A Comparison of Structural Similarity Metrics within Population-Based Structural Health Monitoring

DANIEL S. BRENNAN, ELIZABETH J. CROSS
and KEITH WORDEN

ABSTRACT

Population-based Structural Health Monitoring (PBSHM) aims to gain additional insights into the health of a structure when using data available across a population of similar structures, as compared to the insight available when using only data from a single structure. Before knowledge can be transferred across structures, the similarity between structures (or substructures), within the population must be established.

As a result of recent developments within PBSHM, there are now several methodologies for the comparison of structures, including Graph Matching Networks (GMNs) and the Jaccard index. The recent work on GMNs has highlighted how different methodologies may be required for different structure types or populations.

This paper looks at these aforementioned algorithms and evaluates the performance of each algorithm against different toy and real-world datasets and determines which algorithms have the potential for being used within which scenarios.

INTRODUCTION

Population-based Structural Health Monitoring (PBSHM) [1–3], expands upon the 'traditional' process of Structural Health Monitoring (SHM) [4], by monitoring multiple structures of different types (*the population*). The goal of PBSHM is that by monitoring multiple structures of different compositions and class, additional knowledge on the health of a given structure becomes available, in comparison to available knowledge when using only a single structure's data. To achieve the goals of PBSHM, there are two issues which require addressing; finding which structures (*or components of structures*), are similar [2], and transferring learnt knowledge between those established similar structures.

To establish the similarity of structures, a common language is required to describe structures in a standardised and conformant manner; Irreducible Element (IE) models are such a vehicle used within PBSHM. The recent introduction of the PBSHM database [5], has enabled a unified storage for PBSHM data. With particular relevance to this paper, the inclusion of IE-model data within the PBSHM schema [6], enables the standardised metadata representation required, for structure comparisons at scale.

This paper is the third in a series examining the calculation of similarities within the PBSHM framework. The first paper within the series [7], proposed a potential solution for the inherent issues surrounding expert ambiguity present within IE models; a set of rules (*the Canonical Form reduction rules*) that reduce an IE model down to the Canonical Form (CF), a standardised representation of a structure, free of any ambiguity introduced by a model's author. A Graph Matching Network (GMN) [8] was also introduced as a machine learning method for calculating

Daniel S Brennan, Email: d.s.brennan@sheffield.ac.uk. Dynamics Research Group, Department of Mechanical Engineering, University of Sheffield, Mappin Street, Sheffield, S1 3TD, UK.

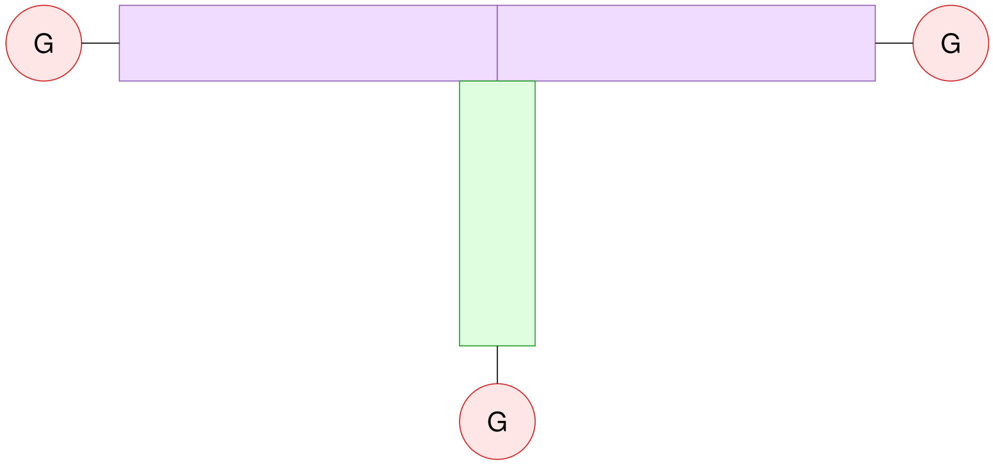

Figure 1. A simplified two-span bridge with two deck elements and one column element. The model interacts with the ground at the left and right of the deck, as well as at the bottom of the column.

the similarities of structures within the PBSHM database.

The second paper in the series, expands upon the idea of using the GMN to determine the similarities. However, instead of calculating the similarities between the structures, the paper explored the use of the CF representation of a structure for calculating the similarities.

This paper builds upon the foundations set in the two previous papers and compares the results of using the Jaccard Index (JI) [2] vs the GMN within the simulated datasets.

BACKGROUND

The objective of Population-based Structural Health Monitoring (PBSHM), is to amass greater knowledge on a structure's health – when viewed within the context of a population – than would be known if the structure was considered in isolation. The vehicle for acquiring this greater knowledge within PBSHM, is *transfer learning* [3], where one can take knowledge gained from structure A and transfer this learnt knowledge to structure B. Whilst the premise of this process may appear straightforward in nature, one must ensure that there is a common characteristic between the structures before transfer of knowledge is executed.

One method in which a common characteristic is determined to be present within PBSHM, is by evaluating the similarity of structures, considering each structures' purpose, composition, and form. However, to facilitate this similarity of structures, each structure must be described in a standardised manner in which the labels assigned to embed structural information, are from a shared space. An Irreducible-Element (IE) model is the aforementioned instrument used to facilitate this shared description within PBSHM.

Each structurally-significant component in the structure becomes an *element* and the interactions between these components become *relationship*s within the model. Knowledge regarding the structural purpose, composition, and form, is embedded via type subdivisions and associated attributes; *ground element*, *regular element*, *perfect relationship*, *connection relationship*, *joint relationship*, and *boun-dary relationship*. In the interests of brevity, only a high-level summary of the IE model types is included here, the interested reader is encouraged to read [6], where an IE-model schema for PBSHM is introduced, including a full breakdown of available types and attributes.

Whilst using a shared set of labels enables a common language description for embedding structural knowledge within PBSHM, there are still topological and attribute variations present within models which must be taken into consideration, when performing comparisons between structures. Distinctions must be made between variations present from variance in the underlying structures and variations present because of differing goals of model authors.

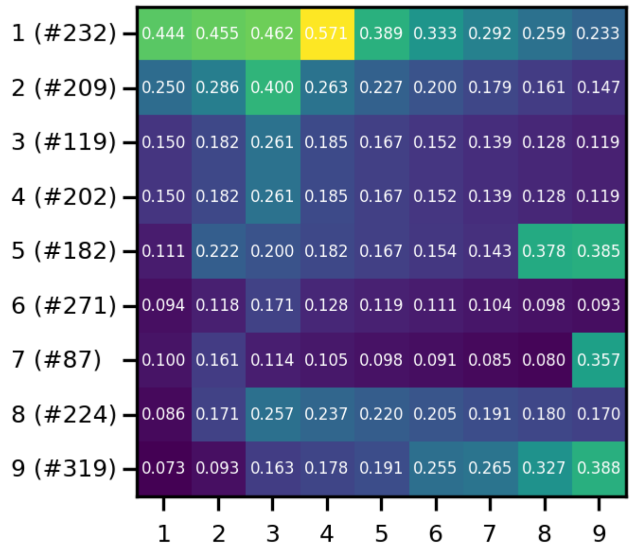

Figure 2. The Jaccard Index results as a heatmap, when using the matching dataset and the 'ideal' Canonical Form dataset.

If one takes the example used in the first paper of this series [7], where there is a simple two-span bridge in the middle of a valley (see Figure 1), there are multiple valid variations on the number of *element*s and *relationship*s used within the model. The beam within the example, could be split into just two *element*s as shown within the figure, or the author may be concerned about damage localisation on the beam, and as such, split the beam into four *element*s, facilitating the damage to be localised to left side of the valley, left of the column, right of the column, or right side of the valley. Both of the aforementioned variations are valid IE models and may be included within the PBSHM database.

The *Canonical Form* (CF) representation, preserves the structural knowledge and engineering decisions embedded within an IE model, whilst providing a common form for structural-similarity calculations. A non-canonical form version (*the detailed version*) of an IE model would be provided by an author and stored within the PBSHM database, a CF version of that IE model would then be generated any time the similarity of structures was to be determined. The generation of a CF from a non-canonical form IE model is achieved by applying the Canonical Form reduction rules against the non-canonical form IE model.

These CF reduction rules, outline how and why particular patterns from within the IE model may be reduced to a simpler pattern, without impairing the key structural knowledge of the model. At the time of writing, there are currently three CF reduction rules; every *boundary relationship* must be to a unique *ground element*, perfect-joint-joint *relationship* loops can be reduced to a single perfect-joint *relationship* and multiple *perfect relationship*s can be reduced to a single *perfect relationship*. For the interested reader, these rules, and the basis for why they can be used, is described in detail in the first paper of this series [7].

The Graph Matching Network (GMN) [8] is a recent development within the Graph Neural Network (GNN) [9] community, where a similarity score can be generated between two graphs via a cross-graph attention-based mechanism. This work was used within the first part of this series [7] to evaluate if the GMN could be trained using pairs of non-canonical IE models and subsequently, if a similarity could be drawn against the simulated simple beam-and-slab bridge example.

This work was then further expanded in Part II of this series [10], by determining if the CF could be used as a common form for the similarity metrics themselves. Previously, the work considered within structure similarities for PBSHM, has been calculating the similarities between every pair of structures within the PBSHM database. Part II proposed, that instead of having to generate similarity metrics between every structure within the database, instead a list of known CF IE models could exist within the database and similarity comparisons be drawn from the non-canonical IE models and the CF IE models. This idea had two main advantages; it would vastly reduce the number of computations required to be performed within the database and, it could potentially provide a natural clustering mechanism for the generation of yet unknown populations within the database. As such, the GMN was modified to learn similarities from non-canonical form IE models to CF IE models in a bid to explore this approach.

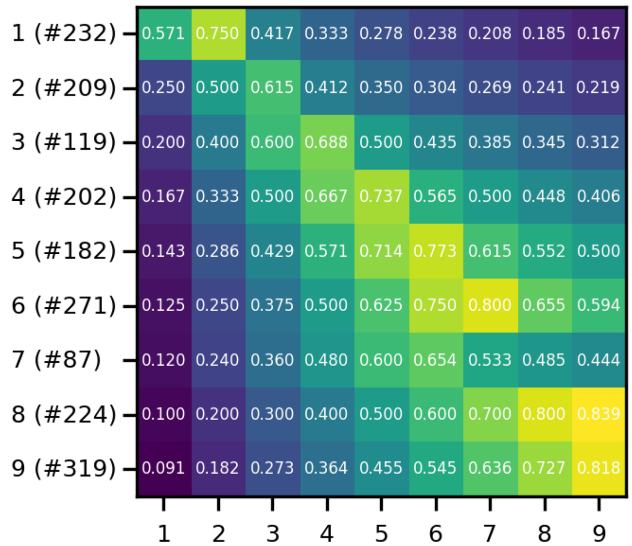

Figure 3. The Jaccard Index results as a heatmap, when using the CF reduction rules on the matching dataset and the 'ideal' Canonical Form dataset.

Previous work within PBSHM, used the Jaccard Index [2, 11] to generate similarity metrics between a small dataset of bridges. This paper takes the Jaccard Index used by Gosliga et al. [2] and evaluates the use of the CF as a common form within the Jaccard Index similarities and compares the similarity results to those from the GMN. Both of the algorithms used within the paper to calculate the similarity of structures, function within a graph space. An IE model can be converted into an Attributed Graph (AG) via each *element* becoming a node and each *relationship* becoming an edge. *Element* and *relationship* attributes are then embedded into attributes within the associated node/edge of the AG.

In the interest of clarity, all the metrics shown within this paper have been scaled between 0 and 1, where 0 is dissimilar and 1 is similar. This is in contrast to the previous papers within this series, however, the final result of Part II of this series, are included in Figure 8a in the Appendix using this papers scaling factors, for the reader's reference.

Like the previous papers in this series, the datasets used contain IE models based upon a simple beam and slab bridge. Each IE model has a varying number of spans (two to ten), with a varying number of nodes per span and varying PJ and PJJ relationships between the beams and the columns. All the datasets used within this paper are freely available via the GitHub repository: https://github.com/dsbrennan/ie-model-learning-data. Within this paper, the datasets used are referenced via a common name, instead of the dataset ID; matching dataset ('learning-slab-bridge-500-span-1-9-column-1-1-node-1-3'), ideal CF dataset ('learning-slab-bridge-canonical-form-span-1-9-column-1-1'), and interim CF dataset ('learning-slab-bridge-interim-canonical-form-span-1-9-column-1-1').

JACCARD INDEX VS GRAPH MATCHING NETWORK

The Jaccard Index works by calculating the Maximum Common Subgraph (MCS) between two Attributed Graphs. Whilst the Jaccard Index allows the embedding of multiple node attributes within a graph for comparison, all the attributes must match for the node to be included within the MCS. Subsequently, the only *element* attributes embedded within the graphs are the contextual attributes, as these are present within the matching IE models and the CF IE models. Figure 2 shows the output of the Jaccard Index similarity matrix, when comparing the standard test subset of the matching dataset to the associated 'ideal' CF dataset. As one can see, the Jaccard Index isn't able to correctly find any pattern of similarities between the input IE models and the reference CF. This is because the MCS method used to generate the Jaccard Index is incredibly sensitive to topological changes present within the matching dataset.

As mentioned previously, the matching dataset contains bridges with varying number of nodes per span and varying relationships between these spans and the columns. However, if one implements the CF reduction rules and reduces the matching dataset via those rules before comparison (see Figure 3), one

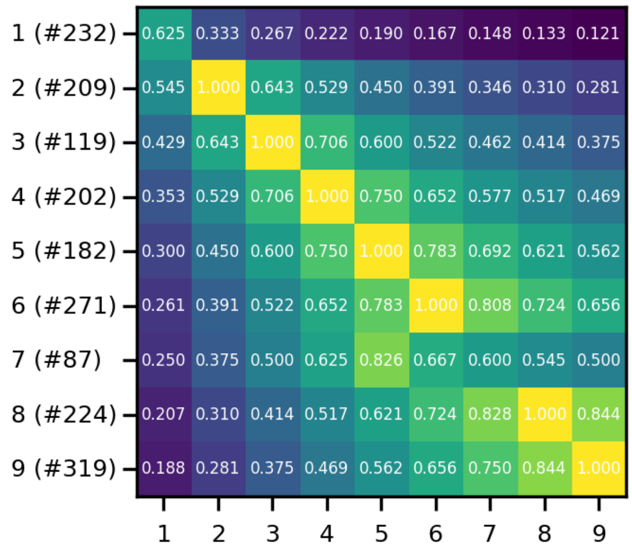

Figure 4. The Jaccard Index results as a heatmap, when using the CF reduction rules on the matching dataset and the 'interim' Canonical Form dataset.

can see that the Jaccard Index is now able to find a pattern of similarity. Whilst the Jaccard Index is now able to notice a pattern of similarity between the matching IE dataset and the 'ideal' CF dataset, one will notice that the pattern isn't yet matching a structure fully – value of 1 – or aligning the matching of a N-span IE model to the corresponding N-span CF IE model.

These algorithmic alignment issues, are a symptom of a incomplete list of CF reduction rules. The CF IE model used within the ideal CF dataset, is the 'ideal' CF IE model; the currently hypothesised eventual version of a CF IE model; however, with the current CF reduction rules available, one is not able to reduce an IE model to this level of graph. As such, an interim CF dataset is required, which has CF IE models at only the level at which one is currently able to reduce with the CF reduction rules. Figure 5 depicts the difference between an IE model, the 'ideal' CF IE model, and the current 'interim' CF IE model, to which we can currently reduce.

If one changes the Jaccard Index, to use the CF reduction rules on the matching dataset and perform the MCS against the interim CF dataset instead of the ideal CF dataset (see Figure 4), one can see that the Jaccard Index is now able to detect a much improved pattern of similarity within the datasets, the alignment of the pattern has now been corrected so that – on the whole – the strongest similarity for an N-span IE model, is the corresponding N-span CF, with the similarity values reducing as the number of spans increase/decrease, and that some of IE models are now providing a full match – i.e. a value of 1 – within the datasets.

The full matrix of options for the Jaccard Index, varying both the 'ideal' and 'interim' dataset and with or without the CF reduction rules are included in Figure 7 for the intrigued reader.

Currently, the GMN available from the results of Part II, embeded both the *element* and *relationship* types as feature attributes within the graph; however, to enable a fair comparison between the Jaccard Index and the GMN, the embedding of the *relationship* features is removed, and the only *element* attribute embedded into the graph is the contextual information. The contextual attributes are encoding into the node features via one-hot encoding, where each potential value within the contextual attributes are given their own dedicated position within a feature array. Each position within the feature array is set to 0/1 depending on if the associated contextual value is true for that position within the feature array. For the interested reader, Figure 8 in the Appendix shows the outputs from the GMN when using ordinal encoding vs one-hot encoding.

Figure 6 shows the results for the GMN when using the matching dataset and both the ideal CF dataset and the interim CF dataset. As one can see, the GMN is able to find the correct pattern for similarity regardless of if the CF is the ideal or interim version. This highlights the power of a machine-learning approach to similarity comparison, as the algorithm is able to learn yet-unknown reductions to be applied to the graph.

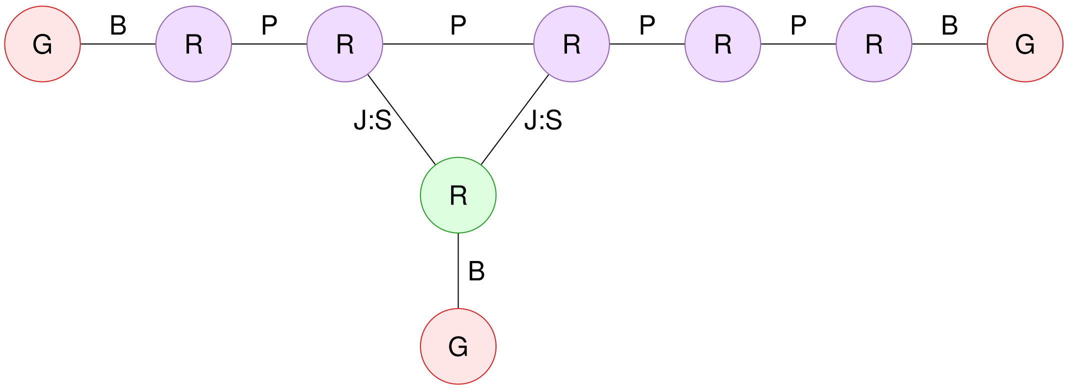

(a) A possible IE model representation of the two-span bridge depicted in Figure 1 which can be reduced to its CF using the PJJ and Perfect Relationship reduction rules.

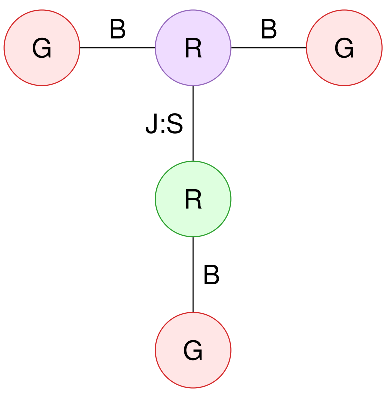

(b) The ‘ideal’ CF representation of the IE model in Figure 5a.

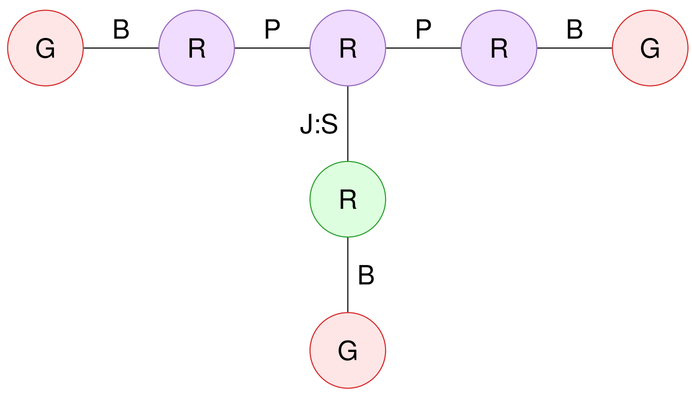

(c) The ‘interim’ CF representation of the IE model in Figure 5a.

Figure 5. The IE model reduction that could be generated from the two-span bridge as depicted in Figure 1. Ground elements are represented by a G in the center of the node, regular elements are represented by a R in the centre of the node. Boundary relationships are represented by a B on the edge, perfect relationships are represented by a P on the edge and static joint relationships are represented by a J:S on the edge.

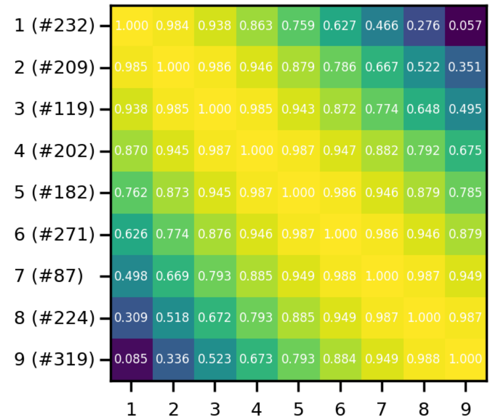

(a) The GMN results when learning the similarity between the matching IE model dataset and the 'ideal' CF dataset.

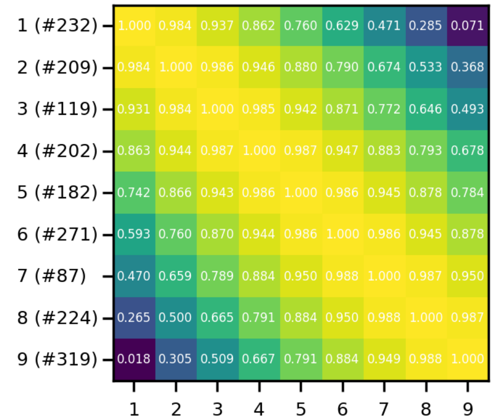

(b) The GMN results when learning the similarity between the matching IE model dataset and the 'interim' CF dataset.

Figure 6. The results of the Graph Matching Network when trained on the matching IE model dataset and both the 'ideal' CF dataset and the 'interim' CF dataset.

The code used to perform the CF reduction rules, is available via the following GitHub repository: https://github.com/dsbrennan/ie-cf-reduction. The code used to generate the results included within this paper, are available from the following GitHub repository: https://github.com/dsbrennan/iwshm-2023-gmn-ji-comparison.

CONCLUSION

In conclusion, this paper had evaluated the use of the Jaccard Index algorithm using the MCS, in comparison to the newly proposed GMN. Without the use of the CF reduction rules, the Jaccard Index is unable to find a similarity pattern within the simulated dataset; however, after implementing the CF reduction rules as part of the similarity comparison, the Jaccard Index is able to find a pattern and correctly identify matching N-span bridges. The GMN – being a machine-learning algorithm – is able to learn the CF reduction rules independently of the currently-known CF reduction rules.

Further research is required into the CF reduction rules; to expand our knowledge of how IE models can be reduced without losing the definitions of the structure within the model. Whilst this paper has used a simulated toy dataset to evaluate the performance, research is required into real world datasets, and comparing the performance of these algorithms within these datasets.

ACKNOWLEDGMENT

The authors of this paper gratefully acknowledge the support of the UK Engineering and Physical Sciences Research Council (EPSRC) via grant reference EP/W005816/1. For the purpose of open access, the authors has applied a Creative Commons Attribution (CC BY) licence to any Author Accepted Manuscript version arising.

REFERENCES

1. Bull, L., P. Gardner, J. Gosliga, T. Rogers, N. Dervilis, E. Cross, E. Papatheou, A. Maguire, C. Campos, and K. Worden. 2021. "Foundations of Population-based SHM, Part I: Homogeneous populations and forms," *Mechanical Systems and Signal Processing*, 148.

2. Gosliga, J., P. Gardner, L. Bull, N. Dervilis, and K. Worden. 2021. “Foundations of Population-based SHM, Part II: Heterogeneous populations – Graphs, networks, and communities,” *Mechanical Systems and Signal Processing*, 148.
3. Gardner, P., L. Bull, J. Gosliga, N. Dervilis, and K. Worden. 2021. “Foundations of Population-based SHM, Part III: Heterogeneous populations – Mapping and transfer,” *Mechanical Systems and Signal Processing*, 149.
4. Farrar, C. R. and K. Worden. 2007. “An introduction to structural health monitoring,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1851):303–315.
5. Brennan, D. S., C. T. Wickramarachchi, E. J. Cross, and K. Worden. 2022. “Implementation of an Organic Database Structure for Population-Based Structural Health Monitoring,” in K. Grimmelsman, ed., *Dynamics of Civil Structures, Volume 2*, Springer International Publishing, Cham, pp. 23–41, series Title: Conference Proceedings of the Society for Experimental Mechanics Series.
6. Brennan, D. S., J. Gosliga, E. J. Cross, and K. Worden. 2021. “On implementing an irreducible element model schema for Population-based Structural Health Monitoring,” :11.
7. Brennan, D. S., T. J. Rogers, E. J. Cross, and K. Worden. “On Quantifying the Similarity of Structures via a Graph Neural Network for Population-based Structural Health Monitoring,” .
8. Li, Y., C. Gu, T. Dullien, O. Vinyals, and P. Kohli. 2019. “Graph Matching Networks for Learning the Similarity of Graph Structured Objects,” *arXiv:1904.12787 [cs, stat]*.
9. Bacciu, D., F. Errica, A. Micheli, and M. Podda. 2020. “A Gentle Introduction to Deep Learning for Graphs,” *Neural Networks*, 129:203–221.
10. Brennan, D. S., T. J. Rogers, E. J. Cross, and K. Worden. “Calculating Structure Similarity via a Graph Neural Network in Population-based Structural Health Monitoring: Part II,” .
11. Gosliga, J., D. Hester, K. Worden, and A. Bunce. 2022. “On Population-based structural health monitoring for bridges,” *Mechanical Systems and Signal Processing*, 173.

APPENDIX

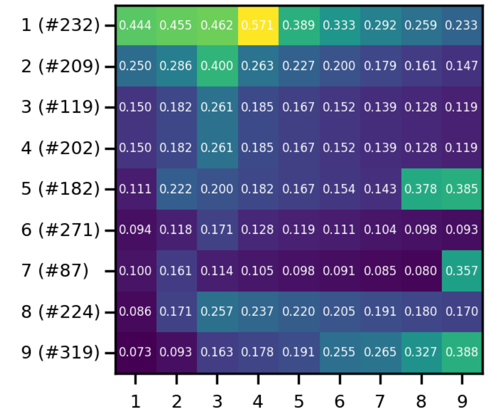

(a) The results of the Jaccard Index on the matching dataset and the ‘ideal’ CF dataset.

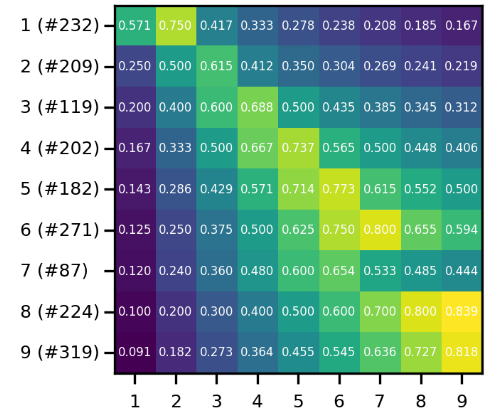

(b) The results of the Jaccard Index on the matching dataset using the CF reduction rules and the ‘ideal’ CF dataset.

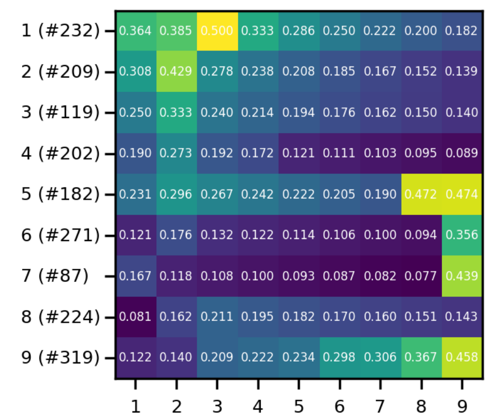

(c) The results of the Jaccard Index on the matching dataset and the ‘interim’ CF dataset.

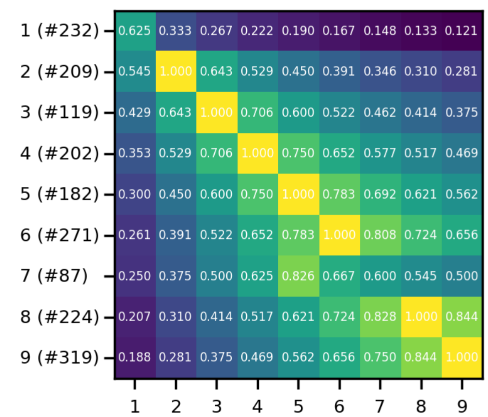

(d) The results of the Jaccard Index on the matching dataset using the CF reduction rules and the ‘interim’ CF dataset.

Figure 7. The results of the Jaccard Index using both the ‘ideal’ and ‘interim’ datasets with and without the CF reduction rules being applied.

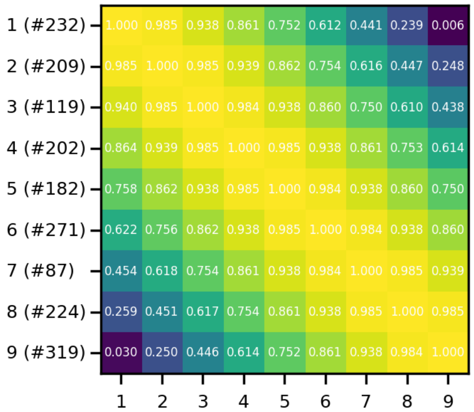

(a) The results of the GMN with embedding the object type against the 'ideal' CF dataset using ordinal encoding.

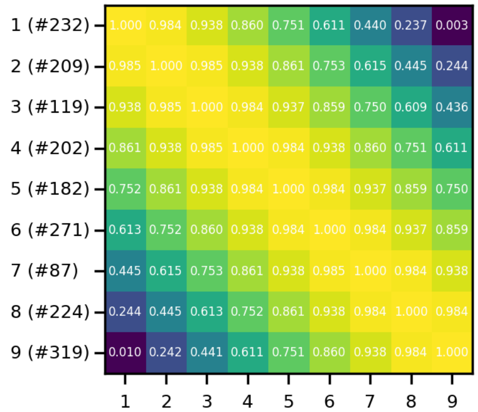

(b) The results of the GMN with embedding the object type against the 'ideal' CF dataset using one-hot encoding.

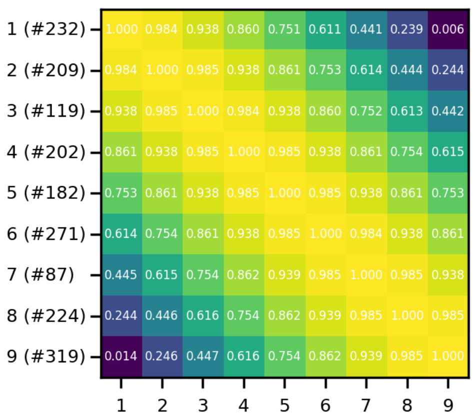

(c) The results of the GMN with embedding only the element's contextual type against the 'ideal' CF dataset using ordinal encoding.

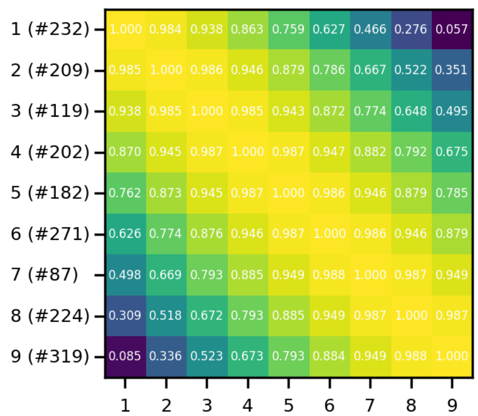

(d) The results of the GMN with embedding only the element's contextual type against the 'ideal' CF dataset using one-hot encoding.

Figure 8. The results of the GMN algorithm with varying the type of data embedded and the method used to encode the data.