

Patrick Philipp\*, Nicole Merkle, Kai Gand, and Carola Gißke

# Continuous support for rehabilitation using machine learning

<https://doi.org/10.1515/itit-2019-0022>

Received June 29, 2019; revised September 30, 2019; accepted November 8, 2019

**Abstract:** Providing a suitable rehabilitation at home after an acute episode or a chronic disease is a major issue as it helps people to live independently and enhance their quality of life. However, as the rehabilitation period usually lasts some months, the continuity of care is often interrupted in the transition from the hospital to the home. Relieving the healthcare system and personalizing the care or even bringing care to the patients' home to a greater extent is, in consequence, the superior need. This is why we propose to make use of information technology to come to participatory design driven by users needs and the personalisation of the care pathways enabled by technology. To allow this, patient rehabilitation at home needs to be supported by automatic decision-making, as physicians cannot constantly supervise the rehabilitation process. Thus, we need computer-assisted patient rehabilitation, which monitors the fitness of the current patient plan to detect sub-optimality, proposes personalised changes for a patient and eventually generalizes over patients and proposes better initial plans. Therefore, we will explain the use case of patient rehabilitation at home, the basic challenges in this field and machine learning applications that could address these challenges by technical means.

**Keywords:** Automated Patient Rehabilitation, Machine Learning, Reinforcement Learning, Semantic Web

**ACM CCS:** Computing methodologies → Machine learning → Machine learning approaches → Rule learning

\*Corresponding author: Patrick Philipp, Forschungszentrum Informatik, Information Process Engineering, D-76131 Karlsruhe, Germany, e-mail: philipp@fzi.de, ORCID: <https://orcid.org/0000-0002-5831-1018>

Nicole Merkle, Forschungszentrum Informatik, Information Process Engineering, D-76131 Karlsruhe, Germany, e-mail: merkle@fzi.de  
Kai Gand, Carola Gißke, TU Dresden, Chair of Wirtschaftsinformatik, esp. Systems Development, D-01062 Dresden, Germany, e-mails: kai.gand@tu-dresden.de, carola.gisske@tu-dresden.de

## 1 Introduction

One out of six people in the European Union has a disability [8], usually caused by an acute episode or a chronic disease. Providing a suitable rehabilitation is the main issue for people as they age. This helps people to live independently and enhance their Quality of Life. However, as the rehabilitation period usually lasts some months, the continuity of care often is interrupted in the transition from hospital to home.

Relieving the healthcare system and personalising the care or even bringing care to the patients' home to a greater extent is, in consequence, the superior need. This is why information technology shall be used to come to participatory design, driven by users needs, and personalisation of clinical pathways enabled by technology [9]. This way, patients can be encouraged and empowered to proceed with a personalised rehabilitation that complies with age-related conditions and enhances the adherence to the care plan and the risk prevention [5].

To allow the desired relief of the healthcare system, patient rehabilitation needs to be supported by automatic decision-making, as physicians cannot constantly supervise the rehabilitation process.

While physicians can setup a first rehabilitation plan (or clinical pathway), the latter is usually not sufficiently personalisable in the beginning. One could, for example, setup a daily schedule for walking 2 km, distinct mobility exercises at home to recover or ensure physical fitness. This could be combined with advises for a healthy living in general and scheduling mechanisms to not forget the proposed measures. Suggesting, and therefore scheduling, social contact to avoid loneliness could be included as well.

At this point, one has to continuously monitor how a patient is following his/her schedule and intervene as soon as possible if the patient feels unmotivated and neglects the plan. As this is not feasible by using physicians, computer-assisted patient rehabilitation is needed which monitors the state of the current patient plan to detect sub-optimality, proposes personalised changes for a patient and eventually generalises over patients and proposes better initial plans.

In the following, the use case of *patient rehabilitation*, the basic challenges in this field as well as machine

learning applications that could address these challenges by technical means will be explained. It is argued that prior work on the problem which exploits standard supervised machine learning (e. g. [45, 44, 23]) does not capture crucial dependencies among individual recommendations and does not allow to successfully explore alternative activities in case interventions are needed. A decision-theoretic model for clinical pathways is thus proposed which can be solved by Reinforcement Learning (RL) [39, 34], where safety of the system is ensured by formalised knowledge of physicians, modeled by using knowledge representations from the Semantic Web. In addition, the potential of using unsupervised machine learning is investigated which can foster efficiency and explainability of the resulting system.

## 2 Use case: Patient rehabilitation

The first part of this section deals with introducing the use case of patient rehabilitation. Based on the latter, the goal of this paper is defined by dwelling on central machine learning-based challenges.

### 2.1 Description & project context

The purpose of home rehabilitation is to reestablish independence and function of the patient, or up to that level, that patients can readily take part in different activities to pursue life and be a productive member of the society [4].

The use case is, in particular, an EU-funded project that addresses the rehabilitation after an acute severe episode causing chronic disabilities (restricting the patient's abilities of movement and physical resilience) such as Parkinson's disease, Stroke, ischemic heart disease and heart failure. Such pathologies demand persistent patient behaviour and a dedicated team of physicians and caregivers to support the patient to maintain a care plan for a long period in order to bypass a relapse of the illness or subsequent deteriorations. The home rehabilitation shall be amended by a Virtual Coach that supports the patients during the rehabilitation at home providing advice for a healthy behaviour. Further, the Virtual Coach provides motivates and monitors the execution of distinct physical exercises (such as serious games) and provides suggestions for alternative actions. The project consists of technological and clinical experts. The latter come from four specialised reference hospital sites that have longstanding expertise with the relevant pathologies. This clinical expertise is used to derive the basic rehabilitation needs as well

as the foreseen in- and outpatient clinical pathways. This way, the classic borders of clinical pathways (namely their sole use within a hospital setting) are broadened as well [19].

Home rehabilitation is approved only when the patient can engage in different suggested activities and where learning and skill development is required for long-period therapy continuation. Meaningful planning, responsive feedback and practice are the key determinants of home rehabilitation to improve Quality of Life (QoL) of the patient [41]. For home rehabilitation there is a tight link between cost and quality and therefore for better cost-effectiveness as compared to the care endured in the hospital under observation [25].

A successful care rehabilitation at home that is aligned to the user's needs without constant assistance from medical staff, requires a user-friendly coaching or guidance, an assessment with real-time feedback strategy and continuous engagement. That can be achieved when the therapy or coach is acting in a personalised manner and provides assistance as required by the user's needs and personal clinical pathways [19]. Analysing, scheduling and recommendation measures then need to be supported by machine learning mechanisms to allow automation, given the medical staff's absence.

### 2.2 Goal

In this paper, the focus is on assessing the usefulness of using machine learning for computer-assisted patient rehabilitation. Therefore, the representative problem of activity recommendation is investigated, where the resulting system needs to autonomously decide when to suggest which activity to the patient, ranging from taking walks of chosen lengths to suggesting to contact friends or family. The goal in this paper is to assess machine learning advances with respect to the following challenges.

**System initialisation:** Recommendation systems have to be initialised in order to retrieve at least average quality recommendations (i. e. better than random) in the beginning. A successful system initialisation is also referred to as *warm-start*, which might be based on manually curated rules or experience from a similar, already trained system.

**System improvement:** When using supervised machine learning algorithms, the system requires training data in the form of labelled examples to generalise novel, unseen situations. Activity recommendation does, however, not fit into the standard machine learning paradigm, where such training data is easy to obtain. The system

needs to learn which activities work well for a given patient, which requires interaction.

**System safety & -trust:** Especially in health care it is highly important to guarantee system safety and -trust. For the patient rehabilitation use case system safety entails that – besides irritating system recommendations such as taking consecutive naps – the patient is never put into situations which might negatively impact his/her health. This is directly connected to ensuring that the system is accepted by both physicians and patients, which requires the building of trust. To enable the latter, the system needs to be able to reason about its decision-making process in order to explain a certain recommendation.

### 3 Machine learning for computer-assisted patient rehabilitation

A central goal of computer-assisted patient rehabilitation is to exploit all relevant information of a patient in a distinct step in a clinical pathway to recommend the optimal activity for the patient with respect to his/her long term rehabilitation goal. Therefore, first a description of the standard supervised learning setting and a discussion to what extent it is suitable for computer-assisted patient rehabilitation is provided. Since it is needed to learn the long term rehabilitation goal, one must not locally assess the utility of an activity but take into account its impact on future activities. It is argued that an adequate formalisation and learning framework for learning long term rehabilitation goals are decision processes, more specifically Markov Decision Processes, which can be solved by Reinforcement Learning. Subsequently, suitable unsupervised machine learning algorithms are addressed, which are of interest to gain insights into demographic- and clinical data of patients independent of clinical pathways as well as to explain the recommended activities. Finally, there is also a focus on exploiting structured information for machine learning, available in the form of domain ontologies for patient-related information. The latter might be able to better design the overall system and support transparency of the decision-making process.

As a central, simplified example for the paper, a single patient is used, referred to as *patient A*, as well as a pool of 4 activities, namely *walking 2 km at 5 km/h*, *walking 5 km at 4 km/h*, *running 1 km at 9 km/h*, *running 2 km at 8 km/h*. The available feature-related information comprise (a) demographic information (e. g. age, place of birth, etc) and

(b) daily measured physiological information (e. g. weight, blood pressure). Finally, the available preference-related information comprise daily information given by the patient after each activity (e. g. Borg assessment).

#### 3.1 Supervised machine learning

In supervised machine learning, access to a ground truth (also referred to as labelled data) is assumed, where for a task input the optimal output (i. e. the solution) is known. In the following, a distinction among standard learning settings and an assessment of the suitability of the latter for the introduced challenges of computer assisted patient rehabilitation are addressed.

##### 3.1.1 Learning settings

Standard problems are classification, where a set of predefined classes is given and one has to choose the correct one, or regression, where one predicts a continuous value. There is a vast amount of successful supervised machine learning algorithms for mentioned settings, such as Support Vector Machines (SVMs) [36], Random Forests [15] and recently (Deep-) Neural Networks (DNNs) [21], which could, in principal, be applied to the problem of recommending activities for patient rehabilitation.

The patient activity recommendation problem would then be modelled as standard (multi-class) classification, where an individual class refers to an individual activity and a *good* activity gets assigned label 1 (i. e. it should be chosen) and a *bad* activity label 0. Existing systems for pathway recommendation already successfully followed this approach by using SVMs and *k*-nearest-neighbours (KNN; a non-parametric supervised learning approach) [45, 44], as well as Bayesian networks (BNs) [23] to learn respective machine learning models to recommend activities within clinical pathways.

For *patient A*, all available feature-related information is modelled as features and all available activities as classes. It is sensible to model additional features, e. g. the average performance of *patient A* for individual activities in the past week, in order to find possible correlations. Assuming that labels for the feature set exist (where a physician potentially defined a mapping from preference-related information to labels for the past), a supervised classifier can be learned to recommend an activity for *patient A*.

While the individual models (i. e. SVM, KNN and BN) can deal with complex relationships among available

patient-related features, the essential problem is that – within a clinical pathway – more than one activity might be beneficial for a patient and, additionally, they have more fine-grained utilities (i. e. continuous values such as in regression). In addition, if the activity recommended to *patient A* was not good, there is no natural way to provide an update to the classifier, as it is unclear which activity might be better.

A better solution would thus be to approach the activity recommendation problem as cost-sensitive (multi-class) classification [43] problem. Here, all classes are assigned a continuous value which refers to the cost of choosing the class (where, other than in standard classification, a value close to 0 would refer to a good class). The learning setting is strongly related to standard regression, more specifically *one-against-all* regression, reducing learning a good solution for it to learning an individual regressor for each class and choosing the output of the one with the lowest cost. As a consequence, all mentioned supervised machine learning algorithms can be directly used for learning the individual regression functions.

A cost-sensitive classifier for *patient A* can be trained with the same labelled corpus of data as a standard classifier. However, when a recommended activity turned out to be bad (or if two activities turn out to be good over time), the model can be directly updated.

While cost-sensitive classification is better suited to model utilities of activities compared to standard classifiers, the individual recommendations remain independent. However, choosing a particular activity might have severe impact on the utilities of subsequent activities. Such dependencies can be modelled within decision processes, which we cover in Sec. 3.2.

### 3.1.2 Learning representations

Independent of the learning setting one usually assumes a vectorised representation as input. Here, one (a) chooses all feature-related relevant information (in cooperation with physicians) for an individual set of alternative activity recommendation (e. g. proposing *patient A* to go for a 1 km run at an average speed of 9 km/h vs. to go for a 5 km walk at an average speed of 4 km/h), (b) discusses the domains of the chosen pieces of information to model them as features and (c) stipulates how the label is to be derived (e. g. solely based on the expertise of the expert) and scored (e. g. a 1 or 0 vs. continuous values for costs).

Step (a) is very important, as it requires a common understanding of the problem. Step (b) then entails to correctly format the data, i. e. distinguishing among categor-

ical (e. g. smoker vs. non-smoker) and numerical features (e. g. heart rate) and correctly transforming them into features for the vector. Step (c) is also central for the eventual *quality* of the activity recommendation solution, as it has to be carefully examined if the expert knowledge of the physician can sufficiently generalise to all patients. To this end, there will be activities which have to be actually performed and then labelled, which renders a problem for standard supervised machine learning.

There are approaches which directly deal with other representations, such as graphs. Here, one either (a) models dependencies among prior mentioned features (within the vector; resulting in a graph of individual features) as input to a probabilistic graphical model (PGM) [35, 16], (b) conducts random walks within the graph [13] to derive node and edge statistics and uses these as input to standard supervised machine learning or (c) chooses a subset of relevant types of nodes and relationships as input to, for example, a lifted PGM [18].

To this end, it is not clear from standard supervised machine learning (independent of the presented learning settings) how appropriate labels for cost-sensitive multi-class classification should be provided, as one can only recommend a single activity to a patient at a distinct point in time. While there are different practical methods to derive the label for the activity (e. g. by modelling expert rules based on a set of (clinical-) patient information or by asking the patient to give feedback), only feedback for a single activity is provided. The resulting problem is referred to as *partial feedback* which is now introduced.

## 3.2 Reinforcement learning

Other than in standard supervised machine learning, now partial feedback is assumed, i. e. the label for all available activities when actually choosing one for a patient is not known. This assumption is well-suited for computer-assisted patient rehabilitation, as each possible activity to a patient at a single point in time cannot be recommended.

For *patient A*, a system could remember that activity *running for 5 km* was chosen before, was not successful and then choose a different one (e. g. walking), but the utility of the activity actually depends on the subsequent activities as well, and the exact situation (i. e. state; where the exact same clinical parameters are measured again) could never occur again. More specifically, the schedule of *patient A* might comprise two outdoor activities a day and he/she must not go for a run twice, where running might be preferred in the morning.

The partial feedback problem occurs for the (*Contextual-*) *Multi-Armed Bandits* problem [7], a one-step decision problem which would occur if activities would be independent of each other, and for Reinforcement Learning (RL) [39], a multi-step decision problem. Since the Bandit problem is an inexact fit for computer-assisted patient rehabilitation (as activities in a pathway have mutual relationships), we focus our study on RL-based algorithms.

### 3.2.1 Learning protocol

RL assumes an underlying decision process, which is referred to as Markov Decision Process (MDP) [30]. In a MDP, the world is modelled as states (which consist of all relevant information a decision needs to be taken; e. g. all patient data, the complete history of medical treatment and activities of a patient, and all sensor data) and actions (e. g. possible activities or dialogues with the patient) is assumed. A single action in a state can be chosen, after a novel state as well as a reward are observed. The latter is a numerical value which quantifies how *good* the action was in that state (i. e. it is the same as the reverse cost in cost-sensitive classification). The reward, however, is only of local utility, i. e. the action might be very good for a short time frame but could have detrimental effects on the future (e. g. eating a lot of ice cream when having the felt positive emotion as reward). The process is repeated over a number of episodes, which depict individual *trajectories* of the whole process (e. g. a day or week of a patient in rehabilitation, where states refer to individual times during the day). The challenge tackled by RL algorithms is now to actively choose which action to take in which state to maximise the so-called expected cumulative reward, i. e. the aggregated and averaged reward estimate over a number of episodes. The resulting function (i. e. the classifier for classification or regressor for regression) is referred to as policy, which recommends an action based on a state.

For a patient in rehabilitation, individual states refer to all relevant information for her/his well-being which might be useful for activity recommendation (e. g. the information for *patient A* plus information about activities chosen before on the same day). The action space of the MDP then refer to all possible activities with possible configurations, such as going for a 2 km walk. After an action was chosen and the respective activity has been recommended to the patient, a numerical reward value can be inferred based on diverse criteria, e. g. the Borg assessment of *patient A*. The learning task is then to strategically find

out which activities yield large cumulative reward scores for individual days, by alternating between *exploring* relatively unknown activities and *exploiting* relatively known, good activities.

### 3.2.2 Learning approaches

In recent years, the problem of RL became increasingly popular, which led to a drastic increase in published approaches with high theoretical and empirical quality. The central application has been robotics, but the increase in popularity has to be attributed to learning to play (computer-) games with superhuman abilities – prominent examples are the game of Go and the general OpenAI<sup>1</sup> platform which supports a variety of video games.

RL approaches are categorised into model-based [10], value-based [27] and policy search [37]. In model-based RL, learning the reward function to subsequently employ planning algorithms to get an estimate of the policy is aimed for. Value-based RL approaches aim to quantify the utility of an action in a state, which is less generalisable than learning the reward itself, but often converges faster to a solution. Finally, policy search approaches directly learn the policy. The latter are often employed in robotics, as the space of available actions (and resulting states) is usually extremely large, and such approaches can find fast solutions for approximate policies.

However, when using RL approaches for activity recommendation one central problem remains: a RL agent assumes to be able to randomly act in the beginning in order to gradually improve its behaviour, which is not possible in computer-assisted patient recommendation, as the patient's well-being is the highest priority. We must not, for example, risk the health of *patient A* by recommending him/her to go for two runs on the same day.

As a consequence, methods to bring the RL agent to a sufficient level of expertise are needed until it takes over the activity recommendation decision process to personalise its behaviour for a particular patient. Such techniques fall under the category of Imitation Learning [34] when dealing with decision processes. In a MDP, it is required to have a set of so-called *expert demonstrations*, which consist of optimal sequential action selections for numerous states. More specifically, physicians need to manually annotate numerous activity recommendations for different patients, which have to be work well with respect to the eventual well-being of the patient. For *patient*

<sup>1</sup> <https://openai.com/research/>

*A*, a demonstration would comprise the current state (i. e. all mentioned information which are recorded) combined with the activity chosen by a physician.

However, such large sets of recommendations are costly and cumbersome to generate, as complete trajectories of the decision process have to be modelled, i. e. each consequence of a recommended activity for the duration of the process (e. g. the goal of 5000 steps within a day). One could exploit general knowledge of physicians by modelling generally applicable rules for individual activity recommendations [22]. Whenever a rule then applies to a patient's current situation, one can choose among a predefined list of activities with known rewards. An agent can then use such rules to adapt its behaviour for a *general patient*, which can be adapted by standard RL for individual patients. Such a general rule for applicable to *patient A* could comprise a subset of the feature-related information we have, e. g. only physiological information, of a specific time during the day combined with the activity *walk for 5 km*. The rule is then applicable to other patients with the same physiological information as well, irrespective of their demographic information.

For computer-assisted patient rehabilitation, two aspects with respect to RL are distinguished. The first aspect simply deals with the fact that accurate claims about the possible performance of RL algorithms on real data without empirically evaluating them are not possible. Thus, while it makes sense to reuse State-of-the-Art RL approaches of similar problem settings, the resulting system sufficiently flexible has to be kept to possibly test different ones.

The second aspect is more important, as it deals with the general challenge of medical assistance systems and automated decision-making. The RL algorithms are adequate for the problem of learning to recommend activities, assuming that they can initially act randomly and have no constraints on choosing activities. First, it has to be ensured that all possible activities the system is allowed to choose are either completely safe (i. e. not critical in the least amount) or that the eventual decision to choose the recommended activity is taken by a medical expert (i. e. he or she is in full control). The set of safe activities for *patient A* after having been running would thus only comprise walking.

In summary, three parallel approaches for using RL for patient rehabilitation are required: (a) Defining, in close work with medical experts, which actions are allowed in which state, which actions need confirmation and which actions are to be ruled out completely. This serves as initial training for the system and can be achieved by modelling

resulting constraints based on structured models. (b) Initialising the system using imitation learning approaches based on expert rule. These rules – similar to mentioned constraints – encapsulate what a physician would generally recommend to a patient with available profile and what available outcomes of a recommendation are good or bad. The resulting rules can then be used to simulate an appropriate learning environment for an agent to warm-start, where artificial reward values are generated in order to reflect general (i. e. non-personalised) feedback. (c) Exploiting each piece of feedback (i. e. each chosen action with reward), which could be enabled by an interactive machine learning process [38]. The latter requires a domain expert to actively train the system with respect to giving it a reward estimate as well as information about useful and not useful features. While reward feedback is standard in both supervised machine learning as well as RL, feature feedback [31] is powerful. It was shown that feature feedback is able to considerably improve the performance and decrease the time until convergence of machine learning and RL models [20].

### 3.3 Unsupervised machine learning

We now discuss the suitability of unsupervised machine learning algorithms for computer-assisted patient rehabilitation. Here, the focus is on finding *general* patterns in data, which are independent of a particular application (e. g. activity recommendation).

Unsupervised machine learning is helpful to take practical decisions with respect to implementing supervised machine learning approaches. One example for such a decision is learning a single global model for all patients (which is able to recommend activities for all patients) or learning individual, specialised models for individual or groups of patients (which are similar in a certain way). Finding such groups is possible via clustering algorithms, which rely on similarity measures and employ different strategies.

Other helpful applications of unsupervised machine learning are vector embeddings, which are expressive representations for text, images or graphs. Although they are learned with supervised machine learning algorithms, they do not require human labelling and are thus considered unsupervised. One could learn such representations to improve the activity recommendation approaches.

Finally, unsupervised learning can also be used to explain available supervised learning algorithms, which helps end-users to interpret the predictions.

### 3.3.1 Clustering algorithms

In the beginning of using a machine learning-based system, one often does not have sufficient amount of labelled data. This is also the case for computer-assisted patient rehabilitation before actual patients are using the system in order to provide feedback for the recommended actions. However, even at this stage important insights into the data can be gained, which can have tremendous impact on the supervised machine learning algorithms for activity recommendation. The practical (i. e. empirical) performance of a machine learning algorithm is often dependent on multiple design choices, which include available parameters (e. g. learning rate for individual updates) but also how specialised the machine learning models should be. This directly applies to RL algorithms (and thus to the activity recommendation problem), as states could be partitioned into individual *clusters* and then learn individual RL policies for each of these. A cluster, here, could be minimised to individual patients, where the result would be extremely specialised (i. e. personalised) policies, which, however, cannot be reused that easily for novel patients (but might perform better for the addressed patient).

Such design decisions can be supported via clustering algorithms, which partition the data into different sets (disjoint or overlapping based on the given approach) based on a chosen similarity metric. For computer-assisted patient rehabilitation, such a metric would either compare selected information about patients or, in later stages when activities have been chosen for actual patients, compare states of the modeled decision process of the pathway.

For the running example, this entails to use the feature-related information of *patient A* as well as of other patients, and to then apply a clustering algorithm on the resulting dataset. The same approach can be applied to states modeled for *patient A* and the residual patients, which would additionally represent a distribution of chosen activities (as the latter are part of the state features).

Two prominent classes of clustering algorithms are link- and density-based. A well-performing link-based algorithm is K-Means [12, 24], which takes a parameter to set the number of eventual cluster and iteratively re-assigns both data points to cluster (based on their distance to the individual centroids of the clusters) as well as the centroids of the cluster themselves. k-Means converges when no data point changes its assigned cluster (and thus getting disjoint clusters). For computer-assisted patient rehabilitation, it is firstly needed to define similarity metrics on the selected information about a patient (e. g. Cosine- or Euclidean similarity which are usually employed to com-

pare vectors and/or their elements) and heuristically set the number of clusters one estimates useful.

In comparison, density-based clustering algorithms do not focus on links, but on dense areas depicting a cluster. DBScan [11], for example, takes as input the minimal amount of data points to form a cluster as well as the minimal similarity between pairs of data points such that a data point can be considered a core point, i. e. similar to a centroid in k-Means. The resulting algorithm creates possibly overlapping clusters, which have dense areas at their core. For computer-assisted patient rehabilitation, again similarity metrics have to be defined, but do not have to constrain the number of clusters. This is useful for initial exploration phases to get an early estimate of how the data is distributed. In addition, the clusters might have arbitrary shapes (other than in k-Means), which could be a better fit (depending on the patient distribution based on the chosen information).

As a consequence, both classes of clustering algorithms are useful to employ. After a number of different parameterised clustering have been generated, one has to examine them with medical experts to assess their meaningfulness. The eventual added value has to be determined in the application of recommending activities, which could yield better results when learning supervised machine learning models for chosen clusters.

### 3.3.2 Embeddings

An embedding for a word, image or element of a graph (e. g. a node) is a vector which is able to generalise the meaning of the respective modality. Taking a single word as an example, one could compare it with another word based on its characters (e. g. apple vs. orange) which would yield zero similarity. However, both words represent fruits which makes them similar to some extent. This could be captured by well-trained embeddings, which are learned on huge amounts of data (e. g. large corpora of text or images, or large knowledge graphs) for finding out in what contexts the respective modality instances occur. Resulting embeddings can also be directly used as representation for data points in mentioned clustering algorithms, as pairwise comparisons of vector embeddings can be conducted with the Cosine similarity.

There are numerous approaches to learn embeddings for text [1, 26] and images [40], usable throughout computer-assisted patient rehabilitation if textual medical reports or scans are incorporated. However, since using an

ontology to structure available patient-related data is advantageous (see Sec. 3.4), another important modality are graphs.

For graphs, one class of approaches can be depicted as translational [3, 17], which generally assume that summing up subject vectors and predicate vectors yields object vectors (which in turn is used as optimisation criterion). Resulting approaches are time and space efficient, and yield good performance (in terms of generalisation when applied to actual tasks). Another class of approaches deals with random walks [14, 42] (analogously to mentioned learning representations), where the resulting walks are used as context to describe nodes and edges. There are numerous other graph embedding approaches, which, for example, are based on tensor factorisation [29] (i. e. relational learning).

In order to learn an embedding for *patient A*, the available feature-related information (or the available states) of all available patients need to be modeled as structured representation in the form of a knowledge graph (see Sec. 3.4). After an appropriate graph embedding technique has been applied to the resulting graph, the embedding for *patient A* can be retrieved and used for a supervised- (or reinforcement-) learning task.

As common in machine learning applications, it cannot be clear which embedding approach works best, as they have to be empirically evaluated for the recommendation task.

### 3.3.3 Explainability

Lastly, unsupervised learning can also be useful to explain the actual reasoning process of a machine learning algorithm (such as a RL algorithm). Here, one tries to uncover the actual subset of features or their *simplified* relationships in order to present the end-user (i. e. the patient or the physician in charge) more than the plain prediction (i. e. the recommended activity).

Explanation components are either completely unsupervised [6] by returning the most important features by repeatedly permuting individual features to uncover correlations or partially supervised [32, 33] by learning a *simpler* model based on available or previously gathered training data. Here, either a list of features is generated [6] or more advanced relationships among individual features are tried to be disclosed [32, 33].

For computer-assisted patient rehabilitation, it is important to personalise the explainability of the activity recommendation process. More specifically, it is important to assess if the end-user is receptive to explanations, how

complex the explanations should be (in terms of the displayed feature relationships) and what kind of visualisation is preferred.

More specifically, the system could either display *patient A* a single list of which individual features have been used to recommend an activity (e. g. *run 1 km* because of features heart rate, weight and age) or a list of features with positive as well as negative influence on the decision respectively (e. g. *run 1 km* because of positive features heart rate and age, and negative feature weight).

## 3.4 Knowledge modeling for machine learning

Irrespective of the added expressivity of graph-based machine learning techniques (see Sec. 3.1.2), modelling available data as graphs has several benefits for computer-assisted patient rehabilitation, especially for system safety and -trust. We first shortly introduce the Resource Description Framework (RDF) as promising graph-based model for computer-assisted patient rehabilitation and then discuss its positive impacts on machine learning applications.

### 3.4.1 Resource description framework & ontologies

RDF<sup>2</sup> is a graph-based data model which represents information as a directed multigraph with nodes and edges with a single label, and is the central data model of the Semantic Web [2].

Nodes consist of IRIs (representing abstract things), literals (i. e. concrete data values) and blank nodes (i. e. dummy convenience nodes). A RDF graph is expressed as a set of (subject, predicate, object) triples, each interpreted as an edge labelled with predicate going from the subject node to the object node. RDF does not support any semantics on its own, other than those carried over from the XML datatype definition. Resulting RDF data can be retrieved using the SPARQL query language,<sup>3</sup> a query language, protocol and API for requesting information from RDF graphs.

By modeling available knowledge as RDF graph one is able to interlink heterogeneous information, such as sensor data or patient-related information. To this end, it is sensible and recommended to reuse available ontologies, which model relevant concepts and properties. Prominent

<sup>2</sup> <https://www.w3.org/TR/rdf-schema/>

<sup>3</sup> <https://www.w3.org/TR/rdf-sparql-query/>

examples for relevant ontologies are IoT light,<sup>4</sup> universAAL,<sup>5</sup> WBAN (Wireless Body Area Network) [28], SmartBAN<sup>6</sup> or SSN (Sensor Network Ontology).<sup>7</sup>

Compared to other graph-based models, such as Labelled Property Graphs (LPGs)<sup>8,9</sup> or Grakn.ai,<sup>10</sup> RDF combined with advanced Semantic Web technologies (e. g. for querying structured data) consists of well-documented, established technologies which offer sufficient expressivity paired with the possibility to extend and interlink the developed schema.

### 3.4.2 Learning with knowledge modelling

There are two general ways to exploit RDF-based data for machine learning. The first option is referred to as *vocabulary-based semantics*. Because you can identify relationships with URIs, you can define agreed-upon vocabularies such as schema.org.<sup>11</sup> The second option is *inference-based semantics*. It is used to express what the data inherently means, to infer facts, to conduct reasoning (e. g. in OWL<sup>12</sup>) or to check the consistency of the graph (e. g. in SHACL (Shapes Constraint Language)).<sup>13</sup>

A combination of both approaches is thus strongly encouraged. Vocabulary-based semantics provide a starting point by enabling to integrate all information to make it queryable for appropriate machine learning technologies. Inference-based semantics are useful when the problem to be tackled by machine learning technologies is challenging (such as applying RL), especially because data about a specific patient has to be gathered over the course of rehabilitation, where no information about the utility of an activity for the patient is available in the beginning.

For the example patient, *patient A*, all available feature-related information need to be modeled in RDF by reusing available concepts and properties of mentioned ontologies. This enables to ensure that, for example, when using the weight of a patient as feature, it is known how to interpret the available value (e. g. when different units are

used), as otherwise machine learning approaches might fail. In addition, the resulting knowledge graph enables to learn graph embeddings, as described in Sec. 3.3.2.

As the patient's well-being is the utmost priority, it is not directly possible to employ machine learning technologies (as their recommendations are close to random behaviour). Therefore, modelling *activity recommendation rules* in close collaboration with medical experts would be possible, which define the initial behaviour of an agent (see Sec. 3.2.2). These rules can be defined based using inference-based semantics, where the extended state the rehabilitation (consisting of clinical state and personal state) is used to find the most appropriate rule. Such rules, which usually consist of implications only, can be defined by Notation3 (N3),<sup>14</sup> an RDF-based vocabulary with appropriate classes and properties for rules, or SHACL, a constraint language for RDF.

SHACL enables to define so-called shapes on top of existing RDF graphs. For example, a patient shape can be defined and list of properties that a patient needs to have can be specified, which is similar to a database schema definition. Each property can be further constrained by a regular expression, cardinality or other more advanced means. Moreover, it is possible to define IF-THEN rules that are applied on specific parts of the graph.

For the running example, a SHACL rule expressing that certain physiological information (e. g. blood pressure smaller than 110) yield a recommendation of activity *run 1 km* could be modelled and then used to learn a baseline policy for *patient A*.

Using inference-based semantics in the form of SHACL shapes and rules is thus a powerful means to establish system safety and -trust, as one is able to control the degree the system is allowed to autonomously decide.

## 4 Summary

The stated challenges for computer-assisted patient rehabilitation with respect to the proposed machine learning-based system are now revisited in the following.

**System initialisation** deals with the warm-start problem of computer-assisted patient rehabilitation, where the system is required to perform well before interaction is possible. As an alternative, an imitation learning approach can be easily integrated into the Markov Decision Process

<sup>4</sup> <https://www.w3.org/Submission/2015/SUBM-iot-lite-20151126/>

<sup>5</sup> <https://github.com/universAAL/ontology>

<sup>6</sup> [https://portal.etsi.org/webapp/WorkProgram/Report\\_WorkItem.asp?WKI\\_ID=50078](https://portal.etsi.org/webapp/WorkProgram/Report_WorkItem.asp?WKI_ID=50078)

<sup>7</sup> <https://www.w3.org/2005/Incubator/ssn/ssnx/ssn>

<sup>8</sup> <https://neo4j.com/developer/graph-database/>

<sup>9</sup> <https://neo4j.com/>

<sup>10</sup> <https://dev.grakn.ai/docs/general/quickstart>

<sup>11</sup> <https://schema.org/>

<sup>12</sup> <https://www.w3.org/OWL>

<sup>13</sup> <https://www.w3.org/TR/shacl/>

<sup>14</sup> <https://www.w3.org/TeamSubmission/n3/>

formalisation and can be implemented by modelling simple expert rules. The latter can be modelled with Semantic Web technologies, such as SHACL or Notation3.

**System improvement** deals with the general challenge of designing a system which can incrementally improve its recommendation capabilities by interaction. At its core, a formalisation as Markov Decision Process is proposed, where one could use Reinforcement Learning to use actual patient feedback to adjust the learned machine learning model in order to better represent the patient's needs.

**System safety & trust** deals with ensuring that the machine learning system is sufficiently constrained, such that the patient is never at risk of being irritated or harmed. By incorporating structured knowledge formalised in a domain ontology (which could be modelled in RDF), it is possible to construct baseline rules for recommending activities, which ensure that the machine learning-based system has only limited autonomy. Finally, explanation-based machine learning can support the transparency of the recommended activities by visualising the used subset of information as well as available correlations.

The defined challenges can thus be met with appropriate advances in machine learning as well as knowledge modelling.

Furthermore, there are several challenges for research projects that aim to perform an innovation uptake. Especially the obstacles how projects usually fail to achieve an actual exploitation need to be distinctly considered. As a technical solution within a medical use case is aimed for, one of the main obstacles (and opportunity at the same time) is the **Medical Device Regulation (MDR)**. This is expected to have a severe impact to all research initiatives and start-ups that aim to provide services by interacting with patients. It has attracted a lot of attention from the industry side and highlighted the need to consider seriously the MDR, which has already been taken into force and currently there is a grace period until 2020. Solutions interacting with patients can easily fall under the MDR regulation and start-ups that are not compliant will be 'out of the game'. On the contrary, start-ups that will consider a different positioning and decide to enter into the process of Medical Device (including CE marking), will have a significant advantage over the others. Within experts' workshop there already have been useful contribution how to address such issues in the upcoming project phases of the project. These suggestion contributed to the project's exploitation strategy. In particular, strategic decisions will be taken to balance the risk of falling under the strict clauses of the MDR (e. g.

by offering a consumer-based lifestyle product), while at the same time to consider the opportunity from the MDR. This opportunity will require to start the processes of becoming a medical device (e. g. Class IIa) in the near future, thus having a significant advantage over the competition.

## 5 Conclusion

In this paper, a survey of different machine learning-based settings with respect to their suitability for the computer-assisted patient rehabilitation has been conducted. Standard supervised machine learning approaches neglect the fact that there might be multiple correct activities (i. e. the classes) for a patient at a particular point in time, that there is no feedback for all activities if it is needed to perform them (i. e. when medical experts cannot be completely sure about the consequences) and that the utility of an activity is dependent on future activities. A learning solution is cost-sensitive classification, but pairing it with an appropriate procedure to choose activities over time is needed, which is the RL problem. It is suitable by modelling all relevant information for a decision as state, all activities as actions and defining the rewards (or labels, i. e. the quantitative assessment of an activity). Unsupervised machine learning methods are applicable to better generalise within RL, to perform clusterings to find patterns among patients or to explain the recommended activities to both patients and physicians. Finally, ontology-based machine learning enables to better control the system by sustainably constraining the available activities per time step.

The running example for *patient A* assumed a simplified pathway with a small set of patient-related information and available activities, but illustrated (i) how the different supervised machine learning settings differ in their use of available information and capability of exploration, (ii) how unsupervised machine learning approaches are applied to available patient-related information and (iii) how a graph-based knowledge representation and structured rules are able to support machine learning approaches.

**Funding:** This work was partially supported by the European Union within the vCare project (funded by Horizon 2020 research and innovation programme under grant agreement No. 769807).

## References

- Bengio, Y., Ducharme, R., Vincent, P., and Janvin, C. A neural probabilistic language model. *Journal of Machine Learning Research* 3 (2003), 1137–1155.
- Berners-Lee, T., Hendler, J., Lassila, O., et al. The semantic web. *Scientific American* 284, 5 (2001), 28–37.
- Bordes, A., Usunier, N., García-Durán, A., Weston, J., and Yakhnenko, O. Translating embeddings for modeling multi-relational data. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*. (2013), pp. 2787–2795.
- Chaiyawat, P., Kulkantrakorn, K., and Sritipsukho, P. Effectiveness of home rehabilitation for ischemic stroke. *Neurology international* 1, 1 (2009).
- Chaudhry, B., Wang, J., Wu, S., Maglione, M., Mojica, W., Roth, E., Morton, S. C., and Shekelle, P. G. Systematic review: impact of health information technology on quality, efficiency, and costs of medical care. *Annals of internal medicine* 144, 10 (2006), 742–752.
- Chen, J., Song, L., Wainwright, M. J., and Jordan, M. I. Learning to explain: An information-theoretic perspective on model interpretation. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018* (2018), pp. 882–891.
- Chu, W., Li, L., Reyzin, L., and Schapire, R. E. Contextual bandits with linear payoff functions. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011* (2011), pp. 208–214.
- Clifford, J. The UN disability convention and its impact on European equality law. *The Equal Rights Review* 6 (2011), 11–25.
- Curry, N., and Ham, C. Clinical and service integration: the route to improved outcomes. *The King's Fund* (2020).
- Doll, B. B., Simon, D. A., and Daw, N. D. The ubiquity of model-based reinforcement learning. *Current opinion in neurobiology* 22, 6 (2012), 1075–1081.
- Ester, M., Kriegel, H., Sander, J., and Xu, X. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining (KDD-96), Portland, Oregon, USA* (1996), pp. 226–231.
- Forgey, E. Cluster analysis of multivariate data: Efficiency vs. interpretability of classification. *Biometrics* 21, 3 (1965), 768–769.
- Gilmer, J., Schoenholz, S. S., Riley, P. F., Vinyals, O., and Dahl, G. E. Neural message passing for quantum chemistry. In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017* (2017), pp. 1263–1272.
- Grover, A., and Leskovec, J. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016* (2016), pp. 855–864.
- Ho, T. K., Random decision forests. In *Third International Conference on Document Analysis and Recognition, ICDAR 1995, August 14–15, 1995, Montreal, Canada. Volume 1* (1995), pp. 278–282.
- Jensen, F. V., et al. *An introduction to Bayesian networks, vol. 210*. UCL press London, 1996.
- Ji, G., He, S., Xu, L., Liu, K., and Zhao, J. Knowledge graph embedding via dynamic mapping matrix. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing of the Asian Federation of Natural Language Processing, ACL 2015, July 26–31, 2015, Beijing, China, Volume 1: Long Papers* (2015), pp. 687–696.
- Kimmig, A., Mihalkova, L., and Getoor, L. Lifted graphical models: a survey. *Machine Learning* 99, 1 (2015), 1–45.
- Kinsman, L., Rotter, T., James, E., Snow, P., and Willis, J. What is a clinical pathway? development of a definition to inform the debate. *BMC medicine* 8, 1 (2010), 31.
- Krening, S., Harrison, B., Feigh, K. M., Jr., Isbell, C. L., Riedel, M., Thomaz, A. Learning from explanations using sentiment and advice in RL. *IEEE Trans. Cognitive and Developmental Systems* 9, 1 (2017), 44–55.
- LeCun, Y., Bengio, Y., and Hinton, G. E. Deep learning. *Nature* 521, 7553 (2015), 436–444.
- Liu, E. Z., Guu, K., Pasapat, P., Shi, T., and Liang, P. Reinforcement learning on web interfaces using workflow-guided exploration. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30–May 3, 2018, Conference Track Proceedings* (2018).
- Liu, R., Srinivasan, R. V., Zolfaghar, K., Chin, S., Roy, S. B., Hasan, A., and Hazel, D. Pathway-finder: An interactive recommender system for supporting personalized care pathways. In *2014 IEEE International Conference on Data Mining Workshops, ICDM Workshops 2014, Shenzhen, China, December 14, 2014* (2014), pp. 1219–1222.
- Lloyd, S. P. Least squares quantization in PCM. *IEEE Trans. Information Theory* 28, 2 (1982), 129–136.
- Mansell, J., Knapp, M., Beadle-Brown, J., and Beecham, J. *Deinstitutionalisation and community living—outcomes and costs: report of a European Study. Volume 2: Main Report*. University of Kent, 2007.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5–8, 2013, Lake Tahoe, Nevada, United States*. (2013), pp. 3111–3119.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M. A., Fidjeland, A., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., and Hassabis, D. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529–533.
- Nachabe, L., Girod-Genet, M., and El Hassan, B. Unified data model for wireless sensor network. *IEEE Sensors Journal* 15, 7 (2015), 3657–3667.
- Nickel, M., Tresp, V., and Kriegel, H. A three-way model for collective learning on multi-relational data. In *Proceedings of the 28th International Conference on Machine Learning, ICML*

- 2011, Bellevue, Washington, USA, June 28–July 2, 2011 (2011), pp. 809–816.
30. Puterman, M. L. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
  31. Raghavan, H., Madani, O., and Jones, R. Active learning with feedback on features and instances. *Journal of Machine Learning Research* 7 (2006), 1655–1686.
  32. Ribeiro, M. T., Singh, S., and Guestrin, C. “Why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13–17, 2016* (2016), pp. 1135–1144.
  33. Ribeiro, M. T., Singh, S., and Guestrin, C. Anchors: High-precision model-agnostic explanations. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th Innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2–7, 2018* (2018), pp. 1527–1535.
  34. Ross, S., and Bagnell, D. Efficient reductions for imitation learning. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13–15, 2010* (2010), pp. 661–668.
  35. Rozanov, Y. A. Markov random fields. In *Markov Random Fields*. Springer (1982), pp. 55–102.
  36. Schölkopf, B., and Smola, A. J. *Learning with Kernels: support vector machines, regularization, optimization, and beyond*. Adaptive computation and machine learning series. MIT Press, 2002.
  37. Schulman, J., Levine, S., Abbeel, P., Jordan, M. I., and Moritz, P. Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6–11 July 2015* (2015), pp. 1889–1897.
  38. Stumpf, S., Rajaram, V., Li, L., Wong, W., Burnett, M. M., Dietterich, T. G., Sullivan, E., and Herlocker, J. L. Interacting meaningfully with machine learning systems: Three experiments. *Int. J. Hum.-Comput. Stud.* 67, 8 (2009), 639–662.
  39. Sutton, R. S., and Barto, A. G. *Reinforcement learning: An introduction*. MIT press, 2018.
  40. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S. E., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7–12, 2015* (2015), pp. 1–9.
  41. World Health Organization, et al. World report on disability 2011.
  42. Yanardag, P., and Vishwanathan, S. V. N. Deep graph kernels. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10–13, 2015* (2015), pp. 1365–1374.
  43. Zadrozny, B., Langford, J., and Abe, N. Cost-sensitive learning by cost-proportionate example weighting. In *Proceedings of the 3rd IEEE International Conference on Data Mining (ICDM 2003), 19–22 December 2003, Melbourne, Florida, USA* (2003), p. 435.
  44. Zhu, M., Cheng, L., Armstrong, J. J., Poss, J. W., Hirdes, J. P., and Stolee, P. Using machine learning to plan rehabilitation for home care clients: Beyond “black-box” predictions. In *Machine Learning in Healthcare Informatics*. 2014, pp. 181–207.
  45. Zhu, M., Zhang, Z., Hirdes, J. P., and Stolee, P. Using machine learning algorithms to guide rehabilitation planning for home care clients. *BMC Med. Inf. & Decision Making* 7 (2007), 41.

## Bionotes

### Dr. Patrick Philipp

Forschungszentrum Informatik, Information Process Engineering, D-76131 Karlsruhe, Germany  
[philipp@fzi.de](mailto:philipp@fzi.de)

Dr. Patrick Philipp studied Wirtschaftsinformatik at the University of Mannheim (M. Sc. 2013) and received his doctoral degree from the Karlsruhe Institute of Technology (KIT) in 2018. As of 2018, he works as project leader at Forschungszentrum Informatik (FZI) in Karlsruhe, Germany.

### Dipl.-Inform. (FH) Nicole Merkle

Forschungszentrum Informatik, Information Process Engineering, D-76131 Karlsruhe, Germany  
[merkle@fzi.de](mailto:merkle@fzi.de)

Dipl.-Inform. (FH) Nicole Merkle studied Informatik at the Hochschule Karlsruhe (diploma 2013) and is a research assistant at the Forschungszentrum Informatik (FZI) in Karlsruhe, Germany.

### Dr. Kai Gand

TU Dresden, Chair of Wirtschaftsinformatik, esp. Systems Development, D-01062 Dresden, Germany  
[kai.gand@tu-dresden.de](mailto:kai.gand@tu-dresden.de)

Dr. Kai Gand studied Wirtschaftsinformatik at the Technische Universität Dresden (diploma 2011, doctorate 2019). Since 2012, he is a member of the scientific staff of Prof. Dr. Werner Esswein, giving lectures on programming, databases and agile project management. As of 2017, he is project manager of the EU-funded project vCare.

### Dipl.-Wirt.-Ing. Carola Gißke

TU Dresden, Chair of Wirtschaftsinformatik, esp. Systems Development, D-01062 Dresden, Germany  
[carola.gisske@tu-dresden.de](mailto:carola.gisske@tu-dresden.de)

Dipl.-Wirt.-Ing. Carola Gißke studied Industrial Engineering at the Technische Universität Dresden (diploma 2018). Since 2018, she is a member of the scientific staff of Prof. Dr. Werner Esswein.