Chapter 01: Getting Started

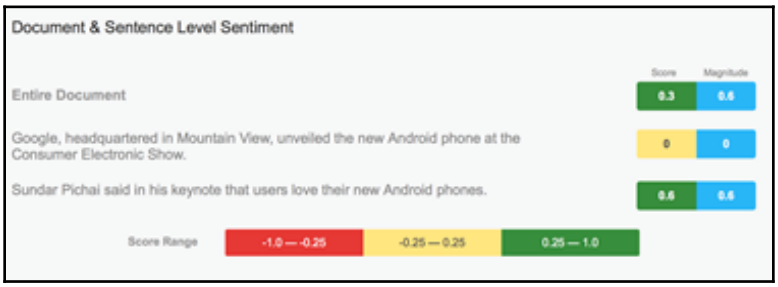

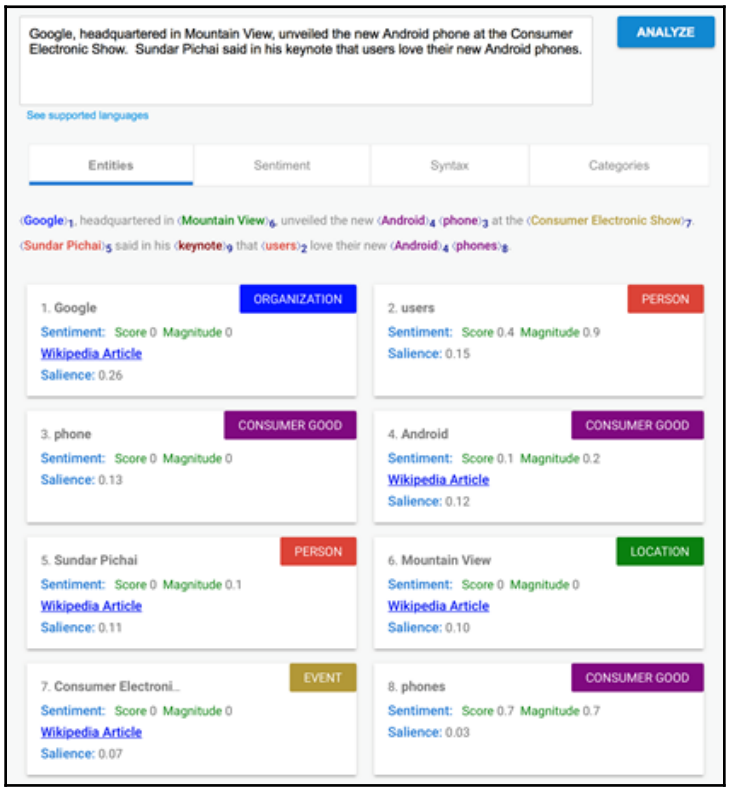

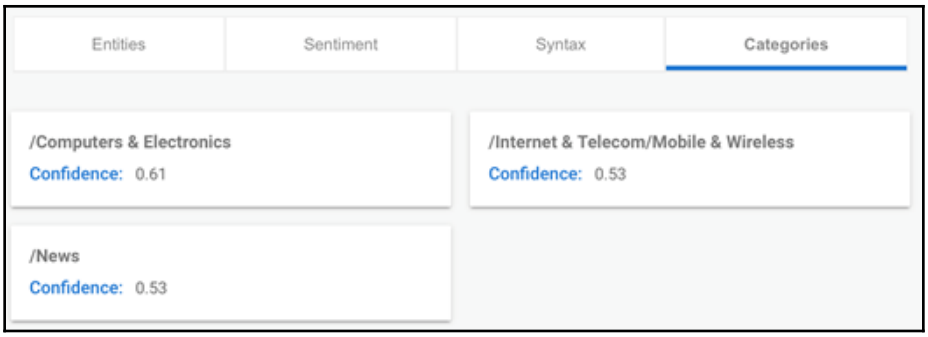
Entities
Sentiment
Syntax
Categories
/Computers & Electronics
Confidence: 0.61
/Internet & Telecom/Mobile & Wireless
Confidence: 0.53
/News
Confidence: 0.53

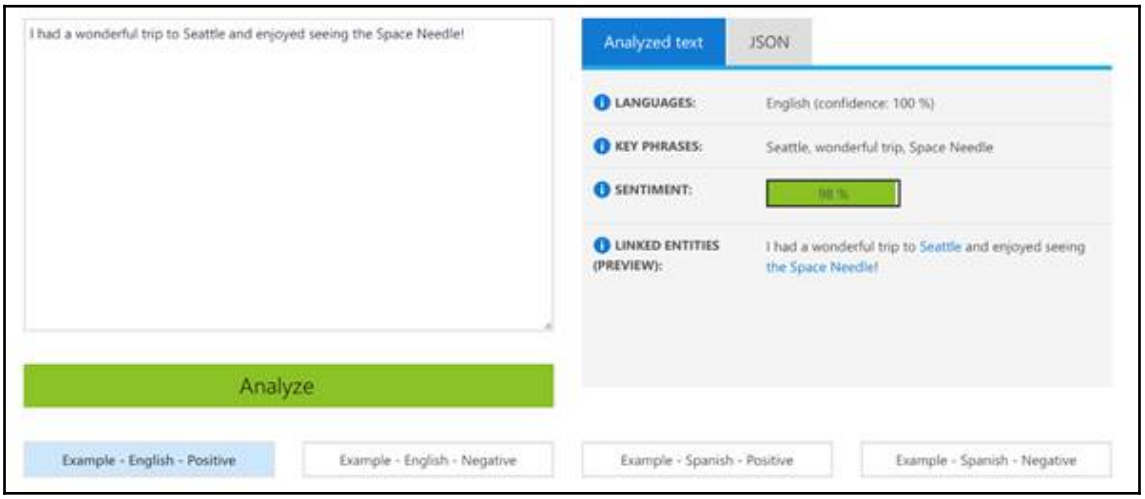
I had a wonderful trip to Seattle and enjoyed seeing the Space Needle!
Analyze
Example - English - Positive
Example - English - Negative
Example - Spanish - Positive
Example - Spanish - Negative
Analyzed text
JSON
LANGUAGES:
English (confidence: 100 %)
KEY PHRASES:
Seattle, wonderful trip, Space Needle
SENTIMENT:
LINKED ENTITIES (PREVIEW):
I had a wonderful trip to Seattle and enjoyed seeing the Space Needle!


Source Language
Chinese (Simplified) (zh)
Target Language
English (en)
Sample text. Enter your own text to translate.
古代書面汉语称为文言文，现代书面汉语一般指使用现代標準漢語語法、詞彙的中文通行文体（又称白话文）。
Ancient written Chinese is called classical Chinese. Modern written Chinese generally refers to the use of modern standard Chinese grammar and Chinese vocabulary (also known as vernacular Chinese).
Translate Text
Request URL
https://translation.googleapis.com/language/translate/v2/?q=%E5%8F%A4%E4%BB%A3%E6%9B%B8%E
JSON Response
{
"data": {
"translations": [
{
"translatedText": "Ancient written Chinese is called classical Chinese. Modern written Chinese generally refers to the use of modern standard Chinese grammar and Chinese vocabulary (also known as vernacular Chinese)."
}
]
}
}

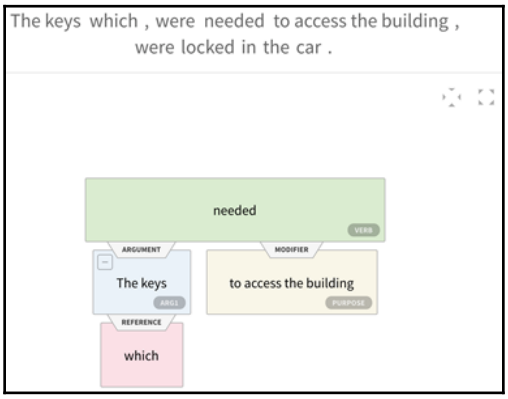
The keys which , were needed to access the building , were locked in the car .
needed
VERB
ARGUMENT
MODIFIER
The keys
ARG1
to access the building
PURPOSE
REFERENCE
which

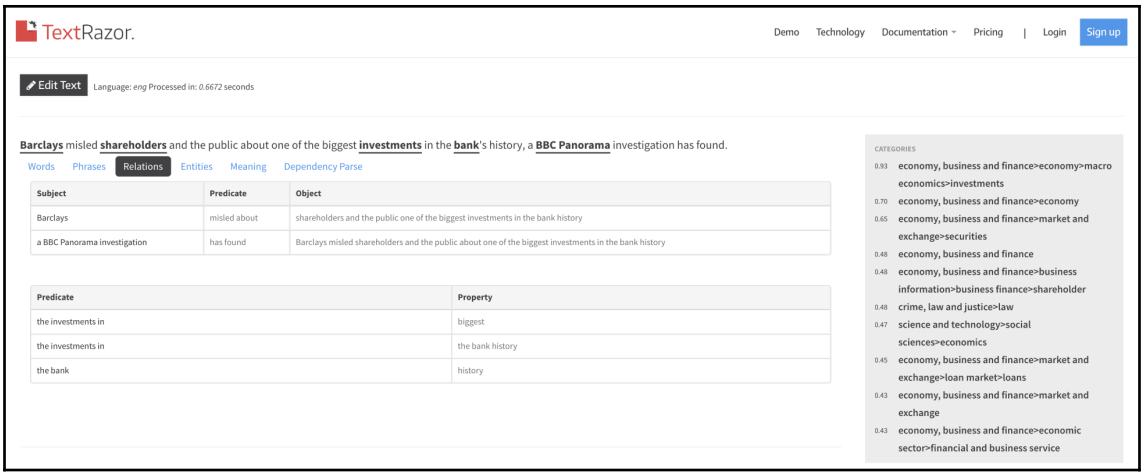
TextRazor.
Demo
Technology
Documentation
Pricing
Login
Sign up
Edit Text
Language: eng Processed in: 0.6672 seconds
Barclays misled shareholders and the public about one of the biggest investments in the bank's history, a BBC Panorama investigation has found.
Words
Phrases
Relations
Entities
Meaning
Dependency Parse
Subject
Predicate
Object
Barclays
misled about
shareholders and the public one of the biggest investments in the bank history
a BBC Panorama investigation
has found
Barclays misled shareholders and the public about one of the biggest investments in the bank history
Predicate
Property
the investments in
biggest
the investments in
the bank history
the bank
history
CATEGORIES
0.93 economy, business and finance>economy>macroeconomics>investments
0.70 economy, business and finance>economy
0.65 economy, business and finance>market and exchange>securities
0.48 economy, business and finance
0.48 economy, business and finance>business information>business finance>shareholder
0.48 crime, law and justice>law
0.47 science and technology>social sciences>economics
0.45 economy, business and finance>market and exchange>loan market>loans
0.43 economy, business and finance>market and exchange
0.43 economy, business and finance>economic sector>financial and business service

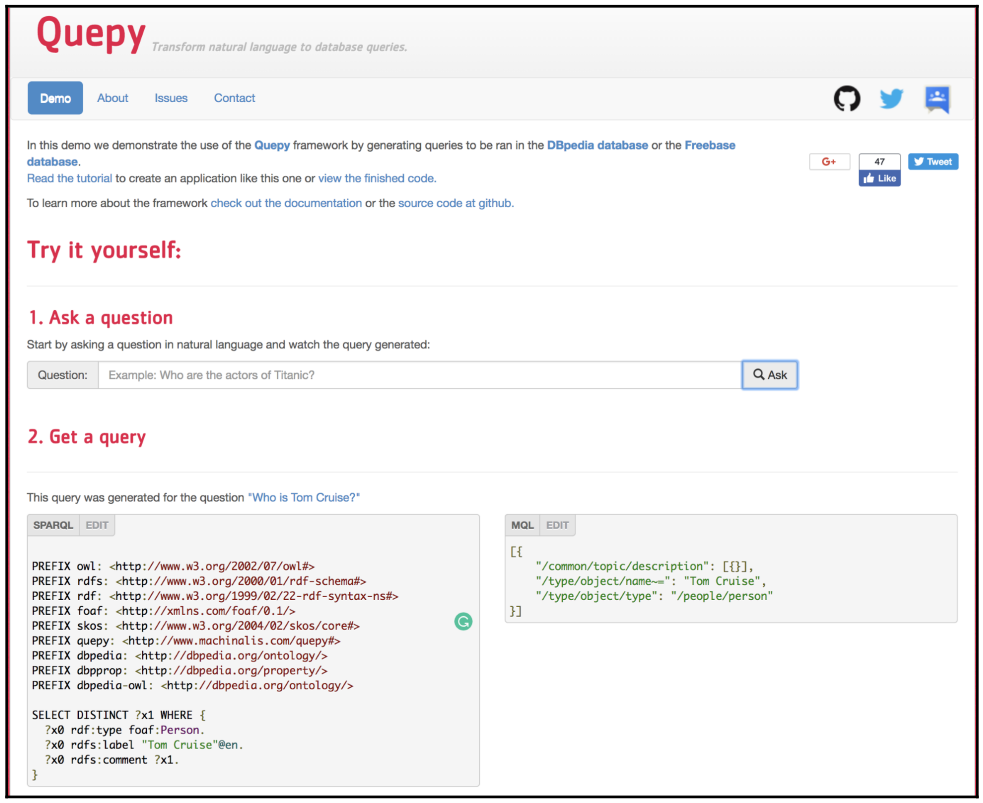
Quepy
Transform natural language to database queries.
Demo
About
Issues
Contact
In this demo we demonstrate the use of the Quepy framework by generating queries to be ran in the DBpedia database or the Freebase database.
Read the tutorial to create an application like this one or view the finished code.
To learn more about the framework check out the documentation or the source code at github.
47
Like
Tweet
Try it yourself:
1. Ask a question
Start by asking a question in natural language and watch the query generated:
Question:
Example: Who are the actors of Titanic?
Ask
2. Get a query
This query was generated for the question "Who is Tom Cruise?"
SPARQL
EDIT
PREFIX owl: <http://www.w3.org/2002/07/owl#>
PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>
PREFIX foaf: <http://xmlns.com/foaf/0.1/>
PREFIX skos: <http://www.w3.org/2004/02/skos/core#>
PREFIX quepy: <http://www.machinalis.com/quepy#>
PREFIX dbpedia: <http://dbpedia.org/ontology/>
PREFIX dbpprop: <http://dbpedia.org/property/>
PREFIX dbpedia-owl: <http://dbpedia.org/ontology/>
SELECT DISTINCT ?x1 WHERE {
?x0 rdf:type foaf:Person.
?x0 rdfs:label "Tom Cruise"@en.
?x0 rdfs:comment ?x1.
}
MQL
EDIT
[{
"/common/topic/description": [{}],
"/type/object/name~=": "Tom Cruise",
"/type/object/type": "/people/person"
}]

Answer

Keanu Reeves, Laurence Fishburne, Carrie-Anne Moss, Hugo Weaving, and Joe Pantoliano

Passage Context

The Matrix is a 1999 science fiction action film written and directed by The Wachowskis, starring **Keanu Reeves, Laurence Fishburne, Carrie-Anne Moss, Hugo Weaving, and Joe Pantoliano**. It depicts a dystopian future in which reality as perceived by most humans is actually a simulated reality called "the Matrix", created by sentient machines to subdue the human population, while their bodies' heat and electrical activity are used as an energy source. Computer programmer "Neo" learns this truth and is drawn into a rebellion against the machines, which involves other people who have been freed from the "dream world."

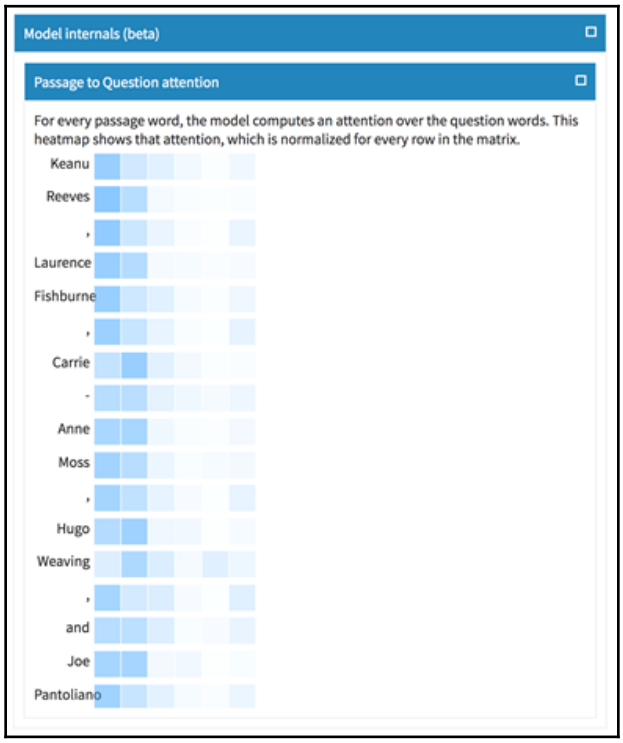

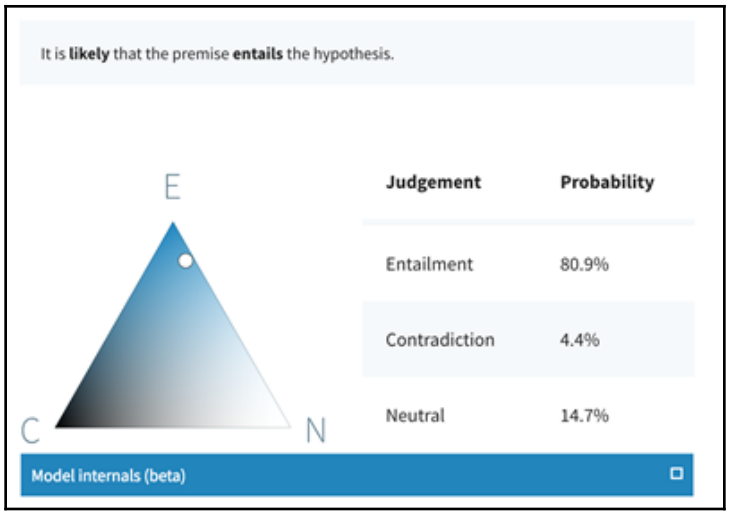
It is **likely** that the premise **entails** the hypothesis.
E
C
N
Judgement
Probability
Entailment
80.9%
Contradiction
4.4%
Neutral
14.7%
Model internals (beta)

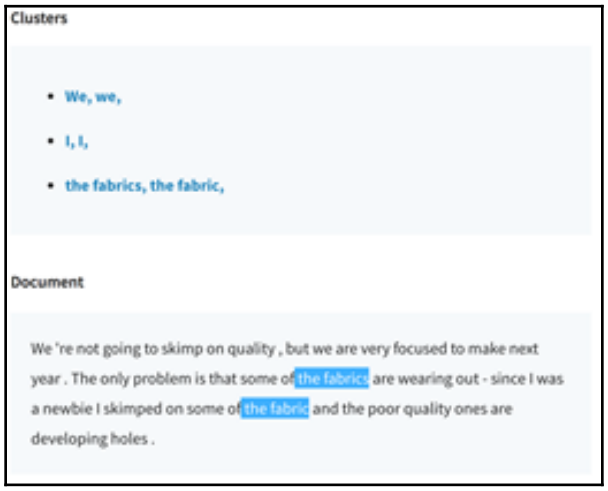
Clusters
We, we,
I, I,
the fabrics, the fabric,
Document
We 're not going to skimp on quality , but we are very focused to make next year . The only problem is that some of the fabrics are wearing out - since I was a newbie I skimped on some of the fabric and the poor quality ones are developing holes .

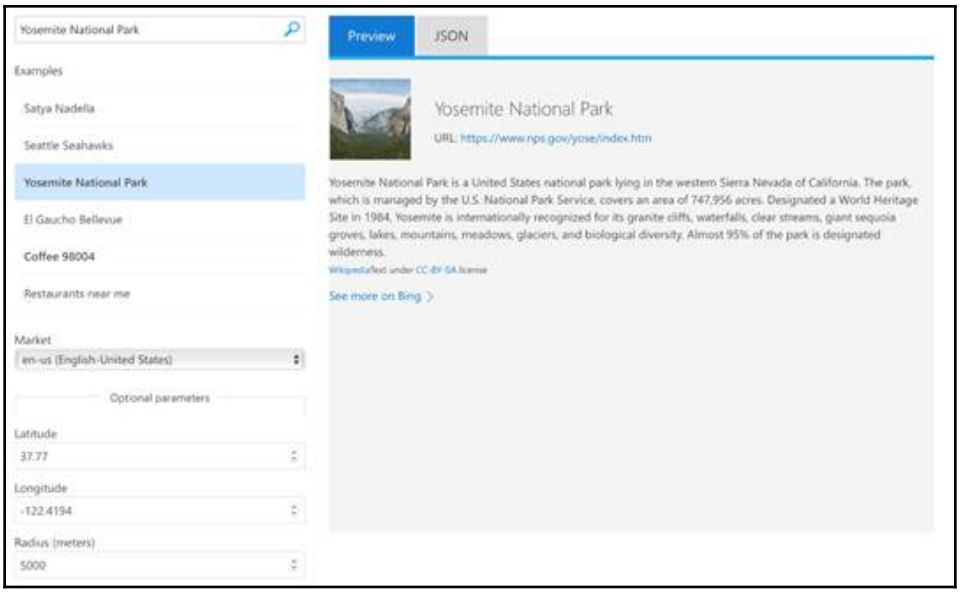
Yosemite National Park
Examples
Satya Nadella
Seattle Seahawks
Yosemite National Park
El Gaucho Bellevue
Coffee 98004
Restaurants near me
Market
en-us (English-United States)
Optional parameters
Latitude
37.77
Longitude
-122.4194
Radius (meters)
5000
Preview
JSON
Yosemite National Park
URL: https://www.nps.gov/yose/index.htm
Yosemite National Park is a United States national park lying in the western Sierra Nevada of California. The park, which is managed by the U.S. National Park Service, covers an area of 747,956 acres. Designated a World Heritage Site in 1984, Yosemite is internationally recognized for its granite cliffs, waterfalls, clear streams, giant sequoia groves, lakes, mountains, meadows, glaciers, and biological diversity. Almost 95% of the park is designated wilderness.
See more on Bing

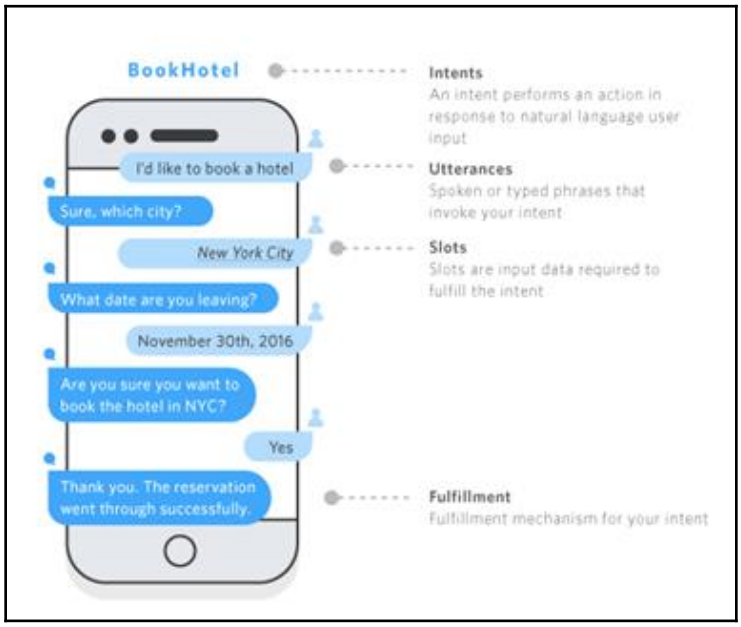
BookHotel
I'd like to book a hotel
Sure, which city?
New York City
What date are you leaving?
November 30th, 2016
Are you sure you want to book the hotel in NYC?
Yes
Thank you. The reservation went through successfully.
Intents
An intent performs an action in response to natural language user input
Utterances
Spoken or typed phrases that invoke your intent
Slots
Slots are input data required to fulfill the intent
Fulfillment
Fulfillment mechanism for your intent

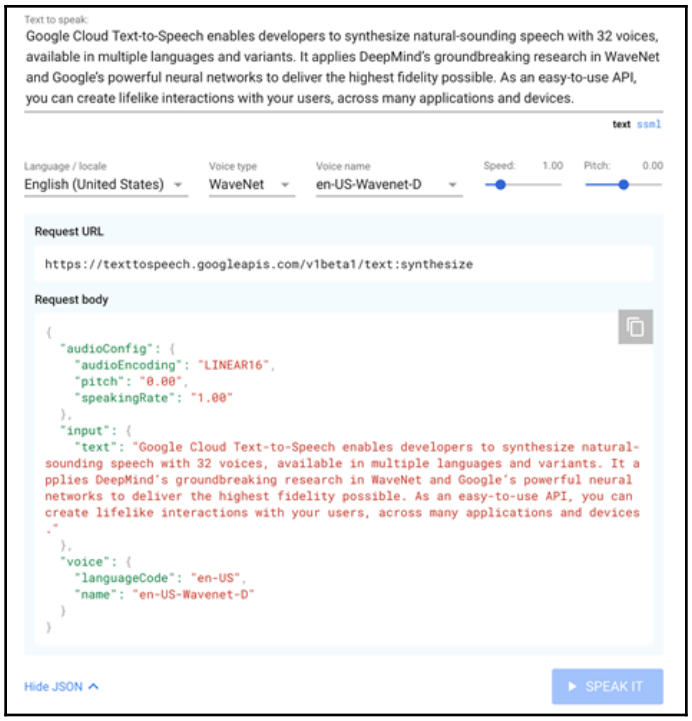
Text to speak:
Google Cloud Text-to-Speech enables developers to synthesize natural-sounding speech with 32 voices, available in multiple languages and variants. It applies DeepMind's groundbreaking research in WaveNet and Google's powerful neural networks to deliver the highest fidelity possible. As an easy-to-use API, you can create lifelike interactions with your users, across many applications and devices.
text ssml
Language / locale
English (United States)
Voice type
WaveNet
Voice name
en-US-Wavenet-D
Speed: 1.00
Pitch: 0.00
Request URL
https://texttospeech.googleapis.com/v1beta1/text:synthesize
Request body
{
"audioConfig": {
"audioEncoding": "LINEAR16",
"pitch": "0.00",
"speakingRate": "1.00"
},
"input": {
"text": "Google Cloud Text-to-Speech enables developers to synthesize natural-sounding speech with 32 voices, available in multiple languages and variants. It applies DeepMind's groundbreaking research in WaveNet and Google's powerful neural networks to deliver the highest fidelity possible. As an easy-to-use API, you can create lifelike interactions with your users, across many applications and devices."
},
"voice": {
"languageCode": "en-US",
"name": "en-US-Wavenet-D"
}
}
Hide JSON
SPEAK IT


Language
English (United States)
Punctuation
Input type
Microphone
File upload
Request URL
https://speech.googleapis.com/v1/speech:recognize
Request body
{
"audio": {
"content": "/* Your audio */"
},
"config": {
"enableAutomaticPunctuation": true,
"encoding": "LINEAR16",
"languageCode": "en-US",
"model": "default"
}
}
Hide JSON
START NOW

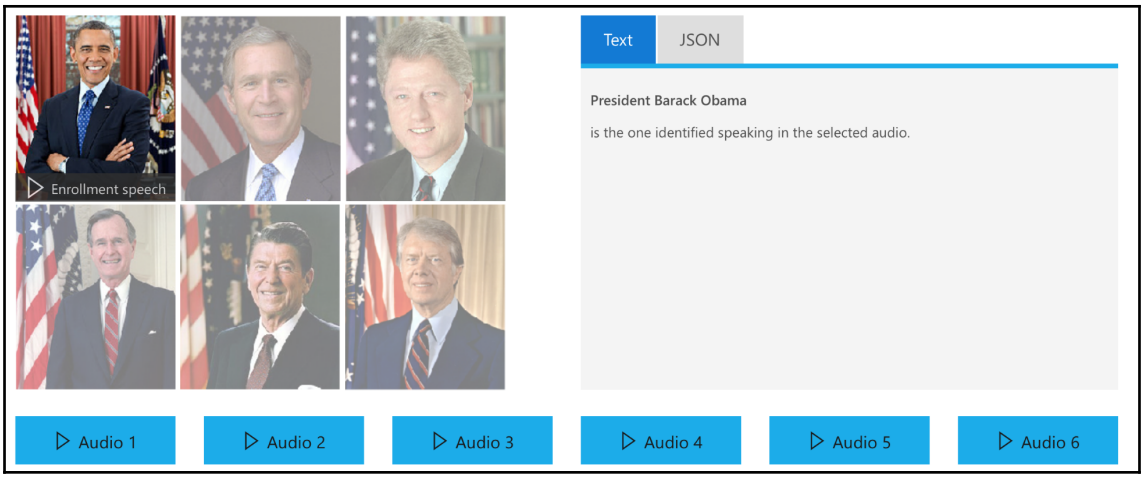
Enrollment speech
Text
JSON
President Barack Obama
is the one identified speaking in the selected audio.
Audio 1
Audio 2
Audio 3
Audio 4
Audio 5
Audio 6

Chapter 02: Text Classification and POS Tagging Using NLTK

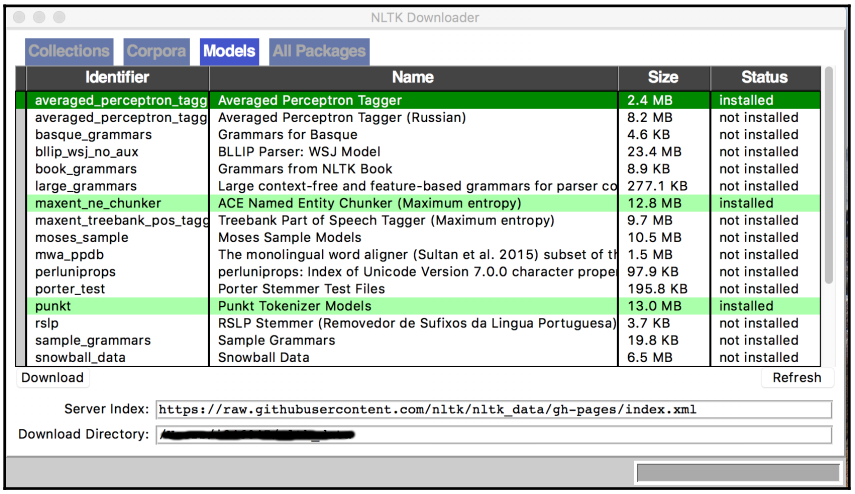

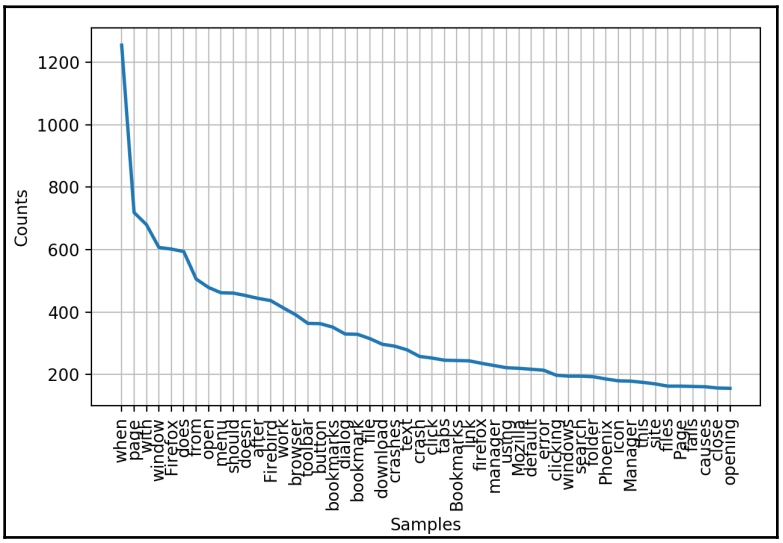


window
with
when
page
does
should
Firefox
Firebird
button
browser
after
from
doesn
menu
open
work
toolbar
dialog
text
download
firefox
bookmarks
crashes
crash


page
window
Firefox
Firebird
browser
dialog
link
open
menu
toolbar
tabs
folder
bookmark
download
file
click
Bookmarks
crashes
work
text
button
bookmarks
crash

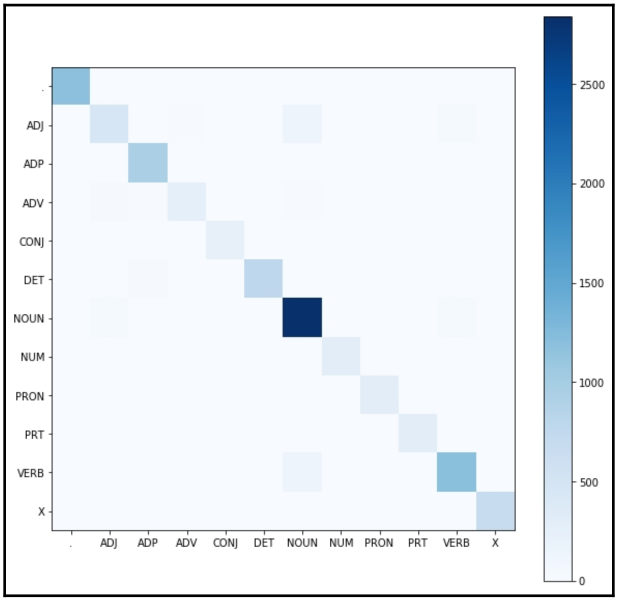
.
ADJ
ADP
ADV
CONJ
DET
NOUN
NUM
PRON
PRT
VERB
X
2500
2000
1500
1000
500
0

Chapter 03: Deep Learning and TensorFlow

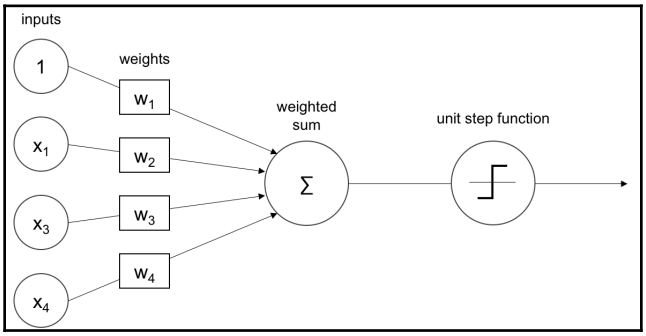

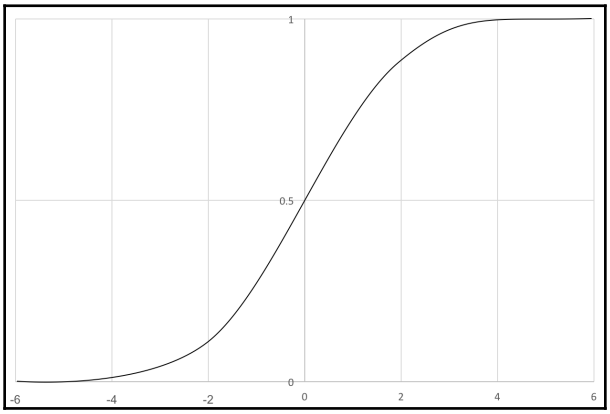

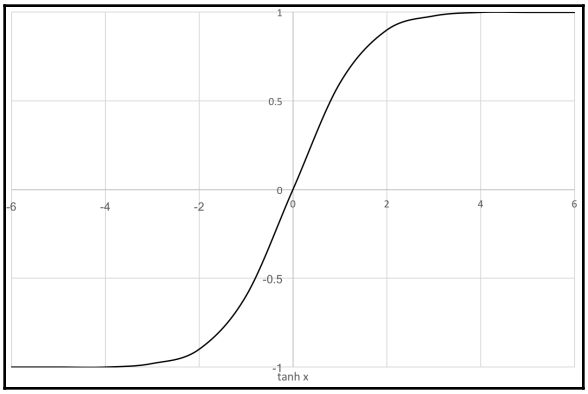

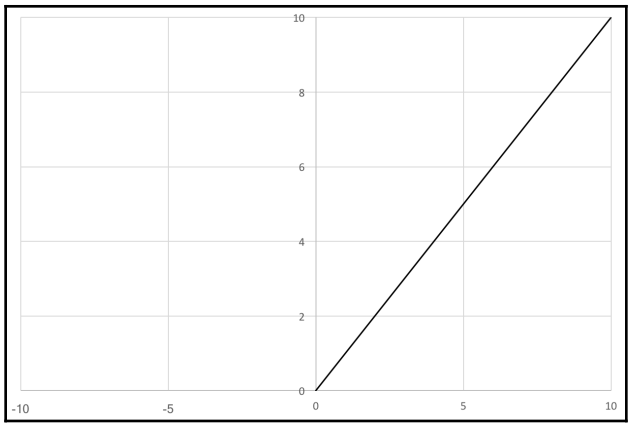
10
8
6
4
2
0
-10
-5
0
5
10

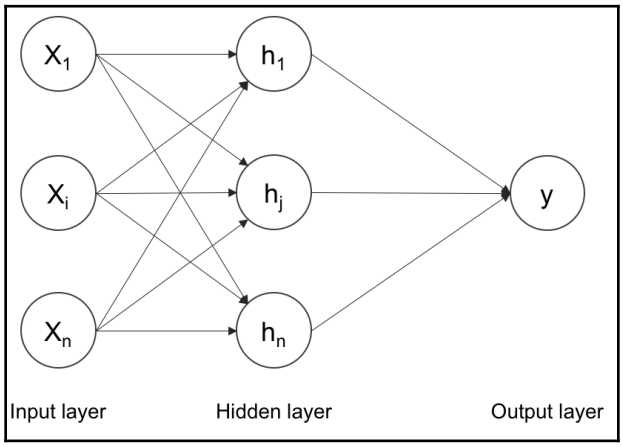
X1
Xi
Xn
h1
hj
hn
y
Input layer
Hidden layer
Output layer

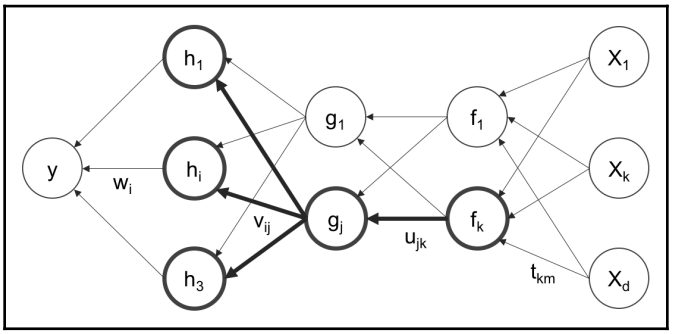
h1
hi
h3
y
wi
vij
gj
g1
fk
f1
ujk
tkm
X1
Xk
Xd

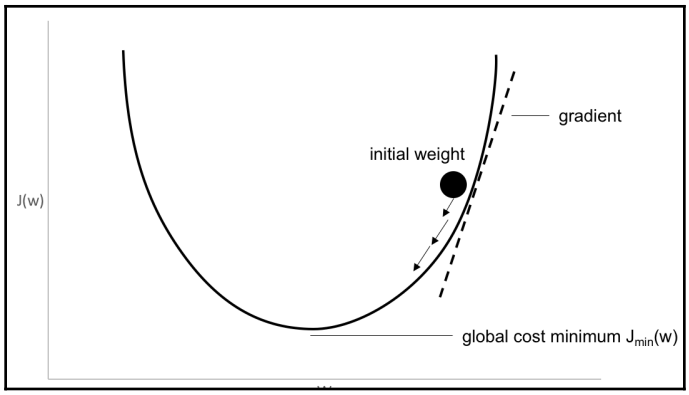
J(w)
gradient
initial weight
global cost minimum $J_{min}(w)$

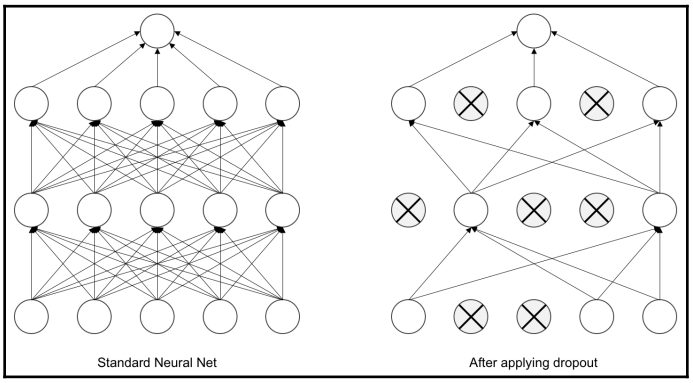
Standard Neural Net
After applying dropout

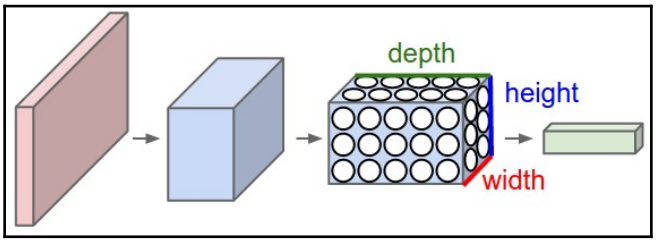
depth
height
width

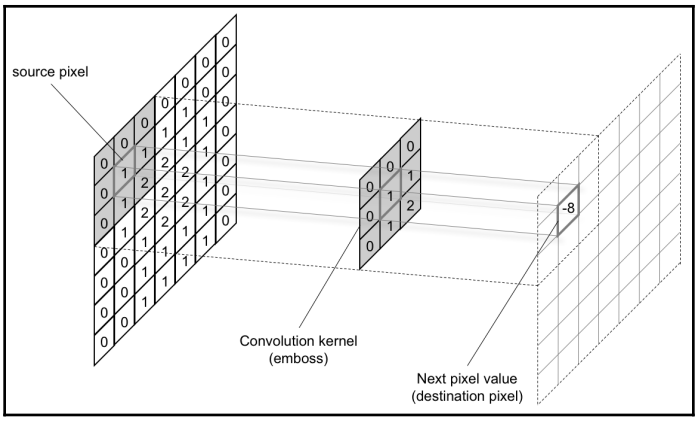
source pixel
Convolution kernel
(emboss)
Next pixel value
(destination pixel)
-8

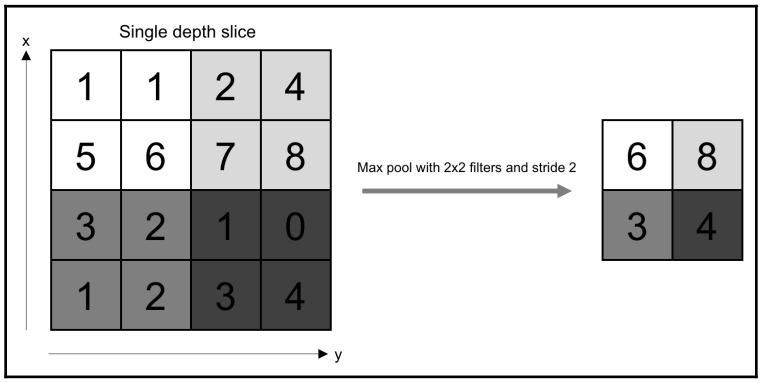
Single depth slice
x
y
1 1 2 4
5 6 7 8
3 2 1 0
1 2 3 4
Max pool with 2x2 filters and stride 2
6 8
3 4

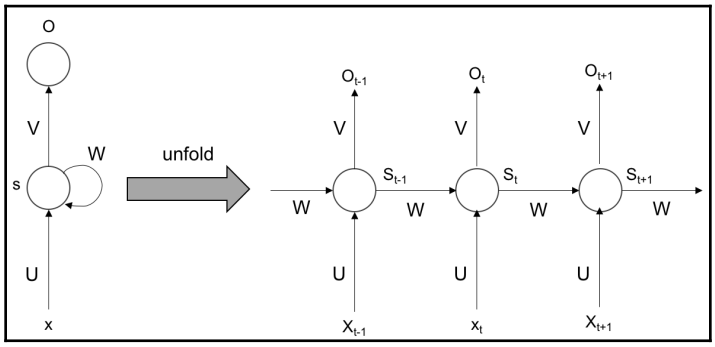
O
V
W
s
U
x
unfold
O_{t-1}
O_t
O_{t+1}
S_{t-1}
S_t
S_{t+1}
X_{t-1}
x_t
X_{t+1}

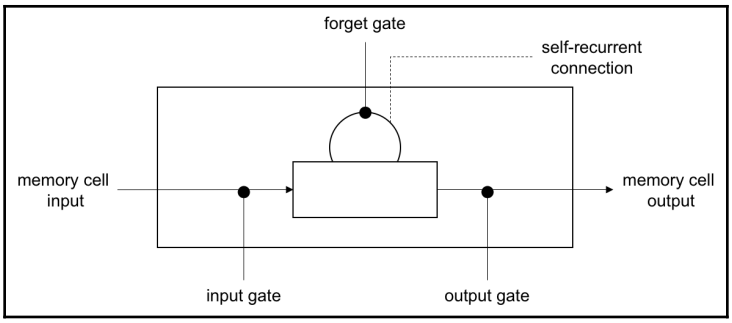

Select Target Platform

Click on the green buttons that describe your target platform. Only supported platforms will be shown.

Operating System	Windows · Linux · Mac OSX
Architecture	x86_64 · ppc64le
Distribution	Fedora · OpenSUSE · RHEL · CentOS · SLES · Ubuntu
Version	17.04 · 16.04
Installer Type	runfile (local) · deb (local) · deb (network) · cluster (local)

Download Installer for Linux Ubuntu 16.04 x86_64

The base installer is available for download below.

Base Installer — Download (2.8 KB)

Installation Instructions:

1. `sudo dpkg -i cuda-repo-ubuntu1604_9.0.176-1_amd64.deb`
2. `sudo apt-key adv --fetch-keys https://developer.download.nvidia.com/compute/cuda/repos/ubuntu1604/x86_64/7fa2af80.pub`
3. `sudo apt-get update`
4. `sudo apt-get install cuda`

Other installation options are available in the form of meta-packages. For example, to install all the library packages, replace "cuda" with the "cuda-libraries-9-0" meta package. For more information on all the available meta packages click here.

The CUDA Toolkit contains Open-Source Software. The source code can be found here.
The checksums for the installer and patches can be found in Installer Checksums.
For further information, see the Installation Guide for Linux and the CUDA Quick Start Guide.

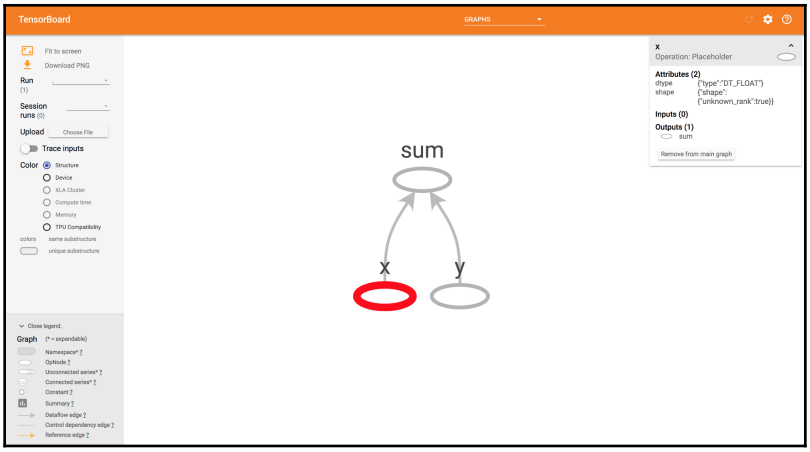
TensorBoard
GRAPHS
Fit to screen
Download PNG
Run
(1)
Session
runs (0)
Upload
Choose File
Trace inputs
Color
Structure
Device
XLA Cluster
Compute time
Memory
TPU Compatibility
colors
same substructure
unique substructure
Close legend.
Graph
(* = expandable)
Namespace* ?
OpNode ?
Unconnected series* ?
Connected series* ?
Constant ?
Summary ?
Dataflow edge ?
Control dependency edge ?
Reference edge ?
sum
x
y
x
Operation: Placeholder
Attributes (2)
dtype
{"type":"DT_FLOAT"}
shape
{"shape":
{"unknown_rank":true}}
Inputs (0)
Outputs (1)
sum
Remove from main graph

Chapter 04: Semantic Embedding Using Shallow Models

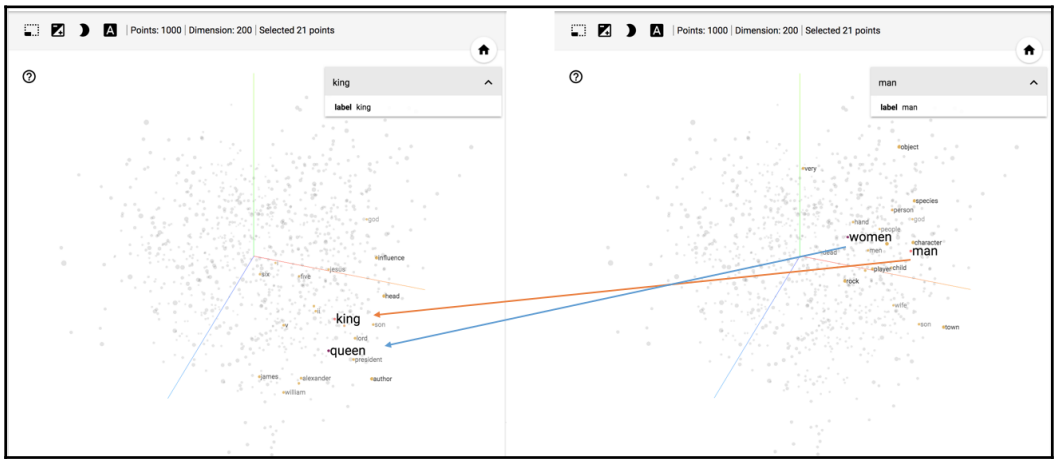

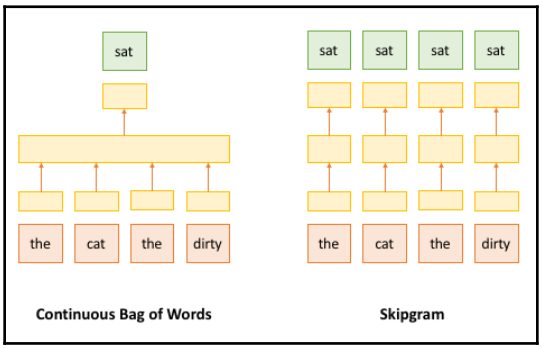

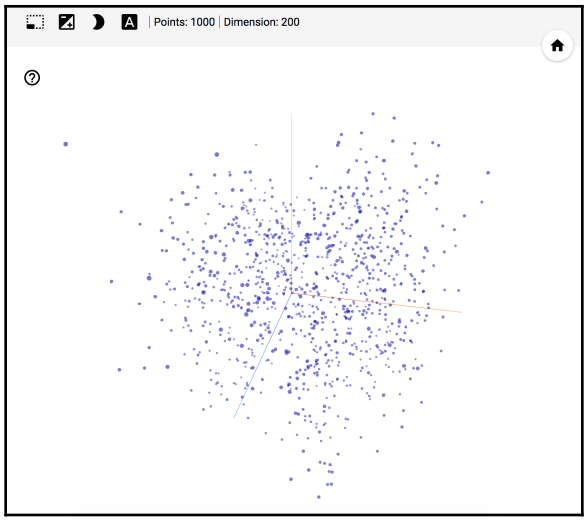
Points: 1000 | Dimension: 200

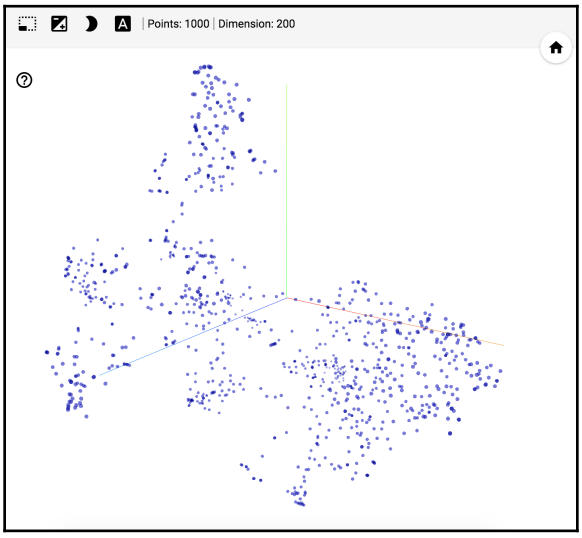
Points: 1000 | Dimension: 200

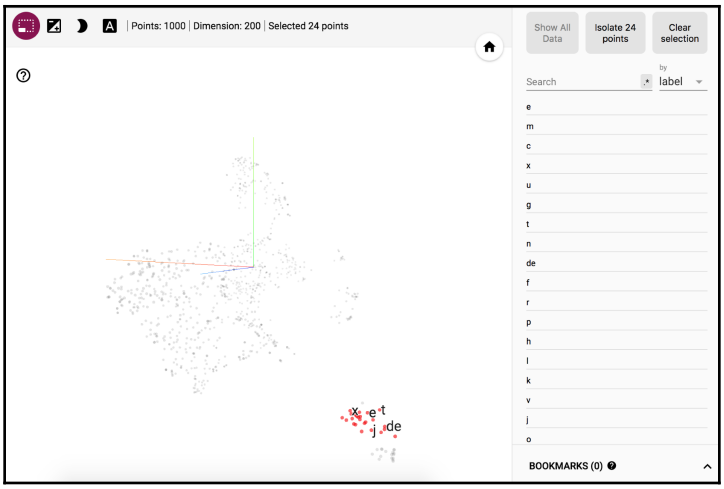
Points: 1000 | Dimension: 200 | Selected 24 points
Show All Data
Isolate 24 points
Clear selection
Search
by
label
e
m
c
x
u
g
t
n
de
f
r
p
h
l
k
v
j
o
BOOKMARKS (0)

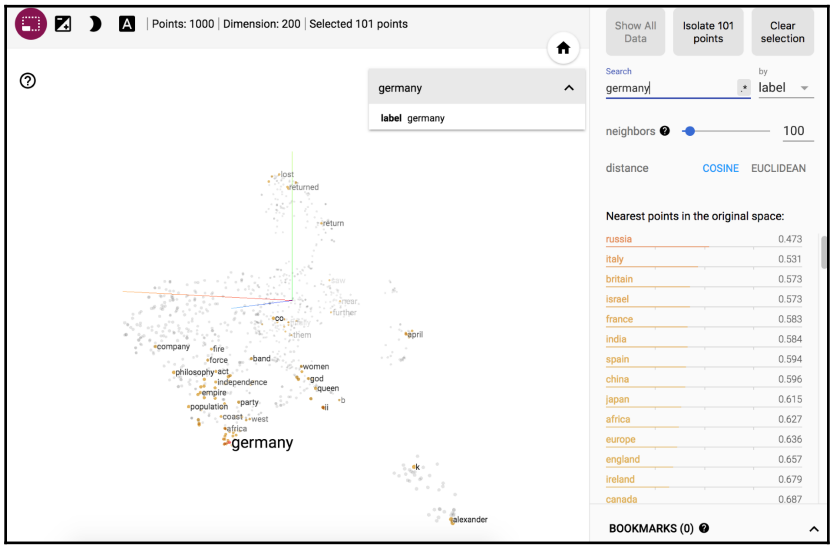
Points: 1000 | Dimension: 200 | Selected 101 points
Show All Data
Isolate 101 points
Clear selection
germany
label germany
Search
germany
by
label
neighbors
100
distance
COSINE
EUCLIDEAN
Nearest points in the original space:
russia 0.473
italy 0.531
britain 0.573
israel 0.573
france 0.583
india 0.584
spain 0.594
china 0.596
japan 0.615
africa 0.627
europe 0.636
england 0.657
ireland 0.679
canada 0.687
BOOKMARKS (0)
germany
alexander

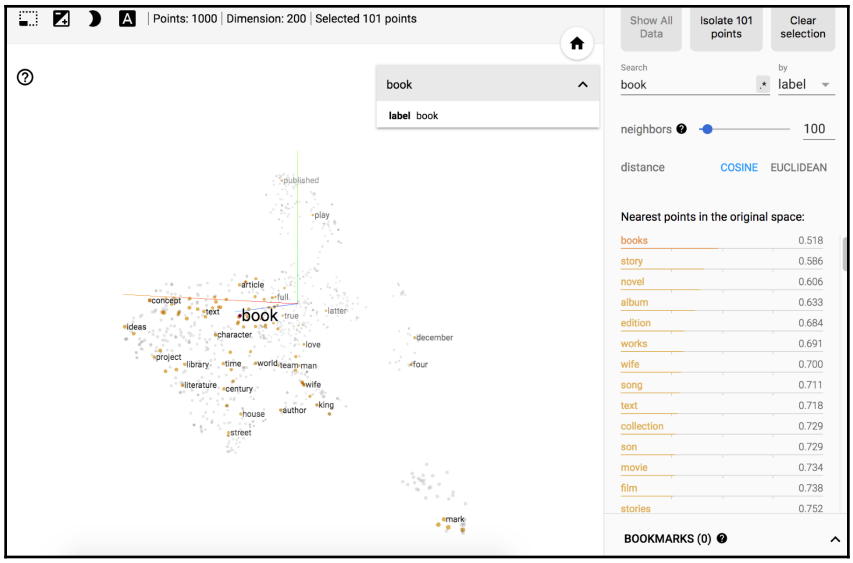
Points: 1000 | Dimension: 200 | Selected 101 points
Show All Data
Isolate 101 points
Clear selection
book
label book
Search
book
by
label
neighbors
100
distance
COSINE
EUCLIDEAN
Nearest points in the original space:
books 0.518
story 0.586
novel 0.606
album 0.633
edition 0.684
works 0.691
wife 0.700
song 0.711
text 0.718
collection 0.729
son 0.729
movie 0.734
film 0.738
stories 0.752
BOOKMARKS (0)
published
play
article
full
concept
text
book
true
latter
ideas
character
love
december
four
project
library
time
world
team
man
wife
literature
century
king
author
house
street
mark

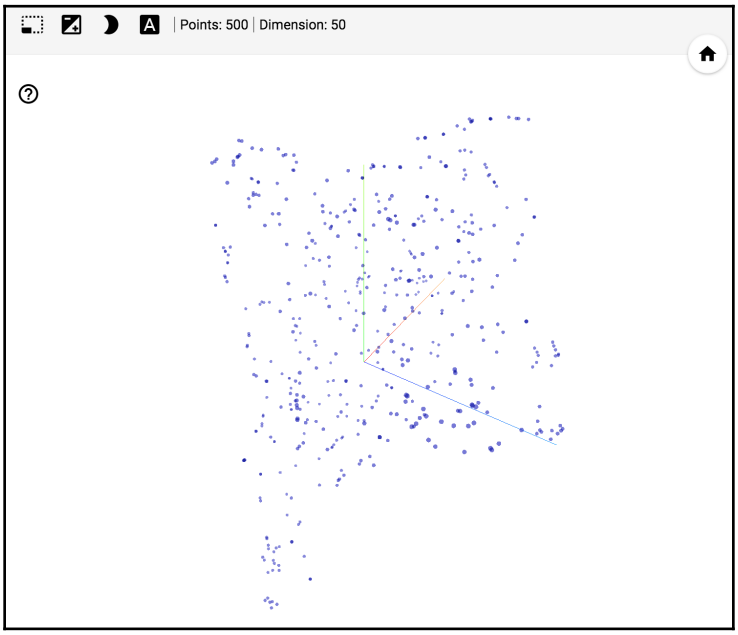
Points: 500 | Dimension: 50

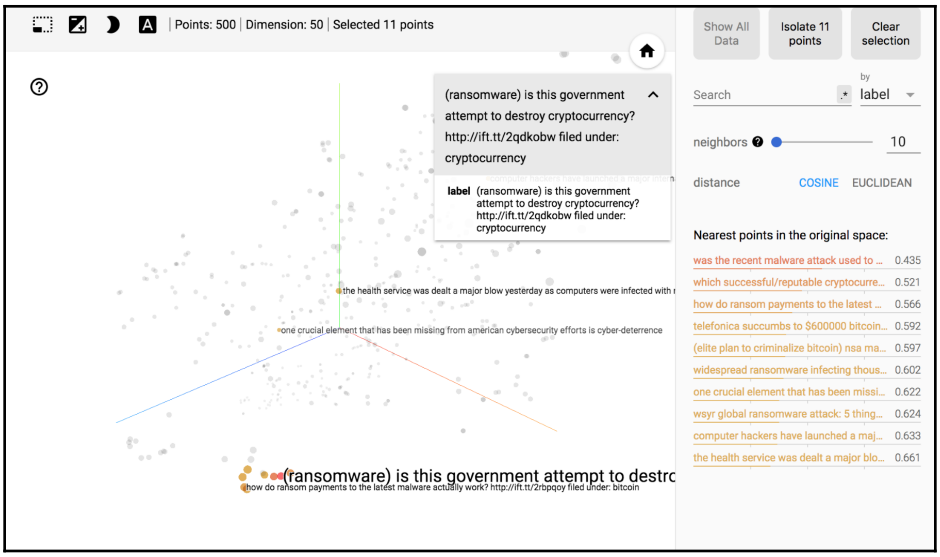
Points: 500 | Dimension: 50 | Selected 11 points
Show All Data
Isolate 11 points
Clear selection
(ransomware) is this government attempt to destroy cryptocurrency? http://ift.tt/2qdkobw filed under: cryptocurrency
label (ransomware) is this government attempt to destroy cryptocurrency? http://ift.tt/2qdkobw filed under: cryptocurrency
Search
by
label
neighbors
10
distance
COSINE
EUCLIDEAN
Nearest points in the original space:
was the recent malware attack used to ... 0.435
which successful/reputable cryptocurre... 0.521
how do ransom payments to the latest ... 0.566
telefonica succumbs to $600000 bitcoin... 0.592
(elite plan to criminalize bitcoin) nsa ma... 0.597
widespread ransomware infecting thous... 0.602
one crucial element that has been missi... 0.622
wsyr global ransomware attack: 5 thing... 0.624
computer hackers have launched a maj... 0.633
the health service was dealt a major blo... 0.661
the health service was dealt a major blow yesterday as computers were infected with
one crucial element that has been missing from american cybersecurity efforts is cyber-deterrence
(ransomware) is this government attempt to destro
how do ransom payments to the latest malware actually work? http://ift.tt/2rbpqoy filed under: bitcoin

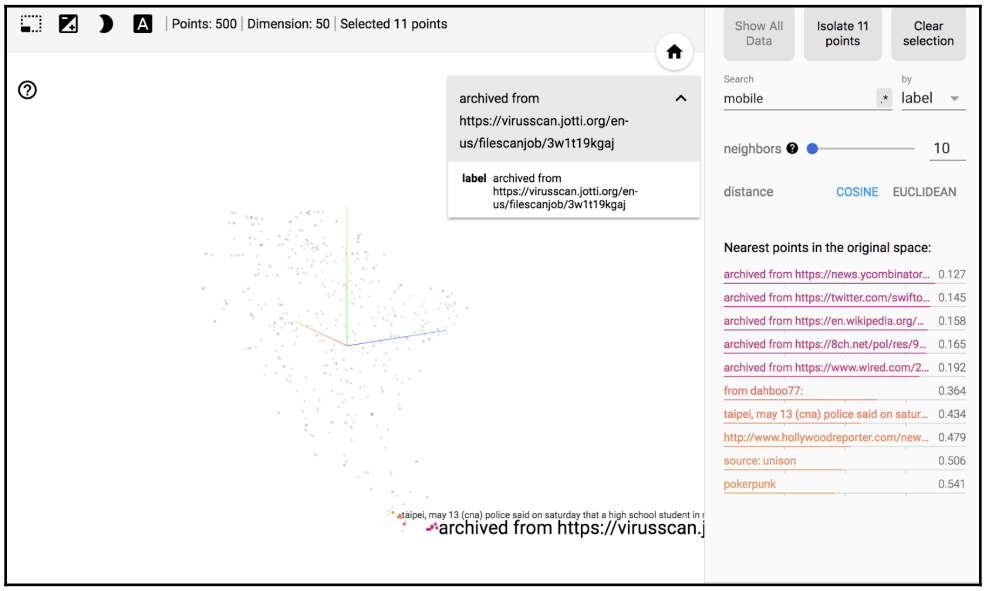
Points: 500 | Dimension: 50 | Selected 11 points
Show All Data
Isolate 11 points
Clear selection
archived from https://virusscan.jotti.org/en-us/filescanjob/3w1t19kgaj
label archived from https://virusscan.jotti.org/en-us/filescanjob/3w1t19kgaj
Search
mobile
by
label
neighbors
10
distance
COSINE
EUCLIDEAN
Nearest points in the original space:
archived from https://news.ycombinator... 0.127
archived from https://twitter.com/swifto... 0.145
archived from https://en.wikipedia.org/... 0.158
archived from https://8ch.net/pol/res/9... 0.165
archived from https://www.wired.com/2... 0.192
from dahboo77: 0.364
taipei, may 13 (cna) police said on satur... 0.434
http://www.hollywoodreporter.com/new... 0.479
source: unison 0.506
pokerpunk 0.541
taipei, may 13 (cna) police said on saturday that a high school student in
archived from https://virusscan.j

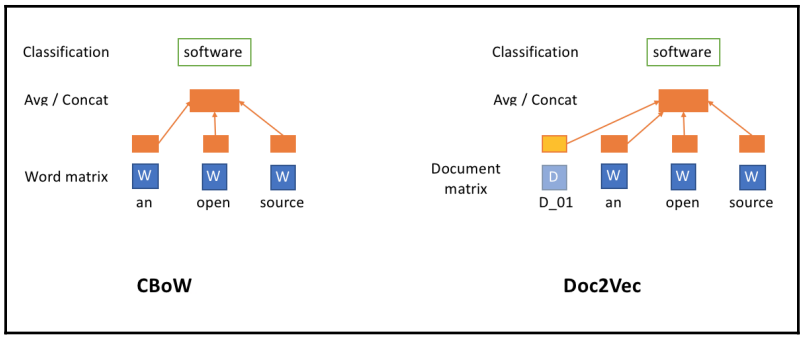
Classification
software
Avg / Concat
Word matrix
W
W
W
an
open
source
CBoW
Classification
software
Avg / Concat
Document matrix
D
W
W
W
D_01
an
open
source
Doc2Vec

Chapter 05: Text Classification Using LSTM

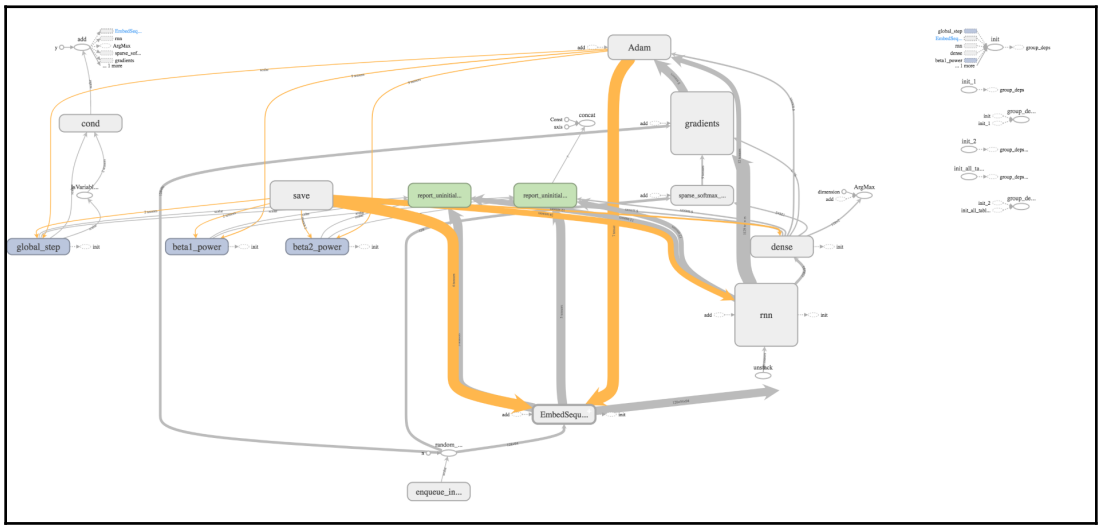

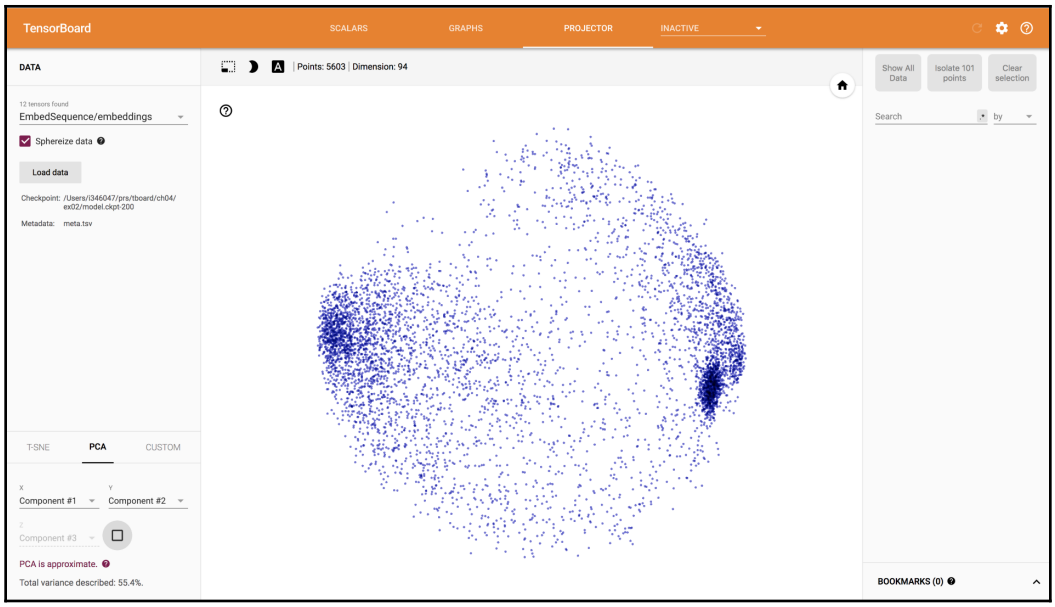
TensorBoard
SCALARS
GRAPHS
PROJECTOR
INACTIVE
DATA
12 tensors found
EmbedSequence/embeddings
Sphereize data
Load data
Checkpoint: /Users/l346047/prs/tboard/ch04/ex02/model.ckpt-200
Metadata: meta.tsv
Points: 5603 | Dimension: 94
Show All Data
Isolate 101 points
Clear selection
Search
by
T-SNE
PCA
CUSTOM
X
Component #1
Y
Component #2
Z
Component #3
PCA is approximate.
Total variance described: 55.4%.
BOOKMARKS (0)

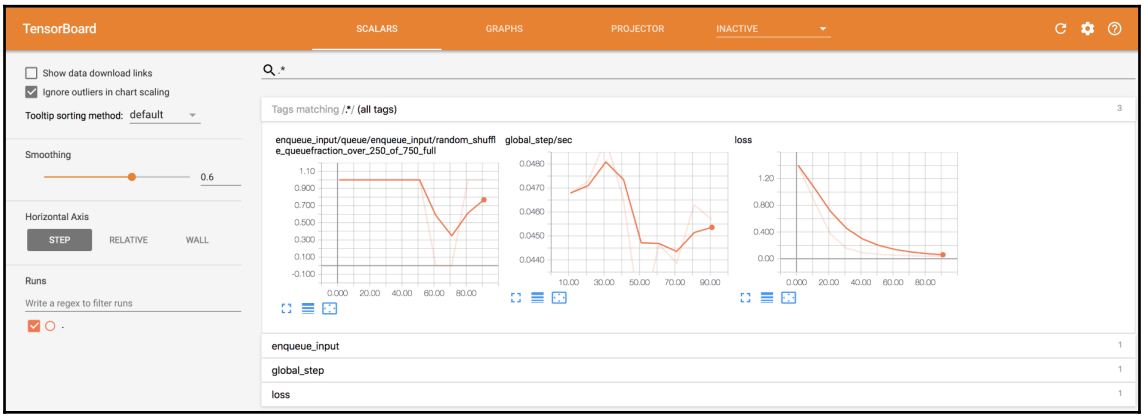
TensorBoard
SCALARS
GRAPHS
PROJECTOR
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method: default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
Tags matching /.*/ (all tags)
3
enqueue_input/queue/enqueue_input/random_shuffle_queuefraction_over_250_of_750_full
global_step/sec
loss
enqueue_input
global_step
loss

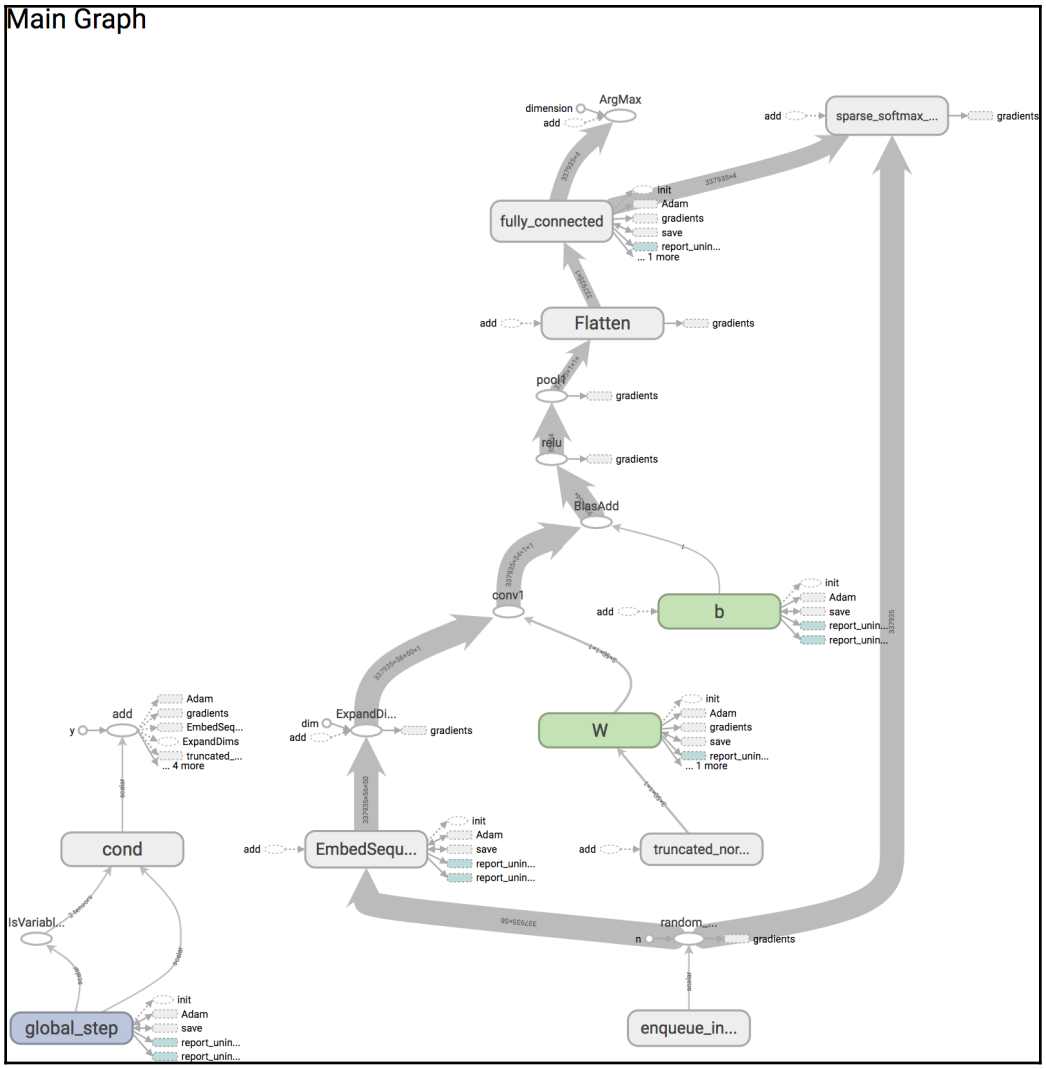
Main Graph
dimension
add
ArgMax
add
sparse_softmax_...
gradients
fully_connected
init
Adam
gradients
save
report_unin...
... 1 more
add
Flatten
gradients
gradients
gradients
BiasAdd
conv1
add
b
init
Adam
save
report_unin...
report_unin...
y
add
Adam
gradients
EmbedSeq...
ExpandDims
truncated_...
... 4 more
dim
add
ExpandDi...
gradients
W
init
Adam
gradients
save
report_unin...
... 1 more
cond
add
EmbedSequ...
init
Adam
save
report_unin...
report_unin...
add
truncated_nor...
IsVariabl...
random_...
n
gradients
global_step
init
Adam
save
report_unin...
report_unin...
enqueue_in...

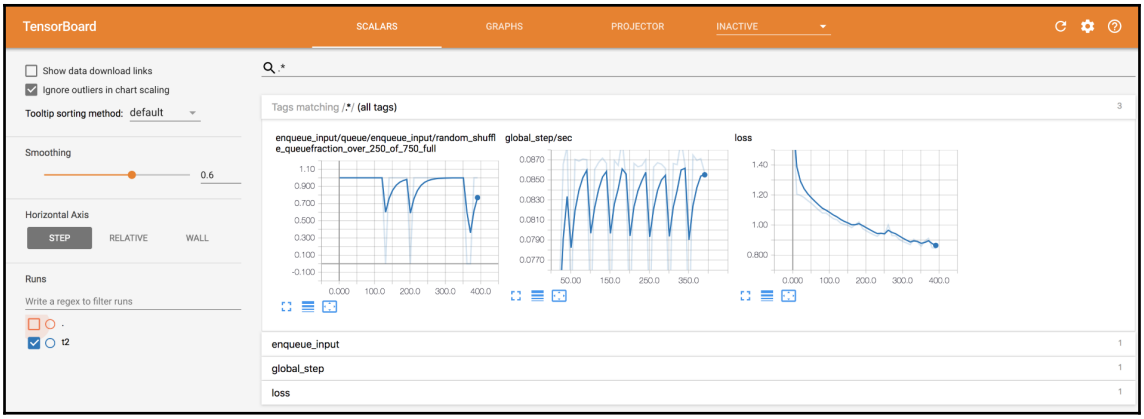
TensorBoard
SCALARS
GRAPHS
PROJECTOR
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method: default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
.
t2
.*
Tags matching /.*/ (all tags)
3
enqueue_input/queue/enqueue_input/random_shuffle_queuefraction_over_250_of_750_full
global_step/sec
loss
enqueue_input
1
global_step
1
loss
1

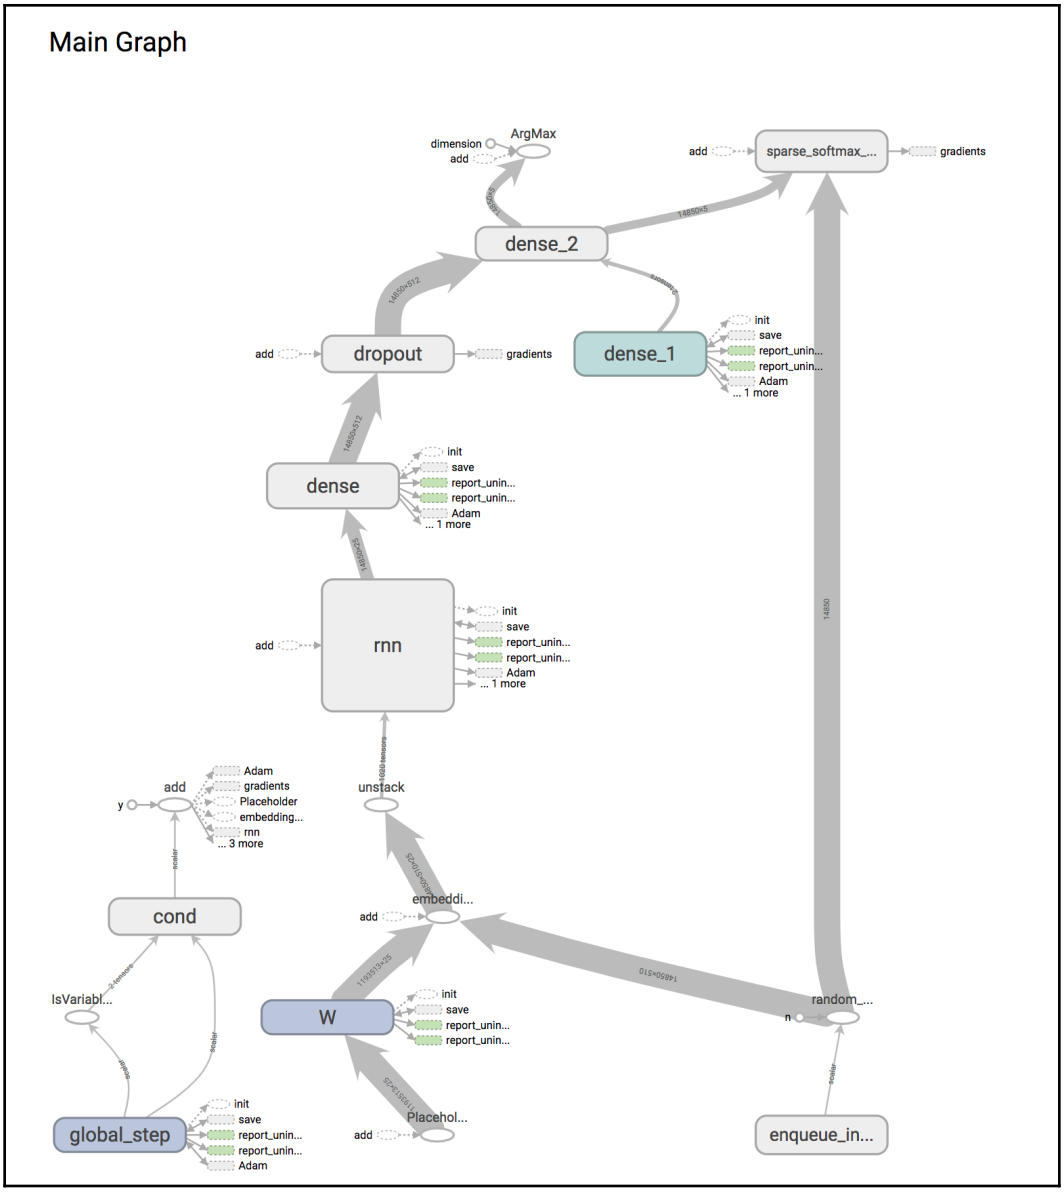
Main Graph
dimension
add
ArgMax
add
sparse_softmax_...
gradients
dense_2
14850×5
14850×512
add
dropout
gradients
dense_1
init
save
report_unin...
report_unin...
Adam
... 1 more
14850×512
dense
init
save
report_unin...
report_unin...
Adam
... 1 more
14850×512
add
rnn
init
save
report_unin...
report_unin...
Adam
... 1 more
14850
add
y
Adam
gradients
Placeholder
embedding...
rnn
... 3 more
unstack
14850 tensors
14850×512
add
embeddi...
cond
1×512×25
14850×510
IsVariabl...
W
init
save
report_unin...
report_unin...
random_...
n
control
control
control
global_step
init
save
report_unin...
report_unin...
Adam
512×25
add
Placehol...
control
enqueue_in...

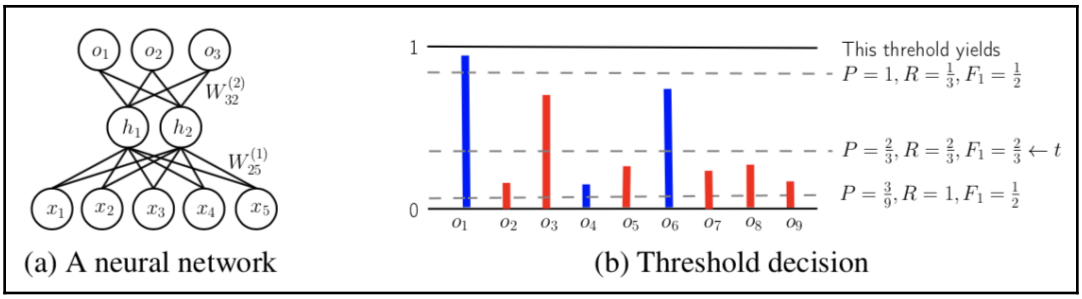
o_1 o_2 o_3
$W_{32}^{(2)}$
h_1 h_2
$W_{25}^{(1)}$
x_1 x_2 x_3 x_4 x_5
(a) A neural network
1
0
This threhold yields
$P = 1, R = \frac{1}{3}, F_1 = \frac{1}{2}$
$P = \frac{2}{3}, R = \frac{2}{3}, F_1 = \frac{2}{3} \leftarrow t$
$P = \frac{3}{9}, R = 1, F_1 = \frac{1}{2}$
o_1 o_2 o_3 o_4 o_5 o_6 o_7 o_8 o_9
(b) Threshold decision

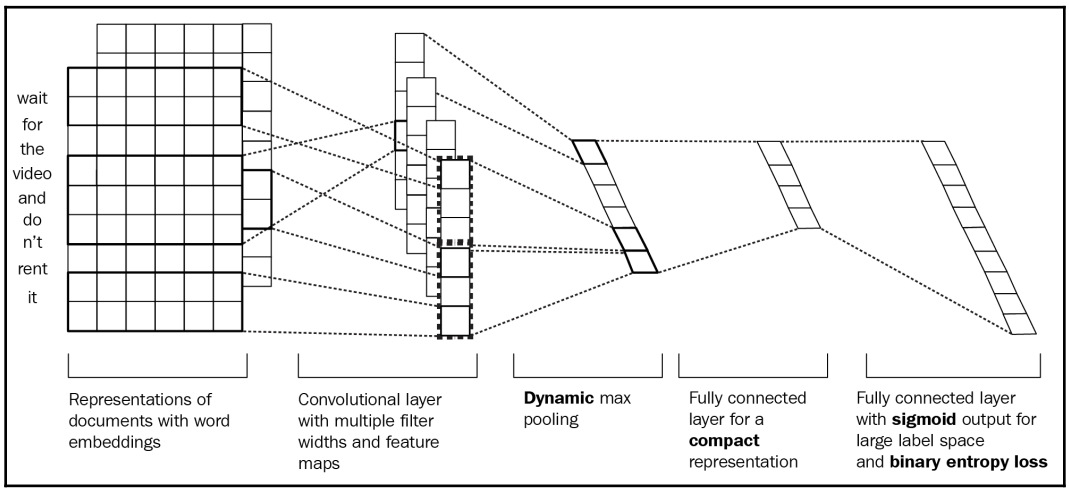
wait
for
the
video
and
do
n't
rent
it
Representations of documents with word embeddings
Convolutional layer with multiple filter widths and feature maps
Dynamic max pooling
Fully connected layer for a **compact** representation
Fully connected layer with **sigmoid** output for large label space and **binary entropy loss**

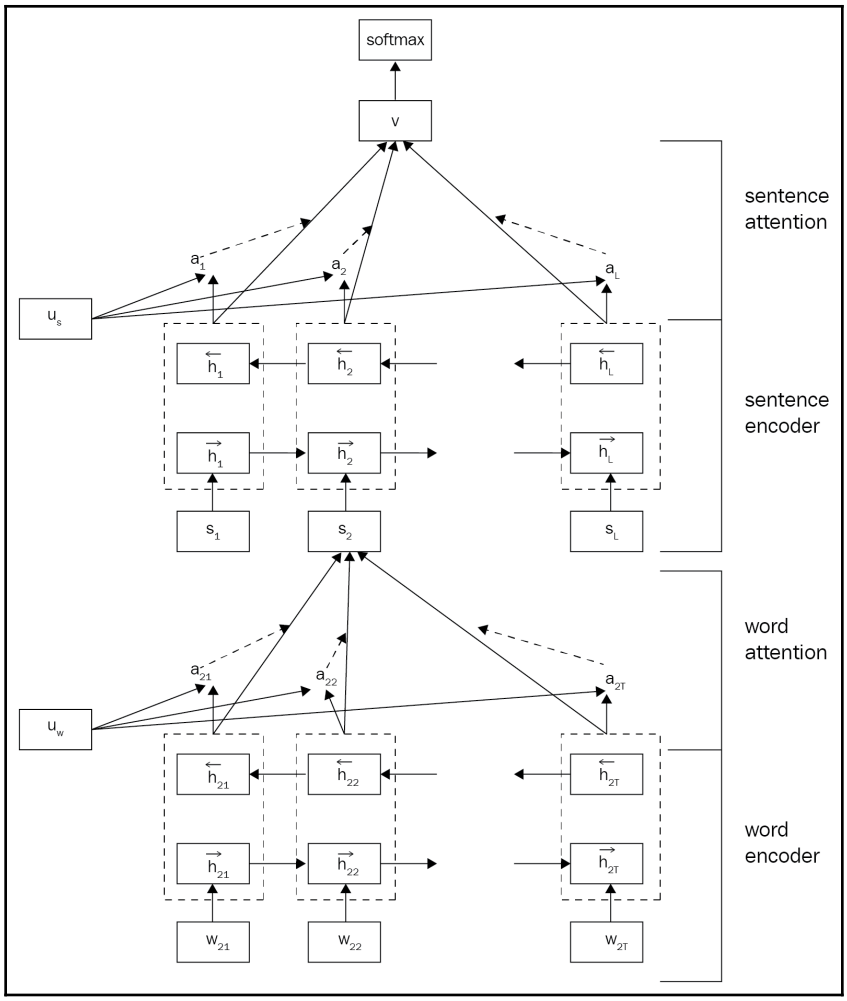
softmax
v
sentence attention
sentence encoder
word attention
word encoder
u_s
u_w
a_1
a_2
a_L
a_{21}
a_{22}
a_{2T}
$\overleftarrow{h}_1$
$\overrightarrow{h}_1$
$\overleftarrow{h}_2$
$\overrightarrow{h}_2$
$\overleftarrow{h}_L$
$\overrightarrow{h}_L$
s_1
s_2
s_L
$\overleftarrow{h}_{21}$
$\overrightarrow{h}_{21}$
$\overleftarrow{h}_{22}$
$\overrightarrow{h}_{22}$
$\overleftarrow{h}_{2T}$
$\overrightarrow{h}_{2T}$
w_{21}
w_{22}
w_{2T}

Chapter 07: Named Entity Recognition Using Character LSTM

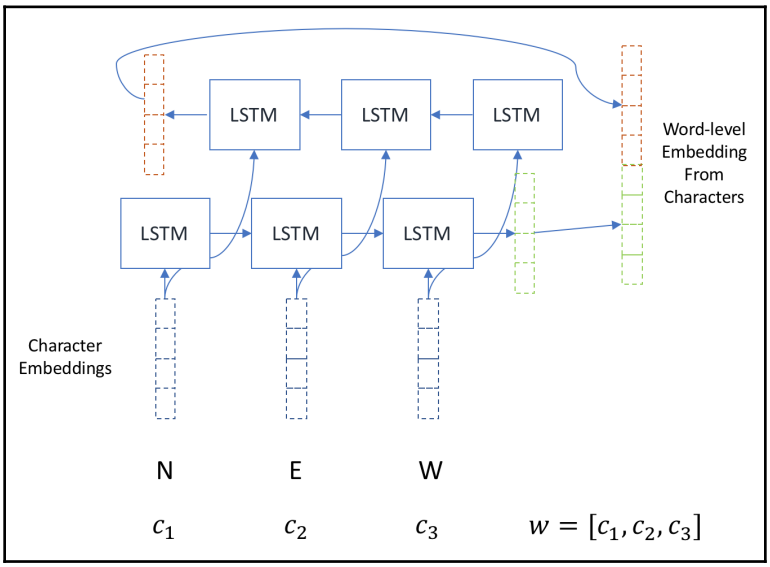

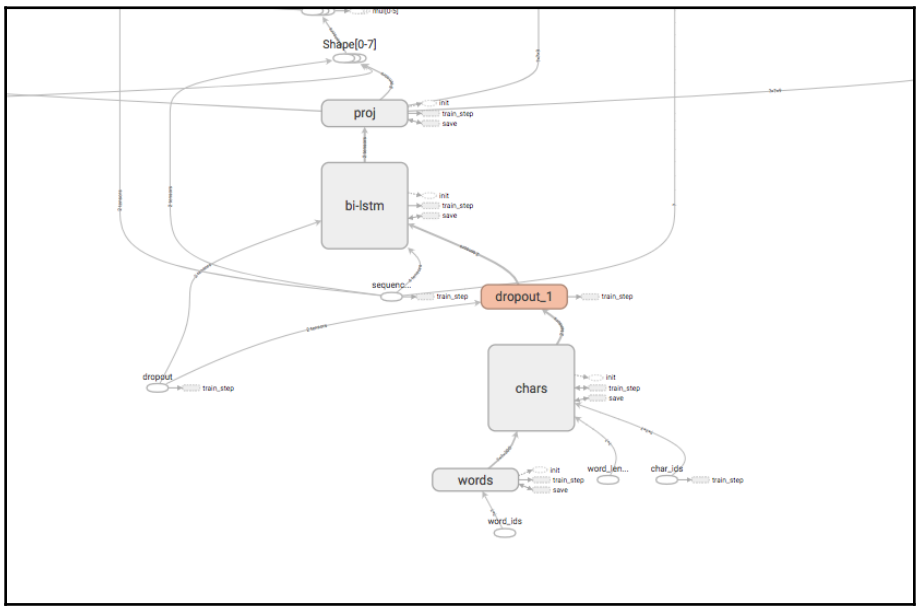

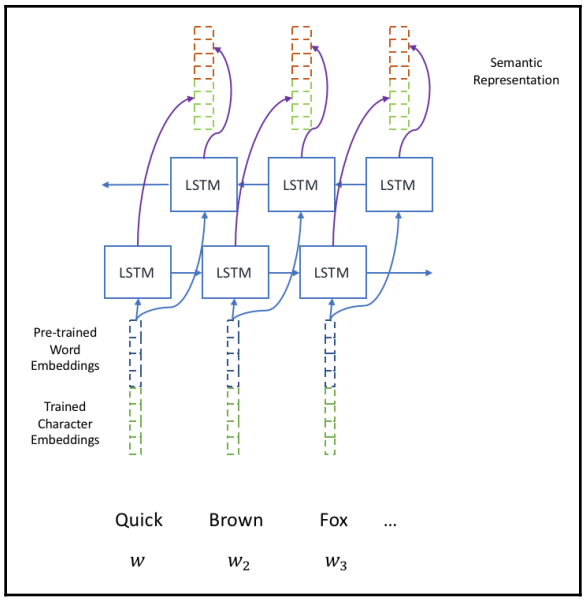
Semantic
Representation
LSTM
LSTM
LSTM
LSTM
LSTM
LSTM
Pre-trained
Word
Embeddings
Trained
Character
Embeddings
Quick
Brown
Fox
...
w
w_2
w_3

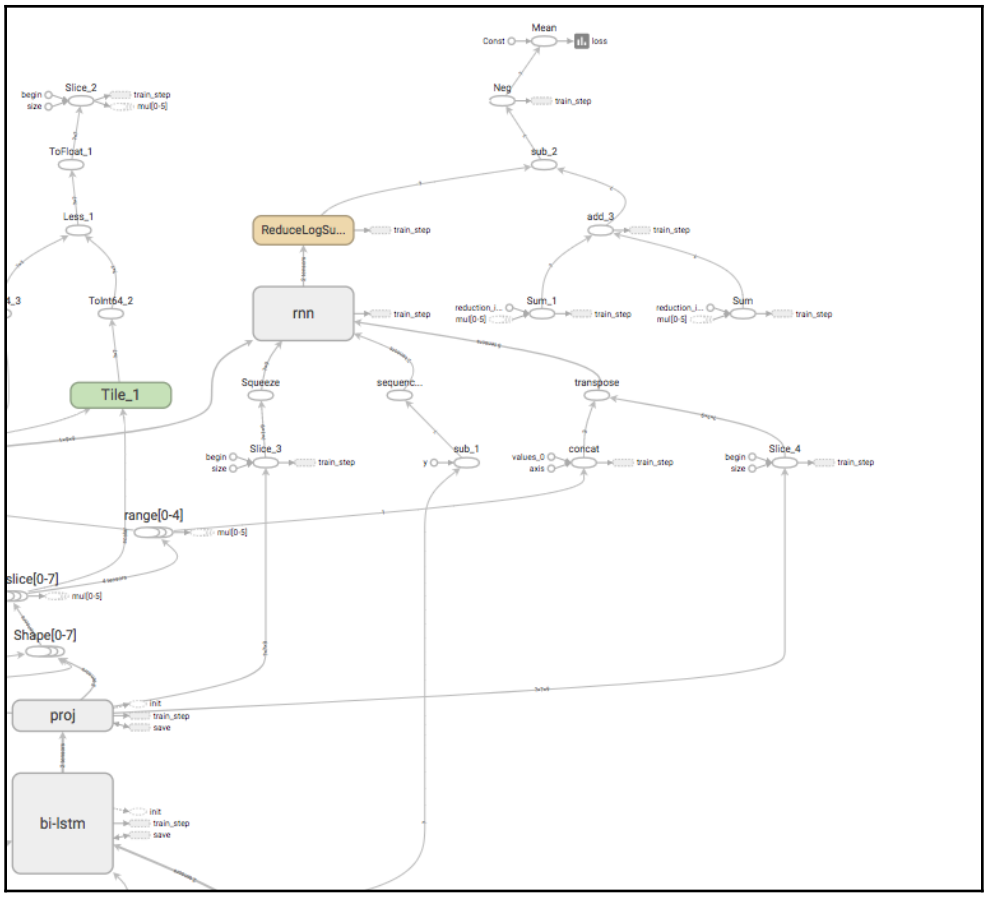
Mean
Const
loss
Neg
train_step
Slice_2
begin
size
train_step
mul[0-5]
sub_2
ToFloat_1
Less_1
ReduceLogSu...
train_step
add_3
train_step
ToInt64_2
rnn
train_step
reduction_i...
mul[0-5]
Sum_1
train_step
reduction_i...
mul[0-5]
Sum
train_step
Tile_1
Squeeze
sequenc...
transpose
Slice_3
begin
size
train_step
sub_1
y
values_0
axis
concat
train_step
Slice_4
begin
size
train_step
range[0-4]
mul[0-5]
slice[0-7]
mul[0-5]
Shape[0-7]
proj
init
train_step
save
bi-lstm
init
train_step
save

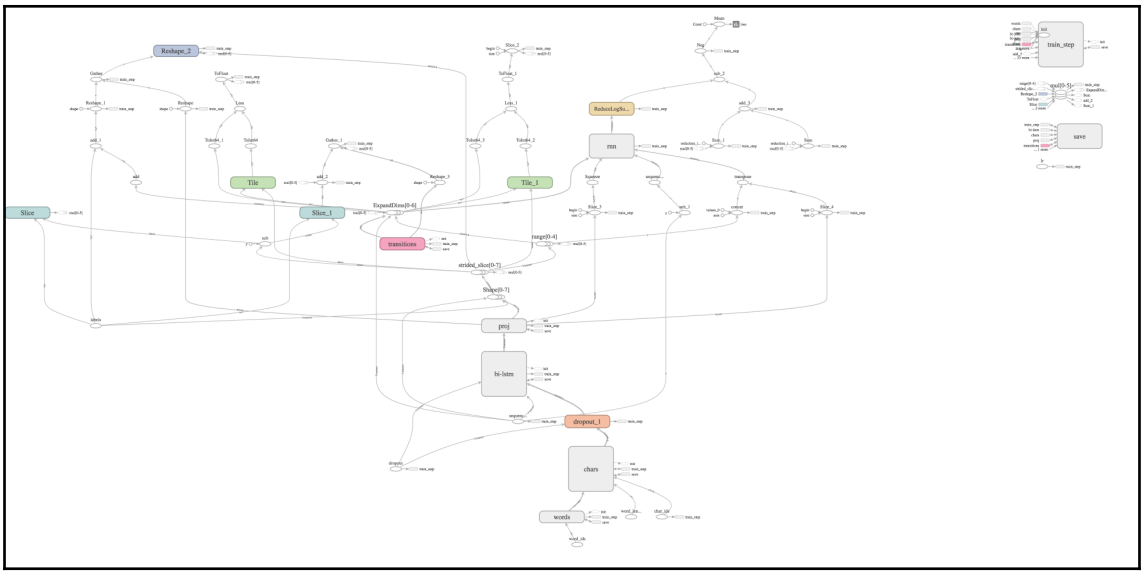
Reshape_3
Tile
Tile_1
Slice
Slice_1
transitions
proj
bi-lstm
dropout_1
chars
words
rnn
train_step
save

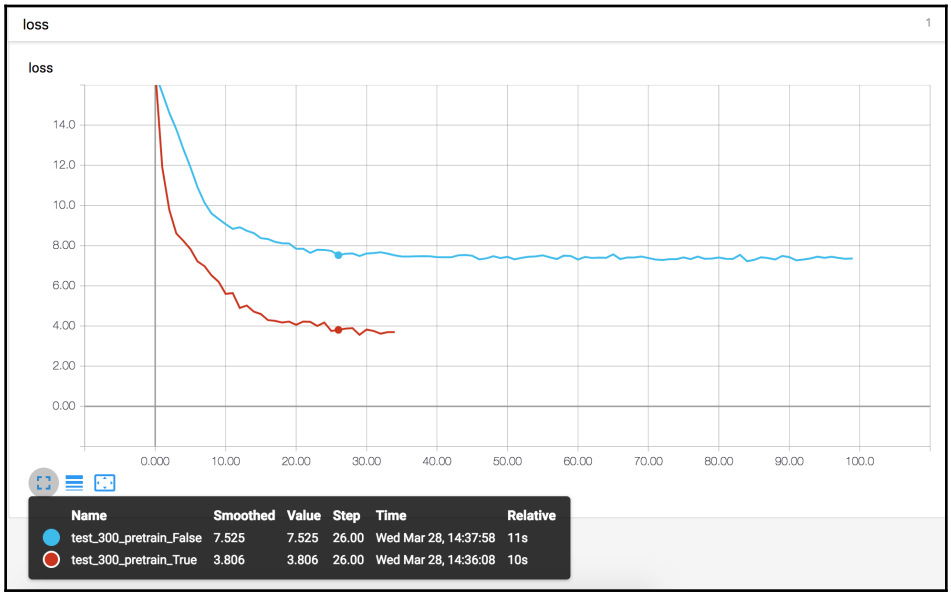
loss
loss
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
0.000
10.00
20.00
30.00
40.00
50.00
60.00
70.00
80.00
90.00
100.0
Name Smoothed Value Step Time Relative
test_300_pretrain_False 7.525 7.525 26.00 Wed Mar 28, 14:37:58 11s
test_300_pretrain_True 3.806 3.806 26.00 Wed Mar 28, 14:36:08 10s

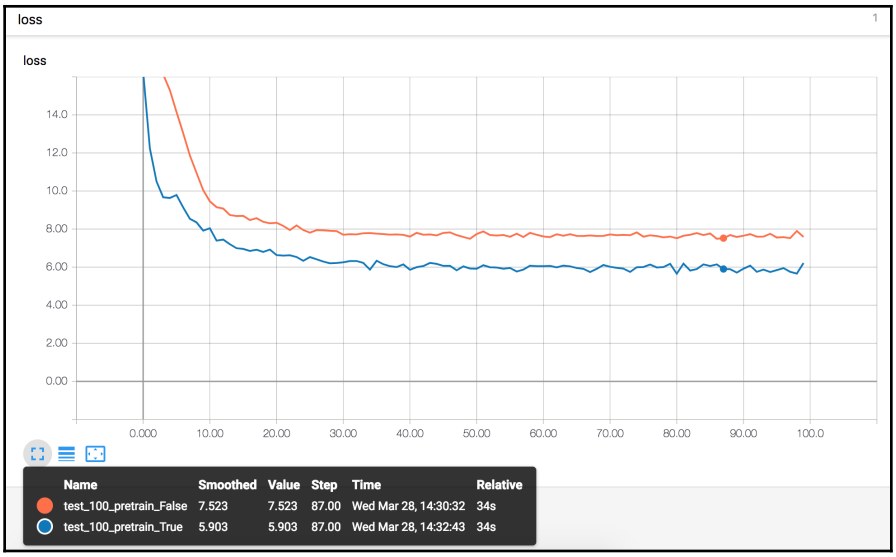
loss
1
loss
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
0.000
10.00
20.00
30.00
40.00
50.00
60.00
70.00
80.00
90.00
100.0
Name Smoothed Value Step Time Relative
test_100_pretrain_False 7.523 7.523 87.00 Wed Mar 28, 14:30:32 34s
test_100_pretrain_True 5.903 5.903 87.00 Wed Mar 28, 14:32:43 34s

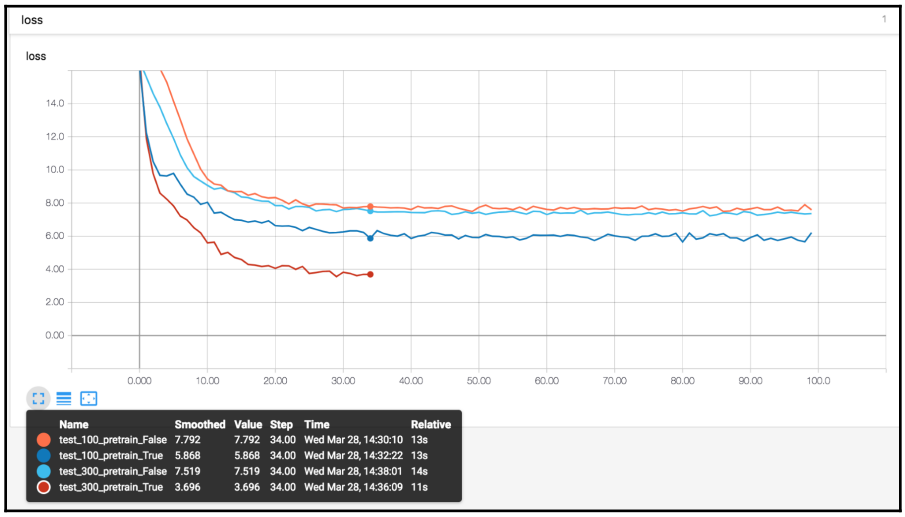
loss
1
loss
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
0.000
10.00
20.00
30.00
40.00
50.00
60.00
70.00
80.00
90.00
100.0
Name Smoothed Value Step Time Relative
test_100_pretrain_False 7.792 7.792 34.00 Wed Mar 28, 14:30:10 13s
test_100_pretrain_True 5.868 5.868 34.00 Wed Mar 28, 14:32:22 13s
test_300_pretrain_False 7.519 7.519 34.00 Wed Mar 28, 14:38:01 14s
test_300_pretrain_True 3.696 3.696 34.00 Wed Mar 28, 14:36:09 11s

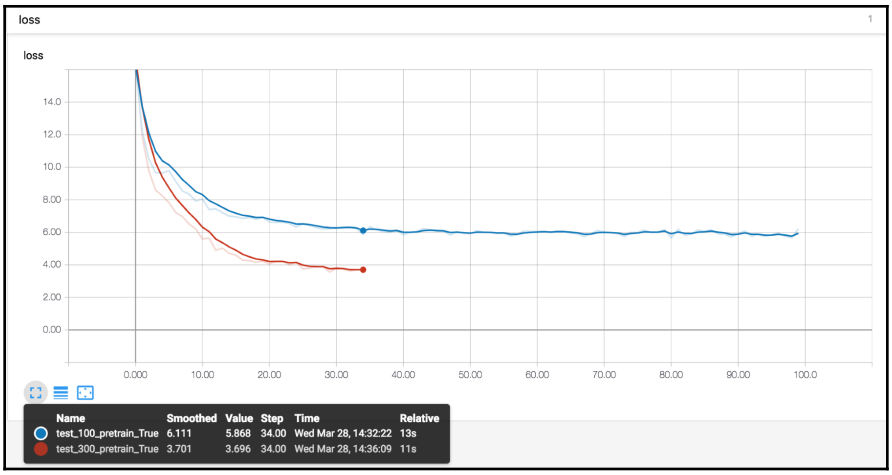
loss
loss
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
0.000
10.00
20.00
30.00
40.00
50.00
60.00
70.00
80.00
90.00
100.0
Name Smoothed Value Step Time Relative
test_100_pretrain_True 6.111 5.868 34.00 Wed Mar 28, 14:32:22 13s
test_300_pretrain_True 3.701 3.696 34.00 Wed Mar 28, 14:36:09 11s

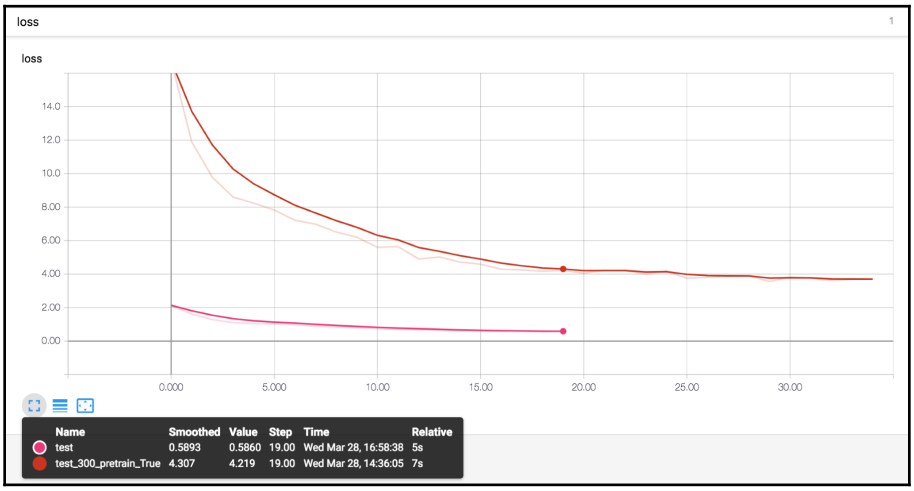
loss
loss
14.0
12.0
10.0
8.00
6.00
4.00
2.00
0.00
0.000
5.000
10.00
15.00
20.00
25.00
30.00
Name Smoothed Value Step Time Relative
test 0.5893 0.5860 19.00 Wed Mar 28, 16:58:38 5s
test_300_pretrain_True 4.307 4.219 19.00 Wed Mar 28, 14:36:05 7s

Chapter 08: Text Generation and Summarization Using GRUs

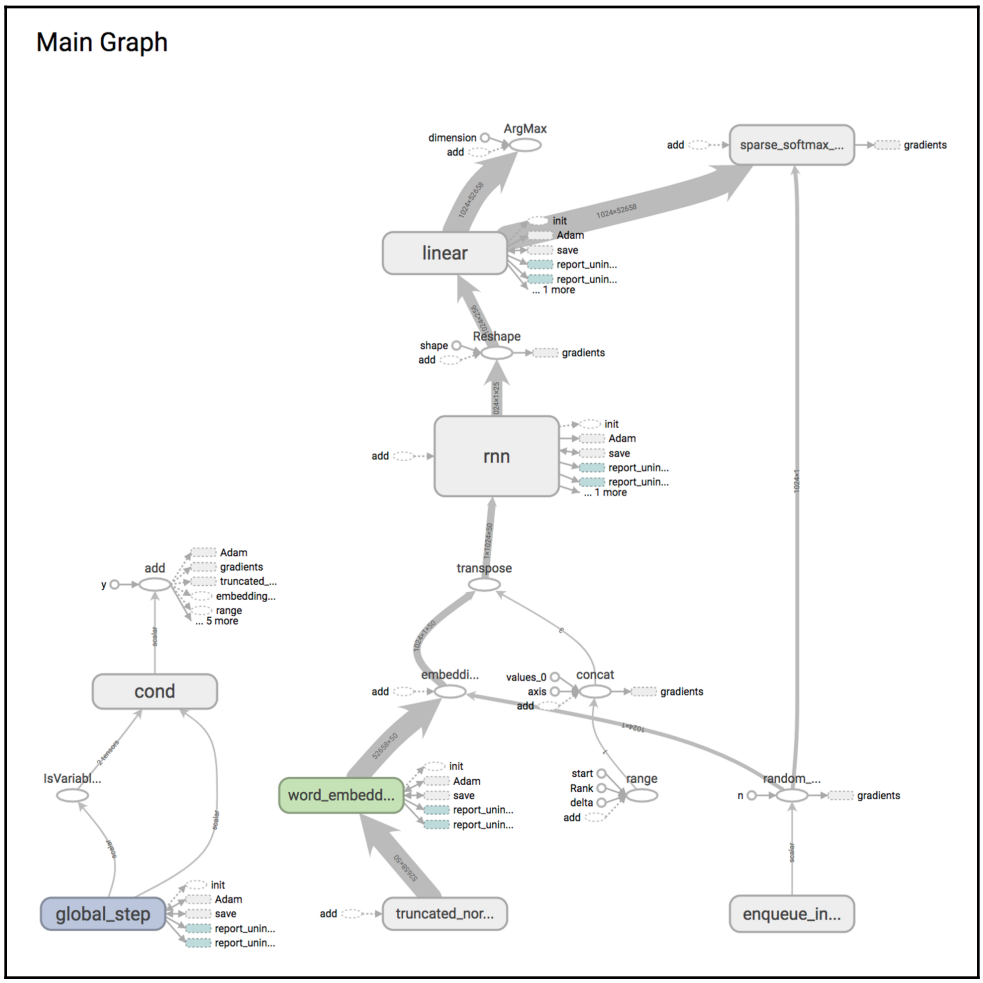

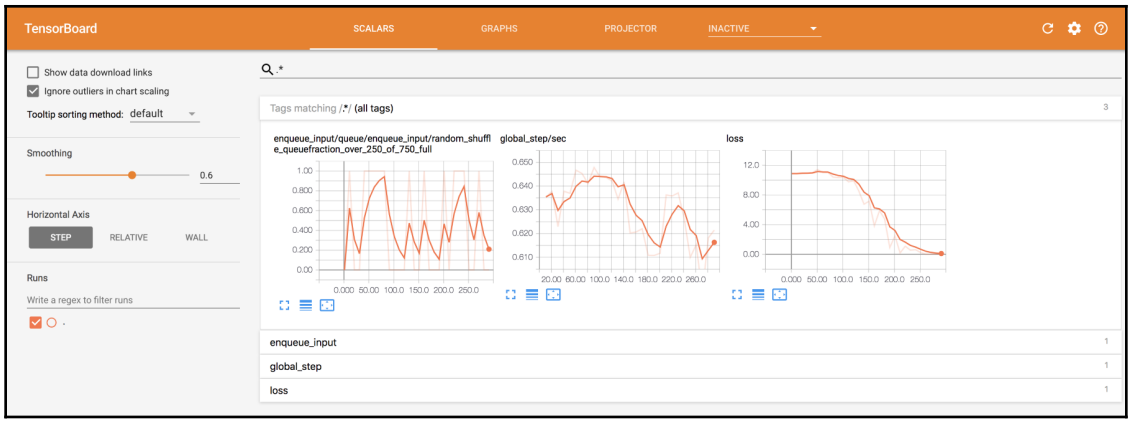
TensorBoard
SCALARS
GRAPHS
PROJECTOR
INACTIVE
Show data download links
Ignore outliers in chart scaling
Tooltip sorting method: default
Smoothing
0.6
Horizontal Axis
STEP
RELATIVE
WALL
Runs
Write a regex to filter runs
Tags matching /.*/ (all tags)
enqueue_input/queue/enqueue_input/random_shuffle_queuefraction_over_250_of_750_full
global_step/sec
loss
enqueue_input
global_step
loss

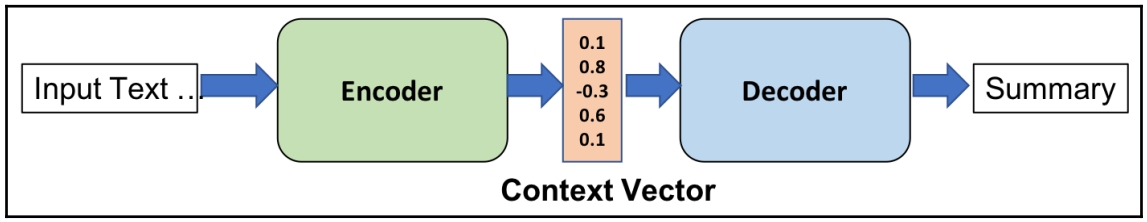
Input Text ..
Encoder
0.1
0.8
-0.3
0.6
0.1
Decoder
Summary
Context Vector

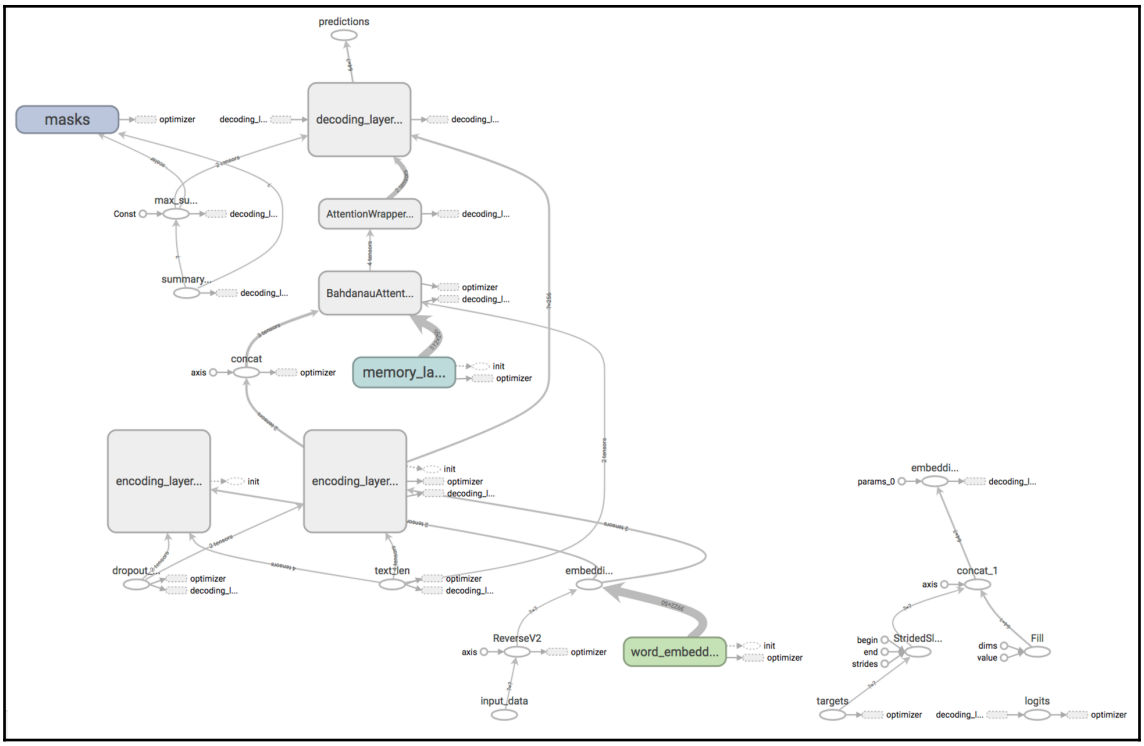
predictions
masks
optimizer
decoding_l...
decoding_layer...
max_su...
Const
AttentionWrapper...
summary...
BahdanauAttent...
concat
axis
memory_la...
init
encoding_layer...
dropout
text_len
embeddi...
ReverseV2
word_embedd...
input_data
params_0
concat_1
StridedSl...
Fill
begin
end
strides
dims
value
targets
logits

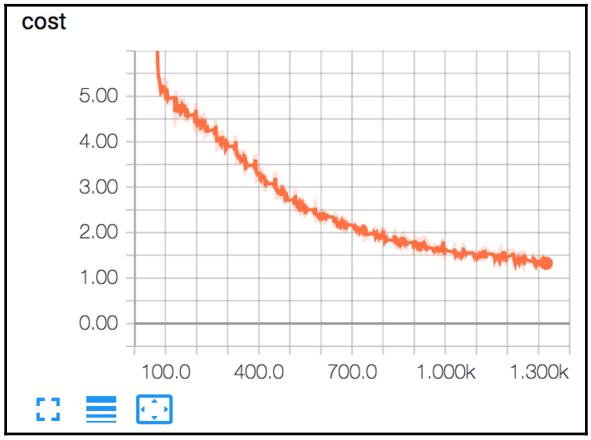
cost
5.00
4.00
3.00
2.00
1.00
0.00
100.0
400.0
700.0
1.000k
1.300k

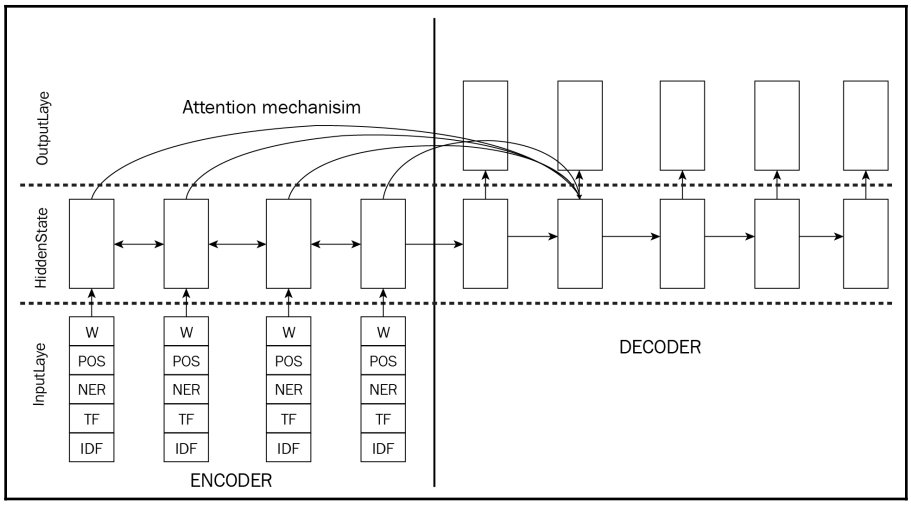
Attention mechanisim
OutputLaye
HiddenState
InputLaye
W
POS
NER
TF
IDF
ENCODER
DECODER

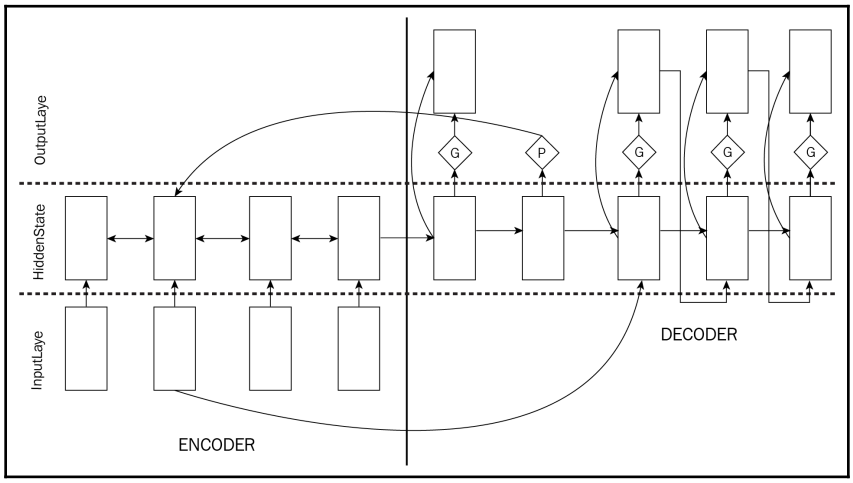
OutputLaye
HiddenState
InputLaye
G
P
ENCODER
DECODER

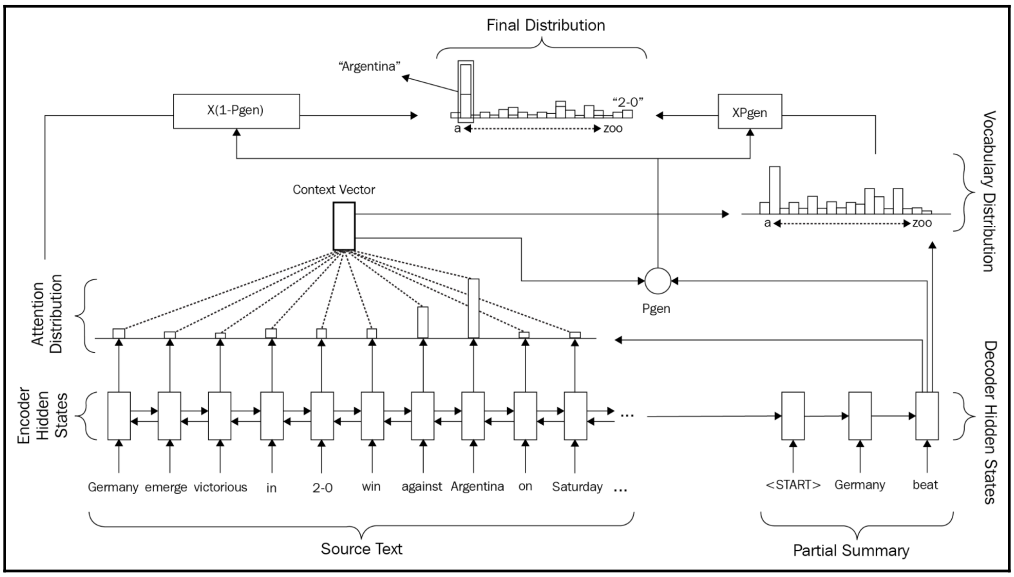
Final Distribution
"Argentina"
"2-0"
X(1-Pgen)
XPgen
a
zoo
Vocabulary Distribution
Context Vector
Attention Distribution
Pgen
Encoder Hidden States
Decoder Hidden States
Germany emerge victorious in 2-0 win against Argentina on Saturday ...
<START> Germany beat
Source Text
Partial Summary

Chapter 09: Question-Answering and Chatbots Using Memory Networks

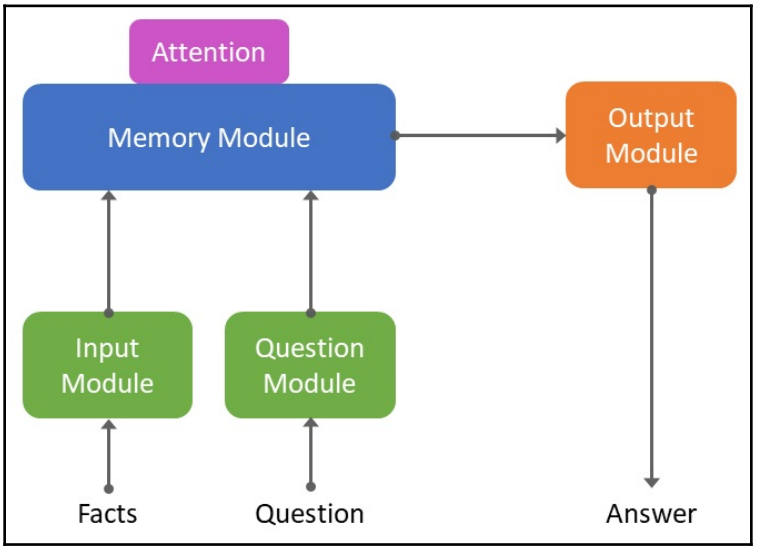

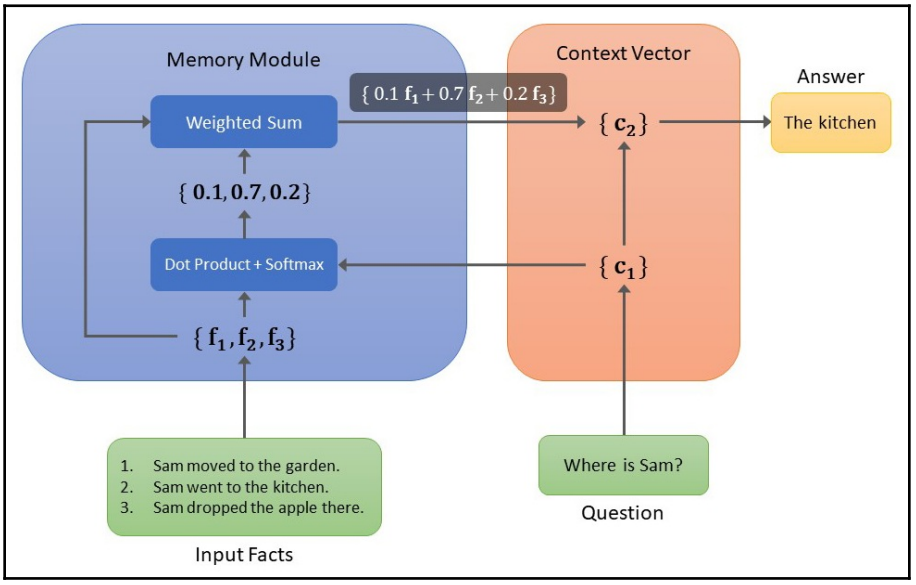

Story (2: 2 supporting facts)	Hop 1	Hop 2	Hop 3
John dropped the milk.	0.06	0.00	0.00
John took the milk there.	0.88	1.00	0.00
Sandra went back to the bathroom.	0.00	0.00	0.00
John moved to the hallway.	0.00	0.00	1.00
Mary went back to the bedroom.	0.00	0.00	0.00
Where is the milk? Answer: hallway Prediction: hallway			

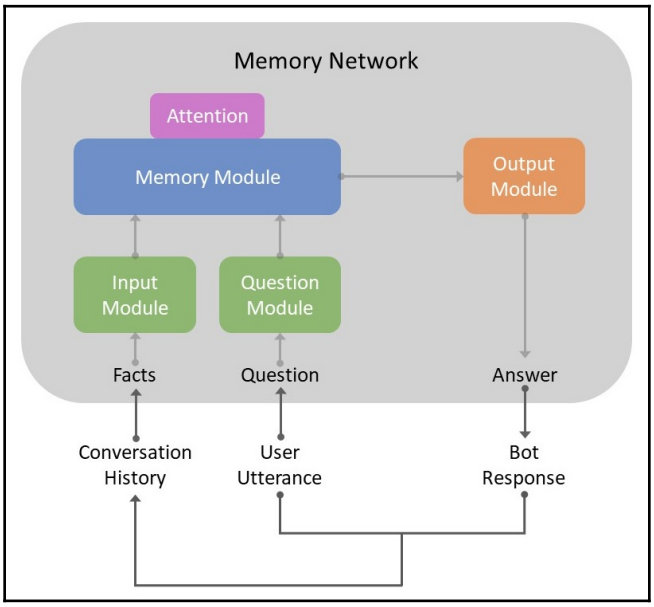

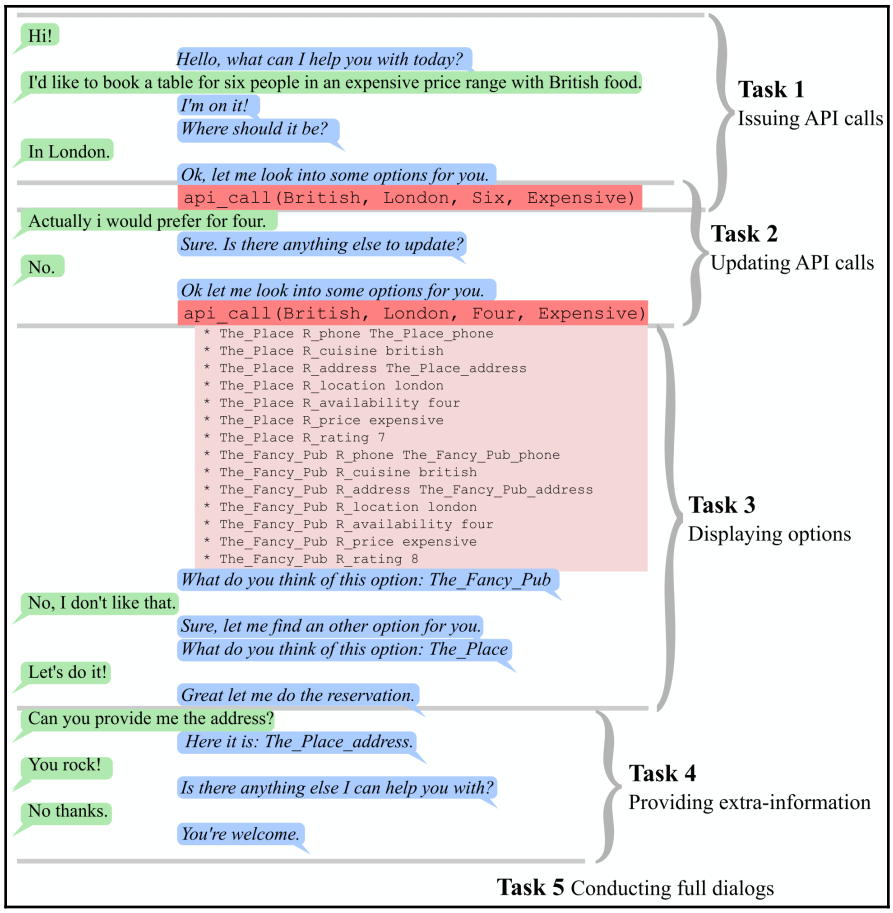
Hi!
Hello, what can I help you with today?
I'd like to book a table for six people in an expensive price range with British food.
I'm on it!
Where should it be?
In London.
Ok, let me look into some options for you.
api_call(British, London, Six, Expensive)
Actually i would prefer for four.
Sure. Is there anything else to update?
No.
Ok let me look into some options for you.
api_call(British, London, Four, Expensive)
* The_Place R_phone The_Place_phone
* The_Place R_cuisine british
* The_Place R_address The_Place_address
* The_Place R_location london
* The_Place R_availability four
* The_Place R_price expensive
* The_Place R_rating 7
* The_Fancy_Pub R_phone The_Fancy_Pub_phone
* The_Fancy_Pub R_cuisine british
* The_Fancy_Pub R_address The_Fancy_Pub_address
* The_Fancy_Pub R_location london
* The_Fancy_Pub R_availability four
* The_Fancy_Pub R_price expensive
* The_Fancy_Pub R_rating 8
What do you think of this option: The_Fancy_Pub
No, I don't like that.
Sure, let me find an other option for you.
What do you think of this option: The_Place
Let's do it!
Great let me do the reservation.
Can you provide me the address?
Here it is: The_Place_address.
You rock!
Is there anything else I can help you with?
No thanks.
You're welcome.
Task 1
Issuing API calls
Task 2
Updating API calls
Task 3
Displaying options
Task 4
Providing extra-information
Task 5 Conducting full dialogs

Chapter 10: Machine Translation Using the Attention-Based Model

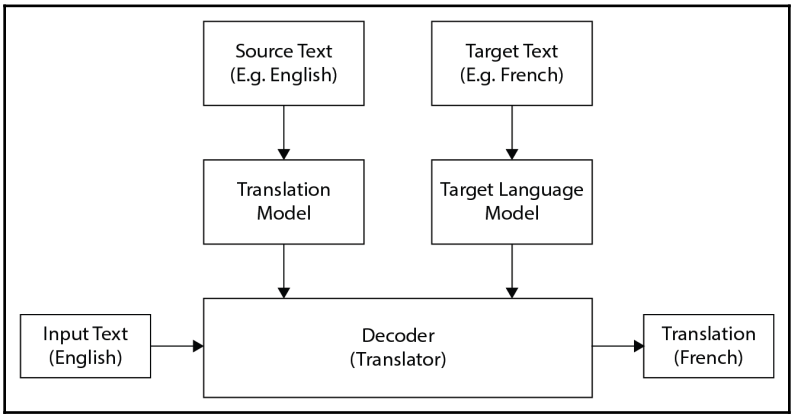

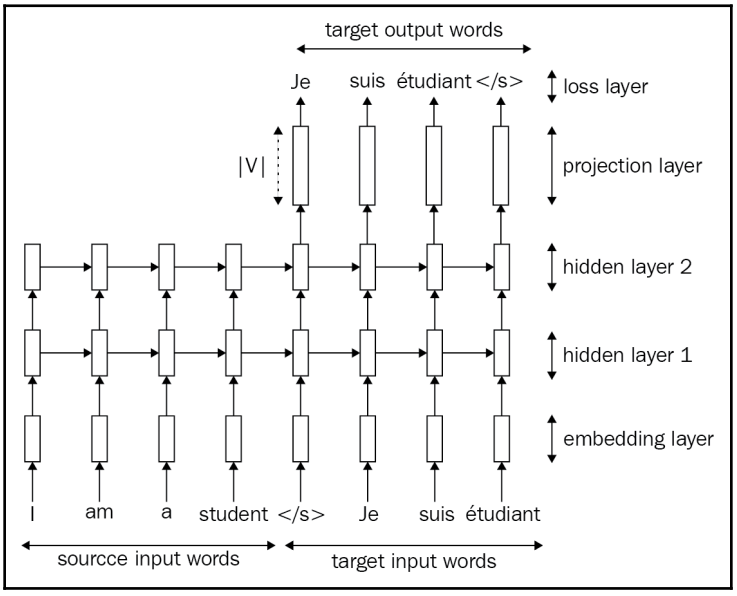

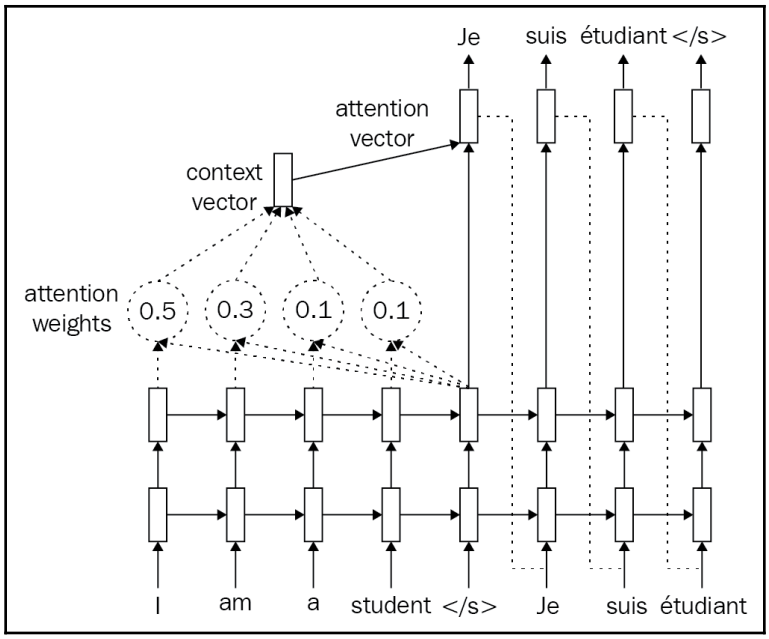
Je
suis
étudiant
</s>
attention
vector
context
vector
attention
weights
0.5
0.3
0.1
0.1
I
am
a
student
</s>
Je
suis
étudiant

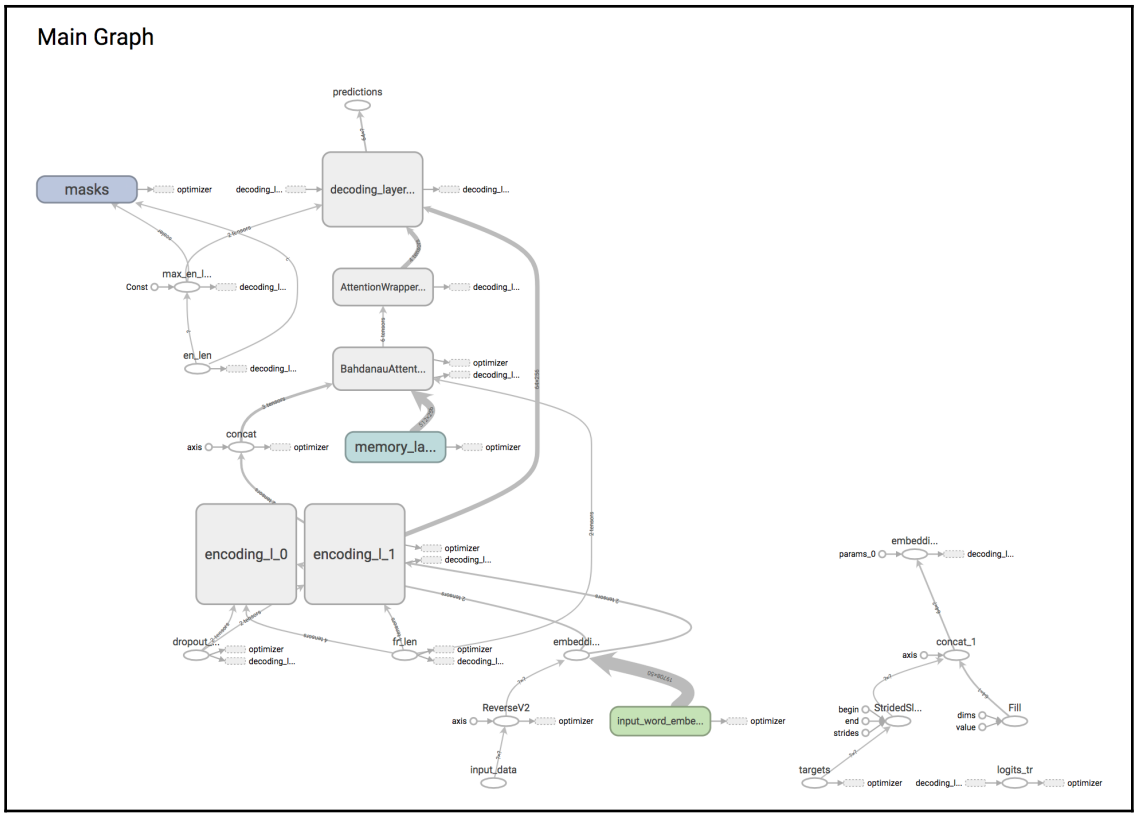
Main Graph
predictions
masks
optimizer
decoding_l...
decoding_layer...
max_en_l...
Const
AttentionWrapper...
en_len
BahdanauAttent...
concat
axis
memory_la...
encoding_l_0
encoding_l_1
dropout_...
fr_len
embeddi...
ReverseV2
input_word_embe...
input_data
params_0
concat_1
StridedSl...
begin
end
strides
Fill
dims
value
targets
logits_tr

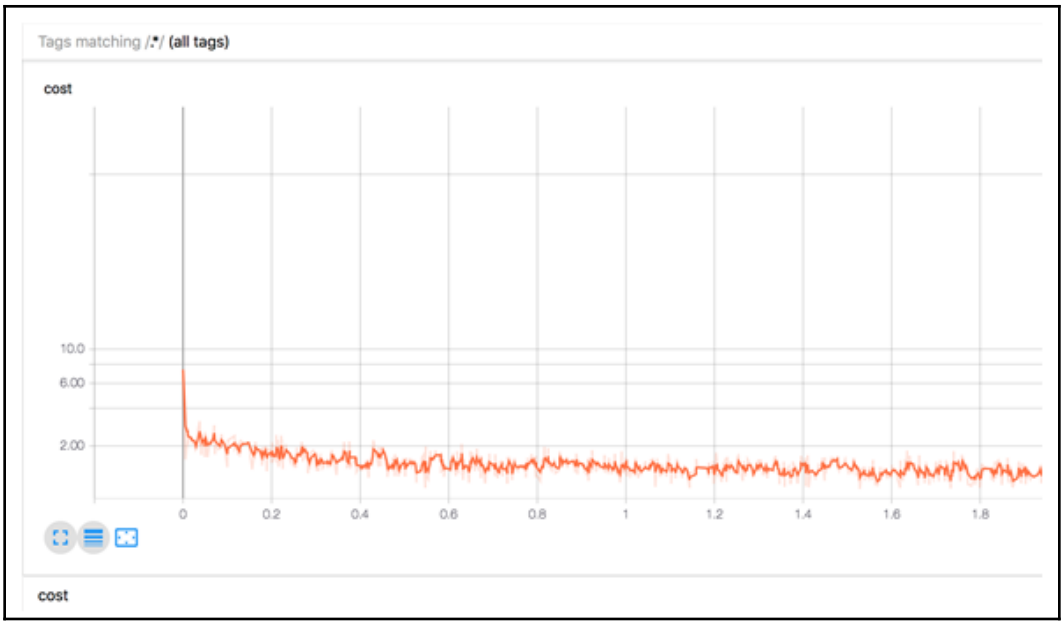
Tags matching /.*/ (all tags)
cost
10.0
6.00
2.00
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
cost

Chapter 11: Speech Recognition Using DeepSpeech

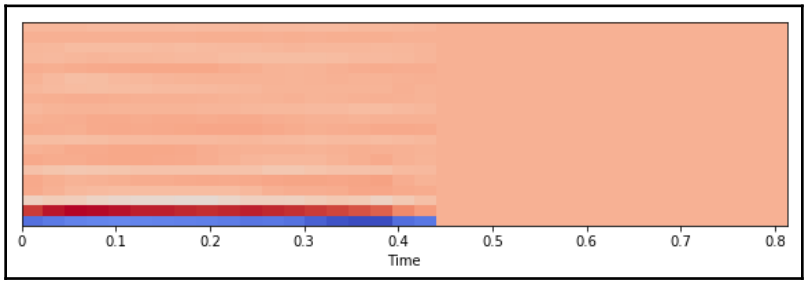

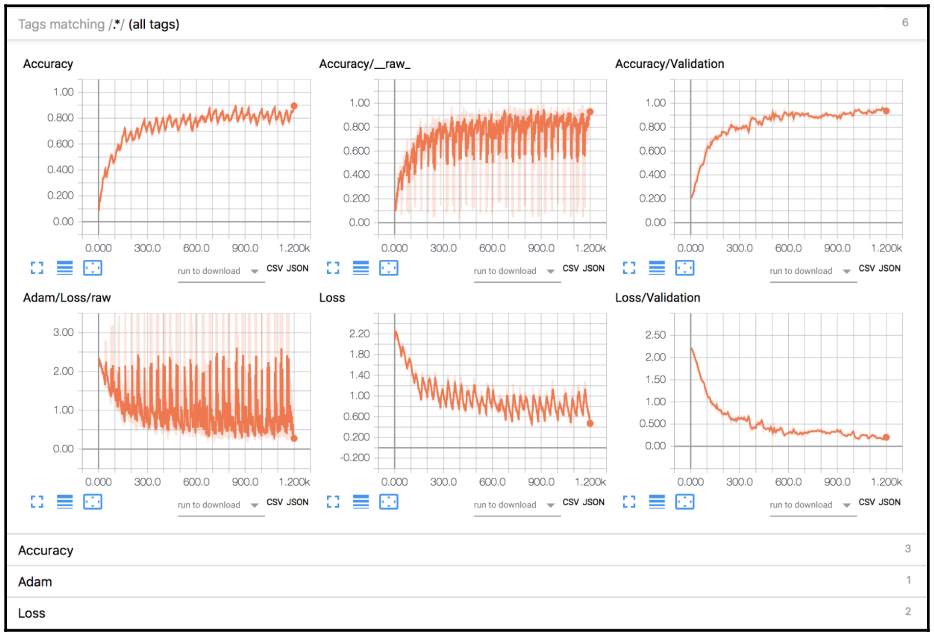

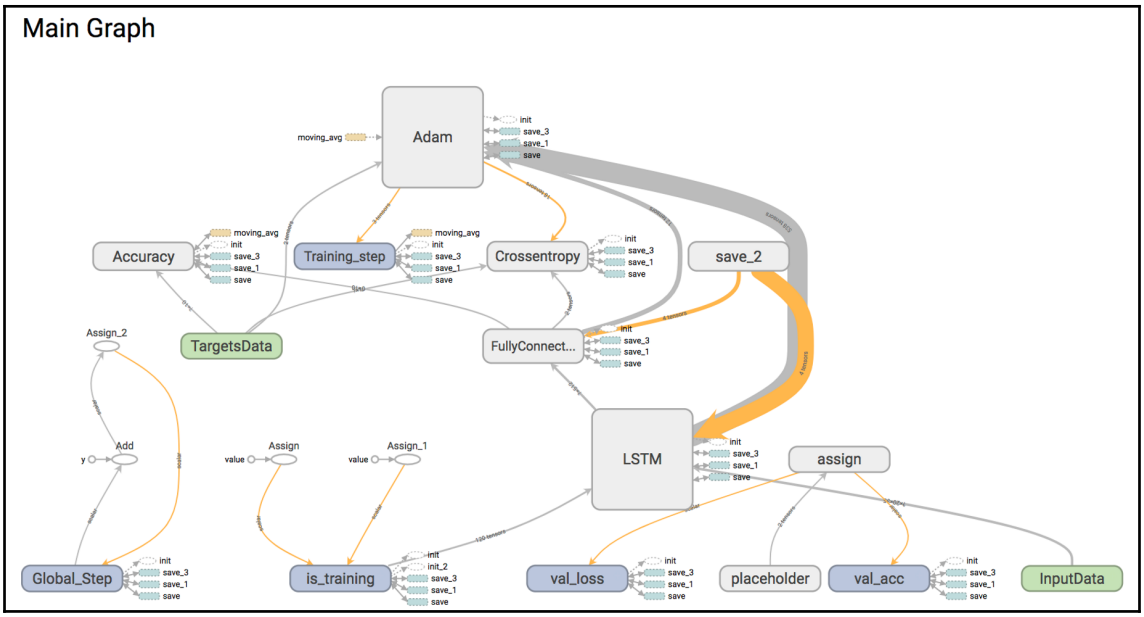
Main Graph
Adam
moving_avg
init
save_3
save_1
save
Accuracy
Training_step
Crossentropy
save_2
TargetsData
FullyConnect...
Assign_2
Add
Assign
Assign_1
LSTM
assign
Global_Step
is_training
val_loss
placeholder
val_acc
InputData

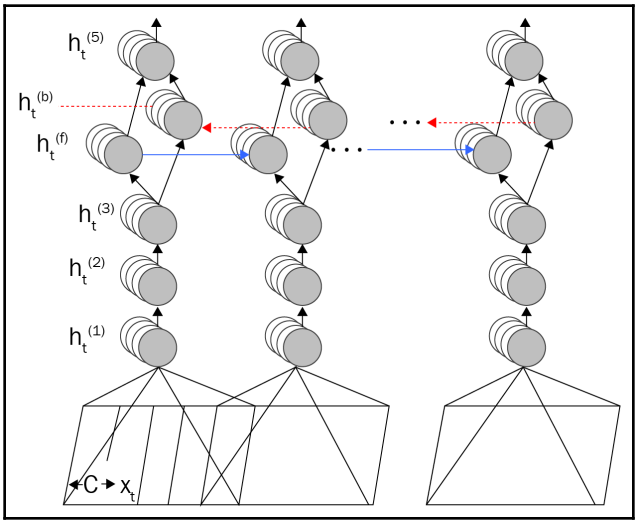
$h_t^{(5)}$
$h_t^{(b)}$
$h_t^{(f)}$
$h_t^{(3)}$
$h_t^{(2)}$
$h_t^{(1)}$
C
x_t

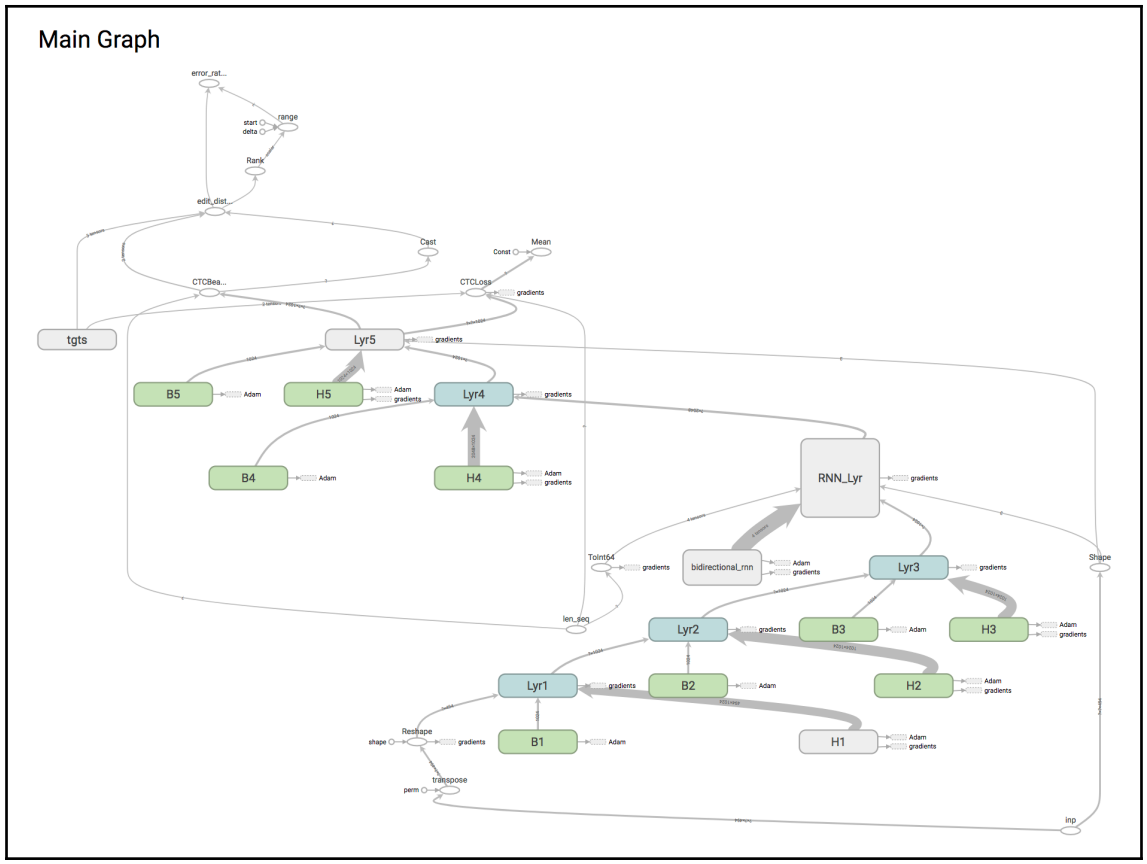
Main Graph
error_rat...
range
start
delta
Rank
edit_dist...
Cast
Mean
Const
CTCBea...
CTCLoss
gradients
tgts
Lyr5
B5
Adam
H5
Lyr4
B4
H4
RNN_Lyr
ToInt64
bidirectional_rnn
Lyr3
Shape
len_seq
Lyr2
B3
H3
Lyr1
B2
H2
Reshape
shape
B1
H1
transpose
perm
inp

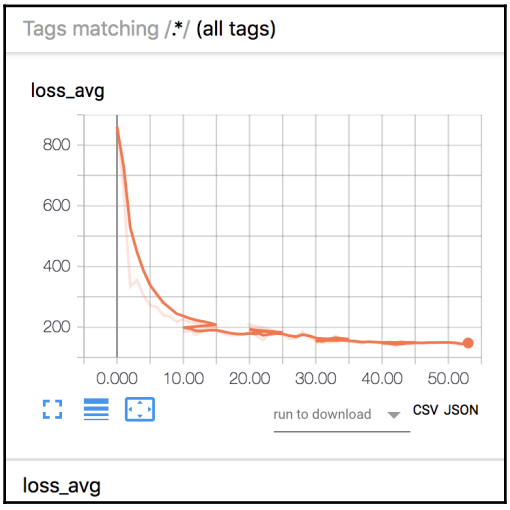
Tags matching /.*/ (all tags)
loss_avg
800
600
400
200
0.000
10.00
20.00
30.00
40.00
50.00
run to download
CSV JSON
loss_avg

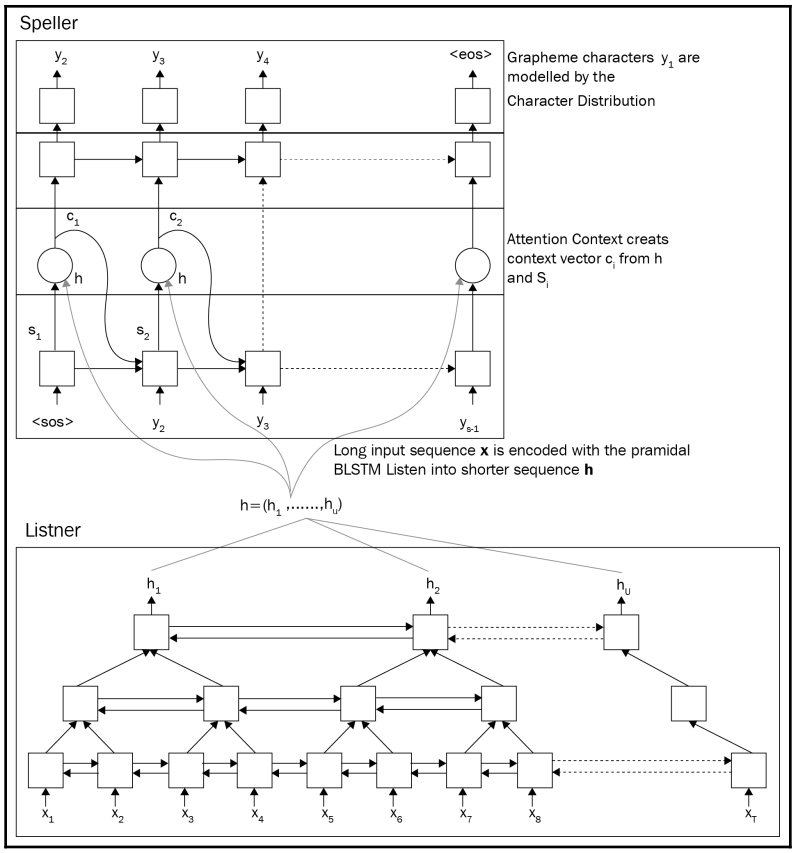
Speller
y2
y3
y4
<eos>
Grapheme characters y1 are modelled by the Character Distribution
c1
c2
h
h
Attention Context creats context vector ci from h and Si
s1
s2
<sos>
y2
y3
ys-1
Long input sequence x is encoded with the pramidal BLSTM Listen into shorter sequence h
h=(h1 ,.....,hU)
Listner
h1
h2
hU
x1
x2
x3
x4
x5
x6
x7
x8
xT

Chapter 12: Text-to-Speech Using Tacotron

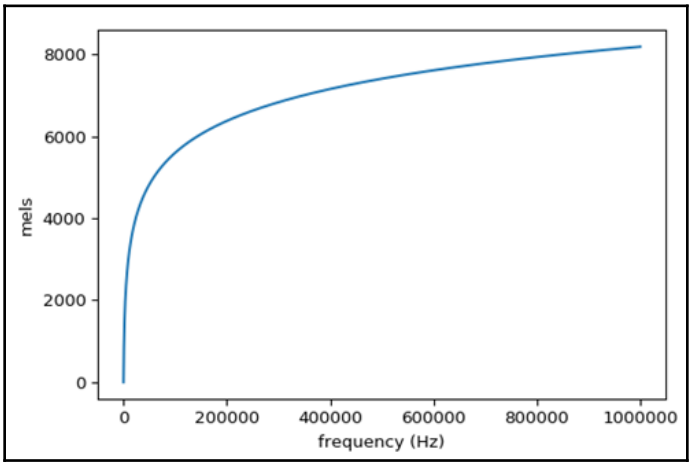

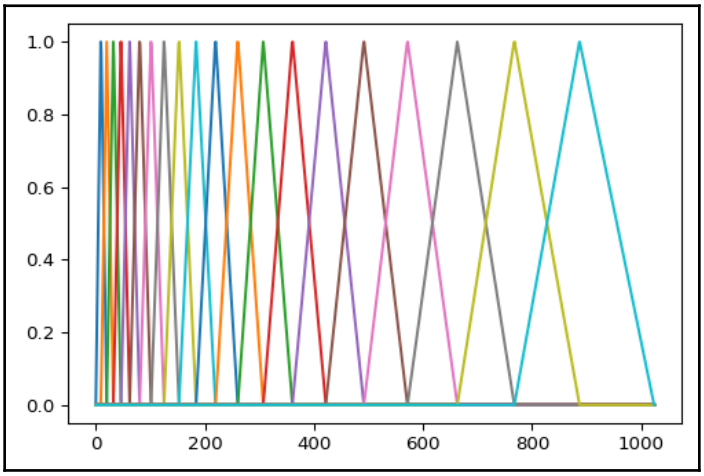

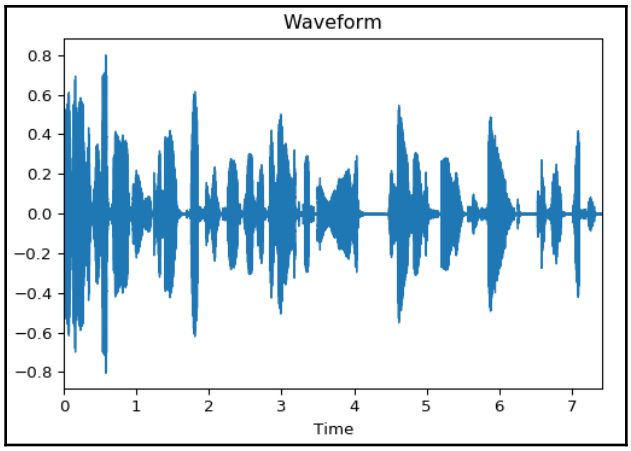
Waveform
0.8
0.6
0.4
0.2
0.0
−0.2
−0.4
−0.6
−0.8
0
1
2
3
4
5
6
7
Time

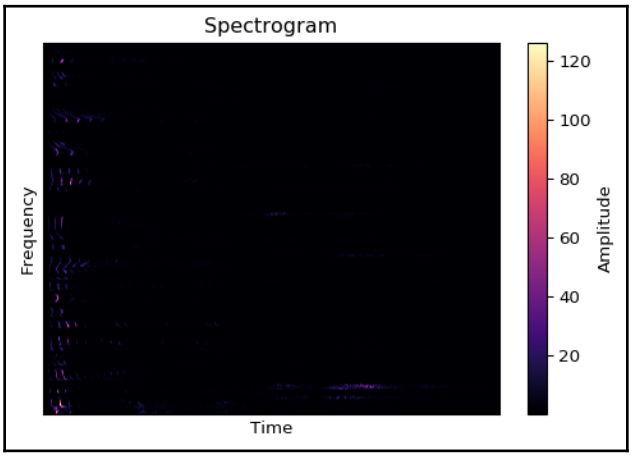
Spectrogram
Frequency
Time
120
100
80
60
40
20
Amplitude

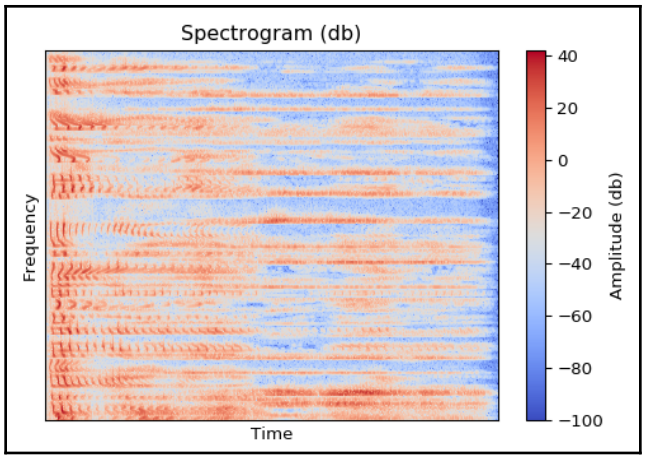
Spectrogram (db)
Frequency
Time
40
20
0
−20
−40
−60
−80
−100
Amplitude (db)

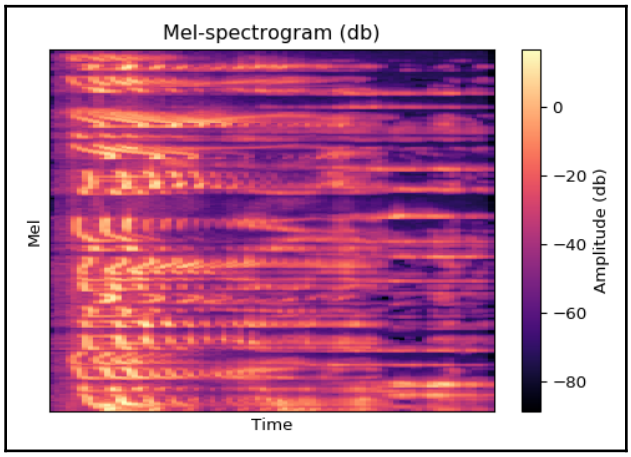
Mel-spectrogram (db)
Mel
Time
0
−20
−40
−60
−80
Amplitude (db)

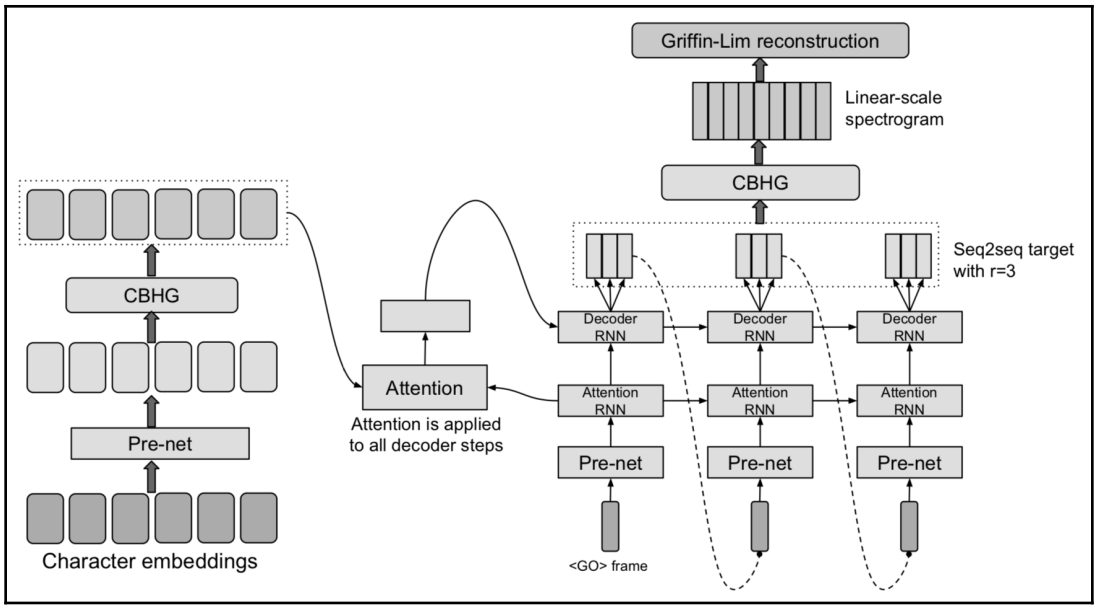
Griffin-Lim reconstruction
Linear-scale spectrogram
CBHG
Seq2seq target with r=3
Decoder RNN
Decoder RNN
Decoder RNN
CBHG
Attention
Attention is applied to all decoder steps
Attention RNN
Attention RNN
Attention RNN
Pre-net
Pre-net
Pre-net
Pre-net
Character embeddings
<GO> frame

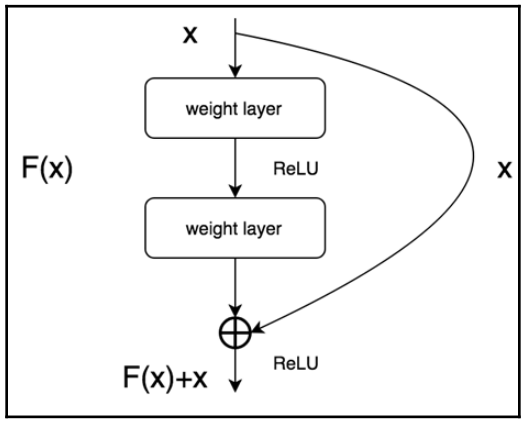
x
weight layer
F(x)
ReLU
x
weight layer
ReLU
F(x)+x

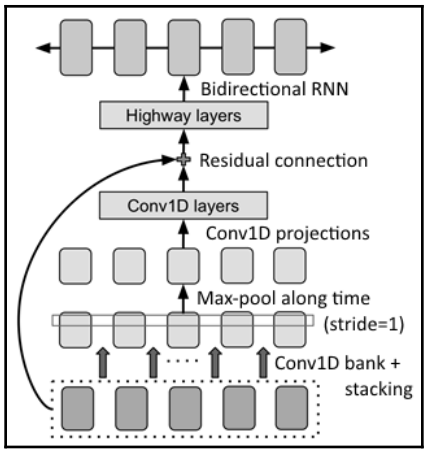

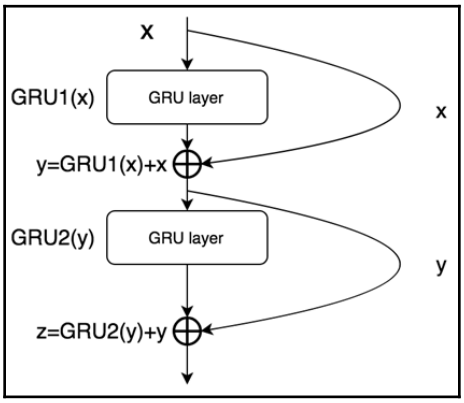

	0	1	2
0	LJ001-0001	Printing, in the only sense with which we are ...	Printing, in the only sense with which we are ...
1	LJ001-0002	in being comparatively modern.	in being comparatively modern.
2	LJ001-0003	For although the Chinese took impressions from...	For although the Chinese took impressions from...

Chapter 13: Deploying Trained Models

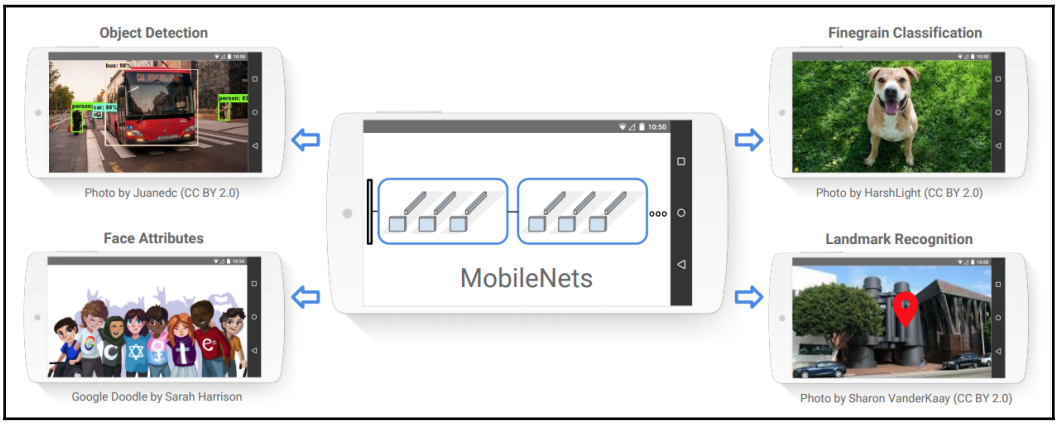

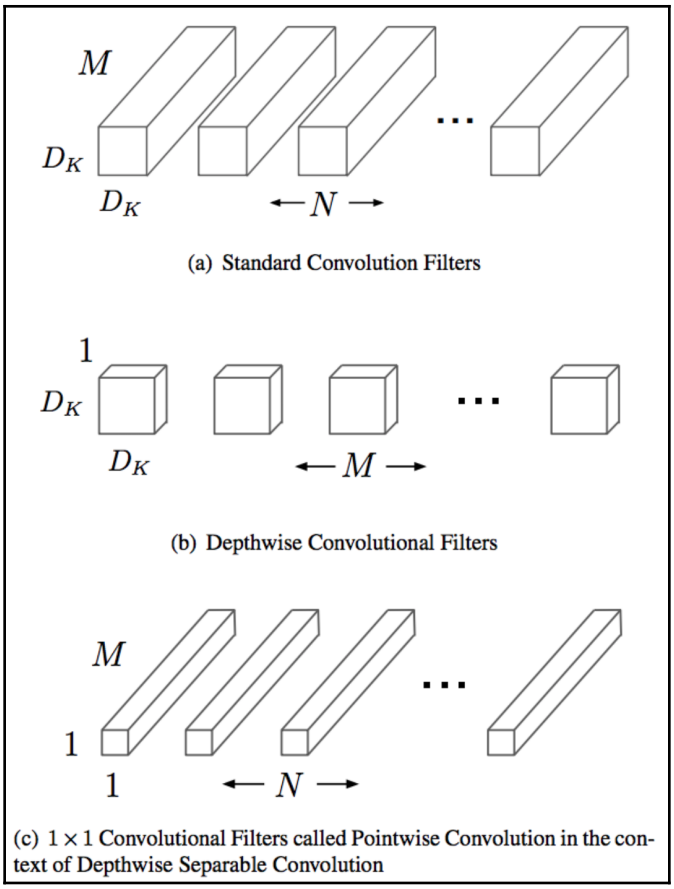
M
D_K
D_K
N
(a) Standard Convolution Filters
1
D_K
D_K
M
(b) Depthwise Convolutional Filters
M
1
1
N
(c) 1 × 1 Convolutional Filters called Pointwise Convolution in the context of Depthwise Separable Convolution

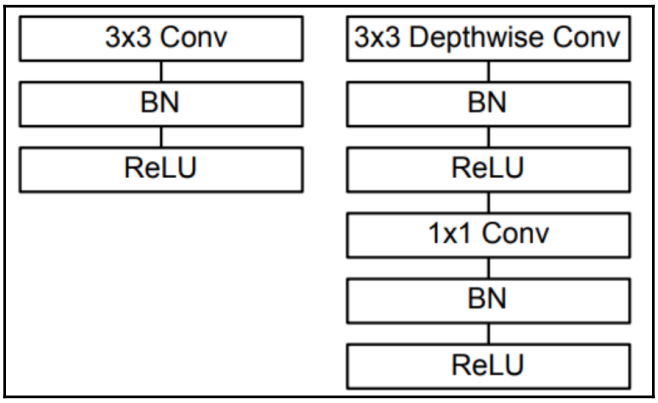
3x3 Conv
BN
ReLU
3x3 Depthwise Conv
BN
ReLU
1x1 Conv
BN
ReLU

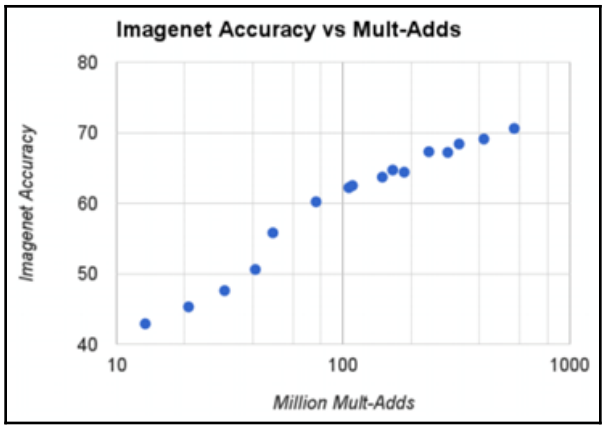
Imagenet Accuracy vs Mult-Adds
80
70
60
50
40
Imagenet Accuracy
10
100
1000
Million Mult-Adds

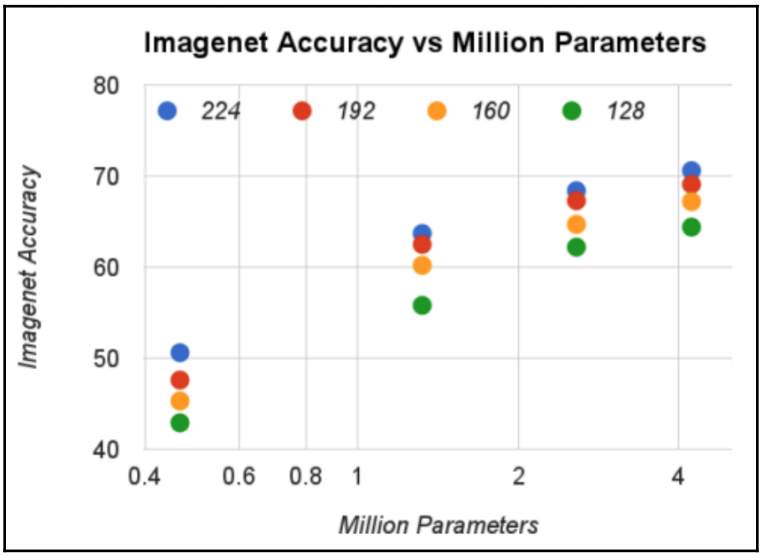
Imagenet Accuracy vs Million Parameters
224
192
160
128
80
70
60
50
40
Imagenet Accuracy
0.4
0.6
0.8
1
2
4
Million Parameters

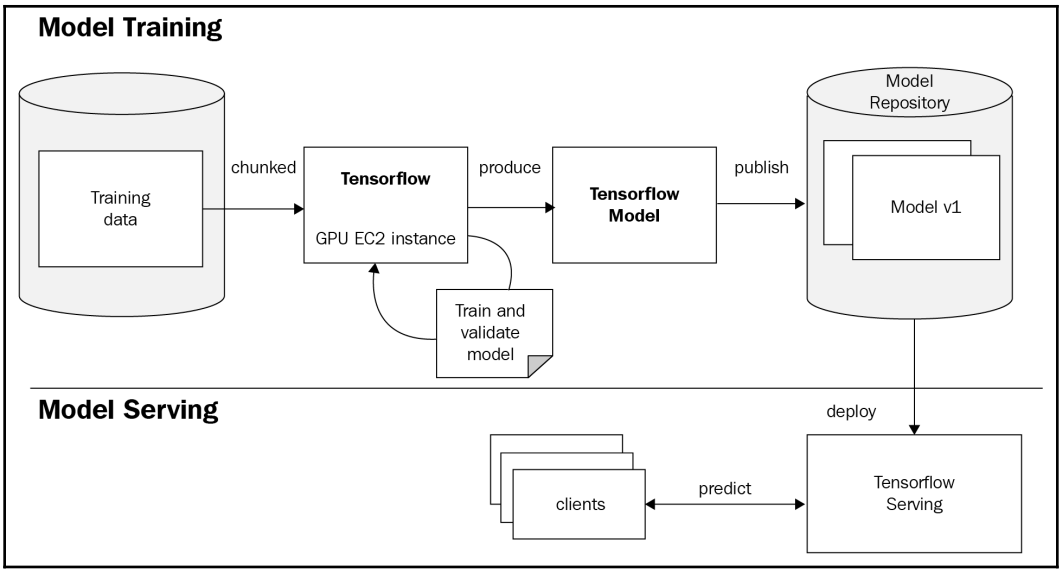
Model Training
Training data
chunked
Tensorflow
GPU EC2 instance
produce
Tensorflow Model
publish
Model Repository
Model v1
Train and validate model
Model Serving
deploy
clients
predict
Tensorflow Serving

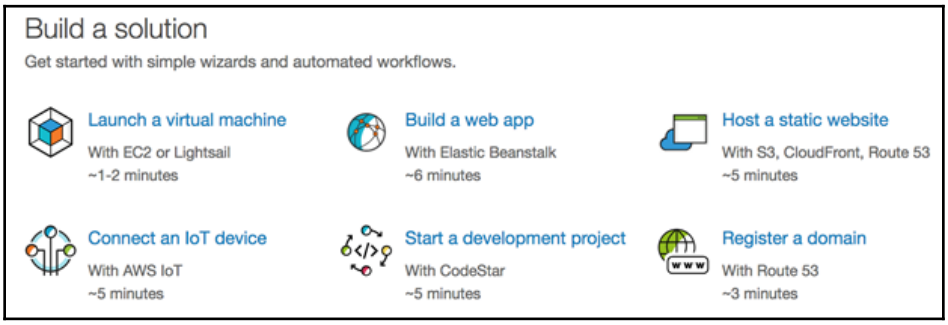
Build a solution
Get started with simple wizards and automated workflows.
Launch a virtual machine
With EC2 or Lightsail
~1-2 minutes
Build a web app
With Elastic Beanstalk
~6 minutes
Host a static website
With S3, CloudFront, Route 53
~5 minutes
Connect an IoT device
With AWS IoT
~5 minutes
Start a development project
With CodeStar
~5 minutes
Register a domain
With Route 53
~3 minutes

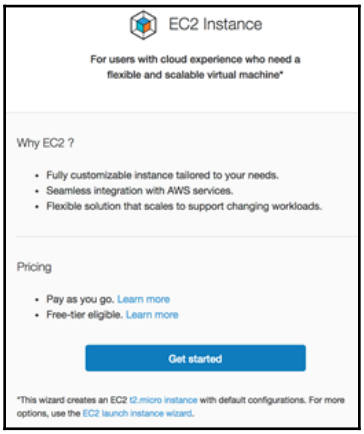
EC2 Instance
For users with cloud experience who need a flexible and scalable virtual machine*
Why EC2 ?
Fully customizable instance tailored to your needs.
Seamless integration with AWS services.
Flexible solution that scales to support changing workloads.
Pricing
Pay as you go. Learn more
Free-tier eligible. Learn more
Get started
*This wizard creates an EC2 t2.micro instance with default configurations. For more options, use the EC2 launch instance wizard.


Name your EC2 instance
This is how you will identify your instance in AWS console. Choose a name that is easy for you to remember.
tensorflow
Use this name

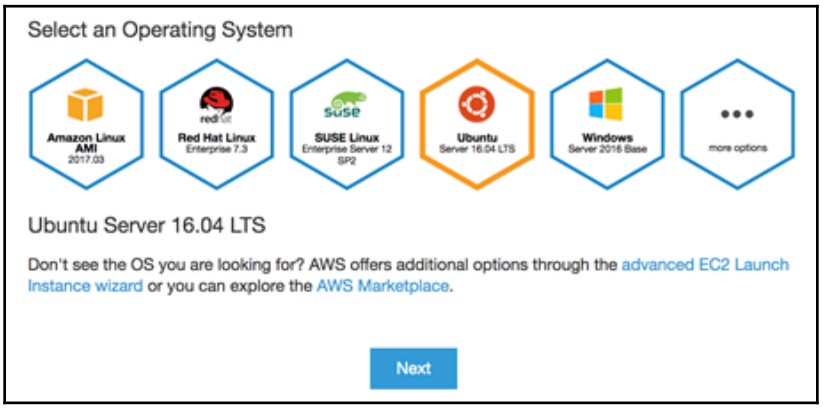
Select an Operating System
Amazon Linux AMI 2017.03
Red Hat Linux Enterprise 7.3
SUSE Linux Enterprise Server 12 SP2
Ubuntu Server 16.04 LTS
Windows Server 2016 Base
more options
Ubuntu Server 16.04 LTS
Don't see the OS you are looking for? AWS offers additional options through the advanced EC2 Launch Instance wizard or you can explore the AWS Marketplace.
Next

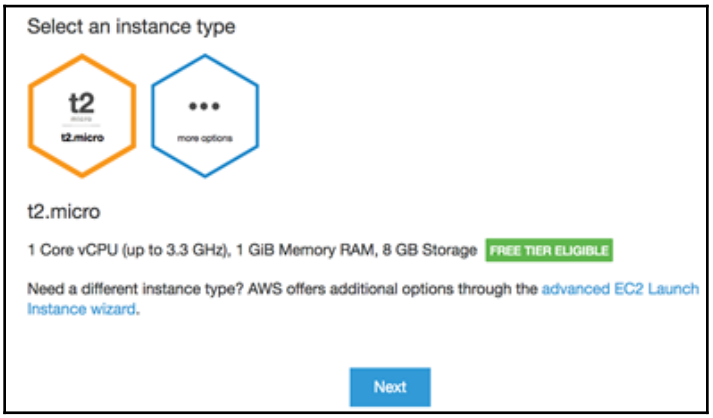
Select an instance type
t2
t2.micro
more options
t2.micro
1 Core vCPU (up to 3.3 GHz), 1 GiB Memory RAM, 8 GB Storage
FREE TIER ELIGIBLE
Need a different instance type? AWS offers additional options through the advanced EC2 Launch Instance wizard.
Next

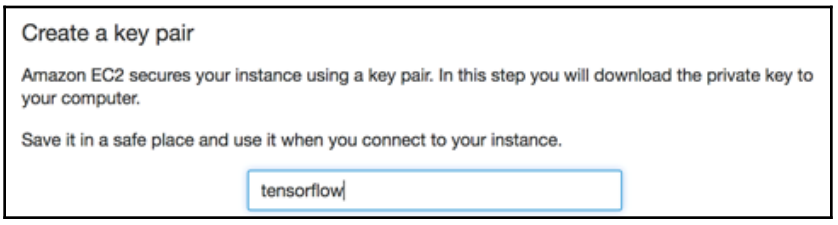
Create a key pair
Amazon EC2 secures your instance using a key pair. In this step you will download the private key to your computer.
Save it in a safe place and use it when you connect to your instance.
tensorflow


tensorflow
Status: Completed!


Proceed to EC2 console

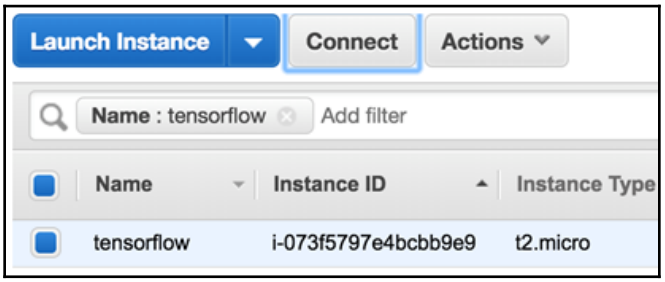
Launch Instance
Connect
Actions
Name : tensorflow
Add filter
Name
Instance ID
Instance Type
tensorflow
i-073f5797e4bcbb9e9
t2.micro

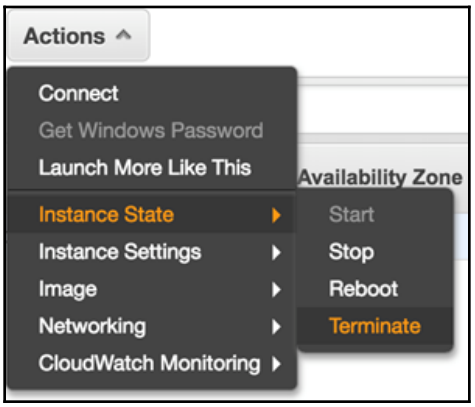
Actions
Connect
Get Windows Password
Launch More Like This
Availability Zone
Instance State
Instance Settings
Image
Networking
CloudWatch Monitoring
Start
Stop
Reboot
Terminate


GO TO CONSOLE

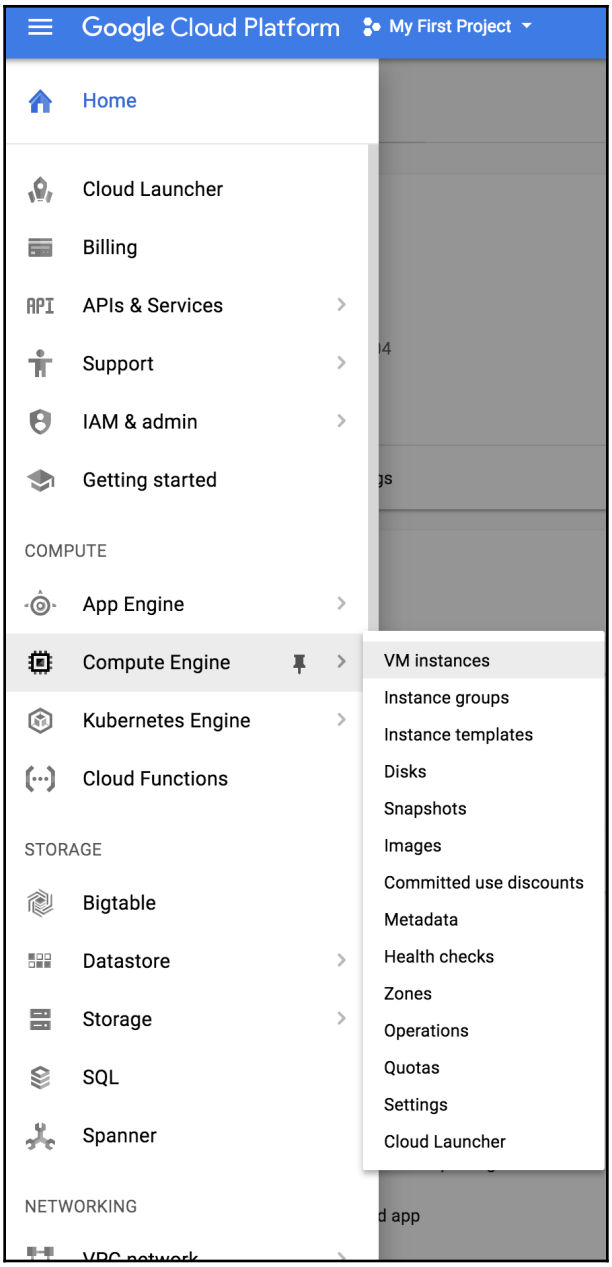
Google Cloud Platform
My First Project
Home
Cloud Launcher
Billing
APIs & Services
Support
IAM & admin
Getting started
COMPUTE
App Engine
Compute Engine
Kubernetes Engine
Cloud Functions
STORAGE
Bigtable
Datastore
Storage
SQL
Spanner
NETWORKING
VM instances
Instance groups
Instance templates
Disks
Snapshots
Images
Committed use discounts
Metadata
Health checks
Zones
Operations
Quotas
Settings
Cloud Launcher
CREATE INSTANCE

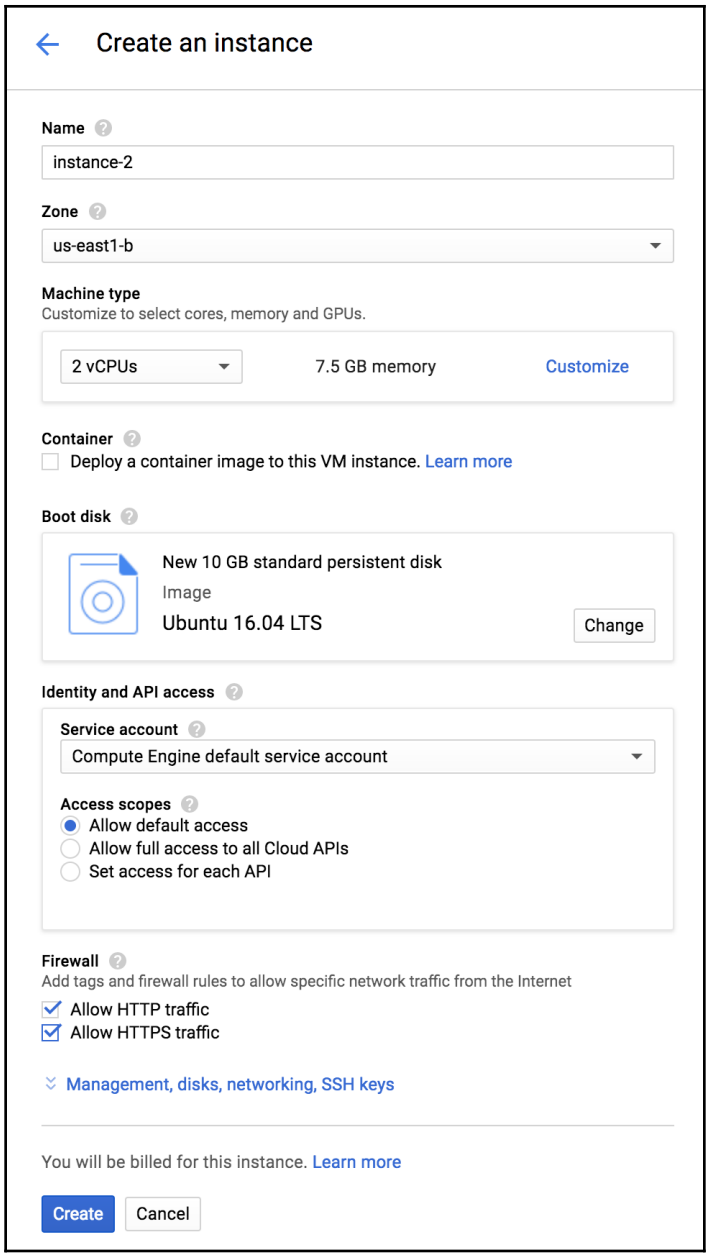
Create an instance
Name
instance-2
Zone
us-east1-b
Machine type
Customize to select cores, memory and GPUs.
2 vCPUs
7.5 GB memory
Customize
Container
Deploy a container image to this VM instance. Learn more
Boot disk
New 10 GB standard persistent disk
Image
Ubuntu 16.04 LTS
Change
Identity and API access
Service account
Compute Engine default service account
Access scopes
Allow default access
Allow full access to all Cloud APIs
Set access for each API
Firewall
Add tags and firewall rules to allow specific network traffic from the Internet
Allow HTTP traffic
Allow HTTPS traffic
Management, disks, networking, SSH keys
You will be billed for this instance. Learn more
Create
Cancel

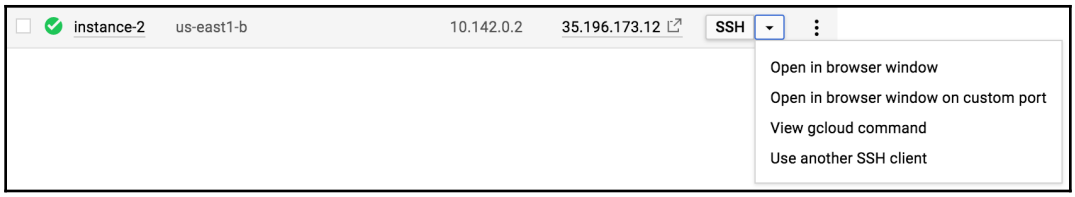
instance-2
us-east1-b
10.142.0.2
35.196.173.12
SSH
Open in browser window
Open in browser window on custom port
View gcloud command
Use another SSH client