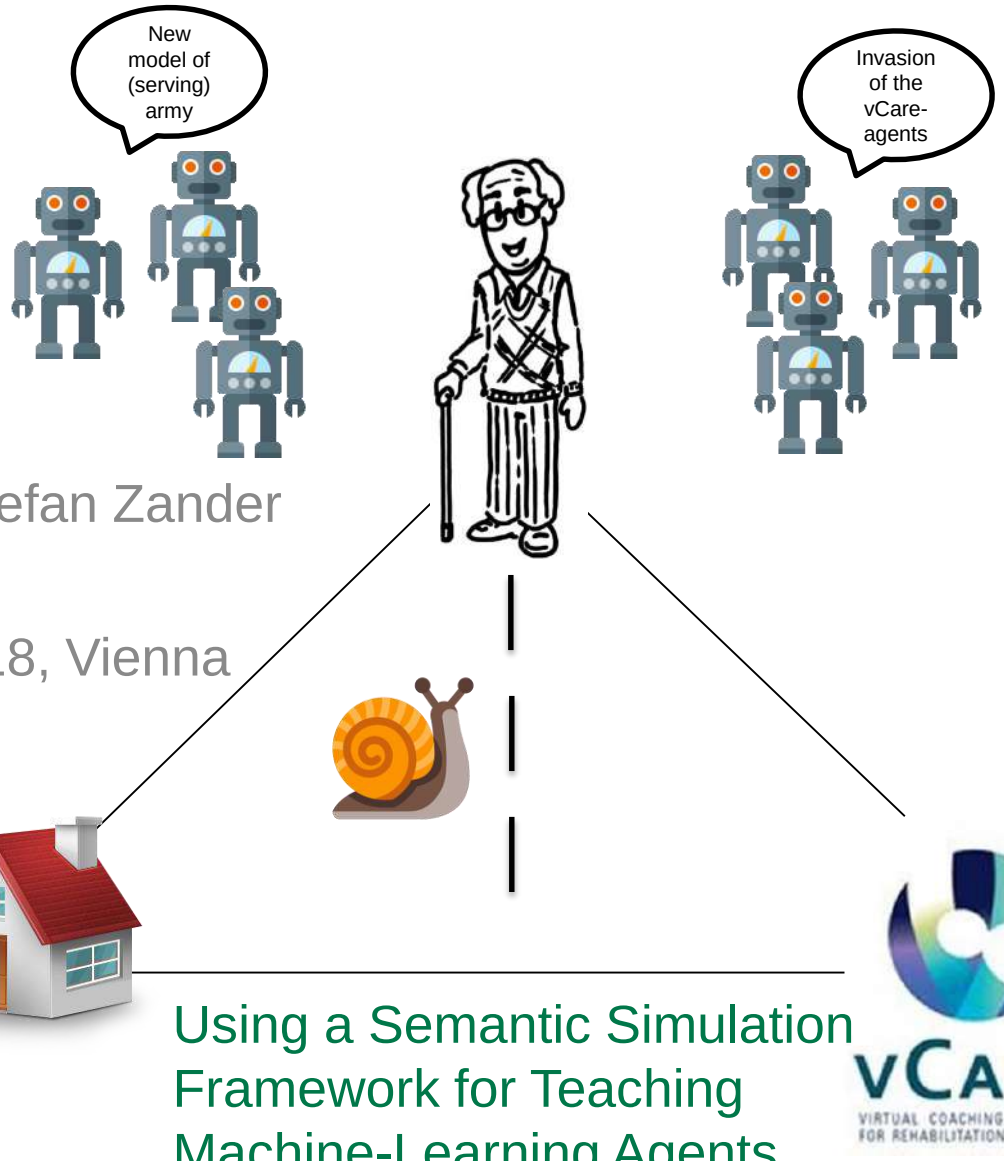

Nicole Merkle, Stefan Zander

12.09.2018

SEMANTICS 2018, Vienna

Using a Semantic Simulation Framework for Teaching Machine-Learning Agents

Motivation

General Idea concerning Machine Learning Tasks

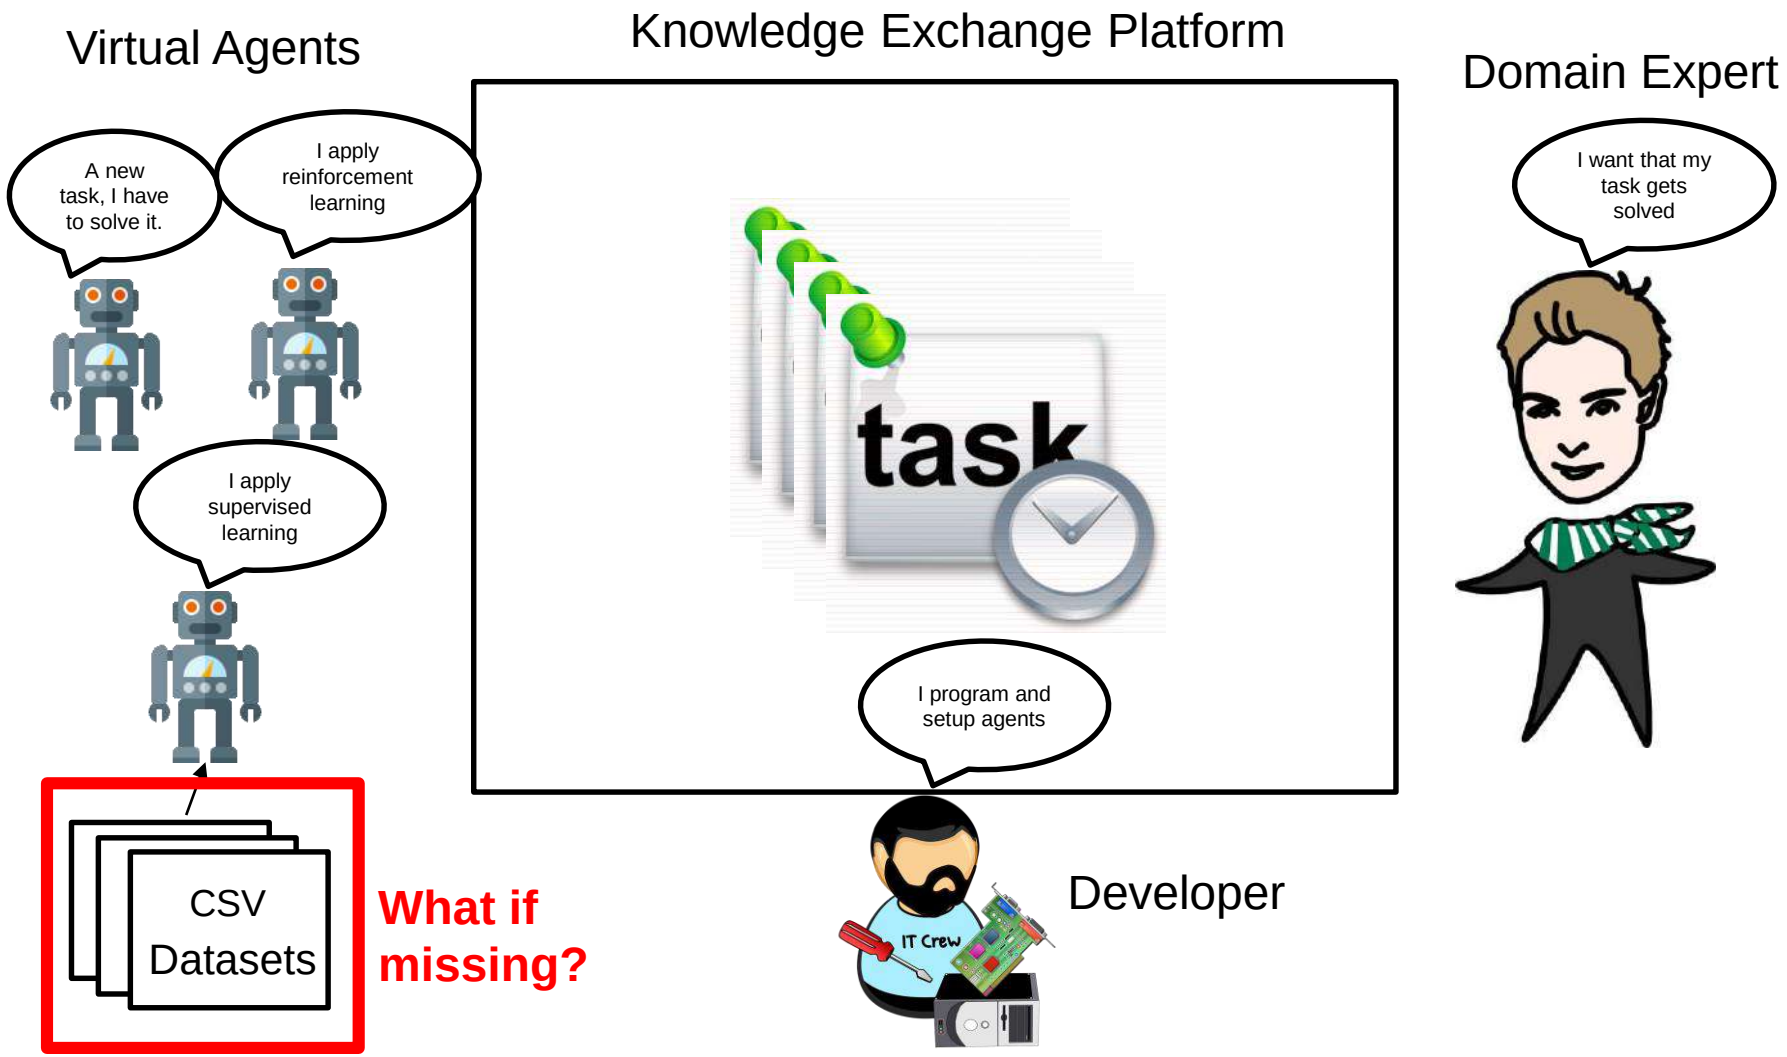

Use Case in the medical domain: Virtual Coaching at Home

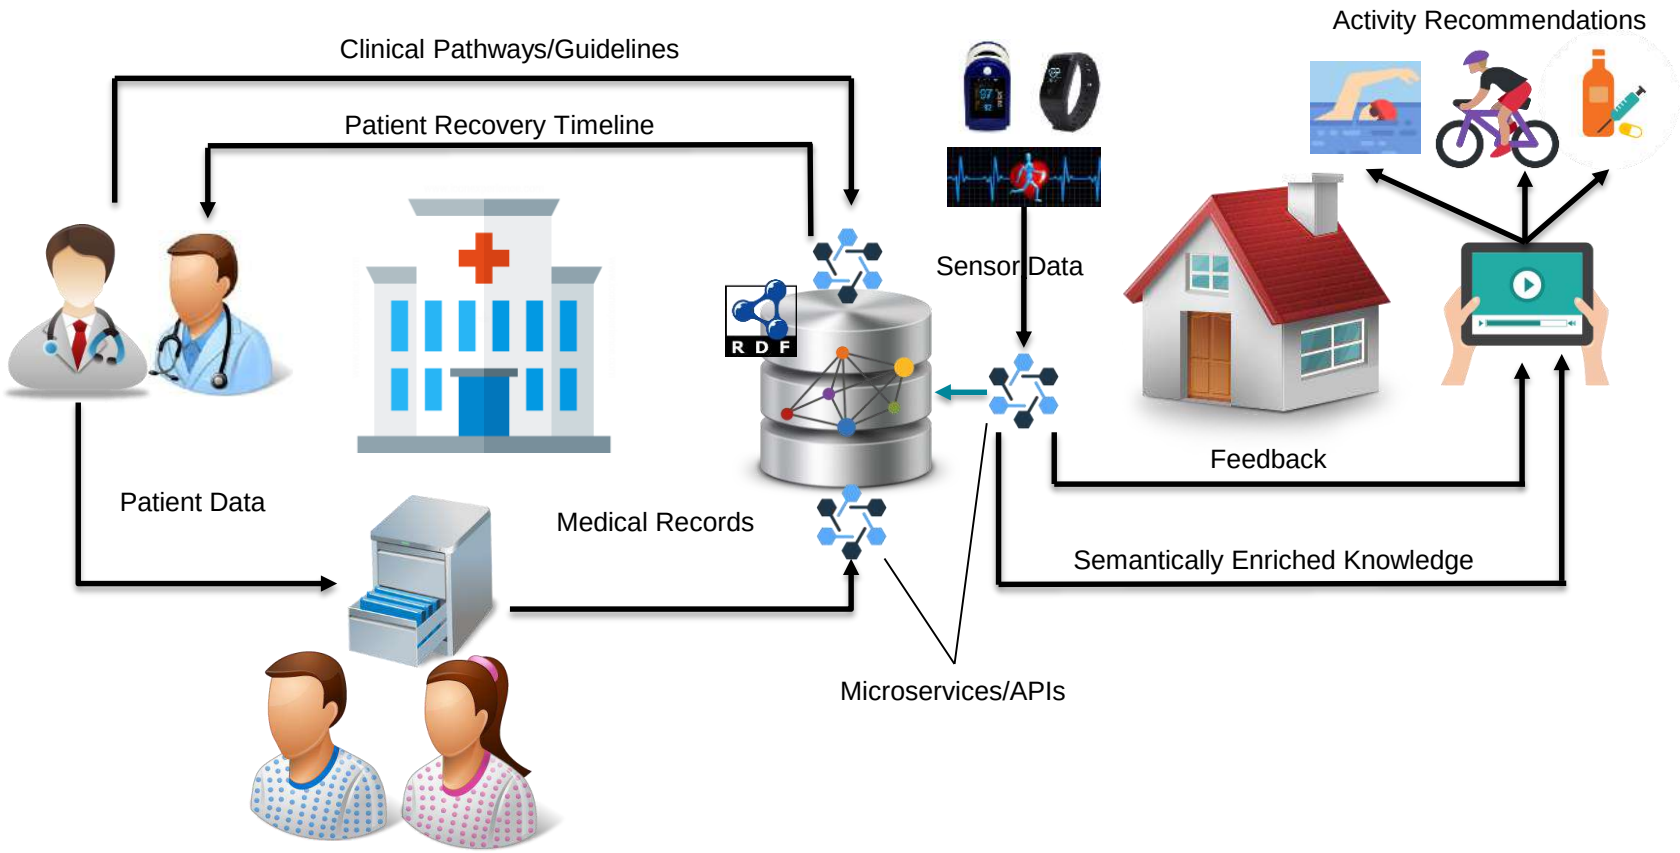

Example Clinical Pathway: Chronic Kidney Disease

Observed vital parameters

Chronic Kidney Disease (CKD) Clinical Pathway

Testing should be targeted for individuals at increased risk of developing CKD.

Who is at Risk

- Hypertension
- Diabetes Mellitus
- Family history of Stage 5 CKD or hereditary kidney disease
- Vascular disease (prior diagnosis of CVD, stroke/TIA or PVD)
- Multisystem disease with potential kidney involvement (e.g. SLE)

Who and How to Test

Decision Points Guideline rules

Patient Profiles

Recommended Activities

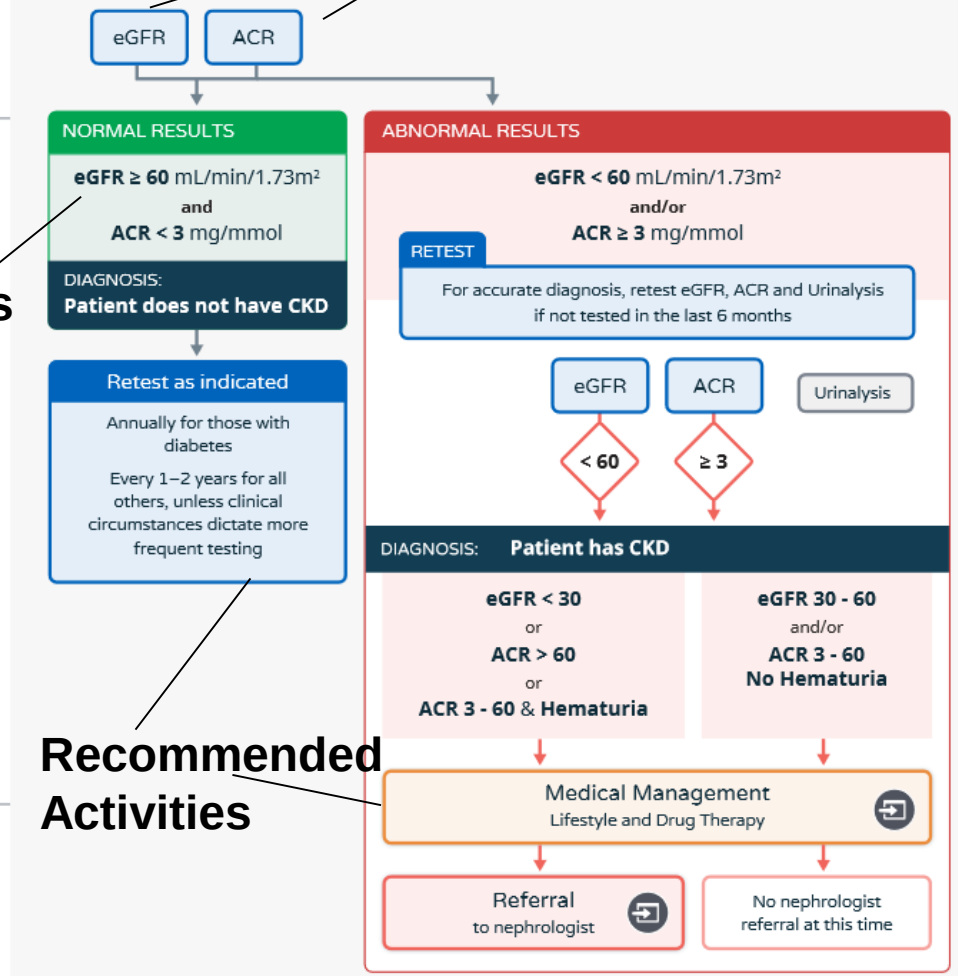

The Chronic Kidney Disease (CKD) Clinical Pathway is a resource for primary care providers to aid in the diagnosis, medical management, and referral of adults with CKD.

www.CKDPathway.ca

Research Questions

- **RQ1: Does the meta-model based simulator provide comprehensive and realistic state representations to virtual agents?**
- **RQ2: Can a reinforcement learning agent-trained by the simulator-achieve the highest expected rewards by making trained action recommendations?**

Approach

What is Reinforcement Learning (RL)?

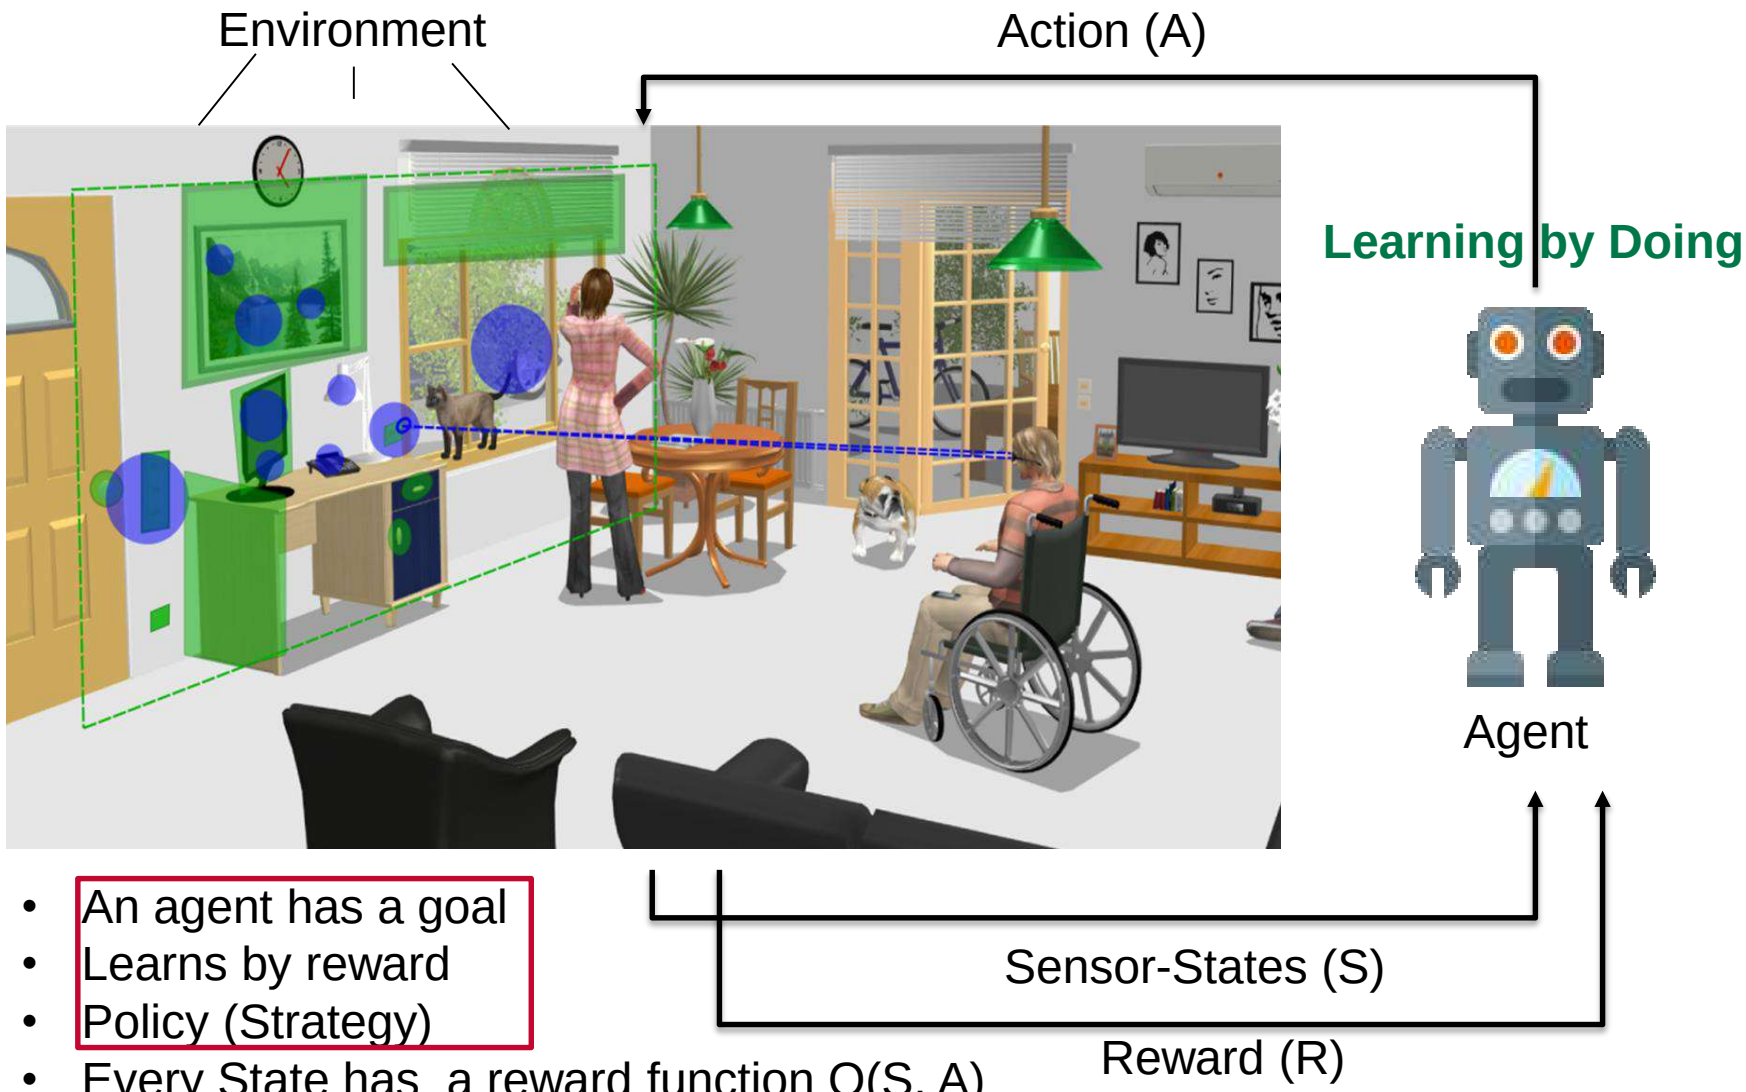

- An agent has a goal
- Learns by reward
- Policy (Strategy)
- Every State has a reward function Q(S, A)

An agent with a goal

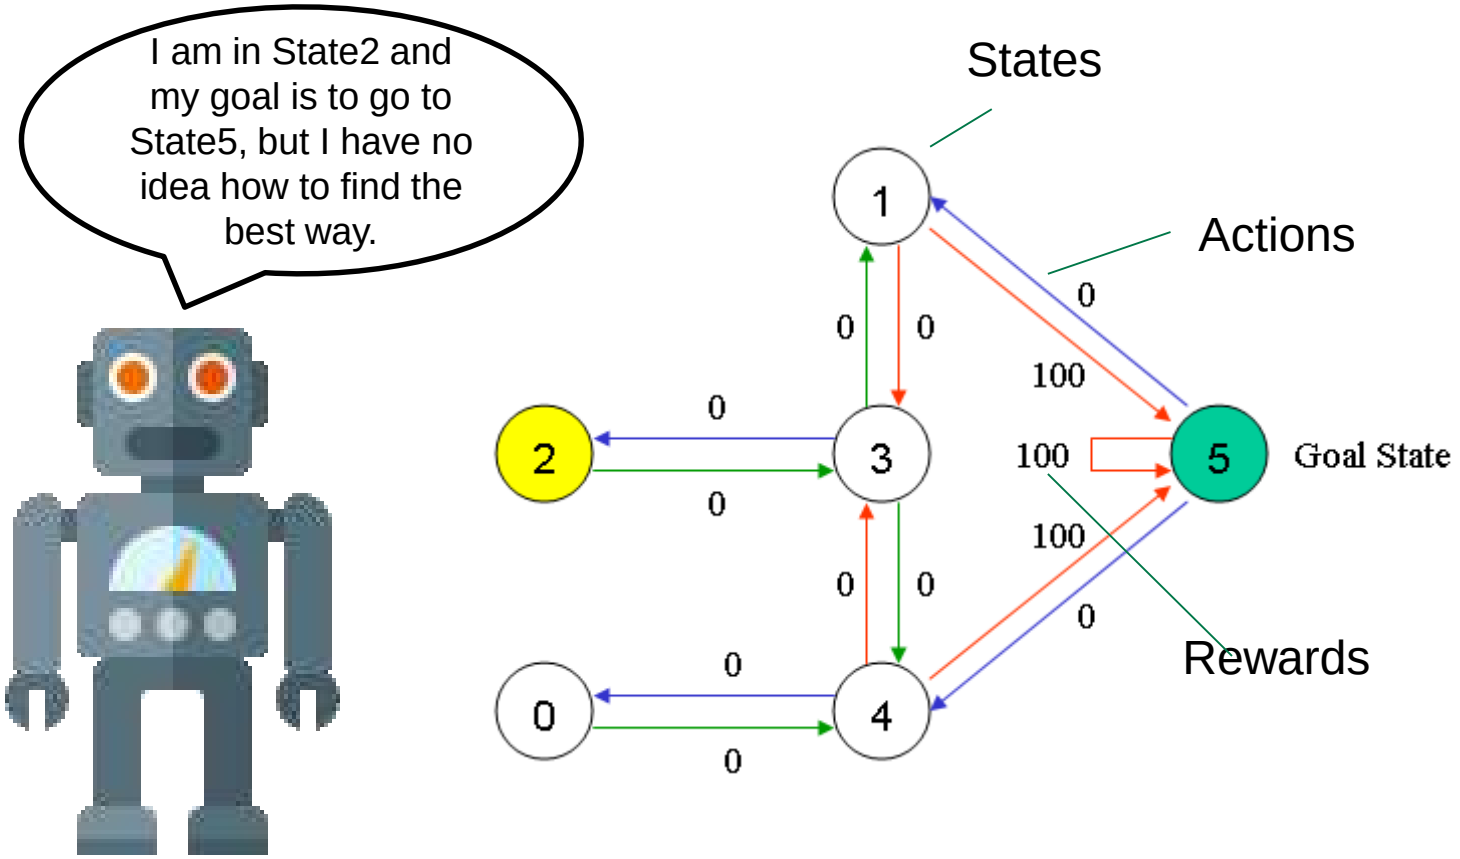

How can the simulator train the agent by simulating states and rewards?

Deep-Q-Network

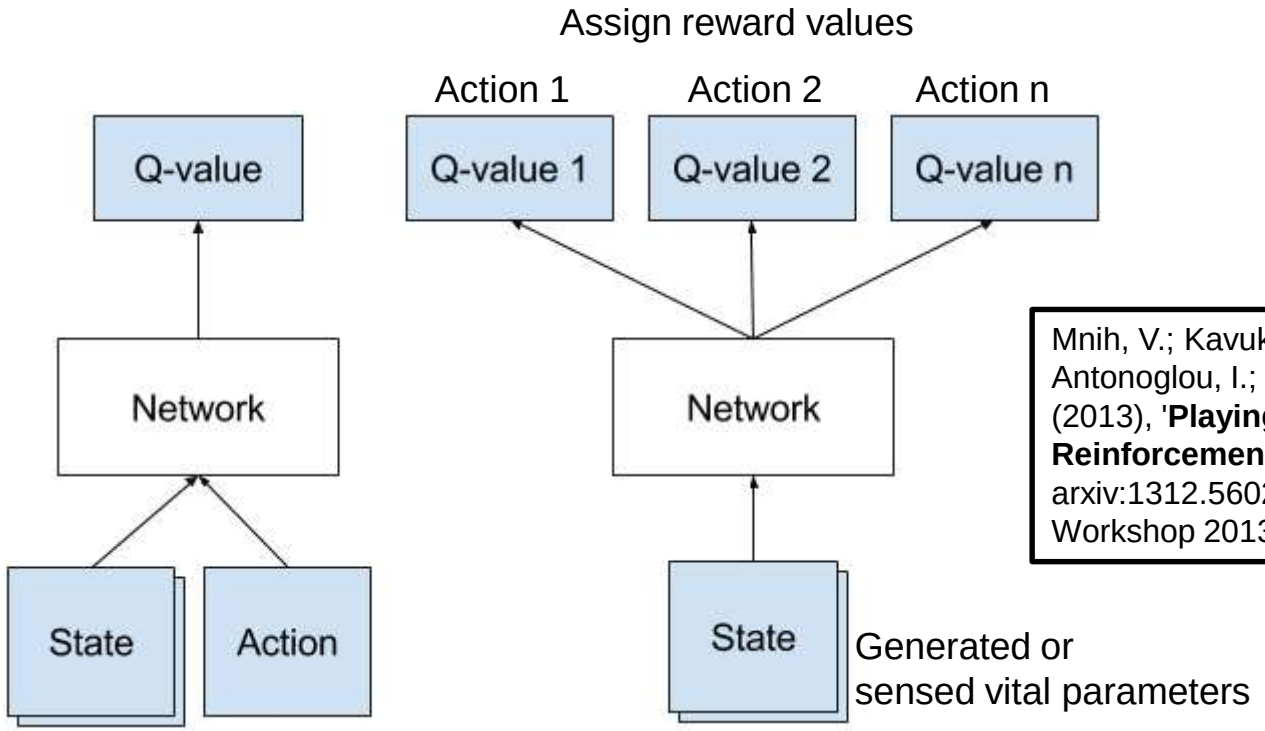

Mnih, V.; Kavukcuoglu, K.; Silver, D.; Graves, A.; Antonoglou, I.; Wierstra, D. & Riedmiller, M. (2013), '**Playing Atari with Deep Reinforcement Learning**' , cite arxiv:1312.5602Comment: NIPS Deep Learning Workshop 2013 .

$$loss = (r + \gamma \max Q(s, a') - Q(s, a))^2$$

Target Prediction

Virtual Agent Platform

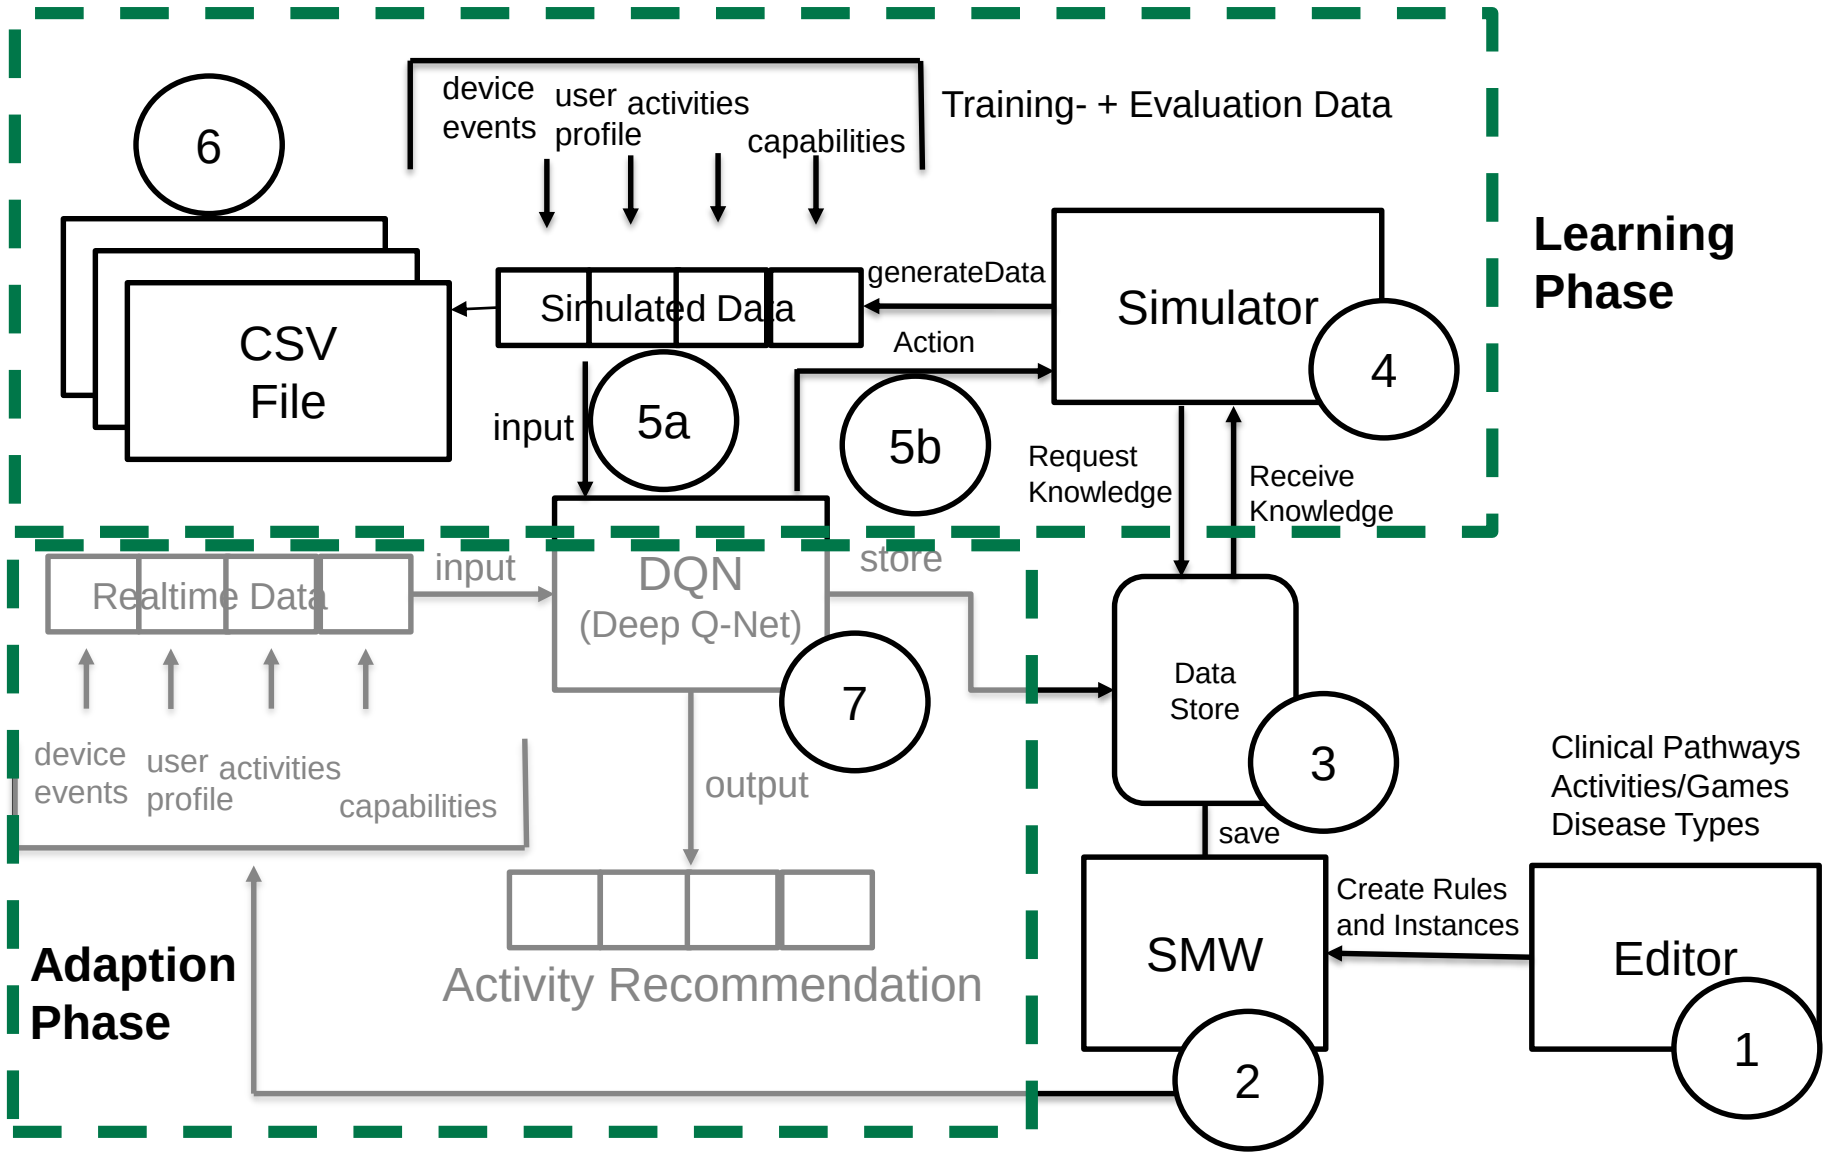

The System Architecture

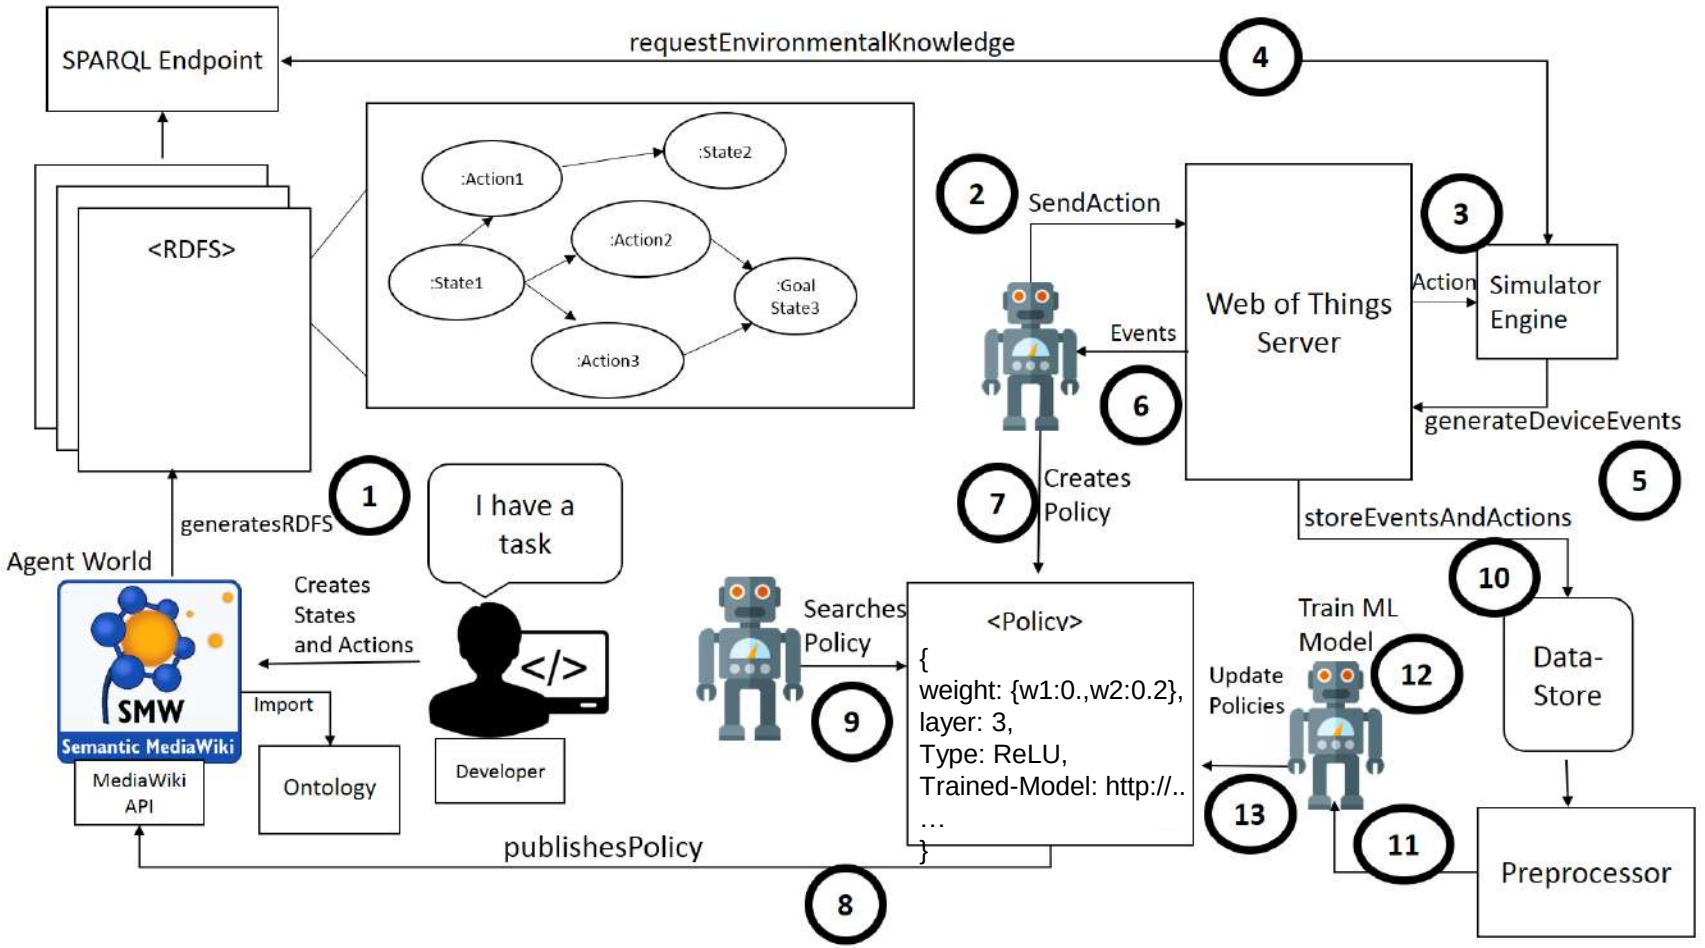

Simulation Process

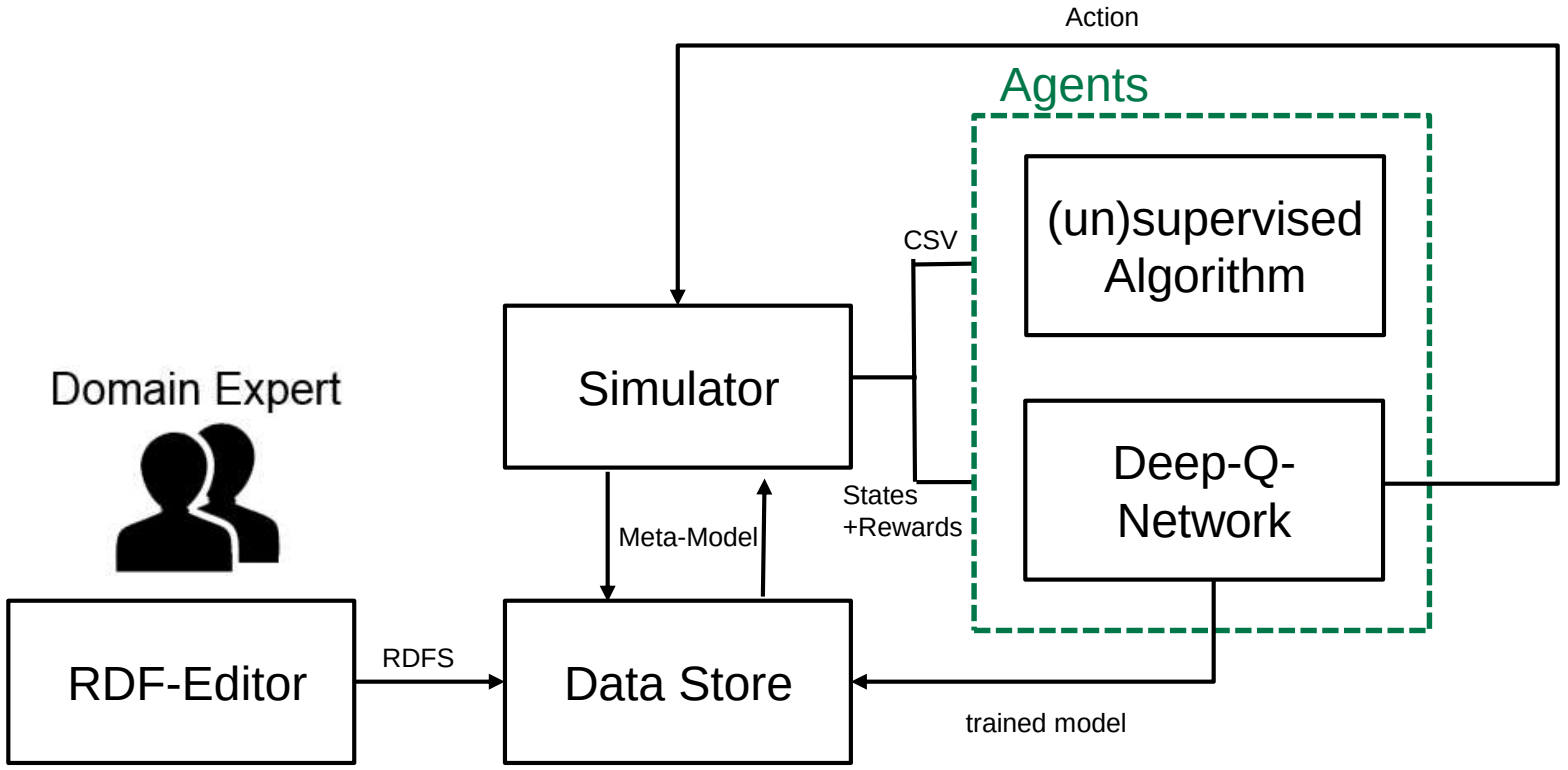

States, Actions and their Effects to Observations

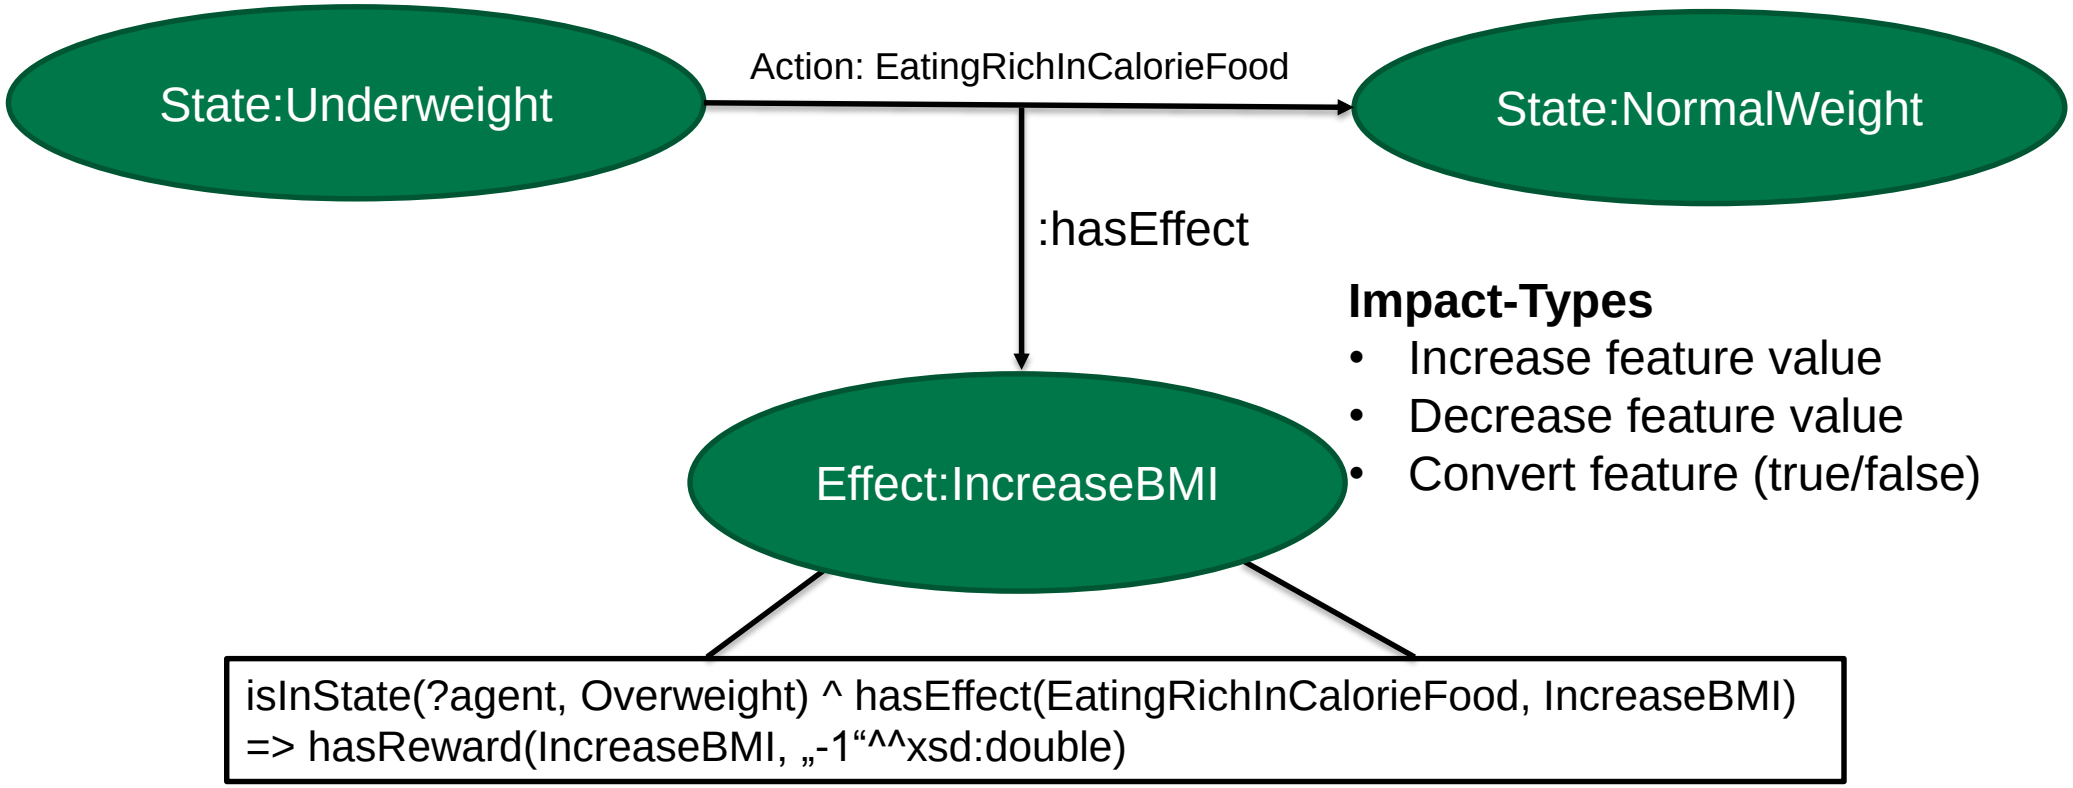

Update of state vectors by activities

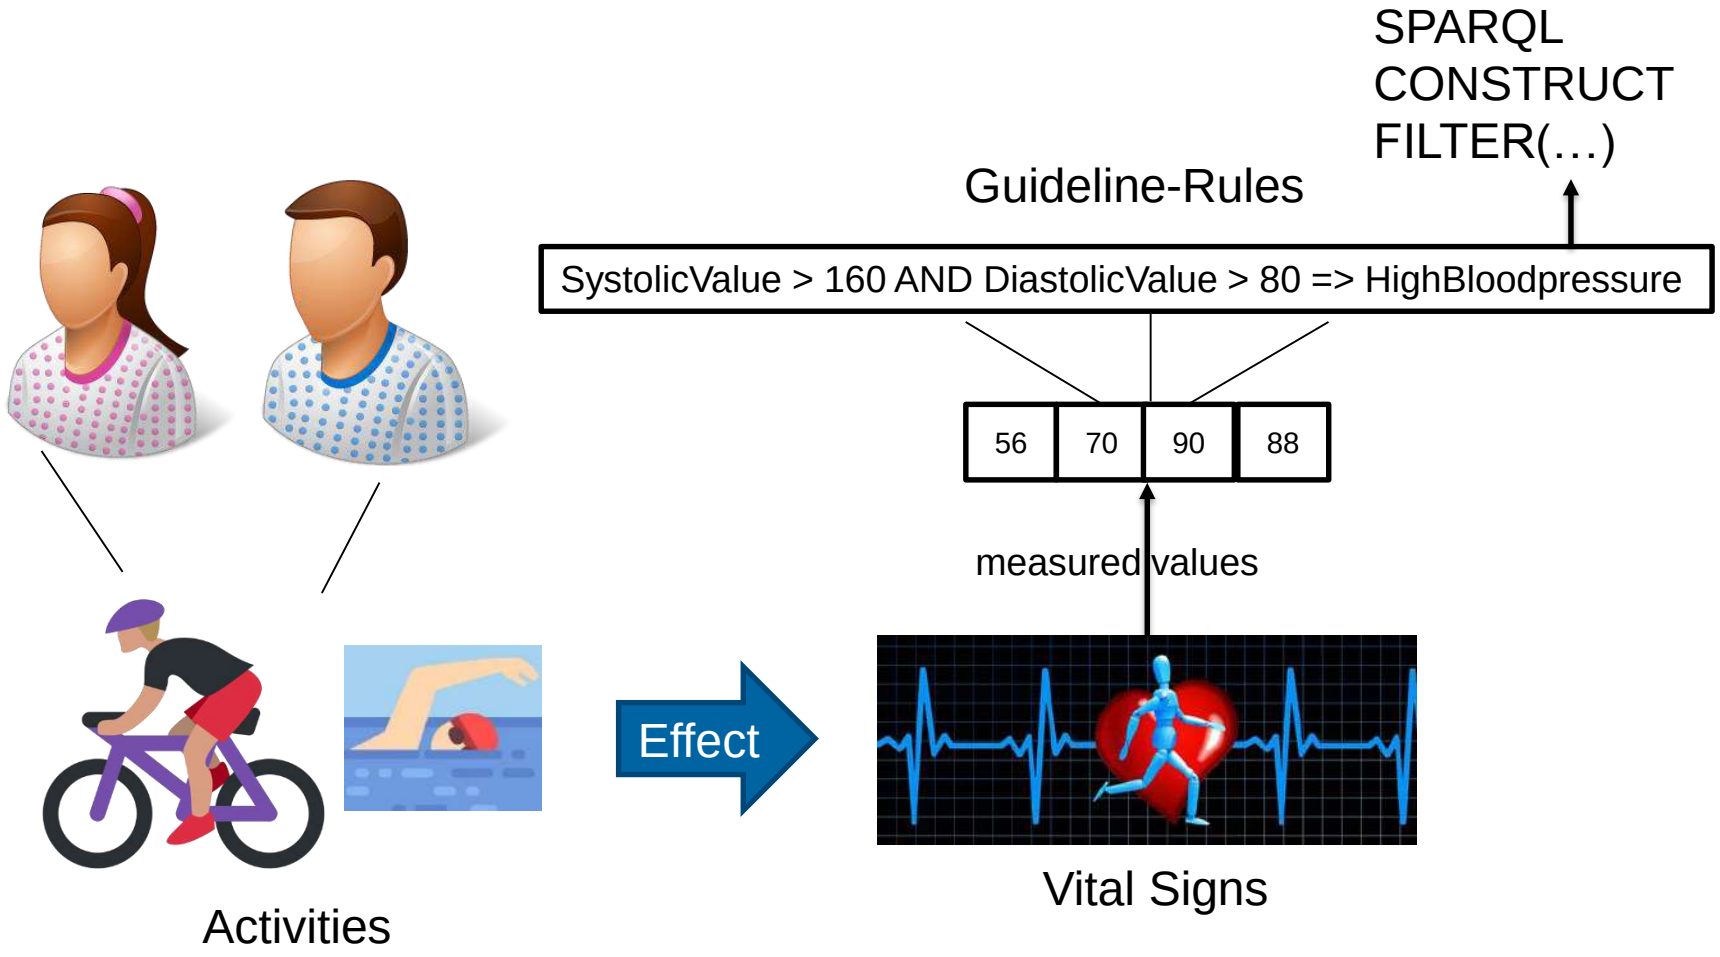

Example of a Numerical State Representation in the CKD Pathway (Generated by a simulator)

Nominal Features „one-hot-encoded“

Numerical Features

Patient-Profile

Sensed Values

Hypertension	Diabetes	Hereditary CKD	Family History Stage 5 CKD	Vascular Disease	Multisystem Disease	Hematuria	Diastolic Value	Systolic Value	BMI	eGFR	ACR
1	1	1	1	0	1	1	90	160	26.2	30	4.3

Observation Features

State: UserHasCKD

Rule: eGFR > 60 AND ACR < 4

Reasoning of States and Rewards

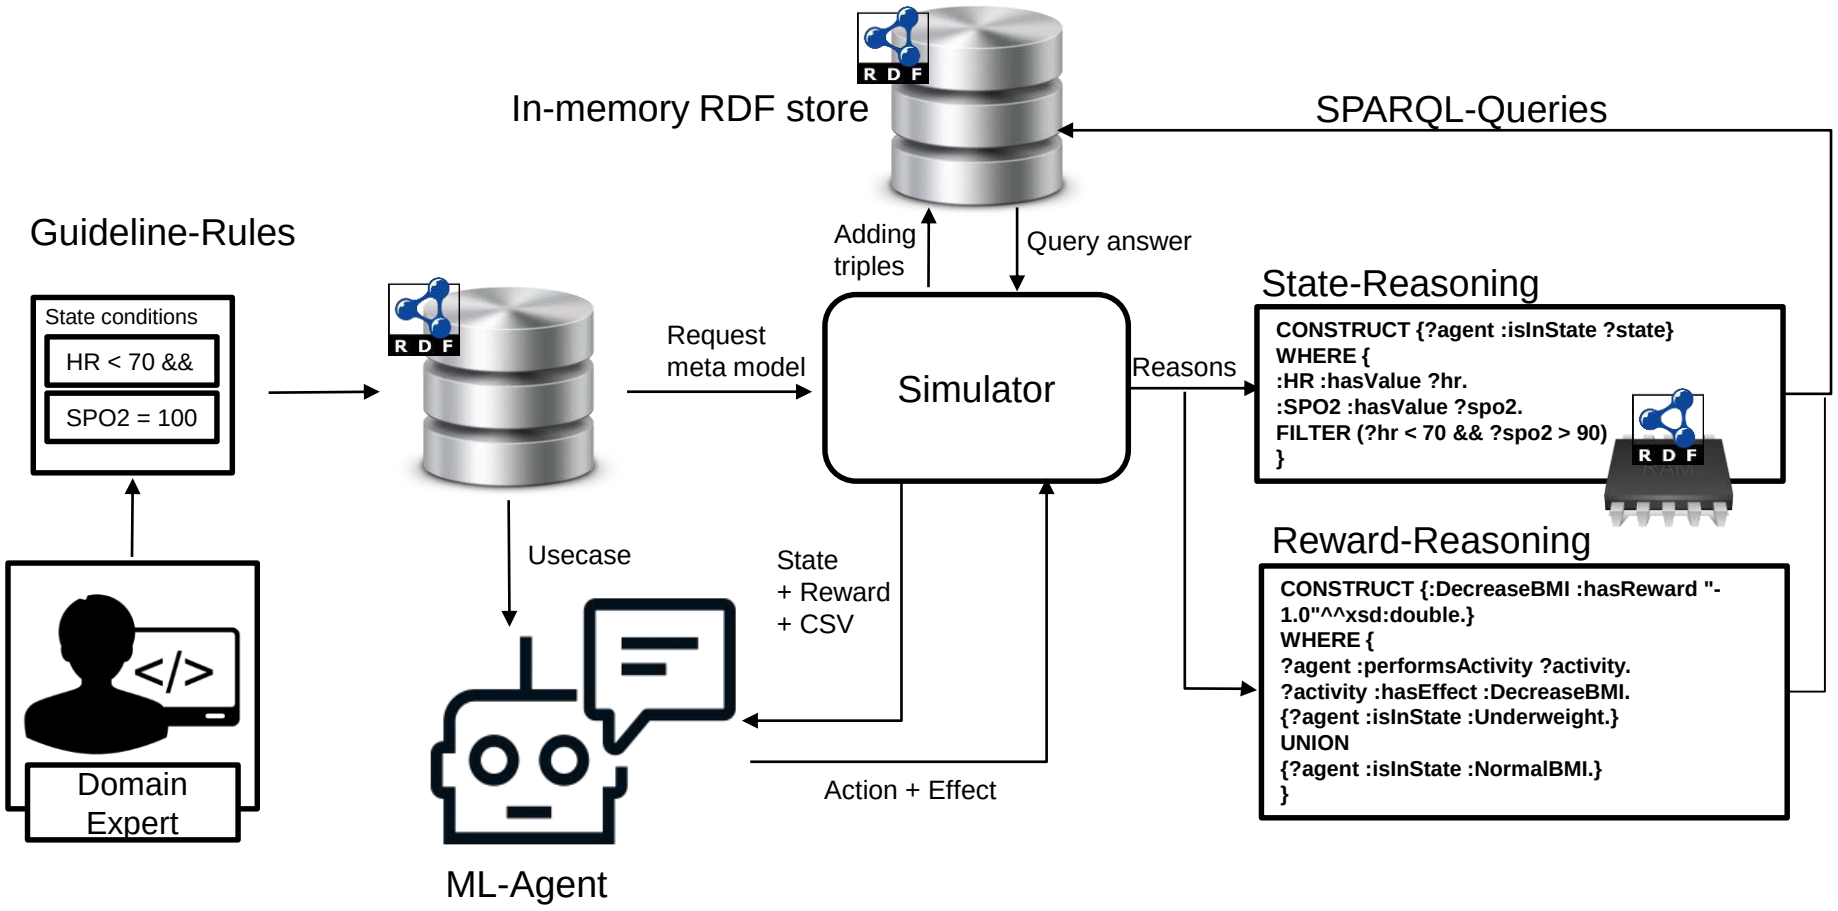

Simulation Ontology Classes (excerpt)

- **Task:** Task for supporting the patient (e.g. clinical pathway, walking cane, smoking cessation, diets, etc.)
- **State:** relevant states for assessing activities.
- **Observation Feature:** sensable parameters (e.g. vital signs)
- **Condition:** prescribes state conditions.
- **Action:** can be performed by patient.
- **Effect:** has an impact to observation features.
- **Agent Profile:** Configuration of the agent (training parameters, e.g. learning rate, discount factor).
- **Virtual Influencer:** represents virtual sensor that changes dynamically varying observation features. Provides some uncertainty in the environment.
- **IoT Device:** sensing and acting devices of the environment.
- **Policy:** the trained model => f(state) = action

Simulation Meta-Model (Excerpt)

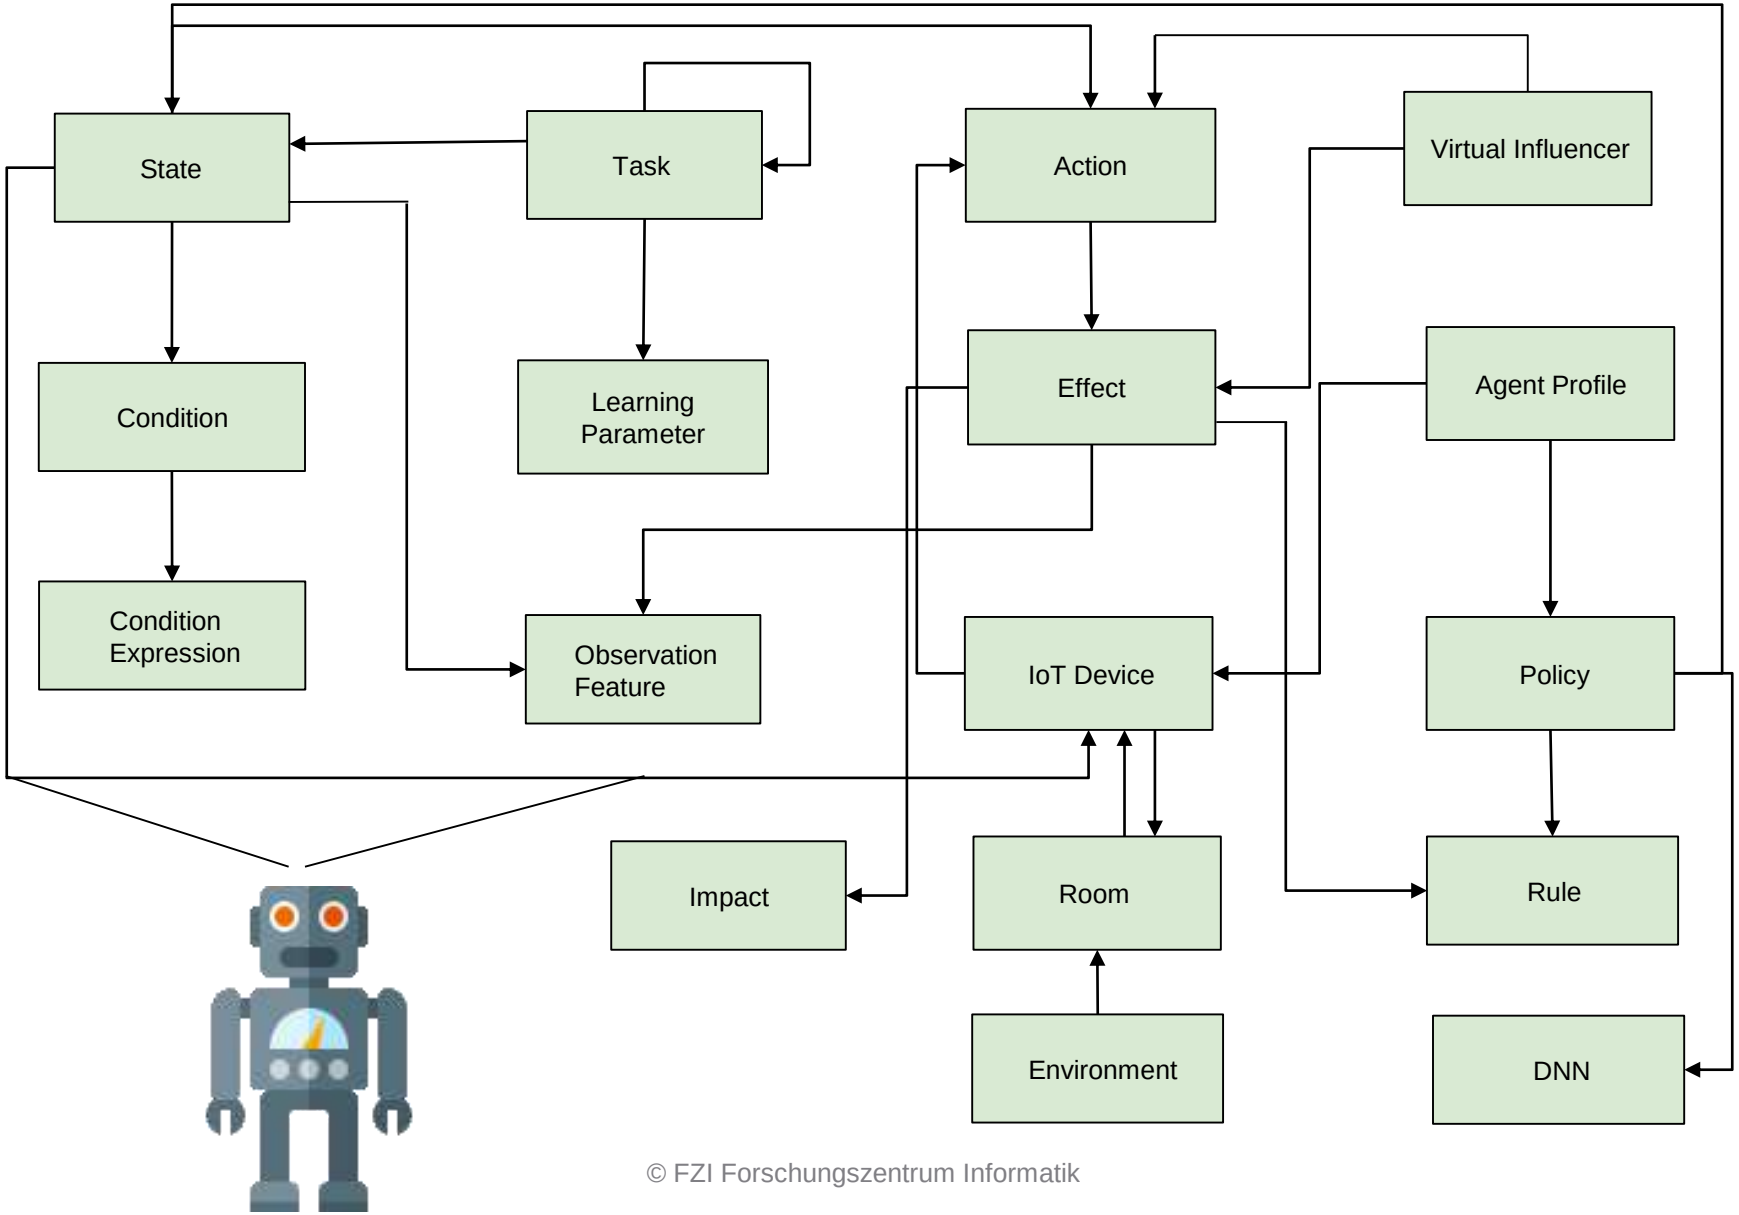

A Task as a loosely coupled set

Task = Set of States, Actions and Effects

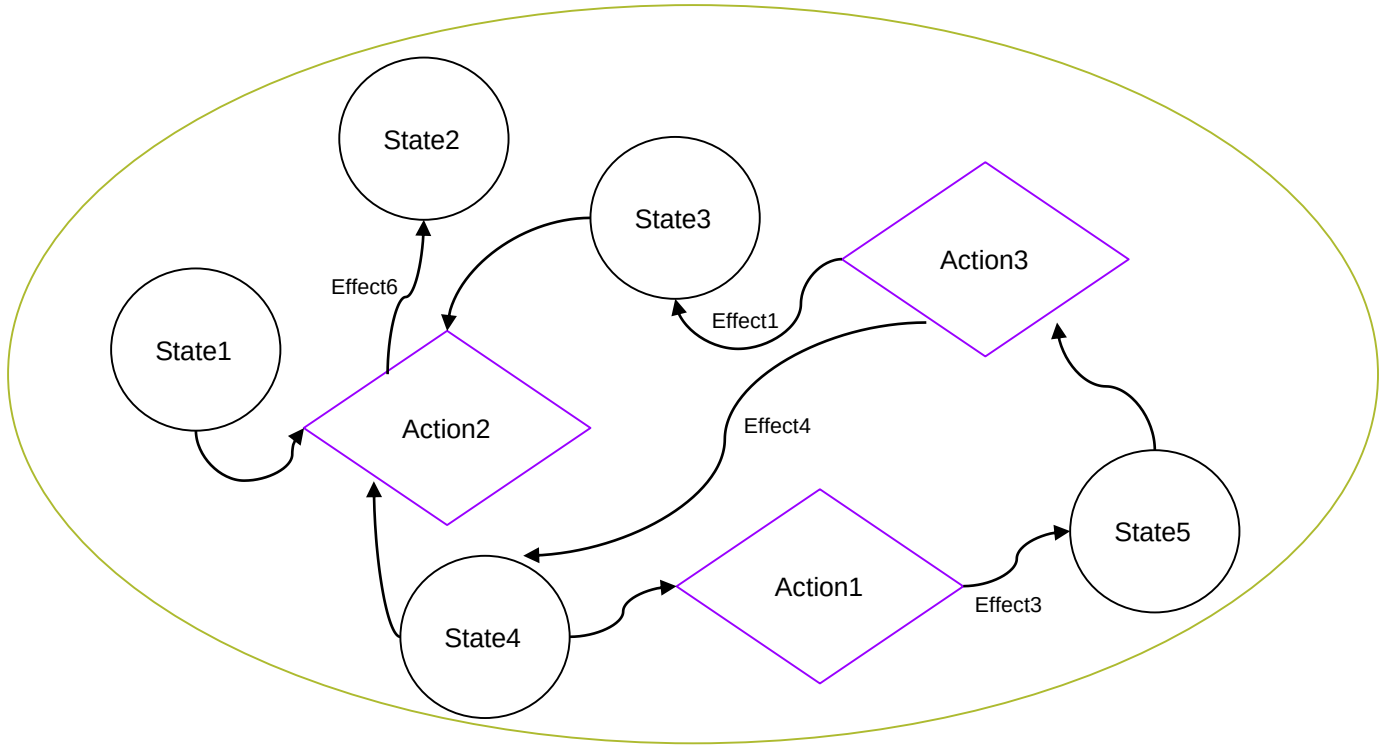

Evaluation

Setting

- **Use Case**: CKDPathway (see: http://ckdpathway.ca)
- **Features**: ACR, eGFR, Hematuria, Potassium
- **Actions**: SportExercise, LowPotassiumDiet, LowSodiumDiet, SmokingCessation, RegulateFluidIntake, ARBIntake, StatinIntake, AntiplateletIntake
- **Javascript-Library**: reinforcejs by Andrej Karpathy
- **Algorithm**: Deep-Q-Learning
- **Discount-Factor**: 0.9
- **Epsilon**: varying => probability of random actions
- **Learning rate**: 0.005
- **Experience-Size**: 10000
- **Learning-Steps-per-iteration**: 5
- **Number of hidden units**: 100
- **Throughputs**: 3 with different epsilon parameters

Results

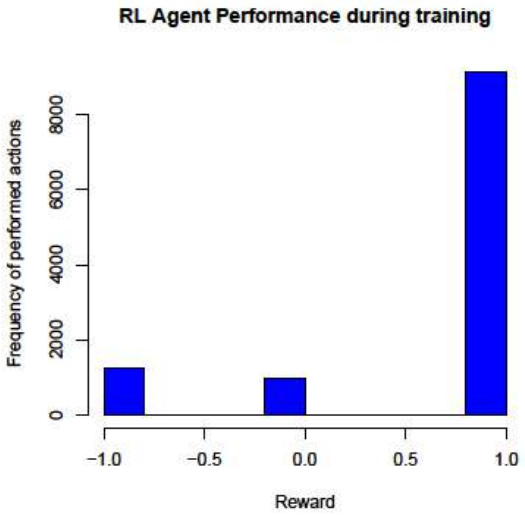

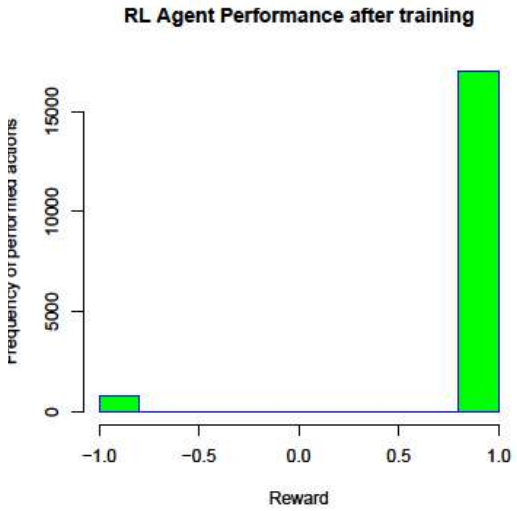

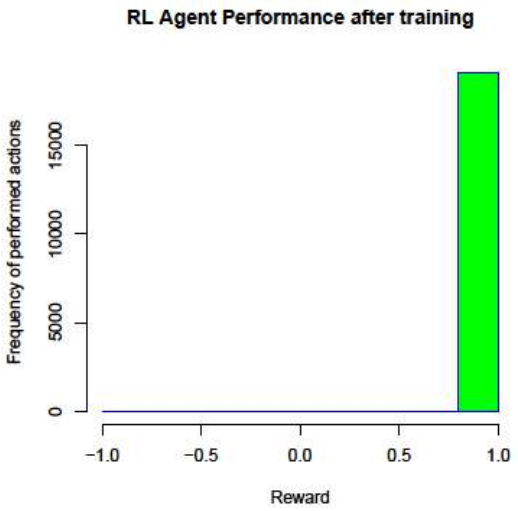

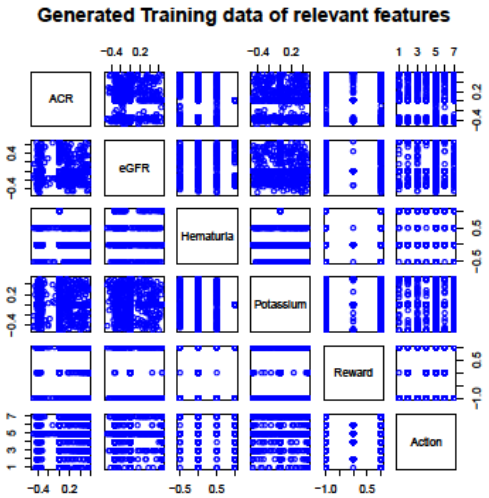

- Figure 1: **epsilon: 0.2** => 20% random actions
- Figure 2: **epsilon: 0.1** => 10% random actions
- Figure 3: **epsilon: 0.0** => no random actions, only learned actions
- Figure 4: Distribution of generated training features

Conclusion and Future Work

- Combining semantic web technologies as well as machine learning improves the integration of new complex tasks and overcomes the cold-start problem.
- Use Case-specific datasets can be created by the simulation framework.

- Agents can train beforehand in a controlled and simulated environment before they act in a real-world environment.

- The approach allows a simplified programming of agents and integration of new tasks.

- **Future Work:**
  - Laboratory studies for evaluating the performance of agents in comparison to humans.
  - Usability tests of the framework.
  - Comparison of rule-based agents with our trained RL agents.
  - Integration of Ensemble learning approaches.

Backup Slides

Roles and Data Flows

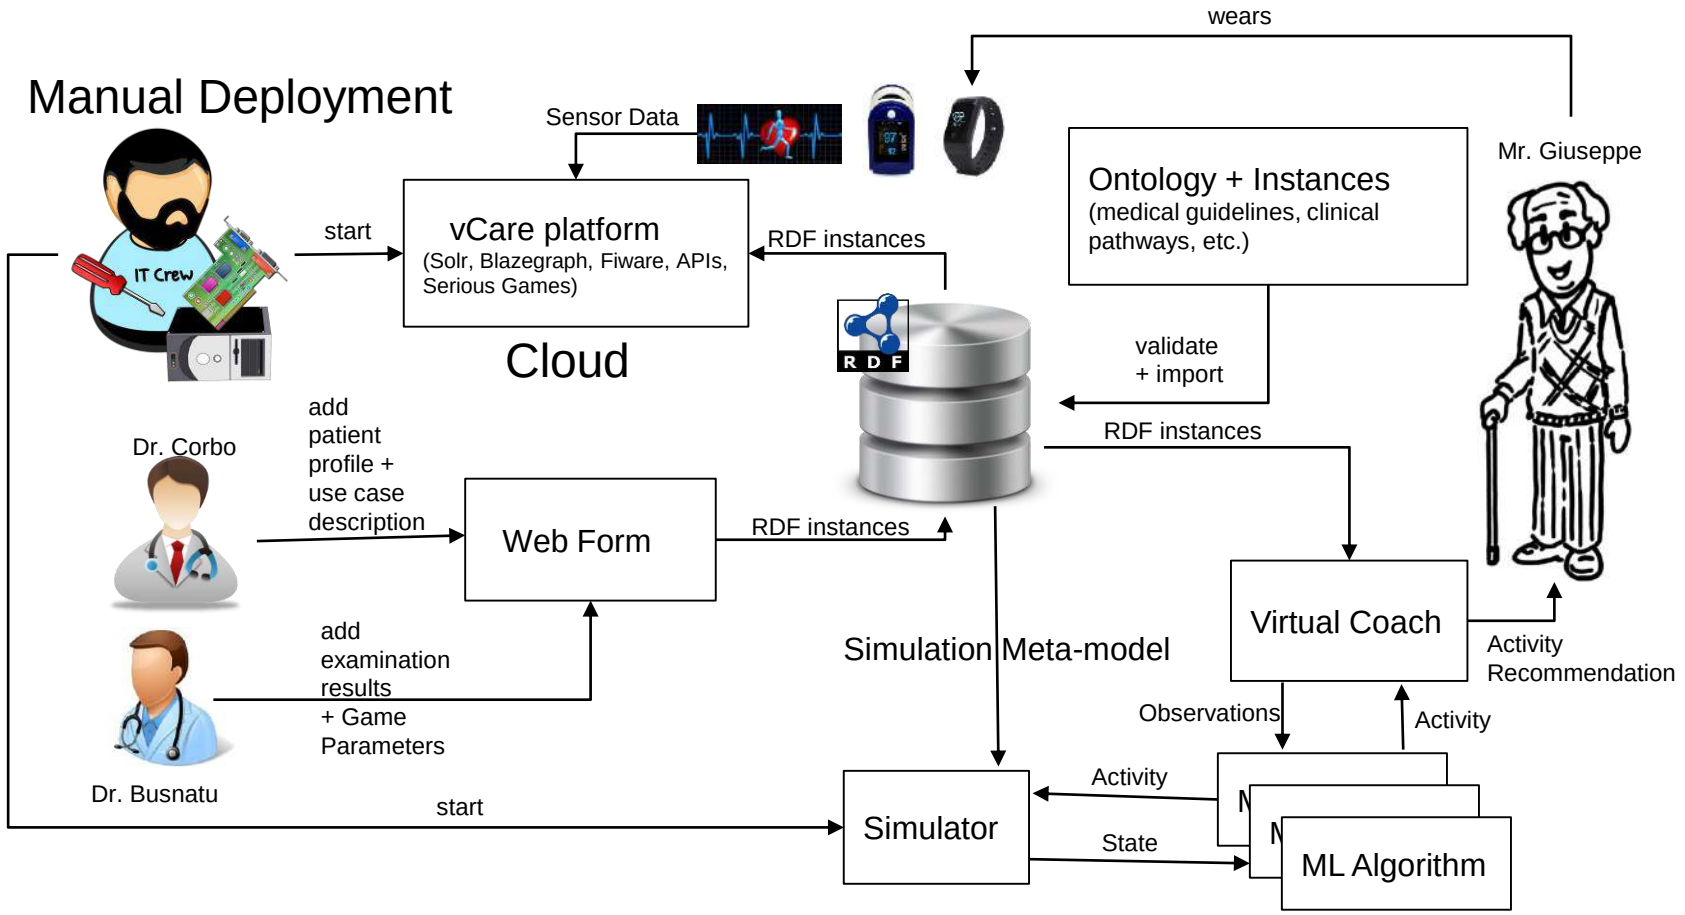

What is a clinical pathway?

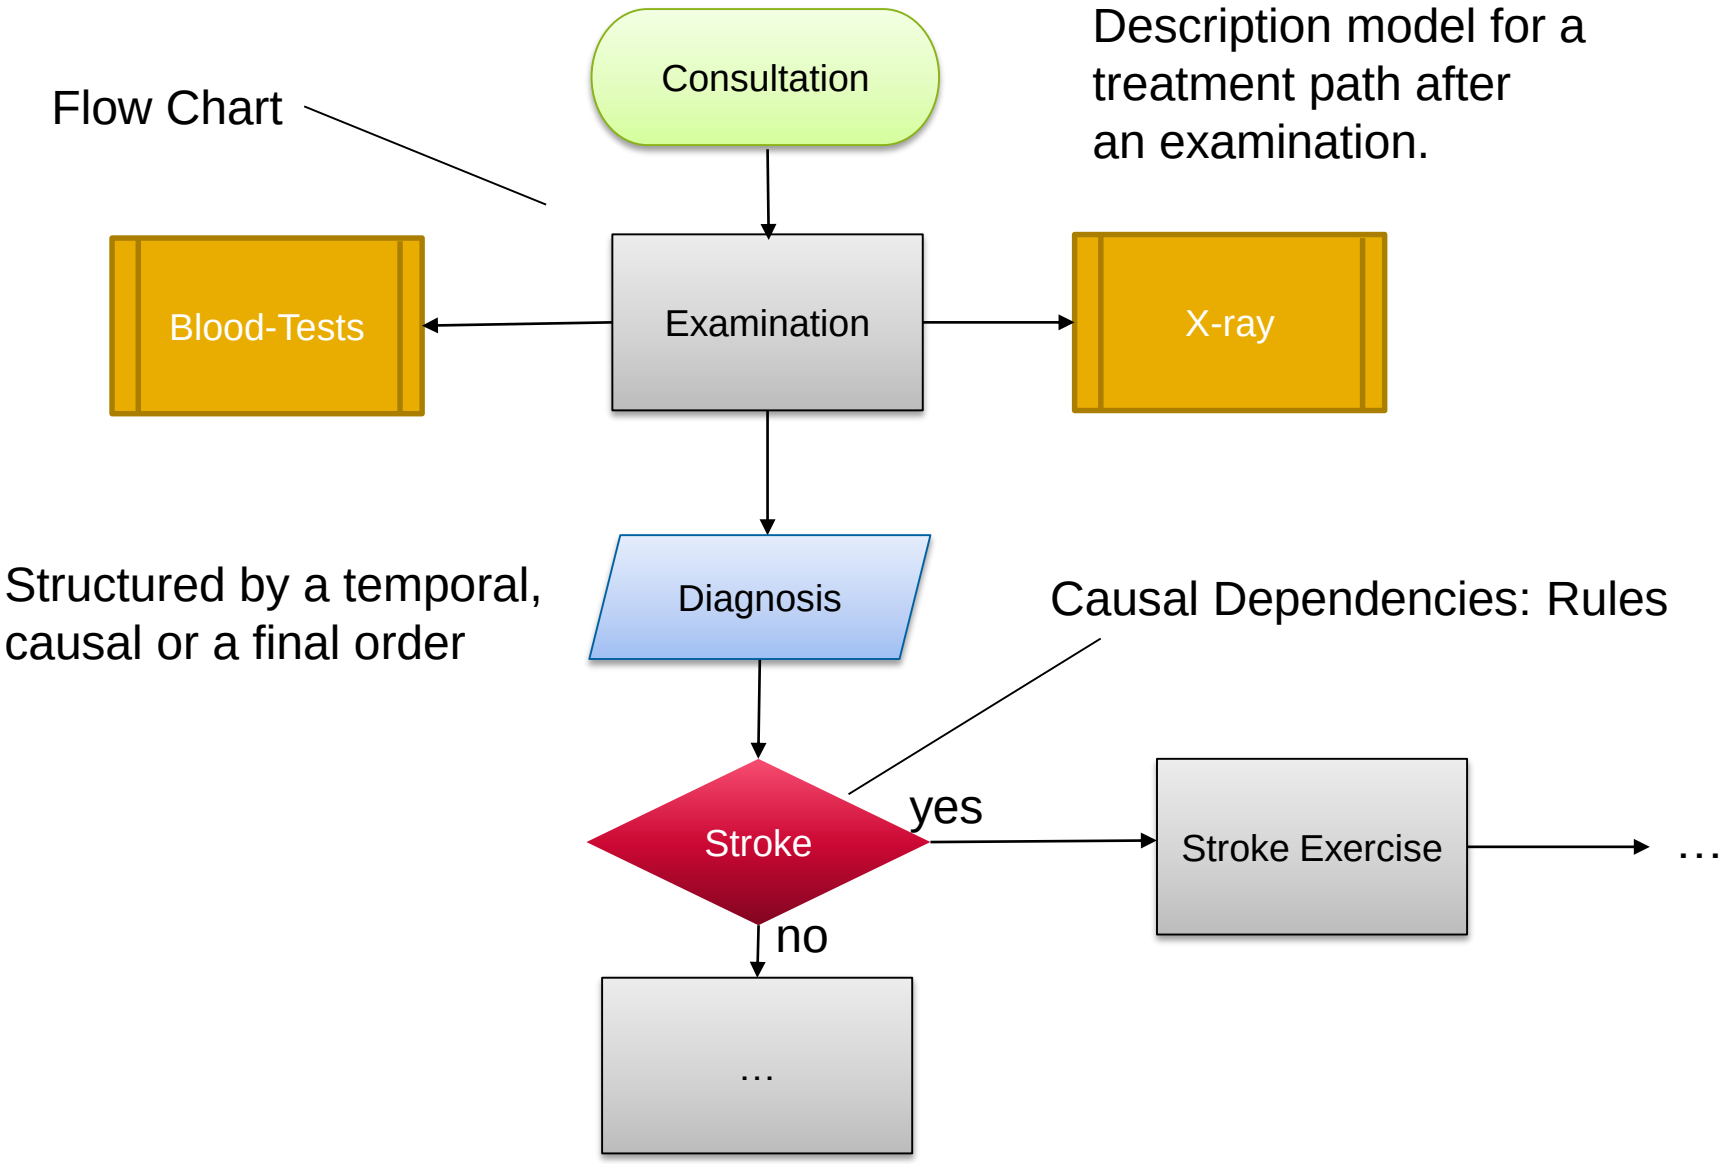

K-Means for the rating of activites

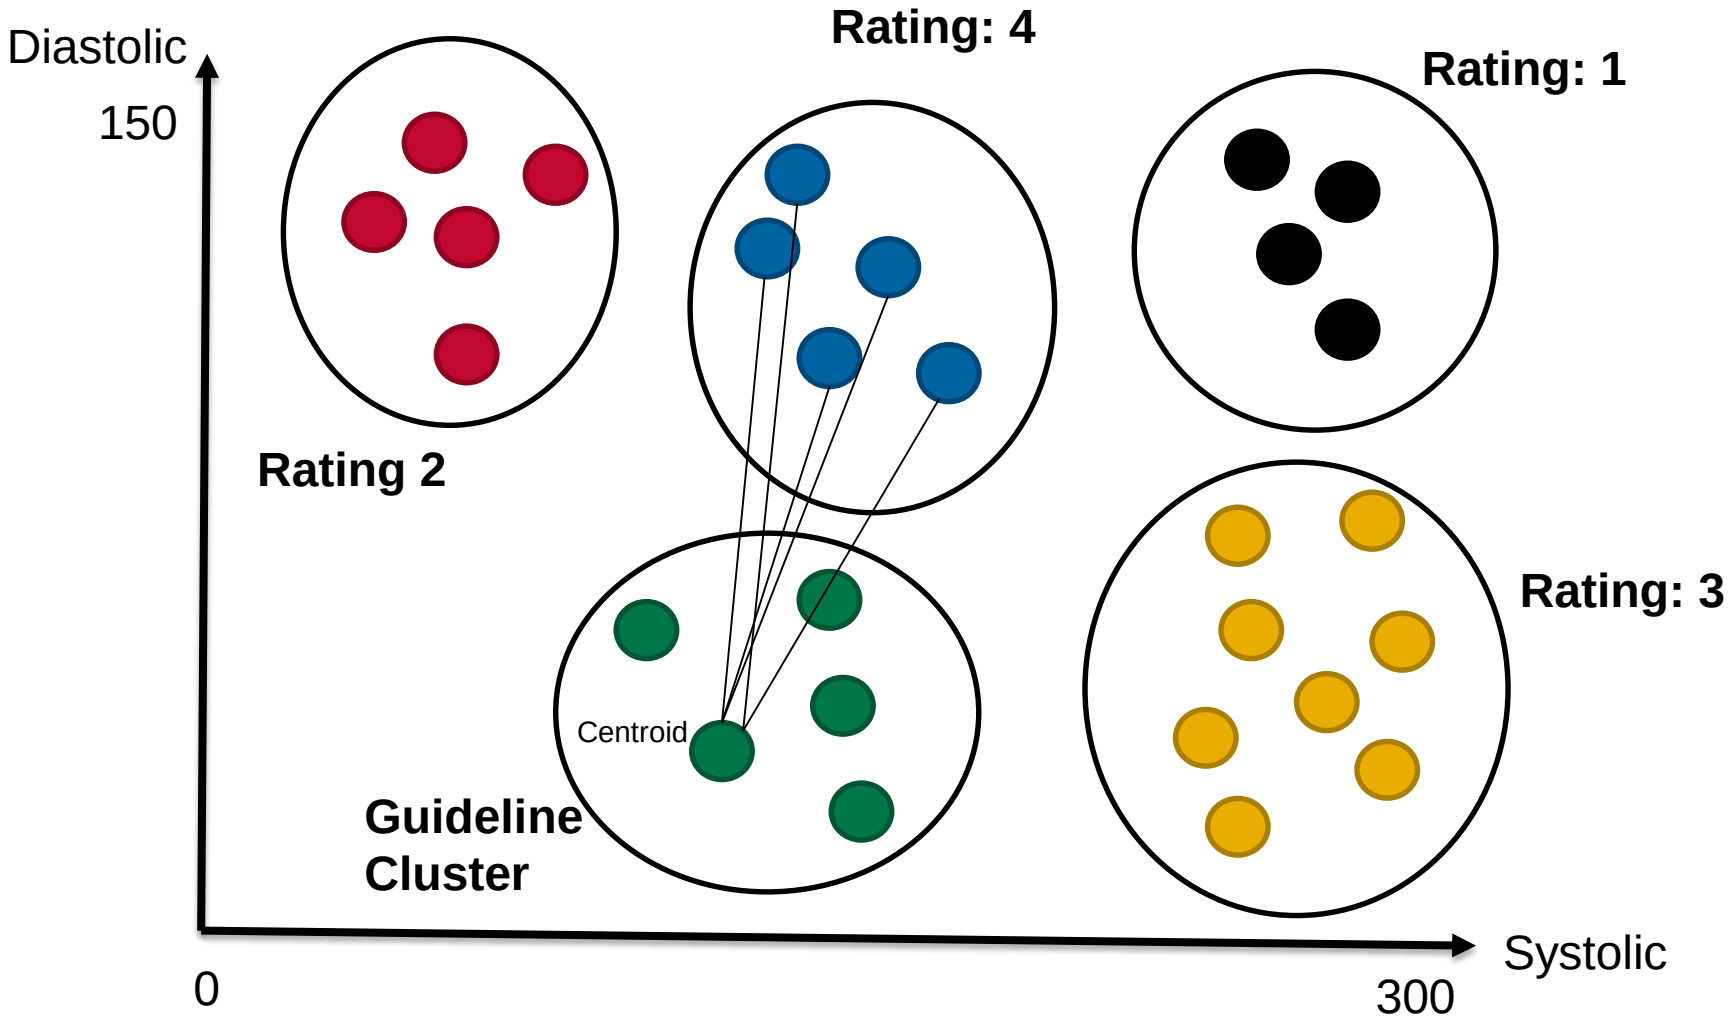

How do we assign the ratings to the clusters?

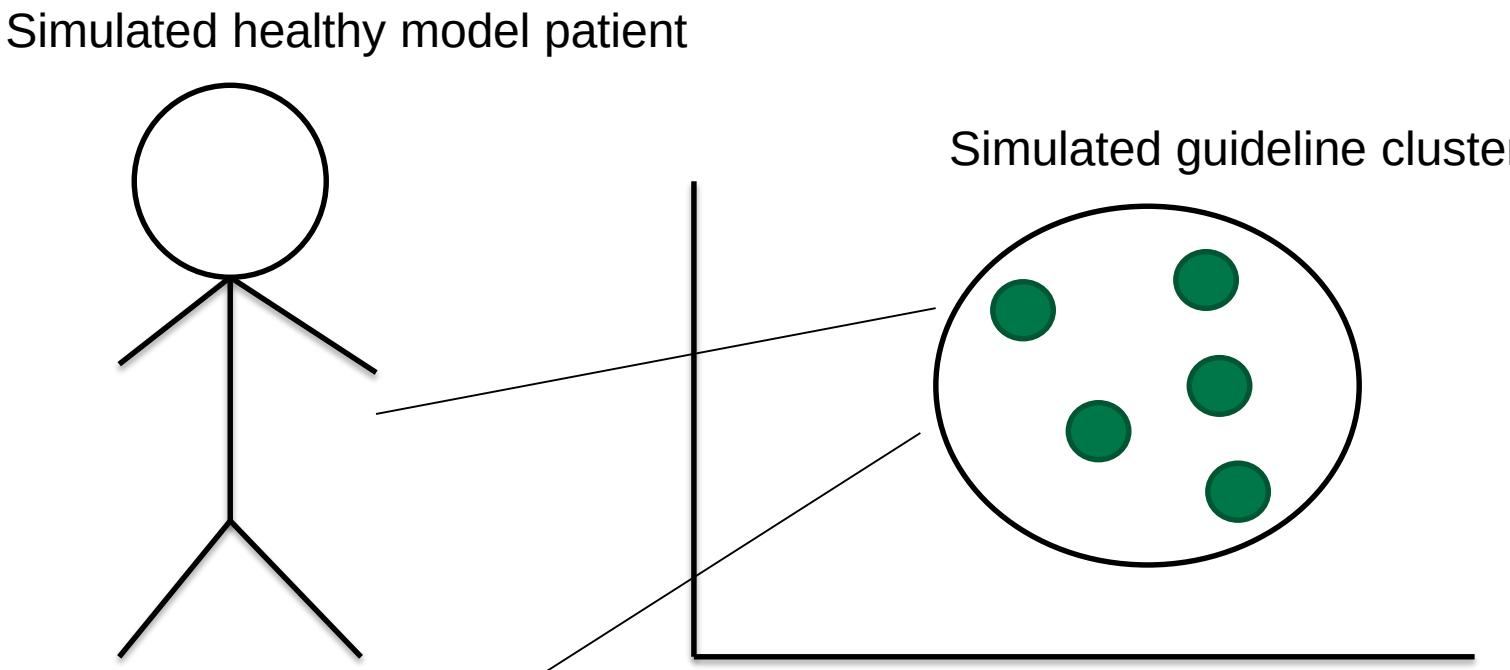

Healthy KPIs defined by domain experts:

Diastolic BP: 100 – 130
Systolic BP: 60 – 80
Heartrate: 40 – 70
BMI: 18.5 – 21
eFGR: < 3

The rating values (1-5) are assigned based on the distance of each cluster to the guideline cluster

Collaborative Filtering

- Assumption: Activities with the highest computed ratings are the best ones for the user with certain disposition.
- Predictions of ratings for activities by vital parameters of the appropriate patient can show the best activity recommendations.
- Activities are considered as items, which shall be recommended by the collaborative filtering algorithm.
- Guidelines are utilized as rating of activities => see previous slide
- Guidelines represent target observations

	Activity B	Activity C	Activity D	Activity E
Patient1	5	3	1	
Patient2		1	5	3
Patient3	1	0		4
Patient4			2	

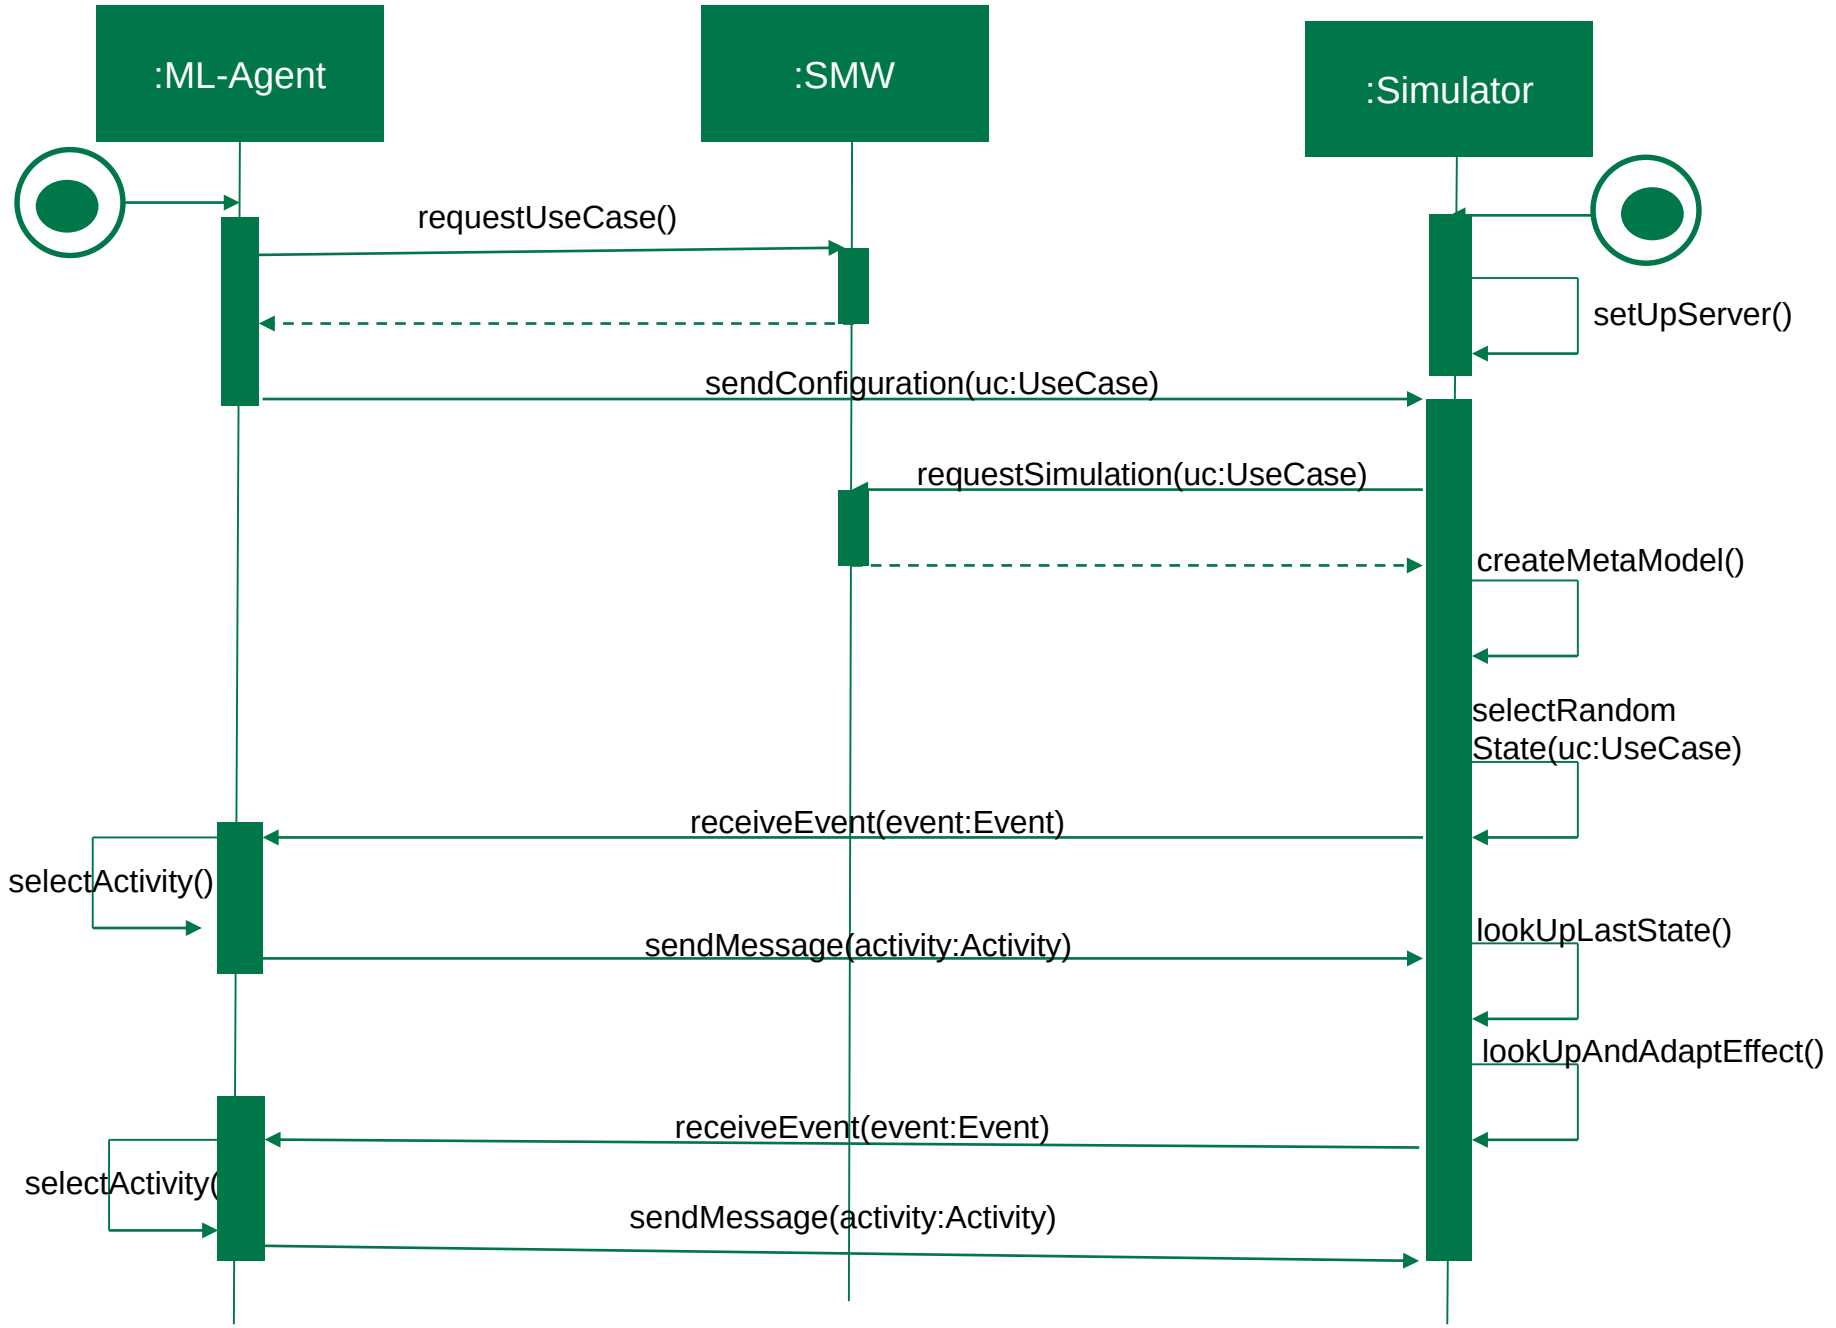
:ML-Agent
:SMW
:Simulator
requestUseCase()
setUpServer()
sendConfiguration(uc:UseCase)
requestSimulation(uc:UseCase)
createMetaModel()
selectRandom
State(uc:UseCase)
receiveEvent(event:Event)
selectActivity()
sendMessage(activity:Activity)
lookUpLastState()
lookUpAndAdaptEffect()
receiveEvent(event:Event)
selectActivity(
sendMessage(activity:Activity)

It‘s time to give feedback

Meta-Model: Simulation of Patients, their vital parameters and activities

- **Problem:** Lack of datasets in the healthcare domain => „cold-start“ problem
- How shall the ML algorithms be trained in order to provide recommendations? => We need models and datasets!
- **Solution:** Simulation of data based on a meta-model of clinical pathways and medical guidelines.

Clustering of Patient Types by Similar Activities with Latent Semantic Analysis (LSA)

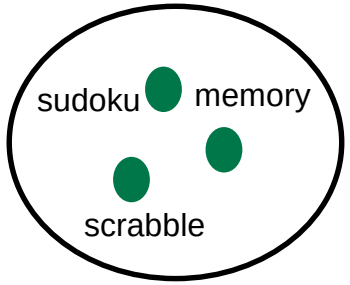

Mental disease patients

	Cardiac Patient	Mental disease Patient
Scrabble	7	40
Sudoku	5	34
Memory	0	60
Walking	80	5
Biking	95	3
Swimming	79	6

Cardiac patients