The future of machine learning: Federated learning and applications

XXXXI Heidelberg Physics Graduate Days

Heidelberg; October 11th, 2018

© d-fine — All rights reserved

Data driven consulting: topics for today and tomorrow

- Monday, 08.10.2018, 14:00 - 17:00: Dr. Anne Kleppe, Dr. Oliver Hein "defining d-fine“ and "From Physics to Finance"
- Tuesday, 09.10.2018, 14:00 - 17:00: Dr. Thorsten Sickenberger, Oliver Wohak, "Traffic simulations for innovative mobility concepts"
- Wednesday, 10.10.2018, 14:00 - 17:00: Dr. Florian Baumann, "From Monte Carlo simulation to volatility filtering: The evolution of simulation methods in market risk"
- **Thursday, 11.10.2018, 14:00 - 17:00: Dr. Patrick Biermann, Dr. Ferdinand Graf, "The Future of Machine Learning: Federated Learning - In Cooperation with DI Lab@TU Munich"**
- Friday, 12.10.2018, 14:00 - 17:00: Dr. Tassilo Christ, Dr. Patrick Sudowe, "Image recognition with machine learning methods - introduction, challenges and example applications from our consulting practice"

d-fine

© d-fine — All rights reserved

Introduction

Ferdinand Graf

- Manager (since 2011-11 working with d-fine)
- PhD in finance, diploma in mathematical finance (both @UNKN), and GARP financial risk manager
- Expert in rating model development and data-science
- Established ‘text analytics’ in d-fine’s project portfolio

Patrick Biermann

- Senior Consultant (since 2016-08 working with d-fine)
- PhD in functional analysis @ Syracuse University
- Project experience in Credit Risk, Recommender Systems and Machine Learning

d-fine

© d-fine — All rights reserved

Agenda

© d-fine — All rights reserved

Federated Learning

© d-fine — All rights reserved

TUM Data Innovation Lab

TUM Data Innovation Lab

- Educational research internship during lecture period (part time)
- Focus: Data-driven methods for interdisciplinary applications
- Target group: Master’s students from any department at Technical University of Munich
- Projects from selected partners from industry or institutions
- Students apply for one or more projects

"The TUM Data Innovation Lab stems from the enthusiasm and curiosity of its participants, students, companies, and researchers. It's an open space where creativity is our common language“ (Prof. Dr. Massimo Fornasier, Head of TUM-DI-LAB)

TUM-DI-LAB Project Schedule

- Students work together in small groups
- Mentoring by employees of the industrial partners
- Project lead by TUM
- Different working packages:
  - Literature research
  - Application to a real-world problem of the industrial partner
  - Composition of a written assignment
  - Presentation of the obtained results

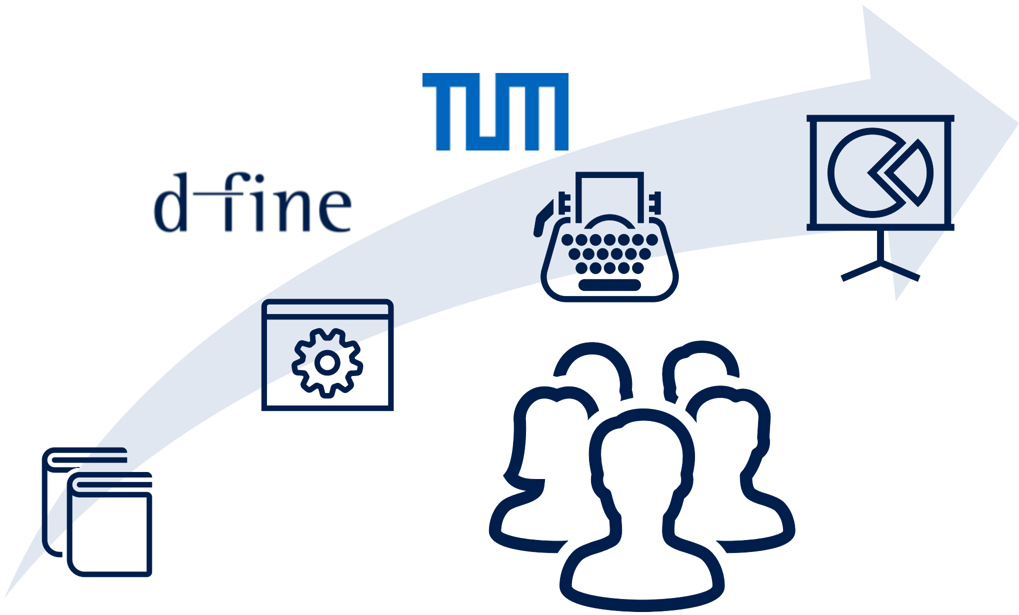

© d-fine — All rights reserved

Federated Learning stands for collaborative machine learning without centralized training data

Motivation

1. Competitive pressure forces corporations to **explore** their **data**
2. Data protection and **privacy** must be guaranteed
3. Complex or even **unstructured data** are challenging to analyse
4. Cost pressure gives incentives to avoid **manual processes**
5. **Rare events**, which may be of interest, are often insufficiently represented

Federated Learning

Promoter

a. Technological developments allow efficient communication
b. Methodological progresses bear sophisticated models
c. Concepts to define and guarantee anonymity
d. Increase in computing power and cloud computing
e. Higher data availability, awareness and openness for data driven products
f. Availability/Acceptance of standard-frameworks

d-fine

© d-fine — All rights reserved

A brief introduction to Federated Learning

1 Key features

- » Keep control of your private data and still extract utility
- » Performance can be similar to that of models trained and developed on a pooled data set
- » Communication between clients and the server is limited to propagating (noisy aggregated) model update vectors from time to time

3 Objective function

number of clients

size of client k's dataset

$$L(w) = \frac{1}{n}\sum_{k=1}^{K} n_k \underbrace{\frac{1}{n_k}\sum_{j=1}^{n_k} L(w; x_{k_j}, y_{k_j})}_{\text{local loss function}}$$

2 Illustrative example

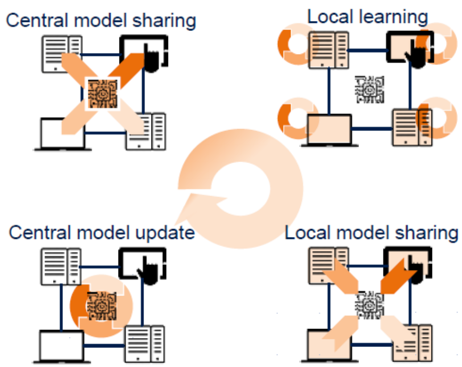

4 Optimization rule

$$w_{t+1} = w_t - \alpha_t \sum_{k=1}^{K} \frac{n_k}{n} \underbrace{\sum_{j=1}^{n_k} \frac{1}{n_k} \nabla L(w; x_{k_j}, y_k)}_{\text{local gradient}}$$

learning rate

The key idea of federated learning is relatively straightforward, most difficulties arise from alleviating the tension field between the three main success drivers: privacy, communication cost and performance.

d-fine

© d-fine — All rights reserved

Understanding the subtle differences between regular gradient descent and federated averaging

1	Gradient Descent	$$w_{t+1} = w_t - \gamma \sum_{k=1}^{K} \frac{n_k}{n} g_k$$ Compute new model weight by taking the old weight and adjusting it according to gradient descent. Gradients are reported by the clients.	
2	Federated Averaging	$$w_{t+1} = \sum_{k=1}^{K} \frac{n_k}{n} w_{t+1}^{k} \quad \text{where} \quad w_{t+1}^{k} := w_t - \gamma g_k$$ Compute new model weights by averaging the new model weights of each client.	

The only difference between these two approaches is a change in perspective.

© d-fine — All rights reserved

Communication cost describe the impact of privacy to performance

Technical costs

1 Access to free (and fast) internet

2 Impact of model updates on device performance and life

3 Depends mostly on explicit use-case

Administrative Roadblocks

I Federated Imputing

II Federated Winsorizing

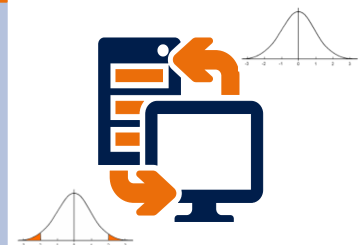

How to solve this issue in a federated setting?

Privacy Costs

Every communication round leaks/reveals information **a**

Noise from anonymizing will creep into central model **b**

Individual data must be noisified more if rounds increase to assure privacy **c**

Where technical costs are mostly a matter of convenience, the privacy cost accrued with more communication rounds can impact model performance severely.

d-fine

© d-fine — All rights reserved

Maintaining data privacy is a key issue for artificial intelligence solutions

1 Differential Privacy

Usually called (ϵ, δ) – Privacy, which gives bounds for the probability of a given sample of a given client participating in model training

» relies not on sharing the whole dataset, but only on answers to certain queries

» there are results for the change in privacy bounds under repeated querying

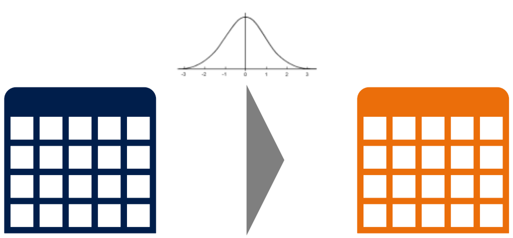

2 k-Anonymity

Refers to every observation being indistinguishable from at least k other observations

» usually obtained by binning or dropping features

» addresses the problem of anonymized datasets not being properly anonymous

» loss of information can impact performance in unpredictably ways

3 Rotation of dataset

Applying a random rotation matrix to a dataset preserves its geometric properties.

» no loss of accuracy

» easy to implement

» limits choice of algorithm, support vector machines and k-nearest neighbour classifier will still work

Of the different privacy notions, differential privacy can be locally controlled and does not restrict the choice of machine learning model

© d-fine — All rights reserved

Optimal model performance under privacy and communication constraints is the ultimate goal in federated learning

- Select a tentative performance goal
- Select adequate performance and satisficing metrics

- Privacy goals need to be set in advance
- Usually the least mutable goal, since requirements must be met

- Execute Federated Learning pipeline and evaluate results

- Pick a machine learning model with a suitable loss function

- Select federated data preparation process
- Decide on how to communicate the necessary meta parameters

- Number of communication rounds need to be determined
- Closely related to threshold of privacy selected and algorithm

Building a federated machine learning model is an iterative process with an interesting interplay of all choices made.

d-fine

© d-fine — All rights reserved

Theoretical and practical challenges are plenty in the context of federated learning

Central Problem

Is there a way to build reasonable data pipelines and set realistic expectations with regards to performance by only looking at the client datasets? What is the minimum amount of cooperation necessary to achieve acceptable performance?

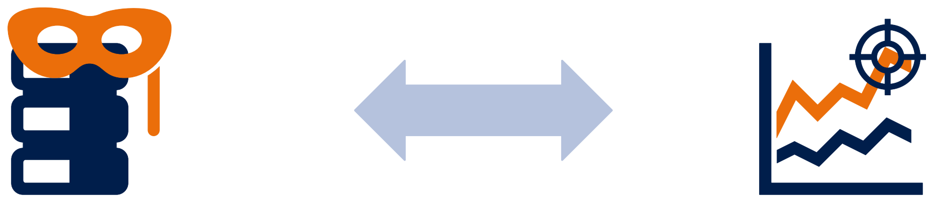

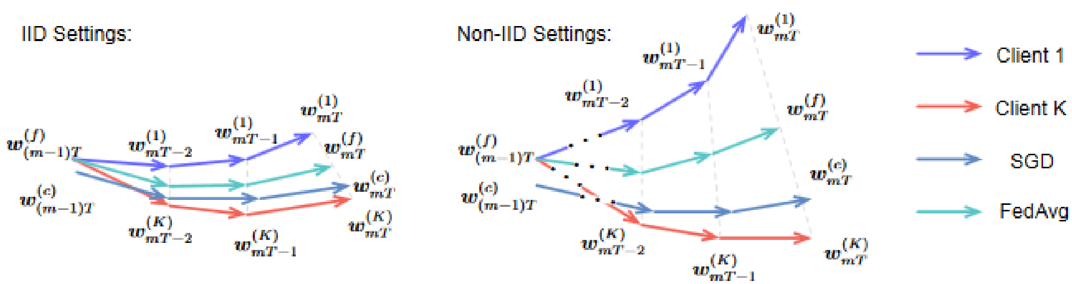

Figure 3: Illustration of the weight divergence for federated learning with IID and non-IID data.

Open Questions

- » Is there an a-priori bound on how much performance will suffer in the face of non-iid data?
- » How should information about data be exchanged?
- » What are the right ways to measure differences between datasets?
- » What are the best tools to address differences between datasets?

Finding a priori bounds on the divergence of the gradients between classic and federated learning settings is an interesting challenge

https://arxiv.org/abs/1806.00582

d-fine

© d-fine — All rights reserved

Federated Learning offers use cases in different fields of banking bussiness

Risk management

- Pooled models for credit risk parameters
  - **PD**
  - LGD
  - EAD
  - CCF
- Early-warning systems
- Fraud prevention models
  - Money laundering
  - Identity theft
  - Credit fraud

Compliance

- **Models for the detection of illegal agreements in chats of trading systems**
- Models for the identification of fraudulent email communcations and phone calls of employees

Marketing

- Recommender systems for existing customers
- Up- und Crossselling models for a more targeted use of marketing campaigns
- User classification in order to increase customer satisfaction:
  - Complexity of homepage in online banking
  - Content and frequency of customer communication (newsletters)

All use cases for federated learning share the property that they concern statistical models with **a low number of** (positive) **observations**, so that **sharing the model will lead to a substantial improvement**. At the same time they **concern highly sensible data**, which should/must not be revealed to competitors.

© d-fine — All rights reserved

Application of Federated Learning

© d-fine — All rights reserved

Building a sentiment dictionary

© d-fine — All rights reserved

KOALA – Our **KO**mmunication **A**na**L**ysis **A**pplication identifies critical communication and reduces compliance risks

Compliance requirements increased

- Shortcomings in measures to prevent insider trading and market manipulation imply huge financial and reputational risks for financial institutions
- Fines imposed by regulatory authorities on banks have been draconic, see e.g. LIBOR or Forex scandal

Support and automatization

- Based on verbal (words and word combinations) and non-verbal information (response time, number of chat participants), **communication is scored according to criticality**
- Different information sources are jointed and analysed (e.g. trader communication. Trader positions, public news and market information)
- Documentation requirements are met efficiently

Internal communication in focus

- Communication channels of traders are manifold (Email, Skype, Lync, …), and the volume exceeds human capacities by far
- Due to humor, sarcasm, abbreviations, different languages, spelling errors etc., internal communication is challenging to analyze

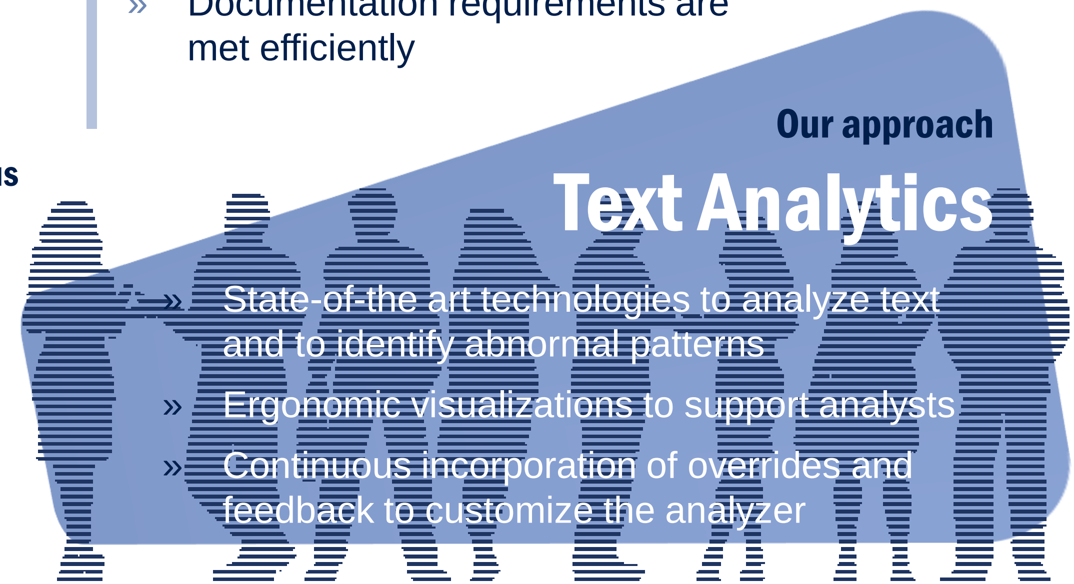

d-fine

© d-fine — All rights reserved

Popular (sentiment) dictionaries are a good starting point but might be too general / not specific enough for some purposes

1 General Inquirer

- General purpose, 182 categories (e.g. Positive, Negative, Hostile, Strong, Power, Weak, Active, Passive)
- the dictionary also contains part-of-speech tags for each word (e.g. Noun, CONJ, DET, PREP)
- Available fee of charge via http://www.wjh.harvard.edu/~inquirer/

2 Sentiment Word Lists

- Financial / economic background, i.e. constructed in 2009 with 10-K fillings
- 6 categories (Litigious, Negative, Positive, Strong, Uncertainty and Weak)
- Available free of charge via http://www3.nd.edu/~mcdonald/Word_Lists.html

3 Subjectivity Lexicon

- General purpose, contains 3 categories (positive, neutral and negative)
- Available free of charge via http://mpqa.cs.pitt.edu/lexicons/subj_lexicon/

4 Diction 5 / 7

- Contains 33 word-categories (e.g. Accomplishment, Aggression, Centrality) and 6 variables based on count ratios in the word categories
- the software is proprietary, see http://www.dictionsoftware.com/

5 Linguistic Inquiry & Word Counts

- Social and psychological background, 64 hierarchical word lists and summary statistics
- the software is proprietary, see http://liwc.wpengine.com/

6 Build your own

- Based on
  - expert knowledge
  - trainings-set, e.g. find the words with the strongest discriminant power
- Use non-dictionary based classification methods like
  - K-nearest-neighbour
  - Support vector machines
  - Naïve Bayes
  - Maximum entropy

d-fine

© d-fine — All rights reserved

Building a sentiment lexicon from labeled newspaper articles

- **Goal**: Create a sentiment lexicon, i.e. map each word in the dictionary to sentiment score
- **Approach**: Train a simple **b**ag-**o**f-**w**ords model on data set of labelled sentences
- **Training data**: Manually labelled newspaper articles. For simplicity, use article titles only.

Examples:

(-) *Analysts slam RIM's latest phone delays.*

(+) *AK Steel reports smaller loss as input costs fall.*

Overview:

Data set size	vocabulary
2794	5845

- **Model**:

$$sent(s) = \frac{1}{|s|}\sum_{x\in s} w(x), \qquad \mathrm{p} = \Pr(y=1) = \frac{1}{1+e^{-sent(s)}}$$

- *x*: word
- *s*: sentence
- *w(x)*: sentiment score of x
- y: predicted sentiment
- t: true sentiment
- D: training data

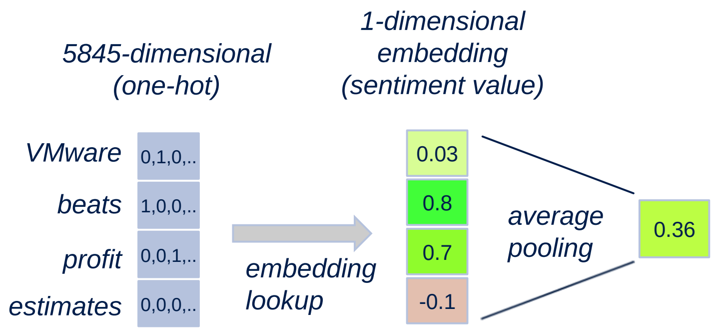

- **Training:** Tune parameters *sent(x)* to increase the prediction accuracy on the training set

Cross-entropy loss:

$$L(D) = \sum_{s\in D} -(t\log(p) + (1-t)\log(1-p))$$

SGD: $w_{t+1} \leftarrow w_t - \alpha_t \nabla L(D;w)$

FedSGD: $w_{t+1}^k \leftarrow w_t - \alpha_t \nabla L(D_k;w), \quad \mathrm{w}_{t+1} = \sum_{k=1}^{K}\frac{n_k}{n} w_{t+1}^k$

© d-fine — All rights reserved

Evaluation of the learned sentiment lexicon

Test set accuracy: 77.2 %

→ Good performance on test set indicates that the model learned sensible word sentiments!

Extract of positive and negative words

Positive words		Negative words	
Financial News	Movie Reviews	Financial News	Movie Reviews
wins, beats, jump, boost, up, buy, tops, order, rise, invest, sells, higher, project, markets	powerful, solid, fun, wonderful, enjoyable, rare, refreshing, best, entertaining	prosecutors, slips, tumbles, sued, sec, over, fine, emissions, downgrades, miss	worst, suffers, flat, stupid, dull, unfunny, lacking, too, pointless, contrived, generic

» The learned dictionary is domain specific.

Sentences where the model performs best and worst

lockheed	wins	117	billion	deal	for	early	work	on	more	f	35	jets
0.3	0.8	0.0	0.2	0.4	0.3	0.0	0.0	0.0	0.0	0.0	0.0	0.4

new	york	times	revenue	misses	as	print	ad	sales	fall	again
0.3	-0.5	0.0	0.1	-0.6	-0.1	0.0	0.0	0.1	-1.0	0.0

best	buy	results	sink	electronics	shares
0.0	0.7	0.0	0.0	0.2	0.0

» For sentences with simple structure, the BOW approach yields good results

» For sentences with complicated structure a more complex model is needed that captures the relationship between words

d-fine

© d-fine — All rights reserved

The size of the sentiment lexicon can be controlled by regularization

- Non-regularized optimization results in a very large sentiment lexicon, i.e. non-zero sentiment scores are assigned to almost all words in the dictionary.
- In reality, however, most words are neutral and should therefore not be part of the sentiment lexicon.
- Solution: Use L1-regularization
  - Has been shown to induce sparsity
  - Easier to optimize than L0-regularization

$$\tilde{L}(D) = L(D) + \gamma||\mathrm{w}||_1$$

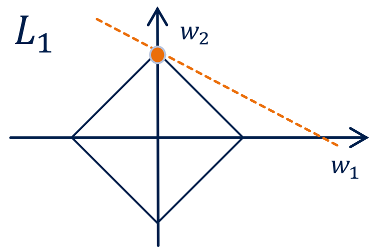

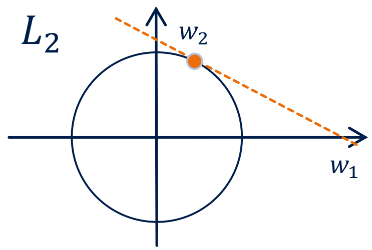

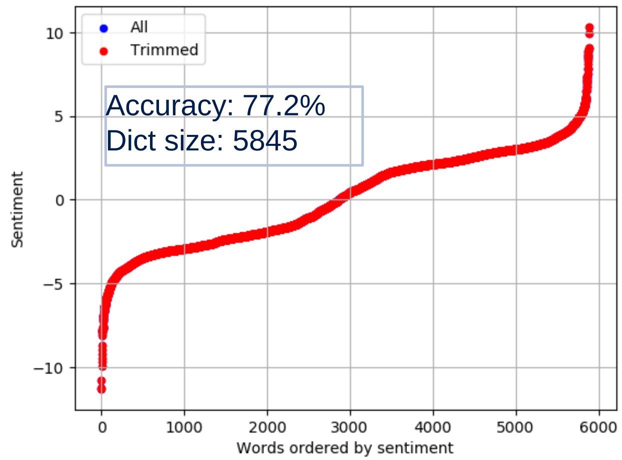

L1-Regularization

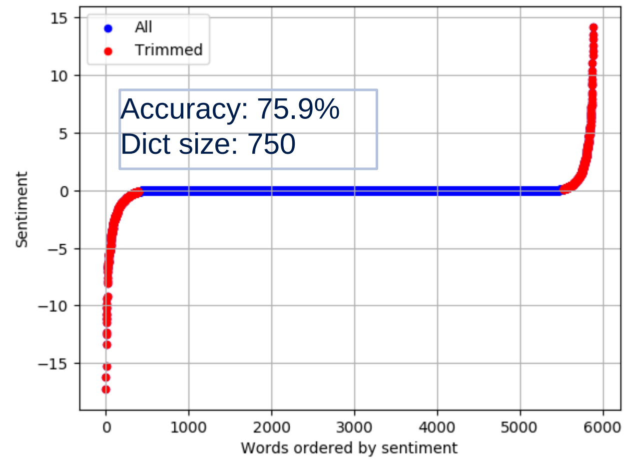

© d-fine — All rights reserved

Credit risk scoring

© d-fine — All rights reserved

d-fine

What is credit risk and why is credit risk modeling needed?
A stylized example (1/3)

A simple business case

- A financial institution finances mortgages for private individuals. Assume that the institution has 100 identical mortgage loans in its portfolio, each of which is worth € 200,000 and is paid back completely after one year with a lump-sum payment.
- Unfortunately, five of the mortgage holders default and the institution is only able to recover € 50,000 of each of the original amounts being extended to the obligors.

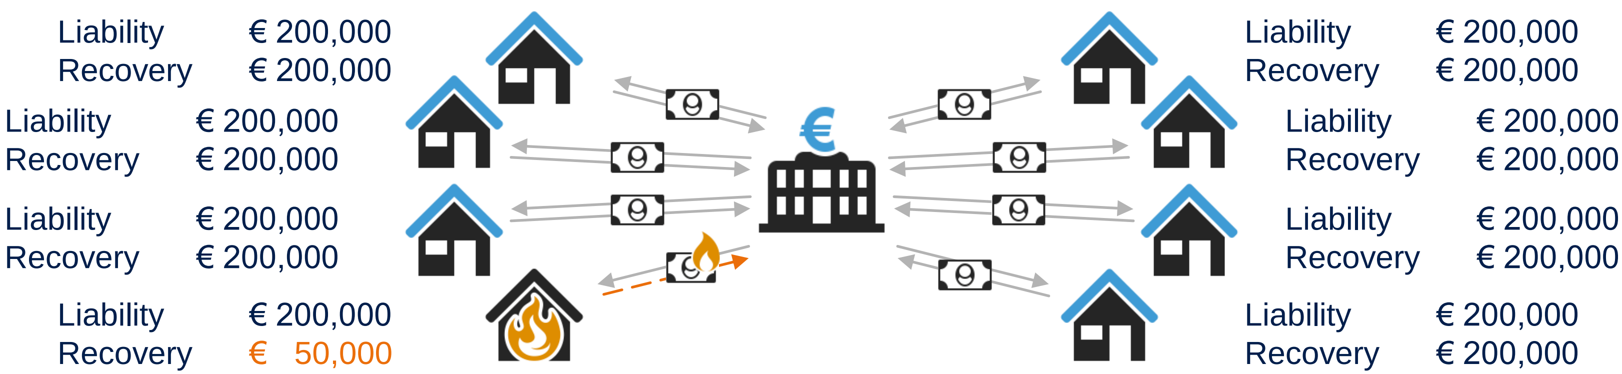

© d-fine — All rights reserved

What is credit risk and why is credit risk modeling needed?
A stylized example (2/3)

A simple solution

- » The institution faces a total loss of € 750,000 (= 5 × € 150,000). This is about 3.95% (= € 750,000 / € 19,000,000) of the total amount that is successfully paid back.
- » To cover its losses, the institution could require every (new) obligor to pay an interest rate (or risk charge) of at least 3.95% (i.e. a lump-sum payment of € 207,895). Since the expected losses due to intermittent defaults is covered by the risk charge, the total loss for the institution is expected to be zero.
- » However, this approach has a huge drawback: „Good“ obligors may refuse to pay such a high interest rate, while this may seem cheap for „bad“ obligors. It may thus lead to an adverse selection of obligors.

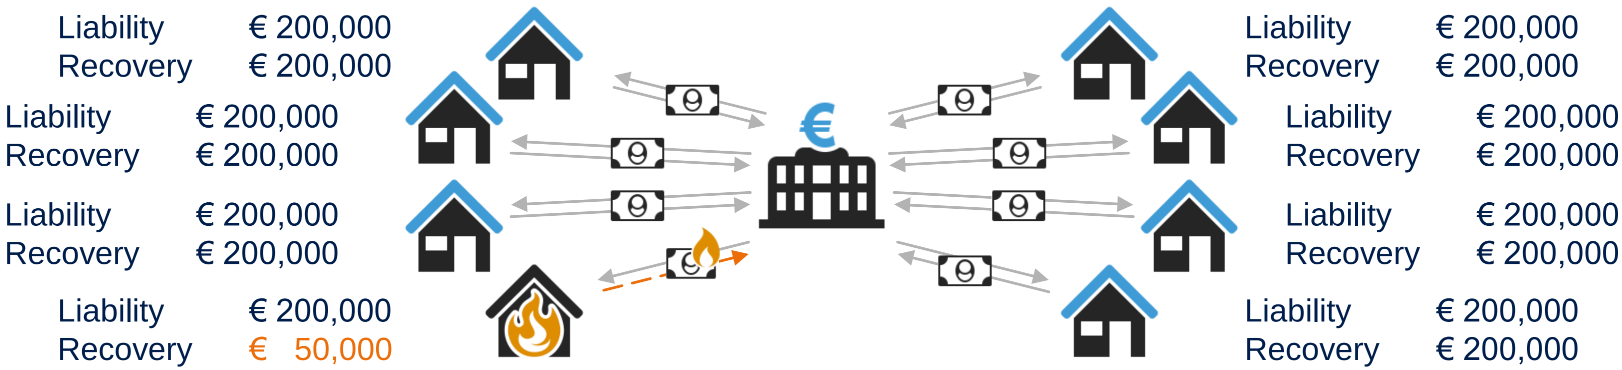

d-fine

© d-fine — All rights reserved

What is credit risk and why is credit risk modeling needed?
A stylized example (3/3)

Credit risk is the risk that a borrower fails to repay his loan, thereby generating a loss for the lender. As a matter of principle, one distinguishes expected (EL) and unexpected losses (UL).

Assessment of the individual obligor

Probability of Default (PD): What is the probability that the obligor will default?

Exposure at Default (EAD): Given that the obligor will default, how much money does (s)he owe at the time of default (e.g. for a credit card with a certain line of credit)?

Loss given Default (LGD): Given that the obligor will default, what is the relative share of money that the institution will be able to recover?

Assessment of the overall portfolio

Expected Loss (EL): How much does the institution need to collect by way of risk charges, i.e. what is the expected loss due to defaults within the given portfolio?

Unexpected Loss (UL): What is the worst case scenario leading to bankruptcy?

Portfolio Loss Distribution

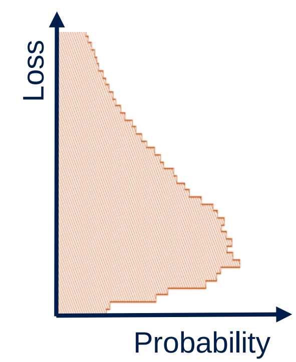

d-fine

© d-fine — All rights reserved

A statistical view (1/2)

Calculating the Expected Loss

For every obligor i, one can define the following random variables:

Default Indicator	*Exposure at Default*	*(Relative) Loss given Default*	*Total portfolio loss (amount)*	*Total loss (amount) for obligor i*
$1_D = \begin{cases} 1 \text{ if obligor defaults} \\ 0 \text{ if obligor does not default} \end{cases}$	EAD	LGD	L	L_i

Generally, if an obligor i defaults, the realized loss is the product of their realized exposure and the loss given their default:

$$L_i = EAD_i \cdot LGD_i \cdot 1_{D_i}$$

Having defined this random variable, the expected loss, too, is a random variable given by:

$$\mathbb{E}[L] = \mathbb{E}\left[\sum_i L_i\right] = \sum_i \mathbb{E}[EAD_i \cdot LGD_i \cdot 1_{D_i}] = \sum_i \mathbb{E}[EAD_i] \cdot \mathbb{E}[LGD_i] \cdot \mathbb{E}[1_{D_i}] = \sum_i \mathbb{E}[EAD_i] \cdot \mathbb{E}[LGD_i] \cdot \mathbb{P}[D_i]$$

Assumption: EAD, LGD, and the default indicator are uncorrelated.

This expression is typically denoted as:

$$EL = \sum_i EAD_i \cdot LGD_i \cdot PD_i$$

Or, for a single obligor:

$$EL_i = EAD_i \cdot LGD_i \cdot PD_i$$

Irrespective of whether or not a split into multiplicative components is justified and regardless of correlations between different obligors, the expected losses of different obligors are additive.

d-fine

© d-fine — All rights reserved

A statistical view (2/2)
Calculating the Unexpected Loss

The **Credit Value at Risk (CVaR)** is the maximum loss that occurs with a certain probability α (often chosen to be 99.9%), i.e. the loss l^* such that:

$$\mathbb{P}[L \leq l^*] = \alpha$$

The **unexpected loss** is defined as the difference between CVaR and the expected loss:

$$UL = CVaR \; - EL$$

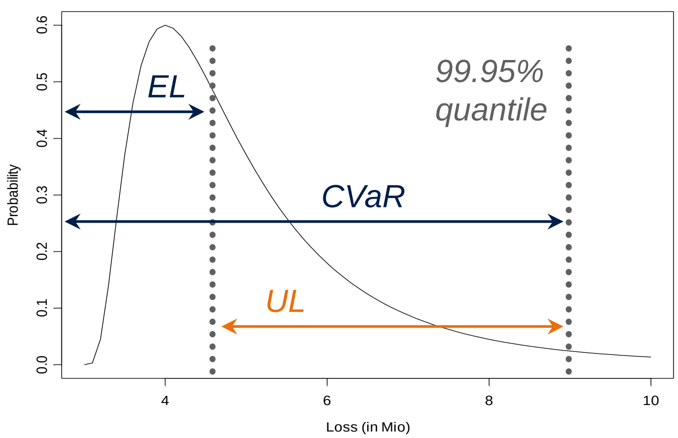

The unexpected loss strongly depends on the correlations between the loans in a portfolio.

A Simple Example:

A portfolio with two loans, each with a PD of 20%, an LGD of 100% and an EAD of € 100 for three different correlations (-100%, 0% and +100%)

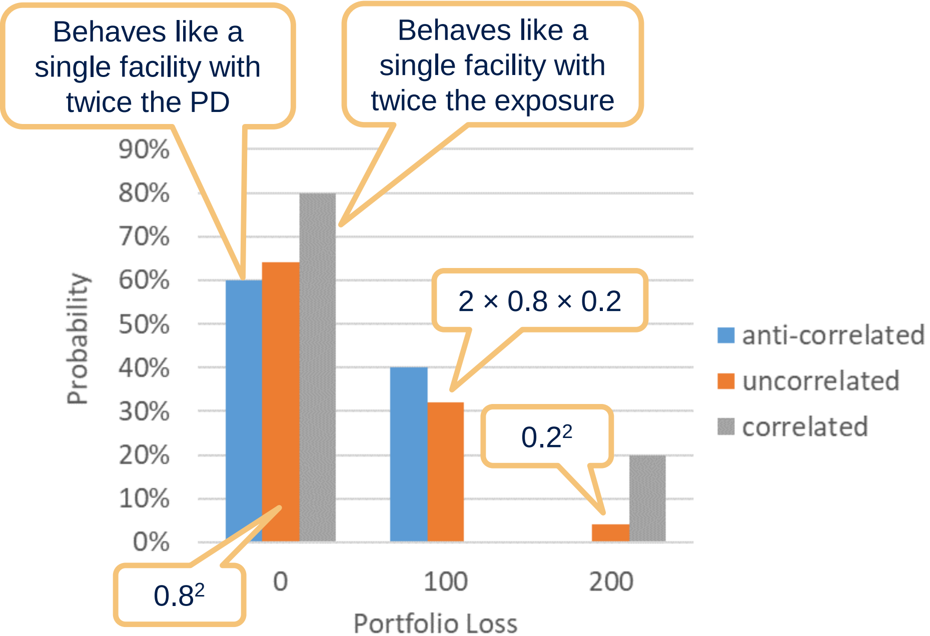

The contribution of a single loan to the total unexpected loss depends on the portfolio and its correlations. It can thus not be calculated without specific knowledge about the portfolio.

d-fine

© d-fine — All rights reserved

To estimate the probability of default of an obligor or exposure, one needs to find appropriate “risk factors”

Idea: *Ex ante* assessment of the quality of potential obligors to be able to tailor the risk charge specifically to an institution’s risk appetite

Estimating the Probability of Default

Analysis of the obligor’s characteristics to determine his/her creditworthiness

► Discovery of “**risk factors**”

- » Risk factors can be any obligor-, contract-, or behavior-related information, either publicly available or subject to confidentiality agreements between the institution and the obligor
- » The perfect risk factor would be a single attribute that clearly discriminates between „good“ and „bad“ obligors
- » Yet, such a risk factors do not exist in practice. Even good risk factors discriminate only partly between good and bad obligors.

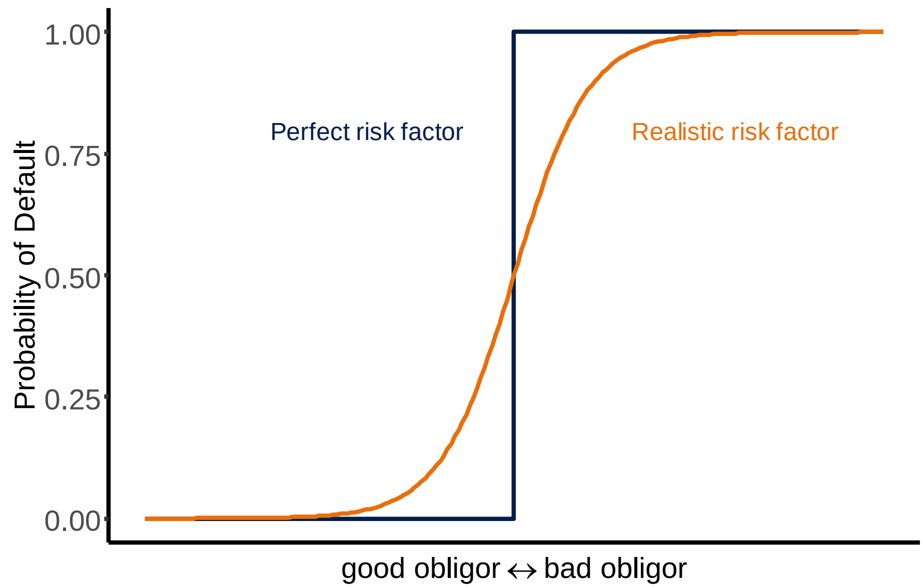

Example:
Monthly income ≥ € 2,000 vs. monthly income < € 2,000

„All obligors with a monthly income below € 2,000 cannot repay their loans.“

The gist of building a meaningful / robust PD model is to identify and combine several “good” risk factors.

d-fine

 © d-fine — All rights reserved

State-of-the-art platforms for (corporate) ratings are usually modular and cover different areas of information

Risk factors are statistically aggregated to a rating grade

specific

Micro enterprises	**Small and mid enterprises**	**Large corporation**
Turnover ≤ 50’’ EUR	Turnover ∈ (50’’ EUR; 1’’’ EUR)	Turnover ≥ 1’’’ EUR

» Statistically optimized combination of module scores
» Manual adjustments due to obligor-specific circumstances if necessary

generic

Static obligor information	Payment history	Accounting figures	Qualitative information	Capital market information	Unstructured information
Founding date	Amount past due	Revenue growth	Management quality	Market value of equity	Reuters, BBG, FT Nachrichten
Legal form	Limit usage	Capital structure	Market position	Volatility.	Twitter
Number of employees	Current balance	Net profit ratio.	Business development.	Trading volume	Facebook
…	…	…	…	…	…

d-fine

© d-fine — All rights reserved

How to develop a (standard) PD model
Step 1

Find risk factors with good discriminatory power → Combine the risk factors to a single measure („score“) → Calibrate the score to a probability of default

Types of Risk Factors

Metric risk factors

- » Values are numbers
- » There is a clear rank order
- » Examples: credit amount, equity ratio of a company, time to maturity of the loan, age of the obligor, etc.

Categorical risk factors

- » Values do not necessarily have to be numbers
- » Typically no rank order
- » Examples: marital status, gender, profession, purpose of financing, zip code, etc.

Choice of Risk Factors (“long list”)

Find all available, appropriate, and economically plausible criteria with sufficient discriminatory power:

- » Application criteria (income, domicile, etc.)
- » Information from financial statements (sales, earnings, etc.)
- » Behavioral criteria (cumulative days past due, current arrears, etc.)

 © d-fine — All rights reserved

How to develop a (standard) PD model

Step 2

Find risk factors with good discriminatory power → **Combine the risk factors to a single measure („score“)** → Calibrate the score to a probability of default

Objective: Improve *ex ante* separation between *ex post* defaulted and non-defaulted obligors

Abstract Concept: Creation of a mapping from the selected risk factor space to the real numbers (score)

$$\begin{pmatrix} \text{risk factor 1} \\ \text{risk factor 2} \\ \vdots \\ \text{risk factor N} \end{pmatrix} \xrightarrow{\text{score function}} \text{score}$$

<u>Illustration</u>:

$$\text{score} = \text{weight}_1 \times \text{value}(\text{criterion}_1) + \ldots + \text{weight}_n \times \text{value}(\text{criterion}_n)$$

Example: Assign points to the answers of a questionnaire and interpret their sum as a classification

Distribution of the risk factors

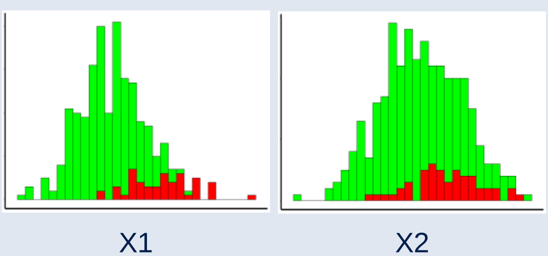

Analyse multivariate distribution

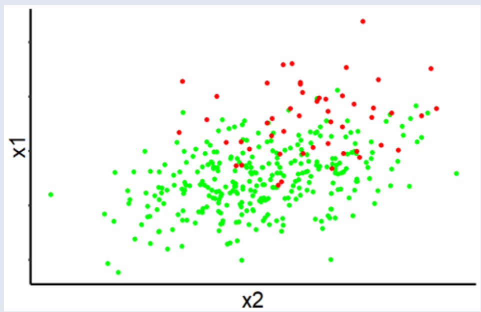

Distribution of the combined score

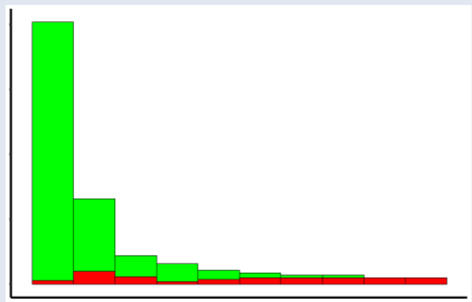

Discriminant analysis can be used to optimize the score. Yet, in practice, logistic regression tends to be preferred, as it yields more robust results and allows for a simpler consideration of categorical risk factors.

d-fine

© d-fine — All rights reserved

How to develop a (standard) PD model
Step 3

Find risk factors with good discriminatory power

Combine the risk factors to a single measure („score“)

Calibrate the score to a probability of default

Why are the optimization of the score and the calibration of the PD different steps?

Optimization of the Score: Introduction of an ordering among obligors

Calibration of the Score: Mapping of the ordering to a economically reasonable level of the probability of default

If the steps are separated, each can be adjusted without affecting the other.

General idea behind the calibration method

- Look at historic transaction data
- Calculate a score for each transaction
- Determine whether or not the transaction defaulted
- Assign a probability of default to a score such that the overall probability of default matches the observed default rate

$$\text{Target PD} = \frac{N_{\text{defaulted}}}{N_{\text{all obligors}}}$$

The PD is calibrated by comparing the score with the historically observed default frequencies.

 © d-fine — All rights reserved

A standard technique to calibrate the score to a PD

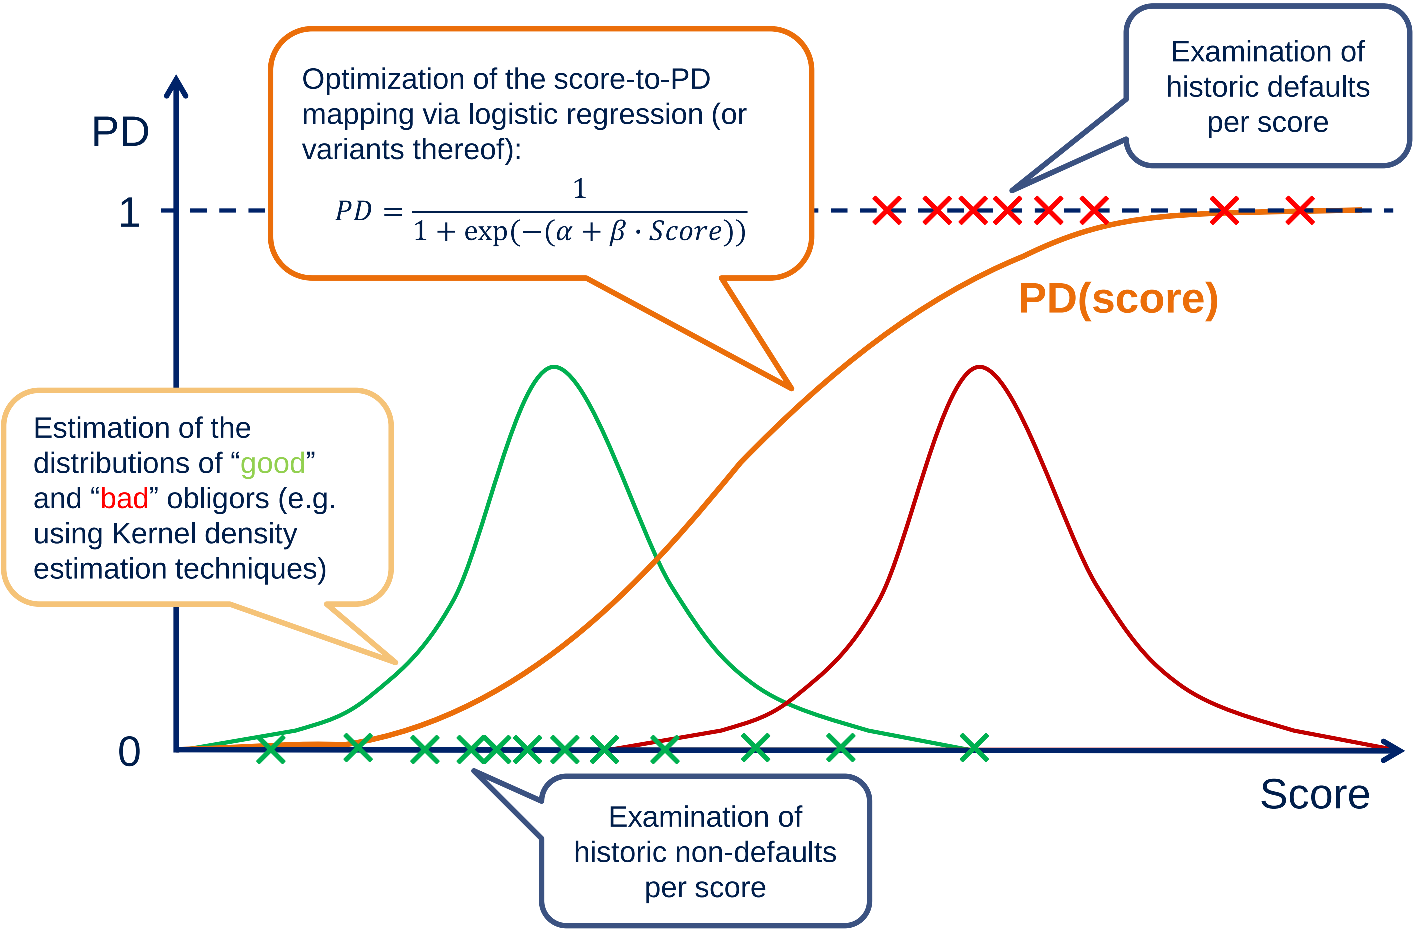

A credit risk model needs to be calibrated properly. Yet, even a well-calibrate model requires a high discriminatory power to be efficient.

d-fine

© d-fine — All rights reserved

Given the global importance of financial stability, banking regulation is developed within and promoted by an international standard setting body

Article 174 (**Use of models**): If an institution uses statistical models and other mechanical methods to assign exposures to obligors or facilities grades or pools, the following requirements shall be met:

- the model shall have good predictive power and capital requirements shall not be distorted as a result of its use. The input variables shall form a reasonable and effective basis for the resulting predictions. The model shall not have material biases;
- […]
- the data used to build the model shall be representative of the population of the institution's actual obligors or exposures;

» *Article 179* (**Overall requirements for estimation**): In quantifying the risk parameters to be associated with rating grades or pools, institutions shall apply the following requirements:

- an institution's own estimates of the risk parameters PD, LGD, conversion factor and EL shall incorporate all relevant data, information and methods. The estimates shall be derived using both historical experience and empirical evidence, and not based purely on judgmental considerations. The estimates shall be plausible and intuitive and shall be based on the material drivers of the respective risk parameters. The less data an institution has, the more conservative it shall be in its estimation;
- […]

PD models that make use of federated learning methods may have several advantages with regard to representativeness, robustness and data protection compared to standard (pool) models.

d-fine

© d-fine — All rights reserved

Live demonstration of a Federated Learning application

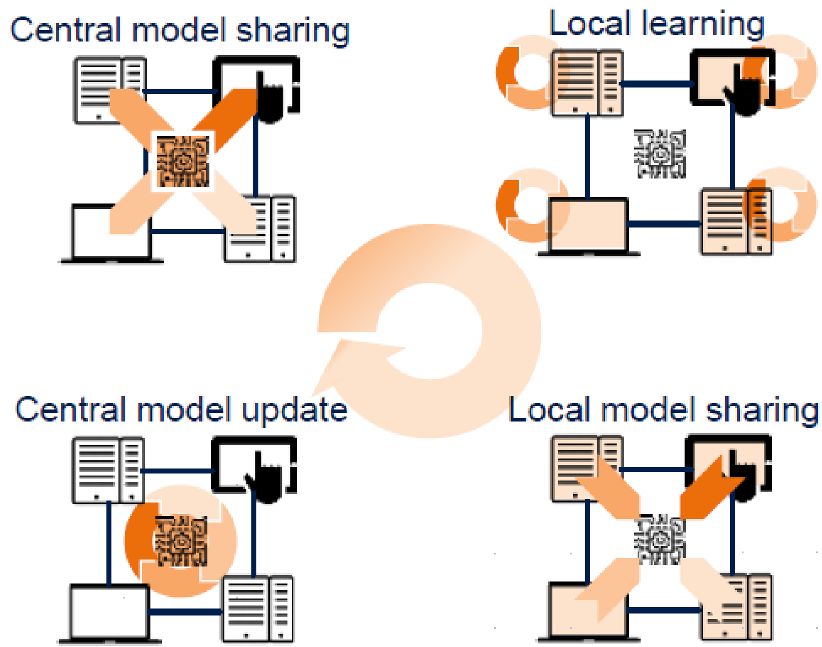

d-fine

© d-fine — All rights reserved

Chat- & Voice Bots

© d-fine — All rights reserved

Landscape of chat-bots

Types of chat-bots | Interfaces | Different Designs

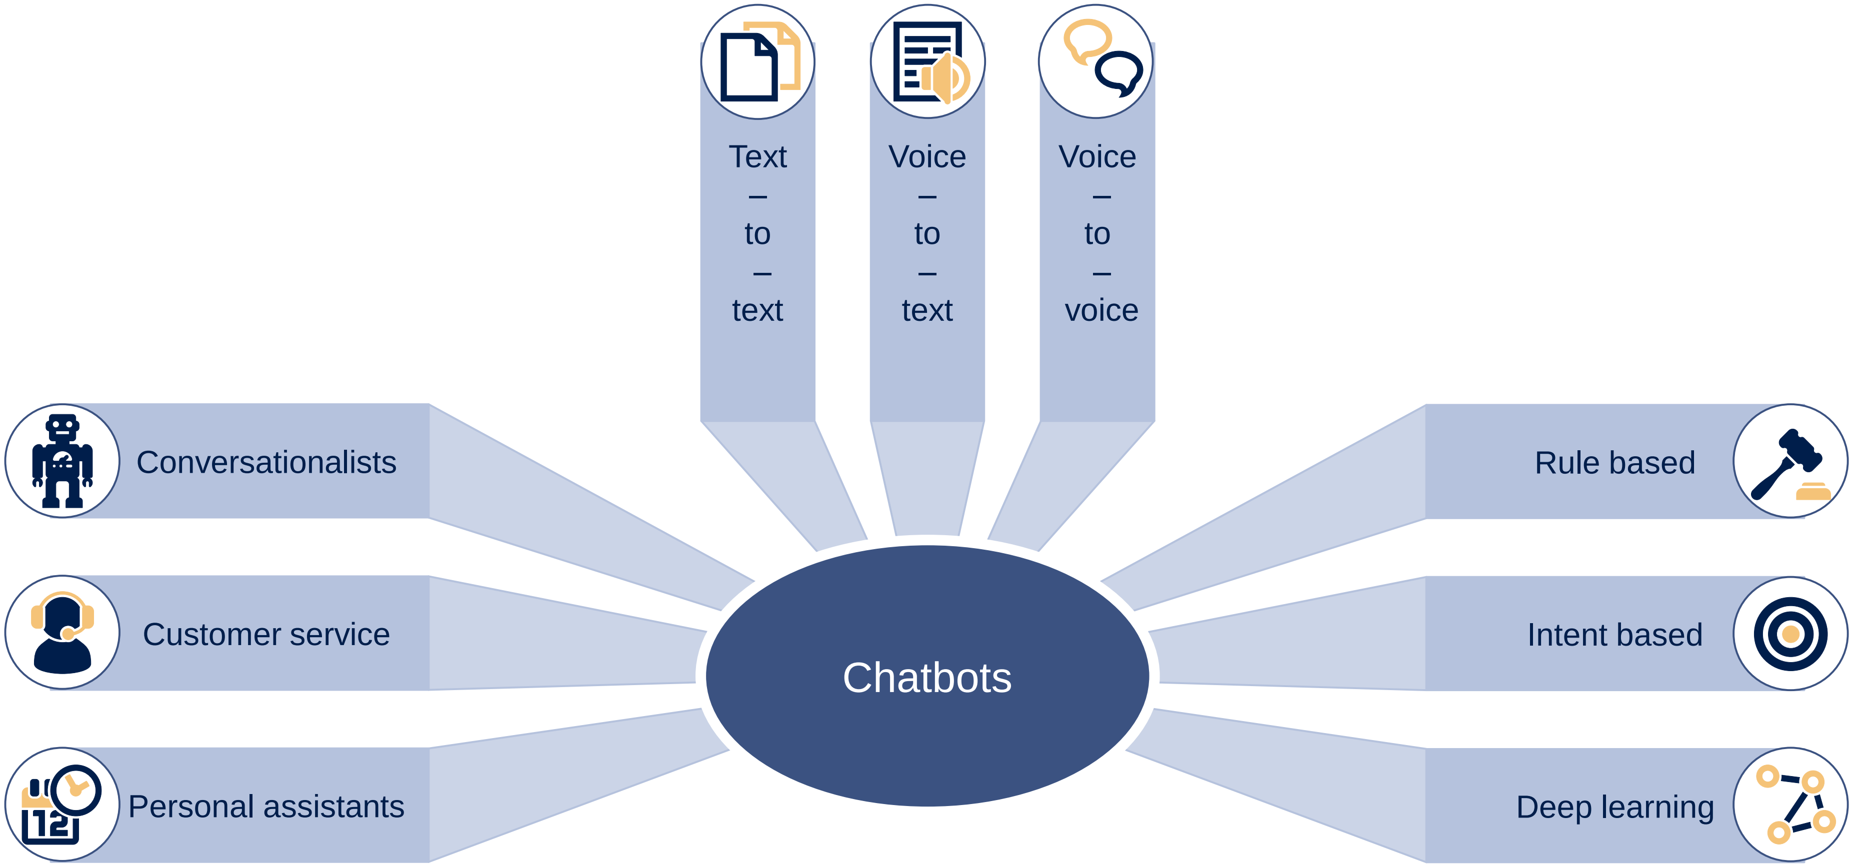

© d-fine — All rights reserved

A brief history of chat-bots: Highlights and milestones

- **1966** – **ELIZA** simulates a therapist & is the first noteworthy NLP computer program
- **1972** – **Parry** simulates a person with paranoid schizophrenia
- **1988** – **Jabberwacky** was designed to pass the Turing-Test, but failed
- **1992** – **Dr. Sbaitso** simulates a psychologist & was designed for Text-to-Speech operations
- **1995** – **A.L.I.C.E** engages in human conversation by applying some heuristical pattern matching
- **2001** – **Smarterchild** was distributed across global instant messaging and SMS networks
- **2006** – **IBM Watson** is a question-answering computer system
- **2011** – **IBM Watson** wins Jeopardy!
- **2012** – Apple and Google launch their personal assistants **Siri** and **Google Now**
- **2015** – **Eugene Groostman** is the first chat-bot to beat the Turing-Test
- **2016** – Amazon and Microsoft bring **Alexa** and **Cortana** to the market

While early chat-bots mainly utilised sophisticated rulesets, current developments leverage natural language understanding (NLU) and deep learning.

d-fine

© d-fine — All rights reserved

A chat-/voice-bot as tailor-made solution may be adapted individually to processes, IT-systems, and requirements

Conformity

- with established processes
- with organizational structure
- with specific requirements

Adaptability

- Seamless integration into existing IT-infrastructure
- Integration and utilization of existing data- and IT-infrastructure

Data Protection

- Full compliance with security standards and requirements during implementation
- Optional full control over data storage

Expandability

- Application is extendible at will
- Simple replacement and update of single modules

Cost-efficiency

- Middle- and long-term lower price than standard solutions
- No license fees, therefore very low running costs
- Low follow-up costs for further developments since this can be carried out by own IT

Self-determination

- Client is owner of the solution
- Autonomy for maintenance and further development
- Development of internal know-how due to participation in technical implementation

© d-fine — All rights reserved

How to start a chatbot project

Preselection of tasks

- » High volume
- » Frequently reoccurring
- » Personnel-intensive
- » Customer-agent interactions

Definition of scope

- » Start with one application that is promising to be accepted by users
- » Improve your solution
- » Adapt to customer’s needs
- » Expand scope

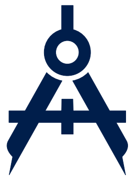

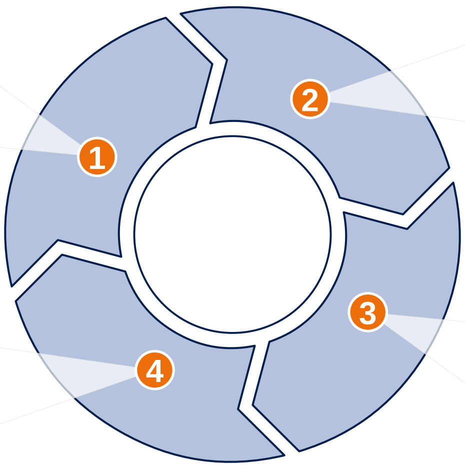

Collection of training data

- » Dialogues and chat protocols
- » FAQs, different formulations of same question, unified answers
- » Customer service agents (to become bot trainers)

Definition of goals

- » Decrease volume for manual processing
- » Turn data into information, and information into insight
- » Increase sales

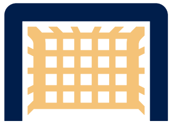

© d-fine — All rights reserved

Business Cases

Case 1: FAQ-Bot

- Answers general questions on products or services
- Easy to train if training data are available (e.g. FAQs)
- Short time to market
- High user acceptance

Case 2: Customer assistant

- **Intent based** bot to answer user-specific questions e.g. about existing contracts, specific payments, etc.
- Use of ***information retrieval*** and ***text summary*** techniques
- Combination of different data sources

Case 3: Navigation bot

- Alternative navigation possibility without clicking through menus
- Particularly useful for mobile devices
- Click-through behaviour can be used for continuous improvement

Case 4: Recommender bot

- **Recommender system**
- **Cluster** customers by behaviour
- Suggest personal *upgrades* or *complementary* products to exploit cross- and up-selling potential
- Add-on for other bots

Common Goals

- **Increase user experience**
  - Make requested information easily accessible to customers
  - Improve customer satisfaction and loyalty
- **Increase efficiency**
  - Decrease amount of requests that have to be processed by human agents
  - Increase employee satisfaction
- **Learn about your customers**
  - Collect information on user behaviour
  - Learn about users' demands
- **Support your marketing department**
  - Explore further data sources
  - Targeted advertising

d-fine

© d-fine — All rights reserved

Our solution provides complete control over the entire training cylce of the AI-engine

Question

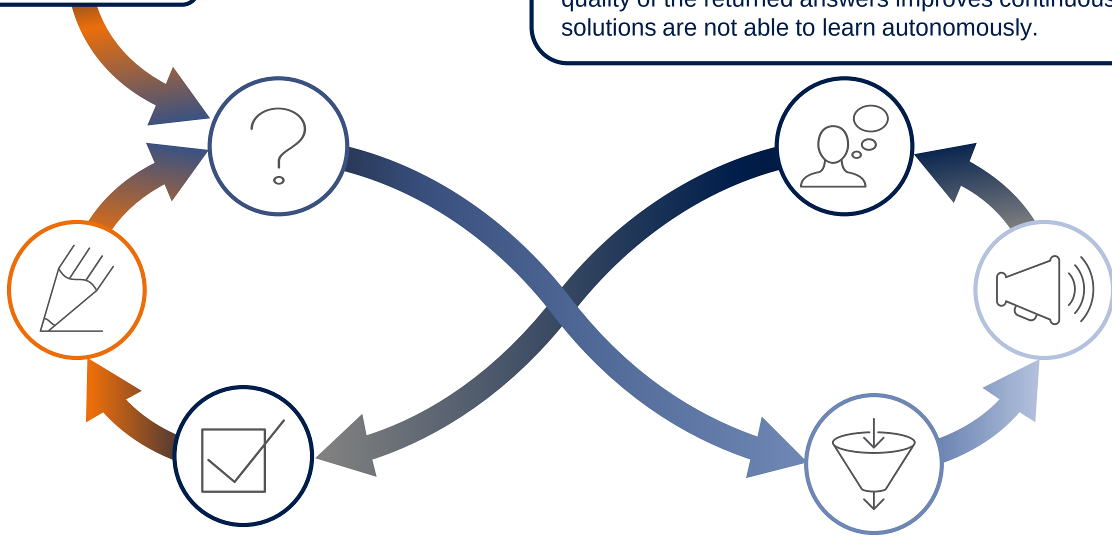

Answer

Feedback

Clients provide feedback for received answers. Based on this feedback, the algorithm learns autonomously. Therefore, the recognition of the intention and the quality of the returned answers improves continuously. Most of the standard solutions are not able to learn autonomously.

Progress

Expansion of the training corpus by adding a validated pair of question/answer in a web application.

Review

An employee checks the received feedback in a web application for monitoring and fraud prevention.

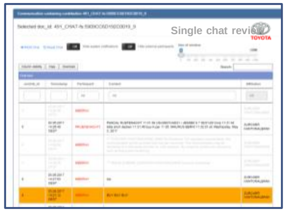

AI answering module

Voice recognition by analysis of syntax, semantics, and pragmatics. Instead of fixed examples and rules: detection of patterns and principles by employing self-learning algorithms for evaluating unknown questions of users.

Standard solutions do not provide feedback based user control over training process – though this is essention for continuous improvement of quality.

d-fine

© d-fine — All rights reserved

The chatbot's learning process leverages on a dynamically growing training set to answer questions

1 Learning process

- Question-Answer pairs serve as training set to answer new questions asked.
- For questions asked we calculate the similarity to all questions in the training set.
- The answer from the best matching Question-Answer pair is used to answer the new question given that the similarity is above some threshold.
- In order to train the chatbot dynamically, it randomly asks for feedback, e.g. "Was this answer helpful?"
  - If the user answers "yes", the new question and the answer is added to the training set.
  - If the user answers "no", a human analyst takes over, answers the question and adds the new question and the new answer to the training set.
- In both cases, a question similar to the one asked before, can be answered the next time it is asked.
- This approach works in almost all languages and can handle standard requests automatically after a short training period.

2 Illustration

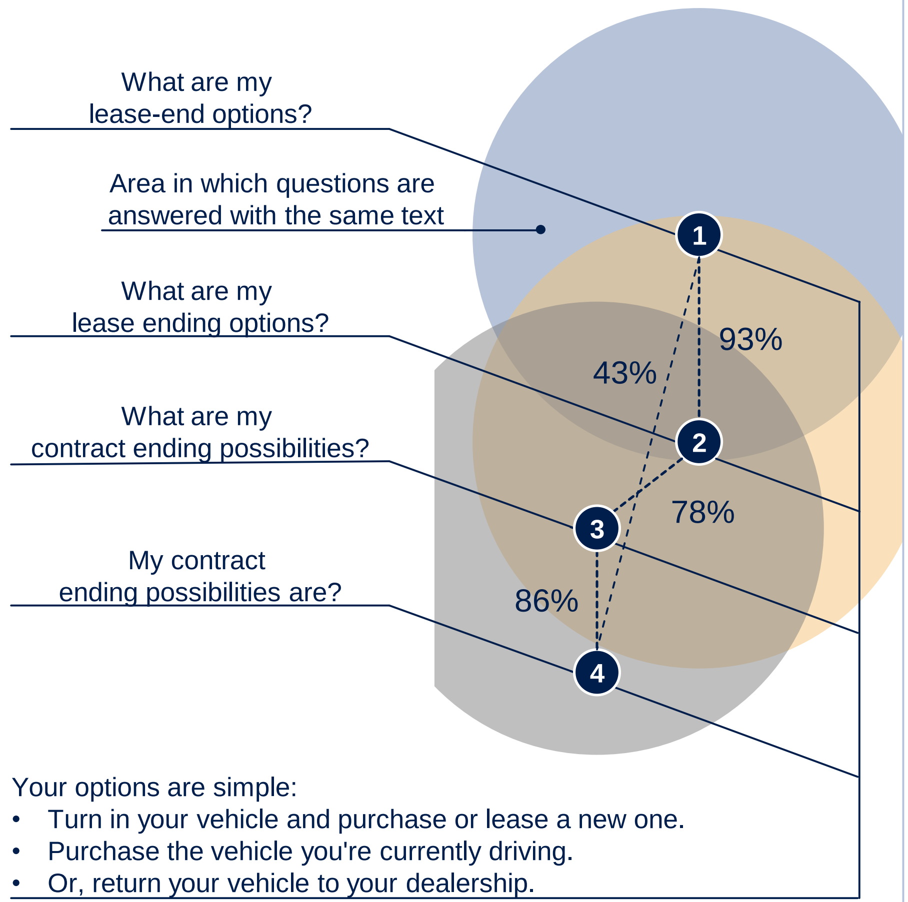

d-fine

© d-fine — All rights reserved

Federated Learning and chat bots

Voice-to-text interface

Make the speech recognition robust against
- › background noise
- › dialects

with federated learning approaches

Intent extraction and information retrieval

- » Distributional semantic approaches to enhance similarity measures, tf/idf, word-embeddings
- » Frame semantic approaches to extent dictionaries and word lists

Text summary and answer building

- » Combining answers to leave the span of formulated (known) answers
- » Gather answers automatically from quality-assured documents

Service offerings

- » Proactive recommendation of products and services based on request
- » Adjustments in the navigation of webpages, wording of contracts or other explanations

d-fine

© d-fine — All rights reserved

Contact

Dr Ferdinand Graf
Tel +49 89 7908617-0
Mobile +49 162 2630080
E-Mail Ferdinand.Graf@d-fine.de

Dr Patrick Biermann
Tel +49 89 7908617-0
Mobile +49 152 57975133
E-Mail Patrick.Biermann@d-fine.de

d-fine

Berlin
Dusseldorf
Frankfurt
London
Munich
Vienna
Zurich

Headquarters

d-fine GmbH
An der Hauptwache 7
D-60313 Frankfurt/Main
Germany

Tel +49 69 90737-0
Fax +49 69 90737-200

www.d-fine.com

d-fine

© d-fine — All rights reserved

d-fine