Towards AI-driven Predictive Modeling of Turbines using Big Data

Sunny Raj*, Steven Lawrence Fernandes†, Andy Michel‡, and Sumit Kumar Jha§
Computer Science Department
University of Central Florida
4328 Scorpius Street
Orlando, FL 32816
Email: {sraj, steven, andy, jha@eecs.ucf.edu}

Artificial Intelligence promises to revolutionize the prognosis of complex systems like gas turbines and facilitate their long-term autonomous operation by enabling both predictive maintenance and recovery from potential faults via a suitable re-configuration of the deployed system. We present a recurrent neural network based approach for probabilistic prediction of the power generated from a gas turbine up to 5 minutes in the future by observing (i) the generated power in the past and (ii) past snapshots of 80 other sensor data, including those monitoring vibrations. Our deep neural network with 367,041 parameters is trained on a system containing 2 NVIDIA GPUs RTX 2080 with 5,888 CUDA cores, 32 CPU cores, and 128 GB RAM, and our neural network achieves a root mean square error of 0.0029 in 10 epochs in about 14.5 minutes. The deep neural network produces a root mean square error of 0.0028 on a validation data set. We also define an event of interest to be a drop of the generated power from 70% of the normal to a value below 50% of the normal value. Our test data set contains 6 such events of interest. Our recurrent neural network is trained with the objective of making predictions up to 5 minutes in the future; we observed that the neural network can predict all of these six events of interest in the test data before their occurrence.

I. Introduction

Our ability to fabricate tens of thousands of processor cores coupled to tens of Gigabytes of memory on a single system has given rise to the current generation of general-purpose Graphics Processing Units (GPUs) with TeraFLOPS of computational capacity available on a single device. The widespread availability of such GPU devices has led to the rise of deep neural networks and a new wave of artificial intelligence (AI) with hitherto unforeseeable capabilities. AI systems based on deep neural networks have reached human-like performance even in those tasks that were earlier considered to be so computationally challenging that Turing tests (such as captchas) separating human being from computers have been designed using these tasks. Examples of such human-centered tasks include character recognition [1], voice recognition [2] and object recognition [3].

Artificial Intelligence promises to revolutionize the prognosis of complex systems like gas turbines and facilitate long-term autonomous operation of such complex systems by enabling both predictive maintenance and recovery from potential faults by a suitable re-configuration of the deployed system. Gas turbines are inherently nonlinear and complex systems, and prior efforts to model them using first principles and statistical approaches have had limited success. Modeling using first principles [4] requires a deep quantitative understanding of the combustion coupled to massively parallel high-performance systems for solving partial differential equations. The generated models can readily be interpreted by human experts but they often cannot achieve the accuracy of the deep learning model. The lack of an adequate number of gas turbine experts required to build a model for every gas turbine is also a challenge that limits the scalability of this approach on its own. Statistical methods used in image recognition [5], natural language processing [6] and other domains have been greatly outperformed by deep learning methods. Hence, it is likely that

*Doctoral Candidate, Computer Science Department, University of Central Florida, Orlando, FL 32816

†Pre-eminent Postdoctoral Program Fellow, Computer Science Department, University of Central Florida, Orlando, FL 32816

‡Doctoral Student, Computer Science Department, University of Central Florida, Orlando, FL 32816

§Associate Professor, Computer Science Department, University of Central Florida, Orlando, FL and Air Force Office of Scientific Research (AFOSR) Young Investigator Awardee 2016 deep learning methods coupled with expert-driven reasoning using first principles may be used to create digital twins of gas turbines in the near future.

To the best of our knowledge, deep recurrent neural networks have not yet been applied to the problem of predicting the performance of gas turbines using sensor data, such as those from vibration sensors and those related to turbine supervisory controllers. In this paper, we make the following contributions:

- We create a recurrent neural network model to predict the short-term power generation performance of a gas turbine by exploiting a data set of $100,000$ data points. Our trained network achieves a root mean square error of 0.0029 in 10 epochs in about 14.5 minutes.
- We use the recurrent neural network to make predictions about events of interest on a hitherto unseen test gas turbine data set of $20,000$ data points and can predict all six events of interest in the data set. We defined an event of interest as a reduction in the generated power from 70% of the normal to less than 50% of the normal power output.

Our work on building deep neural network models for predicting gas turbine performance automatically creates a digital twin for the gas turbine. We anticipate that several important analyses, such as time to failure, mode of failure, creation of data-driven maintenance schedules, and improvements in system efficiency, may be achieved using our deep learning based approach.

II. Related Work

Early prediction plays a vital role for long term sustainability of heavy-duty gas turbine and turboshaft-engines. Pampaloni et al. [7] have used numerical analysis for prediction of pollutant emissions in natural gas with low NOx burners. This model was initially proposed by Klarmann et al. [8]. Results obtained from the dry low-NOx burner, made available from a campaign performed by Baker Hughes, a GE company along with the University of Florence was used for validation of the prediction results. Gu et al. [9] proposed a predictive controller design based on linear parameter-varying control theory and applied it to turboshaft engines. The predictive controller performs better when compared to two loop proportional integral controller due to its lower overshoot time and faster response time.

With the advancements in artificial intelligence, development of intelligent propulsion systems is a widely studied research area. Duan et al. [10] proposed a smart and economical method for measuring the temperature of an aero-engine component using the thermal resistance on the turbine blade surfaces which has thermal barrier coating made up of YSZ material. It is used as a thermal sensor for intelligent monitoring the surface temperature. Stange et al. [11] have presented a review of the sensors developed by propulsion instrumentation working group.

Recently, neural networks are used to automatically extract features from the sensor data and predict anomaly in gas turbine [12]. This approach uses stacked autoencoders which can learn features that are robust to input noise. The features learned are given as input to an extreme learning machine (ELM) for anomaly detection. However, ELM is feed-forward neural networks with only two layers. The first layer is random, and second layer is trained. The randomness of ELM causes uncertainty to the problem in learning and approximation. Also ELM exhibit generalization degradation phenomenon with inappropriate activation functions. In our proposed approach, we have not used ELM. Feature extraction and classification is automatically done by the deep learning network.

Alam et al. [13] proposed a neural network based approach to predict the remaining useful life of an aircraft turbine engine. They used the publicly available simulation platform Propulsion Diagnostic Method Evaluation Strategy (ProDiMES) to generate synthetic fault data to train the neural network. Tsoutsanis et al. [14] used a linear regression based approach to predict the remaining useful life of gas turbine engines under various ambient conditions. They used a two shaft industrial gas turbin model developed in MATLAB/Simulink to generate data to train the regression model. Kiakojoori and Khorasani [15] used a nonlinear autoregressive exogenous model (NARX) neural network to predict compressor fouling and turbine erosion in gas turbine engines. A MATLAB/Simulink model for a single spool gas turbine engine was used to generate data to train the NARX neural network. Daroogheh et al. [16] combined model based and neural network based methods to create a hybrid approach to predict the health condition of gas turbine engines. This hybrid approach has the advantage of not requiring a large dataset to train the model. Again the data used to train the model was obtained by simulating a single spool gas turbine MATLAB/Simulink model. Our work in contrast to these approaches uses real-world gas turbine data to make power output predictions. We use sequences of data obtained from gas turbines to predict turbine power output up to five minutes in the future.

III. Approach

A wide variety of deep learning algorithms rely on the availability of large representative data sets in order to train an AI model with high predictive accuracy. However, AI in predictive maintenance suffers from a critical problem: faults in complex but well-designed deployed systems are relatively rare; hence, most training data represents normal behaviors of the system.

A straightforward solution to the problem is to create physics-based simulations of the system where the effects of various faults in a complex system can be modeled *in silico* and a representative data set of faulty behavior can be obtained. However, this approach is unlikely to scale up as it requires a human expert in the loop. Most success in AI has come from coupling high-velocity data with high-performance computing – both of which arguably operate at frequencies much higher than the human intellect.

In our approach, we resolved this problem by balancing data containing events of interest with data containing normal system behavior through data augmentation. Our initial data set contained about $100,000$ data points with dozens of faults and we created a pruned training data set of $15,995$ data points by rejecting sustained normal behavior of the system with 5% probability. In future, we anticipate augmenting the data set by adding small variations of the data denoting events of interest. We are also considering the use of generative adversarial networks (GANs) as well as variational autoencoders to synthesize new data denoting events of interest or faults.

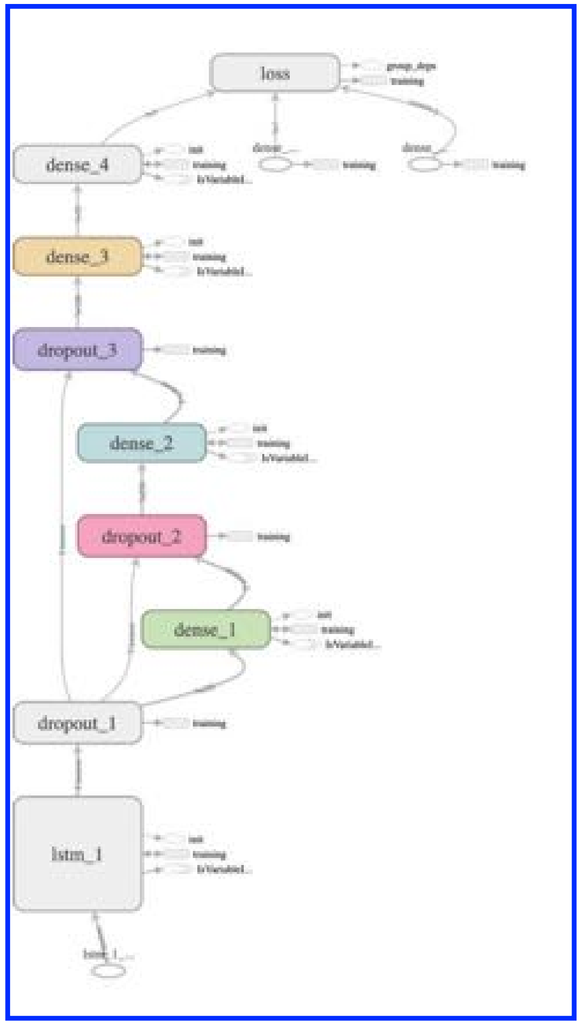

Fig. 1 Recurrent neural network architecture used for predicting power generation of gas turbine engines. The architecture consists of 1 LSTM and 4 fully-connected dense layers with dropouts to prevent overfitting to data.

Our neural network architecture to predict power generation from gas turbine engines is shown in Figure 1. Our architecture consists of eight layers, including 3 dropouts. The first layer consists of long short-term memory (LSTM) nodes that can capture temporal relationship among the 80 variables. This is followed by four fully-connected dense layers of decreasing sizes. The final layer with only one node acts as the output layer. We included a dropout after every layer, and set the dropout probability to 20%.

The complete deep neural network has a total of 367,041 parameters. The output layer of our network predicts the power output of the generator in the future. We used the root mean squared error (RMSE) as our loss function. To train the network, we used the RMSProp optimizer with a learning rate of 0.001. RMSProp optimizer has been shown to train recurrent neural networks faster than many other approaches and helps our neural network to quickly converge to the optimal weights.

Before arriving at the design of the recurrent neural network, we investigated several other candidate designs including those with two or more long short-term memory (LSTM) layers, convolutional layers and deeper neural networks with up to 20 layers. In case of designs with multiple LSTM layers, the performance of the architecture was poorer than our current approach. Convolutional layers were not as helpful in predicting future values as LSTM layers. The use of multiple layers of fully-connected dense networks did not substantially improve the prediction performance but made the training of the network slower. Hence, our recurrent neural network design provides a good trade-off between improved prediction accuracy and the ability to rapidly train on large data sets.

IV. Results

Using our curated training dataset of 15,995 data points with adequate representation of normal and faulty behavior, our recurrent neural network seeks to predict all 6 events of interest 5 minutes before their actual occurrence in a hitherto unseen test data set. We define an event of interest as a reduction in the generated power from 70% of the normal to less than 50% of the normal power output. These predictions are shown in Figure 2.

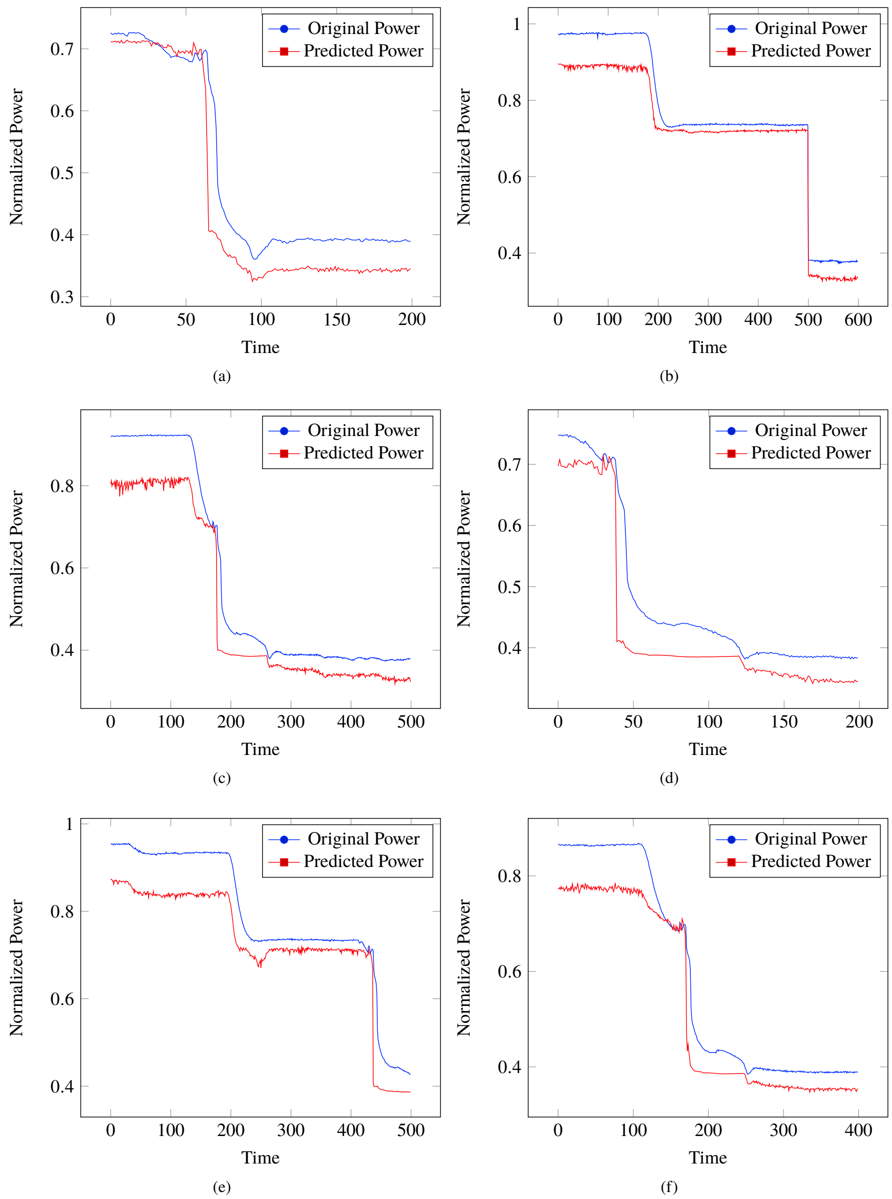

Fig. 2 Prediction from our model trained to predict events of interest 5 minutes before their occurrence.

It is clear from the above plots that the model's predictions (shown in red) occur before the system shows an event of interest (shown in blue). The quality of the predictions varies among the different events of interest, and we have not investigated the reason for such variations.

As a next step, we trained our model to predict events of interest 10 minutes into the future. The results of our prediction are shown in Figure 3.

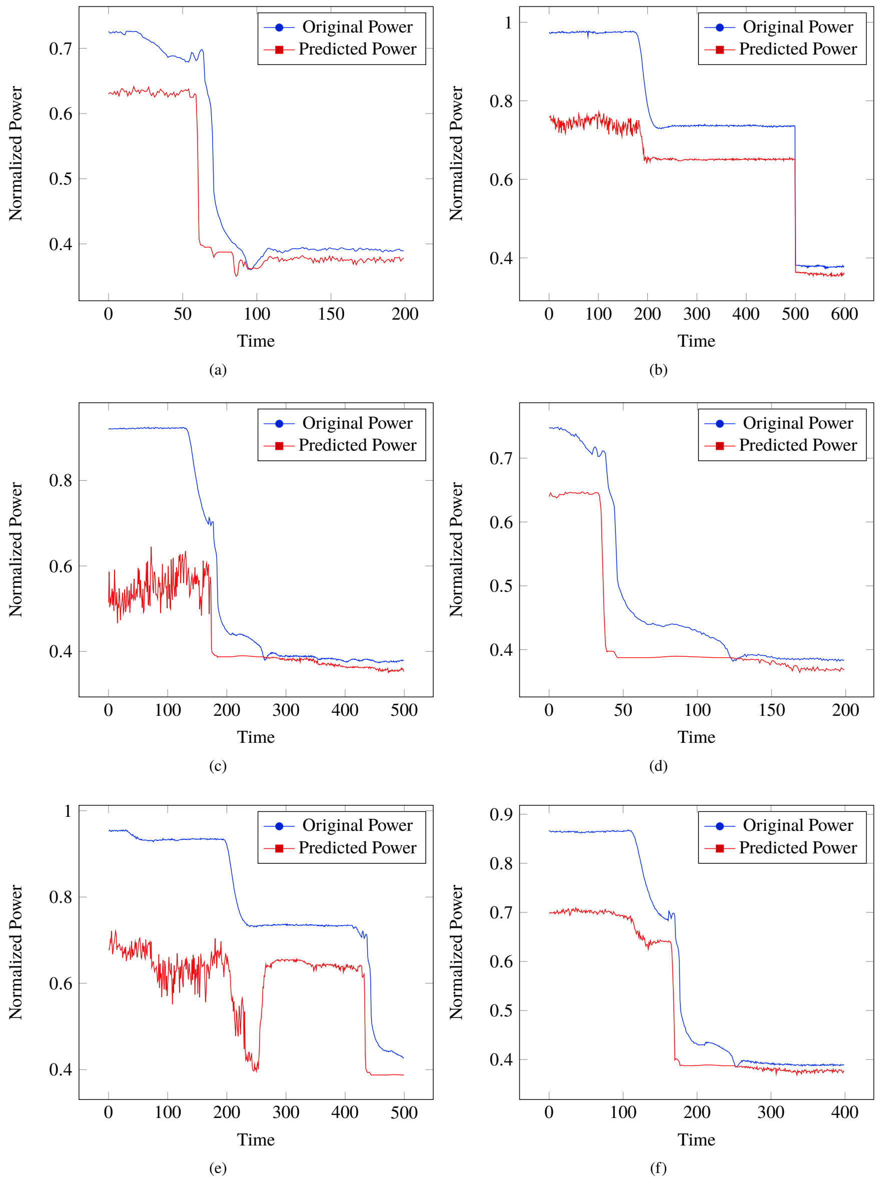

Fig. 3 Prediction from our model trained to predict events of interest 5 minutes before their occurrence.

The overall quality of our predictions 10 minutes into the future is poorer than the quality of our predictions 5 minutes into the future. This is particularly true while predicting the normal behavior of the system. However, we are still able to predict events of interest as the downward trend of the predicted power precedes the actual occurrence of the event. It may be feasible to create a better fit for the normal behavior of the system by providing additional data about the normal behavior of the system.

In order to test the validity of our overall approach, we sought to train a model to predict events of interest 30 minutes into the future. We find that our recurrent neural network model is not very effective at predicting events of interest 30 minutes into the future. Figure 4 illustrates the performance of our recurrent neural model and the predicted power (in red) shows large drifts and deviations from the shape of the observed power (in blue).

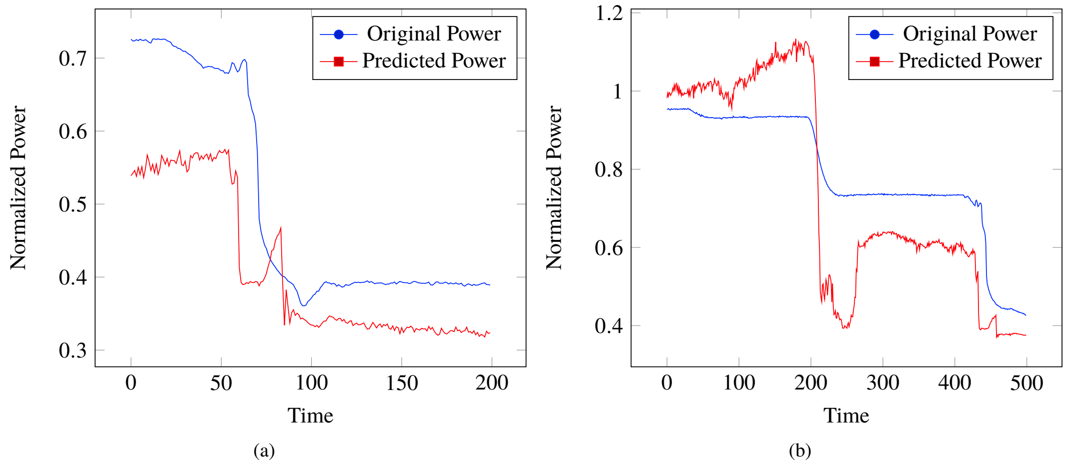

Fig. 4 Our model is not effective at predicting faults 30 minutes into the future.

Our recurrent neural network was used to predict the power output 30 seconds into the future. The results obtained are shown in Figure 5. The predicted behavior of the model shows a tight fit to the observed power output of the gas turbine.

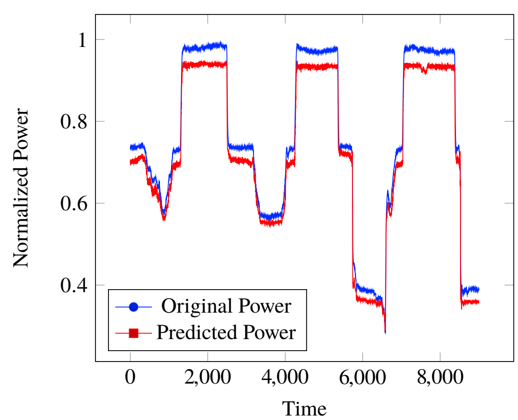

Fig. 5 Overview of the efficacy of our predictions 30 seconds into the future.

V. Conclusions and Future Work

We have designed, implemented and tested a (deep) recurrent neural network based predictive model for predicting the power output of a gas turbine engine. Our AI system achieves a root mean square error of 0.0028 and is effective at predicting all the 6 events of interest in our test data set. As is known from prior studies with neural network models, qualitatively better results may be obtained by using more finely sampled data, such as one sampled according to the Nyquist criterion. Future research in developing digital twins for gas turbines using predictive AI models may benefit from the following suggestions:

- AI-based models may be combined with statistical and physics-based models to enable more long-term predictions. Statistical methods will be used to filter the data and identify those attributes that may not be strongly related to the performance of a complex system, such as a gas turbine. Removal of such extraneous attributes may assist deep learning based approaches. Physics based methods may be used to refit the predicted data to an equation-based or parameterized computational model that can be readily interpreted by human experts. Tight coupling of statistical and physics based approaches with deep learning is likely to create a useful digital twin of a gas turbine.
- Our data set included vibration data provided as a vector of measurements. More interesting results may be obtained by directly analyzing wide-band time-series data from recordings of vibrations (in and out of the audible spectrum) and video data obtained from camera-like sensors at multiple (visible and beyond visual spectrum) frequencies.
- Our current approach has used a relatively simple (deep) recurrent neural network architecture for generating short-term predictions. We anticipate that the use of attention based models and Neural ODEs [17] coupled to training on tens of GPUs is likely to increase the model's ability to predict deeper into the future. Figure 6 shows the efficacy of neural ODEs in predicting the power output of gas turbines. It may also be useful to consider the use of generative adversarial networks (GANs) or variational autoencoders to synthesize new data denoting events of interest or faults.

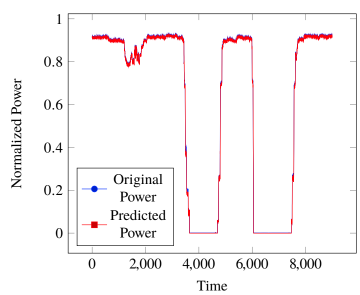

Fig. 6 Predicting power output of a gas turbine using Neural Ordinary Differential Equations

- Analysis of the learned deep learning model using explanations and causality analysis will enable us to refine deep neural networks making them more precise and perhaps more interpretable by human experts. Attribution analysis [18] has been used to develop explanations for deep learning models. Figure 7 shows how attribution analysis identifies the pixels used to detect a taxi in the image that cause a deep learning model to recognize the taxi in the image. Similar attribution analysis can be used to identify sources of faults and develop predictive solutions for complex systems, such as gas turbines.

(a) (b)

Fig. 7 Image classified as taxi and the corresponding attributions obtained using integrated gradients method

References

[1] Cheng, Z., Xu, Y., Bai, F., Niu, Y., Pu, S., and Zhou, S., "AON: Towards Arbitrarily-Oriented Text Recognition," *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[2] Połap, D., and Woźniak, M., "Voice recognition by neuro-heuristic method," *Tsinghua Science and Technology*, Vol. 24, No. 1, 2019, pp. 9–17. doi:10.26599/TST.2018.9010066.

[3] LaLonde, R., Zhang, D., and Shah, M., "ClusterNet: Detecting Small Objects in Large Scenes by Exploiting Spatio-Temporal Information," *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.

[4] Rhinehart, R. R., "Using first-principle models in control," *IEEE Industry Applications Society Dynamic Modeling Control Applications for Industry Workshop*, 1997, pp. 10–19. doi:10.1109/DMCA.1997.603383.

[5] Hiramoto, M., Ogawa, T., and Haseyama, M., "A novel image recognition method based on feature-extraction vector scheme," *2004 International Conference on Image Processing, 2004. ICIP '04.*, Vol. 5, 2004, pp. 3049–3052 Vol. 5. doi:10.1109/ICIP.2004.1421756.

[6] Liu, F.-H., Gu, L., Gao, Y., and Picheny, M., "Use of statistical N-gram models in natural language generation for machine translation," *2003 IEEE International Conference on Acoustics, Speech, and Signal Processing, 2003. Proceedings. (ICASSP '03).*, Vol. 1, 2003, pp. I–I. doi:10.1109/ICASSP.2003.1198861.

[7] Cerutti, M., Riccio, G., Andreini, A., Becchi, R., Bruno Facchini, D., and Picchi, A., "Experimental and Numerical Investigations of Novel Natural Gas Low NO x Burners for Heavy Duty Gas Turbine," *Journal of Engineering for Gas Turbines and Power*, Vol. 141, 2018. doi:10.1115/1.4040814.

[8] Klarmann, N., Zoller, B. T., and Sattelmayer, T., "Numerical modeling of co-emissions for gas turbine combustors operating at part-load conditions," *Journal of the Global Power and Propulsion Society*, Vol. 2, 2018, pp. 376 – 387. doi:10.22261/JGPPS.C3N5OA, URL https://journal.gpps.global/a/C3N5OA/.

[9] Gu, N., and Wang, X., "Model Predictive Controller Design based on the Linear Parameter Varying Model Method for a class of Turboshaft Engines," *2018 Joint Propulsion Conference*, 2018, p. 4620.

[10] Duan, F. L., Hu, M., Zou, B., and Li, J., "An Easy Way of High-Temperature Monitoring of Turbine Blade Surface for Intelligent Propulsion Systems," *2018 Joint Propulsion Conference*, 2018, p. 4619.

[11] Stange, W. A., and Behbahani, A. R., "Modern Instrumentation Channels for Turbine Engine Controls," *2018 Joint Propulsion Conference*, 2018, p. 4623.

[12] Yan, W., and Yu, L., "On Accurate and Reliable Anomaly Detection for Gas Turbine Combustors : A Deep Learning Approach," 2015.

[13] Alam, M. M., Bodruzzaman, M., and Zein-Sabatto, M. S., "Online prognostics of aircraft turbine engine component's remaining useful life (RUL)," *IEEE SOUTHEASTCON 2014*, IEEE, 2014, pp. 1–6.

[14] Tsoutsanis, E., Meskin, N., Benammar, M., and Khorasani, K., "A dynamic prognosis scheme for flexible operation of gas turbines," *Applied energy*, Vol. 164, 2016, pp. 686–701.

[15] Kiakojoori, S., and Khorasani, K., "Dynamic neural networks for gas turbine engine degradation prediction, health monitoring and prognosis," *Neural Computing and Applications*, Vol. 27, No. 8, 2016, pp. 2157–2192.

[16] Daroogheh, N., Baniamerian, A., Meskin, N., and Khorasani, K., "A hybrid prognosis and health monitoring strategy by integrating particle filters and neural networks for gas turbine engines," *2015 IEEE Conference on Prognostics and Health Management (PHM)*, IEEE, 2015, pp. 1–8.

[17] Chen, T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. K., "Neural ordinary differential equations," *Advances in neural information processing systems*, 2018, pp. 6571–6583.

[18] Jha, S., Raj, S., Fernandes, S., Jha, S. K., Jha, S., Brian, J., Verma, G., and Swami, A., "Attribution-driven Causal Analysis for Detection of Adversarial Examples," *Safe Machine Learning workshop at ICLR*, 2019.