



PLEASE NOTE! THIS IS PARALLEL PUBLISHED VERSION /
SELF-ARCHIVED VERSION OF THE OF THE ORIGINAL ARTICLE

This is an electronic reprint of the original article.
This version *may* differ from the original in pagination and typographic detail.

Author(s): Sipola, Tuomo; Puuska, Samir; Kokkonen, Tero

Title: Model Fooling Attacks Against Medical Imaging: A Short Survey

Version: Publisher pdf

Year: 2020

Version: Published version

License: Attribution-NonCommercial 4.0 International

License url: <https://creativecommons.org/licenses/by-nc/4.0/legalcode>

Please cite the original version:

Sipola, T., Puuska, S. & Kokkonen, T. (2020). Model Fooling Attacks Against Medical Imaging: A Short Survey. *Information & Security: An International Journal*, 46(2), 215-224. DOI:

URL: <https://doi.org/10.11610/isij.4615>

Model Fooling Attacks Against Medical Imaging: A Short Survey

Tuomo Sipola (✉), ***Samir Puuska***, ***Tero Kokkonen***

*JAMK University of Applied Sciences, Jyväskylä, Finland,
<https://www.jamk.fi/en/Home/>*

ABSTRACT:

This study aims to find a list of methods to fool artificial neural networks used in medical imaging. We collected a short list of publications related to machine learning model fooling to see if these methods have been used in the medical imaging domain. Specifically, we focused our interest to pathological whole slide images used to study human tissues. While useful, machine learning models such as deep neural networks can be fooled by quite simple attacks involving purposefully engineered images. Such attacks pose a threat to many domains, including the one we focused on since there have been some studies describing such threats.

ARTICLE INFO:

RECEIVED: 08 JUN 2020

REVISED: 31 AUG 2020

ONLINE: 13 SEP 2020

KEYWORDS:

deep learning, artificial neural networks,
adversarial images, machine learning,
medical imaging



Creative Commons BY-NC 4.0

Introduction

Artificial Intelligence (AI) based solutions, especially deep learning based on neural networks, are widely used in the medical domain. For example, AI is used for helping and automatizing cancer diagnosis based on image data. The benefit of this approach is to relieve experts to work on more important tasks while automated systems can inspect images and give initial recommendations.

If an attacker can fool the AI processing, ramifications can be devastating. Such attacks may result in incorrect treatment procedures, causing extreme

circumstances with a worst-case scenario of losing human lives. In addition, wrong diagnoses could undermine the public trust in medical professionals. This paper presents a short survey of model fooling attacks against neural networks in the medical domain.

Fooling neural networks is an important subject because machine learning models are widely used in medicine for automating many processes and for helping with diagnosis. For example, Rai et al. proposed a convolutional neural network for healthcare assistant²⁹ while Rastgar-Jazi and Fernando used neural networks for detecting heart abnormalities from electrocardiogram (ECG) data.³⁰ Similarly, authors of³¹ used neural networks for prediction and prevention of heart attacks from ECG data while Murugesan and Sukanesh used neural networks for detecting brain tumours in electroencephalograms (EEG) signals.²⁴ Syam and Marapareddy discussed three different scenarios of classification problems, where one is skin lesion (cancer) classification from images.³⁴ As can be seen, these machine learning solutions are useful for many medical applications.

Effectiveness of neural network based deep learning is based on the used algorithm and learning data. If learning dataset is inadequate or contains incorrect information, results will be inaccurate. Similarly, if there are known weaknesses in the used algorithm, they can be compromised. In that sense, AI components can be attacked and fooled to behave incorrectly. As an example of a weakness, Afifi and Brown explore how white balance of photography impact the performance of deep neural networks,² while authors of²¹ generated adversarial noise for fooling the neural networks. Gu et al. discussed about gradient shielding method for understanding the vulnerabilities in neural networks.¹³ MOEA-APGA is an algorithm for achieving targeted attacks against neural networks,⁹ and another similar algorithm is called DeepFool implemented for computing perturbations that fool neural networks.²³ In a medical domain, Chuquicusma et.al. studied about fooling radiologists for lung cancer diagnosis.⁸ As can be seen, many such attack vectors exist.

As a powerful machine learning method, deep learning has also been applied to images related to pathology, for example, trying to classify images of cancer whole slide images (WSI). Serag et al. present an overview of the application of artificial intelligence for pathology and tissue analytics.³² As another example, convolutional neural networks have been used for nuclear segmentation, which is an important part of tissue cancer grading.¹⁸ Deliberately produced wrong segmentation could result in wrong diagnoses. Pre-trained convolutional neural networks have been compared to training from scratch using the Kimia Path24 dataset, with results indicating that pre-trained networks are quite competitive.¹⁵ Using such pre-trained models creates a possibility of hidden attacks trained into the model or abusing known deficiencies of such models.

In this study, we collected a list of relevant research papers concerning medical imaging and attacks against neural networks. We queried the publicly available Google Scholar database to identify publications relevant to deep neural network fooling, deep neural networks in medical imaging and deep neural

networks fooled in that domain. The results of this short survey should be useful for anyone trying to understand the vulnerabilities of neural networks in specific domains. Moreover, the use of them in medical imaging raises the question of reliability and robustness when targeted by such attacks. As can be seen, targeted attack against neural networks in medical domain is a realistic scenario. From the attacker perspective, medical domain can be considered as a valuable target because of the critical ramifications of possible attack. In addition, there are known vulnerabilities with neural networks that are highly used in medical domain.

Below, we present the medical imaging domain and then discuss about machine learning regarding that domain. Next sections describe the state of fooling deep neural networks and how it has been applied to the medical domain. Finally, we present our concluding thoughts.

Short Introduction to Whole Slide Images in Cancer Diagnosis

Quick and affordable laboratory cancer diagnosis methods are of great importance. One of the well-established methods is light microscopy with a stain, such as haematoxylin and eosin (H&E). The H&E stain makes various tissue components visible, allowing diagnosis based on e.g. their morphological features.¹⁴

The advent of digital pathology and whole-slide imaging (WSI) have provided a computerized way to analyse and share the results of light microscopy. By digitizing the tissue images, a variety of automated methods can be used to perform image analysis, annotation, and workflow improvements. Turning glass slides to a digital format requires a slide scanner, which digitizes the slide using specialized format that allows e.g. various zoom levels and metadata to be stored in one data file. This data can then be easily shared, further processed using a variety of tools, and even easily used in teaching in a virtualized microscopy environment.¹

Distinguishing between benign and malignant tumours is essential for accurate prognosis. One of the features that separate the two is differentiation and anaplasia. In general, benign tumours consist of cells that resemble the tissue where they originated from. They retain much of the functionality and morphology of their non-transformed counterparts but may invade surrounding tissue. Malignant tumours, on the other hand, lose their resemblance to their normal counterparts and become undifferentiated (anaplastic). This change results in noticeable change in cell morphology, and it is possible to observe this using light microscopy and stains. These observable changes include variations in size and shape, nuclear abnormalities and atypical mitoses. Assigning a value to this differentiation is called grading. The criteria and schemes are dependent on the type of tumor.¹⁹

For breast cancer, observing mitoses has been shown to be a good predictor for tumour development and prognosis. In order to proliferate, tumour cells need to overcome various limitations that prevent ordinary cells from dividing indefinitely. These mutations may result in increased cell-cycle activity, and even majorly affect the mitosis process itself by causing atypical-looking cell

divisions which may be visually observed using light microscopy.²⁷ Figure 1 shows an example of breast cancer WSI from Al-Janabi et al.⁴

Detecting abnormal morphology and quantifying the number of various cell features is a good candidate for automatization via machine learning and computer vision methods. WSI with sufficient quality can be automatically annotated. Digital pathology is expected to improve convenience and quality of the process. Nam et al. provide an introduction to digital pathology aimed at healthcare professionals.²⁵ Furthermore, Komura and Ishikawa made a short review of machine learning methods for histopathological image analysis, listing seven whole slide image datasets and 21 hand annotated histopathological datasets.¹⁶




Figure 1: Example of WSI showing several breast resections with infiltrative ductal carcinoma.

Figure courtesy of Al-Janabi et al.,⁴ distributed under the terms of the Creative Commons Attribution License.

Fooling Deep Neural Networks

Deep neural networks and deep learning in general refers to a field of study, where complex concepts are learned from simpler representations by creating an interconnected network of activation functions and weights.¹¹ Due to their nature, these networks may contain flaws which make them susceptible to various classes of errors. These imperfections may be used maliciously to force the network into making an erroneous prediction. There have been several successful attempts at creating methods to fool deep neural networks.

One approach is to give an adversarial image as input to the classifier. Nguyen, Yosinski and Clune used an evolutionary algorithm optimization method to generate unrecognizable images to the human eye. Those images fooled a neural network to classify them as an object with high certainty, even though it should not have. They describe these images as costly exploits that could be used against deep neural networks.²⁶ Moosavi-Dezfooli, Fawzi and Frossard propose the DeepFool algorithm that efficiently generates adversarial images and quantifies the robustness of image classifiers.²³ An adversarial image is wrongly classified as something else than what the image clearly represents to the human eye. In the paper, a slight perturbation was added to an image of an animal to misclassify a whale as a turtle. Furthermore, it is possible to create adversarial 2D images robust to noise, distortion and affine transformations, and even adversarial 3D printed objects (a turtle).⁶

Adversarial patches are images that can be placed inside another image to fool a neural network classifier. Brown et al. have shown the effectiveness of such images.⁷ It is easy to see that inserting such patches to medical images could yield similar results, resulting in a false classification.

Research has already addressed cases of changing only one pixel of an image to cause it to be classified as another object.³³ It is remarkable that a change of colour in one pixel could fool the neural network. A move towards a more theoretical understanding of one pixel-attacks and incorrect mapping to low dimensional manifold has also been proposed. This makes it easy to find localized areas where one-pixel attacks should be more effective.¹⁷

Backdoored images can be created when attacking the learning stage of a neural network. These malign models can be deployed to production, and the fault is only revealed when the bad image is given as an input, resulting in wrong classification. Outsourced training opens the possibility of creating backdoored neural networks that behave badly on input specified by the attacker.¹² In a similar scheme, called poisoning attack, artificially poisoned data being sent to a model gradually change the model to conform to the attacker's goals. Yang et al. used an autoencoder (instead of the more traditional direct gradient method) to generate poisoned input data for deep neural networks.³⁷

The evident vulnerability of neural networks against several types of attacks is alarming because these methods are being proposed in several real-world domains. See ³ for a survey of adversarial attacks against deep learning in computer vision. The authors not only list several attacks but also include defences. They conclude that there is a threat against safety and security critical applications.

Figure 2 shows a schematic presentation of the possible attack routes described above. In this paper we have identified two parts of the AI process, which could be targeted.

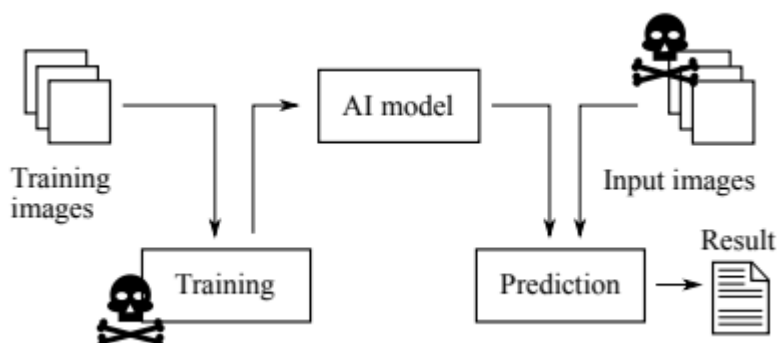


Figure 2: A schematic picture of AI model fooling.

The input images could be altered using adversarial images, patches or one-pixel attacks. In addition, the training process itself could be tampered with.

Fooling Deep Learning in Medical Imaging and Pathology

Although a rather new concern, the vulnerabilities intrinsic to neural network solutions have been identified by the medical community. For evident health reasons the accuracy and robustness of methodology in the medical domain is very important. Tizhoosh and Pantanowitz list challenges and opportunities related to artificial intelligence and digital pathology. One of the challenges concerns adversarial attacks and the shakiness of deep decisions made by neural networks.³⁶ This fundamental lack of robustness could be one avenue of future research.

Adversarial examples in medical imaging can change the behaviour of classifiers and segmentation, illustrating the lack of robustness in the neural network models. Such approach can also be used for model evaluation.²⁸ Vulnerability during segmentation could lead to wrong representation of reality during the following stages of diagnosis. Again, the less understood and erratic boundaries of classification are a concern that enable an attack vector.

Deep learning networks classifying X-ray images are also vulnerable to attacks.³⁵ Being perhaps the most familiar scenario to the public, X-ray image processing is a natural target for automation. However, these kinds of perturbation attacks show that the models can be fooled.

Finlayson et al. successfully use adversarial attacks against medical imaging in three domains: funduscopy, chest X-ray images and dermoscopy. They also present a risk model for the machine learning pipeline.¹⁰ Patch attacks and projected gradient descent both seem to work against real world images, reducing the reliability of the classifier. However, neural networks can be made robust against perturbation attacks by exploiting the structure of the optimization task.²⁰

Not all uses of these methods are harmful. It is possible to use existing medical imaging data to generate more training data and to tackle uneven class balance using various methods, including generative adversarial networks.²² The methods described above can also be used for beneficial inpainting of missing areas in biomedical imaging. Armanious et al. used generative adversarial networks to inpaint missing areas or incomplete medical images.⁵

The identified fooling methods are listed in Table 1. As can be seen, some of the fooling methods have been used in the medical domain. It should be noted that training process tampering is probably more difficult to execute in practice.

Table 1. Fooling methods against deep neural networks and those in the medical domain.

Method	References	Medical domain
Adversarial images	[26],[23],[6]	[28],[35],[10],[22],[5]
Adversarial patches	[7]	[10]
One-pixel attack	[33],[17]	
Training process tampering	[12],[37]	

Conclusions

Although modern neural networks have proven useful for detecting cancerous cell growth, it is possible to mislead these algorithms. There have been research exploits against deep learning methods, even in the field of pathology. Such exploits include specifically engineered adversarial images, adversarial patches put on actual images, one-pixel attacks and attacks focusing on fooling the training process. The scientific studies this short survey inspected include all those attacks. Even medical imaging is not safe from them, which promotes further study of the underlying causes and robustness problems stemming from the structure of neural networks. The expert opinions from the medical community will also broaden the understanding of the effect of these types of attacks.

Acknowledgements

This research is partially funded by the Cyber Security Network of Competence Centres for Europe (CyberSec4Europe) project of the Horizon 2020 SU-ICT-03-2018 program.

References

- 1 Famke Aeffner, Mark Zarella, Nathan Buchbinder, Marilyn M Bui, Matthew Goodman, Douglas Hartman, Giovanni Lujan, Mariam Molani, Anil Parwani, Kate Lillard, Oliver Turner, Venkata Vemur, Ana Yuil-Valdes, and Douglas Bowman, "Introduction to Digital Image Analysis in Whole-slide Imaging: A White Paper from the Digital Pathology Association," *Journal of pathology informatics* 10, no. 1 (2019): 9.

- ² Mahmoud Afifi and Michael S Brown, "What Else Can Fool Deep Learning? Addressing Color Constancy Errors on Deep Neural Network Performance," In: *Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 243–252.
- ³ Naveed Akhtar and Ajmal Mian, "Threat of Adversarial Attacks on Deep Learning in Computer Vision: A Survey," *IEEE Access* 6 (2018): 14410–14430.
- ⁴ Shaimaa Al-Janabi, Henk-Jan van Slooten, Mike Visser, Tjeerd van der Ploeg, Paul J. van Diest, and Mehdi Jiwa, "Evaluation of Mitotic Activity Index in Breast Cancer Using Whole Slide Digital Images," *PloS one* 8, no. 12 (2013).
- ⁵ Karim Armanious, Youssef Mecky, Sergios Gatidis, and Bin Yang, "Adversarial Inpainting of Medical Image Modalities," In: *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2019, pp. 3267–3271.
- ⁶ Anish Athalye, Logan Engstrom, Andrew Ilyas, and Kevin Kwok, "Synthesizing Robust Adversarial Examples," In: *Proceedings of the 35th International Conference on Machine Learning, PMLR, Stockholm, Sweden, vol. 80*, 2018, pp. 284–293.
- ⁷ Tom B. Brown, Dandelion Mané, Aurko Roy, Martin Abadi, and Justin Gilmer, "Adversarial Patch," arXiv e-prints (2018), <https://arxiv.org/abs/1712.09665>.
- ⁸ Maria Chuquicusma, Sarfaraz Hussein, Jeremy Burt, and Ulas Bagci, "How to Fool Radiologists with Generative Adversarial Networks? A Visual Turing Test for Lung Cancer Diagnosis," In: *2018 IEEE 15th international symposium on biomedical imaging (ISBI 2018)*, IEEE, 2018, pp. 240–244.
- ⁹ Yepeng Deng, Chunkai Zhang, and Xuan Wang, "A Multi-objective Examples Generation Approach to Fool the Deep Neural Networks in the Black-box Scenario," In: *2019 IEEE Fourth International Conference on Data Science in Cyberspace (DSC)*, IEEE, 2019, pp. 92–99.
- ¹⁰ Samuel Finlayson, Hyung Won Chung, Isaac Kohane, and Andrew Beam, "Adversarial Attacks Against Medical Deep Learning Systems," *arXiv e-print* (4 Feb. 2019), <https://arxiv.org/abs/1804.05296>.
- ¹¹ Ian Goodfellow, Yoshua Bengio, and Aaron Courville, *Deep Learning* (MIT Press, 2016), <http://www.deeplearningbook.org>.
- ¹² Tianyu Gu, Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg, "Badnets: Evaluating Backdooring Attacks on Deep Neural Networks," *IEEE Access* 7 (2019): 47230–47244.
- ¹³ Zhaoquan Gu, Weixiong Hu, Chuanjing Zhang, Hui Lu, Lihua Yin, and Le Wang, "Gradient Shielding: Towards Understanding Vulnerability of Deep Neural Networks," *IEEE Transactions on Network Science and Engineering*, 2020.
- ¹⁴ Luis Carlos Junqueira and Jose Carneiro, *Basic Histology Text & Atlas* (McGraw-Hill Professional, 2005).
- ¹⁵ Brady Kieffer, Morteza Babaie, Shivam Kalra, and H.R.Tizhoosh, "Convolutional Neural Networks for Histopathology Image Classification: Training vs. Using Pre-trained Networks," In: *2017 Seventh International Conference on Image Processing Theory, Tools and Applications (IPTA)*, IEEE, 2017, pp. 1–6.
- ¹⁶ Daisuke Komura, Shumpei Ishikawa, "Machine Learning Methods for Histopathological Image Analysis," *Computational and structural biotechnology journal* 16 (2018): 34–42.

- ¹⁷ David Kügler, Alexander Distergoft, Arjan Kuijper, and Anirban Mukhopadhyay, “Exploring Adversarial Examples,” In: *Understanding and Interpreting Machine Learning in Medical Image Computing Applications* (Springer, 2018), 70–78.
- ¹⁸ Neeraj Kumar, Ruchika Verma, Sanuj Sharma, Surabhi Bhargava, Abhishek Vahadane, and Amit Sethi “A Dataset and a Technique for Generalized Nuclear Segmentation for Computational Pathology,” *IEEE Transactions on Medical Imaging* 36, no. 7 (2017): 1550–1560.
- ¹⁹ Vinay Kumar, Abul Abbas, and Jon Aster, *Robbins Basic Pathology*, 10th edition (Saunders, Philadelphia: Elsevier, 2017).
- ²⁰ Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu, “Towards Deep Learning Models Resistant to Adversarial Attacks,” *arXiv e-prints*, arXiv:1706.06083 (2019).
- ²¹ Marko Mihajlović and Nikola Popović, “Fooling a Neural Network with Common Adversarial Noise,” In: *2018 19th IEEE Mediterranean Electrotechnical Conference (MELECON)*, IEEE, 2018, pp. 293–296.
- ²² Agnieszka Mikołajczyk and Michał Grochowski, “Data Augmentation for Improving Deep Learning in Image Classification Problem,” In: *2018 international interdisciplinary PhD workshop (IIPhDW)*, IEEE, 2018, pp. 117–122.
- ²³ Seyed-Mohsen Moosavi-Dezfooli, Alhussein Fawzi, and Pascal Frossard, “Deepfool: A Simple and Accurate Method to Fool Deep Neural Networks,” In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2016, pp. 2574–2582.
- ²⁴ Muthukumar Murugesan and R. Sukanesh, “Automated Detection of Brain Tumor in EEG Signals Using Artificial Neural Networks,” In: *2009 International Conference on Advances in Computing, Control, and Telecommunication Technologies*, IEEE, 2009, pp. 284–288.
- ²⁵ Soojeong Nam, Yosep Chong, Chan Kwon Jung, Tae-Yeong Kwak, Ji Youl Lee, Jihwan Park, Mi Jung Rho, and Heounjeong Go, “Introduction to Digital Pathology and Computer-aided Pathology,” *The Korean Journal of Pathology* 54, no. 2 (2020): 125–134.
- ²⁶ Anh Nguyen, Jason Yosinski, and Jeff Clune, “Deep Neural Networks are Easily Fooled: High Confidence Predictions for Unrecognizable Images,” In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, IEEE, 2015, pp. 427–436.
- ²⁷ Ryuji Ohashi, Shigeki Namimatsu, Takashi Sakatani, Zenya Naito, Hiroyuki Takei, and Akira Shimizu, “Prognostic Utility of Atypical Mitoses in Patients with Breast Cancer: A Comparative Study with ki67 and Phosphohistone h3,” *Journal of surgical oncology* 118, no. 3 (2018): 557–567.
- ²⁸ Magdalini Paschali, Sailesh Conjeti, Fernando Navarro, and Nassir Navab, “Generalizability vs. Robustness: Adversarial Examples for Medical Imaging,” *arXiv e-prints* (2018).
- ²⁹ Siddhant Rai, Akshayanand Raut, Akash Savaliya, and Radha Shankarmani, “Darwin: Convolutional Neural Network based Intelligent Health Assistant,” In: *2018 Second International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, IEEE, 2018, pp. 1367–1371.

- ³⁰ Maryam Rastgar-Jazi and Xavier Fernando, "Detection of Heart Abnormalities via Artificial Neural Network: An Application of Self Learning Algorithms," In: *2017 IEEE Canada International Humanitarian Technology Conference (IHTC), IEEE, 2017*, pp. 66–69.
- ³¹ Deepika Ravish, Nayana R Shenoy, K. Shanthi, and S. Nisargh, "Heart Function Monitoring, Prediction and Prevention of Heart Attacks: Using Artificial Neural Networks," In: *2014 International Conference on Contemporary Computing and Informatics (IC3I), IEEE, 2014*, pp. 1–6.
- ³² Ahmed Serag, Adrian Ion-Margineanu, Hammad Qureshi, Ryan McMillan, Marie-Judith Saint Martin, Jim Diamond, Paul O'Reilly, and Peter Hamilton, "Translational AI and Deep Learning in Diagnostic Pathology," *Frontiers in Medicine* 6 (2019): 185.
- ³³ Jiawei Su, Danilo Vasconcellos Vargas, and Sakurai Kouichi, "One Pixel Attack for Fooling Deep Neural Networks," *IEEE Transactions on Evolutionary Computation* 23, no. 5 (2019): 828–841.
- ³⁴ Rajarshi Syam and R. Marapareddy, "Application of Deep Neural Networks in the Field of Information Security and Healthcare," In: *2019 SoutheastCon, IEEE, 2019*, pp. 1–5.
- ³⁵ Saeid Asgari Taghanaki, Arkadeep Das, and Ghassan Hamarneh. "Vulnerability Analysis of Chest X-ray Image Classification against Adversarial Attacks," In: *Understanding and Interpreting Machine Learning in Medical Image Computing Applications* (Springer, 2018): 87–94.
- ³⁶ Hamid Reza Tizhoosh and Liron Pantanowitz, "Artificial Intelligence and Digital Pathology: Challenges and Opportunities," *Journal of Pathology Informatics* 9 (2018).
- ³⁷ Chaofei Yang, Qing Wu, Hai Li, and Yiran Chen, "Generative Poisoning Attack Method against Neural Networks," *arXiv e-prints* (2017).