Visualizing (Veterinary) Medical Data Sets with Jetstream

Alan Nguyen
Purdue University
West Lafayette, Indiana

Yvan Pierre Jr.
Valencia College
Orlando, Florida

Scott Birch
Indiana University
Bloomington, Indiana

Winona Snapp-Childs
Indiana University
Bloomington, Indiana

Evan Suggs
University of Tennessee at Chattanooga
Chattanooga, Tennessee

ABSTRACT

Computed tomography (CT) is a diagnostic imaging test using x-rays to create multiple detailed images of internal organs, bones, soft tissue and blood vessels. It produces a data set of thin, cross-sectional “slices” for viewing and is much more detailed than conventional radiography. Clinicians use CT examination to diagnose cancers, detect abnormal blood vessels, discover disorders of the abdomen, bones, and joints, and to plan surgical interventions such as heart defect or vascular repair. Dedicated visualization workstations allow radiologists to make high-resolution examinations of diagnostic data, but understanding the image stacks can be challenging for clinicians without specialized skills, training, and experience. To aid and enhance diagnostic evaluation, we explored a cloud-based workflow using Jetstream. CT data sets were segmented or translated into regions-of-interest (ROI) and/or volumetric 3D reconstructions which were then exported as polygonal 3D surface models. Using data sets obtained via CT from a variety of animal species, this project focused on the process of compiling a medical imaging/segmentation workstation instance with open source software on Jetstream, importing sample data sets into the imaging software, viewing 2D image sequences volumetrically, setting custom transfer functions based on tissue density, and segmenting the anatomy into multiple ROI for export as stereolithography files. Post-processing and polygon mesh editing techniques such as smoothing, transient reduction, and decimation were employed as the model was optimized for 3D printing or online distribution. Results were rendered into 2D graphical representations, and the 3D models were deployed into interactive or virtual reality environments, or were additively-manufactured (3D printed) into real-world objects for visual and tactile examination. After workflows were verified and vetted, the Jetstream medical segmentation VMs were made available for others to view and/or segment their own volumetric data sets.

Woodstock ’18, June 03–05, 2018, Woodstock, NY

ACM ISBN 978-1-4503-9999-9/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

CCS CONCEPTS

• **Medical Visualization → segmentation**; *Additive Manufacturing*; Radiography; • **Machine Learning** → Convolutional Neural Networks.

KEYWORDS

Machine Learning, Cloud Computing, Neural Networks, CNN, 3D Segmentation, CT scans, DICOM, ParaView, 3D Slicer

ACM Reference Format:
Alan Nguyen, Yvan Pierre Jr., Scott Birch, Winona Snapp-Childs, and Evan Suggs. 2018. Visualizing (Veterinary) Medical Data Sets with Jetstream. In *Woodstock ’18: ACM Symposium on Neural Gaze Detection, June 03–05, 2018, Woodstock, NY*. ACM, New York, NY, USA, 6 pages. https://doi.org/10.1145/1122445.1122456

1 INTRODUCTION

Three-dimensional visualizations such as additive manufacturing of 3D printed models are principle resources in veterinary medicine. When bones are broken or internal organs are damaged, CT scans are usually captured for diagnostic purposes. The CT scans provide a comprehensive map of an organism’s ROI that can be refined further to a specific segment or layer of the body such as the skin, muscle tissue, or bones. As single or nested segments, digital volumetric images and 3D printed models can be derived for enhancing visualizations of medical CT data to optimize the quality of patient care provided by veterinarians[11].

1.1 Background

Conventionally, the 2D image stack of scans from a CT dataset is manually analyzed and segmented, a time-consuming process that can delay patient diagnosis and further action. Particularly with advancements in medical technology, CT scanners are capable of producing data sets of thousands of image slices with high resolution and detail. Such strenuous CT data may require high computer processing power. As a result, through the use of Indiana University’s Jetstream Application Programming Interface (API) [?], we planned to create a virtual machine workstation to help facilitate the management of veterinary CT scans and development of 3D representations of animal segments.

Our workflow comprised two imaging software applications, ParaView and 3D Slicer, which were used to create digital 3D volumetric reconstructions of 2D CT data sets. Those were then segmented for the ROI of a given organism. We also incorporated a convolutional neural network (CNN), a machine learning algorithm for classifying CT image slices of varying animals species. We then evaluated the overall efficacy and applicability of the final workstation for progressing visualizations of veterinary data sets.

1.1.1 *Jetstream.*
Jetstream is an interactive cloud-based computing and data analysis platform. It offers on-demand access to supercomputing power for researchers in a broad range of disciplines, providing support that was previously not available to them for their big-data computational needs[12]. Funded by the National Science Foundation (NSF), Jetstream bridges a connection between its existing users in the Extreme Science and Engineering Discovery Environment (XSEDE) and new users that have not yet utilized high performance computing (HPC) in their research. Jetstream strives to increase diversity in usage and accessibility of HPC, large scale memory, and data storage and visualization resources.

Currently, Jetstream allows researchers and students remote access to its library of various virtual machines (VM) online, and to select the one most tailored to their needs. Because these virtual machines are pre-configured with an operating system, storage allocations, and software environments, Jetstream provides a user-friendly interface for researchers with minimal HPC experience or computer science background.

1.1.2 *Virtual Machine Setup.*
We first began creating our workstation on the IU Jetstream cloud through the Openstack Horizon Dashboard from a pre-configured image, which is an inactive template of a virtual machine. The image we used was a Jetstream API with the CentOS7 operating system. An activated or launched image is known as an instance, in which the the VM or server would be running. In the process of doing so, we selected a flavor, which entails the size specifications of the running instance. Our flavor was the m1.xxlarge that allows for 120 GB of RAM, 44 vCPUs and 60 GB of disk space. A volume, which is a virtual file-system that can be attached to or detached from an instance, was later mounted to provide additional data storage space and enable easy data transfer between instances. Our volume had a size of 60 GB.

2 METHODS
To create 3D models from CT data sets, we began by creating volumetric images and manual segmentation of the species of interest. This was initially done on the virtual desktop of our VM workstation, which was viewed via the console graphical user interface (GUI) on the Openstack Horizon dashboard. The CT image classification CNN was saved and run on the current VM.

2.1 Data
The CT data sets were in the Digital Imaging and Communications in Medicine (DICOM) file format, which is standard for managing medical information. All imaging data and resources were transported between local machines and the VM through Globus, a secure file transfer service via a web browser. An ant (Zasphinctus Obamai) data set[3](see Figure 1) and a canine dataset were used to demonstrate the additive manufacturing workflow. The ant DICOM data set was retrieved from the Dryad Digital Repository[1]. In addition to the ant and dog, data sets of an armadillo, horse, snake, and turtle were used for the CT image classification CNN.

Figure 1: Zasphinctus Obamai, Credit: April Nobile[7]

2.2 ParaView
Our workflow began on ParaView, producing 3D volumetric reconstructions of a dog and ant. In doing so, the DICOM directory of files for both organisms were imported and viewed under the volume representation setting to enable a volumetric rendering of the 2D sequence of image slices. We used the "Color Map Editor" to create custom transfer functions per animal dataset and to map field of dataset values in various brightness, colors, and opacity levels. This feature allowed us to adjust the visibility of different structures of the body and external particles, which all vary in physical densities or appearance in the CT scans. Different densities, quantized on the Hounsfield scale[6], could be assigned variations of colors and levels of opacity to show subtle or extreme contrast between structural features. Screenshots of the volume renderings were saved in a variety of angles on either the sagittal, coronal, or axial anatomical plane. The transfer functions used for editing were saved as presets to be reused or referenced for future analysis.

In Figure 2, a lateral view in the sagittal plane of a canine shows skeletal and cardiovascular features. The canine was scanned due to a defect in the heart. The cardiovascular system is emphasized in the red while the bones are visible in white. The bones are lightly shaded red presumably due to a layer of flesh that has a similar density to that of the heart and blood vessels, which was set for visibility. The transfer function used to visualize these structures is represented in the bottom right corner.

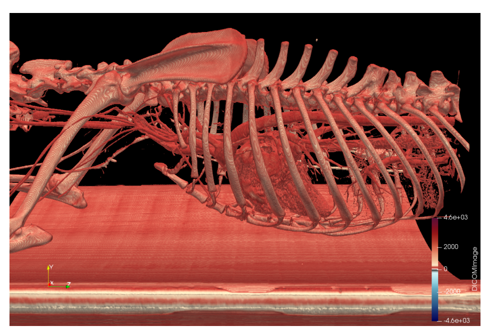

Figure 2: Dog Heart

In Figure 3, a lateral view in the saggital plane of the ant shows the exterior shell of the animal. The renderingâĂŹs corresponding color map editor is captured in Figure 4. On the editor, the density values increase from left to right (x-axis) and the corresponding structures of the chosen density can be made visible be by adjusting the created or selected points on the plot up (higher opacity) or down (lower opacity). Therefore, in respect to the ant, a mixture of light brown, green, black, and gray colors which correspond to lighter densities. Lighter densities of the ant would be representative of the skin covering the bones, which consequently would not be visible despite its corresponding density color (red) being set to full opacity.

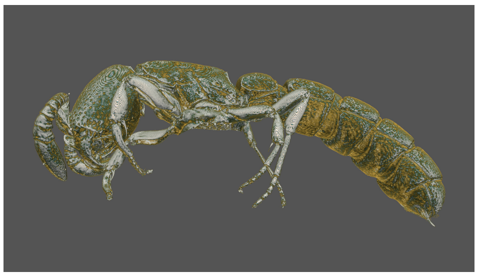

Figure 3: Ant's Full Body Dissection in ParaView

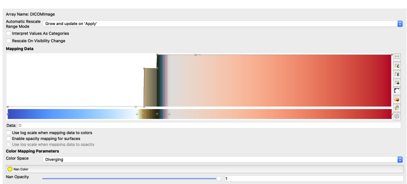

Figure 4: Ant's Map Editor in ParaView

2.3 3D Slicer

After volumetrically viewing the datasets in ParaView, we loaded the data into 3D Slicer for manual segmentation of the organism. Our goal was to segment layers of the body such as the exterior skin (epidermal), muscle tissue (musculature), or skeletal bones. This was accomplished using the segmentation and segment editor modules. Primarily in the segment editor, we used effects such as thresholding, islands, smoothing, erase, and paint to segment the desired layer of the body. Each segment was exported as a stereolithography (STL) file for further optimization on a local machine.

In Figures 4 and 5, the upper-half is the 3D view of the segment's current state in the segment editor. In this view, the segment could be edited, rotated, and magnified. The lower-half shows individual CT image slices in the axial, sagittal, and coronal anatomical plane, left to right respectively. The contrast of structural densities in each view of the CT scans can be adjusted to accent a specific feature of the body. Within the planes of the anatomical views, the presence of green shading depicts the areas and particles of the body that have been thresholded or, in other words, selected to be editable based on density values pertaining to the body's layers. Only the thresholded parts of the organism will be visible in the 3D view and remain as part of the segment unless removed otherwise.

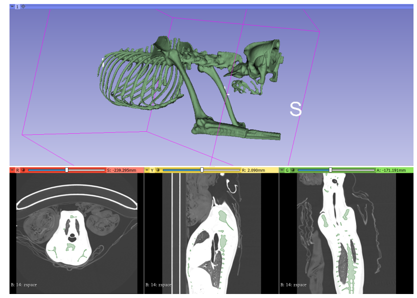

Figure 5: Dog Interior-Skeleton Segment

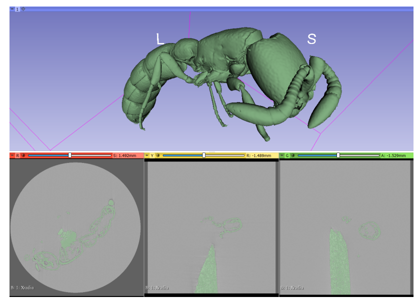

Figure 6: Ant's Exterior-Skin Segment

2.4 Machine Learning

Another approach to visualizing CT datasets involves using image segmentation algorithms. These algorithms can automate the process of segmentation and detection of boundaries to output the isolation of specific ROIs in an organism. Pham et al. [8] discuss a variety of methods for semi- and fully-automated segmentation of anatomical medical images. Due to the growing amount of data and large sizes of medical image stacks, there is a need for higher computer power to process and analyze MRI or CT scans of patients. The study touches upon segmentation algorithms in the realm of machine learning, specifically artificial neural networks. We used a convolutional neural network (CNN), which varies from an ANN structurally in the sense that the last layers of a CNN are fully connected, whereas the neurons are connected to every other one in an ANN[4]. CNNs are most suitable for image processing as the basis for such algorithms were developed for Computer Vision in Deep Learning. The CNN, which focused on image classification of a 2D image slice, we incorporated was founded on the work of Dakila Ledesma. Background information about convolutional neural networking and full documentation on the present algorithm can be publically accessed online[5]. Our working algorithm was developed in a Conda environment written in Python using Keras and Tensorflow libraries. This CNN can be used as a basis for transpiring a medical image segmentation to predominantly automate the process of creating structural segments as opposed to a manual based segmenting process.

Our CNN was trained to predict the CT scan to which animal or species it belongs. The input training and testing images used for data collection were converted DICOM files to PNG. Per each organism dataset, 80 percent of the images were allotted for training and 20 percent was allotted for testing. The figure below represents the specific image allocations and total images in each organism's CT dataset.

CNN IMAGE ALLOTMENTS						
Image Allocation	Tested CT Datasets					
	ANTS	DOGS	ARMADILLO	HORSE	SNAKE	TURTLE
Train (80%)	190	214	107	102	221	140
Test (20%)	48	54	27	26	55	35
Total	238	268	134	128	276	175

Figure 7: CNN Image Input Allotments

The CNN essentially extracts the pixel data of an image into an array and perform a series of matrix operations and transformation to the pixel values of the image. The pixel value describes the brightness or color information in each pixel that makeup the patterns, shapes, and feature boundaries throughout the image. Thus, the manipulation of the pixel values is the CNN's process of isolating different structures of an organism's CT scan by the appearance of different contrasts in brightness levels or shades of gray within the image.

Our CNN consisted of three primary types of layers: convolutional, pooling, and fully connected. These layers and their coinciding functions perform the matrix operations needed to define the unique image features of the CT images associated with the organism's ROIs. The fully connected layers included those from the flatten layer to the the last dense layer. Figure 8 is a depiction of CNN structure used for training:

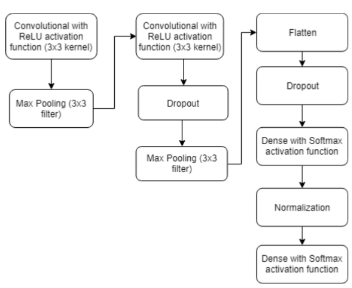

Figure 8: CNN Structure

The CT image classification CNN was tested for training accuracy and certainty of correctness and incorrectness of classifying an individual 2D CT image slice of an organism. For all testing runs, each epoch had a 100 image batch size. Relevant output information of the CNN included a training accuracy for every run of the algorithm as well as classification certainty percentages (6) of each organism whose image data was used for training, per tested image.

3 RESULTS

3D printed models of ant and dog segments were produced after additional rendering and optimization of the modeled segments were done with software on a local machine.

The CNN for CT image classification increased he number of training epochs from 25 to 50, 75, and 100. The 100 images in each batch from an individual CT data set per epoch were shuffled or randomized. We found that as the number of epochs increased, the training accuracy tended to improve as well.

TRAINING ACCURACY				
Number of Epochs	25	50	75	100
Accuracy (%)	93.24	96.575	97.84	98.353

Figure 9: Training Accuracy Results

The certainty results shown below were from the same run as the training accuracy test at 100 epochs. For correctness, the CNN had the highest percentage of certainty in classifying turtles with 95.063 percent, whereas the armadillo was classified with the least percentage of certainty at 14.152 percent. The classification of armadillo CT scans were particularly unique in the sense that the CNN had fairly had certainty percentages that an armadillo CT scan was either dog or turtle with 48.310 percent and 23.905 percent, respectively.

MEAN							
Tested CT Dataset	Certainty (%)						
	Correct	*Incorrect*					
		ANT	DOG	ARMADILLO	HORSE	SNAKE	TURTLE
ANT	88.530	—	3.907	1.756	1.404	2.086	2.318
DOG	77.721	4.457	—	6.514	3.645	4.146	3.512
ARMADILLO	14.152	7.689	48.310	—	5.664	0.280	23.905
HORSE	57.543	3.585	1.772	5.068	—	27.764	4.268
SNAKE	77.684	2.002	3.351	6.612	5.312	—	5.033
TURTLE	95.063	0.915	0.989	2.605	0.244	0.183	—

Figure 10: Certainty Results

Due to the non-normally distributed nature of the results across all of the datasets, a non-parametric ANOVA test was done for further statistical analysis. A Kruskal-Wallis rank sum test was performed on all species models with the following results and parameters: chi-square = 93.991, df = 5, $p < 0.0001$. The results of this test suggest that our organism data sets are significantly different overall. Further more, a Tukey Honest Significant Difference test (HSD) post-hoc comparison procedure was also conducted which indicated that the following data sets were significantly not different from each other: ant and turtle; ant, dog and snake; horse; armadillo.

4 DISCUSSION

In evaluation of the final state of our VM, current 3D imaging software and visualization algorithms can be considered as the groundwork towards cultivating a more effective workstation for creating 3D models and visualizing veterinary medical data sets.

The design and optimization process for generating the 3D printed models of the dog and ant segments was largely completed on a local machine due to inadequate software compatibility with Jetstream. Though ParaView and 3D Slicer were the main software programs used to create the 3D models, additional open-source 3D imaging software (i.e., GIMP, MeshLab, and Blender) were initially downloaded onto the VM as well and were briefly explored as part of the modeling workflow. Further investigation into these and other open-source software could be carried out for improving the current VM workstation on the Jetstream platform for manual visualization and segmentation of veterinary CT datasets.

Previous additive manufacturing work has been done commonly on a local machine as well, but more so with with closed-source software such as OsiriX and Materialise 3-Matic. For example, these programs have been utilized by more technically experienced users, who have visualized and created 3D printed models of canine cardiovascular regions, specifically in the case of a heart defect such as patent ductus arteriosus (PDA). Saunders, et al. conducted a study about analyzing and visualizing this health issue in dogs[10]. The study reported how using the software mentioned above allowed them to study the anatomy of the PDA region and print models to enhance training practice for treating this problem through a particular surgical operation. Based upon our VM, similar advanced open-source 3D modeling software could be implemented and introduced to potential users of the workstation in a manner that adheres to Jetstream's user-friendly precedent. In doing so, a reputable workflow in logic and efficiency for creating advanced 3D visualizations of veterinary medical data sets could be established and thus be utilized among users beyond additive manufacturers or veterinarians.

In regard to our CNN, the provided results and analysis above show our algorithm to be relatively successful in predicting whether a CT scan belonged to an ant, dog, armadillo, horse, snake, or turtle. The success in training accuracy and thus the classification certainties per image can be attributed to limitations on the number of training images and resolution or quality of those images within each organism's data set. The turtle data set had the highest correctness certainty with fairly low training images, which could be a result of having the highest quality resolution across all data sets. The ant, dog, and snake had moderately high correctness certainties in respect to the remaining data sets. This could be due to having the most training data at a margin slightly above or below 200 images. In contrast, the armadillo and horse had the fewest training images and the lowest resolution resulting in the two lowest accuracies of correctness. Furthermore, similarity in CT scan patterns and features from different organism data sets can also be taken into account for low correctness accuracy, which would hold true for the armadillo data set in which its structures have a resemblance of those of a dog or turtle.

Similarly to prior workflows of 3D visualization processes and mediums, machine learning in veterinary medicine has been proven advantageous in the recent past as well. CNN imaging classification of medical data sets have been utilized in deducing specific predictions such as identifying different grades or types of meningiomas, a tumor covering or surrounding the brain and spinal regions in dogs, which was accomplished by Banzato, et al[2]. More generally, CNNs can make predictions on whether an organism has a normal or abnormal lung based on image features or patterns in radiography images, which was an algorithm construct explored by Yoon, et al[14]. Not only have CNN image analyses been useful for classification, but they also have benefited researchers and veterinarians in retrieving quantitative information from medical image data sets. For instance, Vinicki, et al achieved a CNN that was inputted training microscopic sets images of felines' blood smears to test additional blood smear images for counts of immature and mature electrolytes produced based on their identified locations within respective images[13]. These examples of CNN applications in veterinary medicine have aligned with same goal that we had set out for our CT image classification algorithm to meet, which was to employ the advantages of deep learning and thus offer more efficient, accurate, and cost-effective resources for implementation by veterinarians, who may have minimal background in computer science[13].

Moreover, the algorithm on the present VM can be used as a building-block towards a more refined CNN that could perform a fully-automated segmentation process of an organism. Provided more time in exploring and testing machine learning algorithms, we had planned to formulate a variation of a CNN called a Fully Convolutional Network (FCN) in which the neural network would consist of primarily convolutional layers to perform matrix operations and transformations on the pixel data of the images. The ideal FCN pertaining to the intentions of this VM workstation for visualizing veterinary medical data sets would be to automate the segmentation process, which could be done manually on software such as 3D Slicer. This could be accomplished by training the FCN to recognize any structural layer, such as epidermal layer, musculature layer, or skeletal layer, in each 2D scan in the CT dataset and then outputting the 3D volumetric reconstruction of the segment with any ROI. Additionally, in our predicament of limited availability of data images, an FCN would have been more appropriate than a traditional CNN for our project, as an FCN can constructed to be suitable for low amounts of training data[9].

4.1 Conclusion

Our current virtual machine ultimately provides a foundational workstation for visualizing veterinary medical data sets on Jetstream. Future improvements or additions in the 3D modeling software stack of ParaView and 3D Slicer could be explored to enhance the real-time visualization analysis of image data sets. In the same pursuit of progress, our CT image classification CNN could be evolved into an FCN, effectively alleviating the manual labor of segmenting patient-organisms' CT scans and creating 3D printed models. This work on Jetstream has demonstrated how VMs such as the one we had created can of invaluable use to potential users from veterinarians to graphic designers that may have little to no experience with cloud or high-performance computing. Given Jetstream's interactive and on-demand nature, the visualization resources that we have vetted and tested may be made available for as tool in the veterinary medicine file.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1445604. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation. Thank you to Harmony Jankowski and all reviewers of earlier versions of this extended abstract. We would also like to thank everyone at Indiana University UITS Research Technologies, Pervasive Technology Institute provided help and guidance throughout our REU project: Mike Lowe, George Turner, Steve Bird, Sanjana Sudarshan,Jeremy Fischer for Jetstream

Administration and Managment; Eric Wernert for visualization software assistance; Jefferson Davis and Laura Huber for machine learning and neural network advising; Sheri Sanders Bhavya Papudeshi for UNIX and command line guidance.

REFERENCES

[1] [n. d.]. Hita Garcia F, Fischer G, Liu C, Audisio TL, Economo EP (2017) Data from: Next-generation morphological character discovery and evaluation: an X-ray micro-CT enhanced revision of the ant genus Zasphinctus Wheeler (Hymenoptera, Formicidae, Dorylinae) in the Afrotropics. ([n. d.]).

[2] T Banzato, GB Cherubini, M Atzori, and A Zotti. 2018. Development of a deep convolutional neural network to predict grading of canine meningiomas from magnetic resonance images. *The Veterinary Journal* 235 (2018), 90–92.

[3] Francisco Hita Garcia, Georg Fischer, Cong Liu, Tracy L Audisio, and Evan P Economo. 2017. Next-generation morphological character discovery and evaluation: an X-ray micro-CT enhanced revision of the ant genus Zasphinctus Wheeler (Hymenoptera, Formicidae, Dorylinae) in the Afrotropics. *ZooKeys* 693 (2017), 33.

[4] I Gogul and V Sathiesh Kumar. 2017. Flower species recognition system using convolution neural networks and transfer learning. In *2017 Fourth International Conference on Signal Processing, Communication and Networking (ICSCN)*. IEEE, 1–6.

[5] Dakila Ledesma. 2019. cnn-tutorial.md. https://github.com/bgq527/intro-ml/blob/master/cnn-tutorial.md [Online; accessed <7/2/2019>].

[6] MH Lev and RG Gonzalez. 2002. CT angiography and CT perfusion imaging. In *Brain mapping: the methods*. Elsevier, 427–484.

[7] April Nobile. 2017. AntWeb.org image of Order:Hymenoptera Family:Formicidae Genus:Sphinctomyrmex Species:Sphinctomyrmex undet Specimen:casent0178218 View:profile. http://www.antwiki.org/wiki/File:Z_obamai_casent0178218_p_1_high.jpg#file [Online; accessed <7/26/2019>].

[8] Dzung L Pham, Chenyang Xu, and Jerry L Prince. 2000. Current methods in medical image segmentation. *Annual review of biomedical engineering* 2, 1 (2000), 315–337.

[9] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*. Springer, 234–241.

[10] AB Saunders, KR Doocy, and SA Birch. 2019. A pictorial view of the three-dimensional representation and comparative two-dimensional image orientation derived from computed tomography angiography in a dog with a patent ductus arteriosus. *Journal of veterinary cardiology* 21 (2019), 34–40.

[11] AB Saunders, L Keefe, SA Birch, MA Wierzbicki, and DJ Maitland. 2017. Perceptions of transcatheter device closure of patent ductus arteriosus in veterinary cardiology and evaluation of a canine model to simulate device placement: a preliminary study. *Journal of veterinary cardiology* 19, 3 (2017), 268–275.

[12] Craig A Stewart, Timothy M Cockerill, Ian Foster, David Hancock, Nirav Merchant, Edwin Skidmore, Daniel Stanzione, James Taylor, Steven Tuecke, George Turner, et al. 2015. Jetstream: a self-provisioned, scalable science and engineering cloud environment. In *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*. ACM, 29.

[13] Krunoslav Vinicki, Pierluigi Ferrari, Maja Belic, and Romana Turk. 2018. Using Convolutional Neural Networks for Determining Reticulocyte Percentage in Cats. *arXiv preprint arXiv:1803.04873* (2018).

[14] Y Yoon, T Hwang, and H Lee. 2018. Prediction of radiographic abnormalities by the use of bag-of-features and convolutional neural networks. *The Veterinary Journal* 237 (2018), 43–48.