Transfer learning: leveraging human data to build smarter veterinary AI

Nathan Bollig, DVM, MS
Department of Computer Sciences, University of Wisconsin-Madison, Madison, WI
Department of Pathobiological Sciences, School of Veterinary Medicine, University of Wisconsin-Madison, Madison, WI

Amy Nichelason, DVM, DABVP (Canine and Feline Practice)
Department of Medical Sciences, School of Veterinary Medicine, University of Wisconsin-Madison, Madison, WI

Keywords: transfer learning, machine learning, deep learning, artificial intelligence, clinical decision support, predictive analytics, electrocardiogram

ABSTRACT

This lecture provides a non-technical introduction to the popular methodology of transfer learning. In this approach, a predictive model can be trained to perform a task using non-veterinary data and then adapted to a similar task relevant for non-human species. When transfer learning is successful, the resulting predictive models may outperform similar models trained only on veterinary data, and the demand for annotated veterinary data can be significantly reduced. Transfer learning thus may leverage the power of non-veterinary data sets, which are often larger and more readily accessible than comparable veterinary data, to build smarter veterinary AI applications.

The basics of machine learning and transfer learning are discussed. We outline the advantages and important limitations of this technology, referring to examples of veterinary clinical and research applications throughout the scientific literature.

We hope that by explaining transfer learning and exploring the types of veterinary applications that use this methodology, members across the veterinary profession will acquire foundational knowledge to support ongoing advocacy for veterinary informatics.

Learning Objectives

1. Define machine learning and transfer learning in non-technical language.
2. Explain why transfer learning can be used to leverage knowledge from non-veterinary data sets and why this is useful for veterinary medicine.
3. Identify the strengths and limitations of transfer learning as a method for building predictive models.
4. Recognize a range of veterinary clinical and scientific prediction tasks that are supported by transfer learning.
5. Explain how animals and people can benefit from predictive modeling and models developed using transfer learning.

Introduction to Machine Learning

A **machine learning model** is a mathematical function that maps some input A to some output B. An example is a model that accepts a histological photomicrograph of canine mammary tissue and predicts whether it is benign or malignant [1]. The key to machine learning is that instead of trying to tell the computer exactly what constitutes an image of a malignant tumor, the model is *learned* by being presented with a **training set** consisting of many examples of inputs and outputs. Once the model is trained, it can be used to classify images that it has not seen in its training set. In general, a model is designed to carry out a well-defined task, and its possible inputs and outputs are well defined.

Computer vision is a broad area of computer science related to the analysis of visual input, and includes popular predictive tasks like **image classification** (assigning a label), **object detection** (placing a bounding box around detected objects), and **segmentation** (tracing the shape of objects). Vision models are important for veterinary clinical informatics because a variety of clinically significant data are images – ultrasound, MRI, radiographs, photographs, photomicrographs, etc. There are also machine learning predictive models designed to operate on text, audio, genomic sequences, ECG signals, and other diagnostic data.

We will focus on a particular class of machine learning models called **neural networks**. Neural networks work by passing data through a sequence of data transformations called **layers**. The data transformation performed by a given layer can be completely determined by a set of numbers called **weights**, and the goal of learning is to find the optimal values for these weights for all layers in the network. The best weights are those that enable the network to most accurately map inputs to their corresponding outputs.

Since the goal of learning is to make the output of the network close to the true label found in a training set, it is important to measure the distance between the current output and the expected output. For a given input, a **loss function** provides this measurement. It provides a **loss score** by comparing the prediction of the network with the true output. The fundamental idea behind neural network learning is that the loss score is used as a feedback signal to adjust the weights, in order to subsequently reduce the loss score on a particular input. This uses a mathematical process carried out by a module called an **optimizer**. Figure 1 represents these components of a neural network.

Figure 1. Components of a neural network. Reprinted with permission from [2].

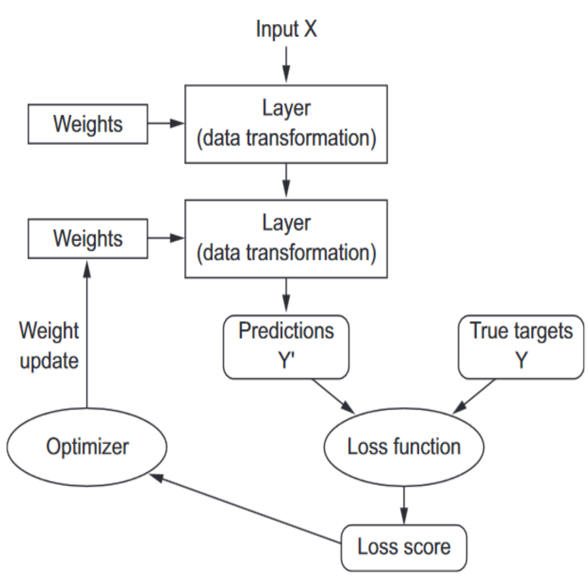

The goal of **training** a network is to iteratively update network weights until we have found a set of weights that minimize the loss function for the images in the training set. The process involves the following steps. First, we start by initializing the network weights randomly. Then we select an image and run it through the network to get a prediction. Let's say that this is an image of a malignant tumor, but the network predicts that it is benign. This results in a higher loss score, because the current prediction is far away from the truth. The optimizer notices that the loss score is high, and compensates by making the appropriate adjustment to the weights for each of the layers in the network. At this point, we are ready to start over by selecting a different image in the training set and repeating this procedure.

The takeaway here is that network training is an iterative process that involves looping over a training data.

Transfer Learning Fundamentals

In **transfer learning**, knowledge acquired by one model is transferred and used to support the development of another model. There is a **source** task, and there is data available to train a model to perform this source task. The source model is used to influence training of a **target** model. The key idea behind transfer learning is that if the source and target tasks are sufficiently related, then the source model might provide some useful information to help improve the target model relative to a model trained from scratch on target data. This is especially important when the target data is scarce and unlikely to be sufficient on its own for training a model.

A basic two-step approach to transfer learning is: (1) Select a source model, and (2) Use the source model (with or without modifications) to continue training on some of the target data. Note that a source model can be developed by training a model on source data, or it can be selected from an assortment of pre-trained models that are publicly available.

In computer vision applications, it is common to use models pretrained on ImageNet, a dataset of 14 million images of everyday objects like animals, plants, cars, and food [3]. Models such as AlexNet, GoogLeNet, ResNet, DenseNet, Inception, and VGG are some of the publicly available choices. One approach to transfer learning involves passing images through some of the initial layers of one such pretrained model, in order to obtain some intermediate data which is then used to train a separate target model. In this context, the intermediate data derived from an image is called a **feature vector** and the process of using part of the source model to convert the image into a feature vector is called **feature extraction**. Even though ImageNet may not contain many clinical images, this approach can be appropriate for clinical tasks because there are visual similarities (edges, shapes, and texture) in the ImageNet data that may be relevant to a specialized target task [4].

Applications of Transfer Learning in Veterinary Medicine

Transfer learning has been applied across a range of tasks and specialty areas in veterinary medicine. Some examples were presented.

<u>Computer Vision</u>

- Radiology
    - Ultrasound
        - Detecting degenerative hepatic disease in canine liver US images [5]
    - MRI
        - Predicting grade of meningioma from canine MRI [6]
        - Distinguishing meningiomas and glioma in canine MRI [7]

    - X-ray
        - Interpreting canine thoracic radiographs [8]
        - Interpreting canine and feline thoracic radiographs [9]
        - Detecting hip joints and hip dysplasia from VD pelvic radiographs [10]
- Animal phenotyping
    - Dog breed classification [11,12]
    - Dog identification [12]
    - Sheep facial expression (pain vs. no pain) classification [13]
    - Estimating BCS in dairy cows [14]
- Ophthalmology
    - Classifying severity of canine ulcerative keratitis [15]
- Histopathology
    - Malignancy classification of canine mammary tumor [1]
    - Detecting neutrophil clusters in mouse lung [16]
    - Classifying normal vs. grade 1 canine soft tissue sarcoma [17]
- Pose estimation
    - Pose estimation in horses [18]
- Conservation medicine
    - Classifying marine animals from underwater images [19]
    - Unique animal identification [20]
    - Aquatic animal image classification [21]

<u>Non-visual tasks</u>

- Audio
    - Bird species classification from song [22]
    - Cat sound classification [23]
- Bioinformatics
    - Predicting antibiotic class to which a gene might confer resistance [24]
    - Predicting deleterious mutations in coding regions of mouse, dog, and cattle [25]
    - Pretrained “Evolutionary Scale Model” [26]
- Text
    - Assigning diagnosis codes to free-text clinical notes [27]
    - Determining indications for antibiotic administration from clinical note text [28]
- Electrocardiogram analysis
    - ECG beat classification in horses [29]

Discussion

Transfer learning has been used to solve a diverse set of problems in veterinary medicine. In the domain of computer vision, many studies leverage neural network models pretrained on ImageNet, but note that such studies may not actually use the words “transfer learning” and they are not always explicit about how the source models were pretrained. It is also not unusual for studies to apply transfer learning without assessing an alternative methodology that does not involve transfer learning [1,5,7,8,10–13,15,16,19–21,23]. This seems to indicate that the described methods are widely accepted in vision applications.

In some domains (e.g. outside of computer vision), it is more typical for a study to compare a transfer learning-based approach to an alternative method such as a model trained from scratch on the dataset of interest [6,9,14,17,18,22,24,25,27–29]. These studies show that sometimes transfer learning is the most effective approach, but occasionally a model trained from scratch is most appropriate [6,14]. The latter may be true when there is insufficient target data to train a larger model typical in some transfer learning scenarios, or when the source data is insufficiently similar to the target data with respect to the desired prediction task. In particular, pretraining on human data can improve learning on veterinary tasks [24,25], or it may bias the model and cause decreased performance on veterinary data [9]. This suggests that transfer learning is an exciting and powerful tool, but that its utility needs to be confirmed on any particular task of interest.

Conclusion

This presentation illustrated how human data can be leveraged to build smarter veterinary AI. However, depending on the task, human clinical data might be too specialized and may reduce the effectiveness of a model intended to be trained on veterinary data. Alternatively, there are large amounts of generic human-generated data or unlabeled veterinary clinical data that may provide utility in these cases.

Because transfer learning is a method for leveraging the power of larger datasets, it can be a powerful tool in veterinary medicine. It should be considered wherever feasible, but it will not be the best approach in every situation. Results should always be subjected to a fair evaluation to determine if transfer learning is beneficial for a particular task.

References

1. Kumar A, Singh SK, Saxena S, Lakshmanan K, Sangaiah AK, Chauhan H, et al. Deep feature learning for histopathological image classification of canine mammary tumors and human breast cancer. Inf Sci. 2020;508: 405–421. doi:10.1016/j.ins.2019.08.072

2. Chollet F. Deep Learning with Python. 1st edition. Shelter Island, New York: Manning Publications; 2017.

3. Deng J, Dong W, Socher R, Li L-J, Li K, Fei-Fei L. ImageNet: A large-scale hierarchical image database. 2009 IEEE Conference on Computer Vision and Pattern Recognition. 2009. pp. 248–255. doi:10.1109/CVPR.2009.5206848

4. Zeiler MD, Fergus R. Visualizing and Understanding Convolutional Networks. In: Fleet D, Pajdla T, Schiele B, Tuytelaars T, editors. Computer Vision – ECCV 2014. Cham: Springer International Publishing; 2014. pp. 818–833. doi:10.1007/978-3-319-10590-1_53

5. Banzato T, Bonsembiante F, Aresu L, Gelain ME, Burti S, Zotti A. Use of transfer learning to detect diffuse degenerative hepatic diseases from ultrasound images in dogs: A methodological study. Vet J. 2018;233: 35–40. doi:10.1016/j.tvjl.2017.12.026

6. Banzato T, Cherubini GB, Atzori M, Zotti A. Development of a deep convolutional neural network to predict grading of canine meningiomas from magnetic resonance images. Vet J. 2018;235: 90–92. doi:10.1016/j.tvjl.2018.04.001

7. Banzato T, Bernardini M, Cherubini GB, Zotti A. A methodological approach for deep learning to distinguish between meningiomas and gliomas on canine MR-images. BMC Vet Res. 2018;14: 317. doi:10.1186/s12917-018-1638-2

8. Banzato T, Wodzinski M, Burti S, Osti VL, Rossoni V, Atzori M, et al. Automatic classification of canine thoracic radiographs using deep learning. Sci Rep. 2021;11: 3964. doi:10.1038/s41598-021-83515-3

9. Boissady E, Comble A de L, Zhu X, Hespel A-M. Artificial intelligence evaluating primary thoracic lesions has an overall lower error rate compared to veterinarians or veterinarians in conjunction with the artificial intelligence. Vet Radiol Ultrasound. 2020;61: 619–627. doi:10.1111/vru.12912

10. McEvoy FJ, Proschowsky HF, Müller AV, Moorman L, Bender-Koch J, Svalastoga EL, et al. Deep transfer learning can be used for the detection of hip joints in pelvis radiographs and the classification of their hip dysplasia status. Vet Radiol Ultrasound. 2021;62: 387–393. doi:10.1111/vru.12968

11. Borwarnginn P, Thongkanchorn K, Kanchanapreechakorn S, Kusakunniran W. Breakthrough Conventional Based Approach for Dog Breed Classification Using CNN with Transfer Learning. 2019 11th International Conference on Information Technology and Electrical Engineering (ICITEE). 2019. pp. 1–5. doi:10.1109/ICITEED.2019.8929955

12. Tu X, Lai K, Yanushkevich S. Transfer Learning on Convolutional Neural Networks for Dog Identification. 2018 IEEE 9th International Conference on Software Engineering and Service Science (ICSESS). 2018. pp. 357–360. doi:10.1109/ICSESS.2018.8663718

13. Noor A, Zhao Y, Koubaa A, Wu L, Khan R, Abdalla FYO. Automated sheep facial expression classification using deep transfer learning. Comput Electron Agric. 2020;175: 105528. doi:10.1016/j.compag.2020.105528

14. Rodríguez Alvarez J, Arroqui M, Mangudo P, Toloza J, Jatip D, Rodriguez JM, et al. Estimating Body Condition Score in Dairy Cows From Depth Images Using Convolutional Neural Networks, Transfer Learning and Model Ensembling Techniques. Agronomy. 2019;9: 90. doi:10.3390/agronomy9020090

15. Kim JY, Lee HE, Choi YH, Lee SJ, Jeon JS. CNN-based diagnosis models for canine ulcerative keratitis. Sci Rep. 2019;9: 14209. doi:10.1038/s41598-019-50437-0

16. Niazi MKK, Beamer G, Gurcan MN. An application of transfer learning to neutrophil cluster detection for tuberculosis: efficient implementation with nonmetric multidimensional scaling and sampling. Medical Imaging 2018: Digital Pathology. International Society for Optics and Photonics; 2018. p. 1058108. doi:10.1117/12.2292249

17. Rai T, Morisi A, Bacci B, Bacon NJ, Thomas SA, Ragione RML, et al. An investigation of aggregated transfer learning for classification in digital pathology. Medical Imaging 2019: Digital Pathology. International Society for Optics and Photonics; 2019. p. 109560U. doi:10.1117/12.2511754

18. Mathis A, Biasi T, Schneider S, Yuksekgonul M, Rogers B, Bethge M, et al. Pretraining Boosts Out-of-Domain Robustness for Pose Estimation. 2021. pp. 1859–1868. Available:

https://openaccess.thecvf.com/content/WACV2021/html/Mathis_Pretraining_Boosts_Out-of-Domain_Robustness_for_Pose_Estimation_WACV_2021_paper.html

19. Liu X, Jia Z, Hou X, Fu M, Ma L, Sun Q. Real-time Marine Animal Images Classification by Embedded System Based on Mobilenet and Transfer Learning. OCEANS 2019 - Marseille. 2019. pp. 1–5. doi:10.1109/OCEANSE.2019.8867190

20. Van Zyl TL, Woolway M, Engelbrecht B. Unique Animal Identification using Deep Transfer Learning For Data Fusion in Siamese Networks. 2020 IEEE 23rd International Conference on Information Fusion (FUSION). 2020. pp. 1–6. doi:10.23919/FUSION45008.2020.9190426

21. Yuan H, Zhang S, Qin E, Zhou H. Aquatic Animal Image Classification Technology Based on Transfer Learning and Data Augmentation. J Coast Res. 2020;105: 129–133. doi:10.2112/JCR-SI105-027.1

22. Ntalampiras S. Bird species identification via transfer learning from music genres. Ecol Inform. 2018;44: 76–81. doi:10.1016/j.ecoinf.2018.01.006

23. Pandeya YR, Lee J. Domestic Cat Sound Classification Using Transfer Learning. Int J Fuzzy Log Intell Syst. 2018;18: 154–160. doi:10.5391/IJFIS.2018.18.2.154

24. Hamid M-N, Friedberg I. Transfer learning improves antibiotic resistance class prediction. bioRxiv. 2020; 2020.04.17.047316. doi:10.1101/2020.04.17.047316

25. Plekhanova E, Nuzhdin SV, Utkin LV, Samsonova MG. Prediction of deleterious mutations in coding regions of mammals with transfer learning. Evol Appl. 2019;12: 18–28. doi:10.1111/eva.12607

26. Rives A, Meier J, Sercu T, Goyal S, Lin Z, Liu J, et al. Biological structure and function emerge from scaling unsupervised learning to 250 million protein sequences. Proc Natl Acad Sci. 2021;118. doi:10.1073/pnas.2016239118

27. Zhang Y, Nie A, Zehnder A, Page RL, Zou J. VetTag: improving automated veterinary diagnosis coding via large-scale language modeling. Npj Digit Med. 2019;2: 1–8. doi:10.1038/s41746-019-0113-1

28. Hur B, Baldwin T, Verspoor K, Hardefeldt L, Gilkerson J. Domain Adaptation and Instance Selection for Disease Syndrome Classification over Veterinary Clinical Notes. Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing. Online: Association for Computational Linguistics; 2020. pp. 156–166. doi:10.18653/v1/2020.bionlp-1.17

29. Van Steenkiste G, van Loon G, Crevecoeur G. Transfer Learning in ECG Classification from Human to Horse Using a Novel Parallel Neural Network Architecture. Sci Rep. 2020;10: 186. doi:10.1038/s41598-019-57025-2

References for the ECG project discussion (omitted from these notes):

Duerr F., Carr A., Duke T., Shmon C., Monnet E., 2007. Prevalence of perioperative arrhythmias in 50 young, healthy dogs. Canada Veterinary Journal 48, 169-177.

Noszczyk-Nowak A., Michalek M., Kaluza E., Cepiel A., Paslawska U., 2017. Prevalence of arrhythmias in dogs examined between 2008-2014. Journal of Veterinary Research 61, 103-110.

Estrada A., Spake A., Kleman M., Blishok-Lapekas D., Margiocco M., Gentile-Solomon J., Piscitelli N., Szlosek D., 2020. Diagnostic accuracy of computer aided electrocardiogram analysis in dogs. Journal of Small Animal Practice 1-5.

Hannun Y., Rajpurkar P., Haghpanahi M., Tison G., Bourn C., Turakhia M., Ng A., 2019. Cardiologist-level arrhythmia detection and classification in ambulatory electrocardiograms using a deep neural network. Nature Medicine 25, 65-69.