Embeddings for Knowledge Graphs and Multimodal Representations

Dr Nitisha Jain
King's College London

November 23, 2023

Knowledge Graphs

A knowledge graph (KG) is essentially a large-scale semantic network that contains entities, concepts as nodes and relationships among them as edges.

- Facts represented in form of triples
- Entities are labeled with attributes (e.g., types)
- Typed edges between two nodes capture a relationship between entities
- Usually based on an underlying schema

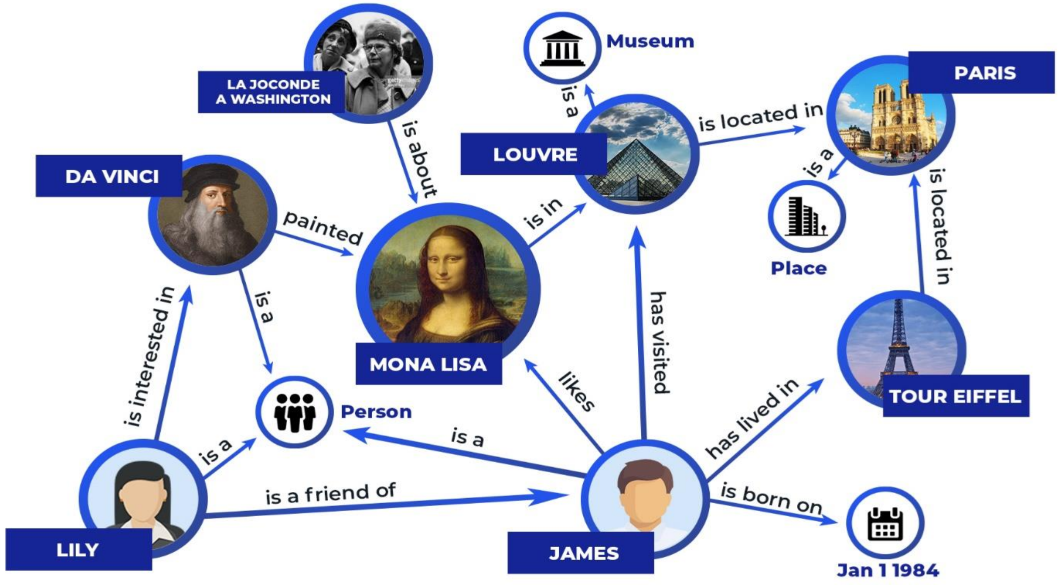

Knowledge Graphs

Examples - DBpedia, Yago, Google Knowledge Graph, Amazon Product Graph, IBM Watson..

Applications are widespread

- Question answering
- Chatbots
- Recommendation systems
- Web search
- ..

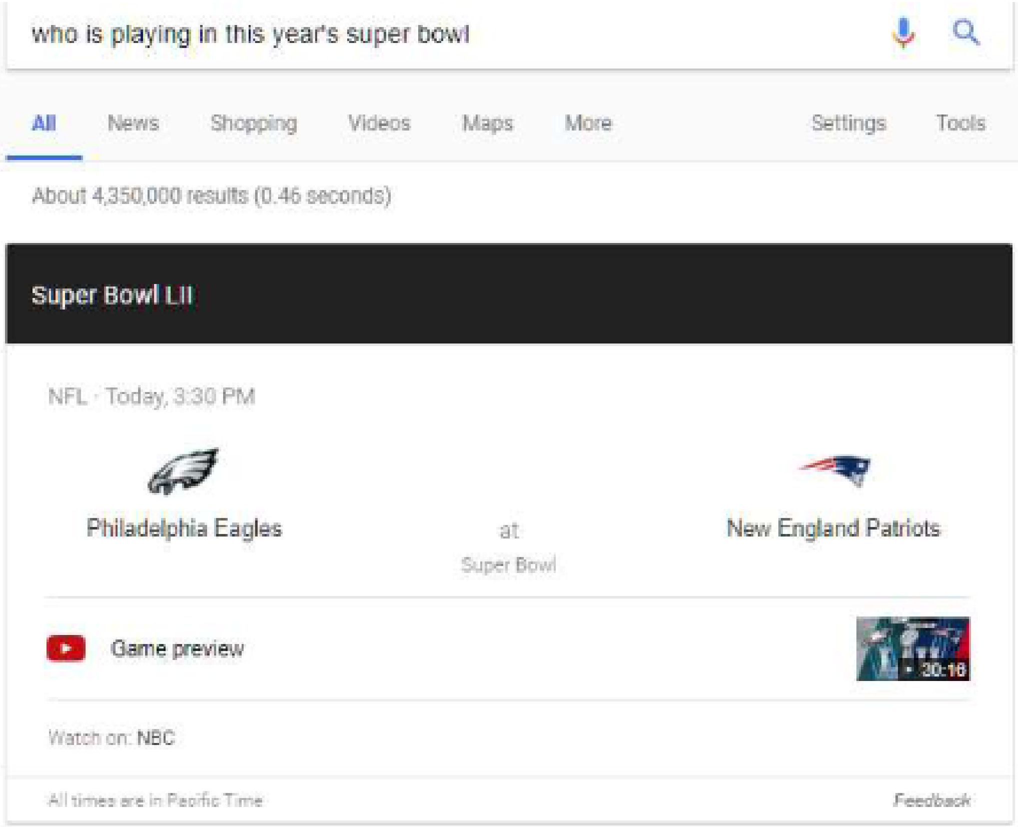

Representation Learning with Embeddings

- Key idea - to generate representations of nodes that depend on the structure of the graph, as well as feature information
- Representations for machine, not human
- Shape of embeddings – vector with floats as an element

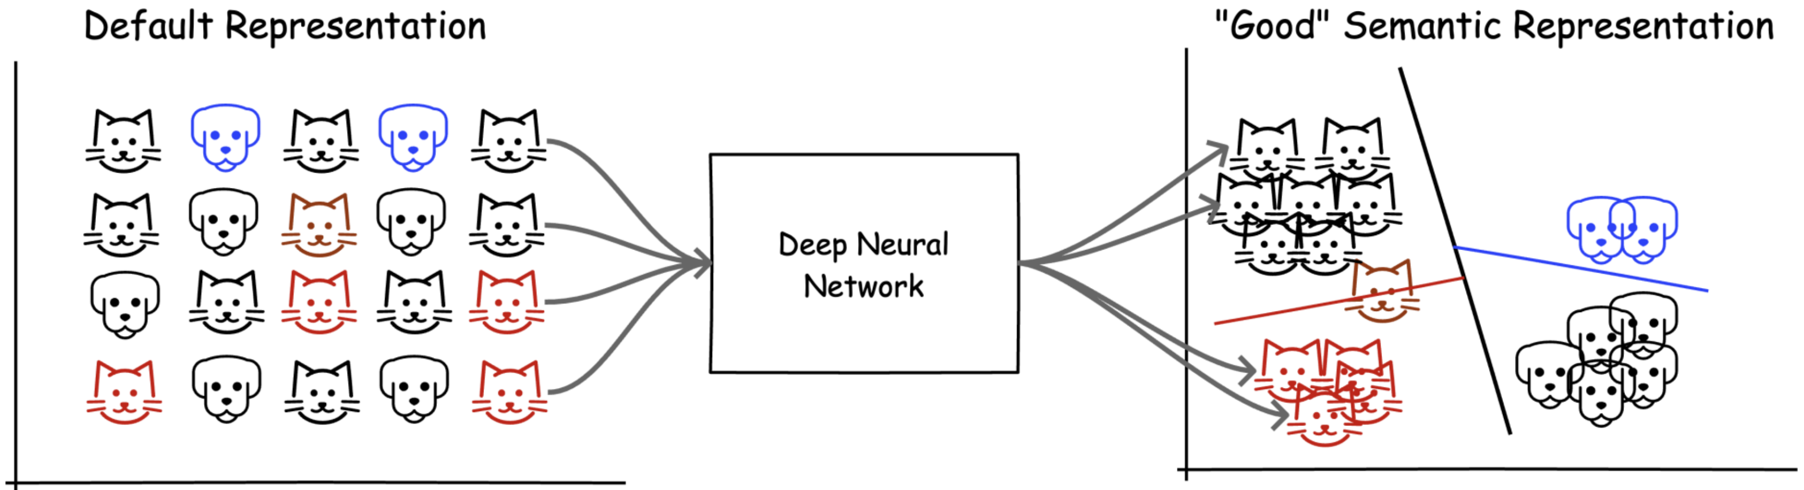

Cat by Martin LEBRETON, Dog by Serhii Smirnov from the Noun Project

Representation Learning with Embeddings

- Embed components of KG (entities, relations) into continuous vector spaces
- Allow easy manipulation of data while preserving inherent structure of KG
- Capture the interactions between entities of KG
- Used for link prediction towards KG completion
  - <v, r, ?> or <?, r, u>

KG triple <v, r, u >

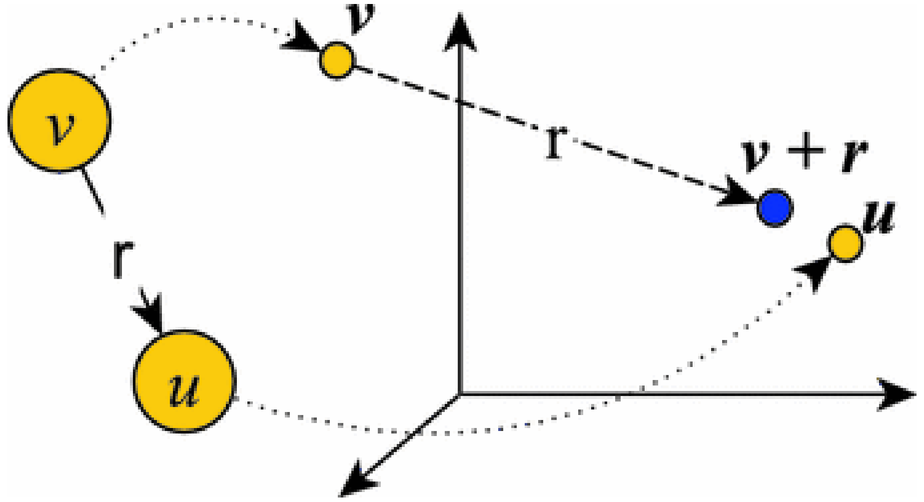

Translation based KG embedding

Popularity of KG embeddings

- Many embedding models
  - TransE
  - RESCAL
  - DistMult
  - ComplEx
  - ConvE
  - …
- Several new models being proposed every year ..

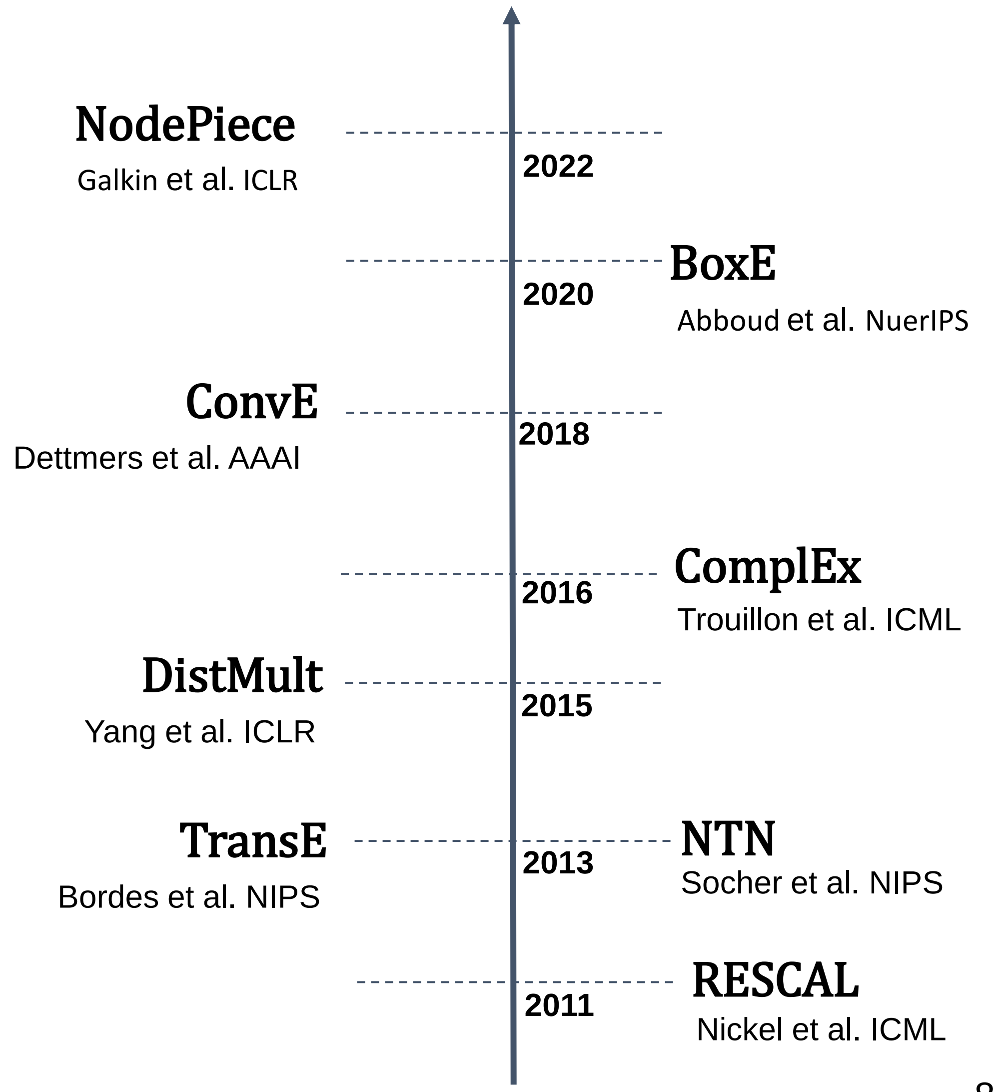

Applications are widespread..

- KG embeddings mainly used for link prediction, KG completion (Wang et al. TKDE 2017)

- KG embeddings are also being explored for various semantic tasks
  - Entity similarity (Sun et al. VLDB 2020)
  - Relation similarity (Kalo et al. ISWC 2019)
  - Conceptual clustering (Gad-Elrab et al. ISWC 2020)
  - Rule-based reasoning (Ho et al. 2018)

- All attempt to leverage semantic knowledge encoded in embeddings

Do embeddings represent semantics ?

Premise: Vectors of entities, relations should reflect their latent semantics

Similar entities **E1, E2, E3**

mapping ↓

Expected to be similar

KG Embedding vectors **e1**, **e2**, **e3**

Expectation vs. Reality

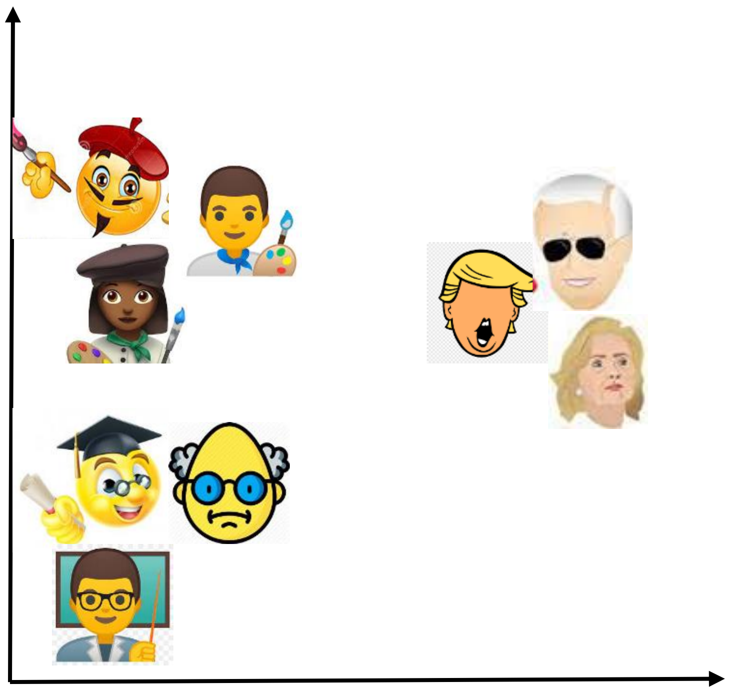

Entity similarity reflected by vectors

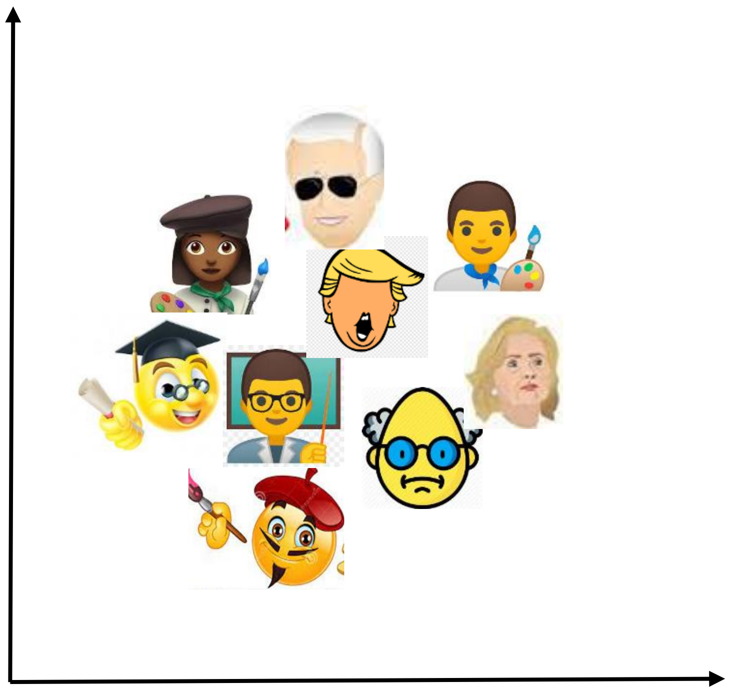

Fine-grained semantics not reflected*

**Nitisha Jain, Jan-Christoph Kalo, Wolf-Tilo Balke, Ralf Krestel: Do Embeddings Actually Capture Knowledge Graph Semantics? ESWC 2021.*

Objectives

- Comprehensive analysis of characteristics of the latent vectors from KG embeddings
- Quantitatively measure their ability for semantic representation
- Explore the reasons for shortcomings

How to measure semantic expressiveness?

- Can vectors express similarities between entities of same type?
  - Classification
  - Clustering

- Entity Classification
  - Assign entities to their types

- Entity Clustering
  - Identify entities belonging to the same type

- Suitable for semantic analysis: simple and can identify salient features of the embeddings, if any

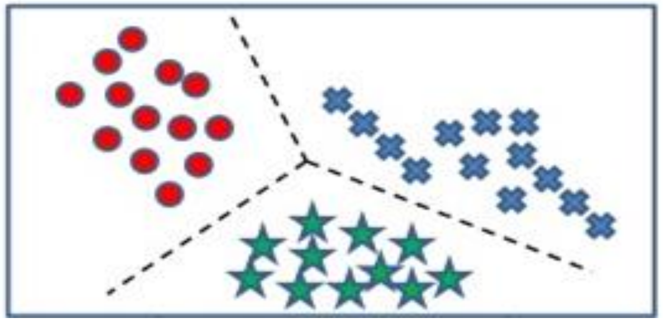

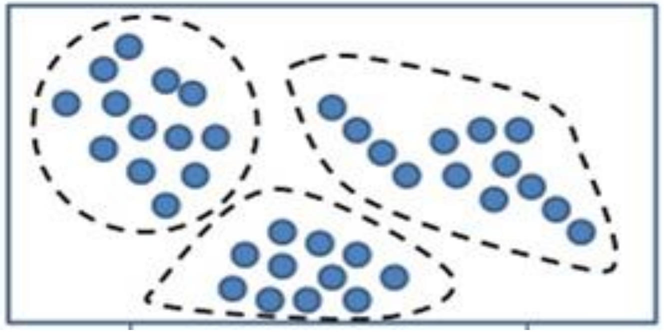

Dataset Preparation Yago3

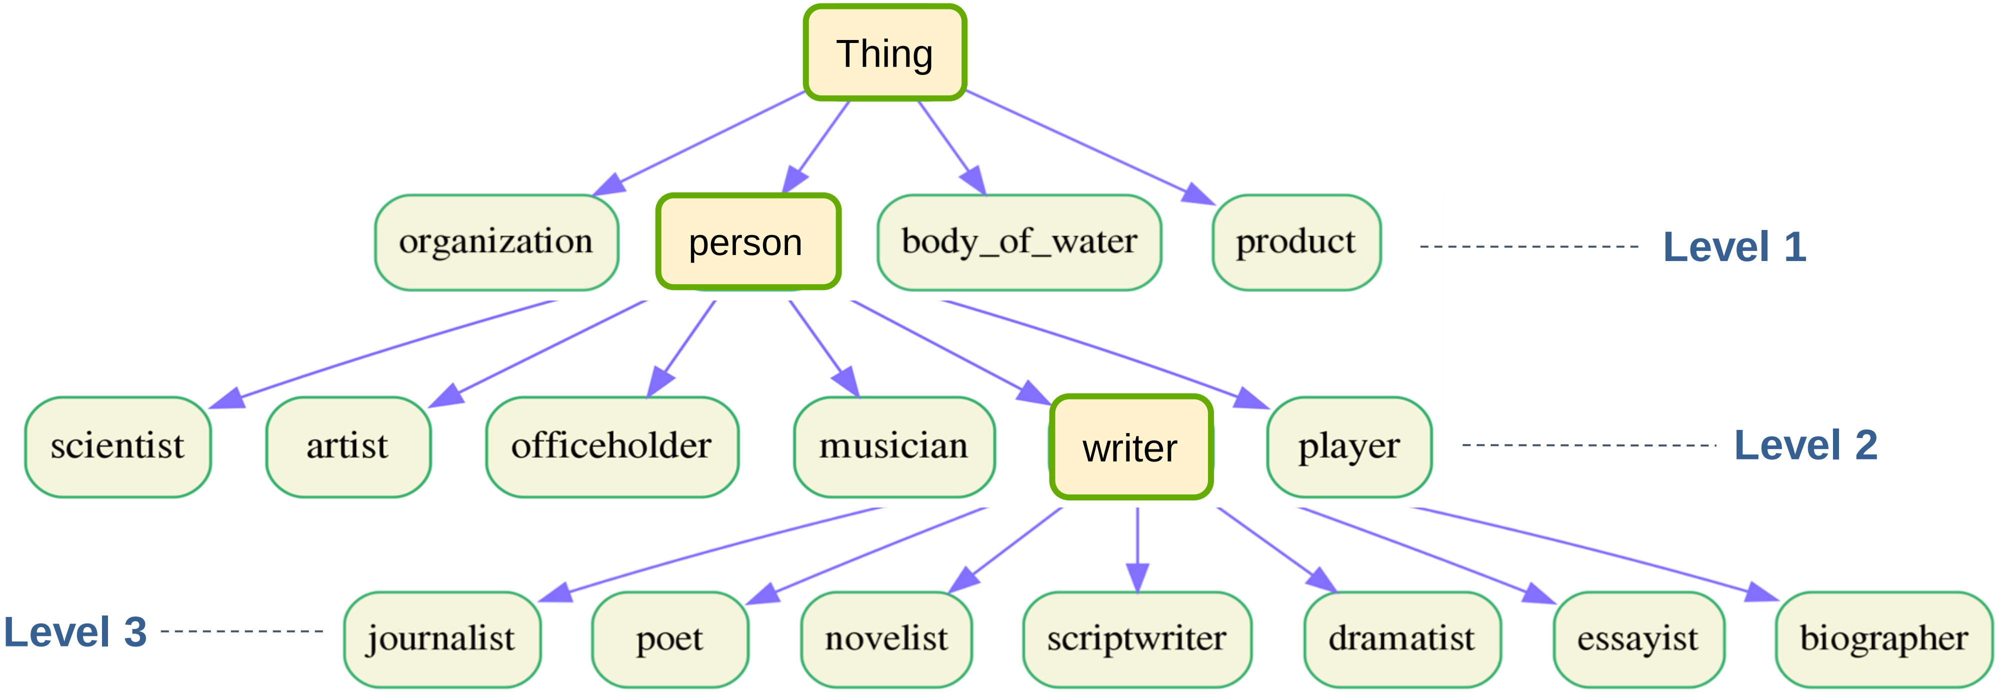

- *Goal*: Extract entities belonging to classes at different levels of ontology
- Sub-trees explored manually, most frequent classes chosen
- Experiments consider entities from same level of granularity for fair comparison

Experiments

- *Embeddings*: TransE, RESCAL, ComplEx, DistMult, ConvE (LibKGE library), RDF2Vec
- *Non-Embedding Baseline*: Entity typing with SDtype [Paulheim et al. 2013]
  - Heuristics based technique relying on statistical distributions of the entity links
  - Robust to noisy facts, agnostic to existing type information
- *Classification*: KNN, Multilayer Perceptron, Random Forest

Classification Results

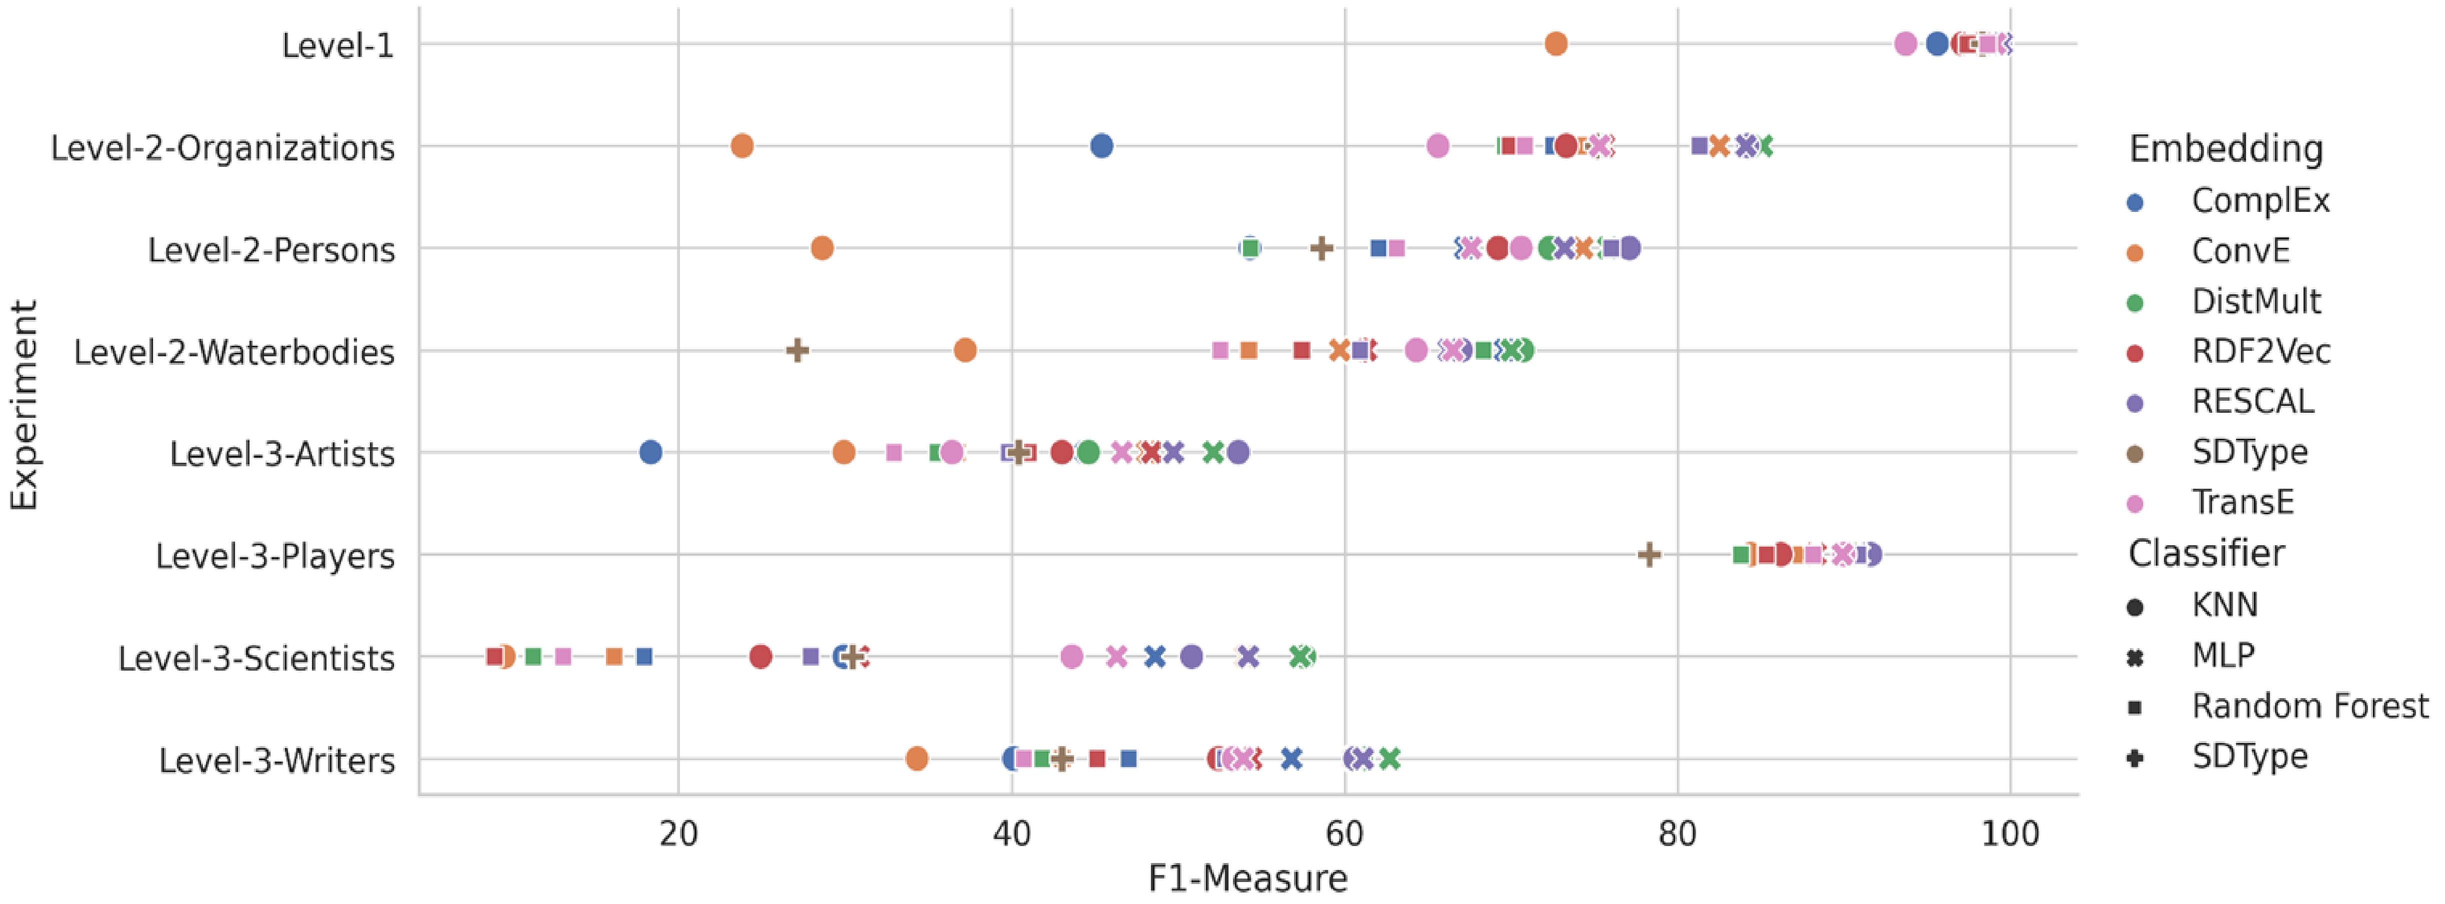

F1 measure for Yago Dataset

Why is it so ? Deepdive Analysis

- Semantic understanding depends on dataset
  - Real-world entities frequently belong to more than one semantic type or class

- Most relations are neither unique nor indicative of any single class in particular: *bornIn, hasSpouse, isCitizenOf ..*

Arnold Schwarzenegger (Q2685)

Austrian-American actor, businessman, bodybuilder and politician

Expectation vs. Reality

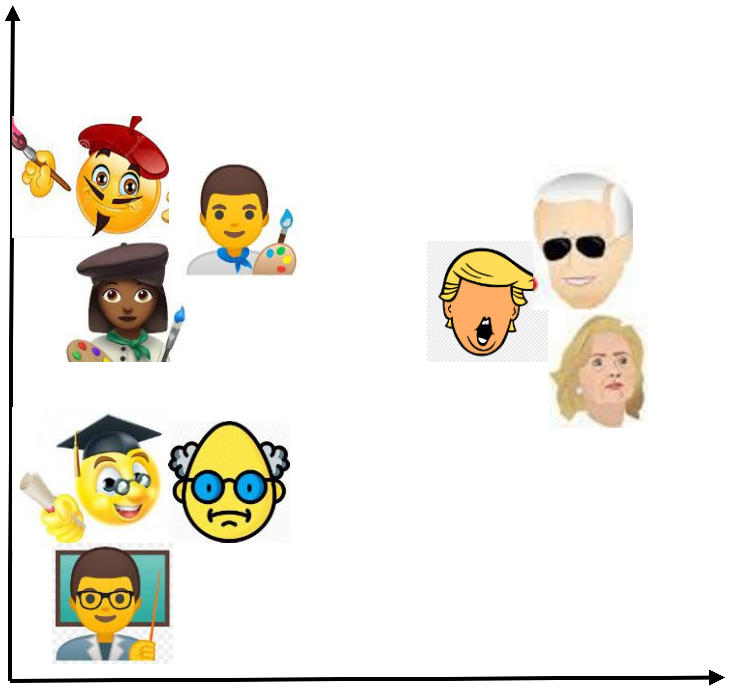

Entity similarity reflected by vectors

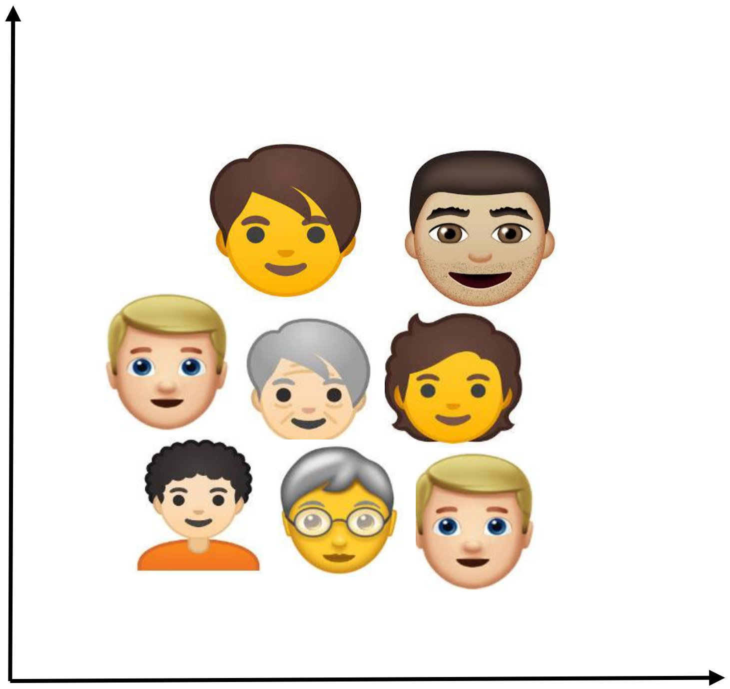

All people are similar !

Major Insights

Assumption: KG embeddings represent the semantics for KG components well

Findings:

- Semantic soundness is restricted to few entities, dependent on the dataset characteristics

- Simple heuristics-based approach can derive the semantics directly from KG triples without any additional information

So where do we go from here..

- Need for wider analysis, proper inspection of the advantages, weaknesses of KG embeddings
- Perform well for link prediction, but generalizability of these models for semantic tasks deserves more attention
- Lack of interpretability, lack of transparency in KG embeddings
- Ontology and reasoning can help with improving the semantics of embeddings

Improving KG embeddings with Ontology

- Lack of ontological knowledge during training, incorrect predictions.
- Proposing *ReasonKGE* to detect semantically incorrect predictions via ontological reasoning, generate targeted negative samples for the next iteration of training.
- Improved link prediction performance, as well as the ratio of semantically consistent predictions for any underlying embedding model.

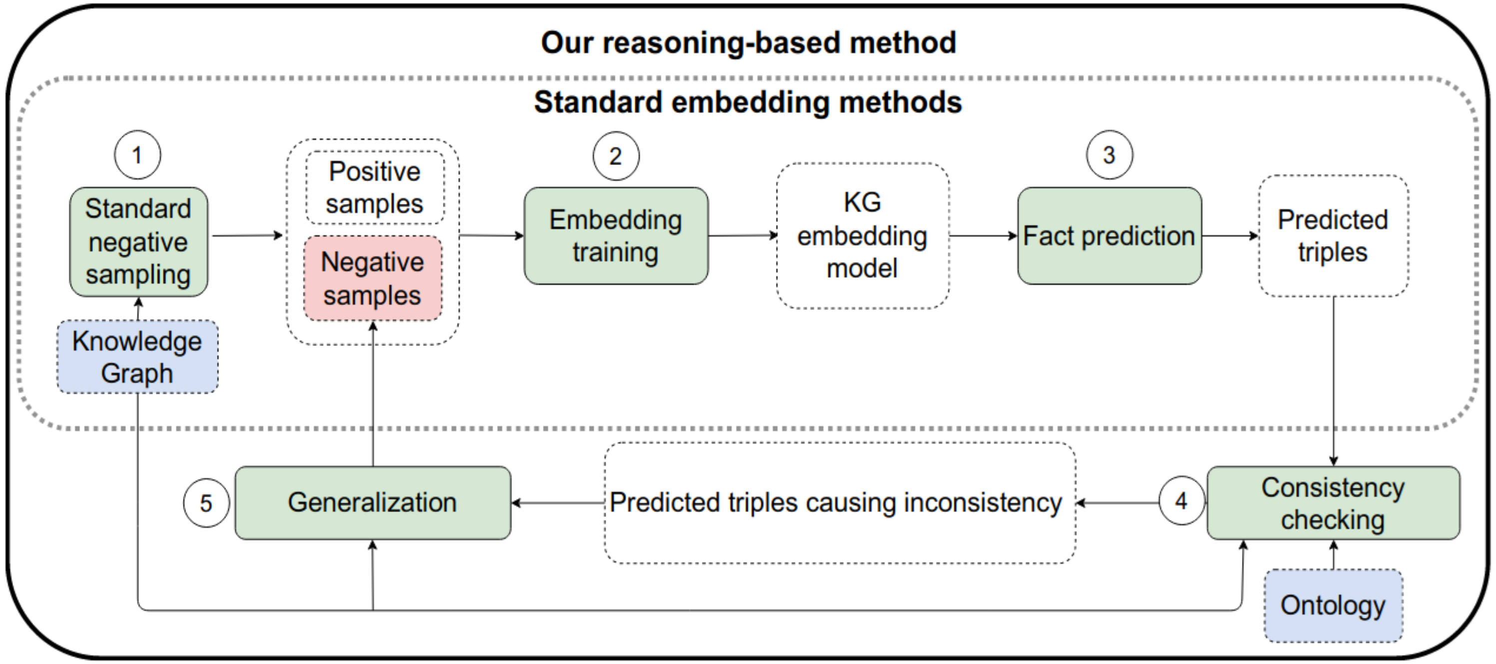

Nitisha Jain, Trung-Kien Tran, Mohamed H. Gad-Elrab, Daria Stepanova: Improving Knowledge Graph Embeddings with Ontological Reasoning. Proceedings of the International Semantic Web Conference (ISWC), 2021.

Multisensory, User-centered, Shared Cultural Experiences through Interactive Technologies

2022-2025

The project has received funding from the European Union’s Horizon Europe Framework Programme (HORIZON) under grant agreement No 101061441.

MuseIT

- Cultural facts can acquire multiple representations simultaneously
- Input: Music, Text, Formal/structured databases, Images
- Mapping into Multisensory experiences: visual, auditory, touch
- Inclusion by access technologies and co-creation by those with challenges to sight and/or hearing (i.e. evaluation)

MuseIT

Objectives

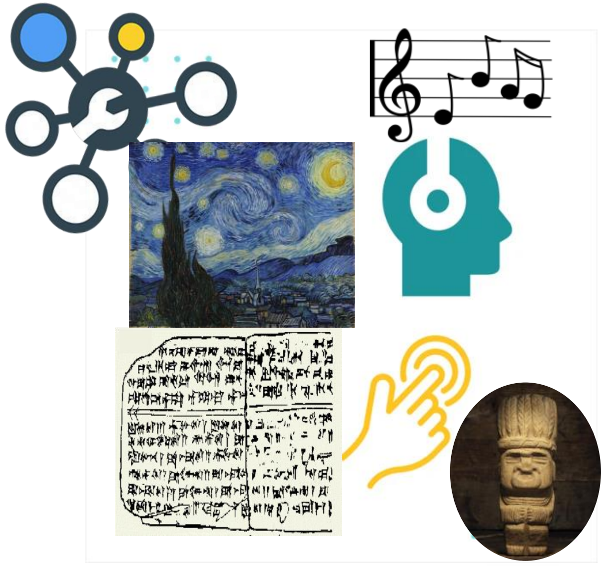

- To investigate statistical and semantic models for multisensory representations of Cultural Heritage (CH).
- To devise innovative methodologies to transform CH modalities and generate them automatically from existing ones with machine learning and crowdsourcing.

Learning Quality Dimensions in Vector Spaces

- The dimensions of learned vector spaces do not normally correspond to semantically meaningful properties.
- This limits the interpretability of learned vector space representations.
- Previous work* on mitigating this issue - identify interpretable directions in learned vector spaces.
- These interpretable directions can then play the role of quality dimensions.

**Joaquín Derrac and Steven Schockaert. Inducing semantic relations from conceptual spaces: A data-driven approach to plausible reasoning. Artif. Intell., 66–94, 2015.*

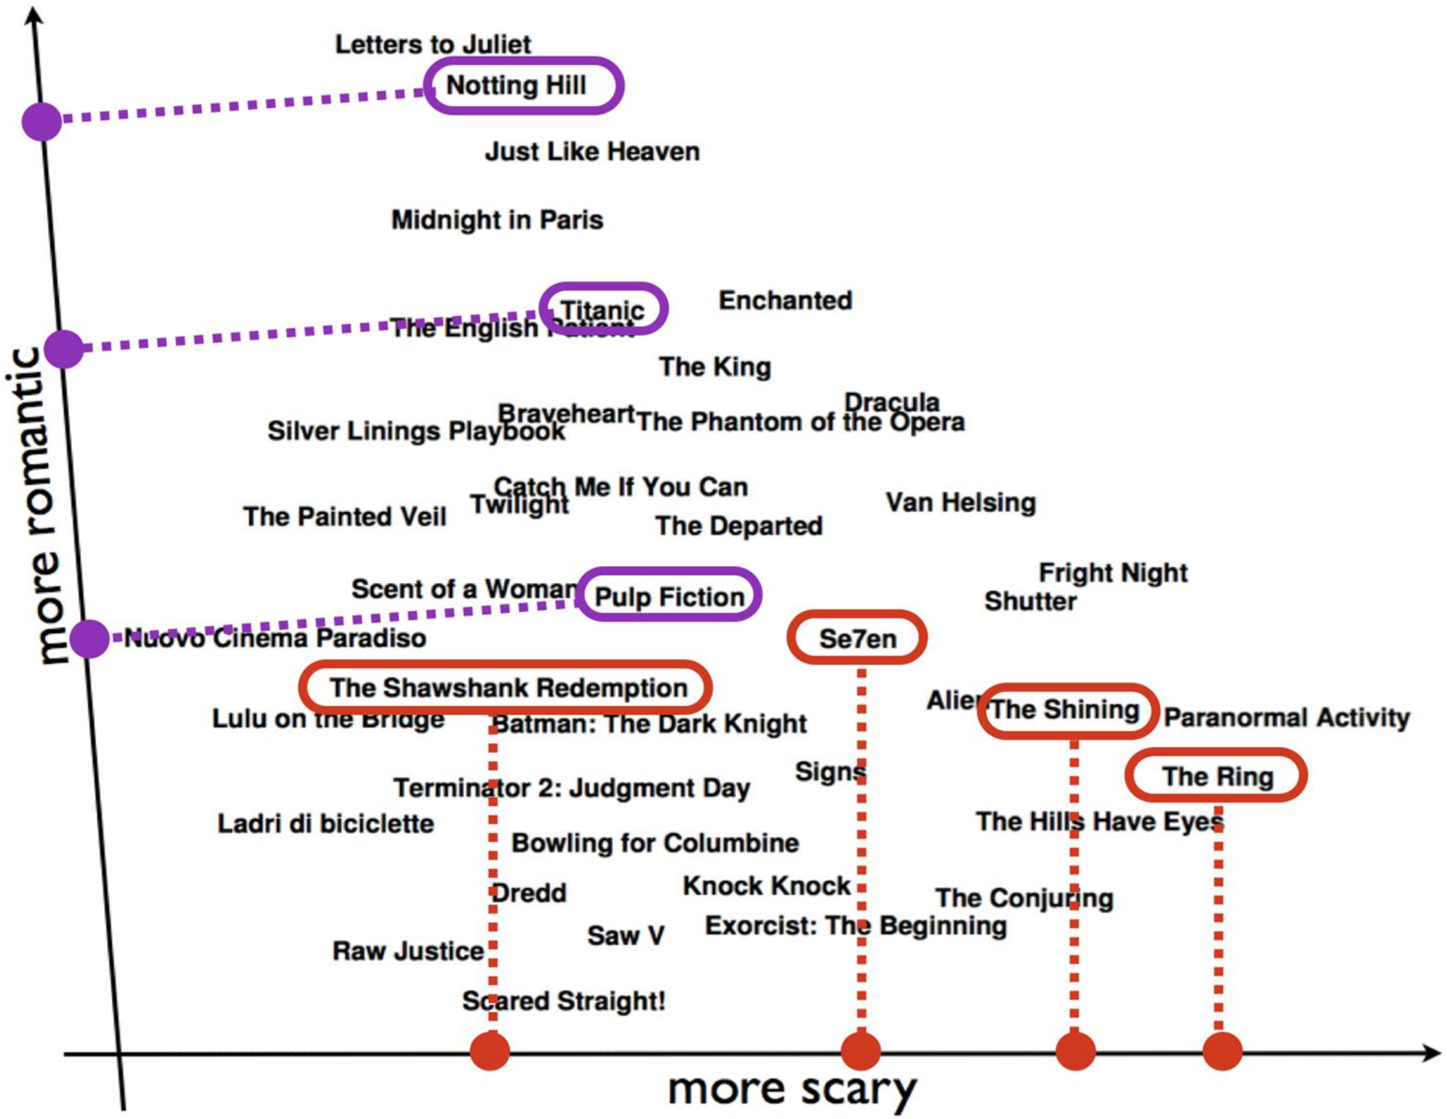

Interpretable directions within a 2d projection of vector space embedding of movies (IMDB)

Learning Quality Dimensions for KG embeddings

- The central aim is to decompose the given vector space into a number of lower-dimensional spaces, each of which captures a different aspect of meaning.
- For KG embeddings, dimensions could correspond to the attributes of entities.
  - E.g., **Movie** entities **-** attributes could be **awards**, **cost** ..
  - **Artist** entities - attributes could be **art style**, **nationality** ..
- This would increase semantic interpretability of the vectors.
- Potential use case for MuseIT - Enable completion of missing attributes of entities, especially helpful for multimodal embeddings.

Multimodal Knowledge Graphs

Multimodal Representations of CH Assets

Multimodal Information

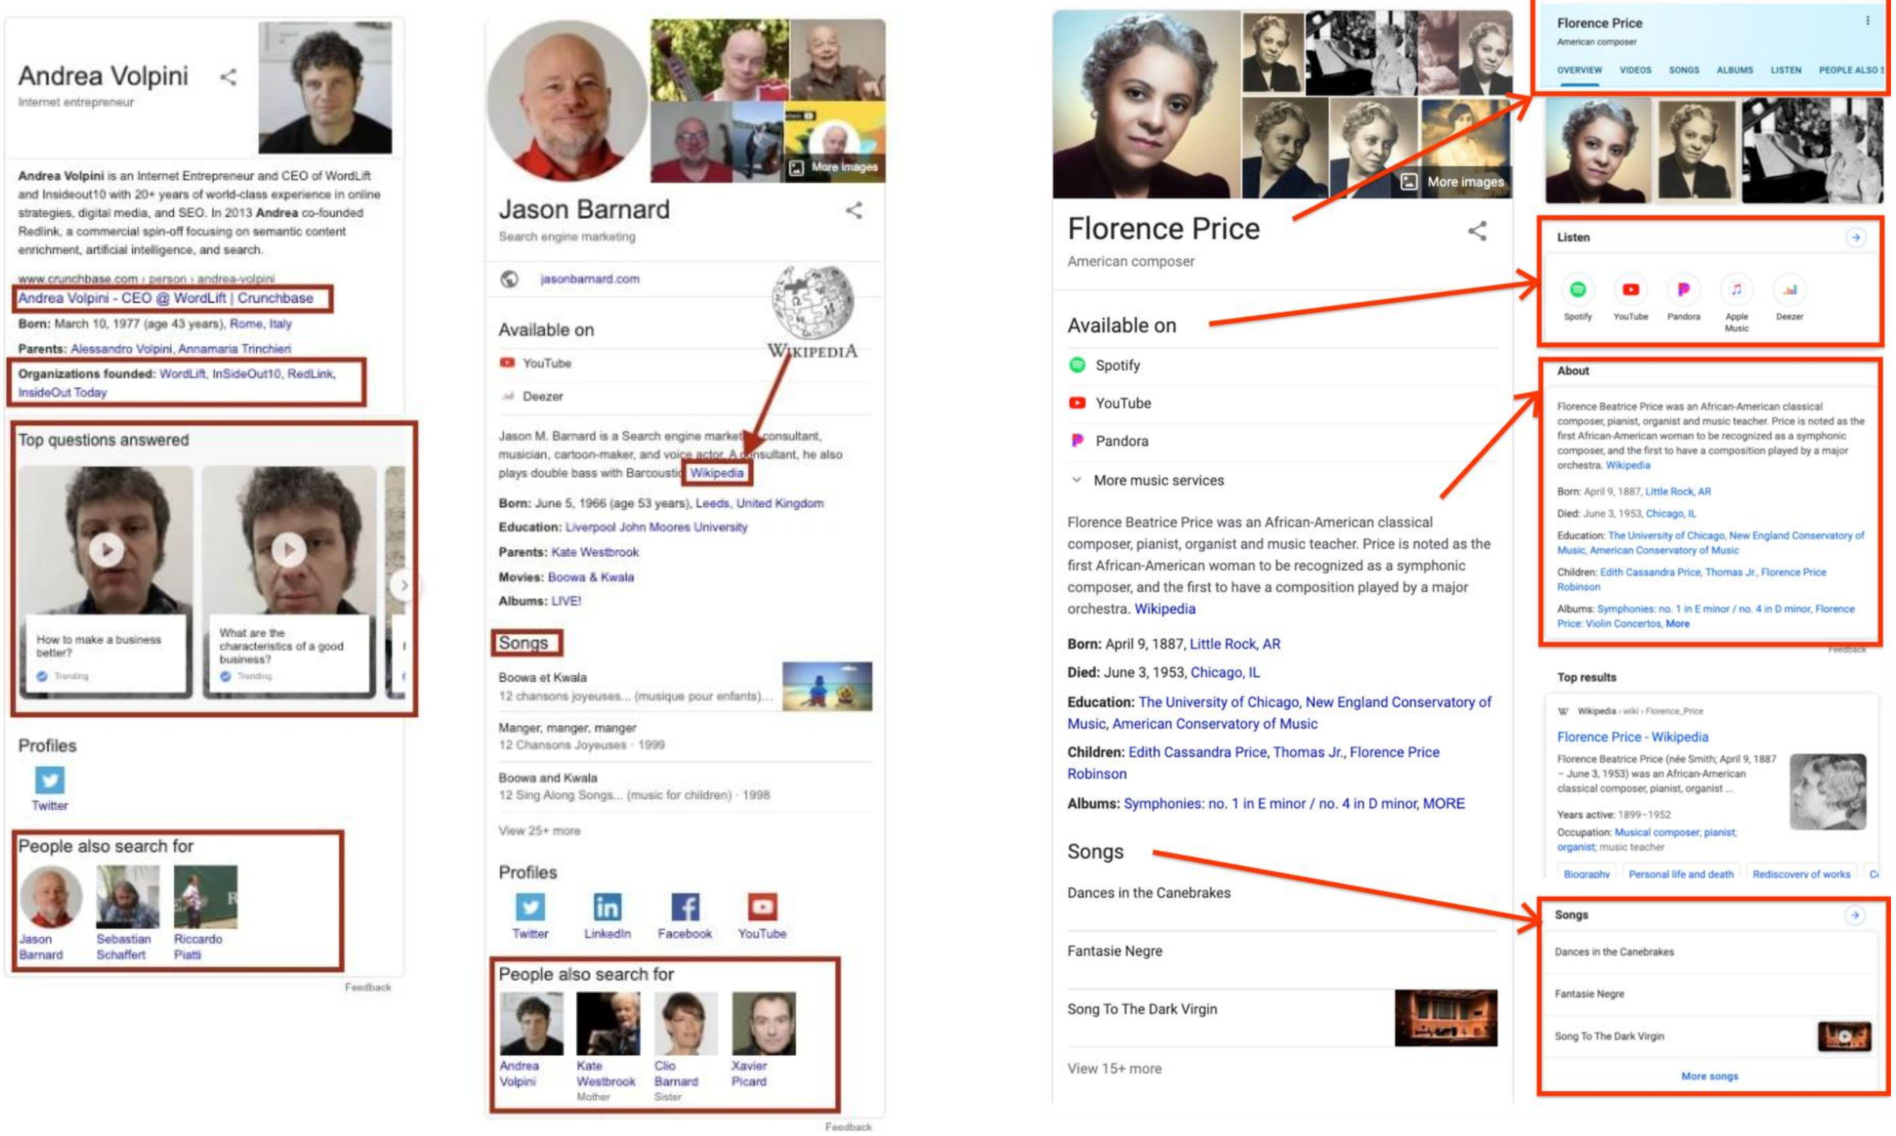

Multimodal Knowledge Graphs

- Real world is inherently heterogeneous.
- Most existing knowledge graphs are represented with pure symbols denoted in the form of text.
- Weakens the capability of machines to describe and understand the real world.
- Establish the connection between the symbol `Dog' and the experience of dogs, i.e. **grounding** a symbol to its physical world meaning.

"Now! *That* should clear up a few things around here!"

Multimodal KGs Benefits

- Relation extraction tasks - additional image helps in attribute extraction
  - partOf (e.g., The keyboard and the screen are parts of a laptop.)
  - colorOf (e.g., A banana is usually yellow or yellowish-green but not blue).

- Text generation tasks - more informative entity-level sentence

  (e.g., Donald Trump is making a speech)

  instead of a vague concept-level description

  (e.g., A tall man with blond hair is making a speech).

MMKGs - Applications

- Multi-modal Entity Recognition and Linking
- Visual Question Answering
- Image-Text Matching
- Multi-modal Generation Tasks
- Multi-modal Recommender System

MMKG with multi-modal data as attribute values

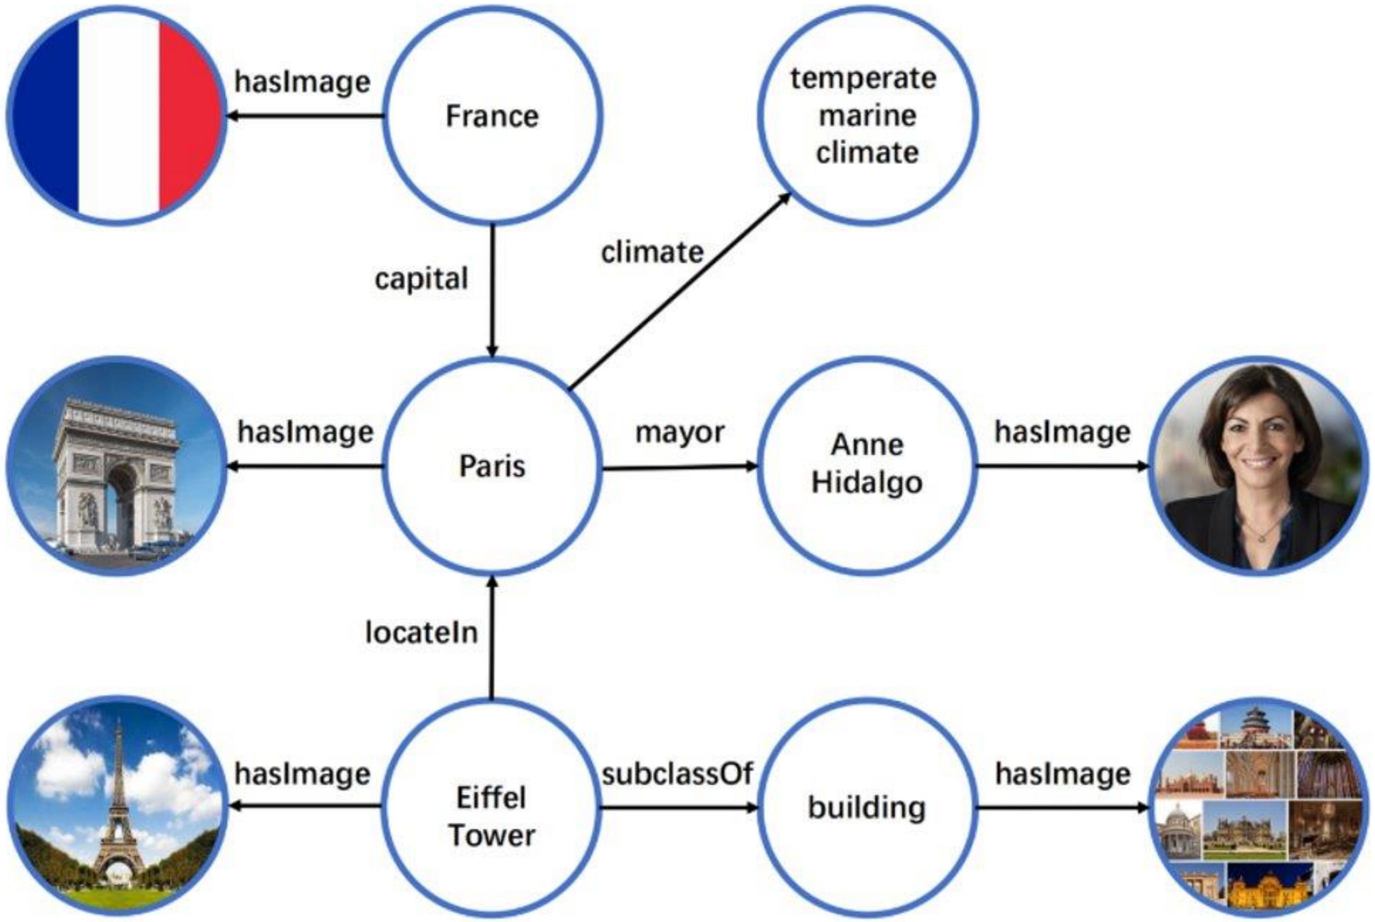

MMKG with multi-modal data as entities

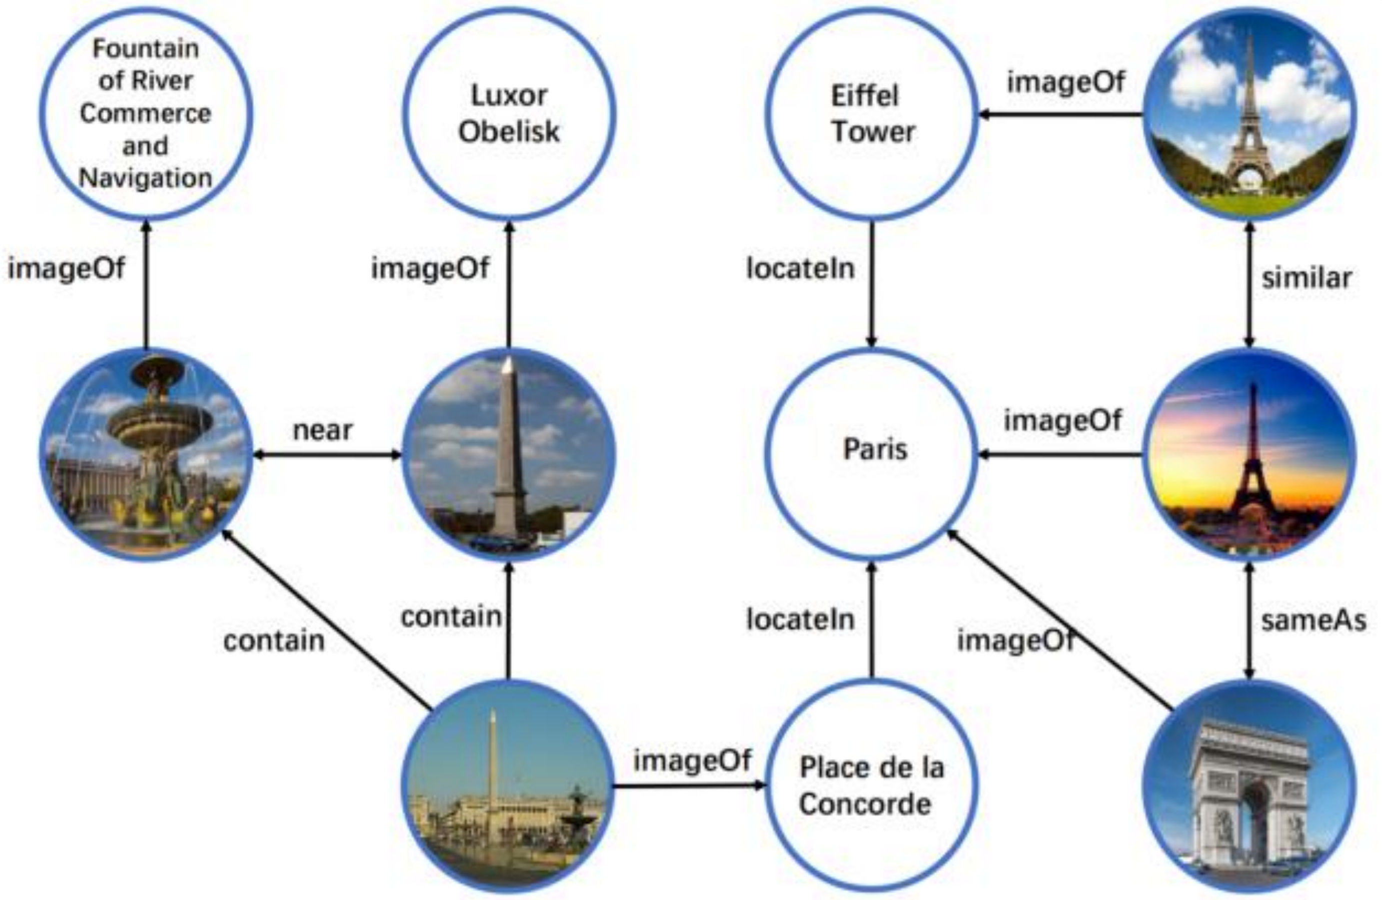

MMKG Construction

Associating symbolic knowledge in a traditional KG, including entities, concepts, relations, etc., with their corresponding images.

(1) labeling images with symbols in KG

(2) grounding symbols in KG to images

MMKG Construction

Labelling images after image segmentation

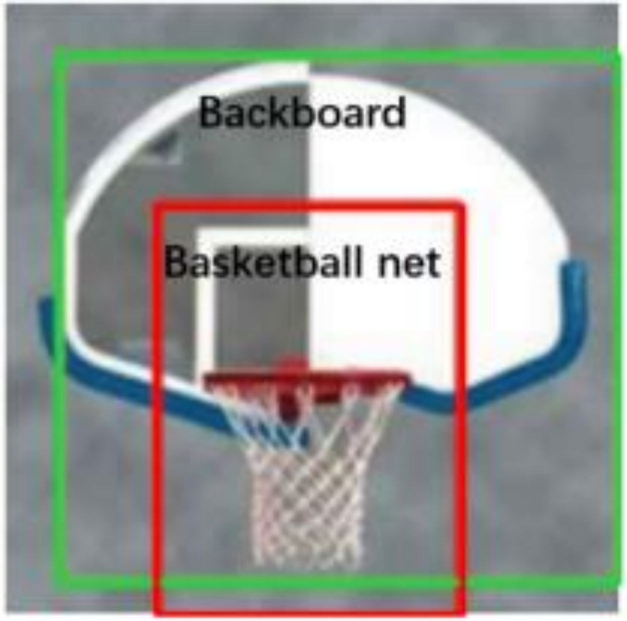

MMKG Construction

Grounding concepts to images

concept type	visualizable concept	non-visualizable concept
example	Surgeon	Physicist
image		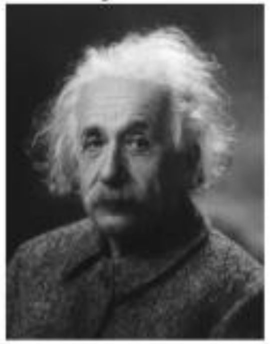

Pre-Trained Models

- Many large companies and research institutions including OpenAI, Microsoft, Huawei trained large very large PTMs based on large-scale unsupervised multi-modal data.
- CLIP trained on 400 million text-image pairs
  - significantly improves the performance of image classification and cross-modal retrieval.

Role of Entity Attributes in multimodal KGs

- Most Knowledge Graph embeddings consider entities and their relations only.
- KGs also house valuable entity information in the form of entity attributes, such as height, weight, or nationality for a person entity.
- Leveraging this attribute information can empower attribute prediction, facilitating the identification of missing features for existing entities.
- Potentially obtain the semantic representations that could facilitate the generation of missing modalities.

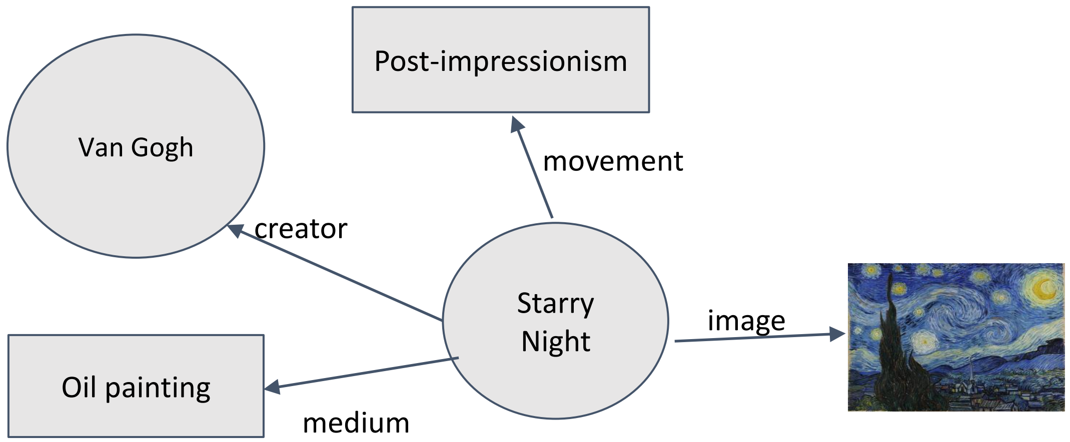

Role of Entity Attributes in multimodal KGs

Are you like me? Entity Attributes for Completion of Multi-modal Cultural Heritage Knowledge Graph

Abstract
Knowledge graphs serve as structured repositories of data, comprising entities, their relations, and attributes. For the cultural heritage domain, where artifacts may belong to different modalities, including textual, audio, visual and haptic, a knowledge graph needs to support multi-modal data. We aim to build such a multi-modal knowledge graph of cultural heritage artifacts for enabling multi-sensory, inclusive and user-centric representations. In this work, we advocate the importance and the merits of leveraging the attributes of cultural heritage entities for the construction and completion of this knowledge graph, which have been largely overlooked by previous works. We present the central role of the entity features for translations between different modalities and prediction of further missing attributes, as well as for the potential generation of entirely novel modalities of the entities. In view of this, we discuss the task of multi-modal attribute prediction in a knowledge graph and outline our strategies for leveraging the predicted attributes for the enrichment of multi-modal representations in the KG.

Keywords
cultural heritage, knowledge graphs, embeddings, multi-modality, entity attributes

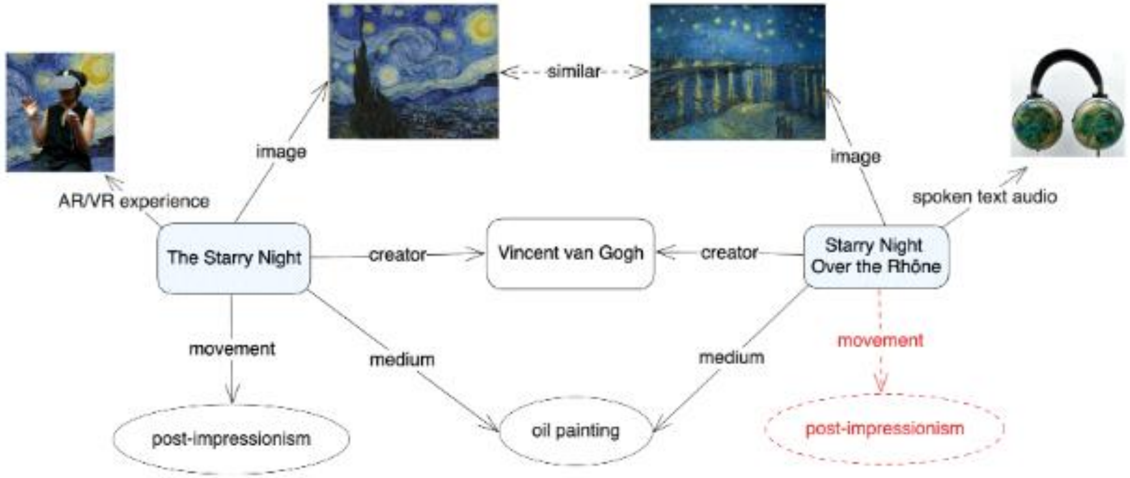

References

- https://github.com/pengfei-luo/multimodal-knowledge-graphb- collection of resources on multimodal knowledge graph, including datasets, research papers and contests.
- https://github.com/ZihengZZH/awesome-multimodal-knowledge-graph - reading list or other resources (datasets, tutorials, etc.), within the research topic of "Multimodal Knowledge Graph".
- Víctor Gutiérrez-Basulto and Steven Schockaert. From knowledge graph embedding to ontology embedding? an analysis of the compatibility between vector space representations and rules. In KR, pages 379–388, 2018.
- Thomas Ager, Ondrej Kuzelka, and Steven Schockaert. Modelling salient features as directions in fine-tuned semantic spaces. In CoNLL, pages 530–540, 2018.
- Rana Alshaikh, Zied Bouraoui, Shelan S. Jeawak, and Steven Schockaert. A mixture-of-experts model for learning multi-facet entity embeddings. In COLING, pages 5124–5135, 2020.
- Rana Alshaikh, Zied Bouraoui, and Steven Schockaert. Learning conceptual spaces with disentangled facets. In CoNLL, pages 131–139, 2019.

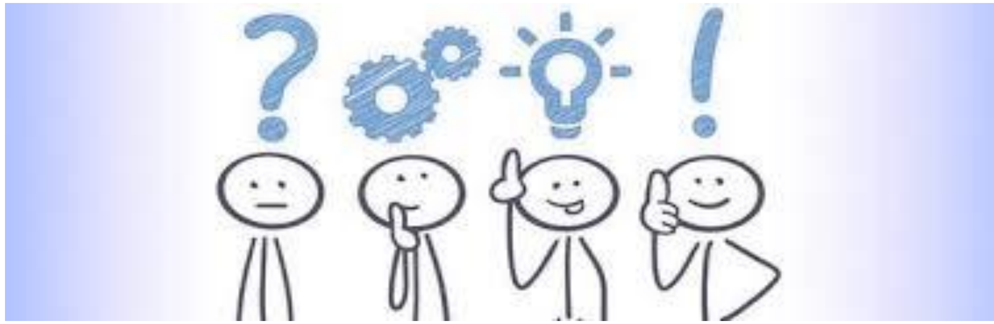

Thank you for your attention.

Questions and comments are welcome !

Contact

Dr Nitisha Jain
Postdoctoral Researcher Associate
Informatics Department
King’s College London (KCL)
nitisha.jain@kcl.ac.uk
https://nitishajain.github.io