

Enhancement of Enterprise Search With Neural Network Language Model

Pengchu Zhang, Laritza Saenz and John Mareda
Sandia National Laboratories
November 30, 2015





Problems in Enterprise Search

- ◆ Search engines are basically matching the query terms and the terms in documents, this may result in a large number of false positives;
- ◆ Some terms may be overly represented in search:
 - E.g., "composite materials"
- ◆ Search engines may fail to deliver if the query terms not exist in the enterprise datasets but the datasets have the relevant information, e.g.,
 - Buckyballs vs. fullerene



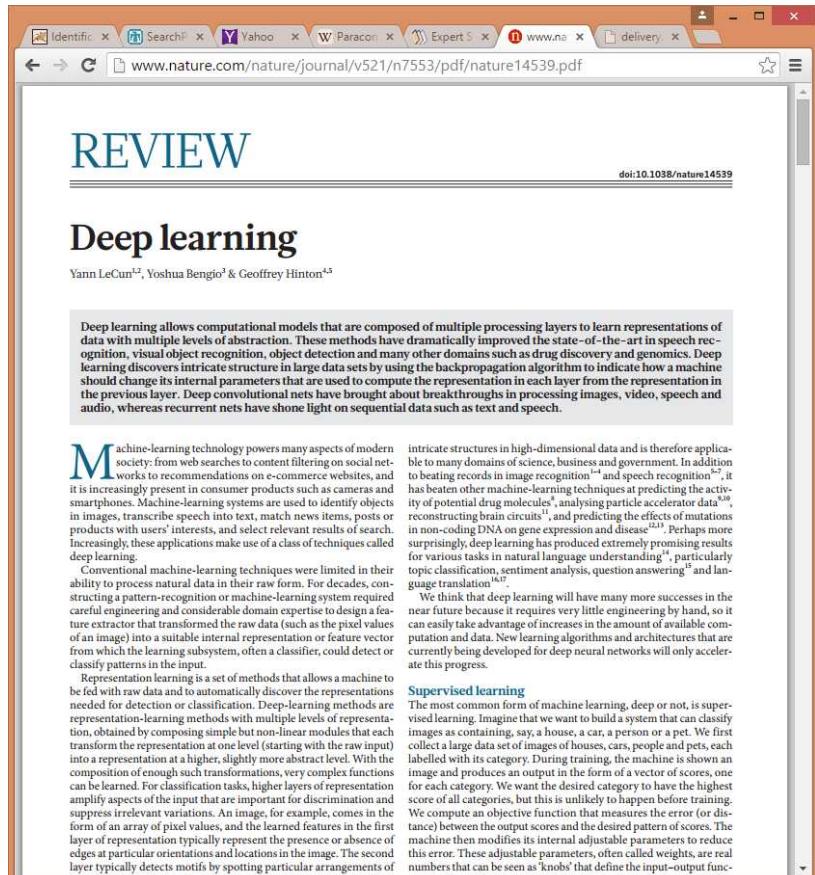
Current Approaches to Enhance Information finding

- ◆ Boosting some terms;
 - E.g., Composite(+2) materials
- ◆ Synonym to expand query terms;
 - Java vs. coffee
- ◆ Others such as applying various similarity algorithms

In this presentation, we propose to improve information findability with Neural Network Language Models



NNLM in NATURE



REVIEW

doi:10.1038/nature14539

Deep learning

Yann LeCun^{1,2}, Yoshua Bengio³ & Geoffrey Hinton^{4,5}

Deep learning allows computational models that are composed of multiple processing layers to learn representations of data with multiple levels of abstraction. These methods have dramatically improved the state-of-the-art in speech recognition, visual object recognition, object detection and many other domains such as drug discovery and genomics. Deep learning discovers intricate structure in large data sets by using the backpropagation algorithm to indicate how a machine should change its internal parameters that are used to compute the representation in each layer from the representation in the previous layer. Deep convolutional nets have brought about breakthroughs in processing images, video, speech and audio, whereas recurrent nets have shone light on sequential data such as text and speech.

Machine-learning technology powers many aspects of modern society: from web searches to content filtering on social networks, to recommendations on e-commerce websites, and it is increasingly present in consumer products such as cameras and smartphones. Machine-learning systems are used to identify objects in images, transcribe speech into text, match news items, posts or products with users' interests, and select relevant results of search. Increasingly, these applications make use of a class of techniques called deep learning.

Conventional machine-learning techniques were limited in their ability to process natural data in their raw form. For decades, constructing a pattern-recognition or machine-learning system required careful engineering and considerable domain expertise to design a feature extractor that transformed the raw data (such as the pixel values of an image) into a suitable internal representation or feature vector from which the learning subsystem, often a classifier, could detect or classify patterns in the input.

Representation learning is a set of methods that allows a machine to be fed with raw data and to automatically discover the representations needed for detection or classification. Deep-learning methods are representation-learning methods with multiple levels of representation, learned by unsupervised learning, and not by hand, that model transform the representation at one level (starting with the raw input) into a representation at a higher, slightly more abstract level. With the composition of enough such transformations, very complex functions can be learned. For classification tasks, higher layers of representation amplify aspects of the input that are important for discrimination and suppress irrelevant variations. An image, for example, comes in the form of an array of pixel values, and the learned features in the first layer of representation typically represent the presence or absence of edges at particular orientations and locations in the image. The second layer typically detects motifs by spotting particular arrangements of

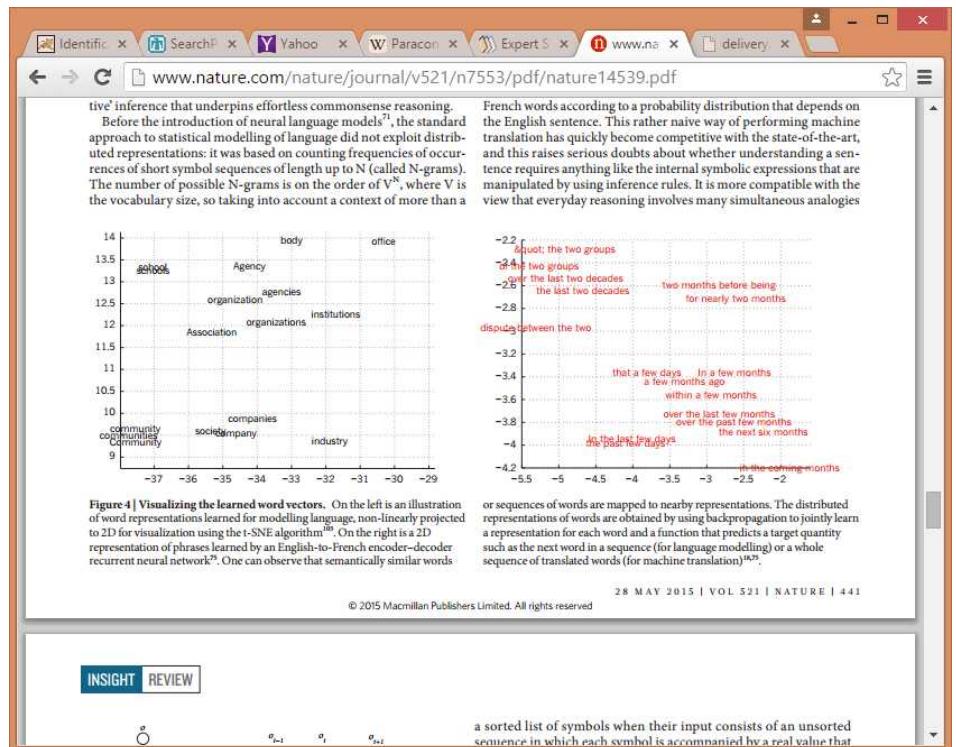
intricate structures in high-dimensional data and is therefore applicable to many domains of science, business and government. In addition to hosting records in image recognition^{1–4} and speech recognition^{5–7}, it has beaten other machine-learning techniques at predicting the activity of potential drug molecules⁸, analysing particle accelerator data^{9,10}, reconstructing brain circuits¹¹, and predicting the effects of mutations in non-coding DNA on gene expression and disease^{12,13}. Perhaps more surprisingly, deep learning has produced extremely promising results for various tasks in natural language understanding¹⁴, particularly topic classification, sentiment analysis, question answering¹⁵ and language translation^{16,17}.

We think that deep learning will have many more successes in the near future because it requires very little engineering by hand, so it can easily take advantage of increases in the amount of available computation and data. New learning algorithms and architectures that are currently being developed for deep neural networks will only accelerate this progress.

Supervised learning

The most common form of machine learning, deep or not, is supervised learning. Imagine that we want to build a system that can classify images of a person, a place, a person at a place. We first collect a large data set of images of these, of people and places labelled with their category. During training, the machine is shown an image and produces an output in the form of a vector of scores, one for each category. We want the desired category to have the highest score of all categories, but this is unlikely to happen before training. We compute an objective function that measures the error (or distance) between the output scores and the desired pattern of scores. The machine then modifies its internal adjustable parameters to reduce this error. These adjustable parameters, often called weights, are real numbers that can be seen as 'knobs' that define the input-output function.

From Nature 521, 4346-444 (28 May 2015)

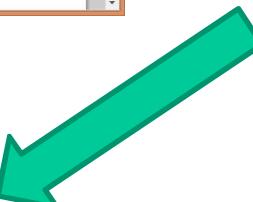
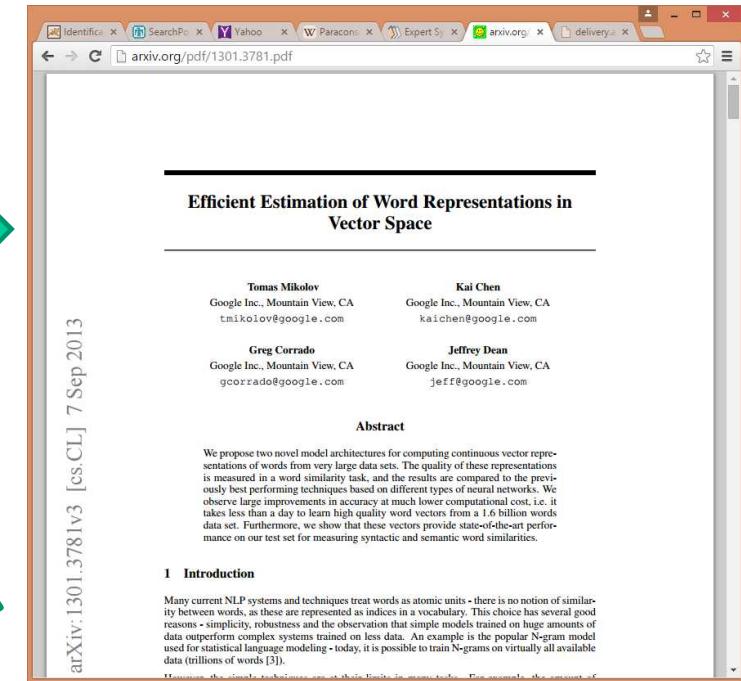
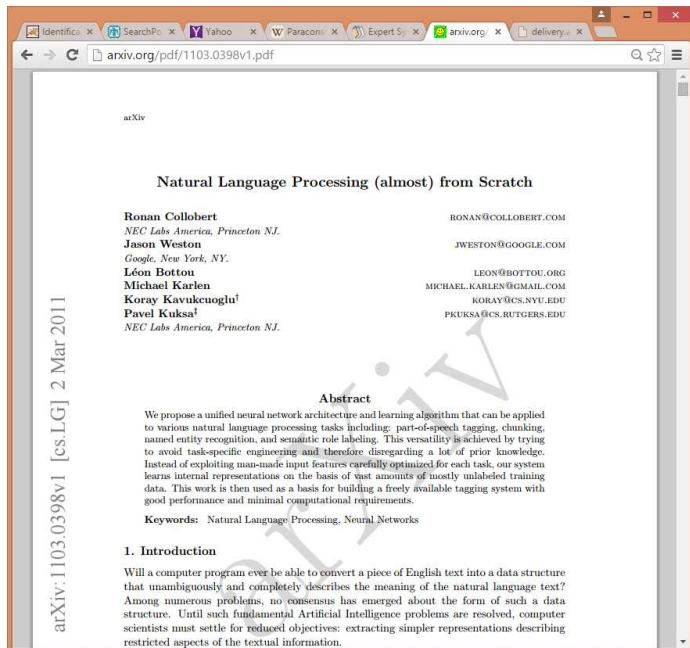


INSIGHT REVIEW

Sandia National Laboratories



From Theory to Application to Codes



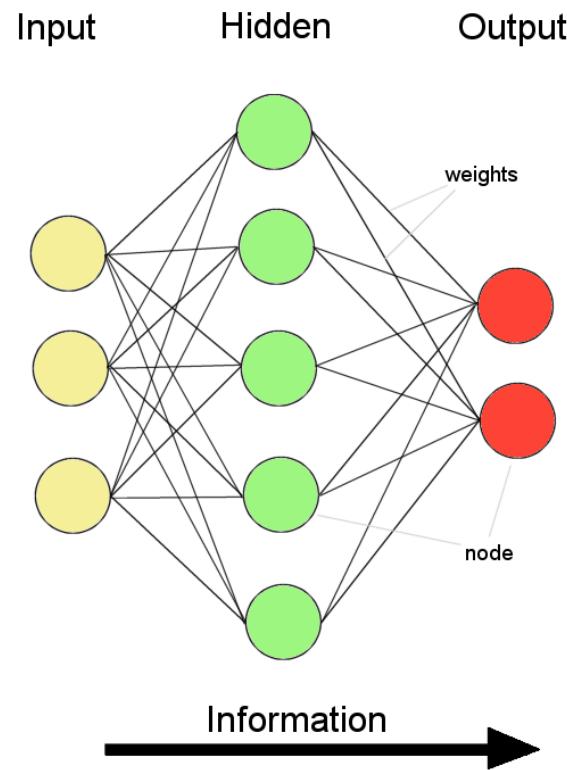
<https://code.google.com/p/word2vec/>

Sandia National Laboratories

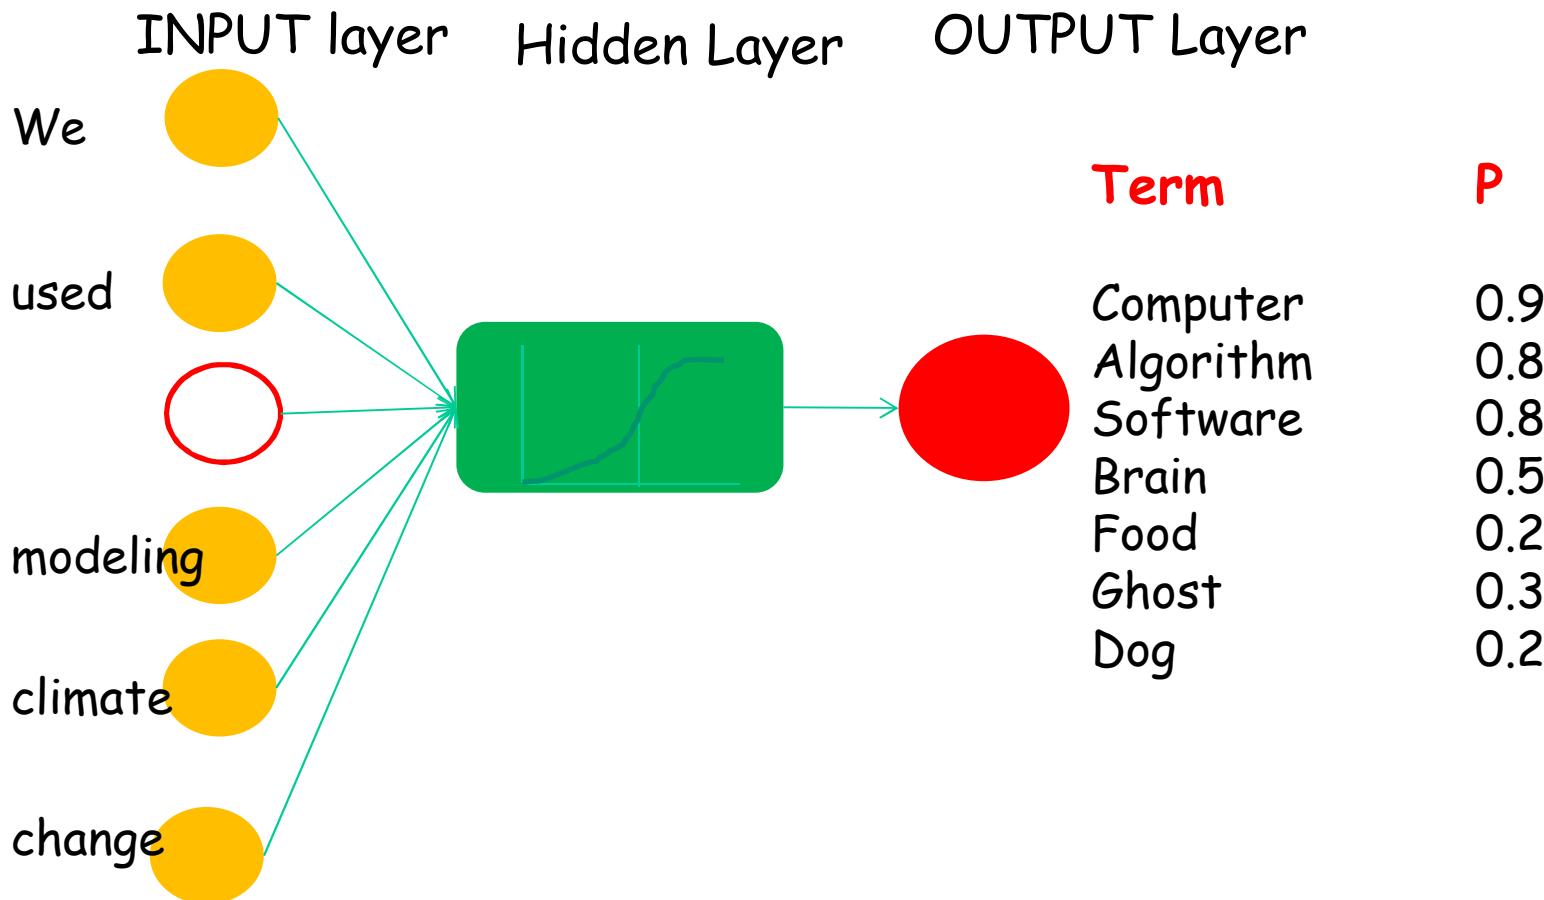


What is Neural Network Language Model?

Neural Network



Concept of Neural Network Language Model

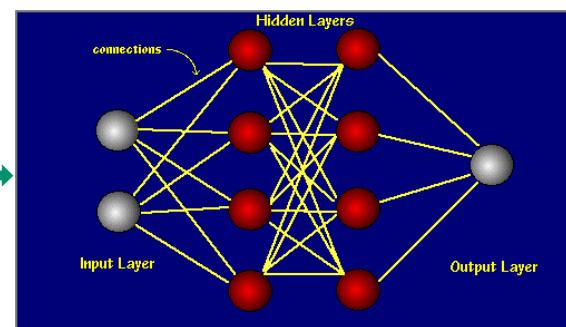


Build and optimize word vectors(Google 2013)

Word in sentence vector

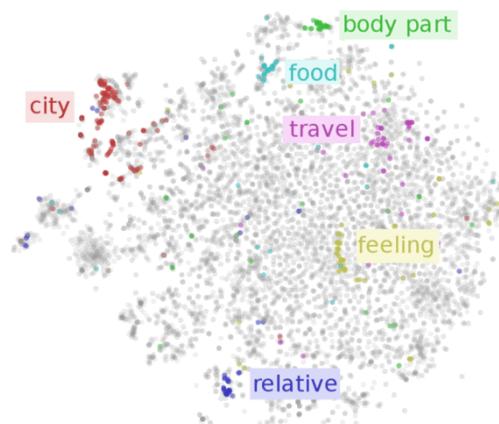
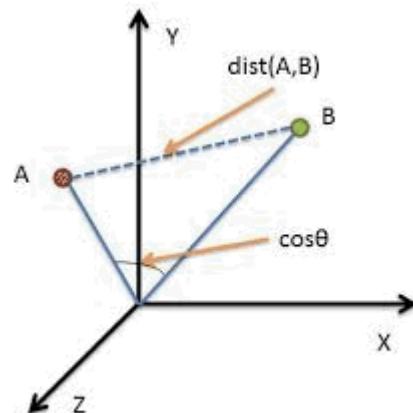
The
dog
is
walking
in
the
room

(0.12, 0.23, 0.22)
(0.32, 0.27, 0.94)
(0.18, 0.88, 0.45)
(0.23, 0.92, 0.23)
(0.77, 0.25, 0.11)
(0.12, 0.23, 0.22)
(0.41, 0.13, 0.29)



New vector

(0.12, 0.23, 0.22)
(0.62, 0.99, 0.14)
(0.18, 0.88, 0.45)
(0.23, 0.92, 0.23)
(0.77, 0.25, 0.11)
(0.12, 0.23, 0.22)
(0.41, 0.13, 0.29)

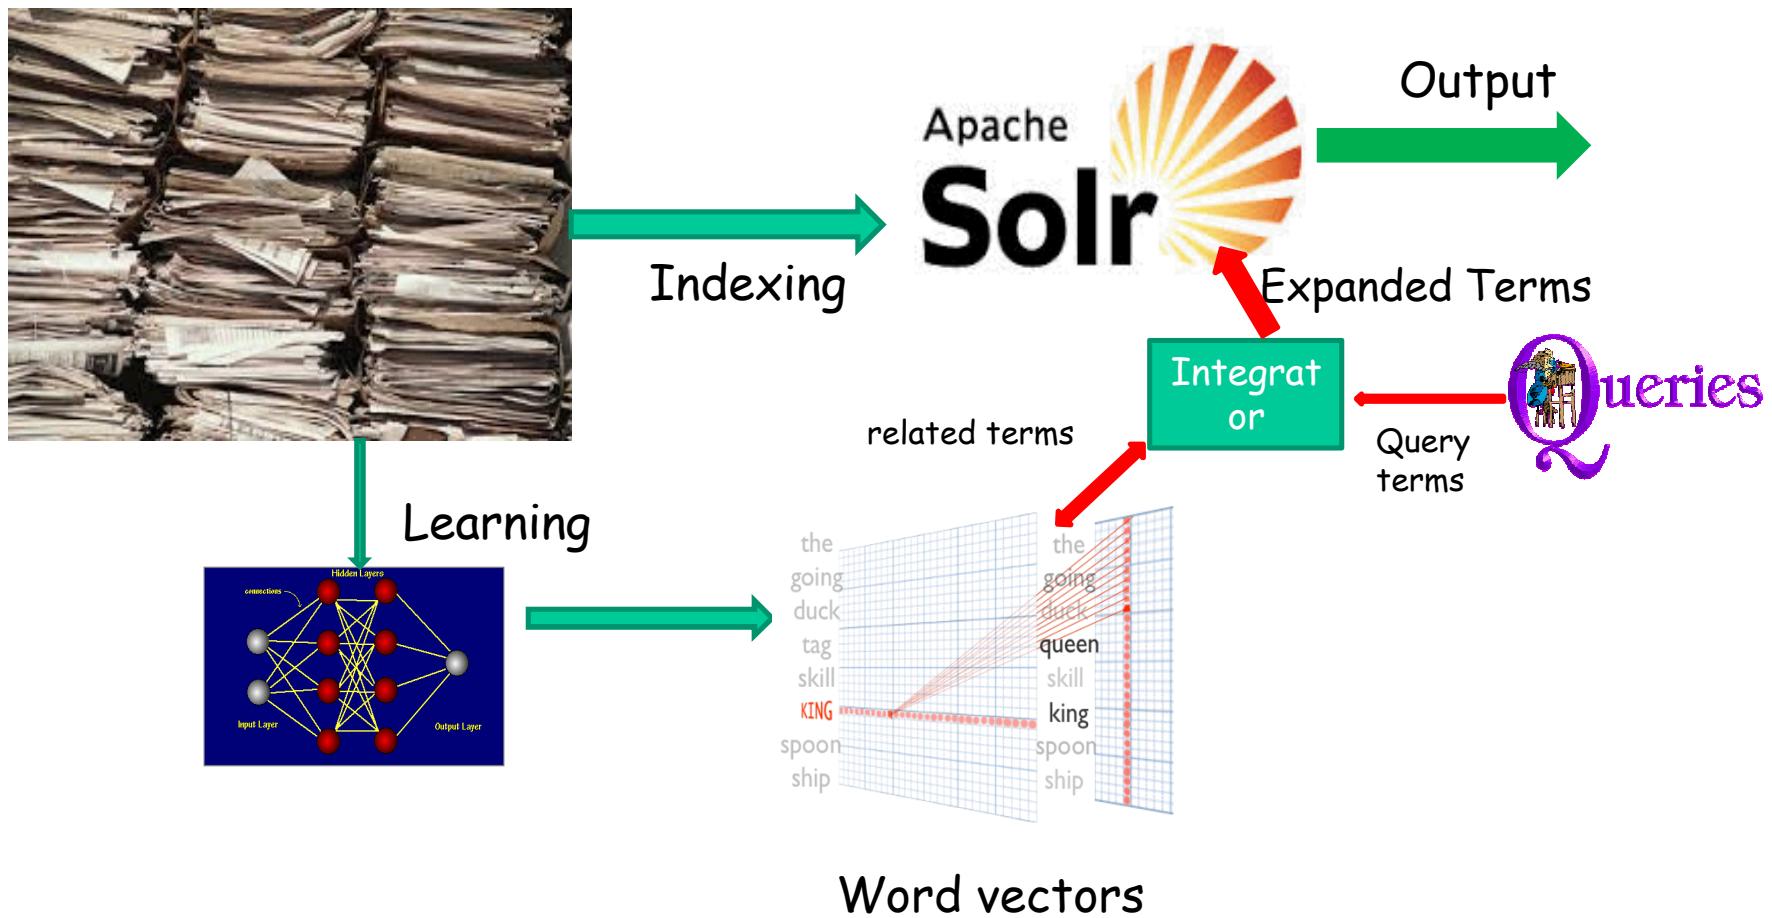




Motivations in Applying Neural Network Language Model (NNLM)

- ◆ NNLM projects the terms into vector space, meaning the words that were used together will stay together at the vector space;
- ◆ NNLM vectorizes the terms that makes it possible to compute the distances between the terms;
- ◆ Most importantly, NNLM learns from natural language semantically and syntactically. With this NNLM meaningfully brings words together without parsing a speech or a document.

Concept for Enhanced Enterprise Search



Note: All images are from online

Sandia National Laboratories



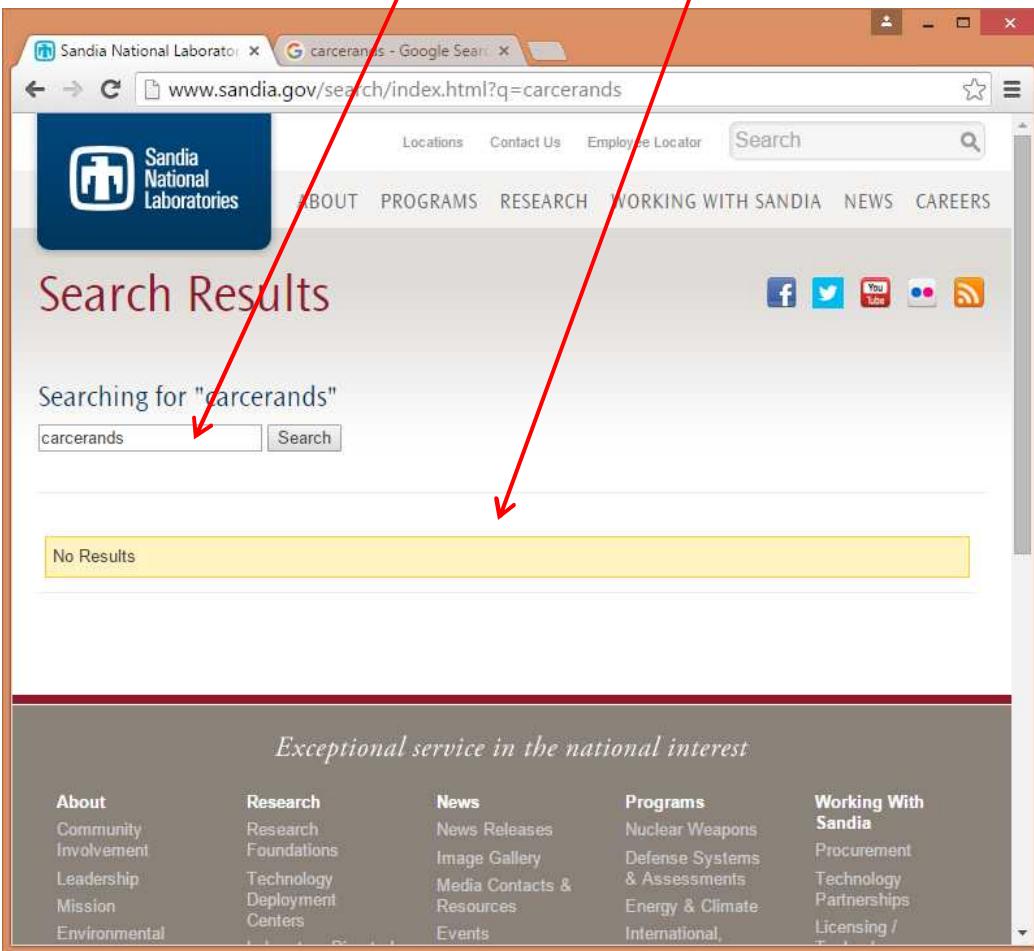


What and How NNLM helps Enterprise Search

- ◆ Deliver the relevant information even the query terms not in enterprise dataset;
- ◆ Increase the quality of information delivered:
 - Relevance
 - Related

- Deliver the relevant information even the query terms not in enterprise dataset

Query term “**carcerands**” not the in enterprise dataset



Sandia National Laboratories - Google Search

www.sandia.gov/search/index.html?q=carcerands

Sandia National Laboratories

Locations Contact Us Employee Locator Search

ABOUT PROGRAMS RESEARCH WORKING WITH SANDIA NEWS CAREERS

Search Results

Searching for "carcerands"

carcerands

No Results

Exceptional service in the national interest

About

- Community Involvement
- Leadership
- Mission
- Environmental

Research

- Research Foundations
- Technology Deployment Centers

News

- News Releases
- Image Gallery
- Media Contacts & Resources
- Events

Programs

- Nuclear Weapons
- Defense Systems & Assessments
- Energy & Climate International

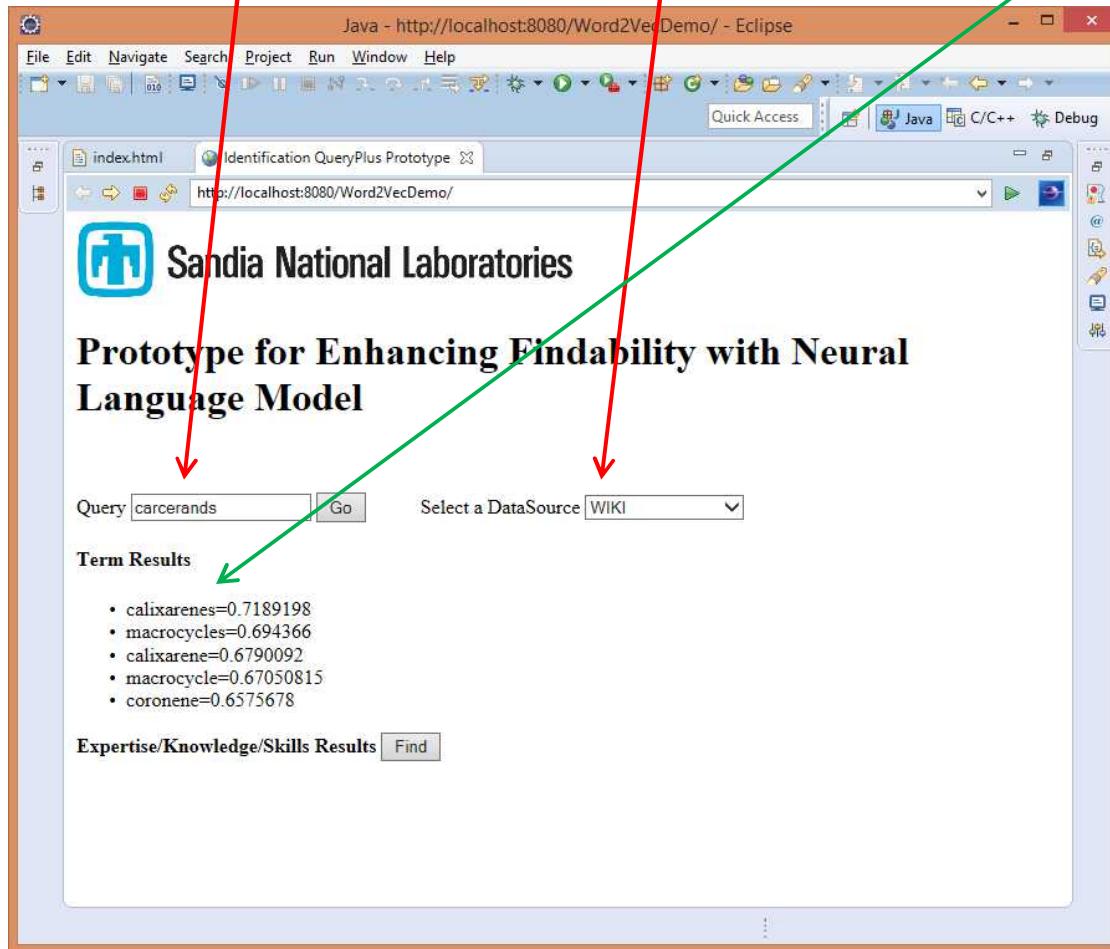
Working With Sandia

- Procurement
- Technology Partnerships
- Licensing /

Sandia National Laboratories



Term “**carcerands**” is found in Wikipedia English and a set of close terms are generated from NNLM



Java - http://localhost:8080/Word2VecDemo/ - Eclipse

File Edit Navigate Search Project Run Window Help

indexhtml Identification QueryPlus Prototype

http://localhost:8080/Word2VecDemo/

Sandia National Laboratories

Prototype for Enhancing Findability with Neural Language Model

Query Go Select a DataSource WIKI

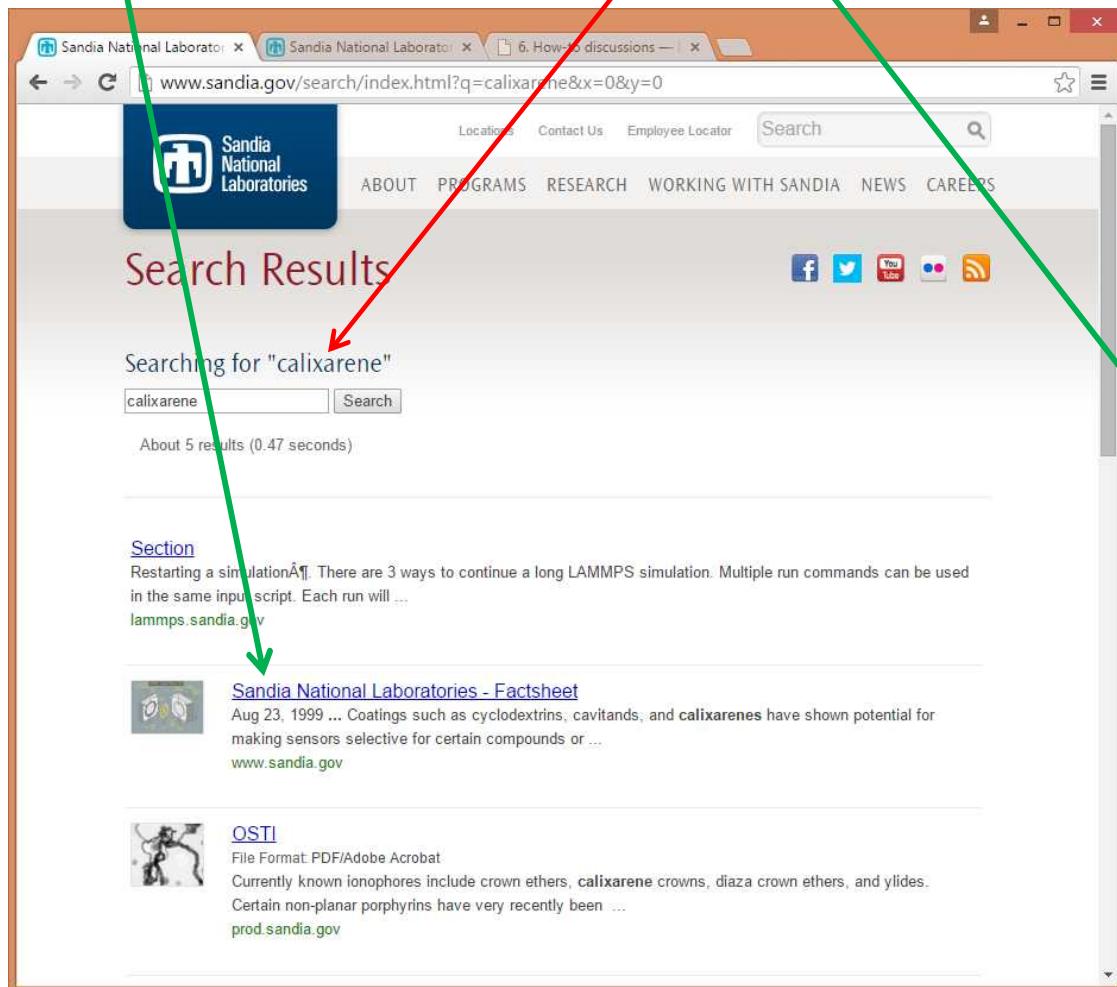
Term Results

- calixarenes=0.7189198
- macrocycles=0.694366
- calixarene=0.6790092
- macrocycle=0.67050815
- coronene=0.6575678

Expertise/Knowledge/Skills Results



Sandia has a number of documents associated with "carcerands" for which it is related to "calixarenes"



Search Results

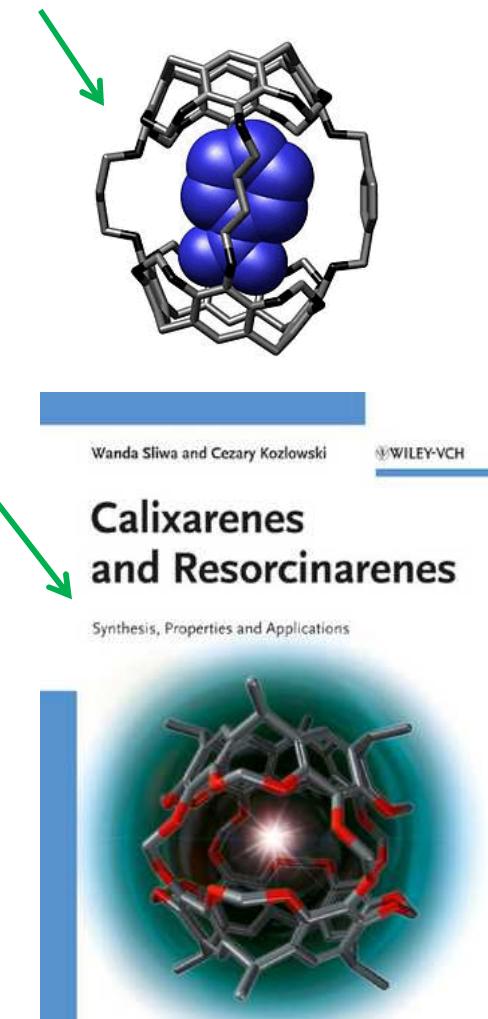
Searching for "calixarene"

calixarene

About 5 results (0.47 seconds)

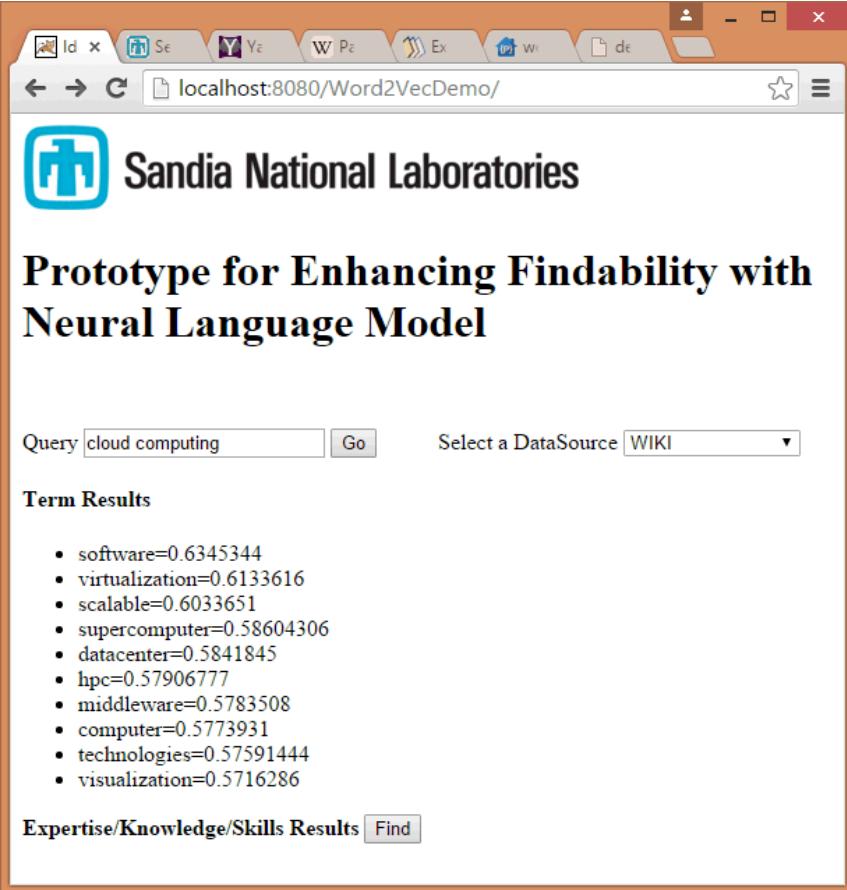
[Sandia National Laboratories - Factsheet](#)
Aug 23, 1999 ... Coatings such as cyclodextrins, cavitands, and calixarenes have shown potential for making sensors selective for certain compounds or ...
www.sandia.gov

[OSTI](#)
File Format: PDF/Adobe Acrobat
Currently known ionophores include crown ethers, calixarene crowns, diaza crown ethers, and ylides.
Certain non-planar porphyrins have very recently been ...
prod.sandia.gov





Increase the quality of retrieved information delivered: **Relevant**



localhost:8080/Word2VecDemo/

Sandia National Laboratories

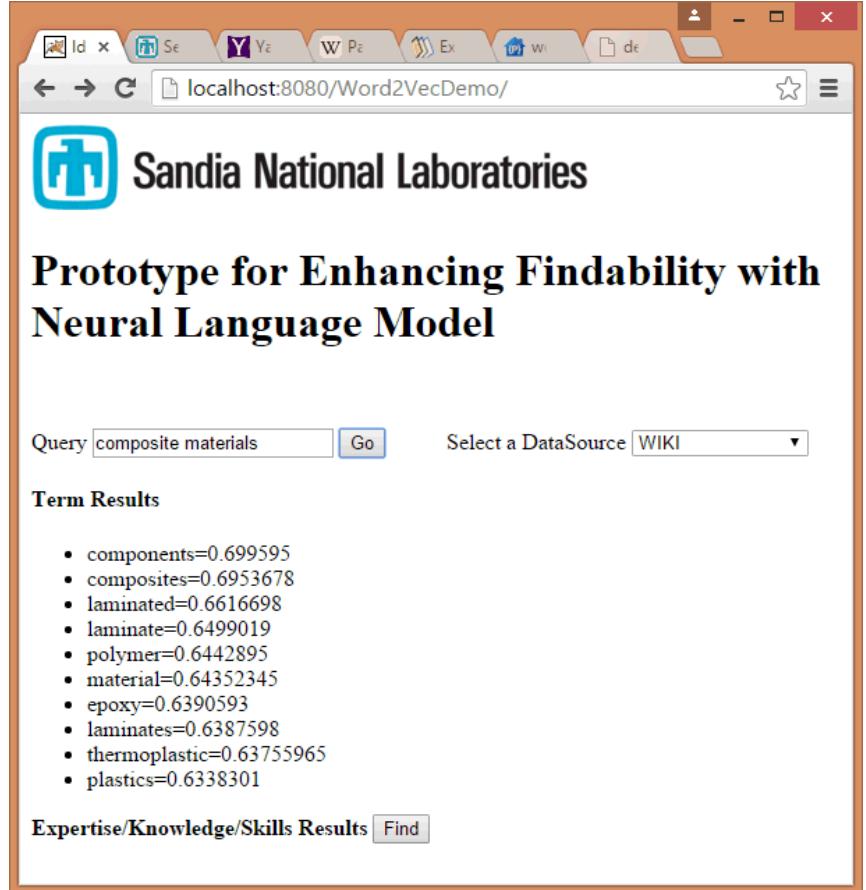
Prototype for Enhancing Findability with Neural Language Model

Query Select a DataSource

Term Results

- software=0.6345344
- virtualization=0.6133616
- scalable=0.6033651
- supercomputer=0.58604306
- datacenter=0.5841845
- hpc=0.57906777
- middleware=0.5783508
- computer=0.5773931
- technologies=0.57591444
- visualization=0.5716286

Expertise/Knowledge/Skills Results



localhost:8080/Word2VecDemo/

Sandia National Laboratories

Prototype for Enhancing Findability with Neural Language Model

Query Select a DataSource

Term Results

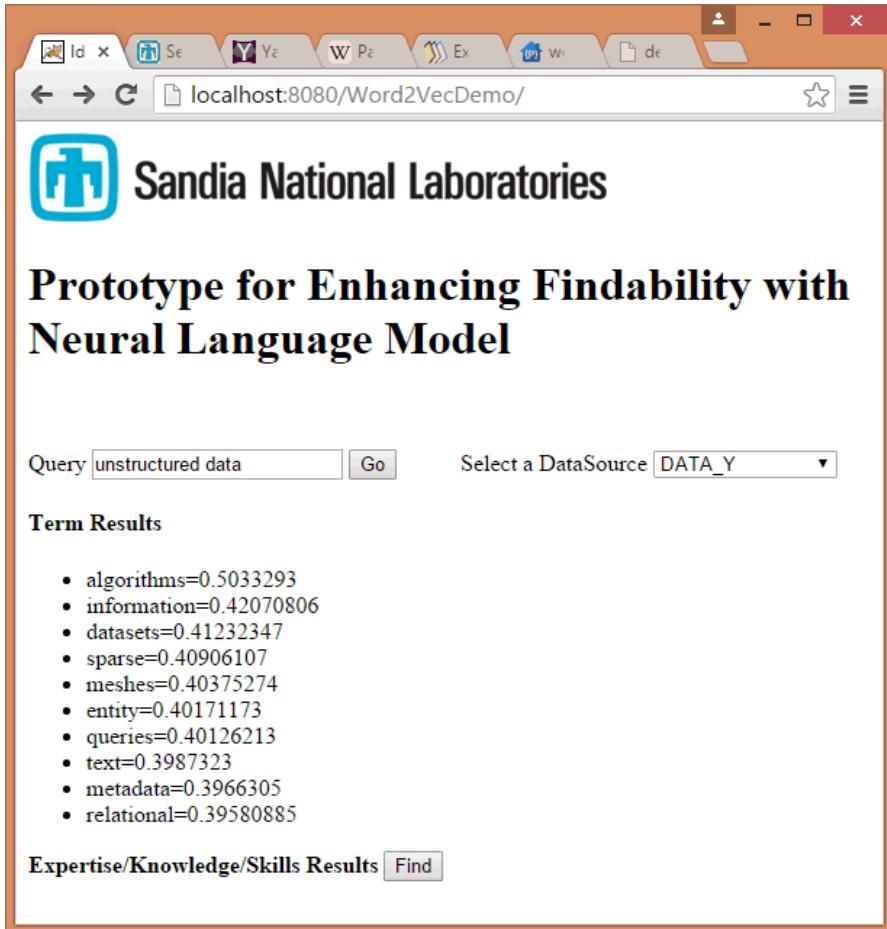
- components=0.699595
- composites=0.6953678
- laminated=0.6616698
- laminate=0.6499019
- polymer=0.6442895
- material=0.64352345
- epoxy=0.6390593
- laminates=0.6387598
- thermoplastic=0.63755965
- plastics=0.6338301

Expertise/Knowledge/Skills Results





Increase the quality of retrieved information delivered: **Related**



localhost:8080/Word2VecDemo/

Sandia National Laboratories

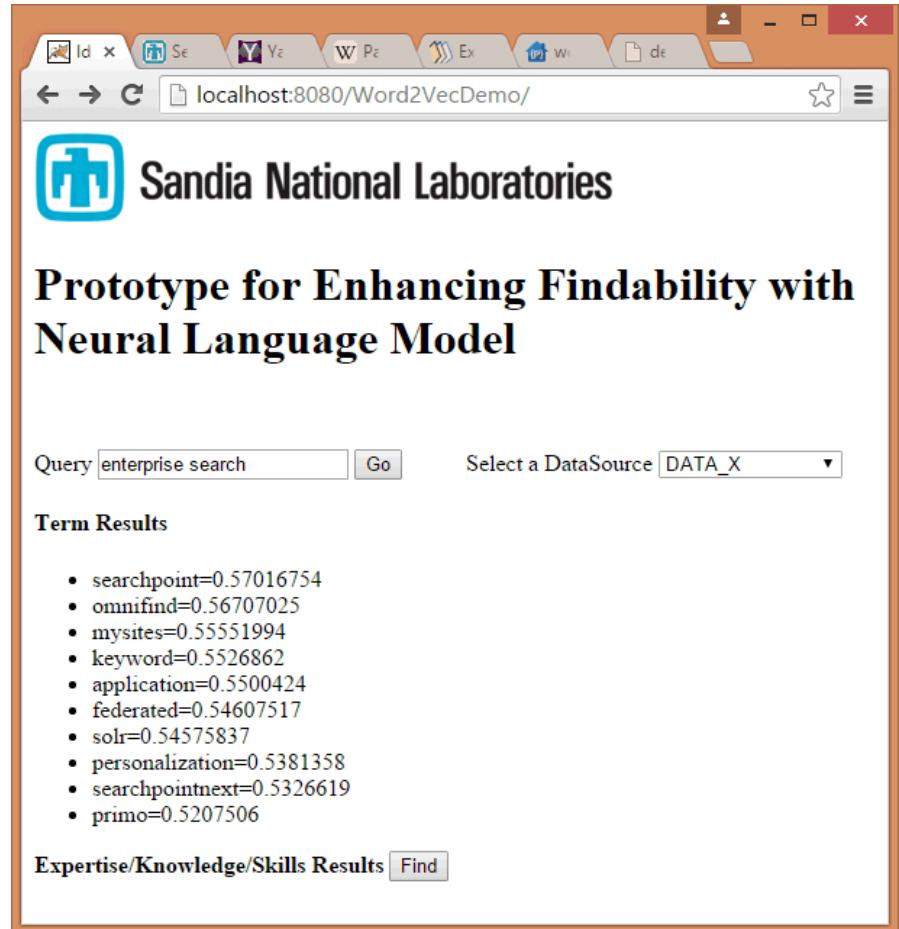
Prototype for Enhancing Findability with Neural Language Model

Query Select a DataSource

Term Results

- algorithms=0.5033293
- information=0.42070806
- datasets=0.41232347
- sparse=0.40906107
- meshes=0.40375274
- entity=0.40171173
- queries=0.40126213
- text=0.3987323
- metadata=0.3966305
- relational=0.39580885

Expertise/Knowledge/Skills Results



localhost:8080/Word2VecDemo/

Sandia National Laboratories

Prototype for Enhancing Findability with Neural Language Model

Query Select a DataSource

Term Results

- searchpoint=0.57016754
- omnifind=0.56707025
- mysites=0.55551994
- keyword=0.5526862
- application=0.5500424
- federated=0.54607517
- solr=0.54575837
- personalization=0.5381358
- searchpointnext=0.5326619
- primo=0.5207506

Expertise/Knowledge/Skills Results

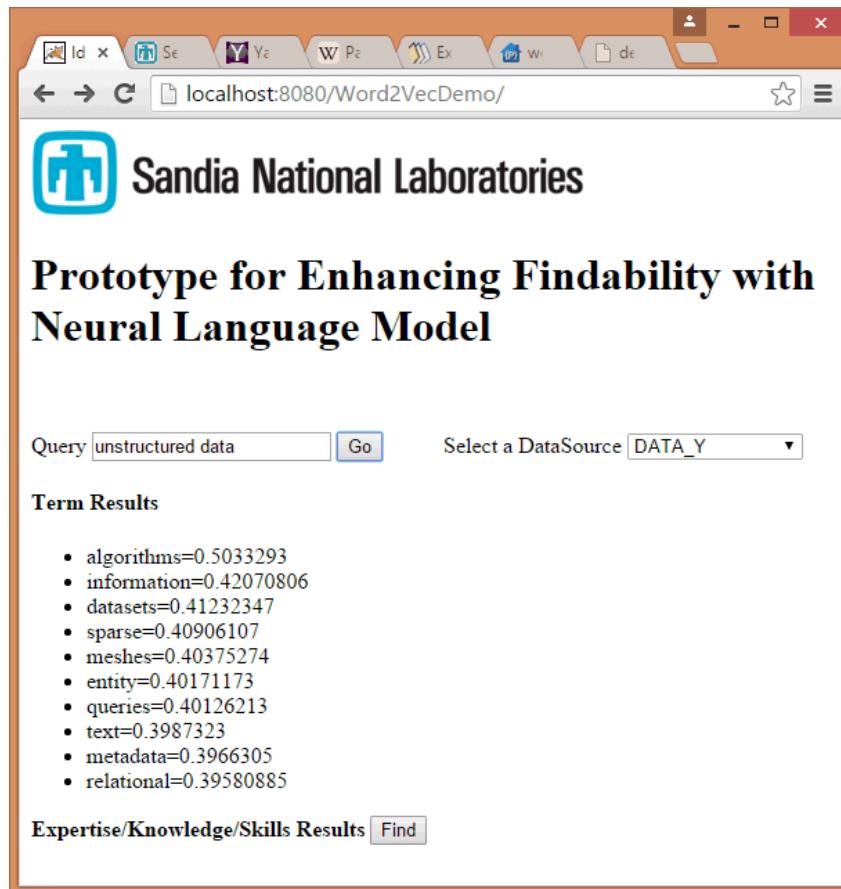




Challenges of Applying NNLM in Enterprise Search

- ◆ Domain Specific?
 - A NNLM trained from a domain may not be directly applied to other domains in a data repository with multiple domains
- ◆ Maintenance?
 - Data in an organization are updated dynamically, there is a need to update NNLM accordingly.

Differences among domains



localhost:8080/Word2VecDemo/

Sandia National Laboratories

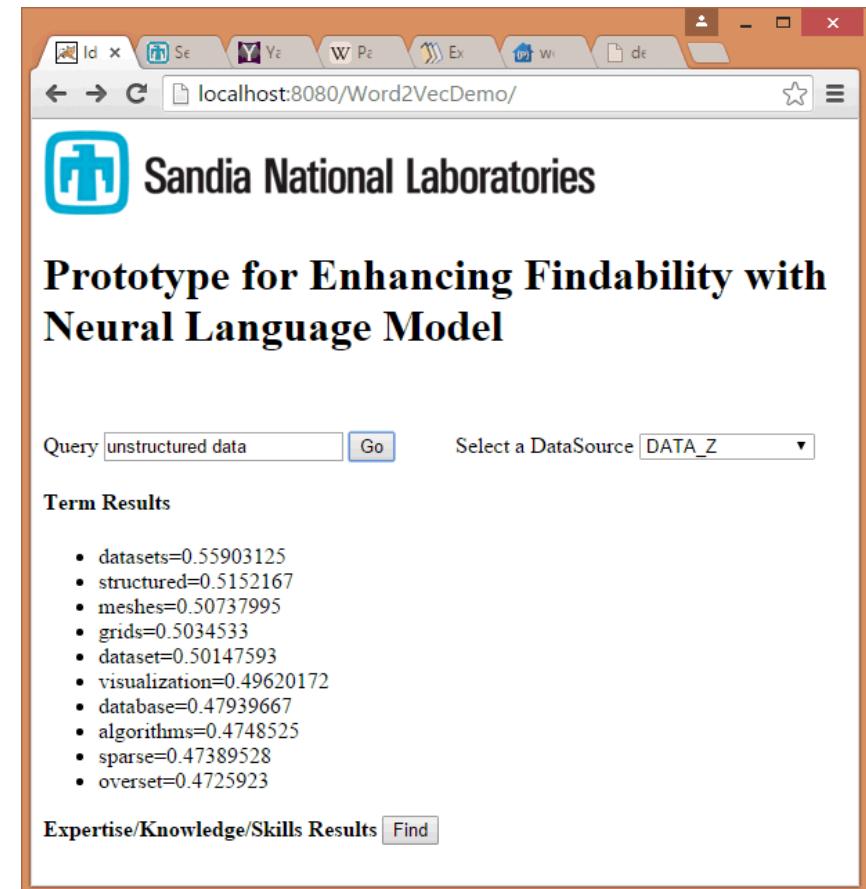
Prototype for Enhancing Findability with Neural Language Model

Query Select a DataSource

Term Results

- algorithms=0.5033293
- information=0.42070806
- datasets=0.41232347
- sparse=0.40906107
- meshes=0.40375274
- entity=0.40171173
- queries=0.40126213
- text=0.3987323
- metadata=0.3966305
- relational=0.39580885

Expertise/Knowledge/Skills Results



localhost:8080/Word2VecDemo/

Sandia National Laboratories

Prototype for Enhancing Findability with Neural Language Model

Query Select a DataSource

Term Results

- datasets=0.55903125
- structured=0.5152167
- meshes=0.50737995
- grids=0.5034533
- dataset=0.50147593
- visualization=0.49620172
- database=0.47939667
- algorithms=0.4748525
- sparse=0.47389528
- overset=0.4725923

Expertise/Knowledge/Skills Results





Conclusions and Further Efforts

- ◆ Enterprise search can significantly improved in:
 - Deliver relevant information even the query terms not appear in the enterprise dataset
 - Increase the quality of retrieved information
 - Expand retrieved information to closely related documents
- ◆ Implementation of NNLM in Enterprise Search requires minor search engine reconfiguration
- ◆ Efforts are needed to improve training a NNLM to cover more domains in an enterprise data repository
- ◆ Further developments are needed to update NNLMs in real time

