Orchestrating a brighter world NEC

Leipzig Data Week

Learning from Temporal Knowledge Graphs

Julia Gastinger, AI Innovation (AIN) Group NEC Laboratories Europe
and University of Mannheim Data and Web Science Group
July 5th, 2022

Content

NEC Laboratories Europe

Advancing Information & Communication

— Through Research Excellence —

KEY R&D METRICS

1000+ Patents

50+ Peer-reviewed publications per year

120+ European Projects

OPERATIONAL AREAS

Research
Leading scientific discovery in Europe

Technology Transfer
Commercializing R&D results in existing and new company business segments

Standards
Defining European technology standards and best practices

RESEARCH AREAS

Data Science
- Multi-modal and relational data (graph)
- Natural language processing
- Continual learning
- Biomedical AI for infectious diseases and immuno-oncology

Security
- Blockchain Privacy, Security and Scalability
- System and Device Security

5G Networks
- Virtual Radio Networks optimized with AI
- Vertical solution enablers
- Network resource optimization

IoT and AI Platforms
- High-performance AI platform for scalable neural networks
- Edge-cloud programming models for IoT
- Distributed AI

Standards
- ETSI
- 3GPP
- IEEE
- IETF
- FIWARE

Overview: Knowledge Graphs (KG)

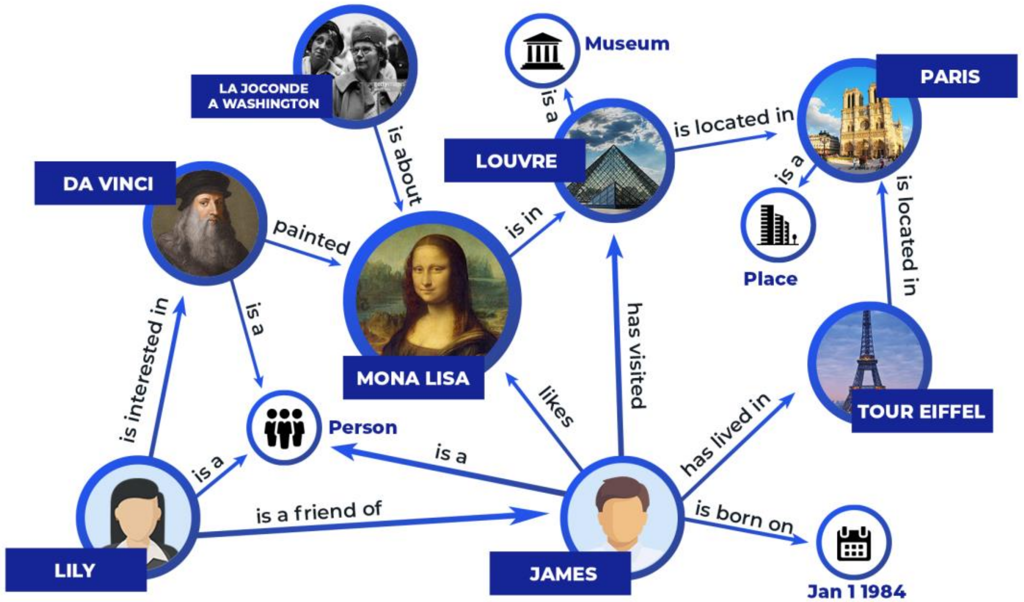

- Consists of:
  - A set of entities (e.g., Lily, Da Vinci) and,
  - A set of relation types or predicates (e.g., is a friend of, has visited).
- One can represent a KG as a set of triples of the form (subject, predicate, object), denoted as (s; p; o).
  - e.g., (Lily, is_interested_in, DaVinci)

Image from: [5]: Atlassin Community, https://community.atlassian.com/t5/Confluence-questions/Knowledge-graph/qaq-p/1565284. Accessed May 19th, 2022.

Overview: Knowledge Graphs (KG)

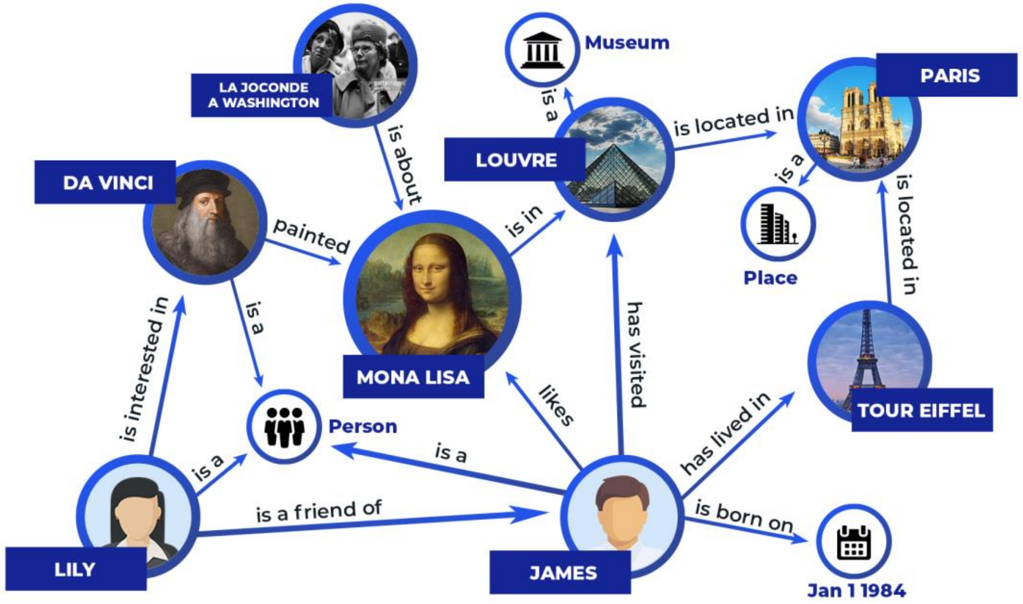

- Typical Tasks:
  - Link prediction: (Lily, is a friend of, ?); (Lily, ?, James)
    - seeks the most probable completion of a triple (s; p; ?) or (?; p; o)
  - Node classification (which class does Paris belong to?)
  - Graph classification (which class does the whole graph belong to?)

➡ Infer hidden links, or other information from the Knowledge Graph structure

Image from: [5]: Atlassin Community, https://community.atlassian.com/t5/Confluence-questions/Knowledge-graph/qaq-p/1565284. Accessed May 19th, 2022.

Overview: Knowledge Graphs (KG)

Knowledge Graphs Change over time

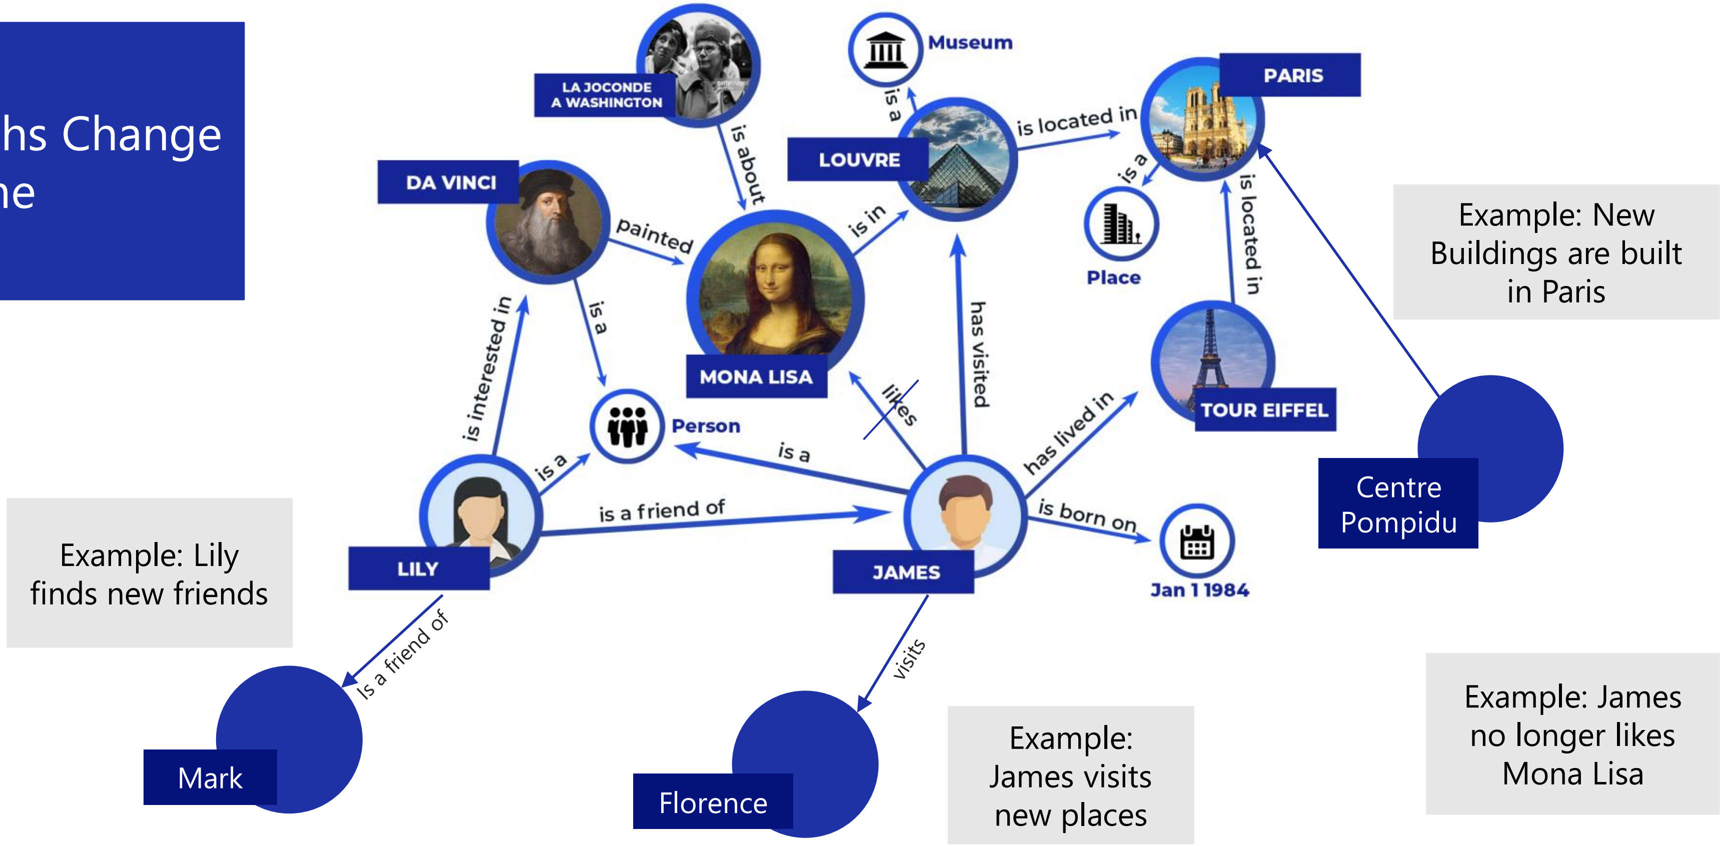

Image from: [5]: Atlassin Community, https://community.atlassian.com/t5/Confluence-questions/Knowledge-graph/qaq-p/1565284. Accessed May 19th, 2022.

Overview: Knowledge Graphs (KG)

Knowledge Graphs Change over time

We want to incorporate this Evolvement over time and learn from it

Temporal Knowledge Graphs
= Knowledge Graphs, where facts occur, recur or evolve over time, and each edge in the graphs has temporal information associated with it [2]

[2] Trivedi, Rakshit, et al. "Know-evolve: Deep temporal reasoning for dynamic knowledge graphs." *international conference on machine learning*. PMLR, 2017.

Overview: Temporal Knowledge Graphs – Definition

Temporal Knowledge Graphs
= Knowledge Graphs, where facts occur, recur or evolve over time, and each edge in the graphs has temporal information associated with it [2]

[2] Trivedi, Rakshit, et al. "Know-evolve: Deep temporal reasoning for dynamic knowledge graphs." *international conference on machine learning*. PMLR, 2017.

Overview: Temporal Knowledge Graphs – Time Representation

Time in nodes or Edges

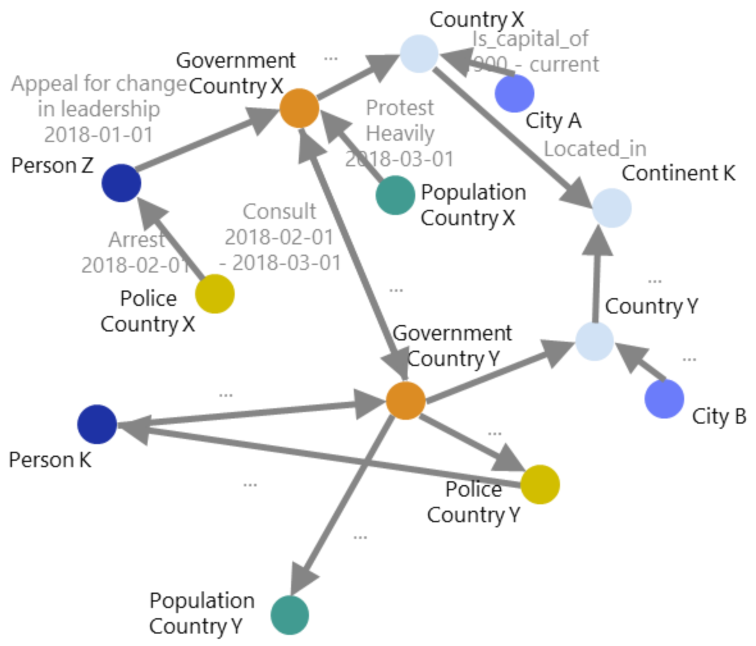

One Graph per Time Step

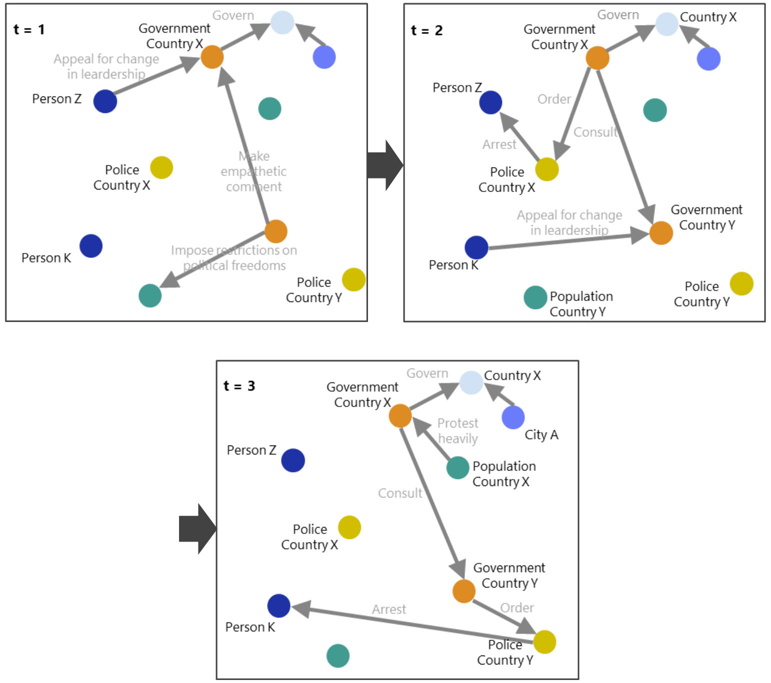

(Time in External Model)

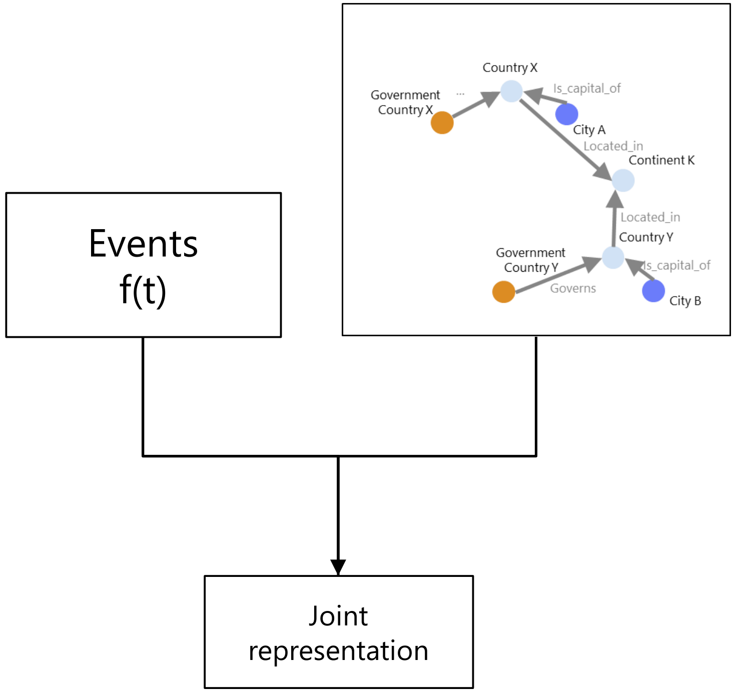

Extend Triples -> Quadruples:

[Police Country X, Arrest, Person Z, **2018-01-01**];
[Population Country X, Protest Heavily, Government Country X, **2018-02-01**]

Triples adapted from Integrated Crisis Early Warning System (ICEWS18) dataset [3]; Event data consists of coded interactions between socio-political actors (i.e., cooperative or hostile actions between individuals, groups, sectors and nation states). This Dataset consists of events from 1/1/2018 to 10/31/2018 (24 hours time granularity).

[3] Boschee, Elizabeth; Lautenschlager, Jennifer; O'Brien, Sean; Shellman, Steve; Starz, James; Ward, Michael, 2015, "ICEWS Coded Event Data", https://doi.org/10.7910/DVN/28075, Harvard Dataverse, V30.

Overview: Temporal Knowledge Graphs – Motivation

Why do we need **relations & graph structure**?

- Model relations in, e.g., a group structure – people influence each other
- Model different types of relations, e.g., is_friend_of, and is_president_of
- Ability to rely not only on simple patterns but can exploit high-order features which are captured by the graph structure

Why do we need **time aspect** in graph (instead of “traditional” link prediction)?

- Model time dependency of changes in graph

Examples:

- Time-dependency in actions (things that happened more often recently: more likely to happen again)
- Temporal fluctuations and seasonality in actions

Goal: Infer hidden links and other information by taking both aspects into account, i.e., present graph structure and past graph embeddings

Overview: Temporal Knowledge Graphs – Research Topics

- Temporal Knowledge Base Completion
  - Link prediction for different time steps
  - Time prediction (s,r,o,?)
- Graph/ Node Classification
- Question-Answering
- Extrapolation to Future Knowledge Graphs: Temporal Knowledge Graph Forecasting

rather new field
(2015 - 2022)

Temporal Knowledge Graph Forecasting

Explanation and Use Cases
State of the Art
Open Issues

Temporal Knowledge Graph (TKG) Forecasting – Explanation

Extrapolation to Future Knowledge Graphs: Temporal Knowledge Graph Forecasting

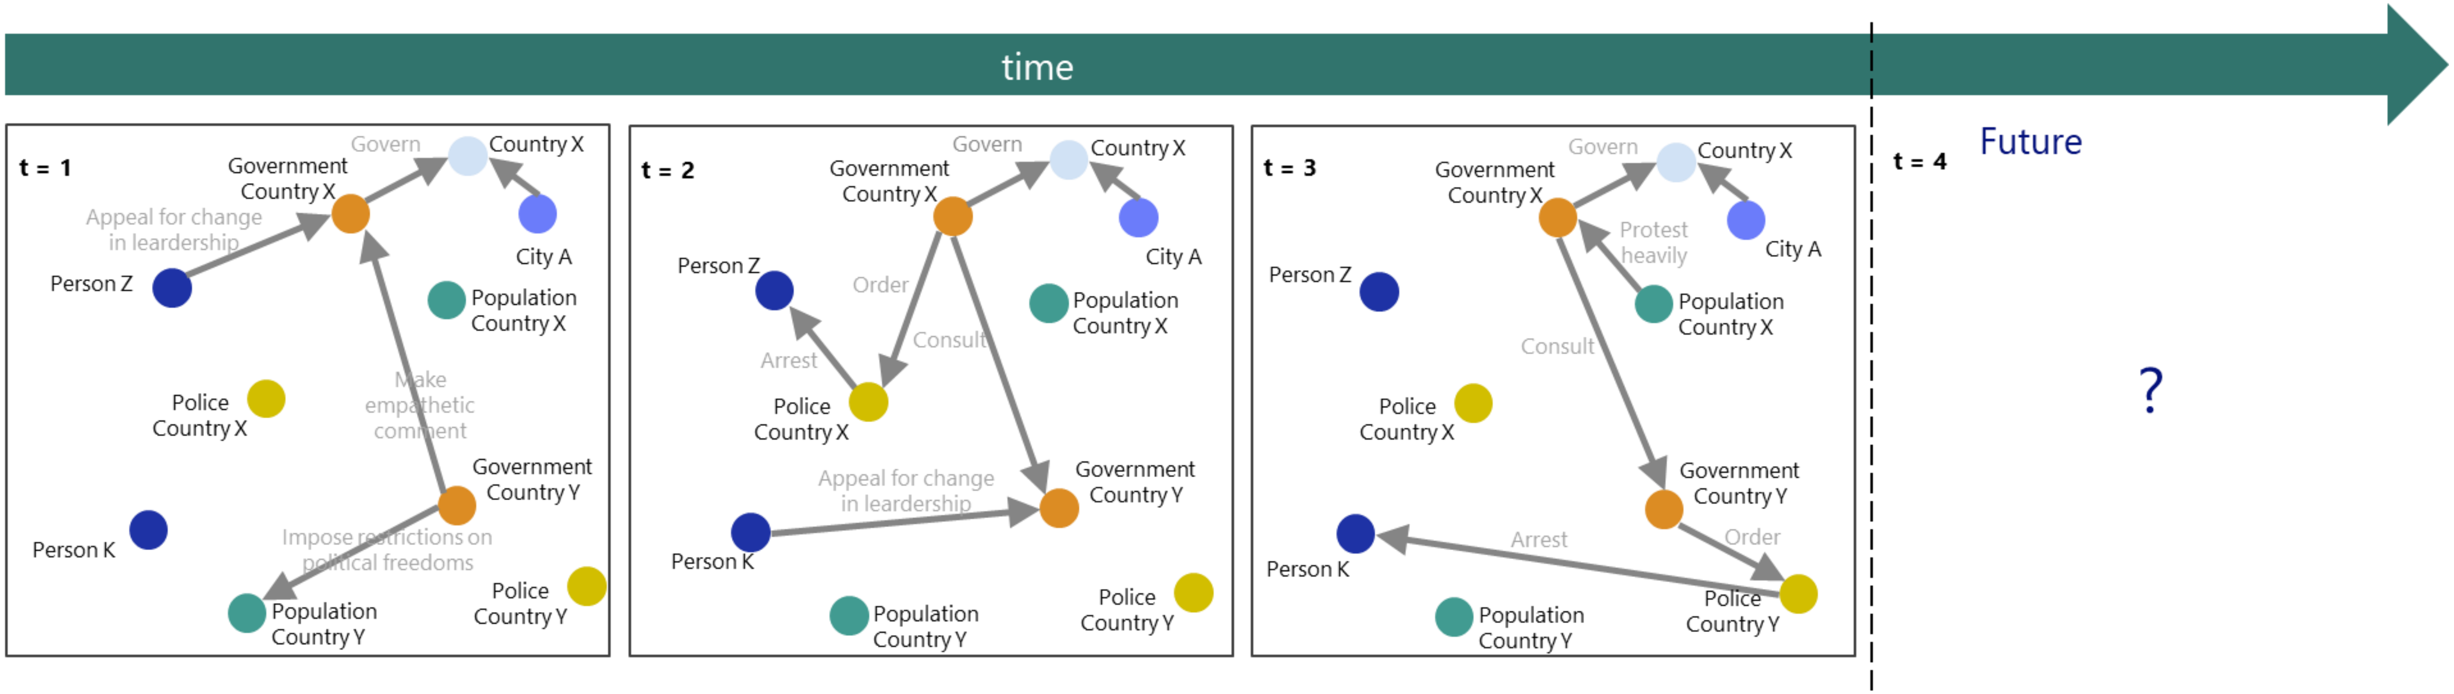

Triples in example adapted from Integrated Crisis Early Warning System (ICEWS18) dataset [3];
[3] Boschee, Elizabeth; Lautenschlager, Jennifer; O'Brien, Sean; Shellman, Steve; Starz, James; Ward, Michael, 2015, "ICEWS Coded Event Data", https://doi.org/10.7910/DVN/28075, Harvard Dataverse, V30.

TKG Forecasting – Use Cases (at NEC)

Crisis Management

Public Safety (Law Enforcement)

Green House Gas emission reporting

What we need:

- Incorporate Information about type of nodes, their attributes and types of links
- Predict future links and entities for a given KG.
- Detect changes in KGs over time (e.g., drift)

Triples in example adapted from Integrated Crisis Early Warning System (ICEWS18) dataset [3];

[3] Boschee, Elizabeth; Lautenschlager, Jennifer; O'Brien, Sean; Shellman, Steve; Starz, James; Ward, Michael, 2015, "ICEWS Coded Event Data", https://doi.org/10.7910/DVN/28075, Harvard Dataverse, V30.

TKG Forecasting – How we represent Time: Graph Snapshots

- Model the evolvement over time at discrete timesteps

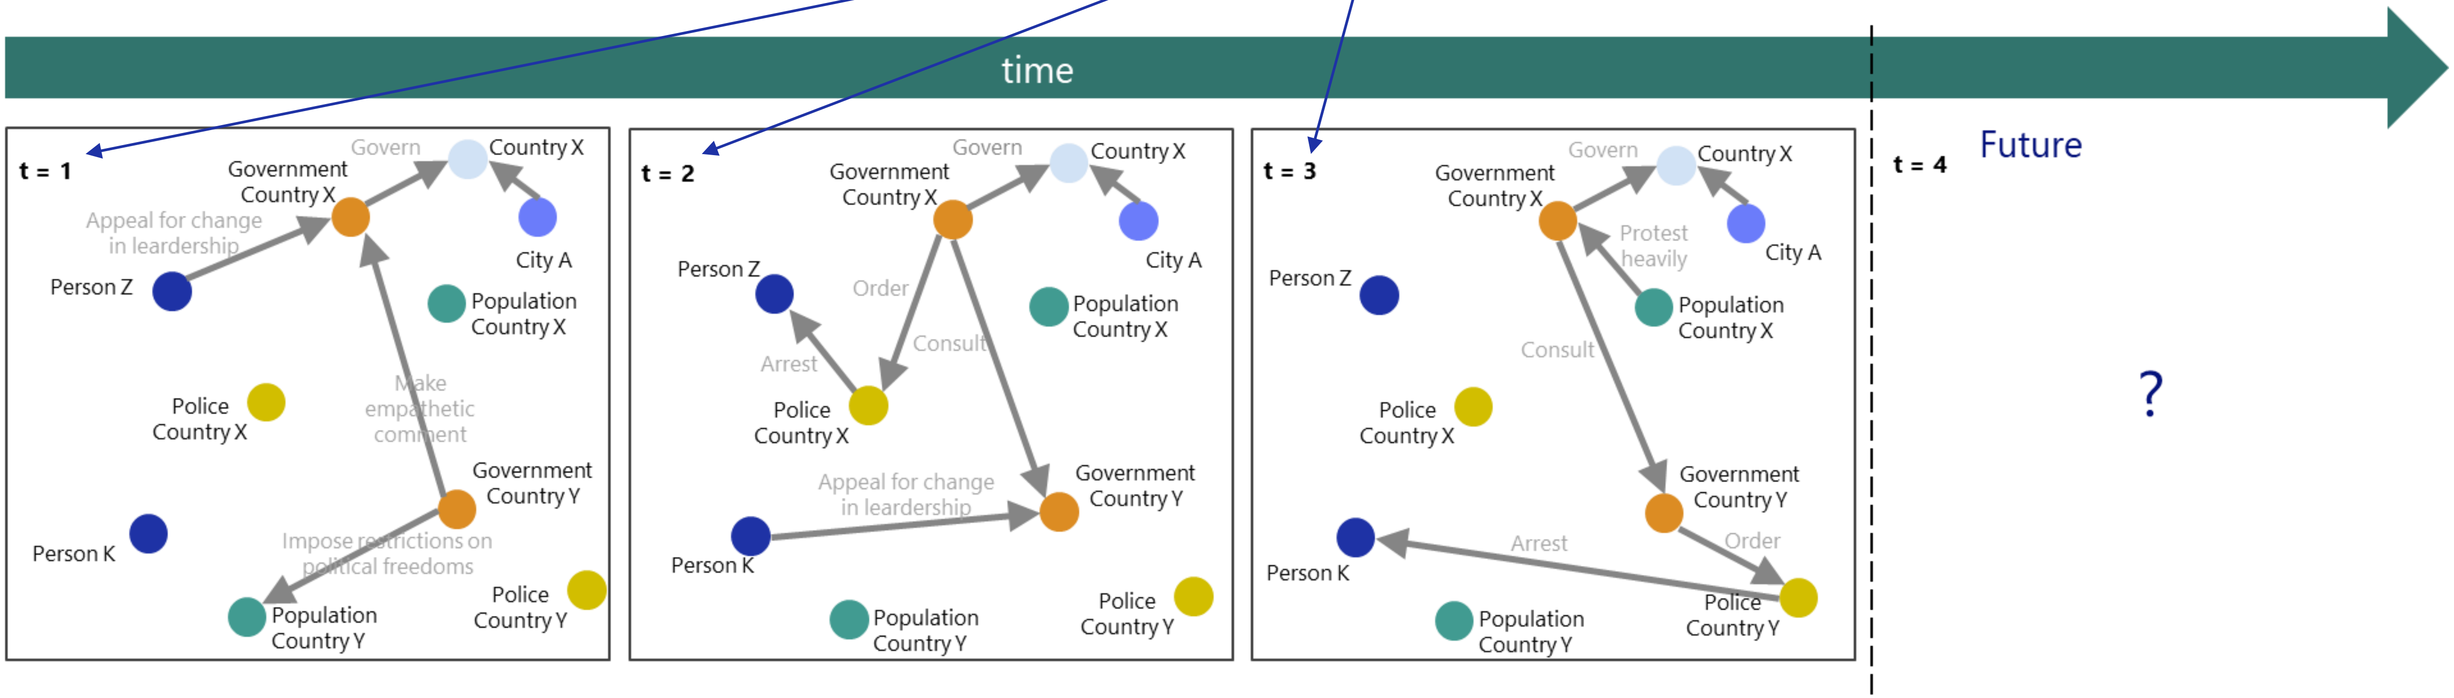

Triples in example adapted from Integrated Crisis Early Warning System (ICEWS18) dataset [3];

[3] Boschee, Elizabeth; Lautenschlager, Jennifer; O'Brien, Sean; Shellman, Steve; Starz, James; Ward, Michael, 2015, "ICEWS Coded Event Data", https://doi.org/10.7910/DVN/28075, Harvard Dataverse, V30.

TKG Forecasting – How we represent Time: Graph Snapshots

- What they can represent:
  - Events at discrete timesteps (Example)
  - Time Intervals (Example)
    - (by showing the respective triples at each timestep)
  - (Static) Facts (Example)
    - (by showing the respective triples at each timestep)

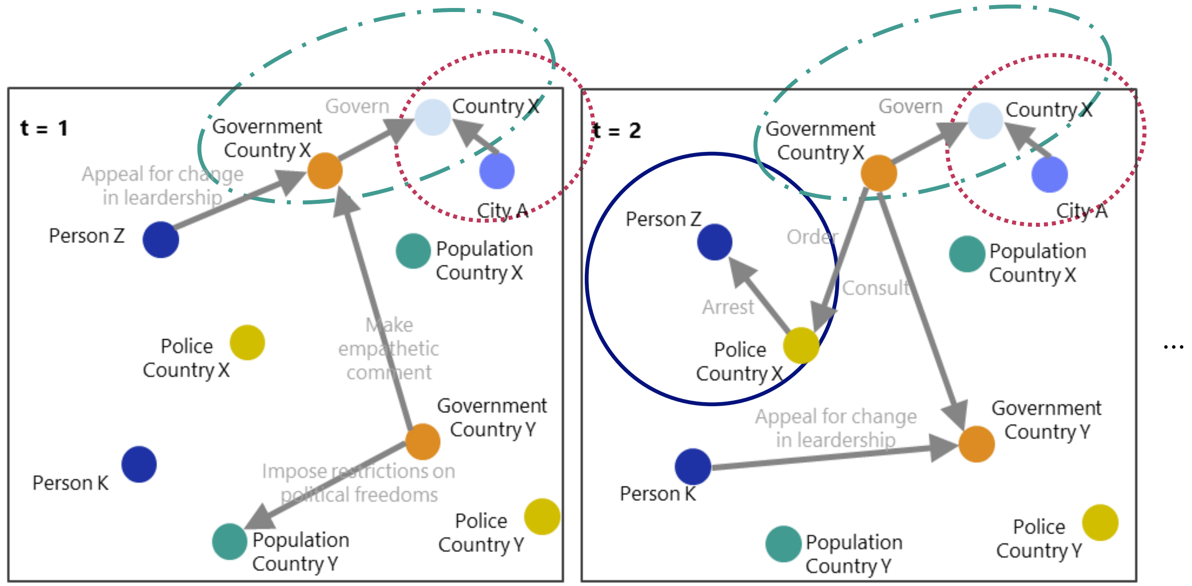

- What they cannot represent:
  - Continuous time
  - Anything more finegrained than the snapshot-stepsize
    - E.g., Obama was no longer president of the US at 2017-01-21, which we cannot exactly capture when using monthly granularity
    - The time granularity needs to be chosen accordingly
  - Time distance between two events and duration
    - E.g., "going to vacation for 3 weeks", or the time that is needed to construct a house
    - Only by (manually) computing the difference between timestep(event1) and timestep(event2)

Triples in example adapted from Integrated Crisis Early Warning System (ICEWS18) dataset [3];

[3] Boschee, Elizabeth; Lautenschlager, Jennifer; O'Brien, Sean; Shellman, Steve; Starz, James; Ward, Michael, 2015, "ICEWS Coded Event Data", https://doi.org/10.7910/DVN/28075, Harvard Dataverse, V30.

TKG Forecasting – State of the Art Summary

Boxes: code available

- Learning Neural Ordinary Equations for Forecasting Future Links on Temporal Knowledge Graphs, Han et al. (EMNLP 2021) - TANGO
- Learning from History: Modeling Temporal Knowledge Graphs with Sequential Copy-Generation Networks, Zhu et al. (AAAI 2021) - CyGNet
- Explainable Subgraph Reasoning for Forecasting on Temporal Knowledge Graphs, Han et al. (ICLR 2021) - xERte
- Search from History and Reason for Future: Two-stage Reasoning on Temporal Knowledge Graphs, Xi et al. (ACL 2021) - CluSTeR
- Temporal Knowledge Graph Reasoning Based on Evolutional Representation Learning, Xi et al. (SIGIR 2021) – RE-GCN
- Recurrent Event Network: Autoregressive Structure Inference over Temporal Knowledge Graphs – Jin et al. (EMNLP 2020) – RENet
- TLogic: Temporal Logical Rules for Explainable Link Forecasting on Temporal Knowledge Graphs – Liu et al (accepted at AAAI 2022) - TLogic

Based on Neural Ordinary Differential Equations (NODE)

Combine predictions from copy mode (historical vocabulary) and generation mode (all objects)

Reasons over query-relevant subgraphs of TKG – node attention score, edge contribution score

Two-stage prediction: clue searching and temporal reasoning

Learn structural and temporal dependencies with different Units, include static graph constraint

Learn temporal dependency from sequence of graphs and local structural dependency from neighborhood

Learns temporal logical rules from TKGs based on temporal random walks

TKG Forecasting – State of the Art Summary

- All mentioned papers have in common that they
  - look into the future of Knowledge Graphs
  - represent time with graph snapshots (one graph per time step)
  - focus on (more or less) the same datasets*

- They differ in their
  - focus (e.g., explainability, speed)
  - approaches/methods (e.g., neural ordinary differential equations, rule-based, gcn-based, using contraints)

* ICEWS14, ICEWS18, GDELT, WIKI, YAGO

TKG Forecasting – State of the Art Open Issues (1/3)

1. Discrepancies in Evaluation

- ICEWS14 dataset: three different versions exist, causing confusion
- Filter-setting: papers report their result on different filter settings for computing the ranks (evaluation metrics)
- Single-Step vs. Multi-step Prediction setting
  - two settings lead to methods being not really comparable
  - Still, methods are compared (unfair)

TKG Forecasting – State of the Art Open Issues (2/3)

1. Discrepancies in Evaluation
2. **Predictions are not yet very good***
   - Example [4]:

Model	ICE18				ICE14				ICE05-15			
	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10	MRR	H@1	H@3	H@10
DistMult	13.86	5.61	15.22	31.26	20.32	6.13	27.59	46.61	19.91	5.63	27.22	47.33
ComplEx	15.45	8.04	17.19	30.73	22.61	9.88	28.93	47.57	20.26	6.66	26.43	47.31
R-GCN	15.05	8.13	16.49	29.00	28.03	19.42	31.95	44.83	27.13	18.83	30.41	43.16
ConvE	22.81	13.63	25.83	41.43	30.30	21.30	34.42	47.89	31.40	21.56	35.70	50.96
ConvTransE	23.22	14.26	26.13	41.34	31.50	22.46	34.98	50.03	30.28	20.79	33.80	49.95
RotatE	14.53	6.47	15.78	31.86	25.71	16.41	29.01	45.16	19.01	10.42	21.35	36.92
HyTE	7.41	3.10	7.33	16.01	16.78	2.13	24.84	43.94	16.05	6.53	20.20	34.72
TTransE	8.44	1.85	8.95	22.38	12.86	3.14	15.72	33.65	16.53	5.51	20.77	39.26
TA-DistMult	16.42	8.60	18.13	32.51	26.22	16.83	29.72	45.23	27.51	17.57	31.46	47.32
RGCRN	23.46	14.24	26.62	41.96	33.31	24.08	36.55	51.54	35.93	26.23	40.02	54.63
CyGNet	24.98	15.54	28.58	43.54	34.68	25.35	38.88	53.16	35.46	25.44	40.20	54.47
RE-NET	26.17	16.43	29.89	44.37	35.77	25.99	40.10	54.87	36.86	26.24	41.85	57.60
$\boxed{\text{RE-GCN}}$	**27.51**	**17.82**	**31.17**	**46.55**	**37.78**	**27.17**	**42.50**	**58.84**	**38.27**	**27.43**	**43.06**	**59.93**

[4] Zixuan Li, Xiaolong Jin, Wei Li, Saiping Guan, Jiafeng Guo, Huawei Shen, Yuanzhuo Wang, Xueqi Cheng: Temporal Knowledge Graph Reasoning Based on Evolutional Representation Learning. SIGIR 2021: 408-417

* I think the scores like MRR can be improved

TKG Forecasting – State of the Art Open Issues (3/3)

1. Discrepancies in Evaluation

2. Predictions are not yet very good

3. **The Methods are not yet applied to any real world Use Cases**

Our Research on Temporal Knowledge Graph Forecasting

Research Questions

Current Tasks

Summary

Our Research – Research Questions

Research Questions:

- How can we improve Temporal Knowledge Graph Forecasting?
- How can we evaluate the methods in this rather new field of research?
- What insights can we get by analyzing the Evolution of Knowledge Graphs over time?

Generate real-world impact :

- How can we apply the developed method(s) and analysis insights to a Real-World Use Case?

Our Research – Current Work

1. Evaluation of existing methods under a joint evaluation protocol
   - Description of settings and evaluation protocol
   - Evaluation of all existing papers on that evaluation protocol
   - Using know research datasets
     - Among them the Integrated Crisis Early Warning Systems (ICEWS) datasets for 2014, 2005-2015, 2018

2. Working on a new method for TKG forecasting
   - Based on Evolving Knowledge Graph Embeddings
   - Taking into account insights from literature research and open issues

Summary: Temporal Knowledge Graphs

Julia Gastinger
Julia.Gastinger@neclab.eu

Temporal Knowledge Graphs: Knowledge Graphs, where facts occur, recur or evolve over time, and each edge in the graphs has temporal information associated with it [2]

Time Representation: Graph Snapshots (1 Knowledge Graph per timestep)

Research Interest:

TKG Forecasting: Extrapolation to Future Knowledge Graphs, based on a history of Knowledge Graphs;

Analysis of Temporal Knowledge Graph Evolution over Time

Currently Working on:

- Evaluation of existing methods under a joint evaluation protocol
- A method for Temporal Knowledge Graph Extrapolations based on Evolving Knowledge Graph Embeddings

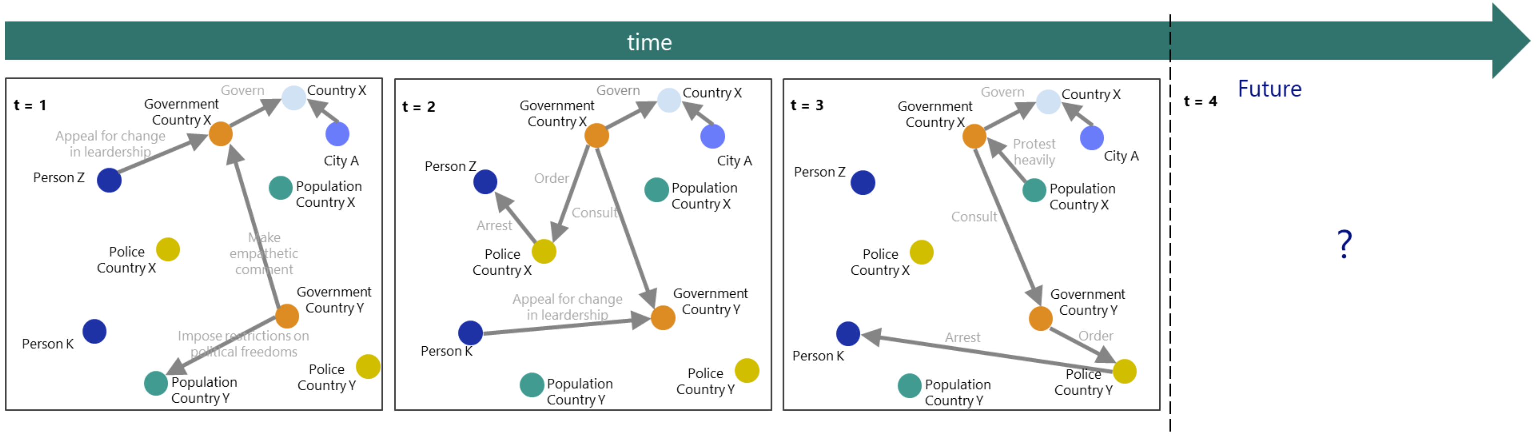

Student Projects Available:

https://jobs.neclab.eu

https://jobs.neclab.eu/jobs/openings/students/NEC-NLE-2204-101-AIN-1-Internship_or_Master_Thesis_on_Graph-based_ML_[2022-04-101-AIN].pdf

Triples in example adapted from Integrated Crisis Early Warning System (ICEWS18) dataset [3];

[3] Boschee, Elizabeth; Lautenschlager, Jennifer; O'Brien, Sean; Shellman, Steve; Starz, James; Ward, Michael, 2015, "ICEWS Coded Event Data", https://doi.org/10.7910/DVN/28075, Harvard Dataverse, V30.
[2] Trivedi, Rakshit, et al. "Know-evolve: Deep temporal reasoning for dynamic knowledge graphs." *international conference on machine learning*. PMLR, 2017.

