

LSTM-based restoration of lossy data to improve anomaly detection algorithms

Anonymous Author(s)

Affiliation

Address

Email

Abstract—Most real-world datasets, and particularly those collected from physical systems, are full of noise, packet loss, and other imperfections. However, most specification mining or anomaly detection algorithms assume, or even require, perfect data quality to function properly. Such algorithms may work in lab conditions when given clean, controlled data, but will fail in the field when given imperfect data. We propose a method for accurately reconstructing discrete temporal or sequential system traces affected by data loss, using Long Short-Term Memory Networks (LSTMs). The model works by learning to predict the next event in a sequence of events, and uses its own output as an input to continue predicting future events. Such a method can reconstruct even long sequence of missing events, and can also help validate and improve data quality for noisy data. The output of the model will be a close reconstruction of the true data, and can be fed to algorithms that rely on clean data. We demonstrate our method by reconstructing automotive CAN traces consisting of long sequences of discrete events. We show that given even small parts of a CAN trace, our LSTM model can accurately reconstruct large portions of the original trace. We separately feed the original, lossy, and reconstructed traces into a specification mining framework to gauge the effect of our method on the performance of state-of-the-art models using these traces for understanding the behavior of complex systems.

I. INTRODUCTION

It is often the case that software and algorithms work very differently in lab conditions than they do out in the field. This is particularly a problem for safety-critical systems ranging from planes to pacemakers, where any malfunction in the field can mean serious - even life-threatening - consequences for users. It is thus important that algorithms designed for such systems work as intended both in the lab and out on the field. The problem is that many such algorithms are developed in the lab using clean, processed, and perfect data. Real data, on the other hand, can contain noise and loss that may radically alter the behaviour of the algorithms processing it. This results in these algorithms working in theory, but failing once put into practice due to their dependence on perfect data quality.

While 'imperfect' real data may appear to be irreparable or different from clean data, the reality is that much of the original structure can still be reconstructed from it, especially when there is sufficient information about similar data that can be used for reconstruction. Supervised learning techniques can be used to learn the general structure of similar datasets, and then reconstruct the lossy dataset of interest.

Many real-world systems produce discrete datasets such as text or event logs during their operation. Automotive Controller Area Network (CAN) traces, for example, consist of a long sequence of discrete events. A CAN bus is a communication system that allows various types of devices like microcontrollers to communicate with each other in real time without needing a host. It is a message-based protocol, designed originally for multiplex electrical wiring within automobiles to save on copper, but is also used in many other contexts. The messages in a CAN trace can be considered as 'words' produced by the CAN bus. Therefore, various supervised learning techniques from natural language processing (NLP) are likely to be effective in analyzing the patterns in sequences - or "sentences"- of these events.

Recurrent Neural Networks (RNNs), are neural networks that have an internal memory. Long Short-Term Networks (LSTMs) are a type of RNN that have a superior ability to learn long-term dependencies in data. LSTMs have been used extensively for work with NLP and event based data [1], they have been used for data compression [2], and they have recently even been used - with varying degrees of success - directly for anomaly detection [3]. The use of LSTMs for specification mining has not yet been explored. We propose a method for using LSTMs to accurately reconstruct discrete temporal traces from the CAN bus affected data loss. The reconstructed traces can then be used by algorithms that rely on 'perfect' data. In our case, we will test the performance of a specification mining framework that uses timed regular expressions and deterministic finite state automata for extracting system behavior.

The model works by learning to predict the next event in a sequence of events from a large database of CAN traces, and uses its own output to continue prediction for multiple points ahead in time, allowing for even large chunks of lost data to potentially be restored. We separately feed the original, lossy, and reconstructed traces into a Timed-Regular Expression Mining (TREM) framework [4], [5], [6] to gauge the effectiveness of our LSTM-based reconstruction approach. We show that given even small parts of a CAN trace, the LSTM model can accurately reconstruct large portions of the original trace thereby permitting the use of algorithms like TREM, that are reliant on 'perfect' data, with 'imperfect' data that would otherwise cause them to perform poorly.

The rest of this paper is divided as follows: Section II will discuss related work in lossy data reconstruction. Section III

will describe the architecture of the LSTM model and discuss the CAN data used for our experiments. Section IV will present a case study with TRE mining, and Section V presents results and conclusion.

II. RELATED WORKS

A. Deep Learning for Sequence Prediction

The application of deep learning techniques has led to dramatic success and development of state-of-the-art solutions to numerous real-world problems in computer vision [7], machine translation [8], robotics [9], etc. In context of sequence prediction, a broad range of techniques have been developed to address problem found in Natural Language Processing (NLP)[10], genetic sequencing[11], stock market prediction[12], music generation[13], and a whole range of other domains. In particular, RNNs and LSTMs have been shown to be very effective for sequence prediction in temporal or sequential data. For example, (Gers and Schmidhuber 2001), demonstrated the significant improvement in performance by using LSTMs on major language benchmarks used for RNNs[14].

Many advances have been made in the closely related problem of using deep learning techniques for generating end-to-end sequences of outputs from a sequence of inputs, beginning with (Sutskever et al. 2014), describing a groundbreaking method for using an LSTM-based encoder-decoder system to learn sequences while minimizing the number of assumptions about the structure of the sequence[1]. (Hong et al. 2017), showed that a modified convolutional sequence-to-sequence autoencoder could be trained to predict visually un-observable weather patterns when given preceding satellite images as input [15]; and (Marchi et al. 2015) devised a method based on LSTM recurrent denoising autoencoders to predict the features of a consecutive audio frame based on the previous ones[16].

B. Deep Learning for Lossy Data Recovery

While numerous studies discuss the use of deep learning for directly solving problems like anomaly detection[17], [18], [19], [3] and intrusion detection [20], [21], [22], only very few explicitly focus on the critical intermediate problem of restoring lossy or noisy data. Often, missing data is considered unusable and is removed completely in pre-processing to avoid having to address it. When missing data are addressed, missing values are typically replaced with global or class means [23], [24] or neighboring values [25], [26]. These approaches are much too coarse and do not take into account the event sequence or order of occurrence.

(Zhou and Huang 2017) discuss this problem and propose a novel LSTM-based approach they call an "Iterative Imputing Network" for restoring missing sensor data in time-series [27]. We draw inspiration from their work and propose a method using LSTMs to restore lossy discrete sequential data.

Some other applications of deep learning to data recovery in a variety of domains include: (Hsieh and Pratt 2001) showing that artificial neural networks could be used to recover lossy field data with high reliability[28]; (Tilk and Alumäe 2015)

showing that LSTMs are effective at restoring punctuation in unpunctuated streams of text[29]; and (Haque, Yousuf, and Rana 2018) showing that a combination of CNNs and LSTMs could be used to successfully de-noise and restore image data [30].

III. LSTM

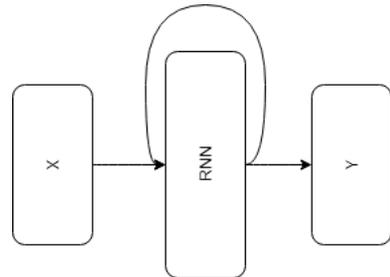
A. RNNs and LSTMs

One of the major limitations of non-recurrent neural networks is that they operate on fixed-size vectors, performing a limited number of transformations to derive another fixed-size vector. This limits the effectiveness of non-recurrent networks in identifying features dispersed over sequences or over time. Recurrent neural networks (RNNs) are a type of neural network that use a feedback mechanism to allow the network to operate on a sequence of inputs as well as outputs by selectively keeping information about previous states. In this manner, the i^{th} output vector y_i of an RNN layer is a function of the i^{th} input x_i to the layer as well as the previous output y_{i-1} of the layer. As with regular layers, the function is generally a non-linear transformation over a weighted sum of the terms involved.

$$y_i = f(y_i, x_i, y_{i-1}) = f(V * y_{i-1} + W * x_i + c)$$

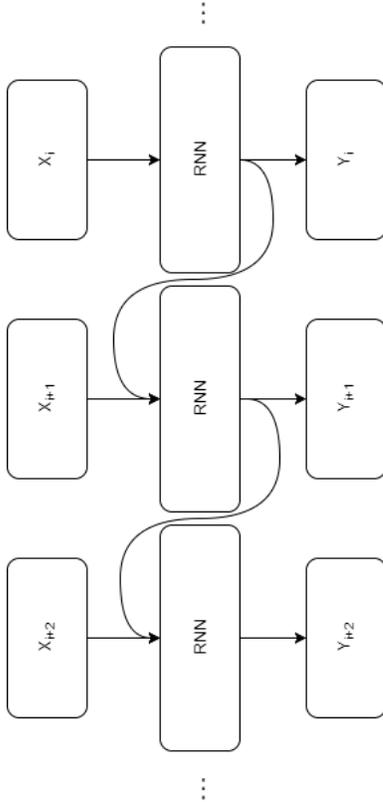
The issue of course, is that in such a recursive function, every previous output and input must be considered when adjusting weights during training, as during backpropagation they would all have an impact on the error gradient. These incredibly long chains of derivatives pose two problems: computational intensity and vanishing or exploding gradients.

Fig. 1. **Single RNN Unit:** A recurrent neural network uses a feedback mechanism to access information about previous states.



One solution to these problems is known as **unrolling** the RNN. This requires considering that if an RNN is essentially a loop with signal flowing through, it can also be equally represented by 'unfolding' the loop as a long series of the same transformations applied to the signal. The structures in Figure 1 and Figure 2 are thus equivalent ways of representing an RNN. Unrolling is the process of using the second, 'unfolded' representation of an RNN and having the option to consider only a fixed number of previous states when predicting the current output. If we use the analogy of short-term memory to describe the feedback mechanism of an RNN, then unrolling would be displaying this memory as a sequence of events

Fig. 2. **Unrolled RNN Unit:** The feedback loop in a recurrent neural network can be unfolded for an alternative, sequential representation of the repeated transformations it performs.



and having the option to cut off the sequence at some event, ignoring any events that came before it.

A second solution to the problem of vanishing or exploding gradients in recurrent networks was the creation of two gates: input and forget. The input gate controls which information from the current state’s input will be used, while the forget gate controls which information will be used from the recurrent outputs of previous states. In such a way, the input and forget gates allow the network to **selectively** remember or forget information about previous and current states. (Hochreiter and Schmidhuber 1997) proposed a new type of RNN that makes use of these two gates in their recurrent units to solve the vanishing gradient problem and gave it the name Long Short-Term Memory network (LSTM) [31].

B. Architecture

We built our LSTM model in Tensorflow. After experimenting with a few different architectures, we found that a smaller model with 5 layers in the hidden portion as seen in Figure 3 was sufficient. Additional hidden layers did not improve performance.

The input layer has one node for each entry in the trace dictionary, and one additional "other" node to account for rare

events that were not seen in the training data that was used to compile the dictionary. Input dropout was set to 0.2.

It is followed by 2 densely-connected hidden layers with double the number of nodes as in the input layer. Each of these hidden layers uses a hyperbolic tan activation function, and has a dropout of 0.4.

The hidden layers are followed by 2 LSTM layers, using Tensorflows LayerNormBasicLSTMCell. These cells perform layer normalization based on (Ba, Kiros, and Hinton 2016)[32]. Each one is unrolled for 40 steps; increasing unrolling further did not seem to increase performance on this dataset, but may be useful for others such as those with higher dimensionality. Both LSTM layers also use the hyperbolic tan activation function, and recurrent dropout of 0.4 based on (Semeniuta, Severyn, and Barth 2016)[33]. Each has four times as many nodes as each of the hidden layers, or 8 times as many as the input layer.

Finally, the last hidden layer is another densely connected layer with a sigmoid activation function and in our final configuration has the same number of nodes as the input layer. However, in an earlier configuration explained below, this layer had n times the number of nodes as the input layer, where n is the number of future events the user wishes to simultaneously predict. The output of this layer is the output of the model.

To increase training speed, LayerNormBasicLSTMCell layers could be replaced with either the peephole[34] or the non-peephole[31] implementation of CoupledInputForgetGateLSTM layers that couple the input and forget gate as described in (Greff et al. 2015)[35] resulting in less computational operations but higher variance in performance. For our CAN trace experiment, as mentioned the vocabulary size was 43, so our input layer had 44 nodes, our hidden layers had 88 nodes each, our LSTM layers had 352 nodes each, and the last hidden layer had 44 nodes whose output was considered the output of the model.

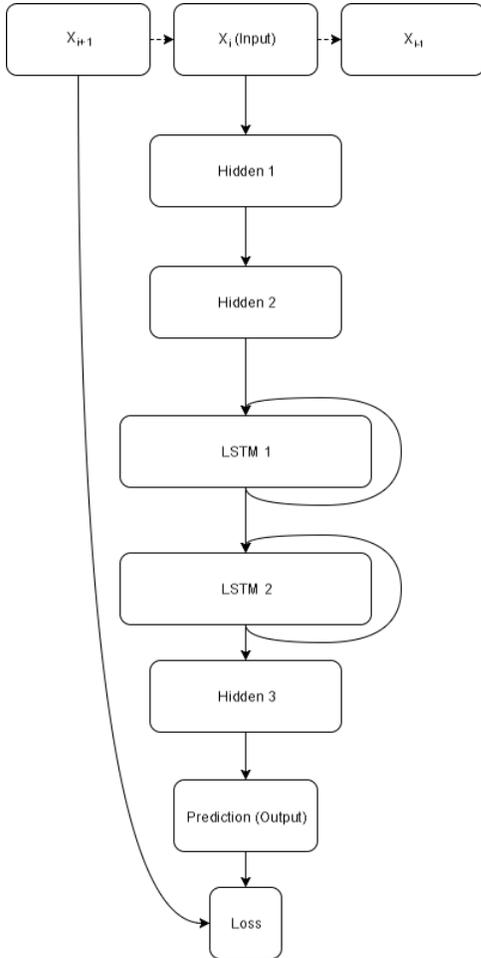
C. Data

Before combining the LSTM approach with the TREM, it was necessary to test whether the LSTM approach works at all. To do so, a simple dataset was needed that was guaranteed to have temporal patterns but is not trivial. Many messages in automotive CAN data occur in a periodic way so temporal patterns are present; however, other messages in CAN data are event-triggered and thus predicting upcoming messages is a non-trivial problem. To keep the dataset simple, we do not attempt to predict each bit of the CAN message payload but only each message ID.

Our CAN data comes from a Lexus RX450h hybrid SUV. The data is split into a number of maneuvers that were repeated multiple times. To reduce some of the variance and reduce training time while testing our LSTM approach, we used traces from a single maneuver, the vehicle driving at 20 km/h and decelerating down to 0 km/h. 20 clean traces of this maneuver were available.

Of the 20 traces, 15 were used for training and validation, while 5 were held out for testing. The traces used for training were examined to compose a dictionary of possible message

Fig. 3. **LSTM Model Architecture:** The LSTM model consists of 2 hidden layers followed by two recurrent LSTM layers and one additional hidden layer. Input is a sequence of events, output is a prediction of next event in the sequence. Loss is calculated as a logloss function comparing the true next event to the predicted one.



IDs. A total of 43 different message IDs were found in the training set. We added a 44th element to the dictionary to designate any other message IDs that may not have been present in the training set. We then one-hot encoded all traces using this dictionary, replacing the one-dimensional message ID with a 44-dimensional vector of indicator variables where 43 elements have the value 0 and 1 element the value 1.

D. Training

The LSTM model was trained using the following procedure:

- 1) Select 2 random traces from set (1 for training, 1 for validation)
- 2) Train the model on these 2 traces for 10 epochs with a learning rate of 0.2 and a logloss loss function

- 3) Train the model on these 2 traces for an additional 20 epochs with learning rate decay of 1/1.1 and a logloss loss function
- 4) Reset learning rate, preserve weights, repeat from step 1 selecting 2 new random traces

Dropout for regularization and randomized input order prevent overfitting, so this procedure should be repeated until accuracy reaches a plateau since early termination due to detection of overfitting is unlikely to occur. Alternatively, the procedure can be repeated a fixed number of times if so desired.

E. Results

Accurately predicting multiple steps forward is a challenging but important problem in lossy data restoration, as multiple consecutive events may have been lost. Two ways of predicting multiple events forward were considered.

The first method was to increase the size of the output layer by a factor of n to directly predict n events forward. In other words, the output layer was modified to have $n \cdot 44$ nodes where n was number of events to simultaneously predict. The output at each step was reshaped into n vectors of 44 nodes each, and each vector predicted 1 future event. Within each vector, the element with the highest value was considered as the prediction. The true positive rate was used as a measure of accuracy. A set of predictions at step i was considered to be a true positive is all n of n predictions made at that step were made correctly. If c_i denotes the number of steps in epoch i where all n predictions were correct, and w_i denotes the number of steps in epoch i where at least 1 of the n predictions was incorrect, then the accuracy for that epoch is denoted by:

$$acc_i = \frac{c_i}{c_i + w_i}$$

One of the issues with this setup is that each time n is changed, a new network needs to be initialized and trained. As such, we retrained the model several times with increasingly large values of n . Table I documents the average performance of the model for different values of n as well as the range 95% of the measured accuracy levels were within. The average accuracy with this method decayed rapidly as n was increased, while the ranges increased significantly. This suggests that the random initialization of the network plays an increasingly large role in the quality of predictions as n increases.

TABLE I. N-FORWARD PREDICTION ACCURACY USING THE DIRECT METHOD

n	Avg. accuracy	95% range
1	0.895	0.88-0.91
5	0.72	0.66-0.76
10	0.58	0.52-0.65
20	0.48	0.4-0.54

A second method was developed to address the shortfalls of the first. The second method predicts only one output at a time; however, the code was altered to allow the model to use its predictions as inputs to itself. In such a way, the model uses its own outputs, one at a time, to make long sequences

of predictions. Using the numbers from Table I, the expected accuracy for this method when predicting n steps forward would be the accuracy for 1 step prediction, to the power of n , since the accuracy metric used above requires every step to be predicted correctly. Table II details the expected accuracy levels for this second method based on this calculation.

TABLE II. EXPECTED N-FORWARD PREDICTION ACCURACY USING THE STEP-BY-STEP METHOD

n	Expected accuracy
1	0.895
5	0.574
10	0.330
20	0.109

While the expected accuracy values for this method were very low, the convergence in logloss in the model suggested that the accuracy metric may be underestimating model performance. In fact, when the model was run on the test trace using this step-by-step method, the results were drastically different from what this accuracy metric would suggest.

As mentioned above, the model uses 40 steps of unrolling for the recurrent portion meaning the model requires 40 inputs to fill its internal feedback sequence. When given the first 40 events of the CAN trace held out for testing, the model was able to step-by-step predict the remaining 3500 events of the trace with only several omissions of rarely occurring events as in Table IV and some localized mistakes in the order of predicted events as in Table V.

Figures 4 & 5 are visualizations of a sequence of true one-hot embedded events and a sequence of predicted one-hot embedded events, respectively, pulled from the same portion of the test trace. It is noticeable that mismatches between the two Figures get increasingly worse as the index increases. This is due to the omission problem: each time an event is omitted in the predictions, the entire sequence of predicted events is shifted to the left, causing increasingly large mismatches. When gauging model performance, each omission was recorded, and then the predicted sequence re-aligned at that point in order to once again match the true sequence. Similarly, each ordering mistake was recorded, and the one point detected to be in the wrong position moved to its correct location, to identify whether other mistakes were made in the same area.

On average, the model omitted 5.058% of points and had one local ordering mistake every 11.1 events. Curiously, both these mismatches and ordering mistakes occur at only a slowly increasing rate throughout all of the 3500 predictions, suggesting that our model’s flexibility makes it at least partially resistant to mistakes in the input, as prediction quality did not decrease even if some local mistakes in output prediction were made and then fed in as input. The LSTM-based approach is thus sufficiently robust to work with and restore not only lossy, but also noisy, data.

IV. CASE STUDY: TIMED REGULAR EXPRESSION MINING

In the domain of real-time systems, temporal behavior of program is of prime importance [36], [37]. The idea of

TABLE III. EXAMPLE OF PROPER ALIGNMENT:

All events in this segment were predicted correctly and match their true counterparts.

	Predicted Event	True Event
1	2C6	2C6
2	5D7	5D7
3	B0	B0
4	224	224
5	B2	B2
6	20	20
7	B4	B4
8	25	25
9	22	22
10	23	23

TABLE IV. EXAMPLE OF OMITTED RARE EVENT:

Rarely occurring event ‘340’ was incorrectly omitted by the model, causing all predicted events beginning from the fourth one to be shifted one up from their true counterparts.

	Predicted Event	True Event
1	B4	B4
2	25	25
3	22	22
4	23	23
5	B0	340
6	320	B0
7	B2	320
8	2D0	B2
9	2C4	2D0
10		2C4

mining temporal properties from system traces is the new trend for extracting commonly occurring behavioral patterns [38]. Quite often these are template based mining frameworks. Many programs lack formal temporal specifications, and mined specifications are therefore valuable since they can be used for a wide variety of activities in the software development life cycle (SDLC) such as testing [39], verification [40], etc.

A vast majority of these frameworks for mining temporal properties use the idea of state machines. These tools learn the behavior of these systems in the form of state machines that can subsequently be used for other tasks such as monitoring, verification etc. In specific case of safety-critical real-time systems the focus is on not only the qualitative notion of time (the specification describe an ordering of events) but also on the quantitative notion of time (the actual duration of time between events). For example, specifications for interrupt handlers, which for real-time systems must complete their executions within a set of predefined deadlines. The time constraints are of importance in such a case since a delayed response to an emitted event may lead to a fault in the system.

With the motivation to address the problem of mining specifications with an explicit notion of time as a use case where recovery of data will play a crucial role in analysis of the system under consideration, the technique is to mine instances of timed regular expression (TRE) templates satisfied by a given system’s traces. TREs [41] was proposed by [4], [5]. The method proposed by Asarin et al. [41] to synthesize a timed automaton for a given TRE is the basis of the proposed approach. The timed automaton is then used as a checker to verify whether traces satisfy the corresponding TRE.

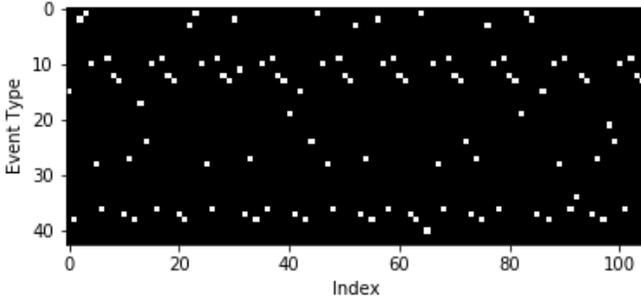
For the experiments, we used a single eight core machine equipped with the Intel i7-3820 CPU at 3.60 GHz and 31.4

TABLE V. EXAMPLE OF LOCAL ORDERING MISTAKE:

'2c4' was incorrectly predicted as the 8th event instead of 4th, causing all of the other events from 4th to 8th to also appear misclassified. In reality, true events 5-8 were shifted up by one and predicted as events 4-7.

	Predicted Event	True Event
1	25	25
2	22	22
3	23	23
4	2C6	2C4
5	B0	2C6
6	320	B0
7	B2	320
8	2C4	B2
9	20	20
10	223	223

Fig. 4. **One-hot Encoded True Events:** Visualization of a sequence of just over 100 true events pulled from the testing trace. White pixels correspond to the one active element in that column.



Gb of RAM. The machine runs Ubuntu 14.04 LTS 64 bit.

A. TRE Templates

The following two TRE templates are used for the evaluation of both real and synthetic traces and parameterized with a time interval of 0 to 1,000. We executed the mining algorithm on all three types of traces for the purpose of evaluation a) Normal traces, b) Lossy traces, and c) Recovered traces.

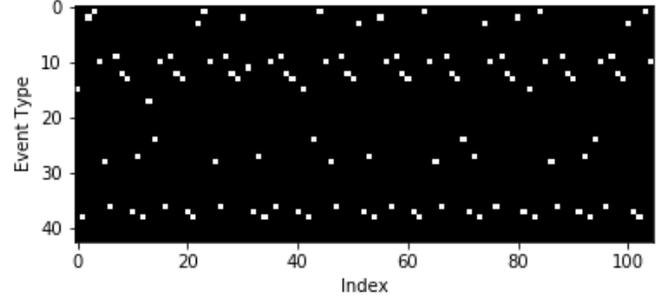
T-1(response): $((P)^* \cdot ((P \cdot (S)^* \cdot S)[0, 1000]) \cdot (P)^*)^+$

T-2(alternating): $((P|S)^* \cdot ((P \cdot (P|S)^* \cdot S \cdot (P|S)^*)[0, 1000]))^+$

V. CONCLUSION

Clean data is a crucial component of training and running numerous machine learning algorithms. Unfortunately, data in the field is often contaminated with noise, loss, and other imperfections. We have developed an LSTM-based approach for restoring lost or noisy data in discrete settings such as CAN traces. This approach can be used to restore or predict arbitrarily large sequences of missing data, using its output as input to predict further into the future. One of the major advantages of this approach is that few assumptions need to be made about the structure of the data, and so it can be applied in any setting where there is discrete data with some form of temporal or sequential dependence.

Fig. 5. **One-hot Encoded Predicted Events:** Visualization of a sequence of just over 100 predicted events pulled from the predictions on the testing trace. White pixels correspond to the one active element in that column.



While the performance of this method was high in the settings we tested it in, there are certainly changes that can be made to further improve it. For example, more work needs to be done to ensure accuracy is just as high when the model is not trained on just a single maneuver but rather on normal driving traces where maneuvers may not be known at time of prediction.

A potential method for increasing the accuracy of lossy data restoration would be using context not only from events preceding the lossy portion, but also from events after it. A bidirectional RNN could be trained to restore lossy sections of data when given some number of events from both before and after those sections. The aforementioned IIN method developed by (Zhou and Huang 2017) does just that, using a forward and backward LSTM to learn the structure of a time series[27]. However, the advantage of using a one-directional LSTM as in our implementation is that once the model is trained, lossy data can be recovered effectively in real-time while a bidirectional recurrent method would require some latency to allow new events to come in.

A second method for increasing accuracy and decreasing dependence on separation by maneuver, would be improved regularization within the network. In particular, zoneout, which stochastically preserves hidden activations instead of dropping them entirely, has recently been shown to be a more effective method for recurrent regularization than recurrent dropout[42].

Our solution of using one-hot encoding and having an "other" node in the dictionary of events works well for data where a majority of event types are known in advance and are consistent between training and test data. Another direction worthy of exploration would be finding a method to avoid having to pre-compile a dictionary of events. This is especially crucial in systems where new events are synthesized, or there are a large number of rare events. One solution could be to use a different embedding for the input data, that can be applied to new events on the fly.

A method that may be able to solve both the issues of performing well on multiple maneuvers and of embedding the input, could be to make use of an encoder network that

would take data as input, perform transformations to embed the data in a different space, and send its output to the LSTM network. This would be more consistent with the sequence-to-sequence networks described by (Sutskever et al. 2014) [1] and we must compare our method with these approaches to determine whether the change in performance merits the additional complexity.

Finally, for training, it is likely that using AdamOptimizer as described in (Kingma and Ba 2014) will lead to faster convergence[43] than using GradientDescentOptimizer with decaying learning rate as we did with our LSTM model. A few limited runs with a smaller model appear to confirm this, but more testing is necessary to determine whether the effect is significant.

Our LSTM-based method is nonetheless an effective standalone solution for denoising or recovering lost discrete data in the field. It can be used to pre-process data intended to be used with any other algorithm in order to improve end-to-end performance.

REFERENCES

- [1] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," *CoRR*, vol. abs/1409.3215, 2014. [Online]. Available: <http://arxiv.org/abs/1409.3215>
- [2] K. Filippova, E. Alfonseca, C. A. Colmenares, L. Kaiser, and O. Vinyals, "Sentence compression by deletion with lstms," in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 2015, pp. 360–368.
- [3] P. Malhotra, L. Vig, G. Shroff, and P. Agarwal, "Long short term memory networks for anomaly detection in time series," in *Proceedings. Presses universitaires de Louvain*, 2015, p. 89.
- [4] G. Cutulenco, Y. Joshi, A. Narayan, and S. Fischmeister, "Mining timed regular expressions from system traces," in *Proceedings of the 5th International Workshop on Software Mining*, ser. SoftwareMining 2016. New York, NY, USA: ACM, 2016, pp. 3–10. [Online]. Available: <http://doi.acm.org/10.1145/2975961.2975962>
- [5] A. Narayan, G. Cutulenco, Y. Joshi, and S. Fischmeister, "Mining timed regular specifications from system traces," *ACM Trans. Embed. Comput. Syst.*, vol. 17, no. 2, pp. 46:1–46:21, Jan. 2018. [Online]. Available: <http://doi.acm.org/10.1145/3147660>
- [6] L. Schmidt, A. Narayan, and S. Fischmeister, "Trem: A tool for mining timed regular specifications from system traces," in *Proceedings of the 32Nd IEEE/ACM International Conference on Automated Software Engineering*, ser. ASE 2017. Piscataway, NJ, USA: IEEE Press, 2017, pp. 901–906. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3155562.3155675>
- [7] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [8] D. Bahdanau, K. Cho, and Y. Bengio, "Neural machine translation by jointly learning to align and translate," *arXiv preprint arXiv:1409.0473*, 2014.
- [9] Y. Duan, M. Andrychowicz, B. Stadie, O. J. Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba, "One-shot imitation learning," in *Advances in neural information processing systems*, 2017, pp. 1087–1098.
- [10] K. Cho, B. van Merriënboer, Ç. Gülçehre, F. Bougares, H. Schwenk, and Y. Bengio, "Learning phrase representations using RNN encoder-decoder for statistical machine translation," *CoRR*, vol. abs/1406.1078, 2014. [Online]. Available: <http://arxiv.org/abs/1406.1078>
- [11] J. M. Paggi and G. Bejerano, "A sequence-based, deep learning model accurately predicts rna splicing branchpoints," *bioRxiv*, 2017. [Online]. Available: <https://www.biorxiv.org/content/early/2017/09/07/185868>
- [12] M. R. Vargas, B. S. L. P. de Lima, and A. G. Evsukoff, "Deep learning for stock market prediction from financial news articles," in *2017 IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications (CIVEMSA)*, June 2017, pp. 60–65.
- [13] M. C. MOZER, "Neural network music composition by prediction: Exploring the benefits of psychoacoustic constraints and multi-scale processing," *Connection Science*, vol. 6, no. 2-3, pp. 247–280, 1994. [Online]. Available: <https://doi.org/10.1080/09540099408915726>
- [14] F. A. Gers and E. Schmidhuber, "LSTM recurrent networks learn simple context-free and context-sensitive languages," *IEEE Transactions on Neural Networks*, vol. 12, no. 6, pp. 1333–1340, Nov 2001.
- [15] S. Hong, S. Kim, M. Joh, and S. Song, "Psique: Next sequence prediction of satellite images using a convolutional sequence-to-sequence network," *CoRR*, vol. abs/1711.10644, 2017. [Online]. Available: <http://arxiv.org/abs/1711.10644>
- [16] E. Marchi, F. Vesperini, F. Weninger, F. Eyben, S. Squartini, and B. Schuller, "Non-linear prediction with LSTM recurrent neural networks for acoustic novelty detection," in *2015 International Joint Conference on Neural Networks (IJCNN)*, July 2015, pp. 1–7.
- [17] A. Taylor, S. Leblanc, and N. Japkowicz, "Anomaly detection in auto-mobile control network data with long short-term memory networks," in *2016 IEEE International Conference on Data Science and Advanced Analytics (DSAA)*, Oct 2016, pp. 130–139.
- [18] L. Bontemps, V. L. Cao, J. McDermott, and N. Le-Khac, "Collective anomaly detection based on long short term memory recurrent neural network," *CoRR*, vol. abs/1703.09752, 2017. [Online]. Available: <http://arxiv.org/abs/1703.09752>
- [19] T. Ergen, A. Hassan Mirza, and S. S. Kozat, "Unsupervised and Semi-supervised Anomaly Detection with LSTM Neural Networks," *ArXiv e-prints*, Oct. 2017.
- [20] A. Javaid, Q. Niyaz, W. Sun, and M. Alam, "A deep learning approach for network intrusion detection system," in *Proceedings of the 9th EAI International Conference on Bio-inspired Information and Communications Technologies (formerly BIONETICS)*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2016, pp. 21–26.
- [21] T. A. Tang, L. Mhamdi, D. McLernon, S. A. R. Zaidi, and M. Ghogho, "Deep learning approach for network intrusion detection in software defined networking," in *Wireless Networks and Mobile Communications (WINCOM), 2016 International Conference on*. IEEE, 2016, pp. 258–263.
- [22] M.-J. Kang and J.-W. Kang, "Intrusion detection system using deep neural network for in-vehicle network security," *PLoS one*, vol. 11, no. 6, p. e0155781, 2016.
- [23] A. R. T. Donders, G. J. Van Der Heijden, T. Stijnen, and K. G. Moons, "A gentle introduction to imputation of missing values," *Journal of clinical epidemiology*, vol. 59, no. 10, pp. 1087–1091, 2006.
- [24] T. Schneider, "Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values," *Journal of climate*, vol. 14, no. 5, pp. 853–871, 2001.
- [25] G. Kalton, "Compensating for missing survey data." 1983.
- [26] P. Royston *et al.*, "Multiple imputation of missing values," *Stata journal*, vol. 4, no. 3, pp. 227–41, 2004.
- [27] J. Zhou and Z. Huang, "Recover missing sensor data with iterative imputing network," *CoRR*, vol. abs/1711.07878, 2017. [Online]. Available: <http://arxiv.org/abs/1711.07878>
- [28] B. B.-N. Hsieh and T. C. Pratt, *Field data recovery in tidal system using artificial neural networks (ANNs)*. U.S. Army Engineer Research and Development Center, 2001.
- [29] O. Tilk and T. Alumäe, "LSTM for punctuation restoration in speech transcripts," in *INTERSPEECH 2015, 16th Annual Conference of the International Speech Communication Association, Dresden, Germany, September 6-10, 2015*, 2015, pp. 683–687. [Online]. Available: <http://www.isca-speech.org/archive/interspeech2015/i150683.html>
- [30] K. Nazmul Haque, M. Abu Yousuf, and R. Rana, "Image denoising and restoration with CNN-LSTM Encoder Decoder with Direct Attention," *ArXiv e-prints*, Jan. 2018.
- [31] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Comput.*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997. [Online]. Available: <http://dx.doi.org/10.1162/neco.1997.9.8.1735>
- [32] J. Lei Ba, J. R. Kiros, and G. E. Hinton, "Layer Normalization," *ArXiv e-prints*, Jul. 2016.
- [33] S. Semeniuta, A. Severyn, and E. Barth, "Recurrent dropout without memory loss," *CoRR*, vol. abs/1603.05118, 2016. [Online]. Available: <http://arxiv.org/abs/1603.05118>
- [34] H. Sak, A. W. Senior, and F. Beaufays, "Long short-term memory based recurrent neural network architectures for large vocabulary speech recognition," *CoRR*, vol. abs/1402.1128, 2014. [Online]. Available: <http://arxiv.org/abs/1402.1128>
- [35] K. Greff, R. K. Srivastava, J. Koutník, B. R. Steunebrink, and J. Schmidhuber, "LSTM: A search space odyssey," *CoRR*, vol. abs/1503.04069, 2015. [Online]. Available: <http://arxiv.org/abs/1503.04069>
- [36] L. Lamport, "Time, Clocks, and the Ordering of Events in a Distributed System," *Communications of the ACM*, vol. 21, no. 7, pp. 558–565, 1978.
- [37] M. B. Dwyer, G. S. Avrunin, and J. C. Corbett, "Patterns in Property Specifications for Finite-State Verification," in *Software Engineering, 1999*.

- Proceedings of the 1999 International Conference on*. IEEE, 1999, pp. 411–420.
- [38] C. Lemieux, D. Park, and I. Beschastnikh, “General LTL Specification Mining,” in *Automated Software Engineering (ASE), 2015 30th IEEE/ACM International Conference on*. IEEE, 2015, pp. 81–92.
- [39] V. Dallmeier, N. Knopp, C. Mallon, S. Hack, and A. Zeller, “Generating test cases for specification mining,” in *Proceedings of the 19th international symposium on Software testing and analysis*. ACM, 2010, pp. 85–96.
- [40] Z. Kincaid and A. Podelski, “Automated Program Verification,” in *Language and Automata Theory and Applications: 9th International Conference, LATA 2015, Nice, France, March 2-6, 2015, Proceedings*, vol. 8977. Springer, 2015, p. 25.
- [41] O. M. Eugene Asarin, Paul Caspi, “Timed Regular Expressions,” *J. ACM*, vol. 49, no. 2, pp. 172–206, Mar. 2002. [Online]. Available: <http://doi.acm.org/10.1145/506147.506151>
- [42] D. Krueger, T. Maharaj, J. Kramár, M. Pezeshki, N. Ballas, N. R. Ke, A. Goyal, Y. Bengio, H. Larochelle, A. C. Courville, and C. Pal, “Zoneout: Regularizing rnns by randomly preserving hidden activations,” *CoRR*, vol. abs/1606.01305, 2016. [Online]. Available: <http://arxiv.org/abs/1606.01305>
- [43] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” *CoRR*, vol. abs/1412.6980, 2014. [Online]. Available: <http://arxiv.org/abs/1412.6980>