

Knowledge Representation, Reasoning and Question Answering

Sargur N. Srihari
srihari@buffalo.edu


Topics in KR & QA

- Embeddings for Representing Knowledge
- Relations
- Link Prediction
- Question Answering


Embeddings for NLP

- Reason for deep learning success in language modeling, MT and NLP
 - Word embeddings
 - Represent semantic knowledge of words/ concepts

Word-to-Vec



Glove: Global Vectors for Word Representation



Unsupervised learning to obtain vector representations for words. Training is performed on aggregated global word-word co-occurrence statistics from a corpus, and the resulting representations showcase interesting linear substructures of the word vector space.

Vector differences *man - woman*, *king - queen*, and *brother - sister* roughly equal.

Embeddings for Knowledge

- Need embeddings for phrases and for relations between words and facts
- Search engines already use ML for this purpose
 - But much more remains to be done to improve these more advanced representations

Distributed Representations for Relations

- For capturing relations between two entities
- Allow formalizing facts about objects and how objects interact with each other

Example of Binary Relation

- Binary relation: set of ordered pairs of objects
 - Pairs in set have the relation while those not in set do not
- Numerical entities
 - Relation less than defined on entities $\{1,2,3\}$ by defining the ordered pairs $S=\{(1,2),(1,3),(2,3)\}$
 - Relation can be used as a verb
 - $(1,2) \in S$ therefore 1 is less than 2
 - $(2,1) \notin S$ therefore we cannot say 2 is less than 1
- Symbolic entities
 - Relation is-a-type defined on entities $\{\text{dog}, \dots\}$, $\{\text{mammal}, \dots\}$ contains ordered pairs (dog,mammal)

Relation in the context of AI

- A relation is a sentence in a syntactically simple and highly structured language
- The relation plays the role of a verb, while two arguments to the relation play the role of its subject and object
- These sentences take the form of a triplet of tokens (subject, verb, object) with values (entity_{*i*}, relation_{*j*}, entity_{*k*})

Definition of Attribute

- We can also define an attribute, a concept analogous to a relation, but taking only one argument:

(entity i , attribute j)

- For example, we could define the `has_fur` attribute, and apply it to entities like `dog`

Reasoning

- Many applications require representing relations and reasoning about them.
- How should we best do this within the context of neural networks?
- Machine learning models of course require training data
- We can infer relations between entities from training datasets consisting of unstructured natural language

Relational Databases

- There are also structured databases that identify relations explicitly.
- A common structure for these databases is the relational database,
 - which stores this same kind of information, albeit not formatted as three-token sentences

Knowledge Base


- Knowledge Base Types
 1. Database conveys commonsense knowledge about everyday life
 - Freebase, OpenCyc, WordNet, or Wikibase, etc
 2. Database has expert knowledge about an application AI system
 - GeneOntology

Representations for Entities/Relations


- Representation can be learned
- Each triplet in knowledge base is a training example
- Maximize a training objective that captures their joint distribution

Tensor Representation


- Binary relational data is represented as a tensor



Fibers of a 3rd order tensor



Slices of a 3rd order tensor



Model Family

- In addition to training data, we also need to define a model family to train
- Extend neural language models to model entities and relations
- Neural language models learn:
 1. A vector that provides a distributed representation of each word
 2. About interactions between words
 - such as which word is likely to come after a sequence of words, by learning functions of these vectors
- Extend approach to entities and relations by learning an embedding vector for each relation

Models from language and KBs

- Close parallel between modeling
 - Language
 - Knowledge encoded as relations
- Representations of entities done using
 - Knowledge bases
 - Natural language sentences
 - Data from multiple relational database

Parameterization

- Many possible parameterizations of models
 1. Highly constrained parametric forms (“linear relational embeddings”), using different forms for relations and entities
 - Vectors for entities
 - Matrices for relations
 - Relation acts like an operator on entities
 2. Relations can be considered as any other entity allowing us to make statements about relations
 - More flexibility is put in machinery that combines them in order to model their joint distribution

Applications of neural knowledge models

- Link prediction:
 - Predicting missing arcs in the knowledge graph
- Word sense disambiguation
 - Deciding which of the senses of a word is the appropriate one, in some context

Link Prediction

- Predicting missing arcs in the knowledge graph
 - A form of generalization to new facts, based on old facts
- Most of the knowledge bases that currently exist have been constructed through manual labor
 - Which tends to leave many and probably the majority of true relations absent from the knowledge base

Evaluating performance of model

- Evaluating link prediction is difficult
- We only have positive examples
 - Facts known to be true
- If model proposes a fact not in dataset, we are unsure whether model has made a mistake or discovered a new, previously unknown fact

Performance Metric

- Metrics are imprecise
 - How model ranks known true positive facts compared to facts that are less likely to be true
- To construct interesting examples that are probably negative (facts that are probably false)
 - Take a true fact and create corrupted versions
 - e.g., by replacing one entity in the relation with a different entity selected at random
- Precision at 10% metric
 - How many times model ranks a “correct” fact among top 10% of all its corrupted versions

Question Answering

- Knowledge of relations combined with a reasoning process and understanding of natural language could allow us to build a general QA system
- A general QA system must be able to process input information and remember important facts, organized in a way that enables it to retrieve and reason about them later
- This remains a difficult open problem which can only be solved in restricted “toy” environments

Memory Networks for remembering Facts

- Current best approach to remembering and retrieving specific declarative facts is to use an explicit memory mechanism
- Memory networks were first proposed to solve a toy question answering task

Explicit Memory Network

- Neural Turing machine
 - Representation part (“task network,” is an RNN)
 - Memory part (cells) can store facts
- Task network
 - Learns to “control” the memory,
 - Deciding where to read from and
 - Where to write to within memory
 - Through reading and writing mechanisms
 - Bold arrows point at reading and writing addresses

