

# Meta-learning Pathologies from Radiology Reports using Variance Aware Prototypical Networks

Arijit Sehanobish Kawshik Kannan\* Nabila Abraham Anasuya Das Benjamin Odry

Covera Health

New York City, NY

{arijit.sehanobish, kawshik.kannan, nabila.abraham,  
anasuya.das, benjamin.odry}@coverahealth.com

## Abstract

Large pretrained Transformer-based language models like BERT and GPT have changed the landscape of Natural Language Processing (NLP). However, fine tuning such models still requires a large number of training examples for each target task, thus annotating multiple datasets and training these models on various downstream tasks becomes time consuming and expensive. In this work, we propose a simple extension of the Prototypical Networks for few-shot text classification. Our main idea is to replace the class prototypes by Gaussians and introduce a regularization term that encourages the examples to be clustered near the appropriate class centroids. Experimental results show that our method outperforms various strong baselines on 13 public and 4 internal datasets. Furthermore, we use the class distributions as a tool for detecting potential out-of-distribution (OOD) data points during deployment.

## 1 Introduction

Pretrained Transformer-based language models (PLMs) have achieved great success on many NLP tasks (Devlin et al., 2019; Brown et al., 2020), but still need a large number of in-domain labeled examples for finetuning (Yogatama et al., 2019). Learning to learn (Lake et al., 2015a; Schmidhuber, 1987; Bengio et al., 1997) from limited supervision is an important problem with widespread application in areas where obtaining labeled data can be difficult or expensive. To that end, meta-learning methods have been proposed as effective solutions for few-shot learning (Hospedales et al., 2020). Current applications of such meta-learning methods have shown improved performance in few-shot learning for vision tasks such as learning to classify new image classes within a similar dataset. Namely, on classical few-shot image classification benchmarks, the training tasks are sampled from

a “single” larger dataset (for ex: Omniglot (Lake et al., 2015b) and miniImageNet (Vinyals et al., 2016)), and the label space contains the same task structure for all tasks. There has been a similar trend of such classical methods in NLP as well (Geng et al., 2019). In contrast, in text classification tasks, the set of source tasks available during training and target tasks during evaluation can range from sentiment analysis to grammatical acceptability judgment (Bansal et al., 2020a,b). In recent works (Wang et al., 2021), the authors use a range of different source tasks (different not only in terms of input domain, but also their task structure i.e. label semantics, and number of labels) for meta-training and show successful performance on a wide range of downstream tasks. In spite of this success, meta-training on various source tasks is quite challenging as it requires resistance to overfitting to certain source tasks due to its few-shot nature and more task-specific adaptation due to the distinct nature among tasks (Roelofs et al., 2019).

However, in medical NLP, collecting large number of diverse labeled datasets is difficult. In our institution, we collect high quality labeled radiology reports (which are always used as held out test data) and use it to train our internal annotators who then annotate our unlabeled data. This training process is expensive and time consuming. Our annotation process is described in section A. Thus a natural question is: if we have a large labeled dataset consisting of a lot of classes, can we use it to meta-train a model that can be used on a large number of downstream datasets where we have little to no training examples? This is a challenging problem as the reports can be structured differently based on the report type and there can be a substantial variation in writing style across radiologists from different institutions. Our main goal is to build out a set of extensible pipelines that can generalize to new pathologies typically in new sub-specialties while also generalizing across different health systems.

\*Equal Contribution

In addition, the exact definition of the pathologies and their severity change can change depending on the clinical use case. This makes fully supervised approaches that rely on large labeled datasets expensive. Having few-shot capabilities allows us to annotate a handful of cases and rapidly expand the list of pathologies we can detect and classify. In addition, we can use our approach to generate pseudo labels for rare pathologies and enrich our validation and test sets for annotation by an in-house clinical team. Lastly our approach can be extended to support patient search and define custom cohorts of patients.

Our contributions in this work are the following: **(1)** We develop a novel loss function that extends the vanilla prototypical networks and introduce a regularization term that encourages tight clustering of examples near the class prototypes. **(2)** We meta-train our models on a large labeled dataset on shoulder MRI reports (single domain) and show good performance on 4 diverse downstream classification tasks on radiology reports on knee, cervical spine and chest. In addition to our internal datasets, we show superior performance of our method on 13 public benchmarks over well-known methods like Leopard. Our model is very simple to train, easy to deploy unlike gradient based methods and just requires a few additional lines of codes to a vanilla prototypical network trainer. **(3)** We deploy our system and use the dataset statistics to inform out-of-distribution (OOD) cases.

## 2 Related Work

There are three common approaches to meta-learning: metric-based, model-based, and optimization-based. Model agnostic meta-learning (MAML) (Finn et al., 2017) is an optimization-based approach to meta-learning which is agnostic to the model architecture and task specification. Over the years, several variants of the method have shown that it is an ideal candidate for learning to learn from diverse tasks (Nichol et al., 2018; Raghu et al., 2019; Bansal et al., 2020b). However, to solve a new task, MAML type methods would require training a new classification layer for the task. In contrast, metric-based approaches, such as prototypical networks (Vinyals et al., 2016; Snell et al., 2017), being non-parametric in nature can handle varied number of classes and thus can be easily deployed. Given the simple nature of prototypical networks, a lot of work has been

done to improve them (Allen et al., 2019; Zhang et al., 2019; Ding et al., 2022; Wang et al., 2021). Prototypical networks usually construct a class prototype (mean) using the support vectors to describe the class and, given a query example, assigns the class whose class prototype is closest to the query vector. In (Allen et al., 2019), the authors use a mixture of Gaussians to describe the class conditional distribution and in (Zhang et al., 2019); the authors try to model an unknown general class distribution. In (Ding et al., 2022), the authors use spherical Gaussians and a KL-divergence type function between the Gaussians to compute the function  $d$  in equation 2. However, the function used by the above authors is not a true metric, i.e. does not satisfy the triangle inequality. Triangle inequality is implicitly important since we use this metric as a form of distance which we optimize, so it makes sense to use a true metric. In this work we replace it by the Wasserstein distance which is a true metric and add in a regularization term that encourages the  $L_2$  norm of the covariance matrices to be small, encouraging the class examples to be clustered close to the centroid. One of our main reasons to work with Gaussians is due to the closed form formula of the Wasserstein distance.

Few shot learning (FSL) in the medical domain has been mostly focused in computer vision (Singh et al., 2021). There are only a few works that have applied FSL in medical NLP (Ge et al., 2022) but most of those works have only focused on different tasks on MIMIC-III (Johnson et al., 2016) which is a single domain dataset (patients from ICU and one hospital system). To the best of our knowledge, ours is the first study to successfully apply FSL on a diverse set of medical datasets (diverse in terms of tasks and patient populations).

## 3 Datasets

All our internal datasets are MRI radiology reports detailing various pathologies in different body parts. Our models are meta-trained on a dataset of shoulder pathologies which is collected from 74 unique and de-identified institutions in the United States. 60 labels are chosen for training and 20 novel labels are chosen for validation. The number of training labels is similar to some well-known image datasets (Lake et al., 2015b; Vinyals et al., 2016; Wah et al., 2011). This diverse dataset has a rich label space detailing multiple structures in shoulder, granular pathologies and their severity

levels in each structure. The relationship between the granularity/severity of these pathologies at different structures can be leveraged for other pathologies in different body parts and may lead to successful transfer to various downstream tasks. The labels are split such that all pathologies in a given structure appear at either training or validation but not both. More details about the label space can be found in section B. The figure 1 and the table 1 shows the distribution of labels and an example of this dataset can be found in figure 4. Our met-




Figure 1: Histogram showing the label distribution in (left) train and (right) validation dataset.

learner is applied to 4 downstream binary classification tasks spanning different sub-specialities (cancer screening, musculoskeletal radiology, and neuro-radiology) that are both common as well as clinically important. The statistics for each task are given in table 2 : (1) High risk cancer screening for lung nodules (according to Fleischner guidelines (Nair et al., 2018) which bucket patients at high-risk of lung cancer and requiring follow up imaging immediately or within 3 months as belonging to Category High Risk ; we consider patients not at high-risk as Low Risk), (2) Complete Anterior Cruciate Ligament (ACL) tear (Grade 3) vs not Complete ACL tear, (3) Acute ACL tears (MRI examination was performed within 6 weeks of injury) and typified by the presence of diffuse or focal increased signal within the ligament vs not Acute ACL tear (Dimond et al., 1998), (4) Severe vs not severe neural foraminal stenosis in the cervical spine as severe foraminal stenosis may indicate nerve impingement, which is clinically significant. Acute tear in ACL refers to the age of the tear/injury whereas the complete tear refers to the integrity of the ligament. Our testing datasets are diverse and sampled from different institutions: the knee data, lung dataset and cervical dataset is sampled

Split	Number of examples	Min	Max	Average
Train	34595	79	6379	567
Validation	5754	44	1138	303

Table 1: Statistics of our meta-training and meta-validation dataset, where the min/max/average refer to min/max/average examples per label.

from 50, 4 and 65 institutions respectively and our annotation process is described in Appendix A. Examples of these datasets can be found in figure 10 (knee), figure 6 (lung), and figure 8 (cervical).

Task	Validation Distribution	Test Distribution
Lung Nodule	Low Risk : 233	Low Risk : 347
	High Risk : 30	High Risk : 46
Knee ACL	Normal: 258	Normal : 439
Acute Tear	Acute Tear: 48	Acute Tear: 93
Knee ACL	Normal : 263	Normal : 429
Complete Tear	Complete Tear : 44	Complete Tear :103
Neural Foraminal	Normal : 215	Normal : 789
Stenosis	Abnormal : 43	Abnormal : 91

Table 2: Statistics of our downstream testing datasets

## 4 Workflow

Our workflow consists of the following parts: A




Figure 2: Overview of our workflow. A report is passed through a report segmenter which splits it into sentences and extracts the relevant portion of the text for downstream classification. The relevant text is passed through our model and we use the pre-computed prototypes and class variances to assign a label to the query point.

report is first de-identified according to HIPAA regulations and passed through a sentence parser (ex. Spacy (Honnibal et al., 2020)) that splits the report into sentences. In the shoulder dataset, each of these sentences is labeled with the appropriate structure and severity label and we filter out sentences that do not have such a label. We first train a meta-learner in an episodic fashion on this dataset and choose the best model based on meta-validation accuracy.

For our downstream tasks, we use a body-part specific custom data processor to collect sentences related to a given structure (ACL in knee, different vertebrae in the cervical spine, the entire impression section for lung reports) and concatenate them

together to create a paragraph describing all the pathologies in the structure of interest. Detailed description of preprocessing for different body parts, is presented in Appendix C. The concatenated text from the validation sets of each task is passed to our trained meta-learner to generate the relevant class statistics (mean and variance). We then perform pathology classification on the test set by using our trained meta-learner and the saved class statistics. The downstream tasks are similar to the shoulder task in the sense that the pathology classification is performed on a sequence of sentences that all pertain to the same anatomical structure. Thus our approach needs to learn the language that describes the severity of the pathology for a specific anatomical structure.

We would like to shed some light on the complexity of the language we encounter. Since our dataset is sourced from multiple health systems, and not all reports follow a standard structure, there is a large amount of variation in the language describing the same diagnosis. For example: a severe tear can be referred to as a rupture, or only the size of the nodule is mentioned without specifying that it is low risk (see Appendix C for more examples). Furthermore, most of our pipelines attempt to classify the different severities for a given pathology and the language describing severity can vary. While it might be possible to construct a rule based system to extract the diagnoses and severities we are interested in, it will be difficult to generalize as we expand to more diagnoses as well as to new health systems.

## 5 Prototypical Networks

Prototypical Networks or ProtoNets (Snell et al., 2017) use an embedding function  $f_\theta$  to encode each input into a  $M$ -dimensional feature vector. A prototype is defined for every class  $c \in \mathcal{L}$ , as the mean of the set of embedded support data samples ( $S_c$ ) for the given class, i.e.

$$v_c = \frac{1}{|S_c|} \sum_{(x_i, y_i) \in S_c} f_\theta(x_i). \quad (1)$$

The distribution over classes for a given test input  $x$  is a softmax over the inverse of distances between the test data embedding and prototype vectors.

$$\begin{aligned} P(y = c|x) &= \text{softmax}(-d(f_\theta(x), v_c)) \\ &= \frac{\exp(-d(f_\theta(x), v_c))}{\sum_{c' \in \mathcal{L}} \exp(-d(f_\theta(x), v_{c'}))} \end{aligned} \quad (2)$$

where  $d$  can be any (differentiable) distance function. The loss function is negative log-likelihood:

$$\mathbb{L}(\theta) = -\log P_\theta(y = c|x).$$

ProtoNets are simple and easy to train and deploy. The mean is used to capture the entire conditional distribution  $P(y = c|x)$ , thus losing a lot of information about the underlying distribution. A lot of work (Ding et al., 2022; Allen et al., 2019; Zhang et al., 2019) has focused on improving ProtoNets by taking into account the above observation. We extend ProtoNets by incorporating the variance (2nd moment) of the distribution and use distributional distance, i.e. 2-Wasserstein metric, directly generalizing the vanilla ProtoNets.

### 5.1 Variance Aware ProtoNets

In this work, we model each conditional distribution as a Gaussian. Now the main question is: *how do we match a query example with a distribution?* The simplest thing here is to treat the query example as a Dirac distribution. With that formulation in mind, recall: the Wasserstein-Bures metric between Gaussians ( $m_i, \Sigma_i$ ) is given by:

$$d^2 = \|m_1 - m_2\|^2 + \text{Tr}(\Sigma_1 + \Sigma_2 - 2(\Sigma_1^{\frac{1}{2}} \Sigma_2 \Sigma_1^{\frac{1}{2}})^{\frac{1}{2}})$$

Given  $(x_i, y_i) \in S_c$ , where  $S_c$  is the support set of examples belonging to class  $c$ , we compute the mean  $m_c$  and covariance matrix  $\Sigma_c$ ; the computation of Wasserstein distance between a Gaussian and a query vector  $q$  (i.e. a Dirac) boils down to

$$d^2 = \|m_c - q\|^2 + \text{Tr}(\Sigma_c) \quad (3)$$

The above formula shows that we can simplify our conditional distribution to be a Gaussian with a *diagonal* covariance matrix. This brings down our space complexity to store this covariance matrix from  $O(n^2)$  to  $O(n)$ . Note, this is a direct generalization of the vanilla prototypical networks as the vanilla prototypical networks can be interpreted as computing the Wasserstein distance (aka simple  $L_2$  distance) between two Dirac distributions (mean of the conditional distribution and the query sample). We also propose another variant of the above called Isotropic Gaussian variant where we average over the diagonal entries of  $\Sigma_c$ , i.e.  $\alpha = \frac{1}{n}(\Sigma_c)_{ii}$  and redefine  $\Sigma_c = \alpha I$ , where  $I$  is the identity matrix, allowing us to just store the scalar  $\alpha$ , further reducing the space complexity. Furthermore, we

Backbone	Methods	Foraminal	Knee (Acute Tear vs Not)	Knee (Complete tear vs Not)	Lung
PubMedBERT	Baseline	0.38	0.44	0.49	0.36
	Multi-Task	0.41	0.47	0.52	0.39
	Vanilla ProtoNet	0.79	0.73	0.60	0.68
	Big ProtoNet	0.58	0.59	0.51	0.64
	Leopard	0.84	0.78	0.80	0.74
	ProtoNet w/ Isotropic Gaussian	0.81	0.74	0.76	0.69
	ProtoNet w/ Isotropic Gaussian + reg	0.83	0.76	0.77	0.73
	Variance Aware ProtoNet (ours)	0.84	0.78	0.79	0.76
	Variance Aware ProtoNet + reg (ours)	0.86	0.81	0.84	<b>0.80</b>
	PubMedBERT w/ Adapters	Baseline	0.42	0.47	0.51
Multi-Task		0.44	0.49	0.53	0.43
Vanilla ProtoNet		0.78	0.71	0.69	0.66
Big ProtoNet		0.59	0.57	0.54	0.67
ProtoNet w/ Isotropic Gaussian		0.83	0.75	0.78	0.72
ProtoNet w/ Isotropic Gaussian + reg		0.89	0.80	0.86	0.77
Variance Aware ProtoNet (ours)		0.87	0.77	0.81	0.74
Variance Aware ProtoNet + reg (ours)		<b>0.91</b>	<b>0.82</b>	<b>0.89</b>	0.78

Table 3: Table showing F1 scores of Few Shot Models in downstream classification tasks.

regularize the negative log likelihood loss to prevent the variance term from blowing up. Our new loss function reads:

$$\mathcal{L}(\theta) = \mathbb{L}(\theta) + \frac{\lambda}{\text{ways}} \|\Sigma_c\|_F \quad (4)$$

where ways are the number of classes in the mini-batch and  $\|\cdot\|_F$  is the Frobenius norm and we average the norm of the variance matrix over all the classes in a given meta-batch. The extra regularization term is designed to encourage the examples to be close to the appropriate cluster centroid. This term can also be seen as an entropic regularization term, i.e. up to a factor as the exponential of  $KL(p||q)$ , where  $p = N(m_c, \Sigma_c)$  and  $q = N(m_c, I)$ . This type of entropy regularized Wasserstein distances is widely studied (Cuturi and Doucet, 2014; Altschuler et al., 2021).

A PyTorch style pseudocode is described in Algorithm 1, where the teal color refers to the changes to a vanilla prototypical networks trainer. We provide detailed motivation for using Wasserstein distance instead of KL divergence in section E.2. This also explains why we compute the Wasserstein distance between the query and the estimated class distribution instead of a simple likelihood.

## 6 Experiments

All our experiments are run on 4 V100 16 GB GPU using PyTorch (Paszke et al., 2019) and HuggingFace libraries (Wolf et al., 2020). Bert-base (Devlin et al., 2019), Clinical BERT (Alsentzer et al., 2019) and PubMedBERT (Gu et al., 2021) are used as our backbone models. Adapters (Pfeiffer et al., 2020) are applied to each of these backbone models.

While training adapter based models, the BERT weights are frozen and only the adapter weights are updated, thus requiring less resources to train. This idea is similar to (Raghu et al., 2019) in the sense that we are reusing the features from these deep pre-trained models. We compare our methods to Leopard (Bansal et al., 2020a), vanilla ProtoNets and big ProtoNets (Ding et al., 2022). Additional results with BERT-base and Clinical BERT backbones can be found in table 6 and table 7. Meta-training is done in an episodic manner using 4-way 8-shot and 16-examples as support. For meta-training on the shoulder dataset, we set the variance regularizer hyperparameter to be .1. It is an important hyperparameter and detailed ablation study is conducted in section E.1. Other hyperparameters and design choices are described in section E.

To prevent overfitting on the test set, we choose the best model from each of these experiments based on the meta-validation accuracy and apply it to our downstream classification tasks. We note that these downstream tasks are significantly different from the few shot regime these models are trained in. Moreover for these downstream tasks, we train BERT models on each task and a multi-tasking model to provide additional baselines.

In all our experiments, PubMedBERT consistently outperforms BERT-base and Clinical BERT by an average of 5 points and 3 points respectively. We believe the reason behind the improved performance is the domain specific vocabulary. Even though Clinical BERT is pre-trained on MIMIC-III (Johnson et al., 2016), it still shares the same vocabulary as BERT-base.

ProtoNet-BERT shows better performance and

faster convergence rates during training and validation (Table 4), but it is outperformed by ProtoNet-AdapterBERT which has fewer orders of magnitude of parameters to learn (Table 3). Like (Wang et al., 2021) we believe that ProtoNet-BERT is more vulnerable to overfitting on the meta-training tasks than the ProtoNet-AdapterBERT. Finally, we note that even though Big ProtoNets work well on meta-validation, they fail on our downstream tasks. We hypothesize that it is due to the fact that big protonets are encouraged to have large radii which has the potential to become a bottleneck where the data distribution is highly imbalanced causing the spherical Gaussians to overlap. In fact, we have found that doing the exact opposite (i.e. constricting the norms of the covariance matrix), tends to produce better results. Finally instead of using

Backbone	Methods	Accuracy
PubMedBERT	Vanilla ProtoNet	89.1 ± 1.1
	Big ProtoNet	90.8 ± 1.2
	Leopard	85.1 ± 9.2
	ProtoNet w/ Isotropic Gaussian	90.2 ± 1.4
	ProtoNet w/ Isotropic Gaussian + reg	92.1 ± .8
	Variance Aware ProtoNet (ours)	91.5 ± 1.3
	Variance Aware ProtoNet + reg (ours)	<b>92.9 ± .9</b>
PubMedBERT w/Adapters	Vanilla ProtoNet	88.3 ± 1.4
	Big ProtoNets	89.4 ± 1.2
	ProtoNet w/ Isotropic Gaussian	89.8 ± 1.4
	ProtoNet w/ Isotropic Gaussian + reg	90.9 ± .7
	Variance Aware ProtoNet (ours)	90.5 ± 1.3
	Variance Aware ProtoNet + reg (ours)	91.2 ± .8

Table 4: Results showing accuracy percentages on the meta-validation dataset. We sampled 1000 tasks with 4-way 8-shot and 16-support classification. We replicate each experiment over 10 random seeds.

the entire validation set to compute the class distribution, we also experiment with choosing a  $k$  shots from the validation set to compute the class distribution (figure 12 in section G).

Our regularized Variance Aware ProtoNets with BERT-base + Adapter is also validated on 13 public datasets. For the models and datasets marked with \* in table 5, we use the results reported in (Bansal et al., 2020a) and for those datasets, we use the code from (Wang et al., 2021) to generate the results for ProtoNet with Bottleneck Adapters while the rest of the results are taken from (Wang et al., 2021). The variance regularization hyperparameter is set to .01 for these experiments. Our method beats Leopard by 5, 3 and 2 points on 4, 8 and 16 shots, respectively. The training details for these experiments can be found in section F.

## 7 Deployment

Based on the results described in table 3, we choose to deploy our regularized Variance Aware ProtoNet with Adapter-PubMedBERT. Our pipeline is deployed on AWS using a single p3.2x instance housed with one NVIDIA V100 GPU. The main pipeline components include (1) body-part specific report segmenter, (2) PubMedBERT backbone with adapters and (3) a dictionary of class prototypes and class variances, for all classes in the datasets. On inference, requests sent to the pipeline include a body part which the pipeline utilizes to load up the relevant report segmenter, class prototypes and variances. A report is then ingested by the pipeline, parsed by a sentencizer, grouped into segments according to its body part specific segmentation, and then passed to the model. Class probabilities and labels are inferred after computing the Wasserstein distance between the text embedding and the appropriate class distributions. These outputs and pipeline metadata are written out to an AWS Redshift database cluster. The entire pipeline is orchestrated in batch mode with a large enough batch size to maximize GPU capacity resulting in an average latency of 68ms/report.

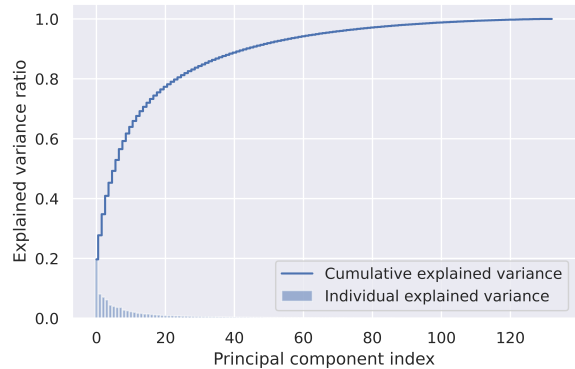


Figure 3: Variance along different directions for the Lung validation set

### 7.1 Monitoring

It is well-known that the BERT embeddings are highly anisotropic (Ethayarajh, 2019). We observe the same phenomenon in our meta-learned models as well (figure 3) which we use to our advantage to monitor OOD cases. For each class in a dataset, we pick top  $k$ -dimensions (a hyperparameter) of maximum variance. We then take the union of these indices that we call the set of dataset indices i.e. the indices that explain the variance among all classes

Shots	Dataset	BERT*	MT-BERT*	Leopard	ProtoNet	ProtoNet+Adapter	Variance Aware ProtoNet (Ours)
4	airline	42.76 ± 13.50	46.29 ± 12.26	54.95 ± 11.81	<b>65.39 ± 12.73</b>	65.33 ± 7.95	62.67 ± 11.18
	disaster	<b>55.73 ± 10.29</b>	50.61 ± 8.33	51.45 ± 4.25	54.01 ± 2.9	53.48 ± 4.76	53.89 ± 3.79
	emotion	9.20 ± 3.22	9.84 ± 2.14	11.71 ± 2.16	11.69 ± 1.87	12.52 ± 1.32	<b>15.15 ± 4.19</b>
	political_audience	51.89 ± 1.72	51.53 ± 1.80	52.60 ± 3.51	<b>52.77 ± 5.86</b>	51.88 ± 6.37	52.5 ± 6.45
	sentiment_kitchen*	56.93 ± 7.10	60.53 ± 9.25	78.35 ± 18.36	62.71 ± 9.53	83.13 ± 0.96	<b>84.16 ± 1.37</b>
	political_bias	54.57 ± 5.02	54.66 ± 3.74	60.49 ± 6.66	58.26 ± 10.42	<b>61.72 ± 5.65</b>	59.39 ± 6.18
	rating_electronics*	39.27 ± 10.15	41.20 ± 10.69	51.71 ± 7.20	37.40 ± 3.72	53.81 ± 6.01	<b>55.49 ± 5.42</b>
	political_message	15.64 ± 2.73	14.49 ± 1.75	15.69 ± 1.57	17.82 ± 1.33	<b>20.98 ± 1.69</b>	19.28 ± .91
	sentiment_books*	54.81 ± 3.75	64.93 ± 8.65	82.54 ± 1.33	73.15 ± 5.85	83.88 ± 0.55	<b>84.95 ± 1.72</b>
	rating_books*	39.42 ± 07.22	38.97 ± 13.27	48.44 ± 7.43	54.92 ± 6.18	59.20 ± 7.26	<b>66.18 ± 7.89</b>
	rating_dvd*	32.22 ± 08.72	41.23 ± 10.98	49.76 ± 9.80	47.73 ± 6.20	50.20 ± 10.26	<b>52.59 ± 14.09</b>
	rating_kitchen	34.76 ± 11.2	36.77 ± 10.62	50.21 ± 9.63	58.47 ± 11.12	55.99 ± 9.85	<b>59.39 ± 8.79</b>
	scitail*	58.53 ± 09.74	63.97 ± 14.36	69.50 ± 9.56	76.27 ± 4.26	77.84 ± 2.61	<b>79.16 ± 2.54</b>
	Average	41.98	44.23	52.11	51.58	56.15	<b>57.29</b>
8	airline	38.00 ± 17.06	49.81 ± 10.86	61.44 ± 3.90	69.14 ± 4.84	<b>69.37 ± 2.46</b>	69.31 ± 2.43
	disaster	<b>56.31 ± 9.57</b>	54.93 ± 7.88	55.96 ± 3.58	54.48 ± 3.17	53.85 ± 3.03	55.19 ± 2.77
	emotion	8.21 ± 2.12	11.21 ± 2.11	12.90 ± 1.63	13.10 ± 2.64	13.87 ± 1.82	<b>15.1 ± 3.58</b>
	political_audience	52.80 ± 2.72	54.34 ± 2.88	54.31 ± 3.95	<b>55.17 ± 4.28</b>	53.08 ± 6.08	53.82 ± 4.13
	sentiment_kitchen*	57.13 ± 6.60	69.66 ± 8.05	<b>84.88 ± 1.12</b>	70.19 ± 6.42	83.48 ± 0.44	84.69 ± .8
	political_bias	56.15 ± 3.75	54.79 ± 4.19	61.74 ± 6.73	63.22 ± 1.96	<b>65.36 ± 2.03</b>	64.09 ± .58
	rating_electronics*	28.74 ± 08.22	45.41 ± 09.49	54.78 ± 6.48	43.64 ± 7.31	56.97 ± 3.19	<b>60.24 ± 2.62</b>
	political_message	13.38 ± 1.74	15.24 ± 2.81	18.02 ± 2.32	20.40 ± 1.12	<b>21.64 ± 1.72</b>	20.44 ± 1.17
	sentiment_books*	53.54 ± 5.17	67.38 ± 9.78	83.03 ± 1.28	75.46 ± 6.87	83.9 ± 0.39	<b>84.68 ± .85</b>
	rating_books*	39.55 ± 10.01	46.77 ± 14.12	59.16 ± 4.13	52.13 ± 4.79	61.74 ± 6.83	<b>65.54 ± 6.78</b>
	rating_dvd*	36.35 ± 12.50	45.24 ± 9.76	53.28 ± 4.66	47.11 ± 4.00	53.25 ± 7.47	<b>53.83 ± 10.46</b>
	rating_kitchen	34.49 ± 8.72	47.98 ± 9.73	53.72 ± 10.31	<b>57.08 ± 11.54</b>	56.27 ± 10.70	56.68 ± 11.21
	scitail*	57.93 ± 10.70	68.24 ± 10.33	75.00 ± 2.42	78.27 ± 0.98	80.41 ± 1.05	<b>80.57 ± .48</b>
	Average	40.97	48.54	56.02	53.8	57.94	<b>58.78</b>
16	airline	58.01 ± 8.23	57.25 ± 9.90	62.15 ± 5.56	<b>71.06 ± 1.60</b>	69.83 ± 1.80	69.9 ± 1.06
	disaster	<b>64.52 ± 8.93</b>	60.70 ± 6.05	61.32 ± 2.83	55.30 ± 2.68	57.38 ± 5.25	60.14 ± 5.36
	emotion	13.43 ± 2.51	12.75 ± 2.04	13.38 ± 2.20	12.81 ± 1.21	<b>14.11 ± 1.12</b>	13.55 ± 3.51
	political_audience	<b>58.45 ± 4.98</b>	55.14 ± 4.57	57.71 ± 3.52	56.16 ± 2.81	57.23 ± 2.77	56.36 ± 2.29
	sentiment_kitchen*	68.88 ± 3.39	77.37 ± 6.74	<b>85.27 ± 01.31</b>	71.83 ± 5.94	83.72 ± 0.30	84.93 ± .49
	political_bias	60.96 ± 4.25	60.30 ± 3.26	65.08 ± 2.14	61.98 ± 6.89	<b>65.38 ± 1.71</b>	63.97 ± 2.49
	rating_electronics*	45.48 ± 06.13	47.29 ± 10.55	58.69 ± 2.41	44.83 ± 5.96	56.62 ± 5.62	<b>61.01 ± 1.54</b>
	political_message	20.67 ± 3.89	19.20 ± 2.20	18.07 ± 2.41	21.36 ± 0.86	<b>24.00 ± 1.39</b>	22.49 ± 1.31
	sentiment_books*	65.56 ± 4.12	69.65 ± 8.94	83.33 ± 0.79	77.26 ± 3.27	83.92 ± 0.41	<b>84.91 ± 0.66</b>
	rating_books*	43.08 ± 11.78	51.68 ± 11.27	61.02 ± 4.19	57.28 ± 4.57	64.75 ± 4.27	<b>67.34 ± 7.52</b>
	rating_dvd*	42.79 ± 10.18	45.19 ± 11.56	53.52 ± 4.77	48.39 ± 3.74	55.08 ± 4.92	<b>56.63 ± 6.11</b>
	rating_kitchen	47.94 ± 8.28	53.79 ± 9.47	57.00 ± 8.69	<b>61.00 ± 9.17</b>	59.45 ± 8.33	58.34 ± 11.72
	scitail*	65.66 ± 06.82	75.35 ± 04.80	77.03 ± 1.82	78.59 ± 0.48	80.27 ± .75	<b>80.89 ± .23</b>
	Average	50.42	52.74	57.97	55.22	59.36	<b>60.04</b>

Table 5: Results on some benchmark text datasets on a wide range of tasks from NLI, sentiment analysis and text classification. For the Variance Aware ProtoNet, we use BERT-base with bottleneck Adapters. For meta-training, WNLI (m/mm), SST-2, QQP, RTE, MRPC, QNLI, and SNLI datasets are used.

in the dataset. For any given query example, we compute the absolute difference ( $\vec{d}_j$ ) between its embedding vector ( $\vec{q}$ ) and class centroids ( $\vec{v}_j$ ), i.e. the  $i$ -th coordinate  $\vec{d}_j : \vec{d}_{j_i} = |\vec{q}_i - \vec{v}_{j_i}|$ . We then select top  $k$  dimensions of the each of these  $\vec{d}_j$ . We propose an OOD metric called Average Variance Indices (AVI\_k) by the overlap between the top-k difference vector indices and the top-k dataset indices, i.e.  $AVI_k := \frac{|\cup_{j=1}^c \text{top-k}(\vec{d}_j)|}{\text{dataset indices}}$ , where  $c$  is the number of classes. For ex: in case of the lung dataset: The text "*The heart is normal in size. There is no pericardial effusion. The pulmonary artery is enlarged.*" shows an AVI\_10 score .79, whereas "*L1L2: There is no disc herniation in lumbar spine.*" gives a score of .31. As part of our monitoring, we threshold reports with an AVI\_10 < .5 to further investigate if the report is OOD.

## 8 Conclusion

We extend Prototypical Networks by using Wasserstein distances instead of Euclidean distances and introduce a regularization term to encourage the class examples to be clustered close to the class prototype. By training our models on a label rich dataset (shoulder MRI reports), we show successful performance on a variety of tasks. Since the model weights are reused for all tasks, a single model is deployed enabling us to cut inference costs. Moreover, adapters are used allowing us to tune smaller number of parameters ( $\sim 10$  million) resulting in huge training cost savings. Our model is also benchmarked on 13 public datasets and outperforms strong baselines like Leopard. Current work is underway to make our training dataset more diverse so that our models are more generalizable.

## Ethical Considerations

Due to various legal and institutional concerns arising from the sensitivity of clinical data, it is difficult for researchers to gain access to relevant data except for MIMIC (Johnson et al., 2016). Despite its large size (covering over 58k hospital admissions), it is only representative of patients from a specific clinical domain (the intensive care unit) and geographic location (a single hospital in the United States). We can not expect such a sample to be representative of either the larger population of patient admissions or other geographical regions/hospital systems. We have tried to address this partially by collecting radiology data for various body parts across multiple practices in the US. However we are always mindful that our work may not generalize to new body parts/pathologies and radiology practices (see Section H). Even though we introduce a simple OOD metric, we realize it is far from perfect. We understand the need to minimize ethical risks of AI implementation like threats to privacy and confidentiality, informed consent, and patient autonomy. And thus we strongly believe that stakeholders should be flexible in incorporating AI technology as a complementary tool and not a replacement for a physician. Thus, we develop our workflows, annotation guidelines and generate actionable insights by working in conjunction with a varied group of radiologists and medical professionals to minimize these above risks. Finally our pipeline as deployed is meant as a pseudo-labeling tool which we expect would cut down on expensive annotation costs but can potentially introduce some bias in our pseudo-labels.

## References

- Kelsey Allen, Evan Shelhamer, Hanul Shin, and Joshua Tenenbaum. 2019. [Infinite mixture prototypes for few-shot learning](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 232–241. PMLR.
- Emily Alsentzer, John Murphy, William Boag, Weihung Weng, Di Jindi, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Jason Altschuler, Sinho Chewi, Patrik Robert Gerber, and Austin J Stromme. 2021. [Averaging on the bures-wasserstein manifold: dimension-free convergence of gradient descent](#). In *Advances in Neural Information Processing Systems*.
- Trapit Bansal, Rishikesh Jha, and Andrew McCallum. 2020a. [Learning to few-shot learn across diverse natural language classification tasks](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5108–5123, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Trapit Bansal, Rishikesh Jha, Tsendsuren Munkhdalai, and Andrew McCallum. 2020b. [Self-supervised meta-learning for few-shot natural language classification tasks](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 522–534.
- Samy Bengio, Yoshua Bengio, Jocelyn Cloutier, and Jan Gecsei. 1997. [On the optimization of a synaptic learning rule](#).
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 632–642, Lisbon, Portugal. Association for Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Marco Cuturi and Arnaud Doucet. 2014. [Fast computation of wasserstein barycenters](#). In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 685–693, Beijing, China. PMLR.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- P Dimond, Paul Fadale, Michael Hulstyn, Glenn Tung, and J Greisberg. 1998. [A comparison of mri findings in patients with acute and chronic acl tears](#). *The American journal of knee surgery*, 11:153–9.



- Ning Ding, Yulin Chen, Xiaobin Wang, Hai-Tao Zheng, Zhiyuan Liu, and Pengjun Xie. 2022. [Few-shot learning with big prototypes](#).
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. [Model-agnostic meta-learning for fast adaptation of deep networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1126–1135. PMLR.
- Yao Ge, Yuting Guo, Yuan-Chi Yang, Mohammed Ali Al-Garadi, and Abeer Sarker. 2022. [Few-shot learning for medical text: A systematic review](#).
- Ruiying Geng, Binhua Li, Yongbin Li, Xiaodan Zhu, Ping Jian, and Jian Sun. 2019. [Induction networks for few-shot text classification](#).
- Yu Gu, Robert Tinn, Hao Cheng, Michael Lucas, Naoto Usuyama, Xiaodong Liu, Tristan Naumann, Jianfeng Gao, and Hoifung Poon. 2021. [Domain-specific language model pretraining for biomedical natural language processing](#). *ACM Trans. Comput. Healthcare*, 3(1).
- Junxian He, Chunting Zhou, Xuezhe Ma, Taylor Berg-Kirkpatrick, and Graham Neubig. 2022. [Towards a unified view of parameter-efficient transfer learning](#). In *International Conference on Learning Representations*.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Timothy Hospedales, Antreas Antoniou, Paul Micaelli, and Amos Storkey. 2020. [Meta-learning in neural networks: A survey](#).
- Alistair E.W. Johnson, Tom J. Pollard, Lu Shen, Liwei H. Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G. Mark. 2016. [Mimic-iii, a freely accessible critical care database](#). *Scientific Data*, 3(1):160035.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021. [Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks](#). In *Annual Meeting of the Association for Computational Linguistics*.
- Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. 2015a. [Human-level concept learning through probabilistic program induction](#). *Science*, 350(6266):1332–1338.
- Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. 2015b. [Human-level concept learning through probabilistic program induction](#). *Science*, 350(6266):1332–1338.
- Arjun Nair, Anand Devaraj, Matthew E J Callister, and David R Baldwin. 2018. [The fleischner society 2017 and british thoracic society 2015 guidelines for managing pulmonary nodules: keep calm and carry on](#). *Thorax*, 73(9):806–812.
- Alex Nichol, Joshua Achiam, and John Schulman. 2018. [On first-order meta-learning algorithms](#).
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc.
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. 2020. [AdapterHub: A framework for adapting transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 46–54, Online. Association for Computational Linguistics.
- Aniruddh Raghu, Maithra Raghu, Samy Bengio, and Oriol Vinyals. 2019. [Rapid learning or feature reuse? towards understanding the effectiveness of maml](#).
- Rebecca Roelofs, Vaishaal Shankar, Benjamin Recht, Sara Fridovich-Keil, Moritz Hardt, John Miller, and Ludwig Schmidt. 2019. [A meta-analysis of overfitting in machine learning](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Jürgen Schmidhuber. 1987. [Evolutionary principles in self-referential learning. on learning now to learn: The meta-meta-meta...-hook](#). Diploma thesis, Technische Universität München, Germany, 14 May.
- Arijit Sehanobish, McCullen Sandora, Nabila Abraham, Jayashri Pawar, Danielle Torres, Anasuya Das, Murray Becker, Richard Herzog, Benjamin Odry, and Ron Vianu. 2022. [Explaining the effectiveness of multi-task learning for efficient knowledge extraction from spine MRI reports](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Industry Track*, pages 130–140, Hybrid: Seattle, Washington + Online. Association for Computational Linguistics.

Rishav Singh, Vandana Bharti, Vishal Purohit, Abhinav Kumar, Amit Kumar Singh, and Sanjay Kumar Singh. 2021. [Metamed: Few-shot medical image classification using gradient-based meta-learning](#). *Pattern Recognition*, 120:108111.

Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 4080–4090, Red Hook, NY, USA. Curran Associates Inc.

Erik E. Swartz, R. T. Floyd, and Mike Cendoma. 2005. [Cervical Spine Functional Anatomy and the Biomechanics of Injury due to Compressive Loading](#). *Journal of athletic training*, 40(3):155–161.

Oriol Vinyals, Charles Blundell, Timothy Lillicrap, koray kavukcuoglu, and Daan Wierstra. 2016. [Matching networks for one shot learning](#). In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.

C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. 2011. Caltech-ucsd birds-200-2011. Technical Report CNS-TR-2011-001, California Institute of Technology.

Jixuan Wang, Kuan-Chieh Wang, Frank Rudzicz, and Michael Brudno. 2021. [Grad2task: Improved few-shot text classification using gradients for task representation](#). In *Advances in Neural Information Processing Systems*.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Dani Yogatama, Cyprien de Masson d’Autume, Jerome Connor, Tomas Kocisky, Mike Chrzanowski, Lingpeng Kong, Angeliki Lazaridou, Wang Ling, Lei Yu, Chris Dyer, and Phil Blunsom. 2019. [Learning and evaluating general linguistic intelligence](#).

Jian Zhang, Chenglong Zhao, Bingbing Ni, Minghao Xu, and Xiaokang Yang. 2019. [Variational few-shot learning](#). In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 1685–1694.

## A Annotation

First we collect data from various sources and a part of the data are annotated by our team of in-house expert annotators with deep clinical expertise, which we use as test and development sets for

our model training. We then use this annotated data to train a larger pool of other annotators who are generally medical students. They are provided clear guidelines on the task and performance is measured periodically on a benchmark set and feedback is provided. As of the writing of the manuscript, the validation and the test sets as described in section 3 are being used to train the annotators. After the completion of their training, the annotators will annotate the remaining unlabeled data that will be used as a training data for our models. The entire process is slow but is designed to generate high quality annotated data. We believe that our few shot models can be used as a source of pseudo-labels and will greatly simplify and quicken our annotation process.

## B Shoulder Dataset

In this section we will briefly describe our label rich shoulder dataset that is used as meta-training and meta-validation sets. There are 80 labels for the shoulder dataset. They range from Clinical history, metadata, Impressions, Finding to various granular pathologies at different structures in the shoulder like AC joint, Rotator Cuff, Muscles, Bursal Fluid, Supraspinatus, Infraspinatus, Subscapularis, Labrum, Glenohumeral Joint, Humeral Head, Acromial Morphology, Impingement: AC Joint. The labels are split such that all pathologies in a given structure appear at either training or validation but not both. We believe that such a split would help a model to learn the key words that may describe the granularity of a pathology in a given structure of interest. The dataset level statistics can be found in figure 1 and table 1. An example of the shoulder data is shown in figure 4.

Text	Labels
The AC joint and anterior acromion show evidence of prior subacromial depression and there may have been a distal clavicle excision as well with widening of the AC fluid in the joint glenohumeral joint/ labrum.	AC Joint: Mild Arthritis with Edema
Type II acromion with hypertrophic changes causing impingement and partial rotator cuff tear of the infraspinatus and supraspinatus myotendinous junction.	Impingement: Acromion
Mild subacromial-subdeltoid bursitis. Findings are age-indeterminate unless otherwise specified.	Bursal Fluid: Small
Acromioclavicular joint: Anatomic alignment. No substantial degenerative change.	AC Joint: Normal
There is fraying of the anterior labrum above the level of the equator.	Labrum: Normal or mild degeneration

Figure 4: Figure showing an example of our shoulder dataset which is used for meta-training. Note that the labels attached to the text have information about the location and severity of a given pathology.

## C Detailed Workflow

We now present a detailed description of various body part specific workflows. All reports, irrespective of body part, are first de-identified according to HIPAA regulations. We then pass the report through a sentence parser to parse the report in sentences.

### C.1 Lung Dataset

For the lung dataset, we use a report segmenter which is a rule-based regex to extract the "Impression" section from the entire report. This section can be thought as the summary of the report and contains all the critical information like number of lung nodules and their sizes and potential for malignancy. This section text is used for final classification task as shown in figure 5. Figure 6 shows examples of the labels in the dataset.

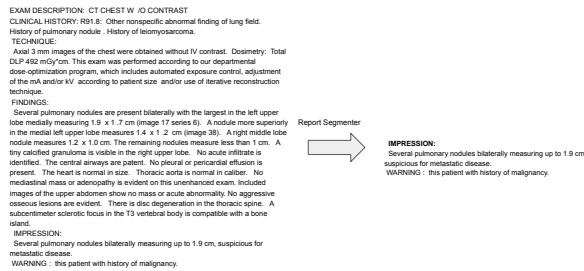


Figure 5: Figure showing the preprocessing of the lung dataset. Our report segmenter extracts the relevant paragraph which is used for downstream classification.

Text	Labels
Dominant 6.7 cm right upper lobe mass, with contiguous extension into the right hilar region, probable adjacent interstitial spread, bilateral pulmonary metastatic nodules, mediastinal lymphadenopathy. 3 cm left adrenal gland lesion suspicious for metastasis.	High Risk
Mild interstitial and ground glass density in the right upper lobe near the apex and probably representing sequelae of radiation treatment. A tiny benign-appearing pulmonary nodule in the right middle lobe is unchanged compared to 2016. No findings suspicious for metastatic disease in the chest. Hepatic steatosis.	Low Risk

Figure 6: Figure showing the labels in the Lung dataset

### C.2 Cervical Dataset

Our task in the cervical dataset is to predict the severity of a neural foraminal stenosis for each motion segment - the smallest physiological motion unit of the spinal cord (Swartz et al., 2005). Breaking information down at the motion segment level in this way enables pathological findings to be correlated with clinical exam findings, and can inform future treatment interventions. A BERT based NER model is used to identify the motion segment(s) referenced in each sentence, and all the sentences containing a particular motion segment

are concatenated together. We also use additional rule-based logic to assign motion segments to relevant sentences that may not mention a motion segment in it. We then predict the disease severity using this concatenated text at *each* motion segment. This data pre-processing mostly follows the ideas and the steps outlined in (Sehanobish et al., 2022). Figure 7 shows our preprocessing steps and figure 8 shows examples in the datasets.

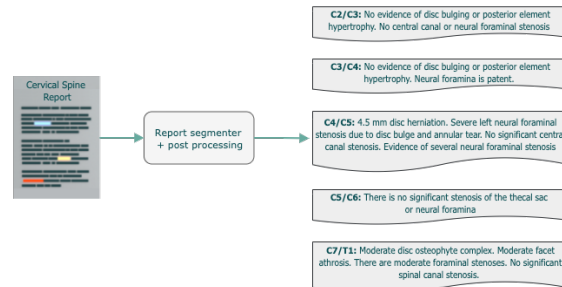


Figure 7: Figure showing the preprocessing of the cervical dataset. Our report segmenter extracts all the motion segments mentioned in the report and groups all sentences belonging to the relevant motion segment. The paragraph belonging to a given motion segment is used for downstream classification.

Text	Label
C3-4: Desiccation mild disc space loss. There is mild disc osteophyte bulge. Uncovertebral hypertrophy left greater than right. Mild facet hypertrophy. In conjunction with short pedicles mild central canal narrowing. There is mild to moderate left and mild right foraminal stenosis.	0
At C3-C4, there is a degenerated bulging disc with osteophytic ridging, facet and uncovertebral hypertrophy and short pedicles combining to cause mild to moderate central stenosis and severe bilateral foraminal stenosis unchanged. There is mild motion artifact on some of the images. There is reversal of the cervical alignment with grade 1 anterolisthesis at C3-4 and C5-6.	1

Figure 8: Figure showing the labels in the Cervical dataset. 0 means absence of severe neural foraminal stenosis and 1 indicates presence of severe neural foraminal stenosis.

### C.3 Knee Dataset

The data processing steps for the knee dataset is similar to the cervical dataset. A BERT based NER model is used to tag sentences that mention the structure of importance, i.e. the anterior cruciate ligament (ACL). We group all the sentences together that mention ACL and we use these grouped sentences to predict our pathology severity as shown in the workflow (figure 9). An example of the labels in the knee dataset can be found in figure 10.

## D Additional Experiments

We also experiment with BERT-base and Clinical BERT as additional backbones. We add adapters




Figure 9: Figure showing the preprocessing of the knee dataset. Our report segmenter selects all the relevant sentences pertaining to the structure of interest, i.e. ACL. We then predict various pathology severities using this paragraph of text.

Text	Acute Tear	Complete Tear
The anterior cruciate ligament is intact and there is a partial tear of the posterior cruciate ligament with a thin residual component of the distal half of the PCL still intact.	0	0
While there is some edema in the ACL, there appear to be intact fibers. This may indicate ACL sprain or partial tear. A complete tear is not identified. Secondary signs of ACL insufficiency are not identified. While there is likely an ACL sprain or mild partial tear, there are intact ACL fibers. The pattern of bone bruises raises some concern for an ACL tear, but the femoral bone bruises slightly more lateral than commonly seen.	1	0
Proximal ACL tear and PCL intact. Left knee MRI demonstrates: Complete ACL tear with bone bruises in the medial tibial plateau, medial femoral condyle and lateral femoral condyle.	1	1
There is no normal anterior cruciate ligament identified. There is diffuse intermediate signal within the posterior cruciate ligament on the proton density images and to a lesser extent on the T2-weighted images compatible with chronic PCL degeneration. While this could be the sequela of old surgery, degenerative tearing of the meniscus including involvement of the root attachment cannot be excluded. Nonvisualization of the ACL compatible with an old ACL tear.	0	1

Figure 10: Figure showing the labels in the Knee dataset. 0 means absence of a given pathology and 1 indicates presence of such.

to these backbones as well. Finally, we choose the best model based on meta-validation accuracy and use it for our downstream tasks. In all our experiments, PubMedBERT-based backbones outperform the BERT-base and the Clinical BERT backbones.

## E Hyperparameters and Additional Experimental Details

In this section, we will describe the hyperparameters used for experiments on our internal and public datasets and explain some of the design choices. Table 8 shows the best hyperparameters used for our experiments. For our internal dataset, we use the Pfeiffer configuration in the adapter implementation from (Pfeiffer et al., 2020), whereas for the public datasets we use the exact implementation and configuration as in (Wang et al., 2021) for a fair comparison to the results reported there. For all vanilla ProtoNet experiments, we use the Euclidean distance as it outperforms the cosine distance. All BERT models without adapters are trained with 8 shots and 8 support due to memory considerations. We choose learning rate and the variance regularizer for each model from  $\{1e-5, 2e-5, 5e-5, 1e-4\}$  and  $\{1e-4, 1e-3, .01, .1, .5\}$  based on the validation performance. For all the experiments, a dropout

Backbone	Methods	Accuracy
BERT-base	Vanilla ProtoNet	86.3 ± 1.2
	Big ProtoNet	87.8 ± .9
	Leopard	81.4 ± 9.7
	ProtoNet w/ Isotropic Gaussian	88.7 ± 1.4
	ProtoNet w/ Isotropic Gaussian + reg	89.5 ± .8
BERT-base w/Adapters	Variance Aware ProtoNet (ours)	88.9 ± 1.5
	Variance Aware ProtoNet + reg (ours)	90.1 ± .9
	Vanilla ProtoNet	85.6 ± 1.3
Clinical BERT	Big ProtoNet	87.1 ± 1.1
	Leopard	87.8 ± .8
	ProtoNet w/ Isotropic Gaussian	88.6 ± .7
	ProtoNet w/ Isotropic Gaussian + reg	88.1 ± 1.2
	Variance Aware ProtoNet (ours)	89.7 ± .8
Clinical BERT w/Adapters	Variance Aware ProtoNet + reg (ours)	87.4 ± 1.3
	Vanilla ProtoNet	88.5 ± 1.1
	Big ProtoNet	82.2 ± 9.8
	Leopard	89.6 ± 1.2
	ProtoNet w/ Isotropic Gaussian	90.1 ± .8
Clinical BERT w/Adapters	ProtoNet w/ Isotropic Gaussian + reg	89.9 ± 1.1
	Variance Aware ProtoNet (ours)	90.9 ± .8
	Variance Aware ProtoNet + reg (ours)	86.8 ± .9
	Vanilla ProtoNet	87.9 ± 1.1
	Big ProtoNet	88.4 ± 1.3
Clinical BERT w/Adapters	ProtoNet w/ Isotropic Gaussian	89.1 ± .9
	ProtoNet w/ Isotropic Gaussian + reg	88.7 ± 1.1
	Variance Aware ProtoNet (ours)	88.7 ± 1.1
	Variance Aware ProtoNet + reg (ours)	89.5 ± .9

Table 6: Results showing accuracy percentages on the meta-validation dataset. We sample 1000 tasks with 4-way 8-shot and 16-support classification. We replicate each experiment over 10 random seeds.

layer is added after the final BERT layer.

For our internal dataset, we also experiment with  $\{2, 3, 4\}$  ways and  $\{4, 6, 8\}$  shots and  $\{4, 6, 8, 12, 16\}$  support. The experiments with 2-way and 3-way produce poor results on our downstream tasks irrespective of the number of shots and support. During training with 4-way, the meta-validation results for lower support show worse performance than the numbers reported in table 4. We believe that it is caused by the high variability between the various groups of samples of a given class. Finally our downstream performance is best for models that are trained on higher number of shots.

In case of ProtoNets, there is no adaptation during testing. The validation set is used to compute prototypes to query the test set. However, in case of Leopard, there is an additional few shot adaptation step that occurs as outlined in (Wang et al., 2021). In this case, the validation set is used for the adaptation and also as the support set for querying the test set.

### E.1 Effect of Regularization on means and variances

Table 9 illustrates the benefits of adding the regularization term. The regularization term not only aids in lowering the variances but also manages

Backbone	Methods	Foraminal	Knee (Acute Tear vs Not)	Knee (Complete tear vs Not)	Lung
BERT-base	Baseline	.24	.29	.32	.19
	Multi-Task	.29	.34	.41	.27
	Vanilla ProtoNet	.75	.71	.66	.65
	Big ProtoNet	.57	.58	.53	.6
	Leopard	.63	.72	.61	.41
	ProtoNet w/ Isotropic Gaussian	.77	.72	.69	.68
	ProtoNet w/ Isotropic Gaussian + reg	.78	.76	.71	.70
	Variance Aware ProtoNet (ours)	.79	.78	.73	.72
	Variance Aware ProtoNet + reg (ours)	.81	.80	.76	.75
BERT-base w/ Adapters	Baseline	.28	.32	.40	.25
	Multi-Task	.32	.35	.44	.29
	Vanilla ProtoNet	.74	.73	.65	.67
	Big ProtoNet	.58	.59	.55	.61
	ProtoNet w/ Isotropic Gaussian	.78	.71	.67	.69
	ProtoNet w/ Isotropic Gaussian + reg	.80	.74	.72	.74
	Variance Aware ProtoNet (ours)	.80	.74	.72	.74
	Variance Aware ProtoNet + reg (ours)	.82	.77	.77	<b>.78</b>
	Clinical BERT	Baseline	.31	.37	.42
Multi-Task		.34	.45	.47	.38
Vanilla ProtoNet		.77	.72	.68	.66
Big ProtoNet		.57	.59	.53	.61
Leopard		.74	.78	.77	.62
ProtoNet w/ Isotropic Gaussian		.78	.74	.71	.68
ProtoNet w/ Isotropic Gaussian + reg		.80	.76	.74	.71
Variance Aware ProtoNet (ours)		.82	.79	.76	.74
Variance Aware ProtoNet + reg (ours)		.84	.81	.79	.76
Clinical BERT w/ Adapters	Baseline	.35	.42	.45	.33
	Multi-Task	.37	.45	.49	.37
	Vanilla ProtoNet	.76	.74	.70	.67
	Big ProtoNet	.58	.60	.57	.62
	ProtoNet w/ Isotropic Gaussian	.79	.76	.72	.70
	ProtoNet w/ Isotropic Gaussian + reg	.81	.77	.73	.72
	Variance Aware ProtoNet (ours)	.83	.81	.76	.73
	Variance Aware ProtoNet + reg (ours)	<b>.85</b>	<b>.82</b>	<b>.81</b>	.77

Table 7: Table showing F1 scores of few shot models with BERT-base and Clinical BERT backbones in downstream classification tasks.

to push the centroids away further which we believe sheds some light on our method’s success in downstream classification tasks. We also carry out various ablation studies by changing the regularization hyperparameter.

## E.2 Metric and other modeling choices

In (Snell et al., 2017), only the sample means (i.e. means of the support vectors) are used to estimate the true population mean. In fact, by the Central limit theorem, we can use the sample variance (after normalization) to get an unbiased estimate of the population variance. Unlike the original work, we sought to use this extra information to better understand the class distribution. The Gaussian assumption is strong but it is motivated by the fact that it allows us to compute Wasserstein distances in a computationally tractable manner. Finally to motivate the choice of using the Wasserstein distance instead of a Bergman divergence like KL diver-

gence, consider the following motivating example,  $N_1(\mu_1, \Sigma_1), N_2(\mu_2, \Sigma_2)$  be 2 Gaussians and for simplicity assume:  $\Sigma_1 = \Sigma_2 = wI$  and  $\mu_1 \neq \mu_2$ . With these assumptions,  $W^2 = \|\mu_1 - \mu_2\|^2$  And  $D_{KL} = \frac{1}{2w}\|\mu_1 - \mu_2\|^2$ . Note that Wasserstein distance does not change if the variance changes ( $w$  can be arbitrarily large) whereas the KL divergence does. In fact, this is pointed out in (Ding et al., 2022) where their goal is to create spherical Gaussians with large radii. However, we found that having large variance produces worse results in our downstream tasks. Finally similar dependence on variance is in play if one computes a simple likelihood of the sample in the class distributions.

## F Public Benchmarks

In this section, we describe the training procedure for the public benchmark datasets. The baseline results are taken from (Wang et al., 2021;

**Algorithm 1:** PyTorch style Pseudocode for Variance Aware ProtoNets

```

/* f: Encoder Network */
/* N: dimension of the representation */
/* c: Number of classes or ways */
/* k: Shots or number of examples per class in the query set */
/* m: Supports per class in the support set */
/* dist: Pairwise squared Euclidean distance function */
/* loss_fn: Cross-Entropy Loss function */
/* λ: regularizer */
Input: Sample a set  $L$  of labels, mini-batch of Support set  $S_L$ , Query set  $Q_L$ 
/* Compute statistics for each class in the Support set */
sorted_labels = torch.sort(support_labels) // sort the labels in the support set
c = len(s.values.unique()) // Number of ways
support_sorted = support[sort.indices]
labels_sorted = labels[sort.indices]
embeddings_support = f(support_sorted) // m * c * N
m = embeddings_support.shape[0] // support per class
embeddings_support = embeddings_support.reshape(c, m, -1) // c * m * N
support_mean = embeddings_support.mean(dim=1) // c * N
support_var = torch.var(embeddings_support, dim=1)**2 // c * N
/* Get embeddings for the query set and compute distances from the support */
query = f(Q_L) // k * N
logits = dist(query, support_mean) + torch.sum(support_var, dim=1) // k * c, adding trace to the
distance matrix
loss = loss_fn(-logits, query_labels) + λ * (torch.norm(support_var, dim=1))/c // Regularizer term

```




Figure 11: Figure showing the effect of changing the regularization hyperparameter. Top: Figure showing the F1 score averaged over our 4 internal datasets. Bottom: Figure showing 16-shot accuracy averaged over 13 public datasets. We see a similar trend for 4 and 8-shot accuracies for these public datasets as well.

Hyperparameter Type	Internal Dataset	Public Datasets
Epochs	30	40
Sequence Length	128	128
Optimizer	AdamW	AdamW
Learning Rate	$3e - 5$	$2e - 5$
Weight Decay	$1e - 4$	$1e - 4$
Gradient Clip	3	2
Early Stopping	Yes	Yes
Learning Rate Scheduler	Linear	Linear
Dropout	.1	.1
Shots	8	8
Number of supports	16	16
Variance Regularizer	.1	.01

Table 8: Hyperparameters used for all our Variance Aware ProtoNet experiments with BERT+Adapter backbones

Bansal et al., 2020a). We have followed the same meta-training procedure as described in (Wang et al., 2021). Specifically, for meta-training, WNLI (m/mm), SST-2, QQP, RTE, MRPC, QNLI, and the SNLI datasets (Bowman et al., 2015) are used. The validation set of each dataset is used for hyperparameter searching and model selection. The models are trained by sampling episodes from the meta-training tasks. The sampling process first selects a

Experiment type	Without regularization			With regularization		
Task Names	Distance	Norm of	Norm of	Distance	Norm of	Norm of
	between centroids	class 0 variance	class 1 variance	between centroids	class 0 variance	class 1 variance
Lung	2.97	1.12	1.31	3.96	1.15	1.08
Foraminal	3.65	1.71	1.59	4.17	1.62	1.38
Knee (ACL complete tear)	4.12	1.97	2.10	5.01	1.87	1.85
Knee (ACL acute tear)	3.95	1.53	1.49	4.32	1.27	1.35

Table 9: Table showing the showing the class statistics with and without regularization. Higher Distance and Lower variance is better.

dataset and then randomly selects  $m$  examples for each class as the support set and another  $k$ -shots as the query set and the probability of a selected task is proportional to the square root of its dataset size (Bansal et al., 2020b). For meta-testing, we use 13 datasets ranging from NLI, text classification and sentiment analysis. For the models and datasets marked with \*, we use the results reported in (Bansal et al., 2020a) and for those datasets, we use the code from (Wang et al., 2021) to generate the results for ProtoNet with Bottleneck Adapters while the rest of the results are taken from (Wang et al., 2021). We reuse their implementation and configuration of their adapters but modify the loss function with the Wasserstein distance along with our variance regularization term. Table 5 shows the superior performance of our method beating all the baselines. For detailed hyperparameters, please see section E. Our method without the variance regularization term shows similar performance to that of the Leopard baselines. For the isotropic variant method, it shows similar performance to Leopard with the variance regularization term and worse without.

## G Stability of the Prototypes

For simplicity, we use our entire validation sets to compute prototypes. In this section we show how our results vary if we choose a subset of our validation set to create the prototypes. The figure 12 shows the F1 scores when a subset of the data is used to compute the prototypes and the variances for a given class.

## H Failure Cases

We also test our models on few additional tasks like (i) predicting the severity of disc herniation in our cervical dataset and (ii) predict the presence of cord compression at various motion segments in our internal dataset on the lumbar spine. Our models achieve an F1 score of .51 and .39 respectively. The figure 13 shows how the classes are distributed.




Figure 12: Figure showing stability of the prototypes. We sample  $k$  examples 50 times to construct the prototypes and the standard deviations.

We attribute the failures to the poor separability between classes and the high variance in the data distribution.

It is an ongoing project to understand what makes our model work for these downstream tasks and why our model works on some tasks and not others. We hope that by simply increasing the diversity of our training data or applying newer adapter architectures like Mix-and-Match Adapter (He et al., 2022) and Compacter (Karimi Mahabadi et al., 2021), our current methods will work on a wide range of downstream pathologies.




Figure 13: T-Sne projections of our Cord and Disc Data. The prototypes for cord classes are very close while the prototypes for disc are well separated. However the large variance in the disc classes causes bad performance.