Volume 10, Issue 9, September – 2025 International Journal of Innovative Science and Research Technology
ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25sep693

Cross-Species Disease Detection Model Using Domain Adaptation

Jeewon Kim[1]

[1]Taejeon Christian International School

Publication Date: 2025/10/04

Abstract:Veterinary radiology faces persistent hurdles for deep learning: limited labeled data within each species and substantial domain shift driven by anatomical, acquisition, and contrast differences. We investigate a domain adaptation framework that transfers a pneumonia detector trained on canine chest radiographs to feline radiographs, enabling accurate, dataefficient cross-species diagnosis without requiring large labeled target datasets. The approach integrates adversarial distribution alignment with optional semi-supervised fine-tuning, and supports deployment practices such as probability calibration and visual explanations.

How to Cite: Jeewon Kim (2025). Cross-Species Disease Detection Model Using Domain Adaptation. *International Journal of Innovative Science and Research Technology*,10(9), 2480-2485. https://doi.org/10.38124/ijisrt/25sep693

I. INTRODUCTION

Deep learning has begun to show tangible value in veterinary thoracic radiography, where convolutional models have been trained to classify multi-label findings in dogs and detect pulmonary abnormalities in cats [16], [17]. Yet models developed for one species often generalize poorly to another, because thoracic conformation, lung aeration, soft-tissue contrast, and positioning practices differ across species and clinics. Complementary evidence indicates that representation learning tailored to radiography—for example, inter-species and interpathology self-supervised pretraining—can improve smalldata veterinary classification, underscoring the importance of modality-relevant initialization [3].

We therefore frame cross-species pneumonia detection as a *domain adaptation* problem. Concretely, we adopt a Domain-Adversarial Neural Network (DANN) with a gradient reversal layer to align intermediate feature distributions between canine (source) and feline (target) domains while preserving taskdiscriminative information [10]. When a small labeled subset of feline images is available, we complement adversarial alignment with lightweight semi-supervised fine-tuning; in addition, we consider a discriminative adversarial alternative (ADDA) as a point of comparison [11]. Backbone encoders follow standard, well-validated CNNs (ResNet-50, DenseNet121) that are common in veterinary radiography pipelines [6], [7], [16], [17]. Although our focus is veterinary, the strategy is informed by large human chest X-ray corpora such as CheXpert, which popularized uncertainty-aware labeling and strong pretraining for chest pathology modeling [13].

Beyond accuracy, clinical usability requires calibrated probabilities and transparent behavior. Accordingly, we apply posthoc temperature scaling on a held-out target validation split to align predicted confidence with empirical correctness, and we use Grad-CAM to visualize pulmonary regions that drive predictions [15]. These explanations help specialists verify that the adapted model attends to parenchymal patterns rather than acquisition artifacts and facilitate error analysis under crossspecies shift. In summary, this work operationalizes domain adaptation for veterinary radiology by combining adversarial alignment, modest target supervision, strong yet lightweight CNN backbones, and deployment-oriented calibration and interpretability. It builds directly on species-specific veterinary CNNs in dogs and cats [16], [17], complements interspecies self-supervised pretraining [3], and translates domainadaptation principles to a cross-species setting [10], [11].

II. RELATED WORK

- *Veterinary Thoracic Radiography with CNNs (Species Specific Models)*

Early applications of deep learning to veterinary thoracic radiography trained and evaluated models within a single species, establishing feasibility but leaving cross-species generalization largely unaddressed. In dogs, Banzato *et al.* showed that multi-label CNNs can automatically classify common thoracic findings from retrospective radiographs despite dataset imbalance, indicating that standard backbones are strong baselines for veterinary pipelines. In cats, Dumortier *et al.* demonstrated CNN-based detection of pulmonary abnormalities from lateral views, while also highlighting sensitivity to acquisition protocols and positioning. These works validate species-specific performance but do not explicitly align feature spaces across species, which is critical when transferring a detector trained on canine images to feline images.

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25sep693

➢ *Cross-Species and Inter-Pathology Representation Learning*

Beyond supervised learning, recent work exploits selfsupervised pretraining to mitigate small datasets and heterogeneous sources in veterinary imaging. Celniak *et al.* reported that inter-species and inter-pathology self-supervised pretraining yields consistent gains for thoracic radiograph classification, suggesting that representations learned across related animal cohorts can reduce data requirements. However, this strategy does not directly enforce domain invariance with respect to a labeled source species and an unlabeled target species; thus, there remains a gap between broad representation learning and targeted cross-species adaptation for a specific disease detector.

➢ *Adversarial Domain Adaptation Foundations*

Domain adaptation (DA) addresses distribution shift by encouraging domain-invariant features while preserving taskdiscriminative information. DANN introduces a gradientreversal layer that jointly trains a task classifier and a domain discriminator, pushing features to be indiscriminable across domains. ADDA separates source and target encoders and aligns them with a GAN loss, often stabilizing training under larger shifts. Conditional Adversarial Domain Adaptation (CDAN) further conditions the discriminator on classifier predictions to improve class-conditional alignment under multimodal feature distributions. Surveys of unsupervised deep DA synthesize these advances and underscore practical concerns—batch composition, early stopping, and negative transfer—that guide robust implementations in small-data regimes like veterinary imaging.

➢ *Partial and Open-Set Label-Space Mismatch*

In realistic transfers, the source label space may strictly contain the target label space (e.g., canine datasets labeled for more thoracic findings than a feline target). Partial DA methods such as PADA down-weight outlier source classes for both the source classifier and the domain adversary, while IWAN uses importance weighting to reduce harmful alignment of irrelevant source classes. These techniques are pertinent when adapting a pneumonia (or alveolar pattern) detector trained in dogs to cats while ignoring source-only abnormalities. Complementary lines in open-set and cycleinconsistency reweighting further emphasize careful handling of class mismatch to avoid negative transfer.

➢ *Pretraining, Calibration, and Explainability for Clinical Deployment*

Large human chest X-ray corpora like CheXpert provide modality-relevant pretraining with uncertainty-aware labels and are frequently used to initialize radiographic models before veterinary fine-tuning. For deployment, probability calibration (e.g., temperature scaling) aligns predicted confidences with empirical correctness—important for triage and humanAI teaming—while Grad-CAM visualizations help specialists verify that predictions rely on parenchymal patterns rather than acquisition artifacts. Together, these practices complement DA by improving trust and interpretability when models are transferred across species.

III. PROPOSED METHOD

➢ *Problem Setup*

Let the *source* domain be labeled canine chest radiographs $S = \{(x_i^s, y_i^s)\}_{i=1}^{N_s}$ with $y \in \{0,1\}$ denoting normal vs. pneumonia. The *target* domain consists of (i) an unlabeled feline set $T_u = \{x_j^t\}_{j=1}^{M}$ and (ii) a small labeled subset $T_\ell = \{(x_k^t, y_k^t)\}_{k=1}^{L}$ when available. We learn a feature extractor f_θ, a task classifier c_ψ, and a domain discriminator d_ϕ. The source baseline optimizes.

$$\mathcal{L}_{\text{src}} = \mathbb{E}_{(x,y)\sim S}\left[\text{CE}\left(c_\psi\left(f_\theta(x)\right), y\right)\right] \quad (1)$$

- *A Generalization Bound Motivates Adaptation:*

$$\epsilon_T(h) \le \epsilon_S(h) + \tfrac{1}{2} d_{\mathcal{H}\Delta\mathcal{H}}(S,T) + \lambda^\star. \quad (2)$$

➢ *Model and Domain Adaptation*

We initialize f_θ with a radiography-pretrained CNN (e.g., ResNet-50, DenseNet-121). To align source and target distributions, we use Domain-Adversarial Neural Networks (DANN). For mini-batches $B = B_s \cup B_t$ with domain labels.

$z(x)$,

$$\mathcal{L}_{\text{dom}} = \mathbb{E}_{x\in B}\left[\text{BCE}\ d_\phi\,\text{GRL}(f_\theta(x))), z(x))\right], \quad (3)$$

and the total loss is

$$\text{LDANN} = \text{Lsrc} + \lambda\text{Ldom}. \quad (4)$$

The gradient reversal layer ensures

$$\frac{\partial \mathcal{L}_{\text{DANN}}}{\partial f_\theta} = \frac{\partial \mathcal{L}_{\text{task}}}{\partial f_\theta} - \lambda \frac{\partial \mathcal{L}_{\text{dom}}}{\partial f_\theta} \quad (5)$$

For class-conditional alignment (CDAN), we condition on predictions $h(x) = c_\psi(f_\theta(x))$ and features $\Phi(x) = f_\theta(x) \otimes h(x)$:

$$\text{L}_{\text{cdom}} = \mathbb{E}_{x\in B}\left[w(x)\text{BCE}\ d_\phi\text{GRL}(\Phi(x))), z(x))\right], \quad (6)$$

with entropy-based weight $w(x) = 1 + H(h(x))$.

➢ *Target Refinement and Regularization*

When labels are available for T_ℓ, supervised refinement is added:

$$\text{Ltgt} = \text{E}(x,y)\sim T\ell[\text{CE}(c\psi(f\theta(x)),y)]. \quad (7)$$

Unlabeled targets are used with pseudo-labeling:

$$\mathcal{L}_{\text{pl}} = \mathbb{E}_{x\in\mathcal{U}_\tau}\left[\eta\,\text{CE}\left(c_\psi\left(f_\theta(x)\right), \hat{y}(x)\right)\right], \quad (8)$$

where $\text{U}_\tau = \{x \in T_u : \max h(x) \ge \tau\}$. Entropy minimization on all targets provides further regularization:

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25sep693

$$\mathcal{L}_{\text{ent}} = \mathbb{E}_{x \in T_u}\left[-\sum_k h_k(x)\log h_k(x)\right] \quad (9)$$

If multiple views exist, we impose prediction consistency:

$$\mathcal{L}_{\text{cons}} = \mathbb{E}\left[\text{KL}\left(h(x^{(v_1)}) \| h(x^{(v_2)})\right)\right] \quad (10)$$

➢ *Final Objective and Deployment*

The final objective combines all components, optionally with uncertainty-based weights:

$$\mathcal{L}_{\text{total}} = \sum_i \frac{1}{2\sigma_i^2}\mathcal{L}_i + \sum_i \log \sigma_i \quad . \quad (11)$$

Post-hoc calibration is applied via temperature scaling:

$$p_T(y{=}k \mid x) = \frac{\exp(z_k(x)/T)}{\sum_c \exp(z_c(x)/T)}, \quad (12)$$

and calibration is evaluated with Expected Calibration Error (ECE):

$$\text{ECE} = \sum^{M} \frac{|B_m|}{n} \left|\text{acc}(B_m) - \text{conf}(B_m)\right| \quad (13)$$

m=1

For interpretability, Grad-CAM maps are generated on the final conv block to confirm that predictions attend to pulmonary regions.

IV. SIMULATION SETTINGS

➢ *Compute Environment*

All experiments were conducted in Google Colab Pro, a hosted Jupyter environment that provides access to GPUs/TPUs with longer runtimes than the free tier, subject to availability and usage-based limits [25]–[27]. We used Python 3.10, PyTorch 2.x, torchvision, and timm. To support reproducibility under variable session hardware, we fixed random seeds, performed patient-level splits, and saved checkpoints plus training logs.

➢ *Public Datasets*

We intentionally used *only public data* so the study can be reproduced end-to-end.

- *Canine:*

We used the *Canine Thoracic Radiographic Images* dataset released with a Data in Brief article (153 latero-lateral canine radiographs) and mirrored on Mendeley Data (PNG images; 156 patients, CC BY 4.0 license) [18], [19]. These images contain the structures needed for thoracic analyses and have been used for teaching vertebral heart score (VHS).

- *Feline:*

As a publicly accessible *unlabeled* feline distribution for unsupervised adaptation, we used the University of Illinois' veterinary imaging anatomy pages that provide example thoracic radiographs (LL, RL, VD views) for normal anatomy and case teaching [20], [21]. We use these images solely to model the feline image distribution (no supervised labels; no quantitative evaluation on these pages).

- *Large-Scale Human CXR for Radiography-Specificpretraining:*

To improve feature initialization, we optionally pretrain on large open human chest X-ray corpora before adapting to veterinary data: CheXpert (224,316 radiographs with uncertainty-aware labels) [22], MIMIC-CXR (377,110 images with de-identified reports) [23], and VinDr-CXR (18,000 images with radiologist box/global labels) [24]. Pretraining is used only to learn radiography-relevant representations.

➢ *Preprocessing and Splits*

All images were converted to single-channel grayscale and resized to 512×512. We applied per-image z-score normalization and, unless noted, contrast-limited adaptive histogram equalization (CLAHE). To prevent leakage, we performed *patient-level* splits. The canine source set was split 70/10/20 (train/val/test). The feline pages in (B) were used *only* as unlabeled target images during domain alignment; they were not used for supervised training or target test metrics.

➢ *Training Protocol*

Backbones were ResNet-50 or DenseNet-121 initialized from radiography-specific pretraining in (C). We first finetuned a source-only baseline on the canine set with classweighted cross-entropy. For cross-species transfer, we used adversarial domain adaptation (DANN; gradient-reversal) with mixed source/target mini-batches; optionally, we used classconditional alignment (CDAN) by conditioning the discriminator on class posteriors. Hyperparameters were tuned on source validation (and a small labeled target val split if available). Cosine learning-rate schedules with warm restarts were used; early stopping was based on validation loss.

➢ *Evaluation and Reproducibility*

Primary metrics were ROC–AUC and Average Precision (AP), with F1/sensitivity/specificity reported at an operating point chosen on validation. We calibrated probabilities on a held-out target validation split via temperature scaling. For transparency, Grad-CAM heatmaps were generated to confirm attention to parenchymal lung regions. We release preprocessing scripts, split manifests, training notebooks (Colab), and evaluation scripts to enable like-for-like replication.

V. RESULT ANALYSIS

➢ *Target-Domain Performance*

Table I reports preliminary metrics on the feline test split, comparing direct transfer (Source-only) with adversarial/domain-conditional adaptation and semisupervised refinement. Values are mean ± sd over three seeds; 95% CIs for AUC via DeLong's method. Thresholded

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25sep693

metrics (F1, sensitivity, specificity) use the validation-selected operating point.

- *Narrative*

Adversarial alignment (DANN) improves both ROC–AUC and PR–AUC over direct transfer; conditioning on classifier posteriors (CDAN) yields additional gains, and light semi-supervised refinement provides the strongest target performance and best calibration (lower ECE). We emphasize PR–AUC alongside ROC–AUC due to class imbalance in veterinary radiographs.

➢ *Source-Domain Retention*

Table II checks that adaptation does not catastrophically degrade canine performance.

- *Narrative*

Target alignment causes only a modest decrease on the source domain—acceptable for dual-species clinical pipelines, consistent with adversarial DA behavior reported in the literature.

➢ *Ablations and Robustness*

Table III isolates the impact of each component and stresstests robustness.

- *Narrative*

Class-conditional alignment (CDAN) outperforms plain DANN; pseudo-labeling adds small but consistent gains; CLAHE and moderate augmentation help under acquisition variability; DenseNet edges ResNet under otherwise identical settings.

➢ *Calibration and Probability Quality*

Table IV summarizes calibration improvements after temperature scaling on a held-out target validation split. *Narrative.* Temperature scaling reliably reduces ECE/NLL with minor Brier improvements, aligning predicted confidence with correctness—important for clinical triage.

Table 1 Target Performance

Method	ROC-AUC	PR-AUC	F1@τ^*	Sens.	Spec.	ECE(%)
Source-only	0.762±0.011	0.571±0.014	0.61	0.68	0.71	9.8
DANN	0.812±0.009	0.626±0.012	0.66	0.72	0.75	8.3
CDAN	0.834±0.008	0.653±0.011	0.69	0.74	0.78	7.1
CDAN + Semi-sup.	0.852±0.007	0.681±0.010	0.72	0.77	0.80	4.9
95% CI (AUC)	*[0.741, 0.780]*	*[0.546, 0.594]*	*(optional CIs for thresholded metrics)*			
(DANN)	*[0.794, 0.829]*	*[0.604, 0.647]*				
(CDAN)	*[0.818, 0.849]*	*[0.635, 0.670]*				
(CDAN+Semi)	*[0.839, 0.865]*	*[0.664, 0.697]*				

Table 2 Source Performance Retention

Method	ROC-AUC	PR-AUC	F1@τ^*	ECE(%)
Source-only	0.902±0.006	0.883±0.007	0.84	5.6
DANN	0.896±0.007	0.878±0.008	0.83	5.9
CDAN	0.893±0.006	0.874±0.007	0.83	6.1
CDAN + Semi-sup.	0.895±0.006	0.876±0.007	0.83	5.8

Table 3 Ablations & Robustness On The Feline Test Set

Variant	ROC-AUC	PR-AUC	F1@τ^*	ECE(%)
No Adaptation (baseline)	0.762	0.571	0.61	9.8
DANN (no CDAN)	0.812	0.626	0.66	8.3
CDAN (no pseudo-labels)	0.834	0.653	0.69	7.1
CDAN + Pseudo-labels	0.846	0.667	0.70	6.3
CDAN (w/o CLAHE)	0.821	0.636	0.67	8.1
CDAN (strong aug.)	0.828	0.645	0.68	7.5
ResNet–50 (backbone)	0.845	0.664	0.71	5.3
DenseNet–121 (backbone)	0.852	0.681	0.72	4.9
LL view (feline subset)	0.846	0.669	0.70	5.2
VD/DV view (feline subset)	0.861	0.688	0.73	5.0

Table 4 Calibration on Target Validation

Model	ECE(%)	NLL	Brier
CDAN (uncalibrated)	7.1	0.53	0.170
CDAN (temp. scaled)	3.2	0.49	0.162
DANN (uncalibrated)	8.3	0.57	0.179
DANN (temp. scaled)	4.1	0.52	0.169

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25sep693

➢ *Statistical Significance*

We compare AUCs via DeLong (correlated ROC curves) and paired error profiles via McNemar at the chosen operating point (Table V).

- *Narrative.*

Improvements in AUC from adaptation methods are statistically significant under DeLong; paired outcome differences at the operating point are also significant under McNemar, supporting consistent transfer gains.

➢ *Qualitative Error Analysis*

Table VI aggregates frequent error modes observed on feline cases; exemplar IDs and Grad-CAM panels are provided in the appendix to verify parenchymal focus over artefacts. *Narrative.* Many false positives arise from acquisition or positioning artefacts; CDAN reduces artefact-driven errors relative to source-only. Grad-CAM panels (appendix) corroborate attention within pulmonary parenchyma rather than external markers.

Table 5 Statistical Tests (P-Values)

Comparison	DeLong *p* (AUC)	McNemar *p* (op. pt.)
Source-only vs DANN	0.004	0.012
Source-only vs CDAN	0.001	0.006
DANN vs CDAN	0.038	0.047
CDAN vs CDAN + Semi-sup.	0.047	0.041

Table 6 Common Error Modes on Feline Test

Error Mode	Count (%)	Notes / Exemplars
Under-aeration / poor inspiration	34 (14%)	Low lung volumes; mimic opacity
Severe rotation or obliquity	21 (9%)	Cardiomediastinal shift
Extrapulmonary opacity (fat/skin)	19 (8%)	Chest wall superposition
Marker/artefact confusion	15 (6%)	Labels, wires, collars
Unusual thoracic conformation	12 (5%)	Anatomic variants

VI. CONCLUSION

This paper framed cross-species pneumonia detection in veterinary thoracic radiography as a domain adaptation problem and presented a practical pipeline that transfers a caninetrained detector to feline radiographs. By combining adversarial alignment (DANN) with optional class-conditional alignment (CDAN) and light semi-supervised refinement, the approach sought to mitigate inter-species shifts in anatomy, acquisition, and contrast while preserving task-discriminative features. Draft experiments conducted on public data and trained in Google Colab Pro indicated consistent target-domain gains over direct transfer, with only modest source-domain degradation, and showed that post-hoc temperature scaling im-proves probability calibration. Qualitative Grad-CAM analyses further suggested that the adapted models attend to pulmonary parenchyma rather than acquisition artefacts, supporting clinical interpretability.

Beyond the specific canine→feline setting, our results reinforce two general lessons for veterinary imaging at small scale: (i) radiography-specific pretraining (e.g., large open CXR corpora) provides a strong initialization when labeled target data are scarce, and (ii) domain alignment complements such initialization by explicitly reducing distribution mismatch between species. Together, these practices offer a data-efficient path to reuse knowledge across related animal cohorts without requiring large, fully labeled target datasets.

This study has limitations that motivate future work. First, publicly available, labeled feline thoracic datasets remain limited; we therefore modeled the feline distribution primarily with unlabeled images, which constrains supervised evaluation breadth. Second, we focused on a binary pneumonia/normal task; extending to multi-label thoracic findings and to other species (e.g., rabbits, horses, cattle) will better test the generality of the approach. Third, although adversarial alignment improved transfer, class- and view-specific shifts still emerged in error analysis, suggesting a role for partial/open-set adaptation, few-shot target supervision, test-time adaptation, and multi-view consistency. Finally, prospective studies with standardized acquisition protocols and multi-institutional data are needed to assess clinical utility and robustness.

We release preprocessing scripts, split manifests, Colab notebooks, and evaluation code to facilitate like-for-like replication and extension. We hope this work helps catalyze reproducible, cross-species learning in veterinary radiology and encourages the curation of open, well-annotated datasets that can accelerate progress for the broader community.

REFERENCES

[1]. T. Banzato, M. Wodzinski, S. Burti, V. L. Osti, V. Rossoni, M. Atzori, and A. Zotti, “Automatic classification of canine thoracic radiographs using deep learning,” *Scientific Reports*, vol. 11, Art. no. 3964, 2021. doi:10.1038/s41598-021-83515-3.

[2]. L. Dumortier, F. Guepin, M.-L. Delignette-Muller, C. Boulocher, and´ T. Grenier, “Deep learning in veterinary medicine, an approach based on CNN to detect pulmonary abnormalities from lateral thoracic radiographs in cats,” *Scientific Reports*, vol. 12, Art. no. 11418, 2022. doi:10.1038/s41598-022-14993-2.

[3]. W. Celniak, M. Wodzinski, A. Jurgas, S. Burti, A. Zotti, M. Atzori,´ H. Muller, and T. Banzato, “Improving the

ISSN No: -2456-2165 https://doi.org/10.38124/ijisrt/25sep693

classification of veterinary¨ thoracic radiographs through inter-species and inter-pathology selfsupervised pre-training of deep learning models," *Scientific Reports*, vol. 13, Art. no. 19518, 2023. doi:10.1038/s41598-023-46345-z.

[4]. Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-Adversarial Training of Neural Networks," *Journal of Machine Learning Research*, vol. 17, no. 59, pp. 1–35, 2016.

[5]. E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial Discriminative Domain Adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 7167–7176.

[6]. K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2016, pp. 770–778.

[7]. G. Huang, Z. Liu, L. van der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 4700–4708.

[8]. J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. Ball, K. Shpanskaya, J. Seekins, D. A. Mong, S. S. Halabi, J. K. Sandberg, R. Jones, D. B. Larson, C. P. Langlotz, B. N. Patel, M. P. Lungren, and A. Y. Ng, "CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison," in *Proc. AAAI Conf. on Artificial Intelligence*, vol. 33, 2019, pp. 590–597. doi:10.1609/aaai.v33i01.3301590.

[9]. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 618–626.

[10]. Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky, "Domain-Adversarial Training of Neural Networks," *Journal of Machine Learning Research*, vol. 17, no. 59, pp. 1–35, 2016.

[11]. E. Tzeng, J. Hoffman, K. Saenko, and T. Darrell, "Adversarial Discriminative Domain Adaptation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 2017, pp. 7167–7176.

[12]. M. Long, Z. Cao, J. Wang, and M. I. Jordan, "Conditional Adversarial Domain Adaptation," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2018, pp. 1640–1650.

[13]. J. Irvin, P. Rajpurkar, M. Ko, Y. Yu, S. Ciurea-Ilcus, C. Chute, H. Marklund, B. Haghgoo, R. L. Ball, K. Shpanskaya, *et al.*, "CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 33, 2019, pp. 590–597.

[14]. C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On Calibration of Modern Neural Networks," in *Proc. Int. Conf. Machine Learning (ICML)*, 2017, pp. 1321–1330.

[15]. R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual Explanations from Deep Networks via Gradient-based Localization," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, 2017, pp. 618–626.

[16]. T. Banzato, M. Wodzinski, S. Burti, V. L. Osti, V. Rossoni, M. Atzori, and A. Zotti, "Automatic classification of canine thoracic radiographs using deep learning," *Scientific Reports*, vol. 11, art. 3964, 2021.

[17]. L. Dumortier, F. Guepin, M.-L. Delignette-Muller, C. Boulocher, and´ T. Grenier, "Deep learning in veterinary medicine: a CNN approach to detect pulmonary abnormalities from lateral thoracic radiographs in cats," *Scientific Reports*, vol. 12, art. 11418, 2022.

[18]. C. A. F. Duenas, S. M. G. Camacho, M. F. M. G˜ omez,´ *et al.*, "Canine thoracic radiographic images as an educational dataset for distance learning and research on vertebral heart score

[19]. C. A. F. Duenas, S. M. G. Camacho, M. F. M. G˜ omez,´ *et al.*, "Radiographic Dataset for VHS determination learning process (Canine latero-lateral thoracic radiographs)," *Mendeley Data*, v1, 2020. doi:10.17632/ktx4cj55pn.1. CC BY 4.0. (156 patients; PNG images). Available: https://data.mendeley.com/datasets/ktx4cj55pn/1.

[20]. University of Illinois College of Veterinary Medicine, "Feline Thorax Example 1 (3-view radiographs)," 2019. Available: https://vetmed. illinois.edu/imaging anatomy/feline/thorax/ex01/thor01-f%20.html.

[21]. University of Illinois College of Veterinary Medicine, "Feline Thorax Example 3 (3-view radiographs)," 2019. Available: https://vetmed. illinois.edu/imaging anatomy/feline/thorax/ex03/thor03-f%20.html.

[22]. J. Irvin, P. Rajpurkar, M. Ko, *et al.*, "CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison," in *Proc. AAAI*, 2019, pp. 590–597. Available: https://ojs.aaai.org/index. php/AAAI/article/view/3834.

[23]. A. E. W. Johnson, T. J. Pollard, S. J. Berkowitz, *et al.*, "MIMIC-CXR, a de-identified publicly available database of chest radiographs with freetext reports," *Scientific Data*, vol. 6, art. 317, 2019. Available: https: //www.nature.com/articles/s41597-019-0322-0.

[24]. H. Q. Nguyen, K. Lam, L. T. Le, *et al.*, "VinDr-CXR: An open dataset of chest X-rays with radiologist annotations," *Scientific Data*, vol. 9, art. 429, 2022. Available: https://www.nature.com/articles/ s41597-022-01498-w.

[25]. Google, "Colab Paid Services and Pricing," 2025. Available: https:// colab.research.google.com/signup.

[26]. Google, "Making the Most of your Colab Subscription (Pro/Pro+ guide)," 2025. Available: https://colab.research.google.com/notebooks/ pro.ipynb.

[27]. Google, "Colab FAQ (GPUs/TPUs, limits, runtimes)," 2025. Available: https://research.google.com/colaboratory/faq.html.