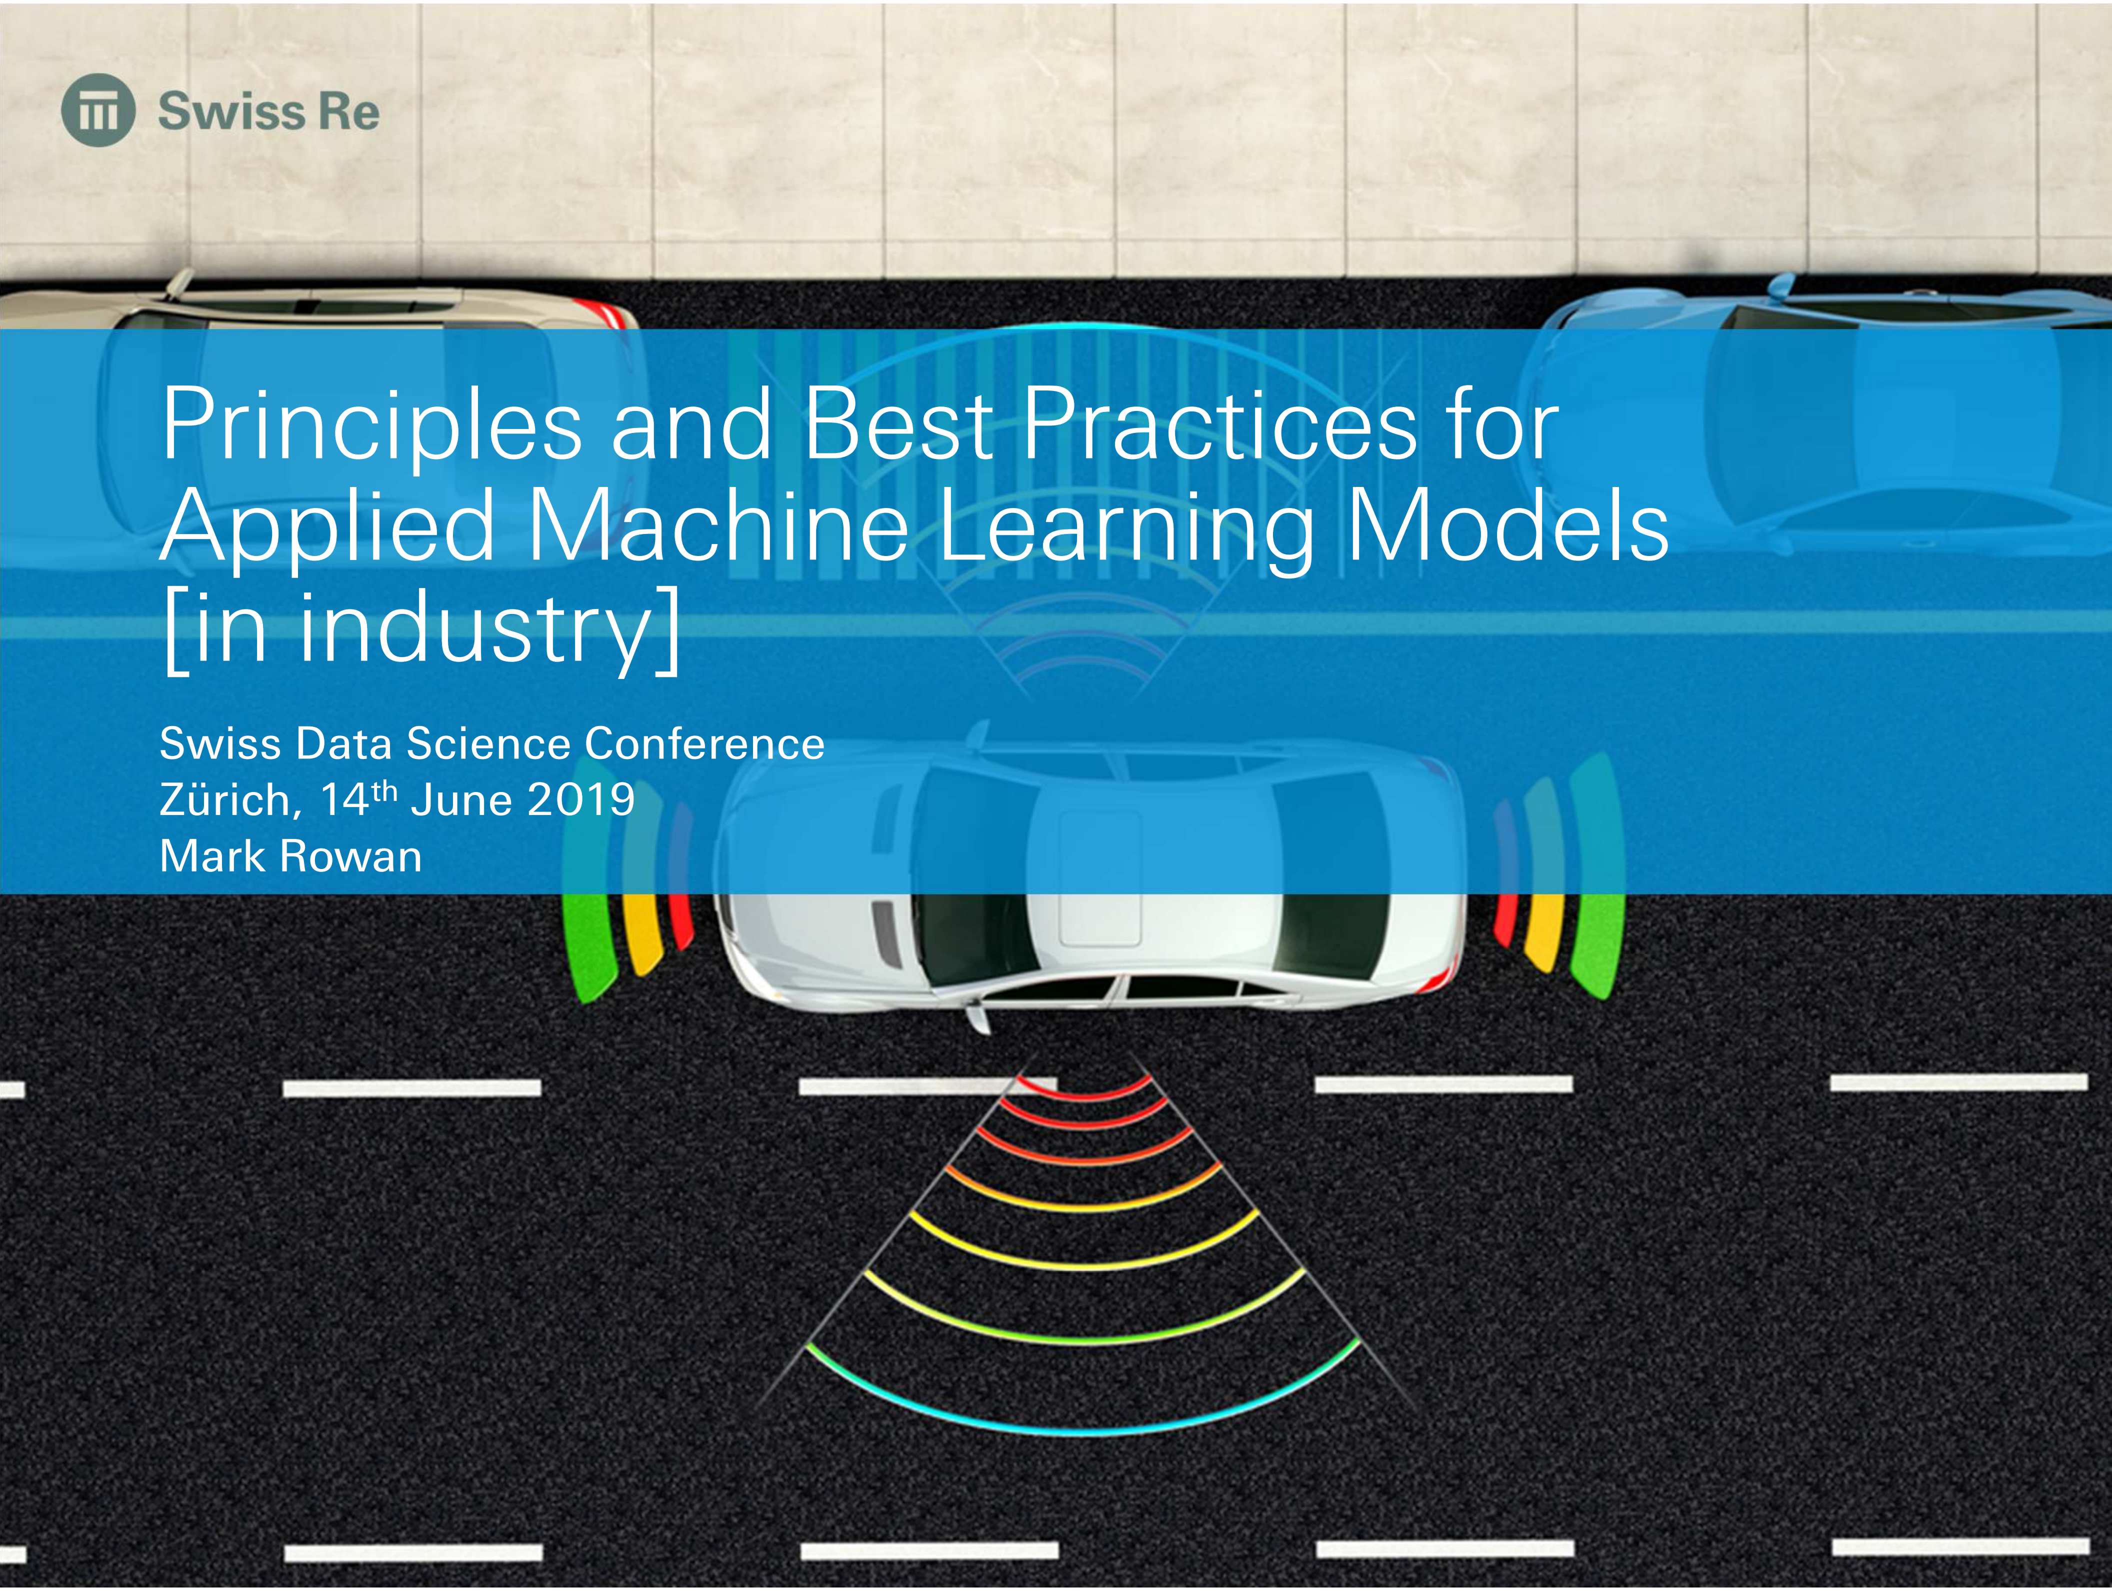

Swiss Re

Principles and Best Practices for Applied Machine Learning Models [in industry]

Swiss Data Science Conference
Zürich, 14th June 2019

Mark Rowan

- World-leading reinsurer
- 14,500 employees worldwide
- Founded in Zürich, 1863
- “We make the world more resilient”
- Many actuaries, scientists, engineers, risk modellers

Motivation (1/2)

Financial crisis 2007/08:

Some replies by researchers:

- (L.C.G. Rogers) The problem is not that mathematics was used by the banking industry, the problem was that it was abused by the banking industry. Quants were instructed to build models which fitted the market prices. Now if the market prices were way out of line, the calibrated models would just faithfully reproduce those wacky values, and the bad prices get reinforced by an overlay of scientific respectability!

r my students

nts will have to use the tools and techniques of QRM wisely in a world where the rules of the game will have been changed.

Always be scientifically critical, as well as socially honest, adhere to the highest ethical principles, especially in the face of temptation … which will come!

From: Prof. Paul Embrechts, «Did a Mathematical Formula Really Blow up Wall Street?», ASTIN Colloquium, 2009

Models were not used properly!

Today, we should be aware that...

Machine learning models can be very powerful, but there are also limitations and fallacies....

...so we need to be very careful how and where they are used.

...we do not want to read «The Neural Network that ....».

Motivation (2/2)

Boeing 737-Max accidents:

April, 4, 2019

We at Boeing are sorry for the lives lo… our hearts and minds, and we extend our sympathies to the loved ones of the p… 2. All of us feel the immense gravity of these events across our company and recognize the devastation of the families and friends of the loved ones who perished.

The full details of what happened in the two accidents will be issued by the government authorities in the final reports, but, with the release of the preliminary report of the Ethiopian Airlines Flight 302 accident investigation, it's apparent that in both flights the Maneuvering Characteristics Augmentation System, known as MCAS, activated in response to erroneous angle of attack information.

From: www.boeing.com

It was a malfunctioning of a system/software!

(it seems to be a manually defined rule-based system, so far no indication about a machine learning model)

Why do we need a set of data science principles?

"So our work has the right quality for decision-making and is ultimately deployed in and trusted by the business."

Modern data-driven techniques (ML) increase expressivity but also complexity compared to well-established methods.

Extra focus on understanding prediction uncertainty, not only classification/regression.

Insurance is an industry which knows a lot about models:

- Oversight and correctness of models is culturally prioritized.
- Industry is very tightly externally regulated
- Wrong predictions can lead to severe implications (e.g. customer churn prevention vs. probability of mortality)

Principles instead of checklists and processes.
Peer-Review Board examines application of principles to projects.

We aim for mitigating model risks by...

Proper use

Ensure machine learning models are used appropriately with awareness of limitations.

Development Best Practice

Establish and build best practice among the data scientist community.

Overseen by...

Internal peer review and governance framework.

Our Principles

„Data-Related“ Principles:

- ✓ Choice of appropriate data features
- ✓ Data Quality & Governance
- ✓ Feature Engineering

„Model Development“ Principles:

- ✓ Performance Metrics
- ✓ Model Validation
- ✓ Model Calibration
- ✓ Model Uncertainty
- ✓ Robustness

„Usage“ Principles:

- ✓ Fit for purpose
- ✓ Explainability
- ✓ Recalibration
- ✓ Change Management

„Governance“ Principles:

- ✓ Reproducibility & Auditability

Data-related principles

Choice of appropriate data features

- Have you ensured that all data used present no regulatory or reputational risk?
- Could any attributes be perceived as discriminatory, for example gender/ age / religion/ ethnicity or highly correlated to them?
- Does the algorithm uses proxies for attributes being perceived as discriminatory?

Data Quality (Outliers, missing data, garbage in garbage out) & Governance

- Are the data complete?, Are they accurate? Are they appropriate?
- Validity: Does the data fully represent the real world problem? Do labels avoid any perceived systemic bias? E.g. ethnicity, social class.
- Veracity: How clean and accurate are the data?
- Volatility: Is the data relevant for the time frame or is it out of date?
- Is there a proper data management process?
- Has the integrity of the data gathering process be ensured?
- Is proper governance of data and process ensured?

Data-related principles

Feature Engineering

- Has feature normalization been done?
- Has the business knowledge for relevant features been discussed and considered?
- Can the features be appropriately tackled by the algorithm?
- Is the target variable independent of all features (no target leakage)?

Model development principles

Performance Metrics

- Are the proposed performance metrics appropriate for the given business challenge?
- Does the proposed evaluation methodology accurately capture the chosen performance metric?
- Has a meaningful benchmark been set up?
- Ex: Accuracy, distance measure, distribution divergence measure, expert user evaluation.

Model Validation

- Is the model appropriate for the business problem and data?
- Are the results brought into perspective? (e.g., do not show a loss ratio forecast without the target loss ratio in parallel)
- Is the model as simple as it can be without affecting performance?
- Have hyperparameters be appropriately selected?
- Are the data properly divided into training, validation and testing (i.e. non-independent records)?
- Have the model(s) been tested? How have the model(s) been tested?
- Are there records which might exclude customers from buying affordable insurance? Is insurability ensured?

Model development principles

Model Calibration (Over/Underfitting)

- Has adequate hold-out or testing be performed to rule out overfitting to the training set?
- Does the model perform poorly on unseen data but well on training data?
- Does the model capture the relationship between data features and labels?
- Does the model perform poorly on training data?

Model uncertainty

- Is the model uncertainty within acceptable limits?

Robustness (model perspective)

- Are the model predictions stable in the presence of noise within the data?
- Do carefully defined changes (i.e. scenarios) by various domain experts in the specific context of the data input result in small changes in the predictions?
- Are the model predictions consistent in relation to previous performance (considering the changes in the data) after re-calibration?

Usage principles

Fit for purpose (data + model)

- Do we understand how business decisions will be made based on the results of the model?
- Are the data and the model adequate to answer the business question?
- Does the business challenge translate appropriately into the corresponding modelling challenge?
- Has the appropriate trade-off between model interpretability and accuracy been reached?

Explainability

- Can the results of the model be fully described and inspected for individual records?
- Are decisions traceable? Can it be a explained why a decision was reached?
- Are simple explanation methods for a complex model available?

Usage principles

Recalibration (process perspective)

- Does there exist an agreed process for updating training data and retraining models on an ongoing basis?
- Are there measures to identify changes in the kind of data delivery (i.e. change in resolution, change in granularity)?
- Is there a fallback plan in case of problems in updating?
- What is the human oversight in production?

Change Management

- Following model development, have the changes to the business processes considered?
- Has the new process been tested and are contingency in place?
- Have the limitations of the model been communicated?

Governance principles

Reproducibility & Auditability

- Are all the data in production be kept?
- Can the training process and results be replicated?
- Is software/platform version tracked?
- Is the model appropriately documented such that a qualified third party can replicate it?
- Is the model and the process fully auditable (IT and procedural audit)?

With special thanks to...

Dr. Jürg Schelldorfer
Senior Analytics Professional

Dr. Mark Rowan
Cognitive Data Scientist

Dr. Claus Horn
Cognitive Data Scientist

Shashikiran Kandade
Senior Analyst Model Validation

Dr. Matthias Leiss
Senior Analytics Professional

Nuno Mesquita
Head DSA Predictive Modelling Solutions

Legal notice

The information and opinions contained in the presentation are provided as at the date of the presentation and are subject to change without notice. Although the information used was taken from reliable sources, Swiss Re does not accept any responsibility for the accuracy or comprehensiveness of the details given. All liability for the accuracy and completeness thereof or for any damage or loss resulting from the use of the information contained in this presentation is expressly excluded. Under no circumstances shall Swiss Re or its Group companies be liable for any financial or consequential loss relating to this presentation.