Positioning Yourself for the Future

: AI from the Backend

Martin Andrews - Google Developer Expert - Machine Learning
martin@reddragon.ai
@mdda

About me

- Google Developer Expert for Machine Learning
- Co-organizer of the SG Deep Learning SG MeetUp
- Background in Finance & ML
- Deep Learning Research
- Red Dragon AI

Outline

RED DRAGON AI

- Where are the jobs?
- Practical AI basics
  - Some key ideas
  - A concrete example
- More image stuff
- Text models
  - Pace of change
- How to get started

Jobs

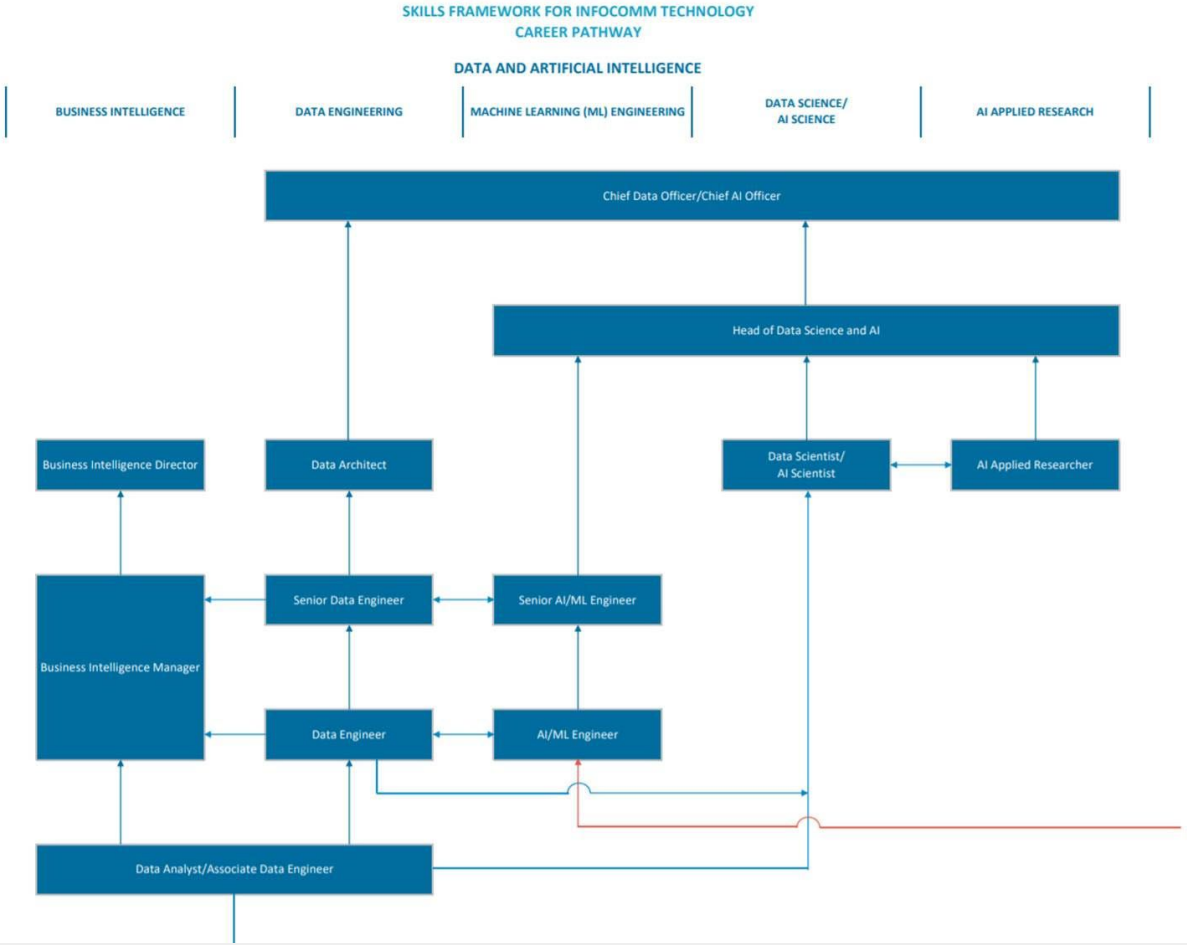

RED DRAGON AI

[Main Infocomm Site]

[Career Pathways Diagram]

Some of the Job Roles

RED DRAGON AI

(non-developer)

- Analyst : Increase value proposition
  - Go beyond spreadsheets, SQL or visualisation
  - Leverage domain expertise
- Manager :
  - Understand how ML thinking is different
  - Managing research is different from implementation

Developer-relevant Job Roles

RED DRAGON AI

- Research :
  - Obvious model creation/research value here
  - But industry usage is different from ‘pure’ datasets
- Developer :
  - Pick up practical AI skill-set = v. doable
  - No need to reinvent the wheel
  - Much ‘AI’ is now an applied science (+art)

Developer-relevant Job Roles

RED DRAGON AI

- Engineers : Understand how these systems are different
  - Cloud / servers
    - Reliability, scaling, versioning, clustering
  - Mobile
    - Making model smaller, lower latency, privacy

Practical AI Basics

ImageNet Competition

RED DRAGON AI

- 1000 different ‘things in images’
  - Breeds of Dogs
  - Airplane, Lamp, Kimono, Tusker
- Training set : Millions of accurately labeled images
- Goal :
  - Make a model to predict on unseen images
  - Measure ‘top-1’ and `top-5’ accuracies

Winning the ImageNet Competition

- Since 2012 : Build a Deep Learning network
- How to improve:
  - Vary building blocks
  - Vary how they are interconnected
  - Use big GPU farms, or create TPU devices...
- Achievements :
  - Models now better than humans (2015/16)
  - ... and are useful for non-ImageNet problems

Leveraging ImageNet

RED DRAGON AI

- Building / Training a full ImageNet model is expensive
  - But trained models are free($) and Free(-dom)
- Key idea :
  - Models have learned to ‘see’
  - On unknown classes, models are consistent
- So ...
  - Can build a new model on top of pretrained one
  - New model = Our task : Our images, Our classes

Quick Example

RED DRAGON AI

- This is from a student for our “Jumpstart Course”
  - (now called “Foundations”)
  - Nicely laid out project
  - Good local flavour
  - (everyone gets to choose their own project)

- WALKTHROUGH

More Image Stuff

Advanced Computer Vision Course

RED DRAGON AI

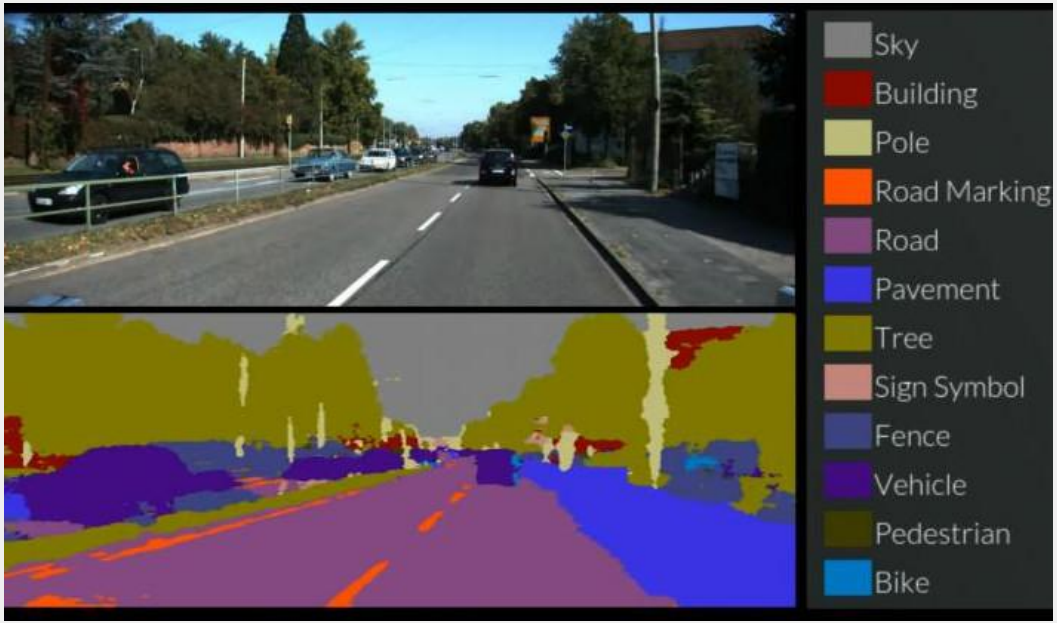

- Object Detection
- Object Tracking
- Counting things
- Image generation
- Colourization
- Facial recognition
- Plus much more...

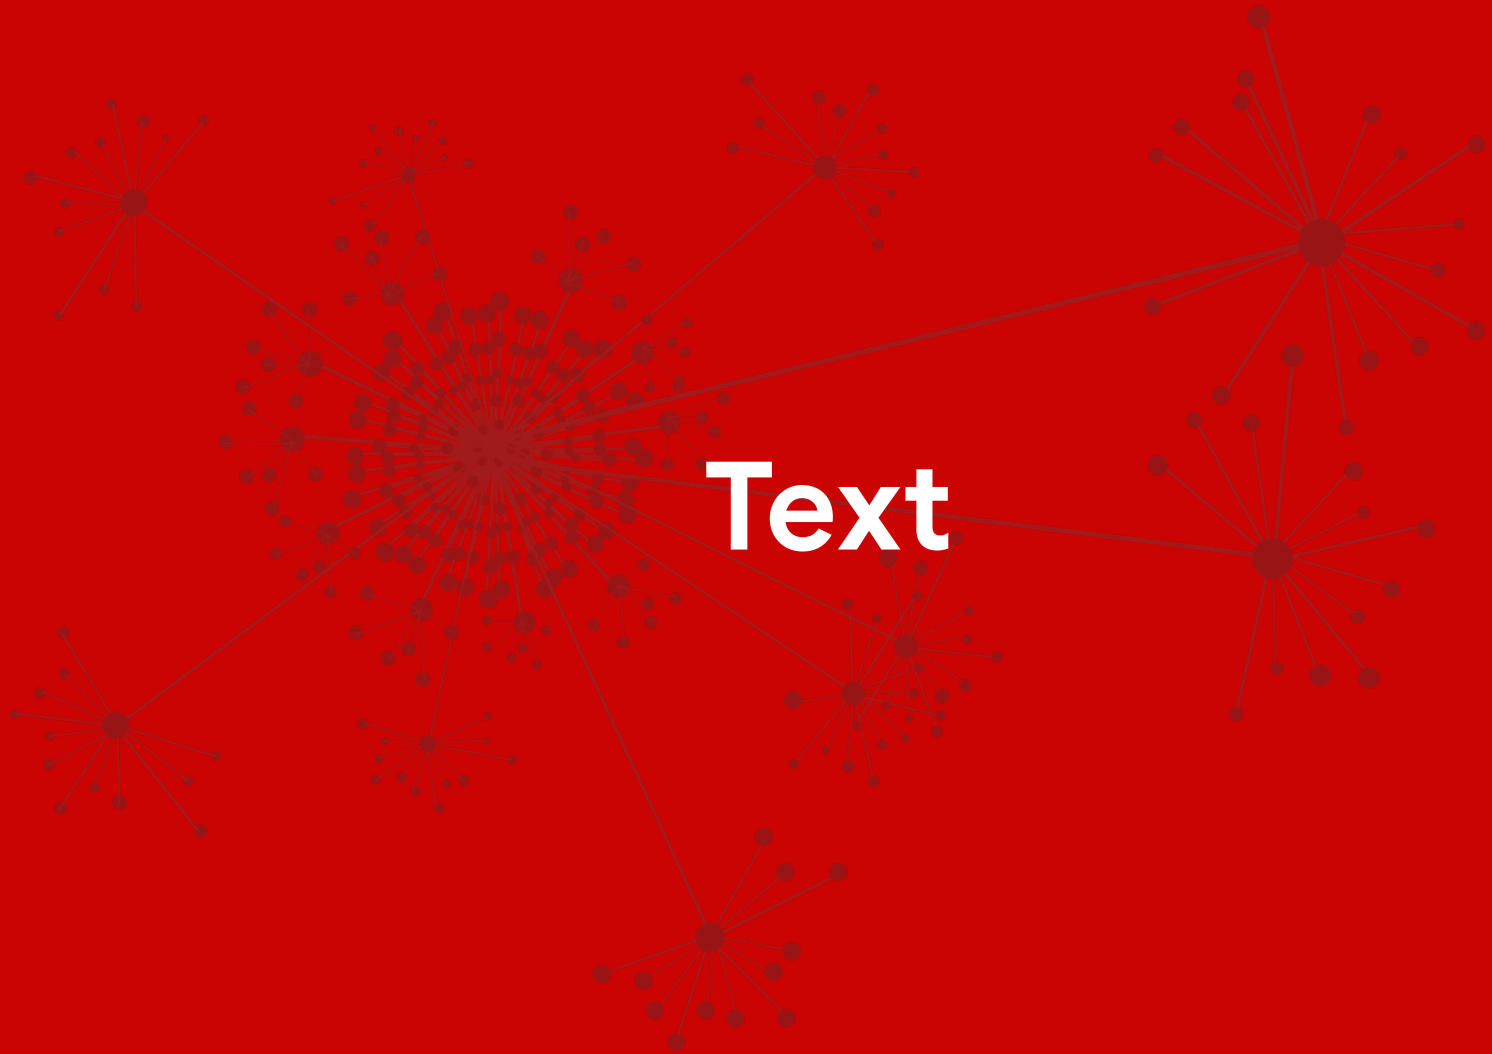
Text

Text Tasks

- Is this product review good or bad?
- Is this comment spam (or offensive)?
- What does this customer want?
- What names are mentioned in this news article?
  - And who are they exactly?

Example Task : Feedback Urgency

RED DRAGON AI

- Suppose we have a stream of incoming feedback
  - Some positive
  - Some mildly negative
  - Some urgently negative
- This is a classification task, so :
  - Think about actual goal (eg: value of each response)
  - Make a dataset (or use existing data creatively)
- Can we leverage other models (like ImageNet before)?

Language Models and BERT

RED DRAGON AI

- Natural Language Processing (NLP) revolutions:
  - Word embeddings (2012)
  - BERT etc (2018...)
- Dramatic improvements in text ‘insights’
- These advances create models that can be re-used

Training BERT (un-supervised)

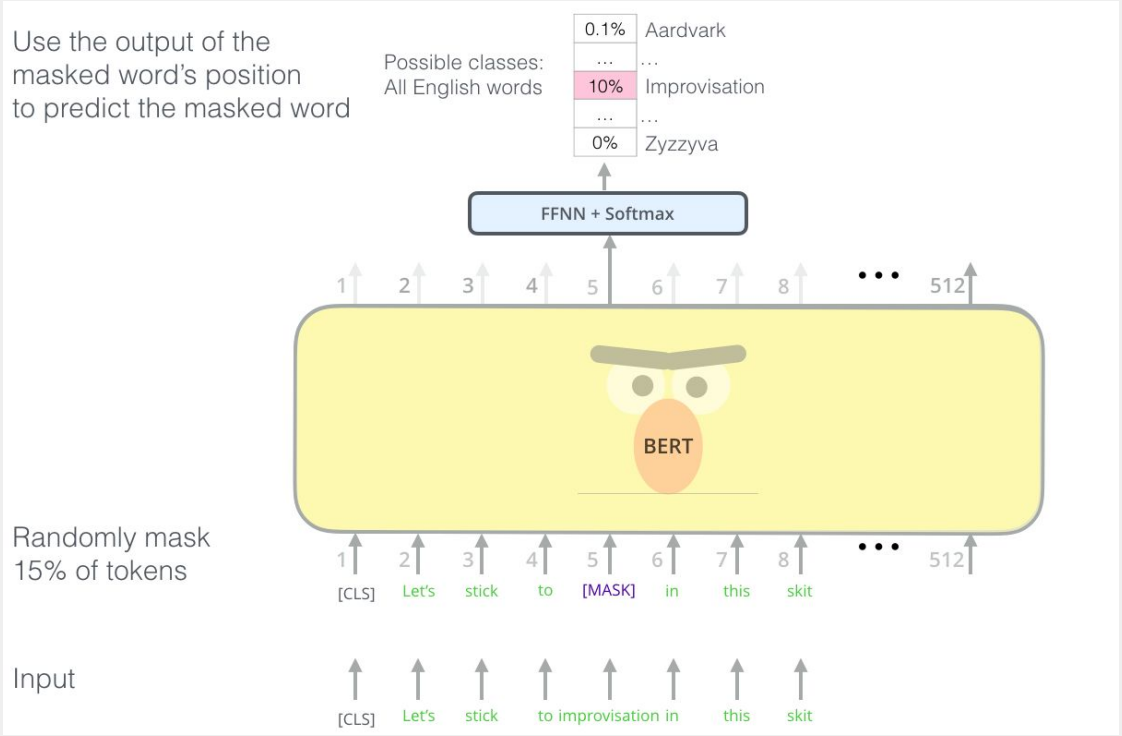

Example Task : Feedback Urgency

RED DRAGON AI

- Take a pre-trained BERT model
  - Use our annotated data
    - (Still need 100s of examples)
  - Train a new ‘head’ for the model
- Deploy combined model as an API
  - Run in a streaming setup
  - ... but also account for updated data, etc...

Pace of Change

RED DRAGON AI

- AI tech scene is moving fast
  - New models -> New capabilities
- Getting latest stuff into production...
  - ... needs people with many different skills
  - ... and the task is never ‘finished’
  - ... but (unlike manual labour) it scales ‘easily’

Our Courses

(Quick Advert)

Red Dragon Courses

RED DRAGON AI

- Currently a set of 5 courses
  - Covering different aspects of AI
- Held in conjunction with SGInnovate
  - Significant funding available
    - for SC/PR from IMDA
  - Link if you're interested :
    - https:// bit.ly / learn-more-ai2

Deep Learning Foundations Course

RED DRAGON AI

- Starting from the basics
  - 3 week-days ‘face-to-face’ + video content
  - Play with real models & Pick-a-Project
  - Knowledge / portfolio building, with Certificate
- Prerequisites :
  - Most participants already do some programming
    - Python element isn’t super-difficult for them
  - Willingness to learn ...

More Advanced Courses

RED DRAGON AI

All independent, but require Foundations

- #2 - **Advanced Computer Vision**
  - Covers all the CV models mentioned so far
- #3 - **Advanced NLP & Sequences**
  - Attention, Transformers : Very up-to-date
- #4 - **Self-Supervised Learning**
  - Super-hot area of research
- #5 - **Models to Production**
  - Serving models on servers, or mobile.
  - Processes for managing Data Science, etc

Conclusions

Conclusion

RED DRAGON AI

- Being a Developer is already a key skill
  - Add Deep Learning knowledge to target future
- There are many skills required to get into production...
  - Engineering roles (keeping models alive)
  - Data roles (training quickly gets ‘tricky’)
  - Strategy roles (understand where AI can be applied)
- Word on the Street :
  - MLOps is better paid than many other roles...

Q&A via Chat

martin@reddragon.ai
sam@reddragon.ai

Learn More : https:// bit.ly / learn-more-ai2