



KNIME Deep Learning Workshop

Christian Dietz & Adrian Nembach
KNIME


Why Deep Learning

- Multi-layer neural networks for learning tasks on images, videos, text and sound
- Learns useful representations from raw data
- Extremely flexible: One framework for classification, regression and unsupervised learning


Outline of this Workshop

- Deep Learning in a Nutshell
- KNIME Deep Learning Integrations
- Use-cases
 - Image classification with Convolutional Neural Networks
 - Language modeling with Recurrent Neural Networks
 - Neural machine translation with Recurrent Neural Networks

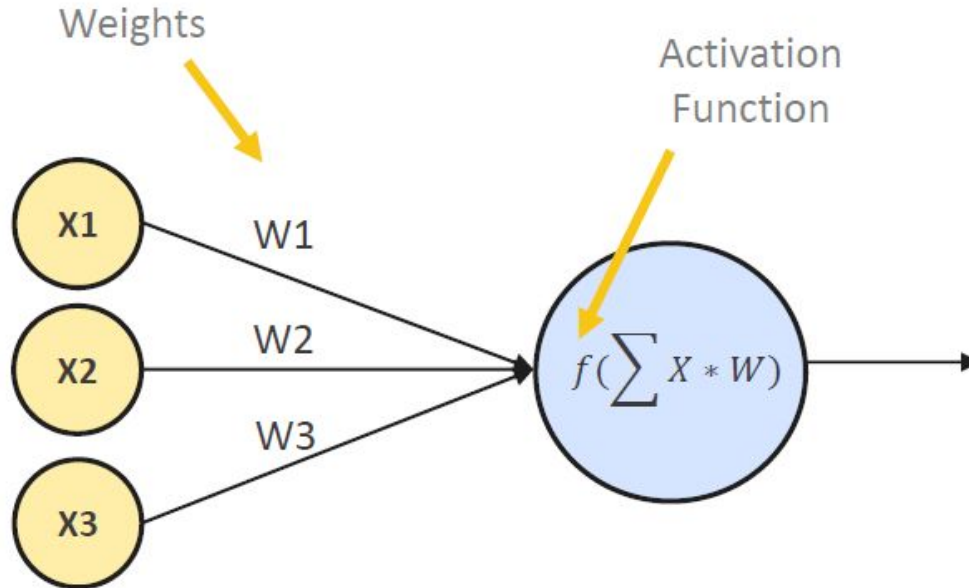
Deep Learning in a Nutshell



Deep Learning in a Nutshell



Multi-Layer Neural Networks



Training of Deep Networks

Row ID	Image	Actual...
Row0	5	5
Row1	0	0
Row2	4	4
Row3	1	1
Row4	9	9
Row5	2	2
Row6	1	1
Row7	3	3
Row8	1	1
Row9	4	4
Row10	3	3
Row11	5	5
Row12	3	3
Row13	6	6
Row14	1	1

Batch 1

Batch 2


Batch 3

Training of Deep Networks

Row10	3	3
Row11	5	5
Row12	3	3
Row13	6	6
Row14	1	1

Row10	3	3
Row11	5	5
Row12	3	3
Row13	6	6
Row14	1	1

Row10	3	3
Row11	5	5
Row12	3	3
Row13	6	6
Row14	1	1



Row 10 is a 3
Row 11 is a 5

...

***Correct?
Adjust weights.***

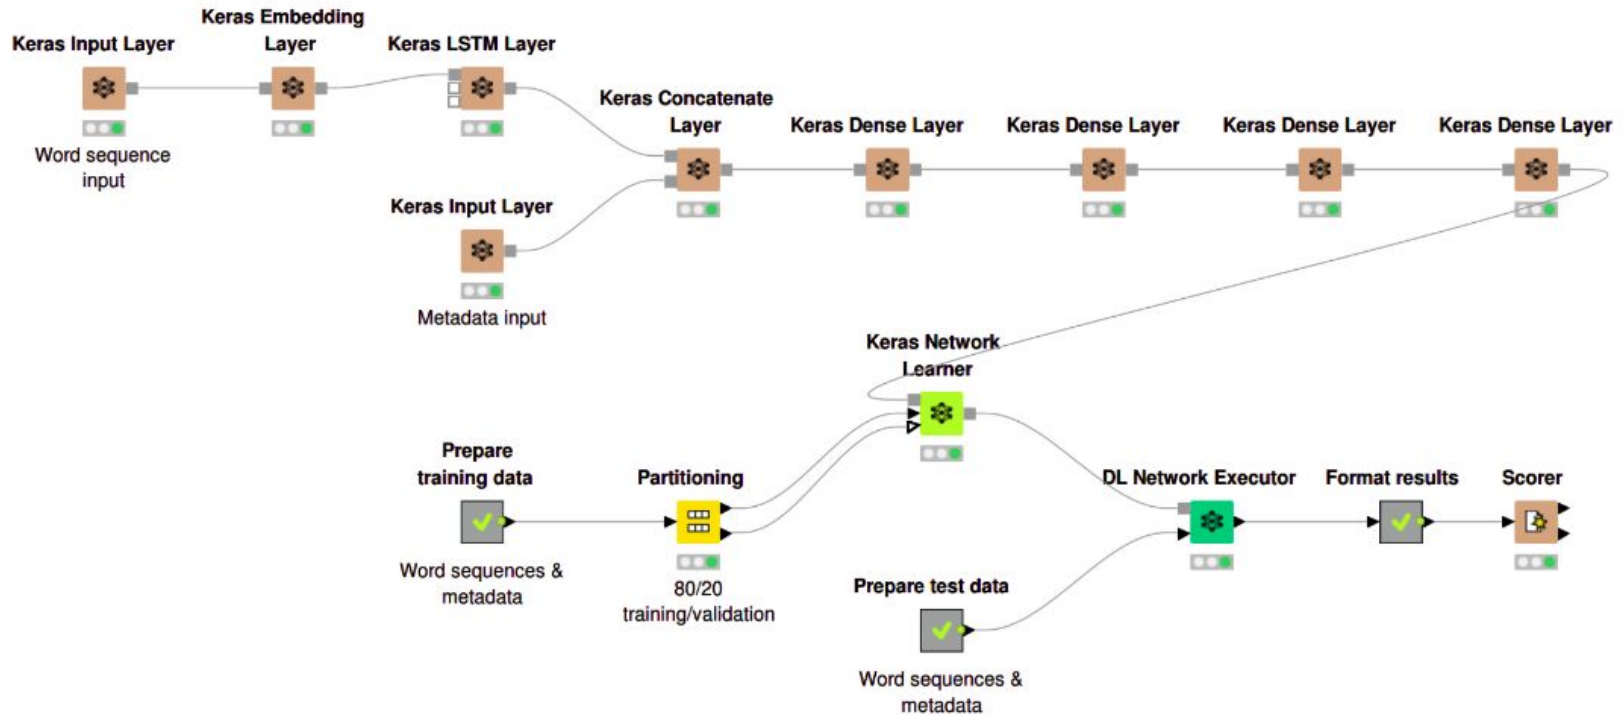
KNIME Deep Learning Integrations



KNIME Deep Learning Integrations

- Keras
 - Network creation with KNIME nodes
 - Training with the Keras Network Learner node
 - Execution with the DL Executor node
 - Save and load with Keras Reader and Writer nodes
- TensorFlow
 - Creation and training with DL Python nodes
 - Conversion of trained Keras networks
 - Fast execution with the DL Executor node
 - Save and load with TensorFlow Reader and Writer nodes

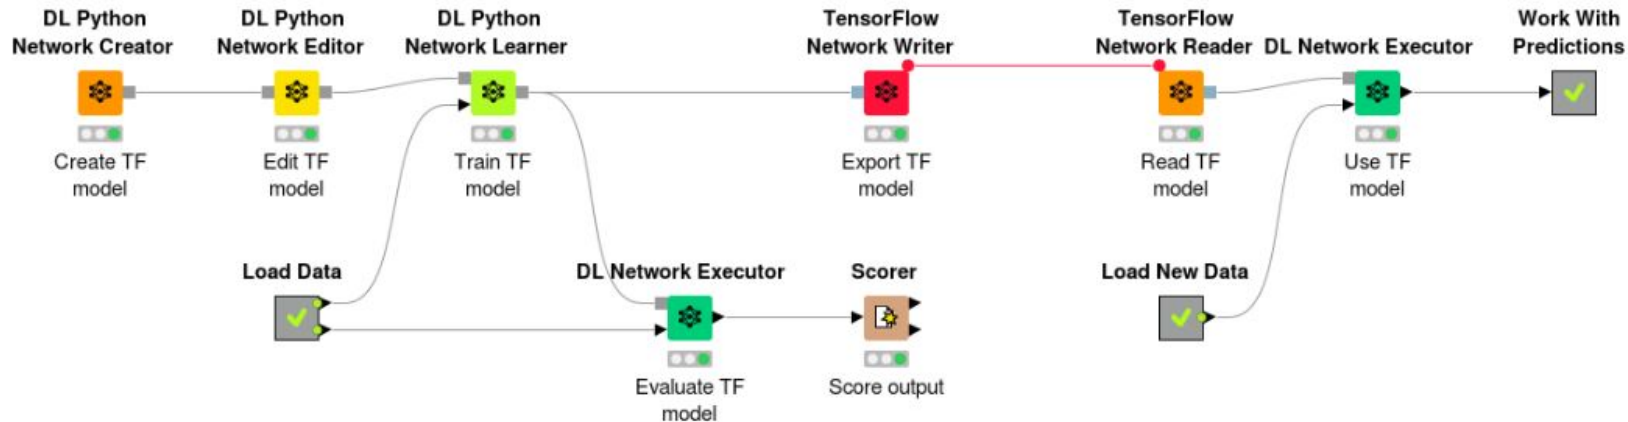
KNIME Keras Integration



KNIME TensorFlow Integration


Training

Deployment




Use-Cases

Image classification



Language modeling



Neural Machine Translation

Image Classification



Image Classification

- Task: Given an image predict a class label
- Supervised Learning
- Realized with Convolutional Neural Networks



Cat

Convolutional Neural Networks

- For high-dimensional inputs like images and video
- Replaces global matrix multiplication with local convolution


Possible Strategies

- Train CNN from scratch
- Fine-tune pretrained CNN to the new task
 - VGG16, Inception, Resnet, ...
- Use pre-trained CNN as feature extractor and build a different classifier on top
 - Native KNIME nodes, H2O, XGBoost, ...

Deep Learning in Action




Language Modeling



Language Modeling


- Task: Given a sequence of characters predict which character is the next in the text
- Unsupervised Learning
- Realized with Recurrent Neural Networks




Recurrent Neural Networks

Useful for sequence data like text, time series, audio and video

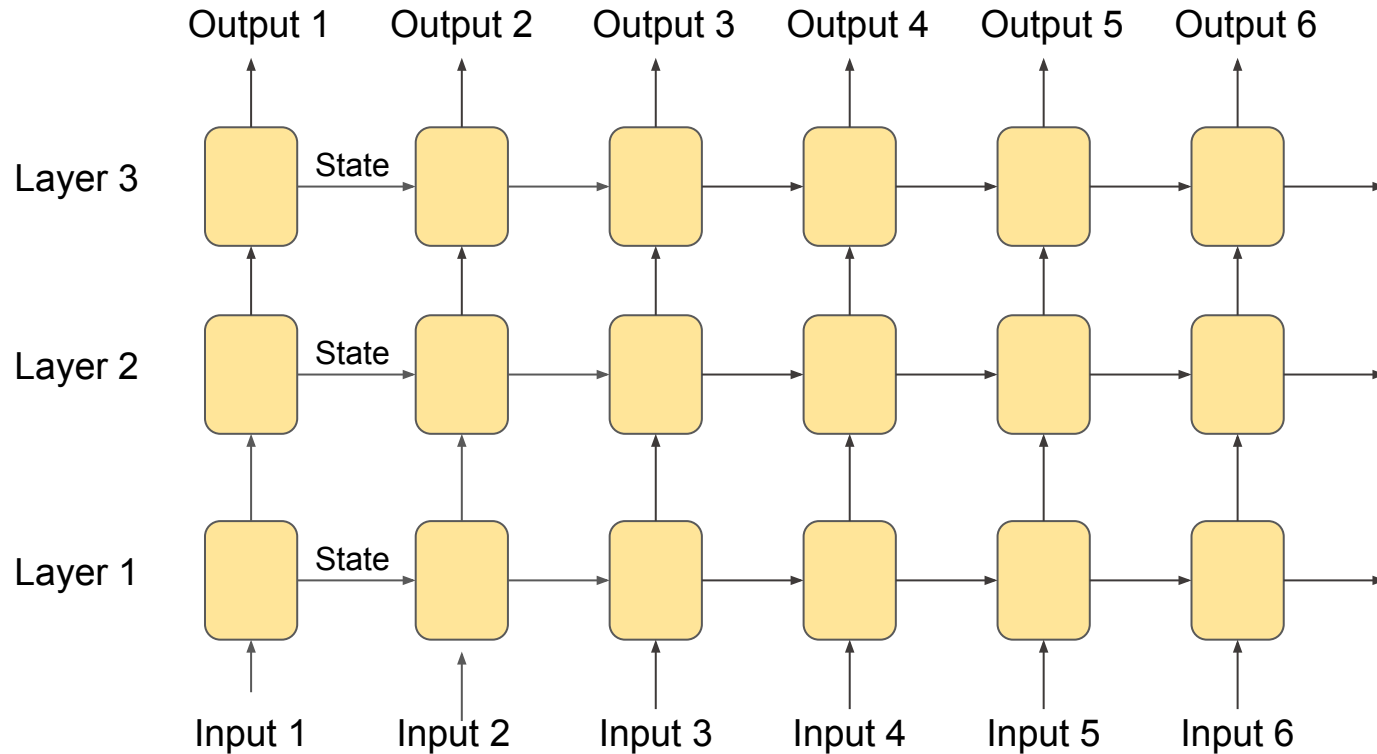
MLP and CNN




RNN



Stacked RNNs



Language Modeling with RNNs



Deep Learning in Action



Neural Machine Translation



Neural Machine Translation

- Task: Given an English input sentence predict a German translation
- Supervised Learning
- Realized with Recurrent Encoder-Decoder networks

How are you?

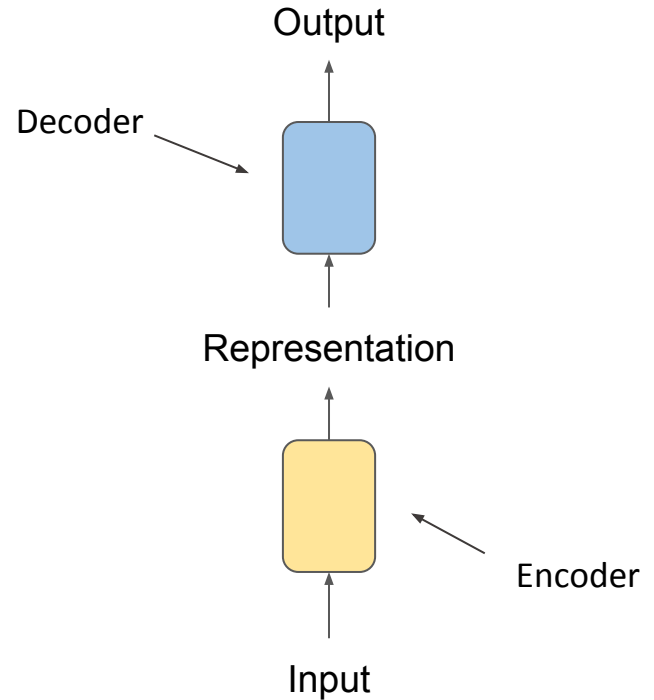


Wie geht es dir?


Encoder-Decoder Networks

Useful for:


- Autoencoders
- Semantic segmentation
- Neural machine translation



Recurrent Encoder-Decoder Networks



Deep Learning in Action



Conclusion

Deep Learning in KNIME allows to

- create and train sophisticated Keras networks
- fine-tune pre-trained networks like Inception
- mix and match Keras and TensorFlow networks with other KNIME nodes and integrations
- Export the networks for deployment with other frameworks

Thank You for coming!
Questions?



The KNIME® trademark and logo and OPEN FOR INNOVATION® trademark are used by KNIME AG under license from KNIME GmbH, and are registered in the United States. KNIME® is also registered in Germany.