

Effective transfer learning for clinical applications

Benjamin van der Burgh (LIACS)



OVERVIEW

1. Transfer learning in NLP
2. Experiments on Dutch data
3. Well-being tracking using clinical journals



PROJECT BACKGROUND

- Physiotherapists keep journals
- Can we quantify well-being from text?
- Not a conventional task, **no labeled data**
- What can we do about it?



1

TRANSFER LEARNING

Learning with a head start

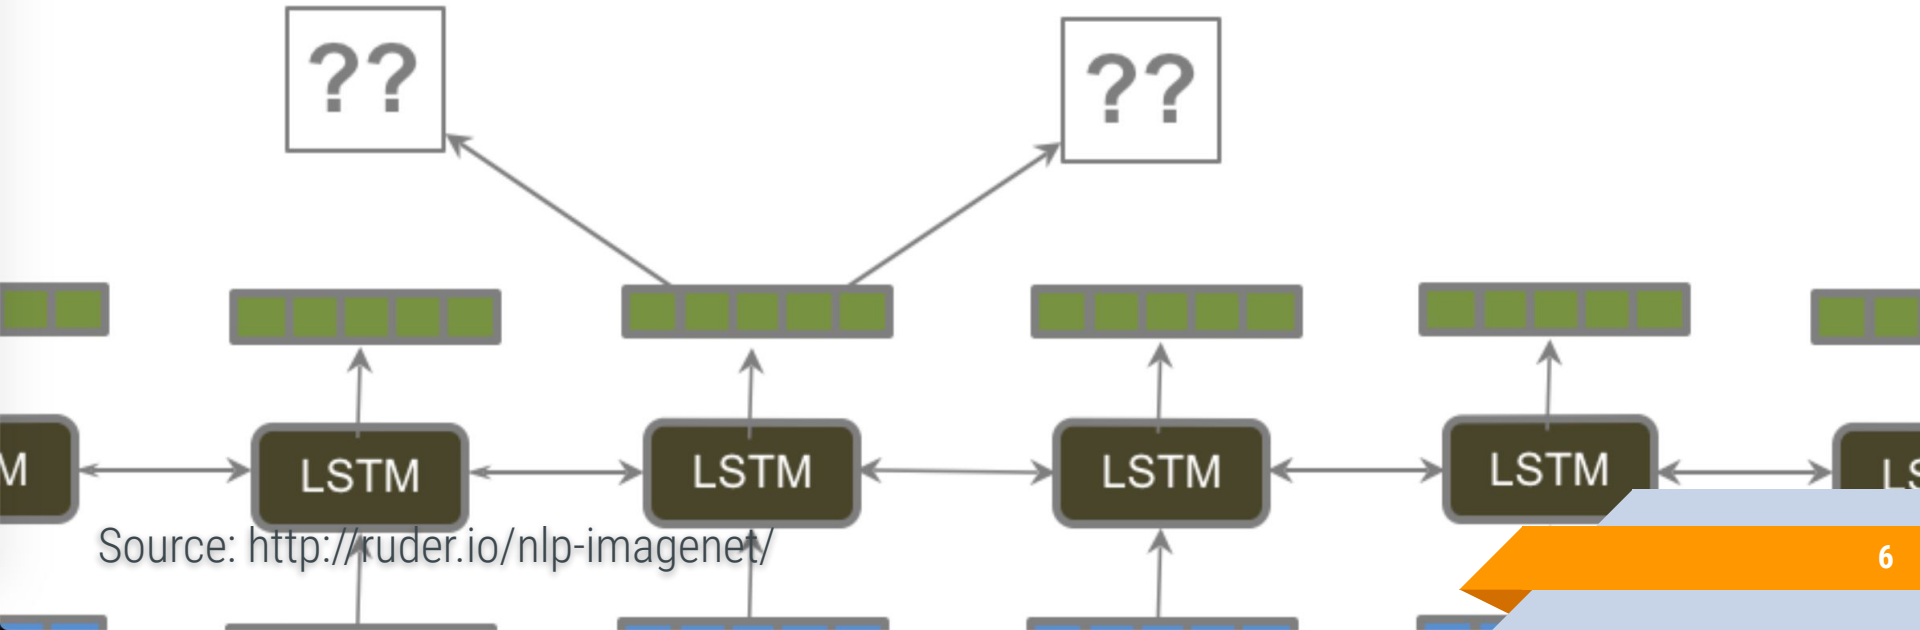


TRANSFER LEARNING

- Deep neural networks
- First train model for **different** but **similar task**
- Learns reusable representation / features
- Replace last layer(s) to adjust to target
- Continue training the model on target dataset

12 JULY 2018 / NATURAL LANGUAGE PROCESSING

NLP's ImageNet moment has arrived



Source: <http://ruder.io/nlp-imagenet/>




TRANSFER LEARNING IN NLP

- Generic task in NLP: **language modelling**
- Example: *"I'm not half the man I ..."*
- Dataset source: Wikipedia, CommonCrawl, etc.
- Suitable architecture
 - ▷ RNN-based: ULMFiT (AWD-LSTM)
 - ▷ Self-attention models: Transformer, BERT

**The
patient
feels**

Input



RNN

[0.12, 0.19, ...]

[0.31, 0.20, ...]

[0.01, 0.11, ...]

Output



FINE-TUNING LANGUAGE MODEL

- Adjust model to idiosyncrasies of target
- Example: “Patient has pain in the ...”
- Use language model as encoder for target



THREE-STAGE PROCESS

Generic
LM

Fine-tuned
LM

Target
Task

2

EXPERIMENTS

Transfer learning on Dutch data




EXPERIMENTS WITH ULMFIT

- Language model trained on Dutch Wikipedia
- Dataset of 110k Dutch book reviews [1]
 - ▷ {1, 2} → negative
 - ▷ {4, 5} → positive
 - ▷ {3} → neutral
- 18836 training examples, 50% pos / 50% neg



EXPERIMENTAL RESULTS

- Training language model took **days**
- Fine-tuning encoder took an **hour**
- Training classifier took **minutes**
- Accuracy 94%
- Off-the-shelf software and hardware





ADVANTAGES

1. Improved **data efficiency**
2. Models can be **shared**
3. Or even collaboratively trained
→ Federated Learning [1]

3

WELL-BEING TRACKING

Learning from subjective data



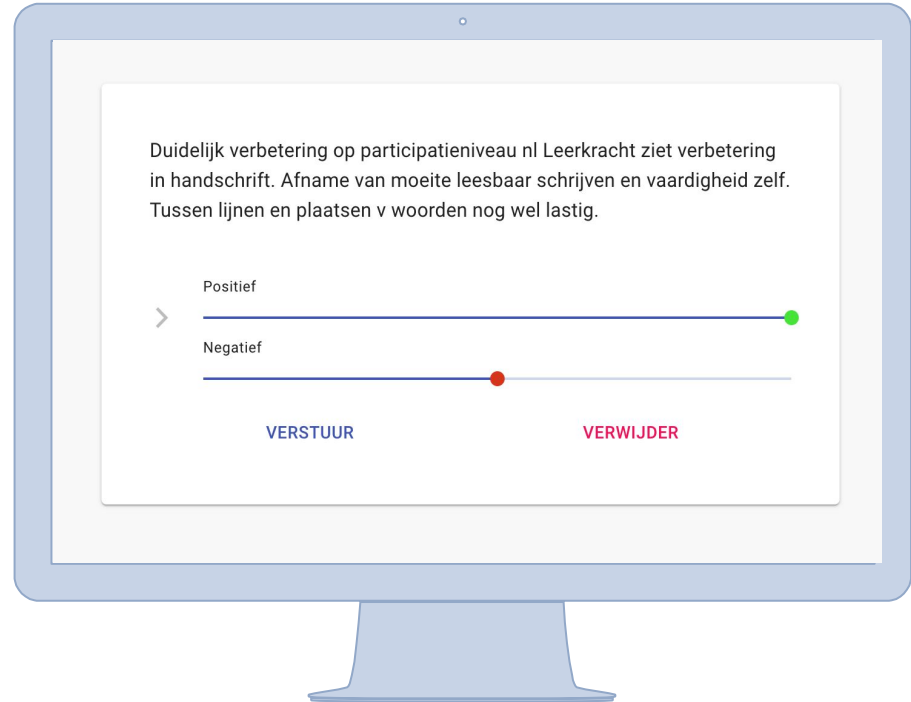
WELL-BEING TRACKING

- Well-being tracking using journal text (SOAP)
- Multivariate regression: positive and negative
- No labeled data available




LABEL DATA

Experts quantify the contents of a journal entry on a positive and negative axis.



Patient 6235



4

TAKEAWAYS



... no, not that kind of takeaway



SUMMARY

- Transfer learning in NLP possible
- State-of-the-art while easy-to-use
- Unlock knowledge in subjective data
- Models can be shared



RELATED WORK

- Bert-as-a-service [1]
- Self-supervised learning for image data [2]
- Sentiment analysis using text in psychiatry [3]

[1] bert-as-a-service: <https://github.com/hanxiao/bert-as-service>

[2] Selfie: Self-supervised Pretraining for Image Embedding: <https://arxiv.org/abs/1906.02940>

[3] Distinguishing Clinical Sentiment: The Importance of Domain Adaptation in Psychiatric Patient Health Records: <https://arxiv.org/abs/1904.03225>



FURTHER RESEARCH

- Can **privacy be preserved** when models are shared?
- How can we make machine learning more accessible?
- What can be learned from subjective data?
- How to explain 'deep results'?




SHARE MODELS

Help patients while preserving privacy

You can download mine from: <https://github.com/benjaminvdb/110kDBRD>





THANKS!

Any questions?

You can find me at

@BenjaminBurgh & b.van.der.burgh@liacs.leidenuniv.nl