Promises and challenges of deep learning for remote sensing image exploitation

Christian Heipke
IPI - Institut für Photogrammetrie und GeoInformation
Leibniz Universität Hannover

President, ISPRS

Table of contents

- **Introduction**
- **Deep learning**
  - **for images**
  - **for remote sensing**
- **Conclusions**

Global ICT trends

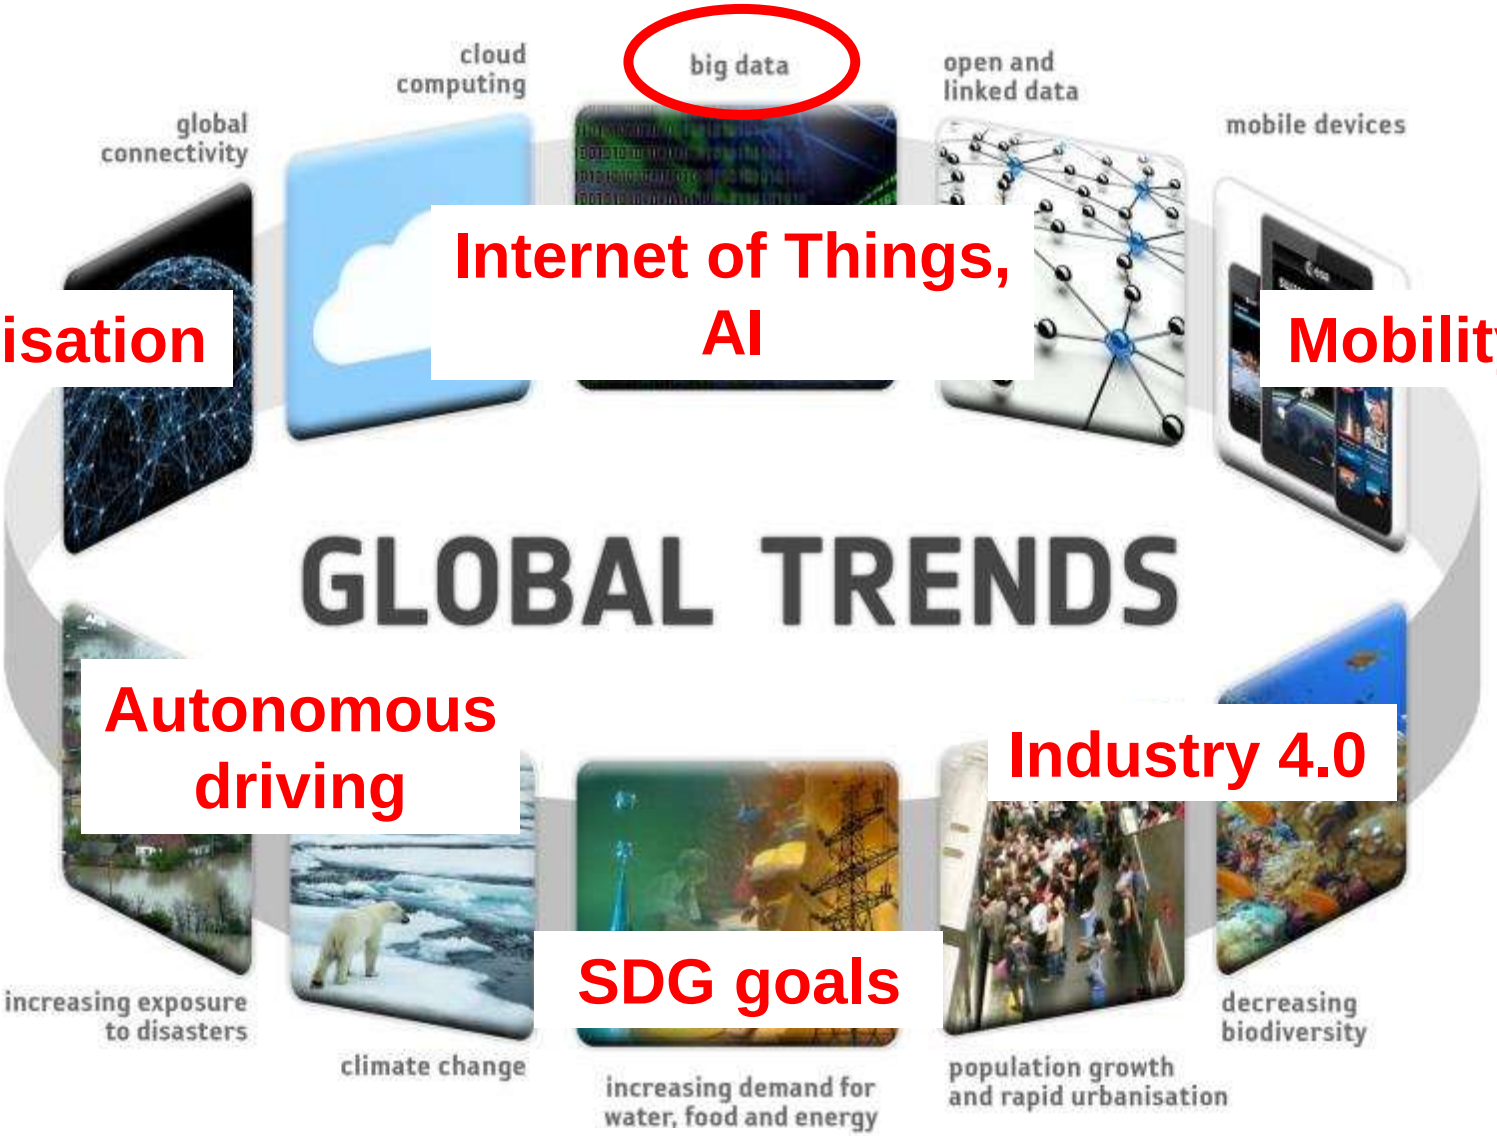

Digitisation

Internet of Things, AI

Mobility

Autonomous driving

Industry 4.0

SDG goals

adapted from Volker Liebig, ESA 2016

HANNOVER MESSE

April 1-5, 2019

ARTIFICIAL INTELLIGENCE

Machine learning: artificial intelligence in Industrie 4.0

ARTIFICIAL INTELLIGENCE

ZF establishes an AI research center

The automobile industry supplier wants to focus its activities in relation to artificial intelligence (AI) in Saarbrücken. One key area of... more

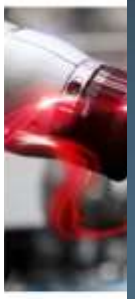

Your Ver

Cobots, d
manufact
solutions
at HANN
find all th

© HAZ, 2.4.2019

AM -

nmit

gence
Platfo
e key
ocesse

© Bamler, 2018

Increased data volume

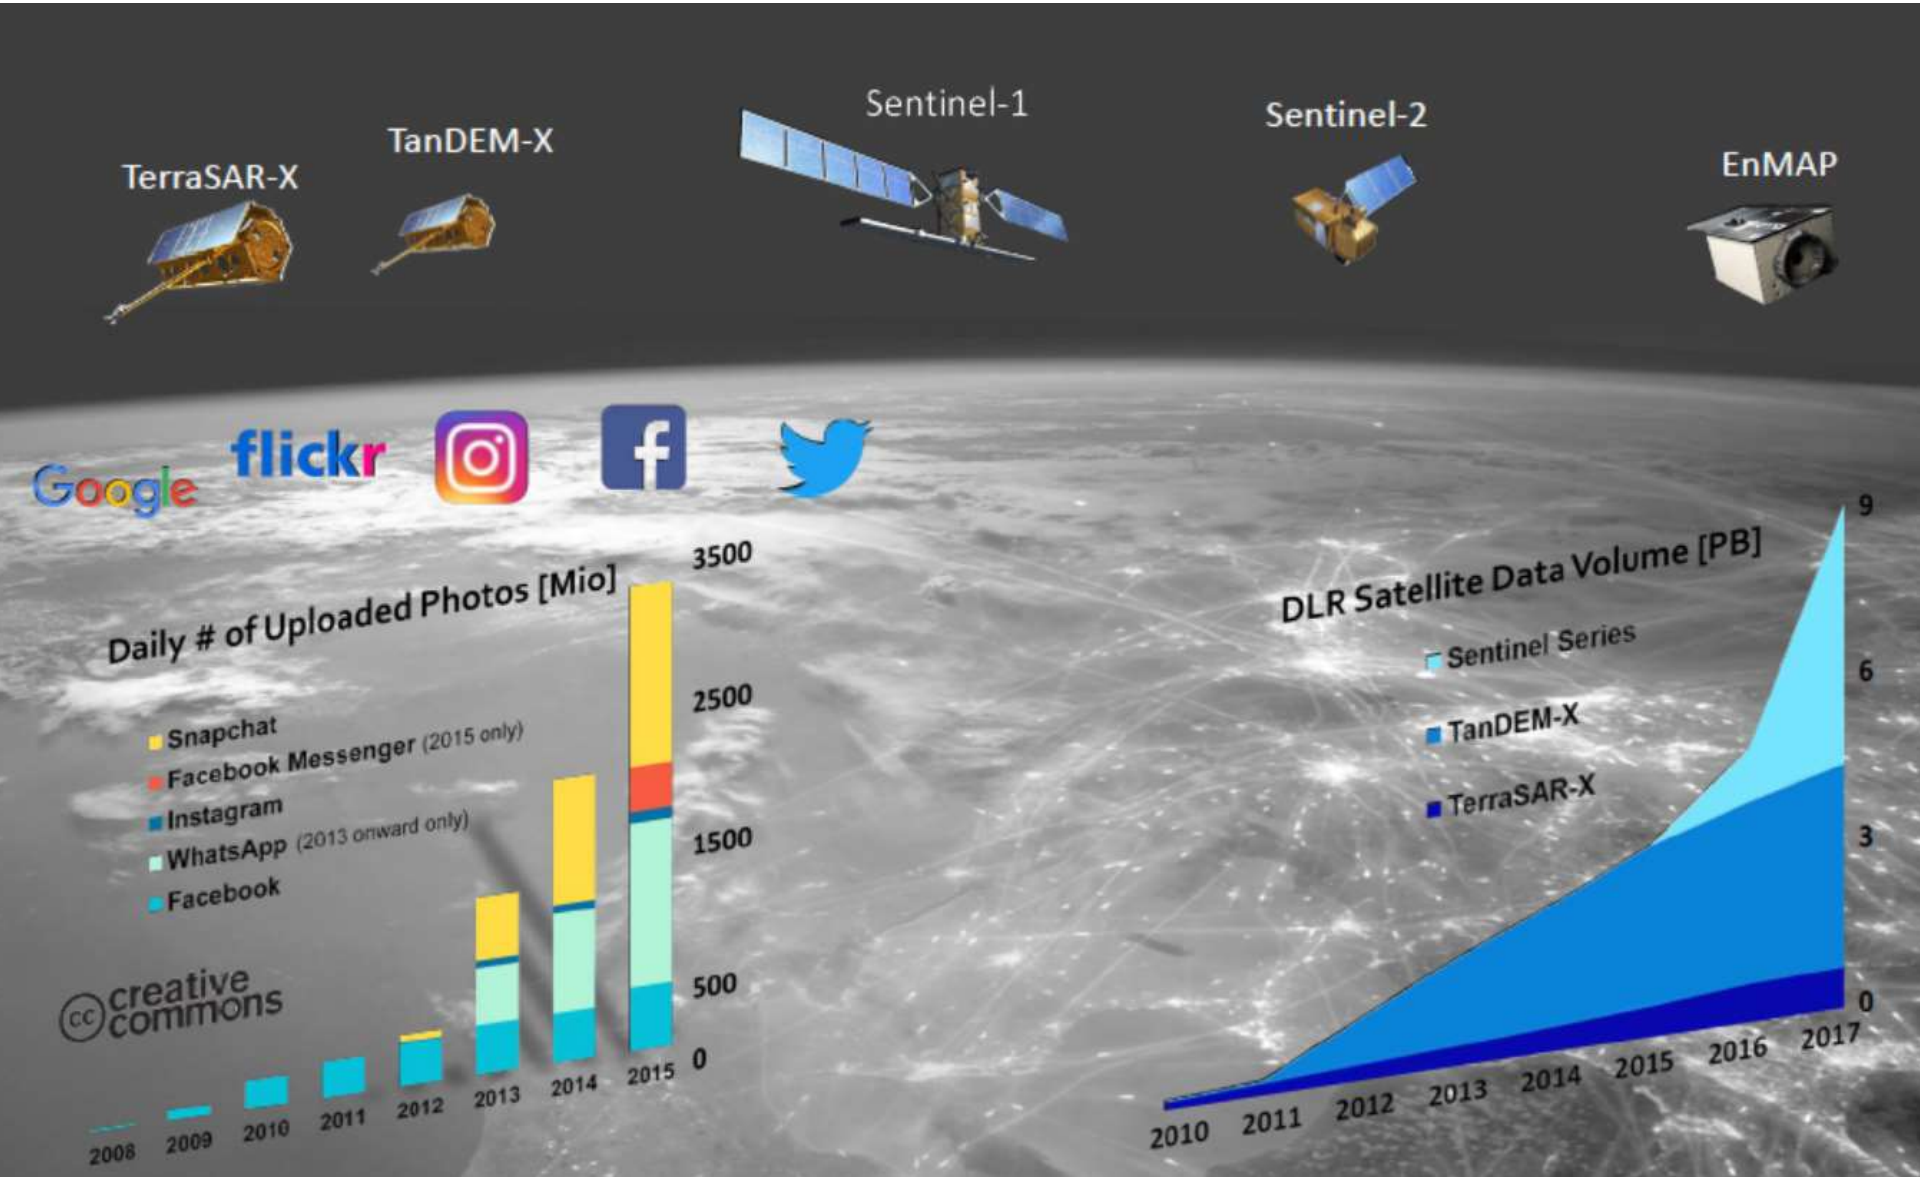

Deep learning

- the answer to automation needs ?

Deep learning for image exploitation

… "learning" an input-output mapping from examples by machines (**supervised classification**) – not so new …

- traditional classification:

most of the "heavy lifting" here, final performance only as good as feature set

- deep learning:

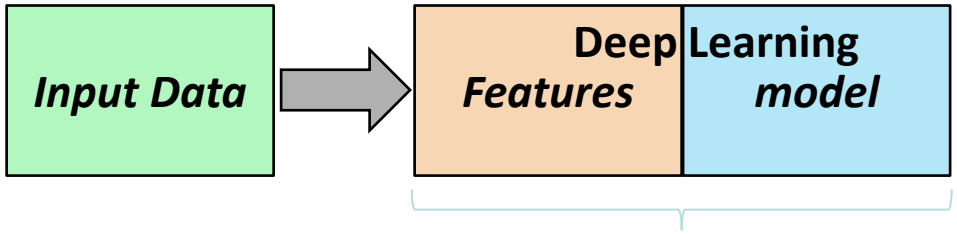

features and model learned together, mutually reinforcing each other

adapted from Michele Catasta

AI / ML / DL

- AI / DL ("data science") recently caught lots of **attention** in science and society alike – big data
  - „2nd best for everything"
  - hardware developments (in part. GPU)
  - abundant (training) data

- one of 10 **breakthrough techno-logies** in MIT Tech. Review 2013
- recently pushed by **Internet companies**
  - Google, Facebook, Microsoft, Baidu, …
- said to rival (even resemble?) **human cognitive ability**
  - victory of AlphaGo software against Go champion Lee Sedol (Oct.'17)
  - „super human" performance in computer vision challenges

Classification (CV terminology), examples

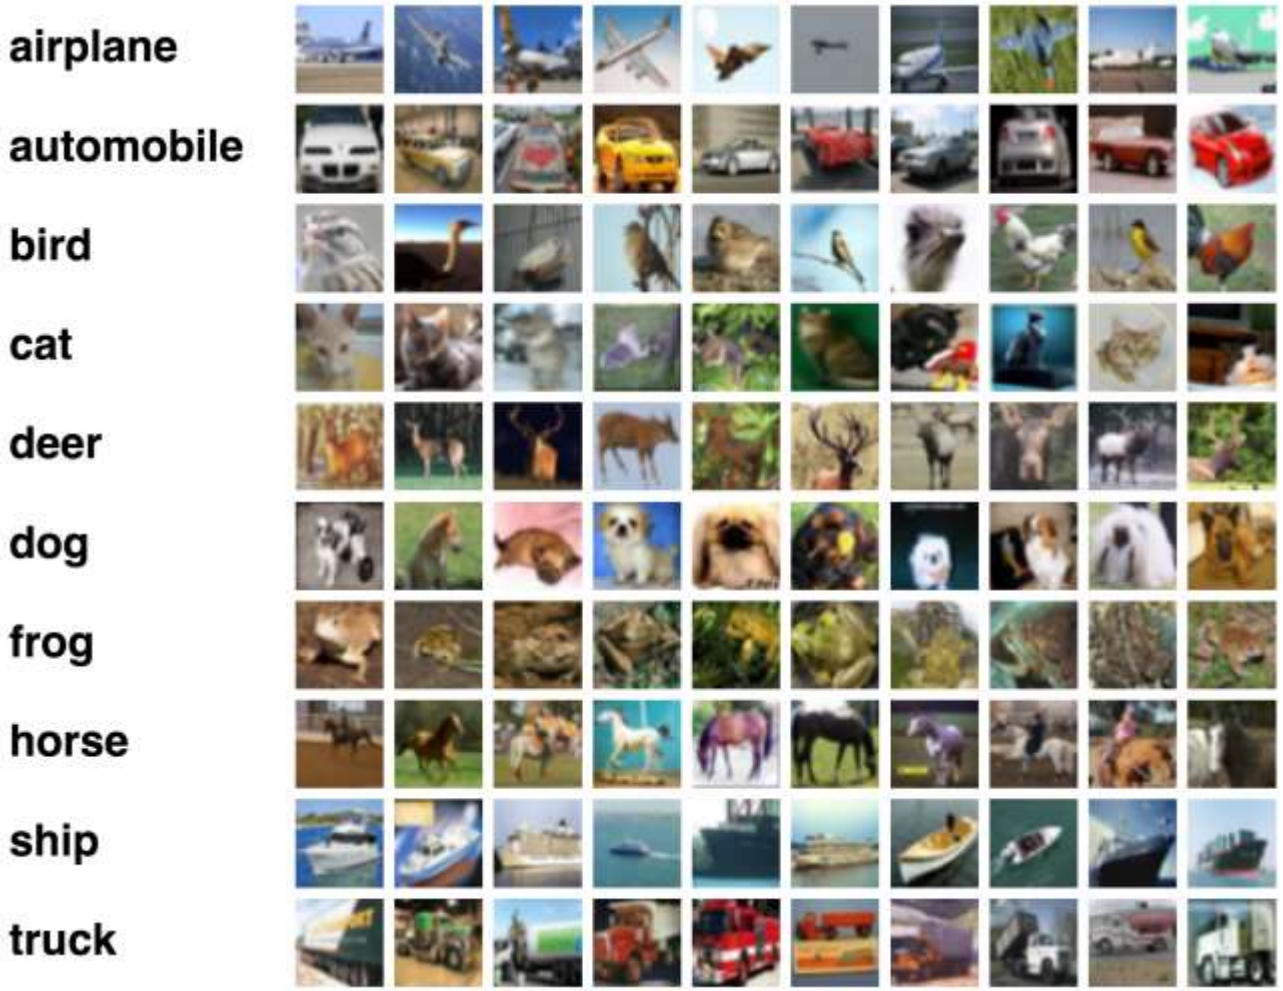

Success story of DL/CNN in computer vision

ImageNet Large Scale Visual Recognition Comp. (INLSVRC):

- **1,2 Mio** test images, **1000** classes
- **one class label** per image
- ever **increasing no. of layers**

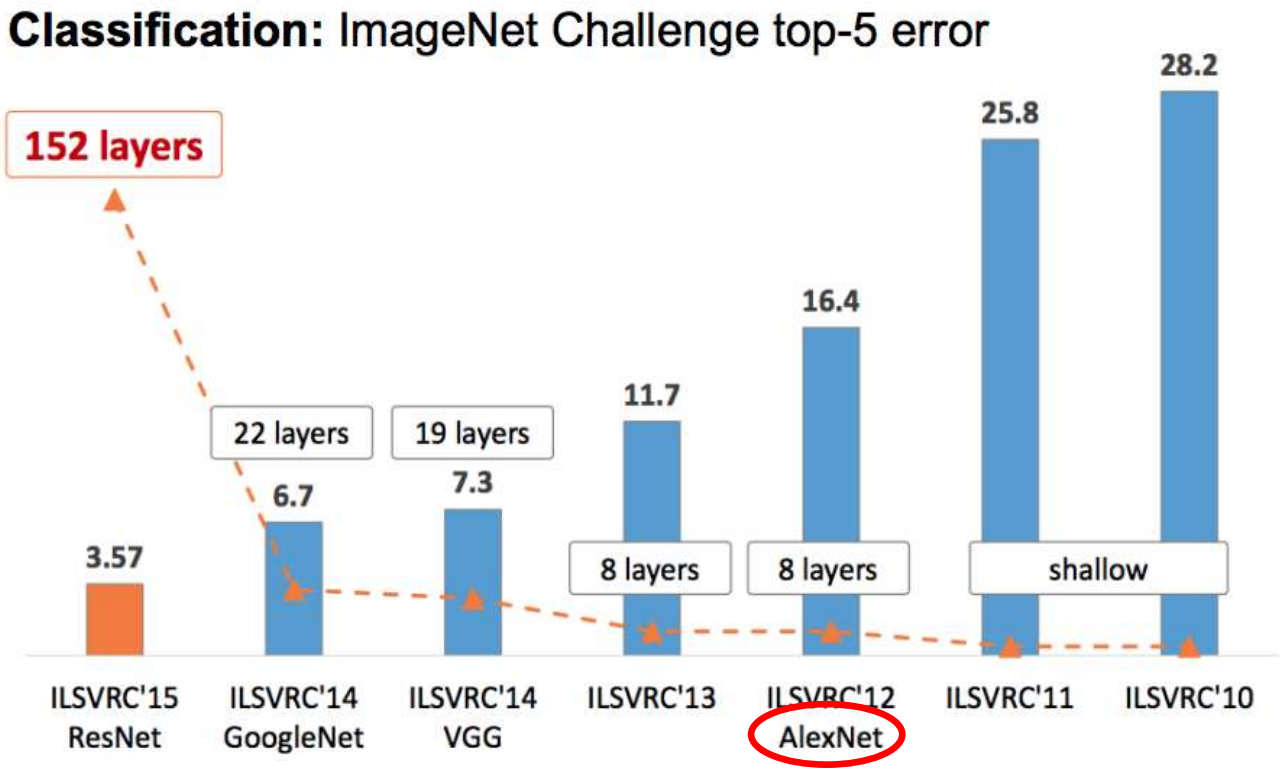

Typical remote sensing data

… are CNNs any good for us also ???

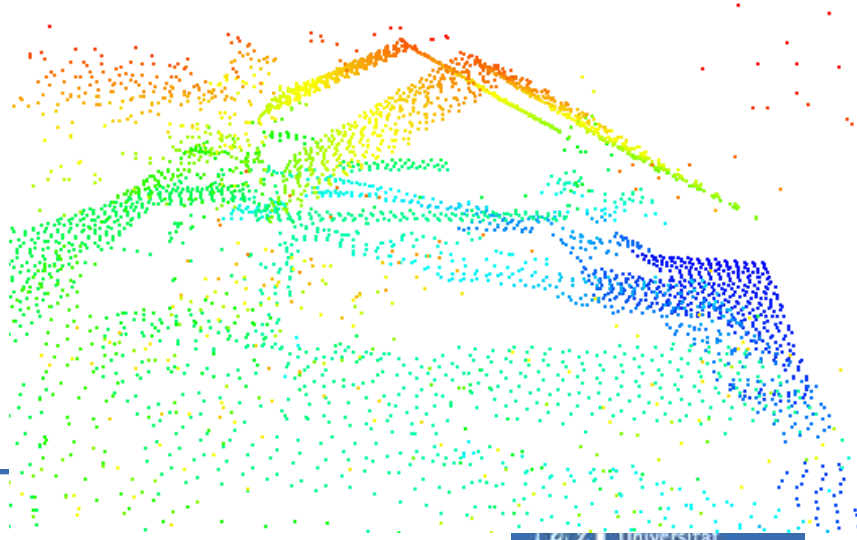

Example I: Road extraction / refinement

- first CNN application using aerial images
- binary classification (road / no road)
- training data: OSM
  - incomplete
  - partly with geometric errors
- models for errors in training data, learning of the related model parameters

Mnih & Hinton, 2012

Mnih V., Hinton G, 2012: Learning to label aerial images from noisy data. In: McCallum, Roweis (eds.), ICML 2012

Example II: land cover / land use update

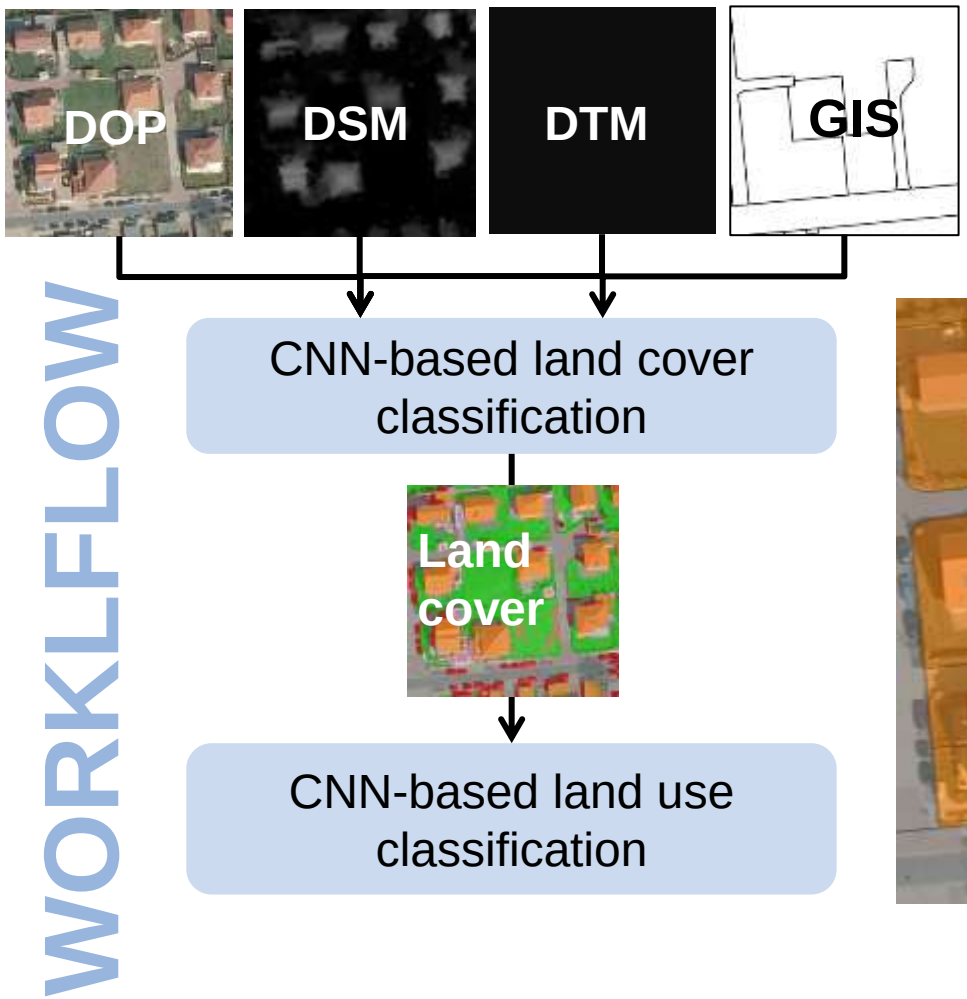

OA (LU) = 89.3 %

OA (LC) = 82.9 %

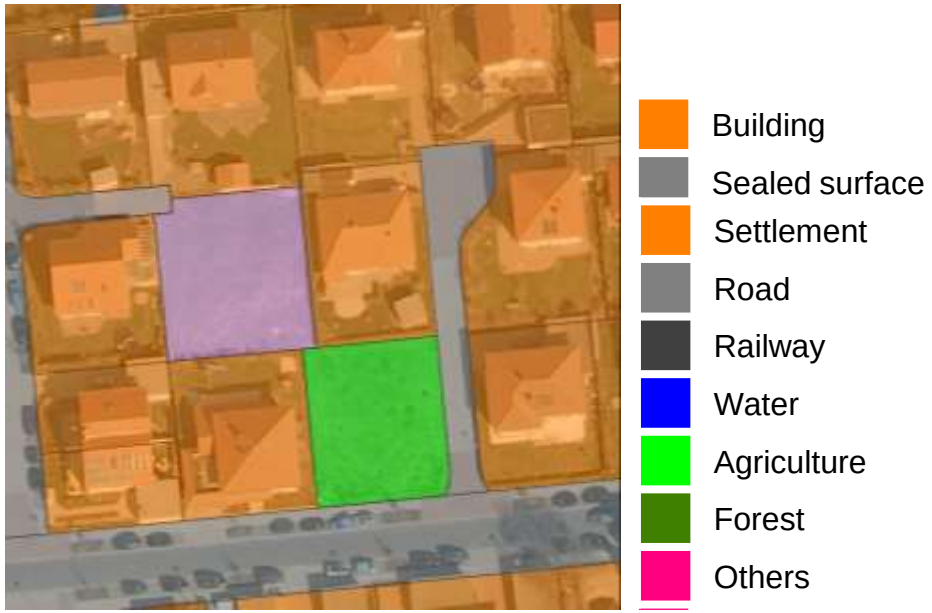

© Yang et al., 2018 - IPI

Yang C., Rottensteiner F., Heipke C., 2018: Classification of land cover and land use based on convolutional neural networks. ISPRS Annals of the Ph, RS IV-3, 251-258.

www2.isprs.org/commissions/comm3/wg4/tests.html

ISPRS Semantic Labelling Challenge

isprs
information from imagery

2018-08-03

Commissions

- Commission I
- Commission II
- Activities
- Working Groups
- WG II/1
- WG II/2
- WG II/3
- WG II/4
- Activities
- Benchmark Test
- Benchmark Test: Results
- Vaihingen Roads
- Vaihingen 2D Semantic Labeling

ISPRS Semantic Labeling Contest (2D): Results

Click here for a description of evaluation measures

Vaihingen: 2D Labelling challenge

All quality measures except for *overall* are F1 scores in [%] using the reference with eroded boundaries. Mouse over the column headings will display more information.

Abbrev.	imp surf	building	low_veg	tree	car	Overall	Strategy	Deta
SVL_1	86.3	90.8	78.2	84.2	56.8	84.7	s	D
SVL_2	82.1	82.8	71.6	81.6	51.9	79.4	s	D
SVL_3	86.6	91.0	77.0	85.0	55.6	84.8	s	D
SVL_4	86.1	90.9	77.6	84.9	59.9	84.7	s	D
SVL_5	86.1	90.3	75.6	84.7	45.9	84.0	s	D
SVL_6	86.0	90.2	75.6	82.1	45.4	83.2	s	D
ADL_1	88.1	92.0	79.0	86.5	59.0	86.1	s	D
ADL_2	89.0	93.0	81.0	87.8	59.5	87.3	s	D
ADL_3	89.5	93.2	82.3	88.2	63.3	88.0	s	D
UT_Mev	84.3	88.7	74.5	82.0	9.9	81.8	u	D

Example III: Dense matching using CNNs

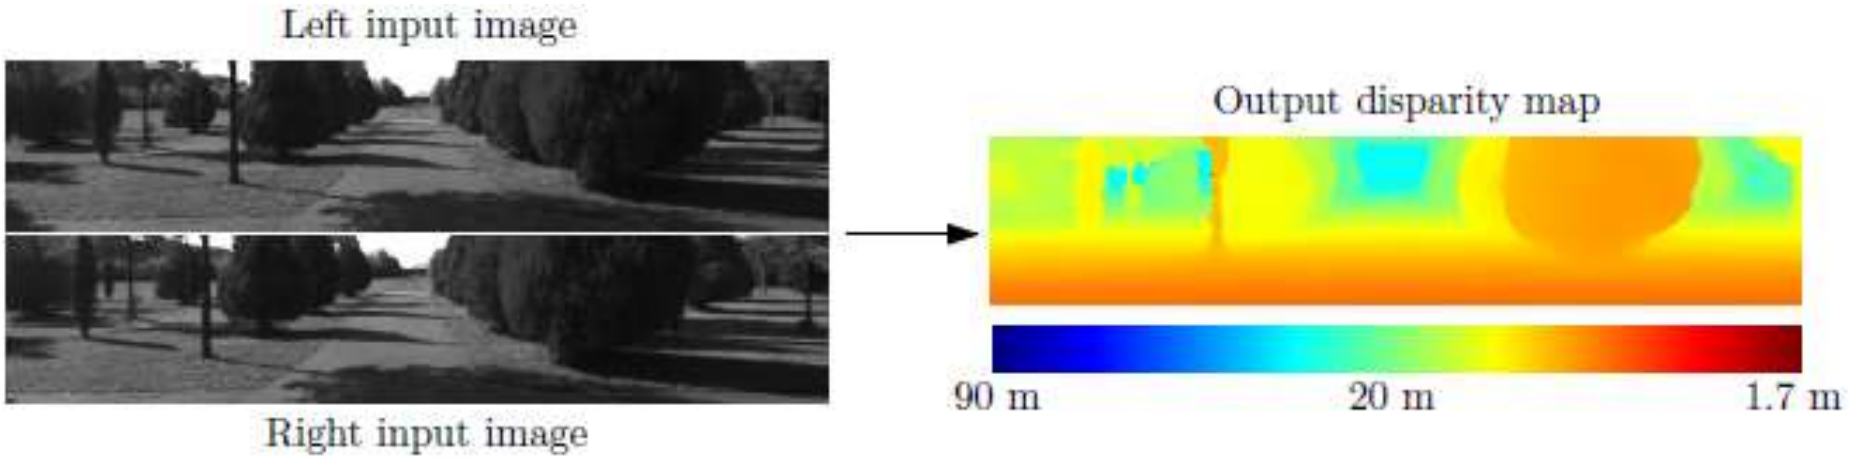

Zbontar J, LeCun Y, 2015: Computing the Stereo Matching Cost with a Convolutional Neural Network, CVPR, pp. 1592-1599, DOI: 10.1109/CVPR.2015.7298767

left image | gt/right image | baseline | Our

Kang J., Chen L., Deng F., Heipke C., 2019: Encoder-Decoder network for local structure preserving stereo matching Illumination. Proc. DGPF, Vol. 28

Conclusions- evaluation

- many other examples in the last few years (see e.g. review Zhu et al., 197 reference entries)
  - **more flexible** wrt. different object characteristics
  - strength: **learning features**, classifier not so important
  - key to good performance: **network depth, no. of training data**
  - has **outperformed just about all classical algorithms** by a large margin, provided enough training data is available
  - is about to do the same in geometric tasks

- no expert-free totally automatic processing chain, but **integration with human** capabilities

Zhu X., Tuia D., Mou L., Xia G., Zhang L., Xu F., Fraundorfer F., 2017. Deep Learning in Remote Sensing: A comprehensive review and list of resources. IEEE GRSS Magazine 5(4): 8-36.

Conclusions - evaluation

- Open-source CNN implementations available
  - Tensorflow (Google): https://www.tensorflow.org
  - CAFFE2 (Facebook) https://caffe2.ai/
  - …

- CNN has arrived in the commercial world
  - Descartes Lab, https://www.descarteslabs.com/
  - Simularity, simularity.com/solutions/satellite-anomaly-detection/
  - Orbital Insight, https://orbitalinsight.com/
  - spire, https://spire.com/
  - NVIDIA (computer games!), https://www.nvidia.com/
  - …

Conclusions - evaluation

- CNN is not magic - fundamentally a classifier
  - based on **correlation** between data sets
  - an enormous **amount of unknowns** to be estimated
    - lots of training data, long training times (days … weeks)
  - sensitive to
    - **overfitting** (“curse of dimensionality”)
    - **non-rep. training data** (incorrect / biased / unbalanced / …)
    - **user choices**, incl. hyper-parameters
      - network architecture, (non-lin.) activation function
      - design of loss function (function to be optimised)
      - training parameters (weight initialisation, learning rate, drop out rate, “momentum”, batch size, regularisation, …)
  - **generalisation and prediction capabilities unclear**
    - the system can’t learn what it never saw

Volvo admits its self-driving cars are confused by kangaroos

Swedish company's animal detection system can identify and avoid deer, elk and caribou, but is yet to work against the marsupials' movements

▲ Kangaroos are responsible for about 90% of collisions between vehicles and animals in Australia – although most are not serious. Photograph: Paul Kane/Getty Images

The Guardian, Int. Edition, July 1st, 2017

Conclusions – to do's

- integrate **prior knowledge**, e.g. physical models
  - don't try to learn what we already know (e.g. laws of nature)
- CNNs for **point clouds** and other irregularly rep. data
- tackle **training data shortage**
  - weakly, semi-, unsupervised learning, reinforcement learning
  - data augmentation and simulation
  - deal with errors in training data, e.g. when using outdated DB
  - combination with crowd sourcing and social media data
  - investigate limits of pre-training + fine tuning
- pay attention to **geometric accuracy** of object delineation
- sequential data, **time series** processing - RNNs
  - streaming (big) data, online CNN learning

Special thanks to

Prof. Franz Rottensteiner

Chair, IPI research group

Photogrammetric image analysis

Chen Lin, Max Coenen, Philo Humberg, Uyen Nguyen, Dennis Wittich, Chun Yang, Junhua Kang

ISPRS Geospatial Week 2019

Enschede (NL), June 10 – 14, 2019

Deadline for full papers: Jan. 17, 2019

ISPRS Congress 2020
Nice, France

Imaging today, foreseeing tomorrow

www.isprs2020-nice.com

June 14 – 20, 2020