



HAL
open science

Achieving maritime situational awareness using knowledge graphs: a study

Jacques Everwyn, Bruno Zanuttini, Abdel-Allah Mouaddib, Sylvain Gatepaille,
Stephan Brunessaux

► To cite this version:

Jacques Everwyn, Bruno Zanuttini, Abdel-Allah Mouaddib, Sylvain Gatepaille, Stephan Brunessaux. Achieving maritime situational awareness using knowledge graphs: a study. 1st Maritime Situational Awareness Workshop (MSAW 2019), Oct 2019, Lerici, Italy. hal-02942856

HAL Id: hal-02942856

<https://hal.science/hal-02942856v1>

Submitted on 18 Sep 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Achieving maritime situational awareness using knowledge graphs: a study

Jacques Everwyn^{a,b}, Bruno Zanuttini^a, Abdel-Allah Mouaddib^a,
Sylvain Gatepaille^b and Stephan Brunessaux^b

^aNormandie Univ., UNICAEN, ENSICAEN, CNRS; GREYC, 14000 Caen, France;

^bAirbus Defence and Space, Elancourt, France

ABSTRACT

Currently, maritime surveillance operators have to monitor by hand the massive amount of data at their disposal to spot the events of interest, thus limiting their capabilities. Maritime data comes from various and heterogeneous sources that can be merged into a dynamic attributed knowledge graph which represents an evolving maritime situation. Using this graph, the automation of alert rising comes through a link prediction task: given some labels from expert knowledge, are there similar situations of interest elsewhere in the graph? In this article, we review link prediction techniques for situation awareness in a maritime context, and draw conclusions on how the addition of attributes in a dynamic graph model could improve results on this task.

Keywords: Dynamic knowledge graph, maritime situation, attributes, machine learning, link prediction

1. INTRODUCTION

The maritime domain is the theater of many unlawful activities that may go unnoticed: terrorism, piracy, smuggling, illegal immigration... For this reason, Maritime Situational Awareness (MSA) is of first importance to maritime security. It is defined by NATO as “The understanding of military and non-military events, activities and circumstances within and associated with the maritime environment that are relevant for current and future NATO operations and exercises where the Maritime Environment (ME) is the oceans, seas, bays, estuaries, waterways, coastal regions and ports”.¹ MSA is often performed by surveillance operators who monitor the flow of data coming from maritime activities. This data is diverse, heterogeneous and coming from several sources: AIS (Automatic Identification System), radars, satellites, intelligence, websites... With more than 50.000 vessels sailing the oceans each day, there is a need for automation in the detection of illicit events.²

A maritime situation implies evolving entities: vessels, ports, countries... Such a situation can be represented by a *dynamic attributed* knowledge graph (DAKG), and understanding how its elements connect and jointly evolve gives valuable information pertaining to MSA. This task is here reduced to a link prediction problem. A *link*, or an *event*, is a relation between two entities at a given time point, for instance (Titanic ; :builtBy ; WhiteStarCompany ; 1909), and *attributed* means that entities have attributes whose values may change over time, e.g. (Titanic ; :passengers ; 2,344 ; April 10th 1912).

Generally, link prediction is performed by learning an embedding for each entity of the graph and predictions are made by ranking the events in the graph using these embeddings. This benefits to MSA in two ways:

- *data completion*: when monitoring an operational situation, the sensors and intelligence services do not always have all the needed information at their disposal. Using link prediction, missing data can be inferred to improve MSA;
- *automated alerts*: link prediction will discover events that a human operator would not have noticed in the massive dataset. Illegal activities could also be anticipated by making prediction in the future and evaluating the risk a ship represents based on its current and past behavior.

In this article, we review (1) two models on a dynamic (but not attributed) knowledge graph, (2) the literature on static/dynamic/attributed knowledge graphs, (3) how to apply DAKGs to MSA.

Further author information:

a: firstname.name@unicaen.fr

b: firstname.name@airbus.com

2. PREVIOUS WORK

The previous work related to this study can be broadly divided into four categories: maritime related work, static graphs, dynamic graphs and attributed graphs.

2.1 MSA

MSA often focuses on anomaly detection.³ It can be tackled with clustering,² bayesian networks⁴, self-organizing maps⁵ and many others techniques⁶. Route estimation is also handled, e.g. with neural networks⁷ or Extended Kalman filter.⁸ These strategies perform well on trajectory analysis but do not take the whole context of a situation to detect an anomaly or a threat. To the best of our knowledge, this is the first attempt of using link prediction on DAKG to improve MSA.

2.2 Static Knowledge Graph

In a static setting, each node is represented by a single vector. This field is largely covered with a broad range of techniques. Translational models evaluate a fact by measuring the distance between the two entities, generally using the relation during the translation. TransE⁹ is its most known representative. Semantic matching models are similarity-based and compare the latent semantics of entities and relations embeddings. RESCAL¹⁰ was the first to do this and has been extended multiple times.^{11,12} Neural network architectures have also been tried with NTN¹³ or VGAE.¹⁴ These models achieve great performances on static knowledge graphs but are not suited to deal with dynamic ones.

2.3 Dynamic Knowledge Graph

In a dynamic setting, each node is represented by a time series of vector modeling its evolution. This topic is emerging and has less contributions but advances have already been made. Leblay et al.¹⁵ predict time validity for unannotated edges using side information in the learning process. Esteban et al.¹⁶ update the knowledge graph using an event graph to add new information, and Trivedi et al.¹⁷ extend the bilinear model (RESCAL) with a LSTM network in order to learn non-linearly evolving entities. Jiang et al.¹⁸ incorporate the valid time of facts using a joint time-aware inference model based on Integer Linear Programming. Self-attention networks were tried by Sankar et al.¹⁹ Although these models are time-aware, they do not include attribute information in the relation prediction task and we will show that they are needed when dealing with MSA.

2.4 Attributes

KR-EAR²⁰ can predict discrete attribute values and find correlation between them. However, they are not included during the learning of relations and relations are not included in the learning of attributes.²¹ propose a model that jointly learns KG^R and KG^A with a neural network and predicts continuous values with a regression task. However, neither model deals with temporal data.

Li et al.²² propose a streaming model (SLIDE) on dynamic attributed networks using a sketching matrix that summarizes the currently observed links and node attributes. They review the challenges pertaining to such networks and real-world data, but they apply it on social networks (Epinions, DBLP, ACM) that have very different kinds of attributes and only a few widely separated timesteps (~20 timesteps from one month to one year each). All these models showed that the addition of attributes improves the results on link prediction.

Table 1 compares representative models from each category and shows that no model currently fits our needs perfectly (see Part 4 for detailed requirements).

| | Static | Dynamic | Attributes | Near Real-Time |
|---------------------------|--------|---------|------------|----------------|
| TransE ⁹ | X | | | |
| Know-Evolve ¹⁷ | | X | | |
| MT-KGNN ²¹ | X | | Continuous | |
| SLIDE ²² | | X | Discrete | ~ |

Table 1: Application domain of models comparison

3. PROBLEM STATEMENT

In this section, we define a knowledge graph structure to represent a maritime situation.

3.1 Knowledge graph

DEFINITION 1 (FRAME). A frame is a quadruple $F = \langle E, R, A, D \rangle$ where E , R , and A are finite sets of elements called entities, relations, and attributes, respectively, and $D : A \rightarrow S_a$ is a function assigning a range $D(A)$ to each attribute and S_a is a set of possible values for $a \in A$ (discrete or continuous).

DEFINITION 2 (DYNAMIC ATTRIBUTED KNOWLEDGE GRAPH). Let $F = \langle E, R, A, D \rangle$ be a frame. A standard knowledge graph on F is a couple $KG = \langle KG^R, KG^A \rangle$, where

- KG^R is a finite subset of $E \times R \times E \times \tau$, with τ the set of time points,
- KG^A is a finite subset of $E \times A \times D(A) \times \tau$ such that for all quadruples $(e, a, v, t) \in KG^A$, $v \in D(a)$ holds.

For $KG = \langle KG^R, KG^A \rangle$, KG^R is called the relational part of KG , and KG^A is called its attributional part.

Intuitively, (e^s, r, e^o, t) is read “entity e^s is in relation r with entity e^o at time t ”, and (e, a, v, u) is read “entity e has value v for attribute a at time u ”. Given a knowledge graph KG , we always write KG^R (resp. KG^A) for its relational (resp. attributional) part. Figure 1 is an example of the previously defined KG .




Figure 1: Example of knowledge graph KG on a frame $F = \langle E, R, A, D \rangle$. The nodes $e_i \in E$ are entities and the nodes $a_j \in A$ are attributes. The relations $r_n \in R$ are annotated on edges between two entities. The values v_k , annotated on the edges between two attributes, belong to the domain $D(A)$ of the attribute they are attached to, and the t_j s are the timestamps of the edges. Attributes of entities can change their value over time (e_3 and a_2) and two entities can have common attributes (e_4 , e_1 and a_1) but not necessarily with the same value. The blue square nodes and edges are KG^R and the red edges with all the nodes are KG^A .

3.2 Prediction problem

We are interested in predicting the missing relations between entities and values of attributes in knowledge graphs. We focus on the case where for some timestamp, they can be predicted from the values of a subset of the relations and attributes at previous timestamps. In this article, we experiment attribute prediction (position) from past attributes.

4. APPLICATION TO MSA

“Real-world” datasets often have more constraints than the academic ones because of their specificities. Maritime datasets are no exception and the following challenges must be overcome.

- *Evolution of attributes*: a maritime situation is a fast evolving world with very little time between two events. In MSA, a good evolutionary model is needed for change detection and the granularity depends on the task. For a change in the position/course/speed of a vessel (dynamic attributes), the information must be given within minutes (e.g. rapid response needed in case of piracy). But to detect a change in a vessel particulars (identifier, name...), the granularity needed can be in hours or days.
- *Event and threat detection*: an event (or quadruple) represents a new relation between two entities or an abovementioned attribute evolution. Using knowledge graphs and machine learning, it could be possible to find events using latent features that cannot be perceived by a human or a rule. If event mining detects facts, threat detection is a task highly related to its context and definition. A nation will not consider a transshipping between two fishing vessels as a threat since they are more likely to exchange fish than warheads, but an NGO for ocean conservation can suspect illicit fishing of an endangered species. Performing this task still requires either expert knowledge or labeled events.
- *Streaming*: MSA requires a constant monitoring of maritime areas, meaning that the model must deal with a continuous flow of data.
- *Uncertainty*: maritime data often results from hard (sensors) and soft (websites, intelligence) data fusion. However, this data is not always 100% certain: an intelligence report may have a typo, sensors have a range and precision (e.g. +/- 500 meters), or collisions may happen when satellites receive signals.
- *Explainability*: link prediction models are often black boxes when it comes to the origin of the prediction. However, a surveillance operator needs to know why a prediction was made in order to understand it and justify any upcoming response to an event. Because operators still do not trust AI-based systems to take decisions, explainability is needed to take DAKG-based decisions for MSA.³⁰

An illustration of all these concepts can be found in Figure 2.

5. EXPERIMENTS

In these experiments, the performed task is the prediction of the position of a vessel in the next time points. Obviously, there are many better fitted methods to do this (like regression or a Kalman filter), but the ultimate goal is to use the full capacities of the knowledge graph i.e. exploit all the relationships, events and attributes in the maritime surveillance ecosystem to perform better link predictions. Position prediction is just a reduction of this task to test knowledge graphs capabilities on MSA. As we could not find a method handling both time and attributes that can be tested on our data, the attribute *:location* is replaced by a relation *:isLocatedIn* between a vessel and an area.

5.1 Dataset

In the absence of publicly available maritime knowledge graph, we created our own in order to evaluate the models.

5.1.1 AIS data

The dataset used in our experiments is based on real maritime data: AIS messages transmitted by vessels. AIS is a short range (37-74km) ship-to-ship and ship-to-shore navigational data exchange system. It is currently the main source of information available in support of maritime surveillance. The satellite version of AIS (S-AIS) gives a broader range (~5000km) but the transmissions are less regular and more subject to signal collision.² AIS provides the following non-exhaustive list of information about ships: the unique identifier of the vessel (called MMSI), its longitude/latitude, its speed and course, the timestamp of the report, the type of ship, the destination.

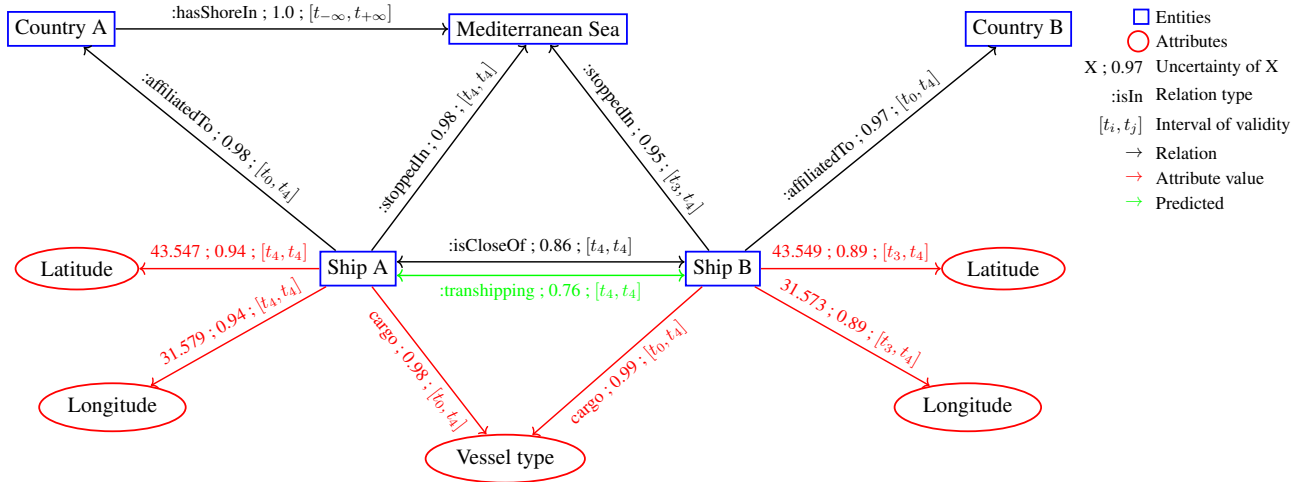


Figure 2: Example of KG^* at time t^* (best viewed in color). Events in red represent KG^A and the other events KG^R . Ship A and Ship B are vessels from different countries, moving in the Mediterranean Sea. They have a static attribute (vessel type) setting them as cargos and dynamic attributes (latitude and longitude) revealing their positions. Before t^* , the two vessels were moving in the Mediterranean Sea and had different positions, but now (t^*) they have stopped and are close to each other. A possible link prediction from $KG^{<t^*}$ is that the two vessels are performing transshipping. Note that the $:\text{isStopped}$ relation can be deduced from the speed attribute going to zero, not represented here for the sake of clarity. If relations such as CountryA having a shore on the Mediterranean Sea are 100% sure, some are more uncertain: the position of Ship B is only 89% sure because the signal was picked up by a satellite in an area with high ship density. More, the uncertainty of predicted relations ($:\text{transshipping}$) depends on the uncertainty of the root events. Finally, to predict the transshipping action in time, the model must be updated with the root causes as soon as they are available, hence the need for streaming link prediction.

5.1.2 AIS to KG

A knowledge graph can be built using these AIS messages, where vessels are entities with attributes. Other entities can be added like nations (flag of the ship) or ports. However, the reviewed methods can only handle time, not attributes, hence the need to consider attributes as entities. In our work, the focus is on the evolution of the positions of vessels. Positions being continuous values, they need to be discretized to be casted as entities in the graph. Therefore, the studied area is converted into a grid made of $1\text{km} \times 1\text{km}$ squares and each square is an entity (further referred to as "areas").

Moreover, AIS messages are on average received every three minutes so it can be a reasonable choice to separate each time point by three minutes, instead of having a time point every second as it happens in the data (different events can be attached to the same time point). Finally, as the chosen models can not always handle entities or relations not encountered during the training phase, the test set is filtered to remove any event involving an entity or relation not present in the train set. Note that only one relation type is considered here: "vessel $:\text{isLocatedIn}$ area" ($|R| = 1$) and each event is represented by a quadruple $(e^s, r, e^o, t) \in KG^R$.

To summarize, we build the knowledge graph consisting of entities = {vessels, areas} and relation = { $:\text{isLocatedIn}$ } over one month, we divide it into train/test sets and run the methods to predict the relation " $:\text{isLocatedIn}$ " between vessels and areas.

The dataset covers the Gibraltar Strait from February 2nd, 2017 to March 2nd, 2017. It holds around 2.5k vessels and one million positions. The two evaluated methods are TransE⁹ and Know-Evolve.¹⁷ They use embeddings (vectors of real values) to represent entities and make link predictions.




Figure 3: One week of maritime traffic in the Gibraltar Strait (best viewed in color)

5.2 Evaluation task

5.2.1 Link prediction

The evaluation is performed on the link prediction task: given a quadruple (e^s, r, e^o, t) , e^o is replaced by every possible entity and the resulting quadruple is evaluated by the model. All the quadruples are then ranked in descending order of plausibility and we record the Mean Average Rank (MAR) and the @Hits10 measure (one of the 10 best ranked quadruples is the true one). A lower rank means that the quadruple is classified better (the best rank being 1 and the worst the number of entities) and @Hits10 is expressed in percentage of correctly ranked quadruples i.e. higher is better. The filtering method of TransE⁹ is applied, i.e. the quadruple is not ranked against corrupted (i.e. modified) quadruples that are true.

5.2.2 Sliding window evaluation

The performance is tested using the sliding window evaluation from Know-Evolve. We divide the test set into 8 different slides, each slide including one day of time (Know-Evolve uses 12 slides of two weeks each). This method is said to *“help to realize the effect of modeling temporal and evolutionary knowledge”*¹⁷.

5.2.3 Static method on dynamic data

As it is a static method, the evaluation of TransE required some modifications of the dataset. All the timestamps t are removed and as a result, multiple occurrences of the same triples (e^s, r, e^o) appear. Those are removed in order to have a unique representant for each triple and the dataset is then comprised of 102,470 (train) and 16,807 (test) events. The test set still only contains entities seen during training.

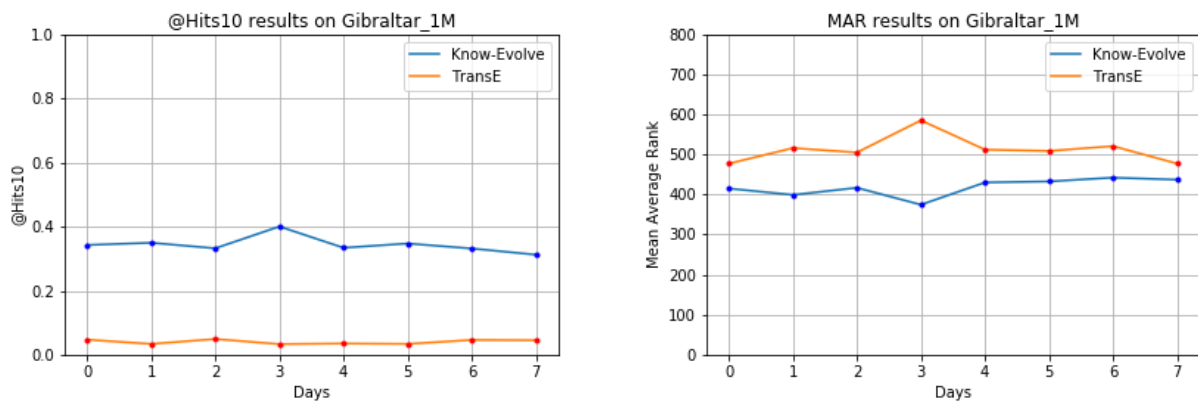


Figure 4: @Hits10 and MAR results of tested models

5.3 Results

5.3.1 Quantitative Analysis

Figure 4 show the results of the reviewed models over the Gibraltar1M dataset. Know-Evolve, being a temporal model, performs way better than TransE which struggle in @Hits10 prediction despite being not so far from Know-Evolve in Mean Rank. The reason is that TransE depend only on static entity embeddings to perform prediction. With an average @Hits10 of 34%, Know-Evolve captured the relationship between the vessels and the areas better than TransE but do not excels at the task.

5.3.2 Contextual analysis

This evaluation was made on a single task: predict in which area a vessel will be next. This task is made harder by the discretization of the positions: areas are independant and the graph does not tell if areas are close to each other or not. The only way to extract proximity is the analysis of a vessel's track (the succession of relations with area entities), meaning that two areas having a relation with a vessel in a short timespan may be close. More, a proximity relationship between two vessels in the same area could not be established because areas are too wide to consider two vessels as close (e.g. enough to perform an exchange of goods). At last, areas not seen in training cannot be predicted as next location due to the limitations of TransE. Know-Evolve somehow managed to find some connections between vessels and areas but the results are very unsatisfactory: a Mean Rank of 400 means that the correct area is on average ranked 400th, against $MR = 20$ on ICEWS.¹⁷ Despite the difficulty induced by the discretization, position prediction is a simple task and the models performed poorly: they are not adequated to address this problem. The use of positions as continuous attributes could solve the abovementionned issues and improve the results on position prediction with knowledge graphs.

5.3.3 Experimental settings

We used the settings reported in the associated article¹⁷ to run Know-Evolve. For TransE, we set batch size=200, learning rate = 0.001 and embedding dimension = 64.

6. CONCLUSION AND FUTURE WORK

In this article, we reviewed two link prediction techniques for a task: the evolution of the positions of vessels using a dynamic knowledge graph for Maritime Situational Awareness. We showed that relational data (KG^R) is not sufficient to modelize the movement of a vessel and that attributional information should be used (KG^A). We also exhibited the challenges that need to be overcome to apply DAKGs on MSA, and formalized the relation and attribute value prediction problem.

We foresee several tasks for future work: (1) make the prediction task more realistic by adding more entity and relation types in the dataset, such as ships going in and out of ports, or encounters between ships, (2) find a model that can handle both KG^R and KG^A for link and attribute prediction in a temporal setting, (3) perform threat and/or anomaly detection on DAKGs. These are the three requirements to fully evaluate the use of DAKGs on operational maritime data.

7. ACKNOWLEDGEMENTS

This work was partially supported by the French National Association for Research and Technology and by Airbus Defence and Space.

REFERENCES

- [1] NATO, "Mc msa draft definition," (2007).
- [2] Le Guillarme, N. and Lerouveau, X., "Unsupervised Extraction of Knowledge from S-AIS Data for Maritime Situational Awareness," *FUSION*, 8 (2013).
- [3] Riveiro, M., Pallotta, G., and Vespe, M., "Maritime anomaly detection: A review," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **8**, 1266 (Sept. 2018).
- [4] Mascaro, S., Nicholso, A. E., and Korb, K. B., "Anomaly detection in vessel tracks using Bayesian networks," *International Journal of Approximate Reasoning* **55**, 84–98 (Jan. 2014).

- [5] Riveiro, M., Johansson, F., Falkman, G., and Ziemka, T., "Supporting maritime situation awareness using self-organizing maps and gaussian mixture models," *Frontiers in Artificial Intelligence and Applications* **173**, 84 (2008).
- [6] Tu, E., Zhang, G., Rachmawati, L., Rajabally, E., and Huang, G.-B., "Exploiting AIS Data for Intelligent Maritime Navigation: A Comprehensive Survey," *arXiv:1606.00981 [cs]* (June 2016). arXiv: 1606.00981.
- [7] Zissis, D., Xidias, E. K., and Lekkas, D., "Real-time vessel behavior prediction," *Evolving Systems* **7**, 29–40 (Mar. 2016).
- [8] Perera, L. P. and Soares, C. G., "Ocean Vessel Trajectory Estimation and Prediction Based on Extended Kalman Filter," *2nd International Conference on Adaptive and Self-adaptive Systems and Applications* , 7 (2010).
- [9] Bordes, A., Usunier, N., Garcia-Duran, A., Weston, J., and Yakhnenko, O., "Translating Embeddings for Modeling Multi-relational Data," 9 (2013).
- [10] Nickel, M., Tresp, V., and Kriegel, H.-P., "A Three-Way Model for Collective Learning on Multi-Relational Data," 8 (2011).
- [11] Nickel, M., Rosasco, L., and Poggio, T., "Holographic Embeddings of Knowledge Graphs," *AAAI* , 7 (2016).
- [12] Liu, H., Wu, Y., and Yang, Y., "Analogical Inference for Multi-relational Embeddings," *ICML* , 11 (2017).
- [13] Socher, R., Chen, D., Manning, C. D., and Ng, A., "Reasoning With Neural Tensor Networks for Knowledge Base Completion," *NIPS* , 10 (2013).
- [14] Kipf, T. N. and Welling, M., "Variational Graph Auto-Encoders," *arXiv:1611.07308 [cs, stat]* (Nov. 2016). arXiv: 1611.07308.
- [15] Leblay, J. and Chekol, M. W., "Deriving Validity Time in Knowledge Graph," in [*Companion of the The Web Conference 2018 on The Web Conference 2018 - WWW '18*], 1771–1776, ACM Press, Lyon, France (2018).
- [16] Esteban, C., Tresp, V., Yang, Y., Baier, S., and Krompaß, D., "Predicting the Co-Evolution of Event and Knowledge Graphs," 8 (2016).
- [17] Trivedi, R., Dai, H., Wang, Y., and Song, L., "Know-Evolve: Deep Temporal Reasoning for Dynamic Knowledge Graphs," *ICML* (May 2017).
- [18] Jiang, T., Liu, T., Ge, T., Sha, L., Chang, B., Li, S., and Sui, Z., "Towards Time-Aware Knowledge Graph Completion," *International Conference on Computational Linguistics* , 10 (2016).
- [19] Sankar, A., Wu, Y., Gou, L., Zhang, W., and Yang, H., "Dynamic Graph Representation Learning via Self-Attention Networks," *arXiv:1812.09430 [cs, stat]* (Dec. 2018). arXiv: 1812.09430.
- [20] Lin, Y., Liu, Z., and Sun, M., "Knowledge Representation Learning with Entities, Attributes and Relations," *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI)* , 7 (2016).
- [21] Tay, Y., Tuan, L. A., Phan, M. C., and Hui, S. C., "Multi-Task Neural Network for Non-discrete Attribute Prediction in Knowledge Graphs," 1029–1038, ACM Press (2017).
- [22] Li, J., Cheng, K., Wu, L., and Liu, H., "Streaming Link Prediction on Dynamic Attributed Networks," in [*Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining - WSDM '18*], 369–377, ACM Press, Marina Del Rey, CA, USA (2018).
- [23] Hoffart, J., Suchanek, F. M., Berberich, K., and Weikum, G., "YAGO2: A spatially and temporally enhanced knowledge base from Wikipedia," *Artificial Intelligence* **194**, 28–61 (Jan. 2013).
- [24] Boschee, E., Lautenschlager, J., O'Brien, S., Shell-man, S., Starz, J., and Ward, M., "Icews coded event data," (2015).
- [25] Leetaru, K. and Schrodt, P. A., "GDELT: Global Data on Events, Location and Tone," *ISA Annual Convention* , 51 (2013).
- [26] Roy, J., "Anomaly detection in the maritime domain," *Proceedings of SPIE - The International Society for Optical Engineering* (2008).
- [27] Zhao, P., Aggarwal, C., and He, G., "Link prediction in graph streams," in [*2016 IEEE 32nd International Conference on Data Engineering (ICDE)*], 553–564, IEEE, Helsinki, Finland (May 2016).
- [28] Rezende, D. J., Mohamed, S., and Wierstra, D., "Stochastic Backpropagation and Approximate Inference in Deep Generative Models," in [*International Conference on Machine Learning*], 1278–1286 (Jan. 2014).
- [29] Chekol, M. W. and Pirro, G., "Marrying Uncertainty and Time in Knowledge Graphs," *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence* , 7 (2017).
- [30] Adadi, A. and Berrada, M., "Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)," *IEEE Access* **6**, 52138–52160 (2018).