4D Point Clouds and Their Machine-Learning-Based Interpretation

Intergeo 2018 Frankfurt, Germany — October, 16 2018

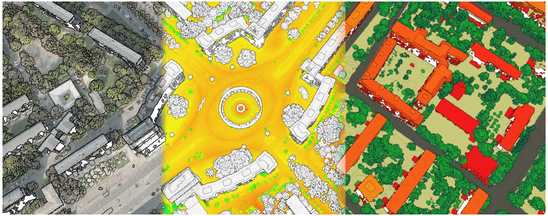

Prof. Dr. Jürgen Döllner

Hasso Plattner Institute — Faculty of Digital Engineering — University of Potsdam, Germany

About This Talk

- Machine Learning (ML)
- Point Clouds — Big Spatial Data
- ML-Based Interpretations of Point Clouds
- ML-Readiness of Point Clouds
- Geodata + AI = GeoAI?
- Conclusions

Machine Learning — Fundamentals

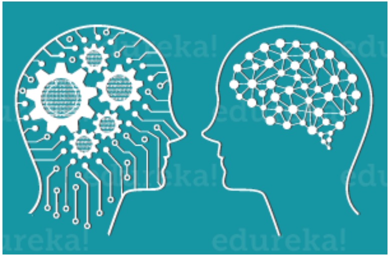

Fundamentals About Machine Learning

Artificial Intelligence, Machine Learning (ML), and Deep Learning (DL)

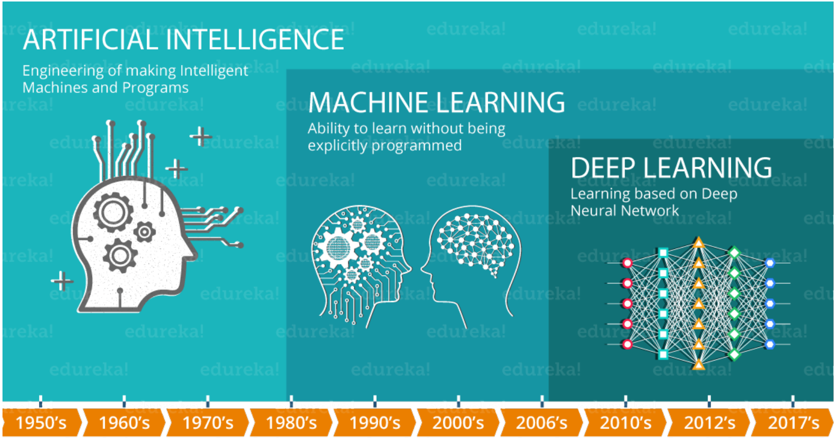

Image: `www.quora.com`

Fundamentals About Machine Learning

Given input x, predict output y

Input x	**Output y**
BUBQJ	B
ZMFXI	Z
NEJOL	N
RVOII	I

What might be the way machine learning thinks about this input-output pairs?

- Hypothesis space: The space of possible mappings from inputs to outputs
- Examples: (Convolutional) Neural Networks, Decision Trees, Nearest-Neighbor Models

Fundamentals About Machine Learning

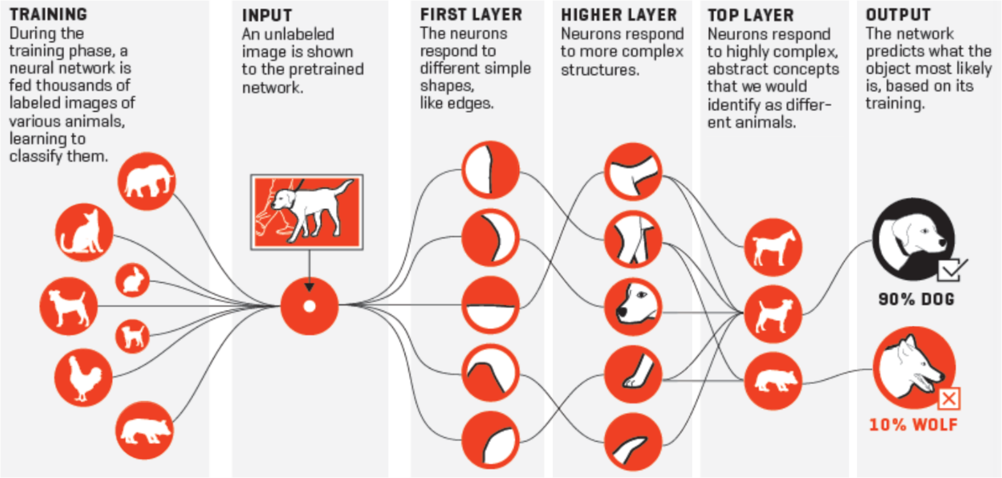

Example: How neural networks recognize a dog in a photo.

http://fortune.com/ai-artificial-intelligence-deep-machine-learning/

Fundamentals About Machine Learning

“The promise and power of machine learning rest on its ability to generalize from examples and to handle noise.”

M. Allamanis et al., “A Survey of Machine Learning for Big Code and Naturalness”, 2017

- The model is trained by training data.
- If you overtrain the model on the training data, then it will be able to identify all the relevant information in the training data, but will fail miserably when presented with the new data.
- The model becomes incapable of generalizing, i.e., it is *overfitting* the training data.

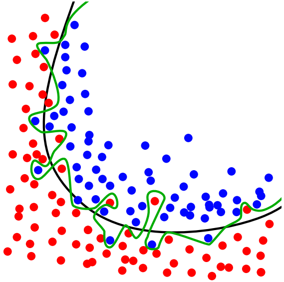

Example of overfitting and loose fitting model:
A loose fitting model leads to mismatching.

Image: wikipedia

Fundamentals About Machine Learning

Big Data and ML

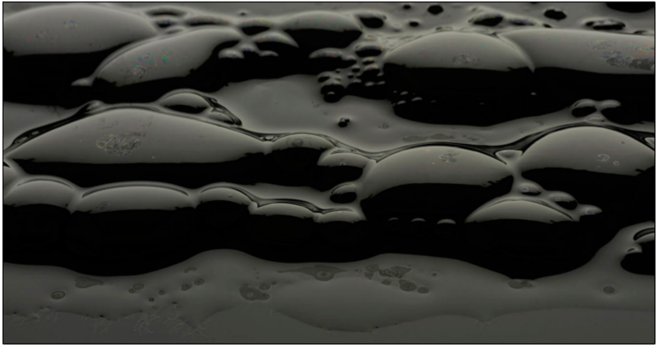

“Data is the new oil. Data is just like crude. It’s valuable, but if unrefined it cannot really be used.” Clive Humby, dunnhumby

Fundamentals About Machine Learning

Big Data and Analytics

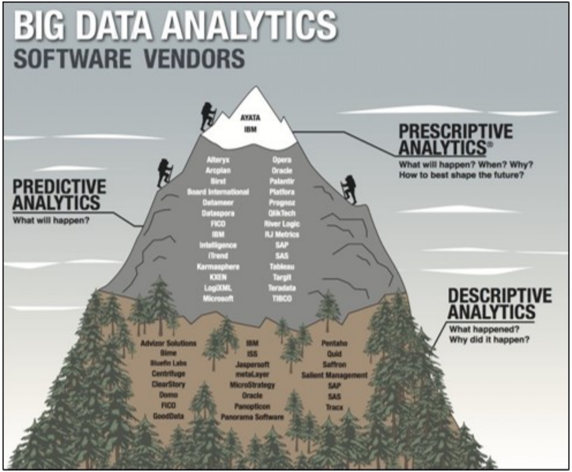

Image: IBM

Point Clouds — Big Spatial Data

3D Point Clouds

Point Cloud Characteristics

- Represent geometry by **discrete, unstructured sets of points** (e.g., in an Euclidian space)
- Approximate any shape, any topology and any geometry
- Points can be attributed by additional measures, e.g., colors

- No assumptions regarding density, regularity, or statistic properties
- Simple and powerful approach to geometric modeling able to approximately represent any physical entity, object, or phenomenon in a unified way

- **Big spatial data** due to compactness and size

3D Point Clouds

Point Cloud Generation

- Remote sensing
- Photogrammetry
- Sensor-based
- Discretization
- Procedural generation
- ...

3D Point Clouds

Example: Point cloud-based as-is building model

Image: HPI — Demo at our Intergeo booth 12.1G.079

3D Point Clouds

Real-Time Immersive Indoor Point Cloud Visualization

Image: HPI, Data: DLR — Demo at our Intergeo booth 12.1G.079

3D Point Clouds

High-Quality Outdoor Point Cloud for City Representations

Image: HPI, Data: Stadt Landshut — Demo at our Intergeo booth 12.1G.079

4D Point Clouds

Spatio-Temporal Point Clouds

- Frequent generation of point clouds (e.g., day-to-day, on-demand, real-time, ...)
- Dense time-variant point clouds, in general with high data redundancy
- Collections of point clouds that have a common geospatial extension are called **time-variant point clouds** or **4D point clouds**

- Availability, accuracy, density, and compactness of 4D point clouds is vastly increasing
- 3D point clouds are big, 4D point clouds are much bigger.

Dataset	Area	Density	# Points
Berlin	890 km^2	100 pts/m^2	80 billion
Baden-Württemberg	36,000 km^2	25 pts/m^2	900 billion
Frankfurt	250 km^2	15-20 pts/m^2	7.1 billion
Netherlands (AHN2)	42,000 km^2	6-10 pts/m^2	>600 billion

4D Point Clouds

Information Cartography

Access point clouds and analytics results

- Explore and present terabytes of point clouds
- Photorealistic and semantic-based illustration
- Real-time and device-independent access

Data Sources

Plug and play – all data sources

- Point clouds with any resolution, density, and size
- Aerial, mobile mapping, and terrestrial point clouds
- LIDAR, photogrammetry, and stereo data
- Manage unlimited amounts of points

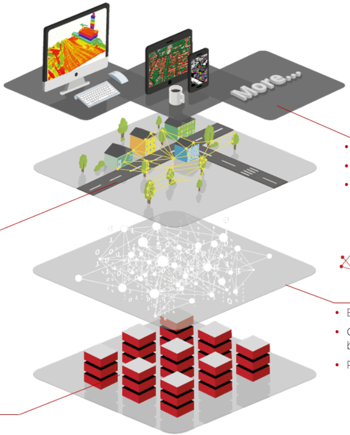

Business Applications

Seamless Integration into Workflows and Platforms

- Documentation, monitoring, and reporting
- Up-to-date models for spatial areas, sites, infrastructures, ...
- Seamless integration into workflows, systems & applications

Point Cloud Analytics

Derive and extract information and added values

- Efficient storage and high-performance processing
- Clustering, classification, feature derivation by Machine Learning & Deep Learning
- Predictive spatial analytics

Image: Point Cloud Technology Stack, `www.pointcloudtechnology.com`

Point Clouds

Change Detection for 4D Point Clouds

- Compare two point clouds
- Identify *differences* to handle redundancy
- Computationally GPU-efficient parallel processing
- Specialized spatial data structures adapted to point clouds
- Caching along the memory hierarchy, i.e., hard disk, main memory, GPU memory, GPU registers

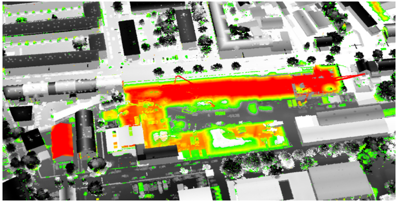

Image: HPI & Point Cloud Technology, **www.pointcloudtechnology.com**

Point Clouds

Change Detection for 4D Point Clouds – Example

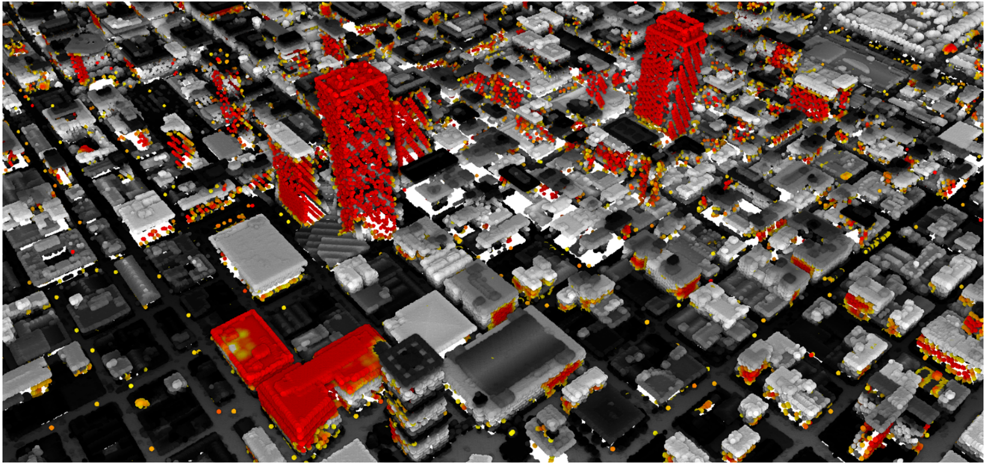

Image: HPI — Demo at our Intergeo booth 12.1G.079

4D Point Cloud Web-Based Platform – ”punctumTube”

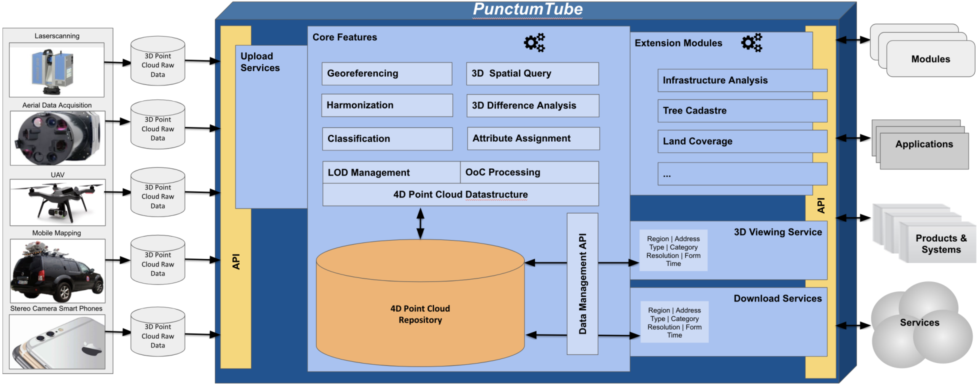

Image: HPI & Point Cloud Technology, `www.pointcloudtechnology.com` — Demo at our Intergeo booth 12.1G.079

ML-Based Interpretations of Point Clouds

ML-Based Interpretation of Point Clouds

”Interpretation” of Point Clouds

- Interpretation distills semantics and derives models, features, or structures out of point clouds.
- Interpretation is required to use point clouds in specific application fields and domains.
- Large amounts of previous work has been created over the last 4-5 decades by explicitly programmed, heuristic-based interpretation algorithms.
- Interpretation techniques can be (re)designed and (re)implemented by means of ML.

ML-Based Interpretation of Point Clouds

Requirements for Approaches of Point Cloud Interpretation

- **Executable on-demand** for a specific target region
- **Configurable** regarding the features as needed by a specific request
- **Real-time computation** to support on-the-fly answers
- **Operating on raw data** without intermediate representations
- **Service-based architecture** to allow for building higher-level services and mashups

Optimal Characteristics

- No intermediate storage required, only one single *4D point-cloud cloud*.
- No duplication of point cloud data by applications and systems.
- Domain-independent, generic 4D point-cloud cloud and ML-based interpretations.
- Modeling and semantics is completely mapped into and covered by interpretation algorithms.
- No pre-selected, pre-defined modeling schemata, no loss of precision regarding raw data.

Examples for Point Clouds Interpretation

Example: Deriving a ground-corresponding area 3D model based on (dense) point clouds

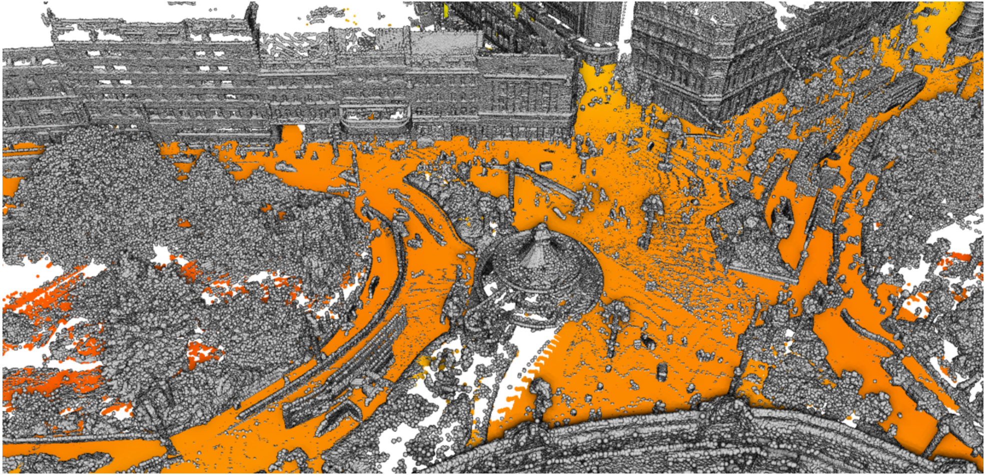

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: Automated classification of airborne point clouds

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: Automated classification of airborne point clouds

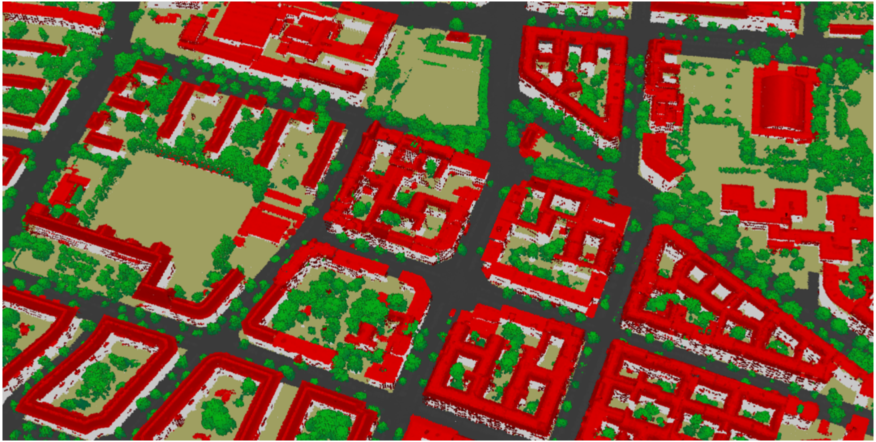

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: ML-based analysis of mobile mapping scans

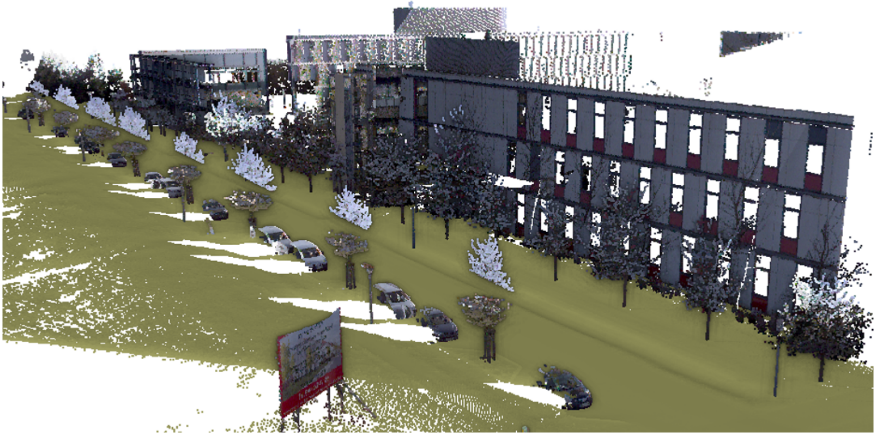

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: ML-based analysis of mobile mapping scans

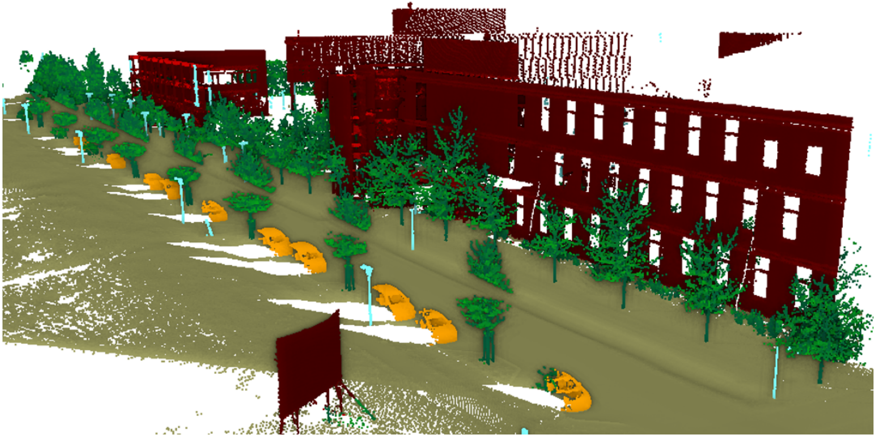

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: ML-Based analysis of railroad scans

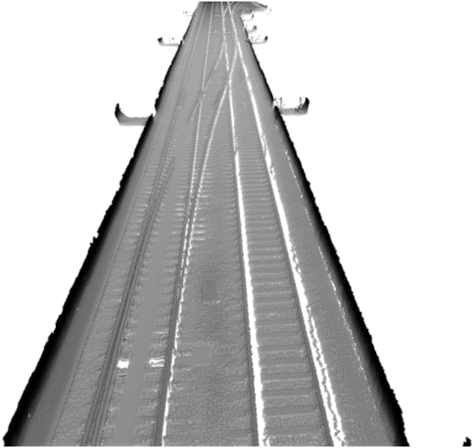

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: ML-Based analysis of railroad scans

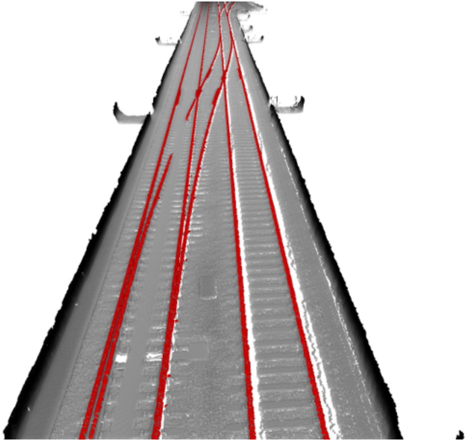

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: ML-Based analysis of railroad scans

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

Example: ML-Based analysis of railroad scans

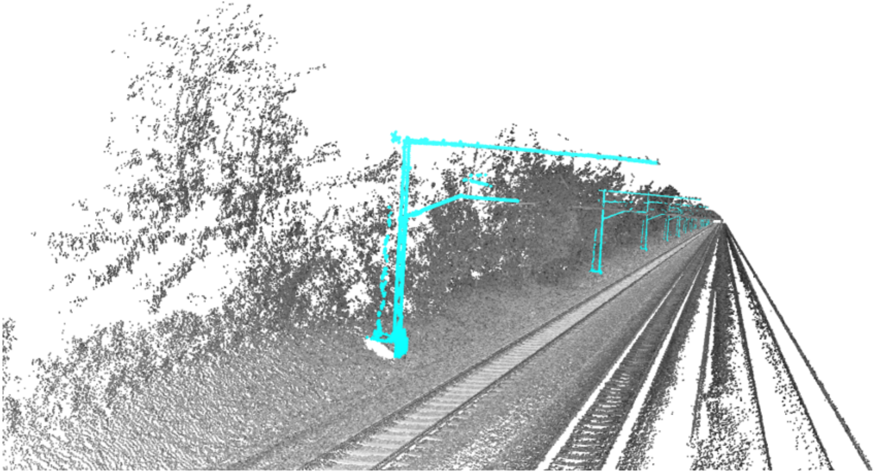

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

ML-Based Asset Detection using Ground Penetration Radar

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

ML-Based Asset Detection using Ground Penetration Radar

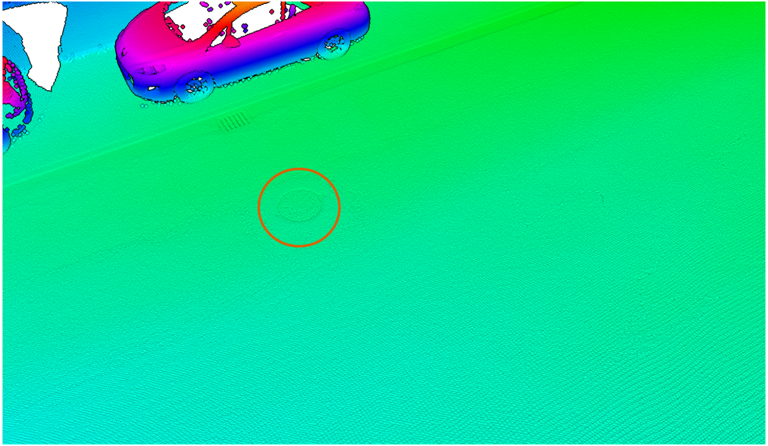

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

ML-Based Asset Detection using Ground Penetration Radar

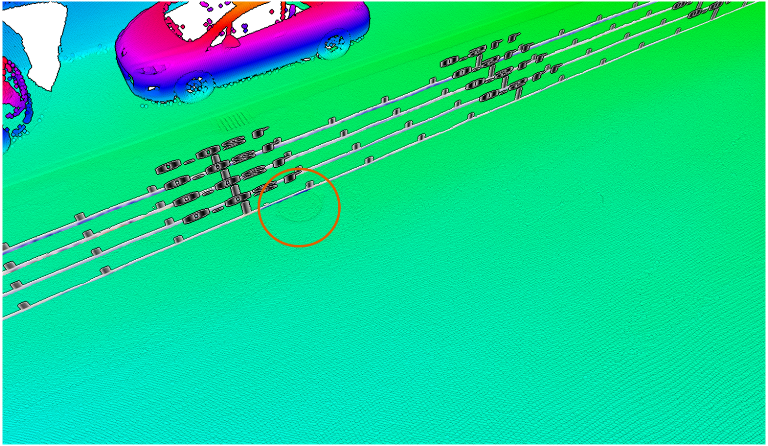

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

ML-Based Asset Detection using Ground Penetration Radar

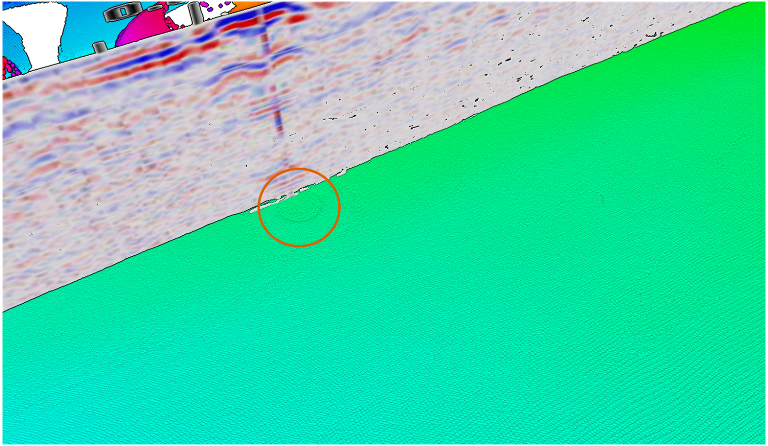

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

ML-Based Vegetation Detection

- Classify, extract, evaluate, and visualize information about vegetation—based on only point-cloud data.
- Extract high-level data, e.g., about isolated large trees, strongly inclinated trees close to buildings or streets, annual growth, etc.

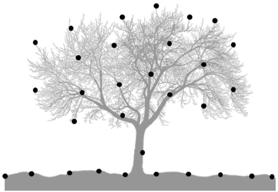
LiDAR Point Cloud

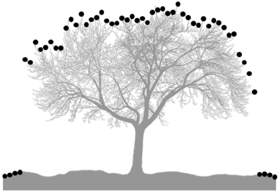
2.5D Point Cloud


RGB Image


CIR Image

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Examples for Point Clouds Interpretation

ML-Based Vegetation Detection – Case Study (Berlin, Fully Automated Tree Map)

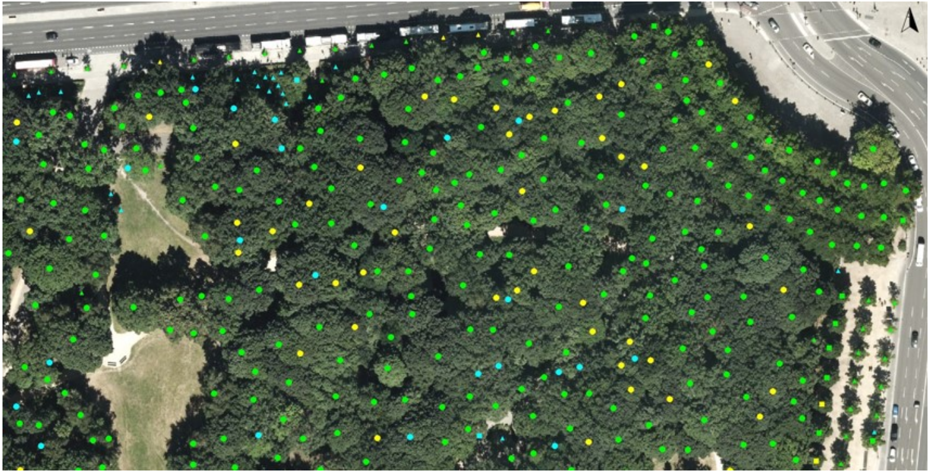

Image: HPI & PCT — Demo at our Intergeo booth 12.1G.079

Point Clouds

The Vanishing Need for Derived Models - The Power of Real-Time Raw Data Analysis

- Deriving secondary models (e.g., virtual 3D city models) can be time-consuming, lead to a loss of precision, large parts may never be used, …

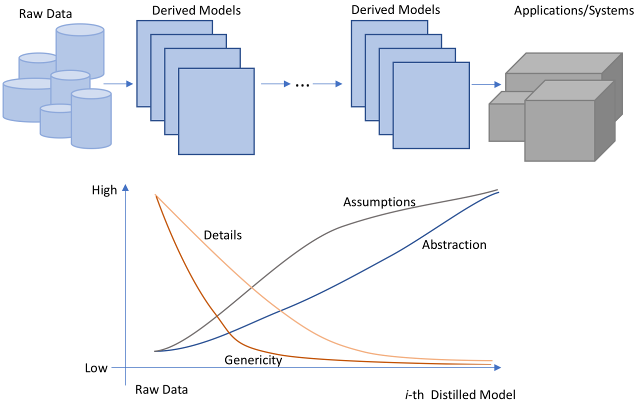

ML-Readiness of Point Clouds

ML-Readiness of Point Clouds

ML Criteria

- **Big data**: 4D point clouds—due to feasible and cost-efficient acquisition approaches—generate ”big spatial data”.
- **Inherent degree of fuzziness**: Point clouds of captured objects at different points in time or by different capturing techniques, show inherent fuzziness.
- **Defined semantics**: Semantics can be assigned by interpreting point clouds based on training data.
- **Efficiency and Effectiveness**: Compared to traditionally, procedural-like, heuristically or empiric-based algorithms, ML-based approaches show significantly less implementation complexity and much higher stability and robustness.

ML-Readiness of Point Clouds

Naturalness Hypothesis

- One key approach to ML is to find out whether a given (different) problem domain corresponds to or has similar statistically properties as **large natural language corpora**.
- ML-based approaches have shown extraordinary success in natural language recognition, natural language translation, question-answering, text mining, text comprehension, etc.—in particular, compared to all other approaches of the past.

Key insight:

Most human utterances are far simpler and much more repetitive and predictable.

ML-Readiness of Point Clouds

Naturalness Hypothesis—Example: Computer programs seen as text corpora

The Naturalness Hypothesis[1]

“ Software is a form of human communication; software corpora have similar statistical properties to natural language corpora … In fact, these utterances can be very usefully modeled using modern statistical methods. ”

“ Programming languages, in theory, are complex, flexible and powerful, but the programs that *real* people *actually* write are mostly simple and rather repetitive, and thus they have usefully predictable statistical properties that can be captured in *statistical language models* ”

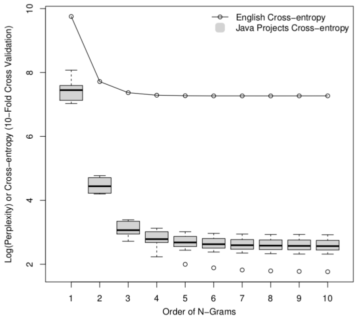

Figure 1. Comparison of English cross-entropy versus the code cross-entropy of 10 Java projects.

“The measure *perplexity* used in its log-transformed version is called *cross-entropy*—roughly speaking, it tells how surprised a model is for a document.”

[1] A. Hindle, “On the Naturalness of Software”, ACM, 2016

ML-Readiness of Point Clouds

Hypothesis: ”*Point clouds are similar to natural language text corpora*” (HPI)

A Naturalness Hypothesis for Point Clouds

“ *Point clouds* is a form of *natural* communication; *point cloud corpora* have similar statistical properties to natural language corpora; and these properties can be exploited to build better geospatial tools. In fact, these ”utterances” can be very usefully modeled using modern statistical methods. ”

A very likely future strong insight:

Most point clouds are far simpler and much more repetitive and predictable.

Geodata + AI = GeoAI?

Attempts for Definitions of Artificial Intelligence

- "*the ability of a digital computer or computer-controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience.*" (https://www.britannica.com/technology/artificial-intelligence)
- "*the capability of a machine to imitate intelligent human behavior.*" (https://www.merriam-webster.com/dictionary/artificial%20intelligence)

- All definitions lack a concise definition of "intelligence".
- Frequently, the terms *machine learning*, *deep learning*, and *AI* are confused with each other or are used interchangably (e.g., see Intergeo ads and press articles).

Artificial Intelligence (AI)

```
 Welcome to
                    EEEEEE  LL      IIII  ZZZZZZZ   AAAAA
                    EE      LL       II       ZZ   AA   AA
                    EEEEE   LL       II     ZZZ    AAAAAAA
                    EE      LL       II    ZZ      AA   AA
                    EEEEEE  LLLLLL  IIII  ZZZZZZZ  AA   AA

 Eliza is a mock Rogerian psychotherapist.
 The original program was described by Joseph Weizenbaum in 1966.
 This implementation by Norbert Landsteiner 2005.


ELIZA: Is something troubling you ?
YOU:   Men are all alike.
ELIZA: What is the connection, do you suppose ?
YOU:   They're always bugging us about something or other.
ELIZA: Can you think of a specific example ?
YOU:   Well, my boyfriend made me come here.
ELIZA: Is it important to you that your boyfriend made you come here ?
YOU:   He says I'm depressed much of the time.
ELIZA: I am sorry to hear that you are depressed.
YOU:   It's true. I am unhappy.
ELIZA: Can you explain what made you unhappy ?
YOU:
```

ELIZA—A Computer Program For the Study of Natural Language Communication Between Man And Machine

JOSEPH WEIZENBAUM
Massachusetts Institute of Technology, Cambridge, Mass.*

ELIZA is a program operating within the MAC time-sharing system at MIT which makes certain kinds of natural language conversation between man and computer possible. Input sentences are analyzed on the basis of decomposition rules which are triggered by key words appearing in the input text. Responses are generated by reassembly rules associated with selected decomposition rules. The fundamental technical problems with which ELIZA is concerned are: (1) the identification of key words, (2) the discovery of minimal context, (3) the choice of appropriate transformations, (4) generation of responses in the absence of key words, and (5) the provision of an editing capability for ELIZA "scripts". A discussion of some psychological issues relevant to the ELIZA approach as well as of future developments concludes the paper.

Eliza, created by Joseph Weizenbaum (1966)—a first example of how artificial intelligence could incarnate. For the first time, a programmer had attempted a human-machine interaction with the goal of creating the illusion (however brief) of human-human interaction.

Artificial Intelligence

- Artificial intelligence, by its very nature, involves *diverse aspects of intelligence*, e.g., as common to "intelligent beings".
- In computer science, sometimes the positive aspects are seen but the negative ones stay invisible.
- Think of AI aspects such as an algorithm that is greedy, lying, selfish, malicious, perfidious, ...

Hieronymus Bosch: The Seven Deadly Sins. Ca. 1505-1510; Gluttony, Sloth, Lust, Pride, Wrath, Envy, Greed

GeoAI

Artificial Intelligence in the Geospatial Domain

- Geospatial information is key to most processes and events—if big spatial data is available, most technologies apply perfectly.
- Geospatial information is mostly associated with semantics concepts and spatial ontologies—AI can take advantage here to an extraordinary degree.
- Geospatial applications are close to the user, are close to real life—this is a perfect place where AI comes into play.
- Machine learning algorithms for geospatial information open a large door towards ”intelligent” applications.

GeoAI

"DemonGo" Study at HPI (2018)

- Interactive, multi-user augmented-reality game with the objective to invoke and catch demons placed in the near-by environment.
- DemonGo apps collects information about the physical environment by video streams captured by mobile built-in cameras.

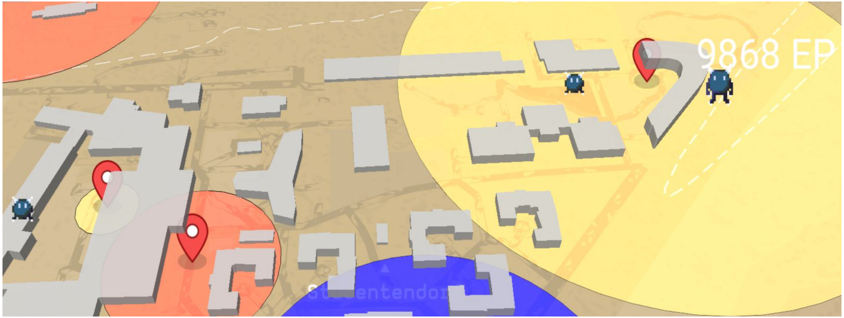

GeoAI

"DemonGo" Study at HPI (2018)

- Interactive, multi-user augmented-reality game with the objective to invoke and catch demons placed in the near-by environment.
- DemonGo apps collects information about the physical environment by video streams captured by mobile built-in cameras.

”DemonGo” Study at HPI (2018)

- 3D reconstruction (on-the-fly) for outdoor and indoor environments

"DemonGo" Study at HPI (2018)

- Various analysis techniques are applied on the captured video streams and reconstructed environments, e.g., detecting image areas with visual noise, with large number of edges, and texts.

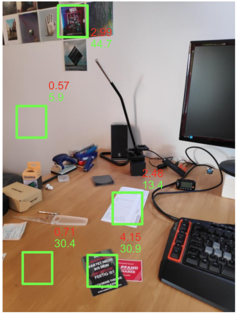

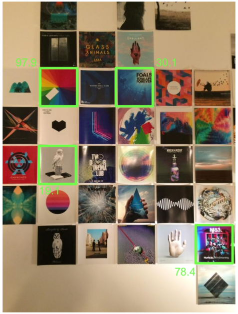

”DemonGo” Study at HPI (2018)

- ML-based outcome and statistical analysis lead to ranked points of interest within the currently viewed/seen spatial user environment.

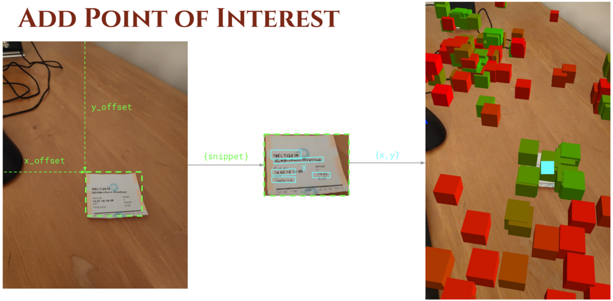

DemonGo! Experiment at HPI - and what about AI?

- The DemonGo app shows how to intrude and spy the environment, most likely without the user's awareness—with astonishing technical precision and robustness (we didn't expect this).
- To this end, points of interests are placed where most promising objects have been detected—the user will intend to show these places to catch the demons.
- Extremely interesting, for example, are all objects nearby texts that can be OCR-scanned in real-time.
- With a user photo, after bio-metric analysis, and the user's email address, a relatively good base is reached to systematically "investigate" the user's environment by placing demons exactly at places where most promising information can be found.

⇒ "Intelligence" can always involve bad capabilities—enabled by the combination of image vision, VR, and AR based on ML data processing.
⇒ This app forms part of a study and will not be released.

Conclusions

Conclusions

ML + Point Clouds?

- **Point clouds** serve to represent geospatial entities in a unique, consistent and simple way; in a sense, they are prototypes of **big spatial data**.
- Captured or generated frequently, we get **4D point clouds**, which require specialized, optimized data handling due to their inherent redundancy.
- ML-based point cloud approaches **handle perfectly these incomplete and fuzzy discrete point sets** and will *eat up* many GIS implementations and solutions (like they do in other domains).
- ML serves to **reconstruct semantics within point clouds** on-demand/on-the-fly/in real-time, that is, ML heals one of the biggest weaknesses of point clouds.
- Almost scaring but: yet low-density point clouds appear to allow us to **interpret robustly assets and objects** in point clouds.
- ML-based point-cloud approaches will drastically **simplify software engineering** of complex GI solutions.

Contact Data

Visit our Intergeo 2018 booth 12.1G.079 (joint booth with GeoKOMM e.V.)

Online Demos: www.pointcloudtechnology.com

Contact:

- Prof. Dr. Jürgen Döllner juergen.doellner@hpi.de
- Rico Richter rico.richter@hpi.de
- Sören Discher soeren.discher@hpi.de
- Johannes Wolf johannes.wolf@hpi.de