View Reviews

Paper ID
6337

Paper Title
Relation Inference among Sensor Time Series in Smart Buildings

Reviewer #1

Questions

1. [Summary] Please summarize the main claims/contributions of the paper in your own words.
This paper proposes a triplet network architecture for learning relational structure on spatially-oriented time series sensor data.The model can learn both "functional" and spatial relationships between sensor pairs, and is evaluated on a real-world dataset in the domain of smart buildings. The model demonstrates good accuracy, but also transfer between buildings, and better accuracy on less available data.

2. [Relevance] Is this paper relevant to an AI audience?
Relevant to researchers in subareas only

3. [Significance] Are the results significant?
Significant

4. [Novelty] Are the problems or approaches novel?
Novel

5. [Soundness] Is the paper technically sound?
Technically sound

6. [Evaluation] Are claims well-supported by theoretical analysis or experimental results?
Sufficient

7. [Clarity] Is the paper well-organized and clearly written?
Excellent

8. [Detailed Comments] Please elaborate on your assessments and provide constructive feedback.
This paper is overall very well-written, has a very extensive evaluation methodology, and shows very good results against several baselines. While this is primarily an applications paper with some narrow scope, it is a very good application study which should meet the threshold for acceptance.

Strong points:
1. The proposed method is very straightforward and very well presented (the best in my group in terms of clarity). The authors use a common architecture but show very good empirical results.

2. The ablation and transfer learning studies are very impressive. The authors have some very interesting take-aways that triplet networks can be trained easier than siamese and may be more appropriate for small and medium sized data. Second, the authors show some transfer generalizability between buildings (Figure 4).

A few points of improvement:

1. Explicit contributions statement. The authors seem to combine several aspects of existing work. It should be clearly stated that their contribution is combining for this specific problem.

2. Greater generalizability: the authors may want to consider similar sensor network datasets outside of the smart

buildings domain. Also, it may be interesting to present an ablation model with minimal/no frequency-domain preprocessing such as the frequency thresholding. Presumably the network can still handle this noise.

3. Figure 4 could be greatly improved by showing boxplots over all pairs of train/test buildings. In the current presentation this looks somewhat like cherry-picking because why these two buildings in particular?

9. [QUESTIONS FOR THE AUTHORS] Please provide questions for authors to address during the author feedback period.

1. How can the model be adapted to categorical and/or sparse stream data? In many domains (e.g image object detection, anomaly detectors) the output may be a label or discrete (and sparse) counts rather than a signal with a fixed sampling rate. Have you considered mixed-domain signal pairs and does the non-STSF variants work as-is here?

2. RelatedI to (1), might you include sensitivity analysis wrt sampling rates? This is analogous to Table 4, but where the sampling *rate* is adjusted in the STSF setting, e.g. using frequency thresholding. This would give an idea of the suitability of the model for low-frequency signals. Furthermore, in many low-resource environments, a lower sampling rate is very advantageous for the sensor network maintenance.

10. [OVERALL SCORE]

7 - Accept

Reviewer #2

Questions

1. [Summary] Please summarize the main claims/contributions of the paper in your own words.

This manuscript presents a novel methodology for identifying relationships between sensor time series in smart buildings. It solves a problem which otherwise need manual intervention. Two kinds of relationships are analyzed, namely, functional and spatial. The methodology is based on three main steps. First, time series readings are transformed from time domain to frequency domain via Short-Time Fourier Transformation (STFT). Then Fourier coefficients are composed by a Deep Metric Learning Triplet Network to generate an embedding which effectively represents functionally and spatially connected sensors. In other words, the Deep Metric Learning Network produces an embedding vector for each sensor, where sensors having functional or spatial relationships have smaller distances among related embedding vectors than sensors having no relationship among each other. Finally, relationships among sensors are detected from sensor embedding vectors by solving a graph cut problem in which nodes represent sensors and edges represent relationships between pairs of sensors (with weights proportional to vector embedding similarity). Experiments performed on two datasets (i.e., six buildings for functional relationship inference and eight buildings for spatial relationship inference) show that the proposed method outperforms three supervised deep learning methods and two unsupervised methods. Moreover, to inspect the influence of STFT a pure Triplet network (without STFT variable preparation) is tested, and to inspect the influence of the triplet loss architecture an approach based on STFT+Siamese Network is tested. Last analyses evaluate the effect of sensor pair selection, cross-building learning and amount of data on performance.

2. [Relevance] Is this paper relevant to an AI audience?

Of limited interest to an AI audience

3. [Significance] Are the results significant?

Significant

4. [Novelty] Are the problems or approaches novel?

Somewhat novel or somewhat incremental

5. [Soundness] Is the paper technically sound?

Technically sound

6. [Evaluation] Are claims well-supported by theoretical analysis or experimental results?
Sufficient

7. [Clarity] Is the paper well-organized and clearly written?
Good

8. [Detailed Comments] Please elaborate on your assessments and provide constructive feedback.
The main strength of the approach here presented is, from my point of view, the performance it achieves on real world datasets with respect to baseline methods. This makes the methodology appealing for practical usage. On the other hand, the main weaknesses is that there is no significant theoretical progress, since state-of-the-art techniques (i.e., STFT, Deep Metric Learning based on triplet loss, standard algorithms for graph min-cut) are only composed together. Moreover, sections Introduction and Related work are very application-oriented, with cited literature quite out-of-scope for the AAAI community (only a few references concern the development of original artificial intelligence and machine learning methods). From my perspective, in order to be published in the AAAI proceedings, the methodology should be further extended with original theoretical elements and proposed as a general approach for time series clustering, instead of as a specific method for inferring relationships among sensors in smart buildings.

The structure of the paper is good. Its clarity can be improved, some sentences are vague (e.g., page 2 "...we appeal to different approximate optimization algorithms..."). Some notation is not clear (e.g., page 3 "the direct current (DC)", please define it). Triplet networks are not formally defined. The elements of the convolution networks (beginning of page 5) are not motivated.

Finally, but very important, I'm not sure to have understood if the performance of Tables 1 and 2 are computed on the training set or on a separate test set (see question below).

9. [QUESTIONS FOR THE AUTHORS] Please provide questions for authors to address during the author feedback period.
I'm not sure to have understood if the performance of Tables 1 and 2 are computed on the training set or on a separate test set. More precisely, how did you use the ground truth to train the model? In page 6 you say "As our algorithm still requires labeled training data to obtain an effective embedding of related sensors..." and in page 5 you say "The ground-truth of the VAV and AHU connection is obtained from the vendor of these buildings" but how did you use this ground truth in the learning process, and in particular in learning the triplet network. If you use it to define pairs of positive and negative sensors isn't it possible that the performance improvement of your method is only due to overfitting (i.e., this could happen if the network you trained is larger than those trained for other methods)? Thank you in advance for your feedback.

10. [OVERALL SCORE]
6 - Marginally above threshold

15. Please acknowledge that you have read the author rebuttal. If your opinion has changed, please summarize the main reasons below.
I have read author rebuttal. It clarifies my doubts about model evaluation (in both the unsupervised and supervised cases) and possible overfitting problems. For this reason I have changed my overall score to "Marginally above threshold".

Reviewer #3

Questions

1. [Summary] Please summarize the main claims/contributions of the paper in your own words.
The paper presents a method for reasoning about relations between smart sensors in buildings.

The method is based on transferring time-series data to the frequency domain using Short-time Fourier Transformation and learning which time-series are related using deep metric learning. The idea is that related time-series will be recognized even though events that could be used for recognition manifest themselves differently in the different time-series and not synchronously. The method handles both spatial and functional relationships.

The paper investigates a very important topic as the deployment of IoT sensors accelerates quickly. Many different providers provide sensors and the meta data that describe them is sparse and spread over different closed systems with diverse meta data structures, so automatic detection of locations and functional relations are extremely relevant.

2. [Relevance] Is this paper relevant to an AI audience?
Likely to be of interest to a large proportion of the community

3. [Significance] Are the results significant?
Highly significant

4. [Novelty] Are the problems or approaches novel?
Novel

5. [Soundness] Is the paper technically sound?
Technically sound

6. [Evaluation] Are claims well-supported by theoretical analysis or experimental results?
Very convincing

7. [Clarity] Is the paper well-organized and clearly written?
Excellent

8. [Detailed Comments] Please elaborate on your assessments and provide constructive feedback.
This paper is well-written, and it is an interesting and nice read. The literature review is nice.

The experiments are sound and support the claims. I trust the evaluation and results.

The results presented in the experiments are very promising and indicates that the method works, although the types of sensors are limited to ventilation systems. The number of time-series and sensor readings in the data set is impressive.

9. [QUESTIONS FOR THE AUTHORS] Please provide questions for authors to address during the author feedback period.
Do you think that this method would work if you included several different types of sensors that not only related to ventilation?

10. [OVERALL SCORE]
7 - Accept

15. Please acknowledge that you have read the author rebuttal. If your opinion has changed, please summarize the main reasons below.
I have read the author rebuttal, which answered the questions in a good manner. I would like to see the added experiments that were mentioned in the rebuttal. I stand by my overall score to accept this paper to AAAI 2020.