $i$ -th largest  $\hat{p}_{(y)}(x)$ . Define  $j$  to be the index in the sorted  
 389 order that corresponds to class  $y$ , i.e.,  $\hat{p}_{(j)}(x) = \hat{p}_y(x)$ . Then,

$$s_{\text{APS}}(x, y) = \sum_{i=1}^{j-1} \hat{p}_{(i)}(x) + \text{Unif}([0, \hat{p}_{(j)}(x)])$$

390 • RAPS: One problem with APS is that the resulting prediction sets are often very large. *Reg-*  
 391 *ularized Adaptive Prediction Sets* [3] modifies APS by introducing an additive regularization  
 392 term designed to reduce the prediction set sizes:

$$s_{\text{RAPS}}(x, y) = s_{\text{APS}}(x, y) + \max(0, \lambda(o_x(y) - k_{\text{reg}}))$$

393 where  $o_x(y)$  denotes the ranking of  $y$  among the values of  $\hat{p}_k(x)$  for all classes  $k$  (e.g.,  
 394  $o_x(y) = 1$  if  $\hat{p}_y(x)$  is larger than all other  $\hat{p}_k(x)$ ), and  $\lambda$  and  $k_{\text{reg}}$  are user-chosen parameters.  
 395 In our experiments, we use  $\lambda = 0.01$  and  $k_{\text{reg}} = 5$ , which Angelopoulos et al. found to  
 396 work well for ImageNet [3].

## 397 B.2 Model training

398 An important consideration when training our models is that we need to reserve sufficient data for  
 399 evaluating the class-conditional coverage of the conformal prediction methods. In practice, this  
 400 means we should aim to exclude at least 250 examples per class from the model training dataset so  
 401 that we can use those untouched examples for validation (i.e., applying the conformal methods and  
 402 computing coverage and set size metrics).

403 For all datasets except ImageNet, we use a ResNet-50 as our predictive model. We initialize to the  
 404 IMAGENET1K\_V2 pre-trained weights, then fine-tune all parameters by training on the dataset-specific  
 405 data. We apply the model to the validation data to obtain softmax scores.

406 **ImageNet.** Setting up ImageNet for our setting is a bit tricky because we want sufficient data  
 407 for performing validation, but we also need this data to be separate from the model training data.  
 408 Unfortunately, the ImageNet validation set only contains 50 examples per class, which is not enough  
 409 for validation in our setting. Fortunately, we have access to more labeled data from the ImageNet  
 410 training set, which has roughly 1000 examples per class. However, if we want to use this data for  
 411 validation, we cannot use our ResNet-50 initialized to the IMAGENET1K\_V2 pretrained weights, as  
 412 these weights are obtained by training on the ImageNet training set and would violate the assumption  
 413 of independence of the validation and model training datasets. To approximately satisfy this inde-  
 414 pendence assumption, we instead use SimCLR-v2 [5], which is trained on the ImageNet training  
 415 set *without labels*, to extract feature vectors of length 6144 for all images in the ImageNet training  
 416 set. We then use 10% of these feature vectors for training a linear head (i.e., a single fully connected  
 417 neural network layer). After training for 10 epochs, the model achieves a validation accuracy of 78%.  
 418 We then apply the linear head to the remaining 90% of the feature vectors to obtain softmax scores.

419 **CIFAR-100.** In total, there are 600 images per class (500 from the training set and 100 from the  
 420 validation set). We combine the data and randomly sample 50% for model training, leaving the  
 421 remaining data for testing our procedure. After training for 30 epochs, the validation accuracy is  
 422 60%.

423 **Places365.** This dataset contains more than 10 million images of 365 classes. Each class has 5000  
 424 to 30000 examples. We randomly sample 90% of the data for model training and use the remaining  
 425 data for testing our procedure. After training for one epoch, the validation accuracy is 52%.

426 **iNaturalist.** This dataset has class labels of varying specificity. At the species level, there are  
 427 6414 classes with 300 examples each (290 training examples and 10 validation examples) and a  
 428 total of 10000 classes with at least 150 examples. We operate at the family level, which groups  
 429 the species into 1103 classes. We randomly sample 50% of the data for model training and use the  
 430 remaining for testing our procedure. After training for one epoch, the validation accuracy is 69%.  
 431 However, due to class imbalance and the randomness of the model training set construction, some  
 432 classes have insufficient validation samples. We filter out classes with fewer than 250 validation  
 433 examples, which leaves us with 633 classes. The entries of the softmax vectors that correspond to  
 434 rare classes are removed and the vector is renormalized to sum to one.

### 435 B.3 Choosing clustering parameters

436 In order to perform CLUSTERED, there are two parameters that must be chosen:  $\gamma$ , the probability  
 437 that a calibration example will be assigned to the clustering dataset, and  $M$ , the number of clusters  
 438 that will be requested when performing  $k$ -means.

439 As mentioned in Section 3.1, we make use of two intuitive heuristics to choose these parameters. We  
 440 restate these heuristics in more detail here.

- 441 • First, to distinguish between more clusters (or distributions), we need more samples from  
 442 each distribution. As a rough guess, to distinguish between two distributions, we want at  
 443 least four samples per distribution; to distinguish between five distributions, we want at least  
 444 ten samples per distribution. In other words, we want the number of clustering examples per  
 445 class to be at least twice as large as the number of clusters. This heuristic can be expressed  
 446 as

$$\gamma\tilde{n} \geq 2M, \tag{3}$$

447 where  $\gamma\tilde{n}$  is the expected number of clustering examples for the rarest class not assigned to  
 448 the null cluster.

- 449 • Second, we want enough data for computing the conformal quantiles for each cluster. We  
 450 translate this into asking for at least 150 examples per cluster on average. This heuristic can  
 451 be expressed as

$$(1 - \gamma)\tilde{n} \frac{K}{M} \geq 150, \tag{4}$$

452 where  $\frac{K}{M}$  is the average number of classes per cluster and  $(1 - \gamma)\tilde{n}$  is the expected number  
 453 of proper calibration examples for the rarest class not assigned to the null cluster.

454 Changing the inequalities of (3) and (4) into equalities and solving for  $\gamma$  and  $M$  yields

$$M = \frac{\gamma\tilde{n}}{2} \quad \text{and} \quad \gamma = \frac{K}{K + 75}.$$

455 **Varying the clustering parameters.** Although our method for choosing parameter values is  
 456 arguably ad-hoc, we find that it does not really matter what parameter values are used, as long as they  
 457 fall into a reasonable range. As the heatmaps in Figure 3 illustrate, the performance of CLUSTERED  
 458 is not very sensitive to  $\gamma$  and  $M$ . When  $n_{\text{avg}} = 10$ , the heuristic chooses  $\gamma = 0.89$  and  $M = 4$ .  
 459 When  $n_{\text{avg}} = 50$ , the heuristic chooses  $\gamma \in [0.88, 0.92]$  and  $M \in [7, 12]$  (since the calibration  
 460 dataset is randomly sampled, and  $\gamma$  and  $M$  are chosen based on the calibration dataset, there can be  
 461 some randomness in the chosen values). However, there are large areas surrounding these chosen  
 462 values that would yield similar performance. We observe that the heuristics do not always choose the  
 463 parameter values that yield the lowest CovGap. The heatmaps show that the optimal parameter values  
 464 are dependent not only on dataset characteristics, but also on the score function. Future work could  
 465 be done to extract further performance improvements by determining a better method for choosing  $\gamma$   
 466 and  $M$ .

### 467 B.4 Measuring dataset class balance in Table 1

468 The class balance metric in Table 1 is defined as the number of examples in the rarest 5% of classes  
 469 divided by the expected number of examples if the class distribution were perfectly uniform. This  
 470 metric is bounded between 0 and 1, with lower values denoting more class imbalance. The metric  
 471 is computed on the validation datasets, which are sampled uniformly at random from the publicly  
 472 available versions of each dataset.

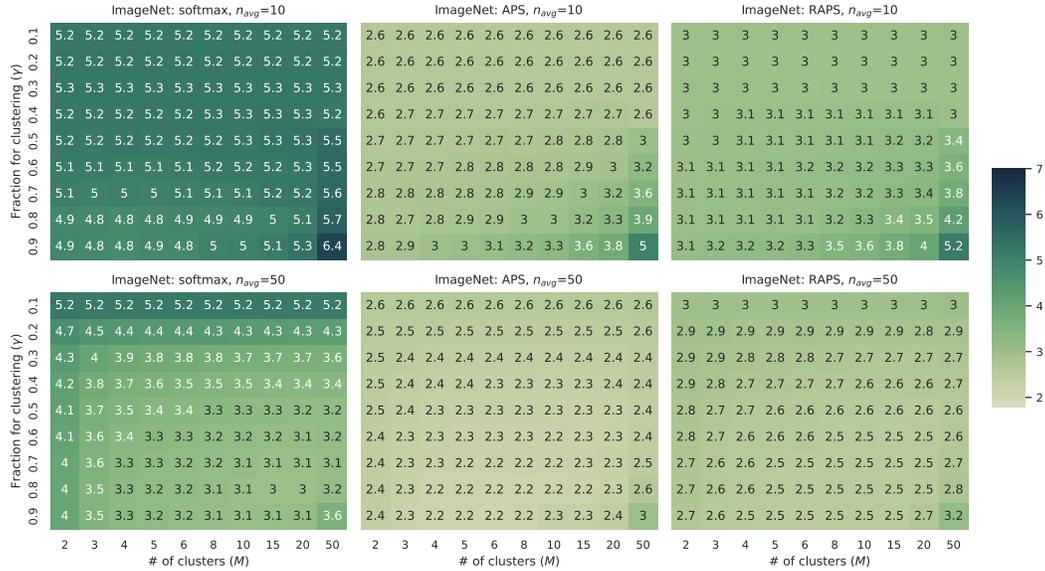


Figure 3: The average class coverage gap on ImageNet for  $n_{avg} = 10, 50$  using softmax, APS, and RAPS as we vary the clustering parameters. Each entry is computed across 10 random splits of the data into calibration and validation sets.

## 473 C Additional experimental results

474 We present additional experimental results in this section. As in the main text, shaded regions in plots  
 475 denote  $\pm 1.96$  times the standard errors.

### 476 C.1 RAPS CovGap results

477 Figure 4 shows the CovGap on all datasets when we use RAPS as our score function.

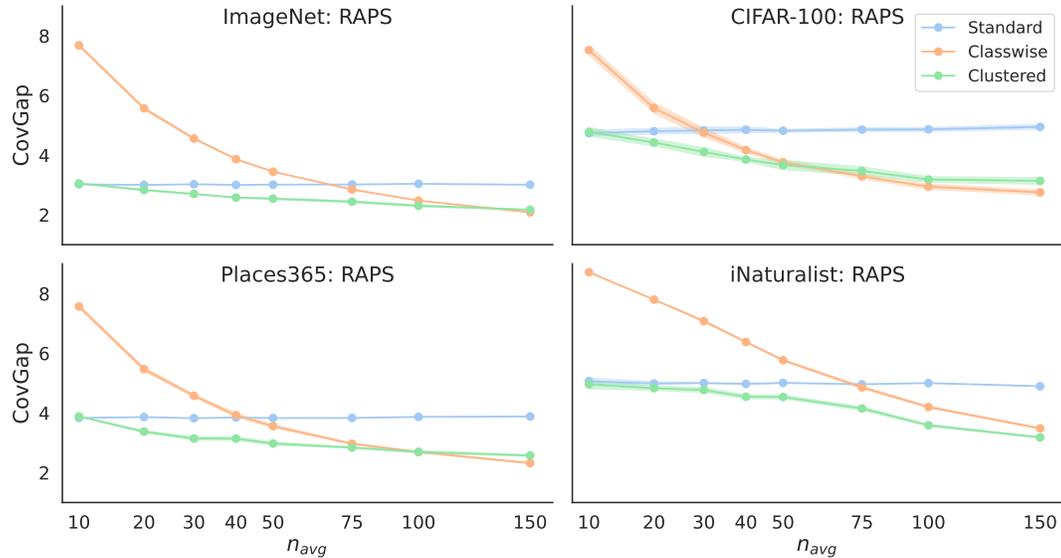


Figure 4: Average class coverage gap for ImageNet, CIFAR-100, Places365, and iNaturalist using RAPS scores, as we vary the average number of calibration examples per class.

478 **C.2 Additional metrics**

479 **Average set size.** To supplement Table 2 from the main text, which reports AvgSize for four values  
 480 of  $n_{avg}$ , Figure 5 plots AvgSize for all values of  $n_{avg}$  that we use in our experimental setup. Note that  
 481 RAPS sharply reduces AvgSize relative to APS on ImageNet and also induces a slight reduction for  
 482 the other three datasets. This asymmetric reduction is likely in large part due to the fact that the RAPS  
 483 hyperparameters, which control the strength of the set-size regularization, were tuned on ImageNet.  
 484 The set sizes of RAPS on other datasets could likely be improved by tuning the hyperparameters for each  
 485 dataset.

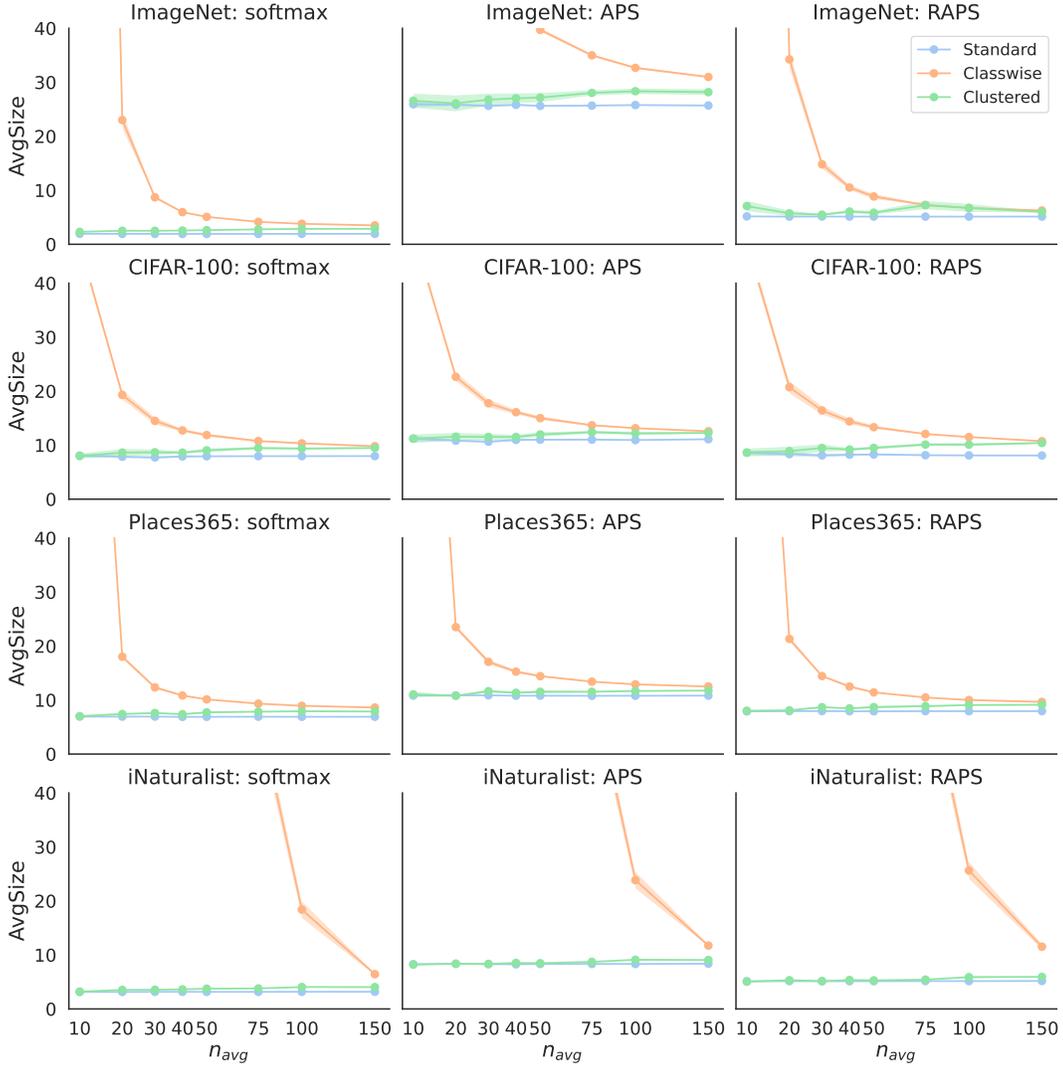


Figure 5: Average set size for ImageNet, CIFAR-100, Places365, and iNaturalist using softmax, APS, and RAPS scores, as we vary the average number of calibration examples per class.

486 **Fraction under-covered.** In many practical settings, we want to limit the number of classes that  
 487 are severely under-covered, which we define as having a class-conditional coverage that is more than  
 488 10% below the desired coverage level. We define  $\text{FracUnderCov}$  to be the fraction of classes that are  
 489 severely under-covered:

$$\text{FracUnderCov} = \frac{1}{|\mathcal{Y}|} \sum_{y=1}^{|\mathcal{Y}|} \mathbb{1}\{c_y \leq 1 - \alpha - 0.1\},$$

490 recalling that  $c_y$  is the class-conditional coverage for class  $y$ .

491 Figure 6 plots  $\text{FracUnderCov}$  for all experimental settings. Comparing to the  $\text{CovGap}$  plots in Figure  
 492 2 and Figure 4, we see that the trends in  $\text{FracUnderCov}$  generally mirror the trends in  $\text{CovGap}$ .  
 493 However,  $\text{FracUnderCov}$  is a much noisier metric, as evidenced by the large error bars. Another  
 494 flaw of  $\text{FracUnderCov}$  as a metric is it is unable to penalize uninformatively large set sizes. This is  
 495 best seen in the performance of CLASSWISE on iNaturalist: for every score function, CLASSWISE  
 496 has very low  $\text{FracUnderCov}$ , but this is achieved by producing extremely large prediction sets, as  
 497 shown in the bottom row of Figure 5. On the other hand,  $\text{CovGap}$  is able to impose a slight penalty  
 498 on this kind of behavior since unnecessarily large set sizes often lead to over-coverage, and  $\text{CovGap}$   
 499 penalizes over-coverage.

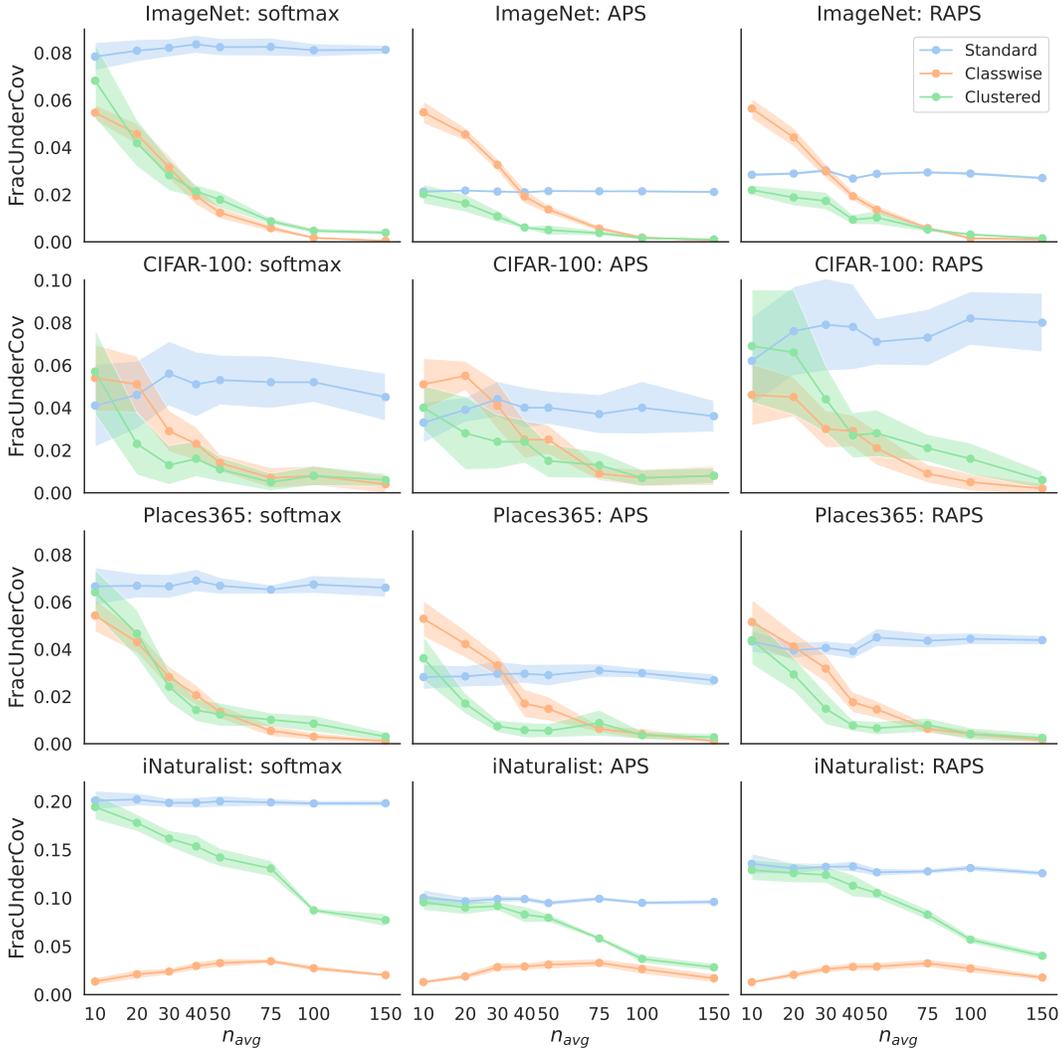


Figure 6: Fraction of very under-covered classes for ImageNet, CIFAR-100, Places365, and iNaturalist using softmax, APS, and RAPS scores, as we vary the average number of calibration examples per class.