Image2Vec: Learning image representation for reasoning

Lerrel J. Pinto, Gunnar A. Sigurdsson

Word2Vec

- Learning the meaning of words along with semantic relationships
- Vector representations encode linguistic regularities and patterns
- Sparse representations > Dense representations
- Learning enforces contextually similar words get encoded similarly

Want to play around with Word2Vec? : https://www.kaggle.com/c/word2vec-nlp-tutorial

So What?

So What?

Vector Arithmetic holds!

‘Rome’ ~ ‘Paris’ - ‘France’ +‘Italy’

‘Queen’ ~ ‘King’ - ‘Man’ + ‘Woman’

‘Tiger Woods’ ~ ‘Michael Jordan’ - ‘basketball’ + ‘golf’

Check out more word arithmetic :http://www.thisplusthat.me http://nlp.stanford.edu/projects/glove/

Can we do this with images?

Can we do this?

=

-

+

=

+

Image2Vec

Learning a ‘Deep’ Dictionary of attributes

Image2Vec

Learning a ‘Deep’ Dictionary of attributes

CNNs to learn interpretable dictionary

Image2Vec

Learning a ‘Deep’ Dictionary of attributes

CNNs to learn interpretable dictionary

Can we supervise the learning of attributes?

How do we learn this?

How do we learn this?

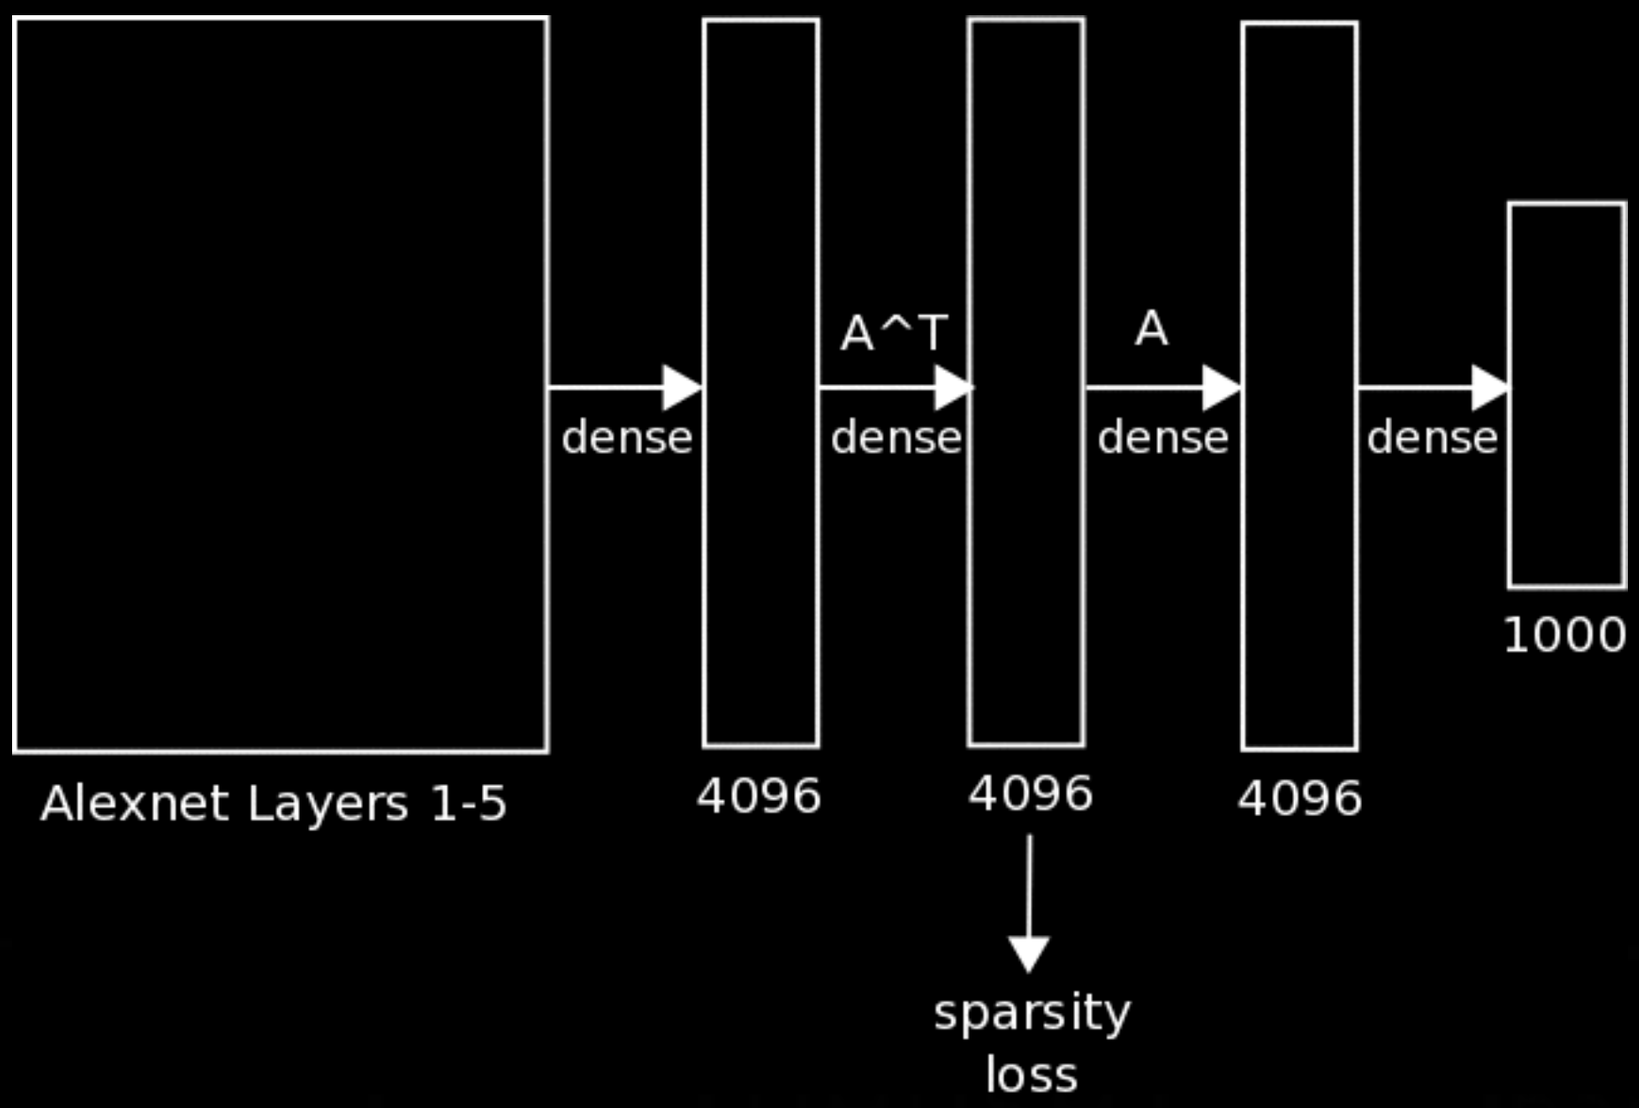

How do we learn this?