

Deep Learning for Everybody

Elliot English and the MetaMind team

elliott@metamind.io



MetaMind

Unstructured visual and textual data

- Enormous growth of images and text
 - 1.8B images shared / day
 - 100B business emails sent / day
- They span all industries and their analysis is valuable
 - Advertising, ad-optimization based on content
 - Medicine, Radiology images: early cancer detection
 - Insurance, Satellite images: building risk analysis
 - Finance, Sentiment analysis for trading
 - Customer Relationship Management, churn prediction
- Their analysis requires machine learning


Machine Learning Used to Require a Ph.D.



IN CS, IT CAN BE HARD TO EXPLAIN
THE DIFFERENCE BETWEEN THE EASY
AND THE VIRTUALLY IMPOSSIBLE.


<http://xkcd.com/1425/>

Why is that?



Feature Engineering is hard! – Real NLP Example


- Task: Predict quality of a radiology report
- Features:




Parsing



Named Entities




Char n-grams



POS tags

His father, Nick Begich, won an election posthumously, only they didn't know for sure that it was posthumous because his plane just disappeared. It still hasn't turned up. It's why locators are now required in all US planes.


Coreference




Taxonomy

Deep Learning can replace all of these: NLP!

- Word vectors:



- Recursive structures




Feature Engineering is hard! – Real Vision Example


- **Task: Predict class of object in image**
- **Features:**



http://en.wikipedia.org/wiki/Edge_detection
Canny Edges




http://docs.opencv.org/trunk/doc/py_tutorials/py_feature2d/py_feature_harris/py_feature_harris.html
Harris Corners



<http://www.cs.toronto.edu/~jcv/papers/ijcv04.pdf>

SIFT



<http://www.vision.ee.ntu.edu.tw/~psurf/pccv06.pdf>

SURF

Deep Learning can replace all of these: Vision!


- Bottom level features from a convolutional neural network:




<http://www.cs.toronto.edu/~fritz/absps/imagenet.pdf>

Deep Learning has revolutionized the industry

- Speech recognition systems of Google, Microsoft, Baidu all use DL
- Google+, Microsoft and others use DL for very accurate image classification, e.g. results for: seat belt, boston rocker, archery, shredder
- Let's take a look at how we're doing on the latter by examining a popular benchmark.






ImageNet Large Scale Visual Recognition Challenge




State-of-the-art rapidly improving

- Convolutional neural networks now the de facto standard for image classification
- LeCun, Yann, et al. "Gradient-based learning applied to document recognition." (1998).
- Krizhevsky, Alex, Ilya Sutskever, and Geoffrey E. Hinton. "Imagenet classification with deep convolutional neural networks." (2012).
- Szegedy, Christian, et al. "Going deeper with convolutions." (2014).




ImageNet Large Scale Visual Recognition Challenge

- We're now at human accuracy! (not really)
- Deep learning still limited to select companies




ImageNet Large Scale Visual Recognition Challenge

- We're now at human accuracy! (not really)
- Deep learning still limited to select companies
- MetaMind makes state-of-the-art deep learning readily usable



MetaMind: Deep learning for everybody

- We take care of the details:
 - Machine learning algorithm selection
 - Hyper parameter tuning
 - Efficient training procedures
 - Computational resource management
 - you don't need to worry about owning your own GPU machines
 - Scalable inference infrastructure
- We constantly improve your performance



Demos!

Language Demo: Twitter Sentiment

 MetaMind

Products ▾ Demos ▾ Developers ▾ Company ▾ Account ▾

Provide Feedback!

Twitter Predictions

TRY IT YOURSELF!

MetaMind lets you search data from Twitter feeds or Twitter hashtags, and classify it automatically using one of our state-of-the-art classifiers. Find out what sentiments your favorite users, politicians or artists express in their tweets!

Classify the sentiment of tweets about

Go!


DAILY TWITTER TRENDS, ANALYZED

Every day, observe a sentiment analysis made by MetaMind on the most popular Twitter topics and hashtags.

RESULTS FOR TRENDING: #ASKJAX

↻ Load next trending

Predictions Distribution



Sentiment	Percentage
positive	33%
neutral	55%
negative	12%

Top 5 predictions

neutral

- 98% boots for men and women to fit a lot of 2015 ARISE from here=> http://t.co/OPaek9pxRi #AskJax #DWTS 4434
- 97% @JAX when you hit 30k on Twitter, can you record a video of you singing and post it? #AskJax
- 95% @Jax what do you think of this vine? https://t.co/M9QI2Zy5gr #AskJAX
- 95% RT @cassyy_x: @Jax I saw this and immediately thought of you #AskJAX http://t.co/yfKqFKNww0
- 100% RT @usahelmaria · "@cassyy_x: @Jax I

negative

- 99% @Jax Worst nightmare? #AskJax Pls. Answer :) :(
- 97% Why did Sportsnet let @CamStewartLive go ? Bad mistake @Jackie_Redmond #askjax
- 96% #askjax Whats your worst personality trait ????????
- 95% @Jax #AskJAX whoa! why u did not answer my questions? Too many fans out there? ????
- 94% RT @Jax: "@sierraaroseee: @Jax #AskJAX why does my lotion smell like something died in it" I don't know how to answer this ????

Language Demo: Semantic Similarity

Products ▾Demos ▾Developers ▾Company ▾Account ▾[Provide Feedback!](#)

Relatedness of Sentences (beta)

This deep learning algorithm is trained to score the semantic similarity between two sentences.

- A score of 1 means the sentences have little to do with each other.
- A score of 5 means they are perfect paraphrases.
- The gray area in between is where things get interesting.

You can type two sentences and press test in order to have the algorithm score their similarity.

Test >

Relatedness: **2.86**

1 - not related at all
5 - an almost perfect paraphrase

Language Demo: Train Your Own Classifier



Products ▾Demos ▾Developers ▾Company ▾Account ▾

Provide Feedback!

Train text classifier

If there is a classification problem for which we do not have a classifier, you can upload labeled training data and, with a just a few clicks, you can train a classifier to predict on new datasets. We will give you an estimate of how well it works on new data.

UPLOAD TEXT DATASET

☒ This dataset is private☒ This classifier is private

 Single TSV

 Multiple Files

 Single ZIP

Select one single text file with one labeled example per line in the following format:

```
sell    the stocks are falling
noAction    the company reported nothing unexpected
sell    revenue has plunged
buy     The company is doing amazingly well. They have just acquired their biggest competitor.
noAction    the company reported is what analysts expected
sell    The employees are leaving
buy     Big breakthroughs in AI are enabling this company to predict outcomes more accurately.
```

Select or drop file

 Upload and train

Vision: General Image Classifier

Products ▾Demos ▾Developers ▾Company ▾Account ▾

Provide Feedback!

General image classifier

Upload a picture to classify it between 1000 image classes. [View classes list](#)



[Change image](#)

- or -


<https://www.metamind.io/static/images/cl>

[Classify >](#)

Please make sure that this URL is an image (PNG or JPEG) and not a web page

- or -

Choose an example



King Penguin, Aptenodytes Patagonica

76% ☒

Ptarmigan 3% ☒

Albatross, Mollymawk 2% ☒

Goose 1% ☒


Sulphur-Crested Cockatoo, Kakatoe Galerita, Cacatua Galerita 1% ☒

Did we make a mistake?

Select the correct label for this image



Vision Demo: Food Classifier



 Change image

- or -

[Classify >](#)

Please make sure that this URL is an image (PNG or JPEG) and not a web page

- or -

Choose an example



Hamburger 99% ☒

Pulled Pork Sandwich <1% ☒

Falafel <1% ☒

Hot Dog <1% ☒

Club Sandwich <1% ☒

Did we make a mistake?

Select the correct label for this image



Vision Demo: Train your own classifier

[Products ▾](#)[Demos ▾](#)[Developers ▾](#)[Company ▾](#)[Account ▾](#)[Provide Feedback!](#)

Train your image classifier

MetaMind allows you to automatically label your images with state-of-the-art classifiers based on convolutional neural networks.

UPLOAD IMAGE DATASET

Datasets with 400 or more images give a better classifier. MetaMind will train for more time on those datasets.


☒ This dataset is private ☒ This classifier is privateLabel 1 Label 2

Select or drop files:
Label 1 (0)

Select or drop files:
Label 2 (0)

Vision Classifier Use Cases

Language Classifiers




Classifiers
Twitter Searches
Other Datasets

CLASSIFIERS

167 results Search				
Name	#examples	(#) Classes	Accuracy	Actions
12 class home improvement 1000 datum: per class	12k	(12) Appliances, Cooking, Crafts, Electrical, Flooring, Garden & Outdoor, Heating & AC, Home Repair, Interior Design, Painting, Plumbing, Pools	81%	Details >
Sentiment (tweets)	331k	(3) neutral, negative, positive	81%	Details >
Informações do Futuro Possível Dono de um Terra Nova	4	(3) , theme, _rels	-	Details >
Copia di TAGS v6.0ns - Archive	3k	(3k) id_str: 575349254912409600, 575349252486962432, 575349251926061056, 575349250908292958	-	Details >
Copia di TAGS v6.0ns - Archive	3k	(3k) id_str: 575349254912409600, 575349252486962432, 575349251926061056, 575349250908292958	-	Details >
Test	2	(2) Cuperlo, Civati	-	Details >
materials for finance	2	(2) materials for finance, test	-	Details >

Vision Classifiers



CLASSIFIERS

216 results Search				
Name	#examples	(#) Classes	Accuracy	Actions
	6k	(71) 1-AB  2-BDP  3-BG  4-BG  5-BB 	59%	Details >
Twitter-brand-Flickr	2k	(38) adidas, aldi, becks, bmw, carlsberg, chimay, Citroen, corona, erdinger, essig, fedex, ford, fosters, guinness, Intel, McDonalds, milka, Mini, Nbc, Nike, nvidia, paulaner, Porsche, Puma, RedBull, rittersport, shell, singha, Sprite, stellaarots, tavaro, teinatan, Unifal, uns, Vorfalme, Yahn, anile, rrvavala	59%	Details >
small_sample	970	(3) acantharia_protist_halo  acantharia_protist  acantharia_protist_big_center 	98%	Details >
plandon_sample	304	(7) tunicate_salp  stomatopod  protist_star 	85%	Details >


API

- Pre-trained classifiers:
 - <https://www.metamind.io/api-quick-start>
- Train your own classifier tutorial:
 - <https://www.metamind.io/api-tutorial-fit>

```
Roberts-MacBook-Pro-2:~ eenglish$ python
Python 2.7.9 (default, Jan 7 2015, 11:50:42)
[GCC 4.2.1 Compatible Apple LLVM 6.0 (clang-600.0.56)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
>>> from metamind.api import ClassificationData, ClassificationModel, set_api_key
>>> set_api_key("FZYw2pAAozbnSdWKMh0zGNvYRGms20HeUTrd07DQsraa0iPMp2")
Hello, elliot_testing last_name
>>> training_data = ClassificationData(private=True, data_type="image", name="my training images")
You are creating dataset 'my training images', with id: 145979
>>> training_data.add_samples([('http://thumbs.dreamstime.com/z/oatmeal-raisin-cookie-13727424.jpg',
...                           ('http://www.mountainmamacooks.com/wp-content/uploads/2011/05/oatmeal-raisin-cookies.jpg', 'oatmeal raisin'),
...                           ('http://whippedtheblog.com/wp-content/uploads/2011/01/oatmeal-raisin-best-cookie.jpg', 'oatmeal raisin'),
...                           ('http://imgls.com/wp-content/uploads/2014/04/P1050813.jpg', 'chocolate chip'),
...                           ('http://thequotablekitchen.com/wp-content/uploads/2012/05/IMG_0732.jpg', 'chocolate chip'),
...                           ('http://media.philly.com/images/chocolate-chip-cookie-app-600.jpg', 'chocolate chip')],
...                           input_type='urls')
Starting data upload...
Finished uploading 6 of 6 samples...
>>> classifier = ClassificationModel(private=True, name="my classifier")
>>> classifier.fit(training_data)
Your MetaMind image model is now training on the server!
Your model is done training!
Model name: my classifier
Model id: 15339
>>> print classifier.predict('http://img.foodnetwork.com/F00D/2012/12/13/HE_Gift-Cookies-baked_s4x3_lead.jpg', input_type='urls')
[{'u'user_value': u'http://img.foodnetwork.com/F00D/2012/12/13/HE_Gift-Cookies-baked_s4x3_lead.jpg', u'probability': 0.9520224728174005, u'label': u'oatmeal raisin'}]
>>> █
```


Also doing research

- Developing new models to improve accuracy
- Improving both training and inference speed
- Addressing new problems involving multimodal systems





MetaMind's Vision

Breakthrough AI for Everybody





Image-Sentence Demo

