

EMBEDDINGS

E.g. collaborative filtering:

Input: 1mil movies watched by 500k users

Task: Recommend movies to users

→ Need to determine movie similarity

ID of similarity is insufficient to properly capture this.

2D also probably not (nice for visualization though)

→ Similarity = dist. between nodes in the D-dim. space you embed them into

But how you need to compute the embeddings, i.e. compute the D-dim. vectors.

How? Neural net, supervised by "user watched"

Representation?

- Per-user vec $\langle 1, 0, \dots, 1 \rangle$ encoding whether user watched? Slow + big input
- Sparse encoding: list of movie IDs (but see below this is just an optimization for your data)

Applications: ① Predict home sales price and use the home description text as an input to the regression model (e.g. a NN)

So we need to learn an embedding layer to produce D -dim. vectors, then use that as inputs to the regression model.

② Handwriting recognition

→ Learn embedding for the raw bitmap


③ Movie recommendations

→ Learn embedding for user movies sparse encoding

How to learn it?

① Just add a hidden layer in the NN between the $\left\{ \begin{array}{l} \text{of dim. vocab size} \\ \text{nodes} \end{array} \right\}$ input and the rest of the model! $\left\{ \begin{array}{l} \text{of dim. } D \\ \text{nodes} \end{array} \right\}$ Back propagation will just "learn" it! i.e. in passing.

What is D ? Hyperparameter - Tune it. → home price output



→ Back propagation just figures out everything.

- Feature crosses
- Embeddings
- Weights to combine them, etc.

② Standard dimensionality reduction methods
(e.g. principal component analysis)

③ word2vec: similarity = semantic based
on word collocations. Train model.
(Invented at Google)

(See tutorial at [tensorflow.org/tutorials/text/word2vec](https://www.tensorflow.org/tutorials/text/word2vec))