What is AI?

An introduction to the basics of artificial intelligence in the context of NMCAs

apl. Prof. Dr. Franz Rottensteiner

Institute of Photogrammetry and GeoInformation
Leibniz Universität Hannover
rottensteiner@ipi.uni-hannover.de

Contents

What is Artificial Intelligence?

- Oxford English Dictionary (https://www.oed.com/):

  "*The capacity of computers or other machines to exhibit or simulate intelligent behaviour.*"

- Very broad field, including
  - Perception (visual / speech / …)
  - Decision making, planning

- Often associated with Machine Learning, but: AI is broader

- Colloquially, "AI" refers to applications of (deep) neural networks (**deep learning**) for complex tasks

- **Relevance for NMCAs: Automation of mapping processes from remote sensing data based on deep learning**

Artificial Intelligence vs. Deep Learning

- Deep learning: one specific but very powerful strategy of AI

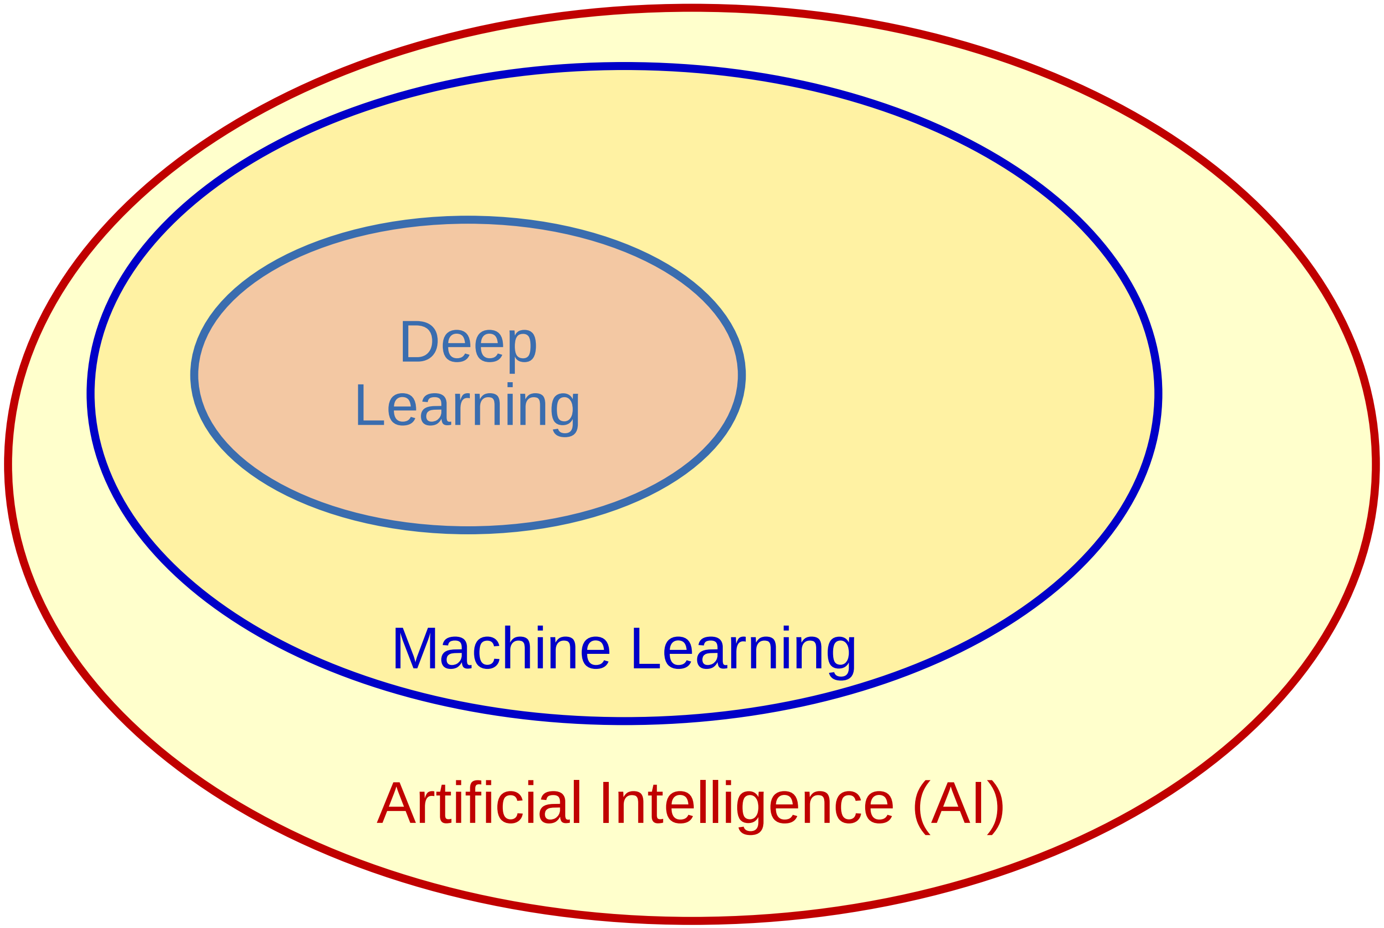

adapted from [Goodfellow et al., 2016]

What‘s Different in Deep Learning?

- Machine Learning, classical approach:

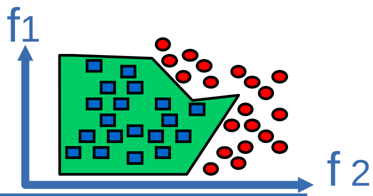

Input → Compute cleverly defined “hand-crafted” features → $f_1 \ldots f_n$ → Machine Learning, Classification → Output

performance only as good as feature set

- **Deep Learning**: Joint learning of features and classification model

Input → Deep Learning [Feature extraction | classification] → Output

  – “End-to-end learning“, based on artificial neural networks (ANN)

Contents

- Introduction: What is AI?
- Deep Learning
- Applications for NMCAs: classification of land cover and land use
- Stumbling blocks and potential solutions
- Conclusions

Development of Deep Learning

- Cybernetics (Perceptron)
  - *f*: step function; linear decision boundary [McCulloch & Pitts, 1943] [Rosenblatt, 1958]

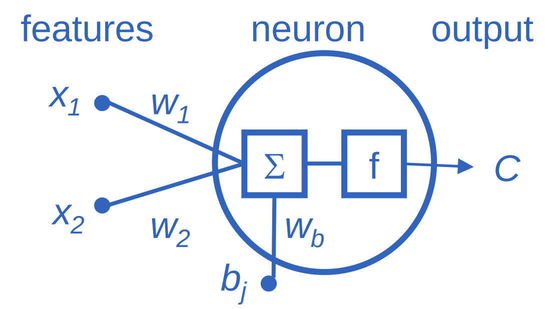

- ANN, Multi-layer Perceptron (MLP)
  - *f*: sigmoid/tanh; non-linear decision b.
  - back propagation for training [Rumelhart et al., 1986]

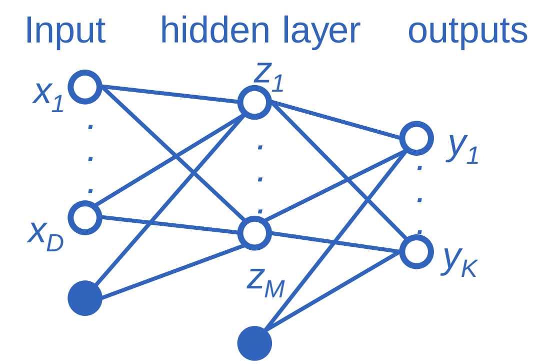

- **Deep learning:** learn representation (features)
  - Fast hardware (GPU), vast amount of training data (ILSVRC)
  - Minor algorithmic changes, e.g. ReLu activations for *f*
  - LeNet-1 [LeCun et al., 1989]; AlexNet [Krizhevsky et al., 2012]

Convolutional Neural Networks (CNN): Components

- Convolutional Layers

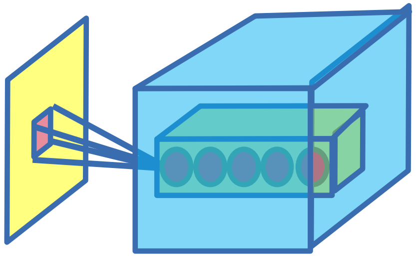

- (max-)pooling

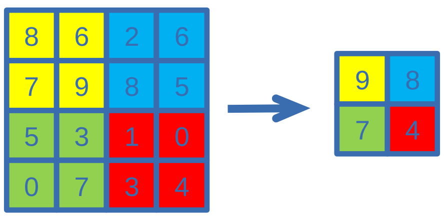

- activation function

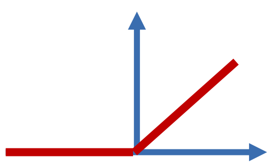

e.g. ReLU(x) = max(0,x)

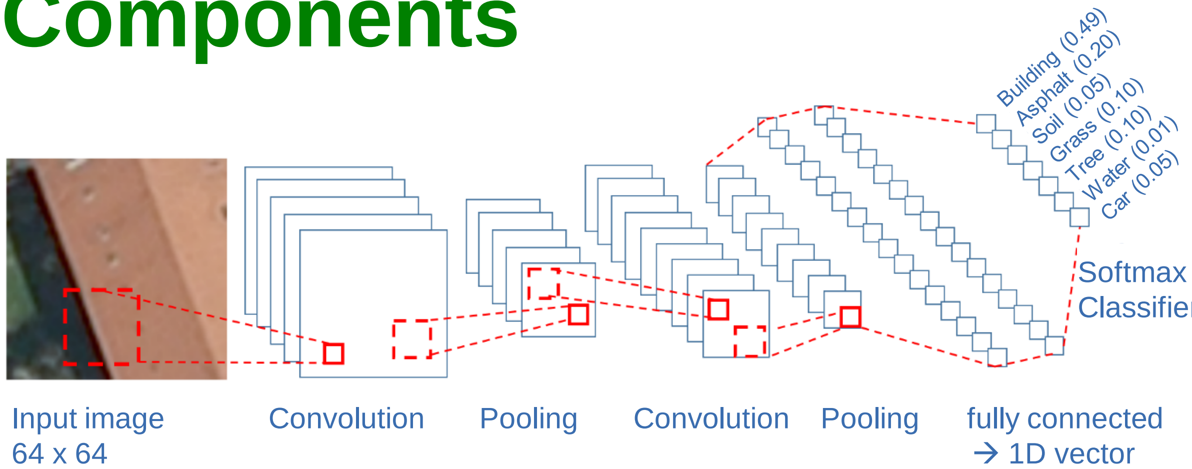

- For regularly structured data
- Repeated execution of those steps based on “some“ architecture
- Training: determine filter coefficients by minimizing a loss function
- Classification: one label per patch

CNN and Depth

- The depth of CNN has increased considerably over time
- ILSVRC: number of layers and top-5 errors for image classification (one label per image)

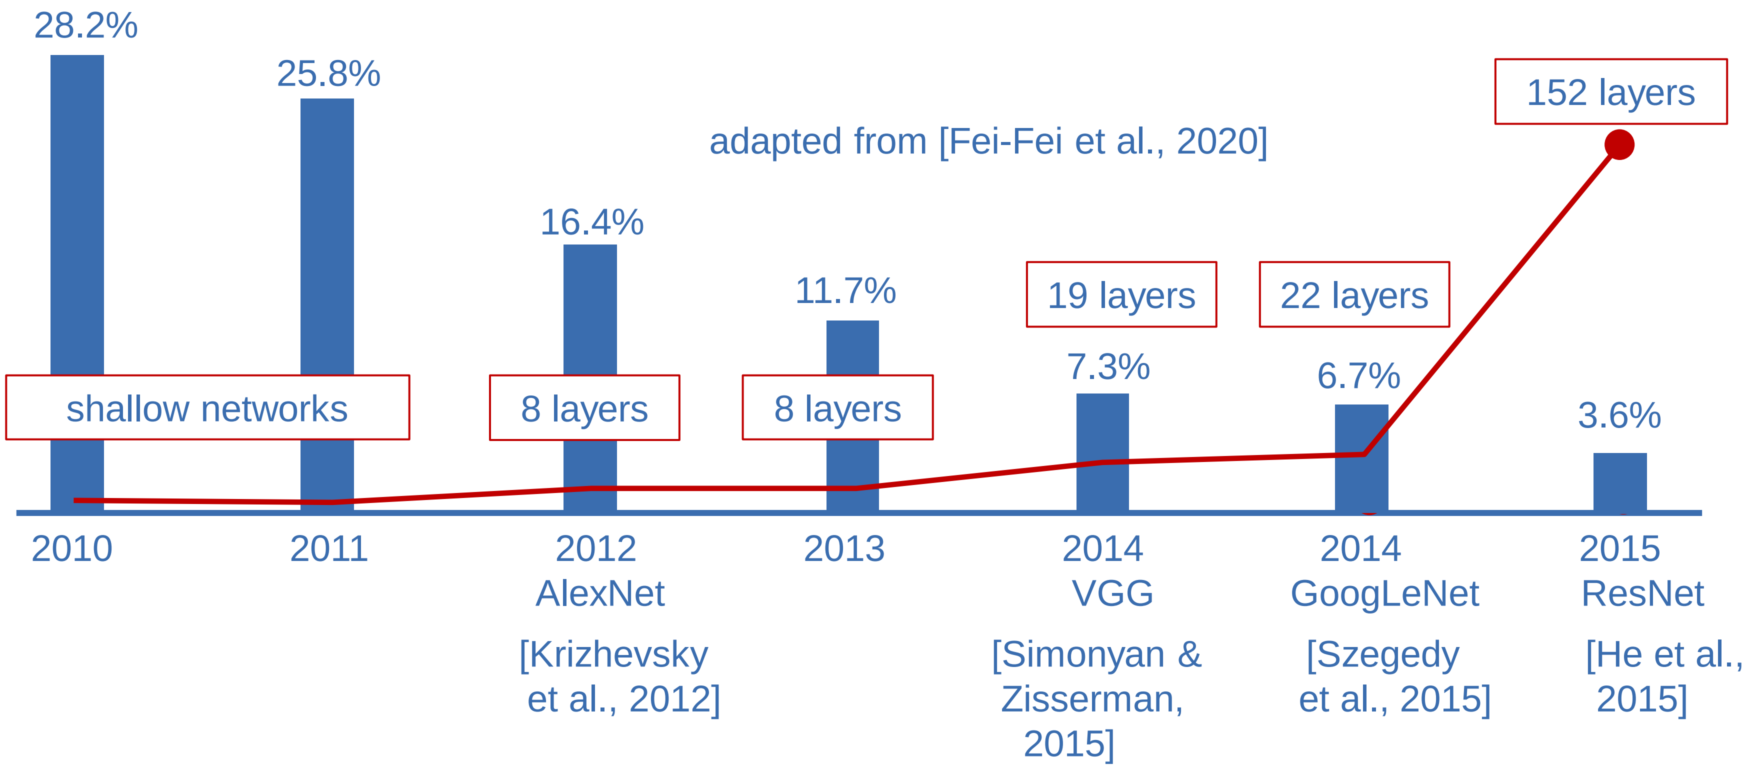

Classification (CV Terminology), Examples

airplane
automobile
bird
cat
deer
dog
frog
horse
ship
truck

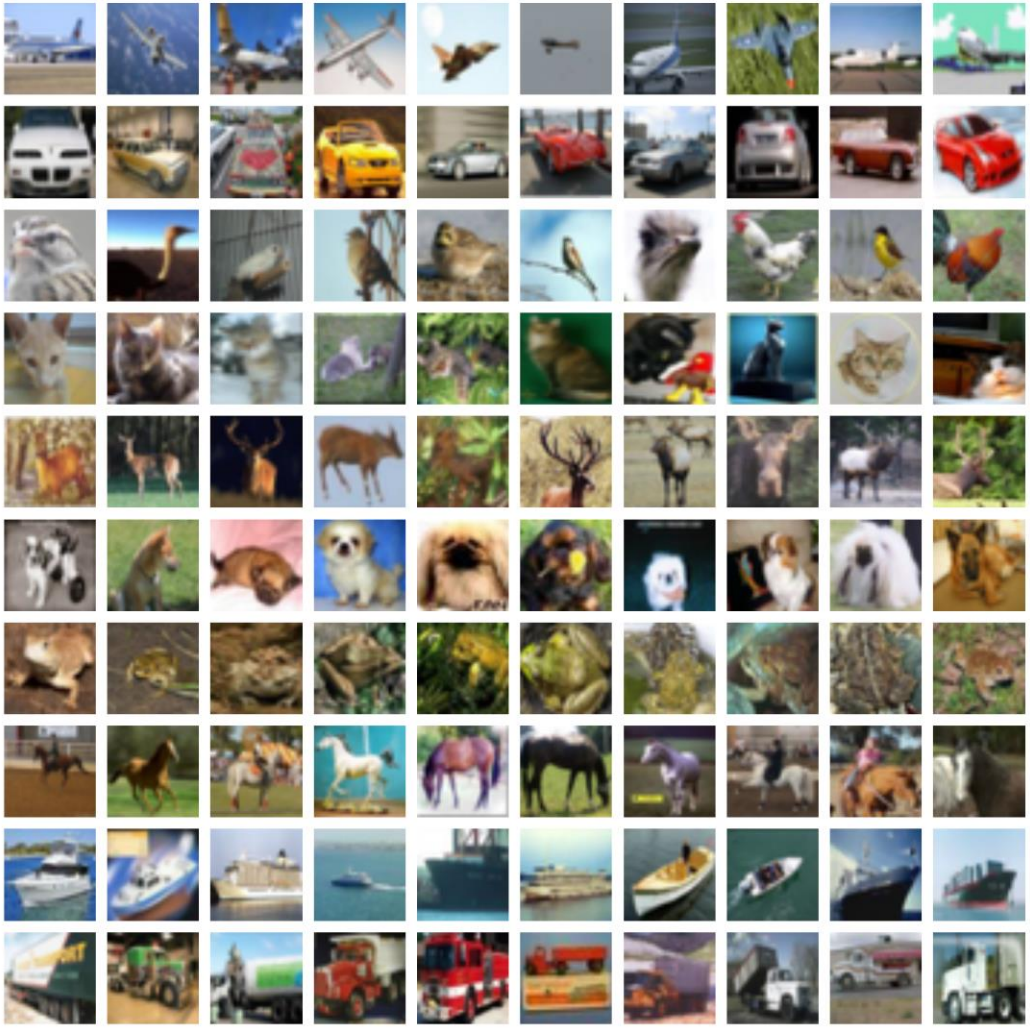

Pixel-wise classification: FCN

- **Fully convolutional network (FCN)** [Long et al., 2015]:
  - Apply convolutions and pooling to entire images
    → representation at reduced resolution
  - **Upsampling for pixelwise predictions**

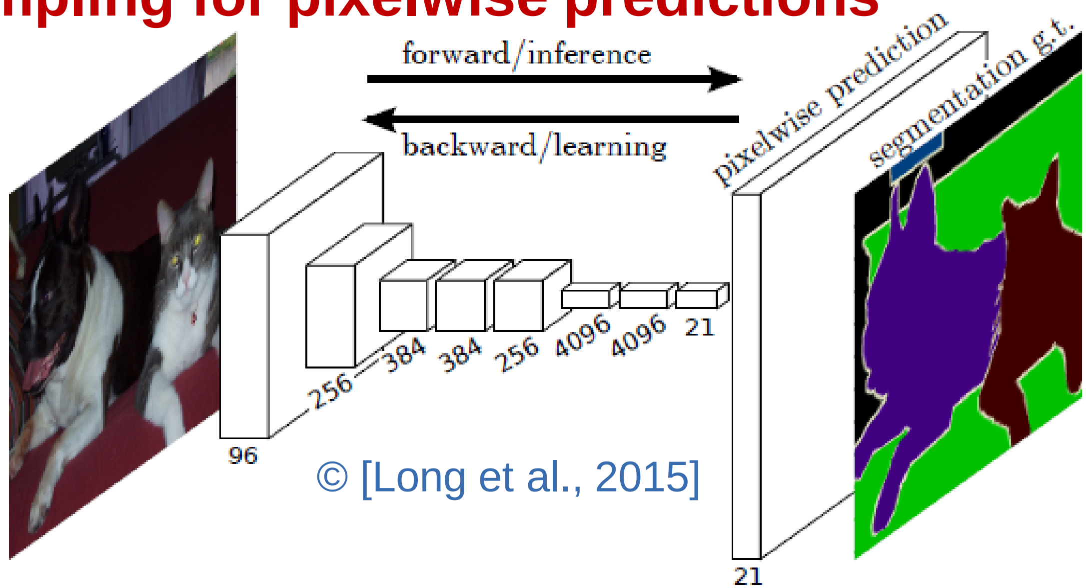

© [Long et al., 2015]

- **Encoder-Decoder networks**, e.g. **U-Net** [Ronneberger et al.,2015]:
  - Symmetrical structure of downsampling and upsampling layers

Contents

- Introduction: What is AI?
- Deep Learning
- Applications for NMCAs: classification of land cover and land use
- Stumbling blocks and potential solutions
- Conclusions

Example: Land Cover Classification

- **Goal:** assign a class label to every pixel of the input data
- Land cover: material / type of the object “seen” by the sensor
  - Examples: building, grass, tree, asphalt, car

- Input: remote sensing data, e.g.
  - Aerial / satellite images
  - Digital surface models (DSM)

CNN for Land Cover Classification

- CNN for aerial imagery (20 cm GSD), DSM [Yang et al., 2020a]

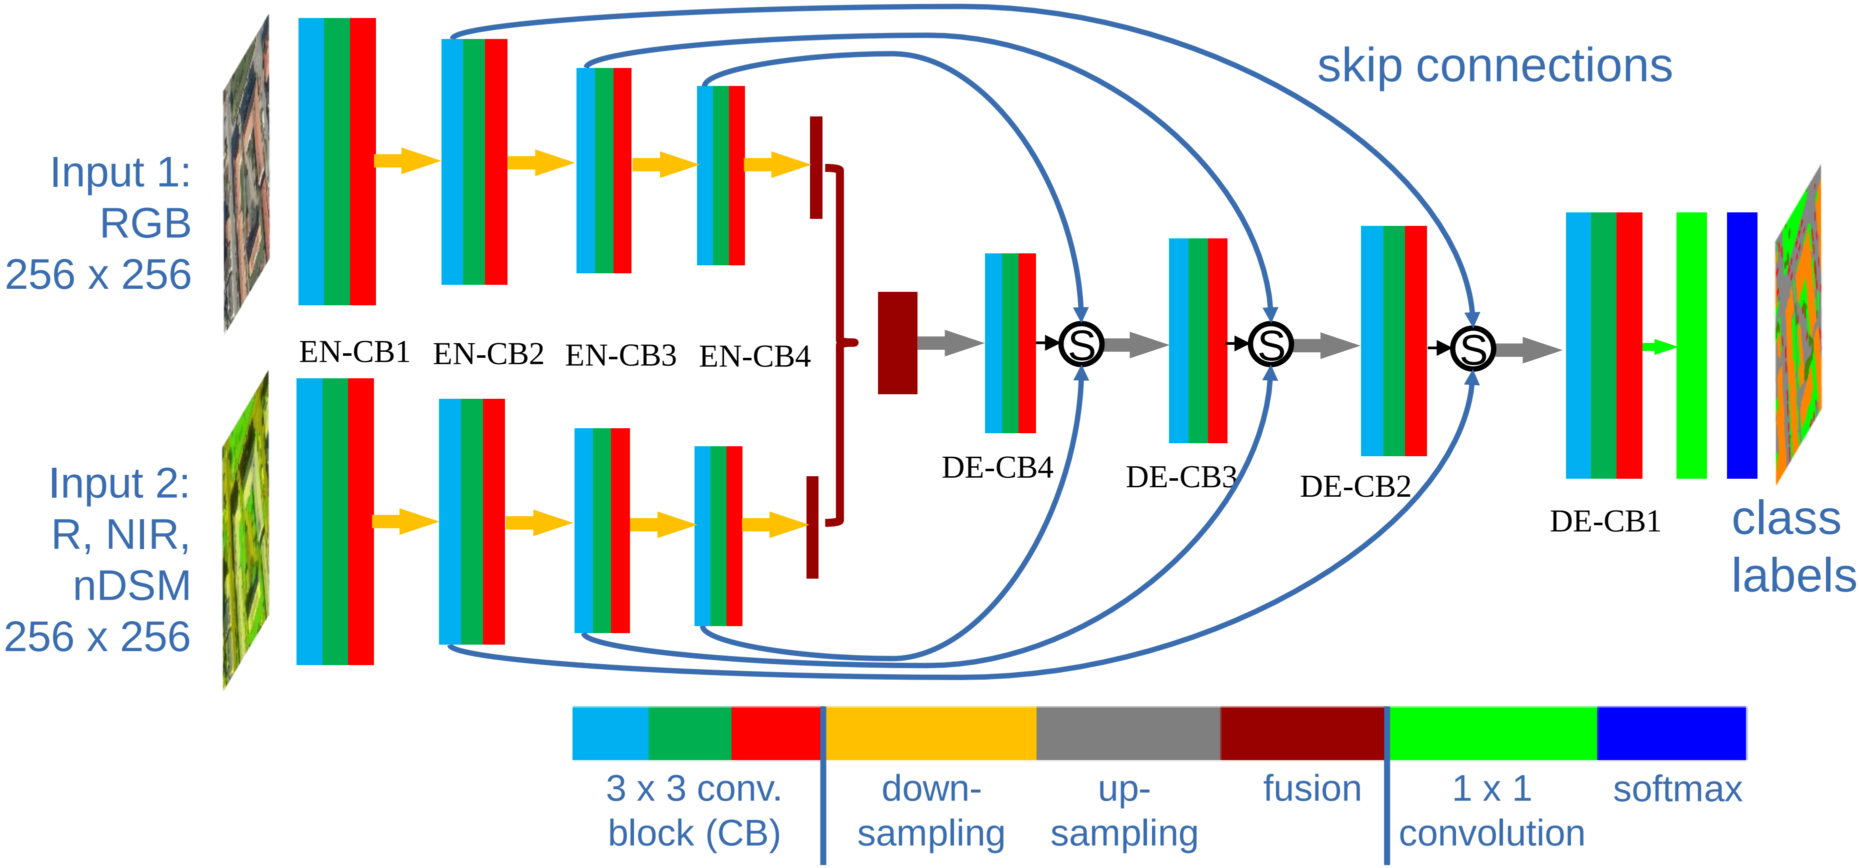

- **Overall accuracy: ~85%-90%**

Example: Land Use Classification for Database Verification

- **Goal:** predict land use for every polygon of an existing database
  - Examples: residential, park, traffic

- **Input:**
  - Image and height data
  - **Pixel-wise land cover information**
- **Problem:** Very detailed object catalogues (> 150 classes!)
  - → Predict class labels at multiple semantic levels
  - → Enforce consistency with the hierarchical catalogue

Hierarchical Land Use Classification

- **Land Use Network** [Yang et al., 2020b]

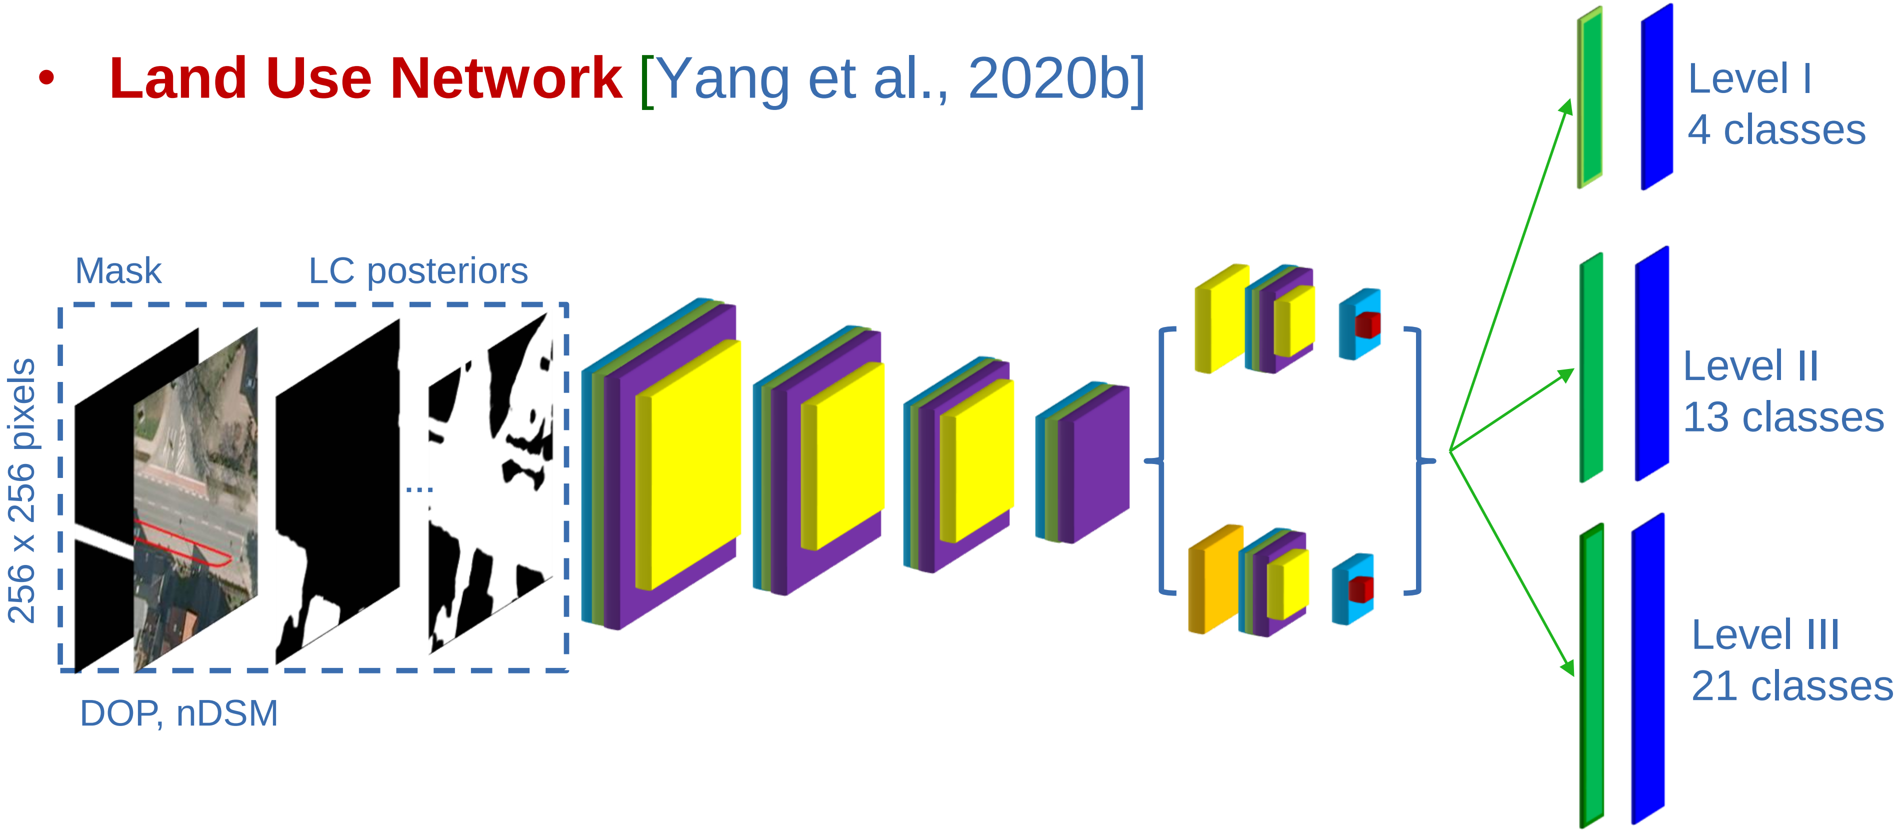

convolution blocks	max-pooling	average pooling	ROI location	Multi-task classification head	softmax

- **Overall Accuracy: 92% (4 cl.) / 77% (13 cl.) / 73% (21 cl.)**

Contents

Deep Learning in Remote Sensing: Problems and Potential Solutions

- **Imbalanced class distribution in the training samples**
  - Adapt loss functions [Yang et al., 2018, 2020a]
- **Structure of object catalogues**
  - Reconsider class structures and make “pragmatic choices” (© George Vosselman) to simplify automation
- **Shortage of training samples**
  - **Data augmentation**: use synthetic samples [Wittich, 2020]
  - **Domain adaptation**: use training samples from another domain (e.g. other regions) [Tuia et al., 2016; Wittich, 2020]
  - **Use existing map for training**: classifier has to cope with **label noise** [Maas et al., 2018, 2019; Voelsen et al., 2002]

Example: Training Based on Existing Map

- Land Cover classification based on Sentinel-2 (GSD: 10 m), based on [Voelsen et al., 2020]
- Data covering the **entire state** of Lower Saxony (47,600 km^2) at 10 m
  - 14 epochs (different seasons)
  - Class labels from the German landscape model (6 classes)
- Estimated level of label noise: 8%
- **Preliminary** results:
  - Overall accuracy: 90%
  - Hardly any impact of label noise IF all epochs are used

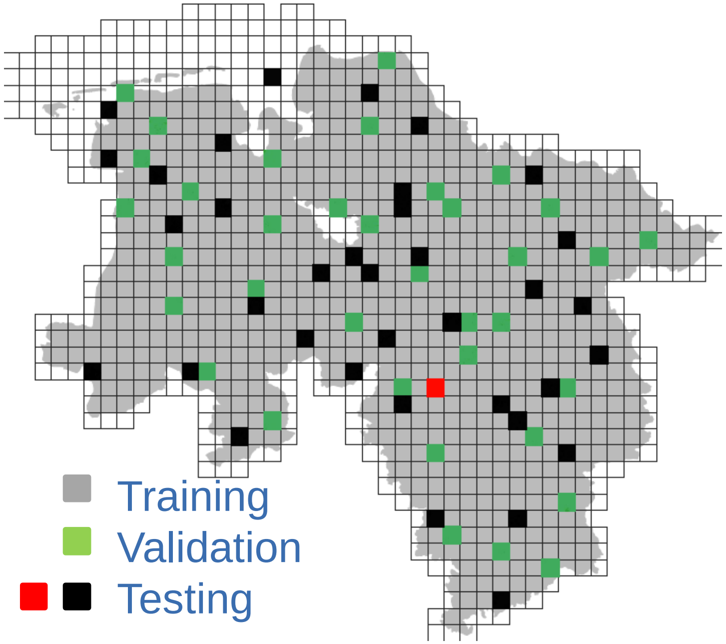

Contents

Artificial Intelligence for NMCAs?

- **Deep Learning for automating the update process of geospatial databases!**
- Deep learning has been adapted for such applications in our domain for some time [Zhu et al., 2017]
  - It has outperformed just about all classical algorithms by a large margin, provided enough training data are available
  - Strength: learning features, classifier not so important
  - Key to good performance: network depth
- Consider techniques to reduce the requirements w.r.t. training data
- Current work focusses on **classification** only; hardly any work on **vectorization**

References I

- Fei-Fei, L., Johnson, J., Yeung, S.: Convolutional Neural Networks for Visual Recognition. Lecture, Stanford University, http://cs231n.stanford.edu/syllabus.html (accessed 3rd February 2021), 2020.
- Goodfellow, I., Bengio, Y., Courville, A.: Deep Learning. MIT Press, Cambridge, MS, USA, 2016.
- He , K., Zhang, X., Ren, S., Sun, J.: Delving Deep into rectifiers: Surpassing human-level performance on ImageNet classification. *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 1026-1034, 2015.
- Kirilov, A., He, K., Girshick, R., Rother, C., Dollár, P., 2019: Panoptic segmentation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 9396-9405.
- Krizhevsky, A., Sutskever, I., Hinton, G. E.: ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems* 25 (NIPS'12), Volume 1, pp. 1097-1105, 2012.
- LeCun, Y., Boser, B., Denker, J. S., Henderson, D., Howard, R. E., Hubbard, W., Jackel, L. D.: Handwritten digit recognition with a back-propagation network. *Proc. 2nd International Conference on Neural Information Processing Systems (NIPS'89)*, pp. 396-404, 1989.
- Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- Maas, A., Rottensteiner, F., Alobeid, A., Heipke, C.: Multitemporal classification under label noise based on outdated maps. *Photogrammetric Engineering & Remote Sensing* 84(5):263-277, 2018.

References II

- Maas, A., Rottensteiner, F., Heipke, C.: A label noise tolerant random forest for the classification of remote sensing data based on outdated maps for training. *Computer Vision and Image Understanding* 188, paper 102782, 2019.
- McCulloch, W. S., Pitts, W.: A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics* 5: 115-133, 1943.
- Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, Springer, LNCS, Vol. 9351, pp. 234-241, 2015.
- Rosenblatt, F.: The Perceptron: A probabilistic model for information storage and organization in the brain. *Psychological Review* 65(6): 386–408, 1958.
- Rumelhart, D. E., Hinton, G. E., Williams, R. J.: Learning representations by back propagating errors. *Nature* 323(6088): 533-536, 1986.
- Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *Proc. 3rd International Conference for Learning Representations*, 2015.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erha, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- Tuia, D., Persello, C., Bruzzone, L.: Domain adaptation for the classification of remote sensing data: an overview of recent advances. *IEEE Geoscience and Remote Sensing Magazine* 4(2): 41-57, 2016.

References III

- Tzeng, E., Hoffman, J., Saenko, K., Darrell, T.: Adversarial Discriminative Domain Adaptation. Proc.IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2017.
- Voelsen, M., Bostelmann, J., Maas, A., Rottensteiner, F., Heipke, C.: Automatically generated training data for land cover classification with CNNs using Sentinel-2 images. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLIII-B3-2020: 767–774, 2020.
- Wittich, D.: Deep domain adaptation by weighted entropy minimization for the classification of aerial images. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* V-2:591–598, 2020.
- Wittich, D., Rottensteiner, F.: Adversarial domain adaptation for the classification of aerial images and height data using convolutional neural networks. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* IV-2/W7:197–204, 2019.
- Wu, R., Shan, Y., Dang, Q., Sun, G.: Deep image: scaling up image recognition. ArXiv preprint, arXiv:1501.02876v5, 2015a.
- Yang, C., Rottensteiner, F., Heipke, C.: Classification of land cover and land use based on convolution neural networks. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* IV-3:251-258, 2018.
- Yang, C., Rottensteiner, F., Heipke, C.: Investigations on skip-connections with an additional cosine similarity loss for land cover classification. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* I V-3:339–346, 2020a.

References IV

- Yang, C., Rottensteiner, F., Heipke, C.: Exploring semantic relationships for hierarchical land use classification based on convolutional neural networks. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences* V-2:599–607, 2020b.
- Zhu, X., Tuia, D., Mou, L., Xia, G.-S., Zhang, L., Xu, F., Fraundorfer, F.: Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine* 5(4): 8-36, 2017.