Explaining Drug Recommendations with Deep Learning

Wednesday 5/4/2023, University of Magdeburg

UNIVERSITY OF THE AEGEAN



SCHOOL OF ENGINEERING

DEPARTMENT OF INFORMATION AND COMMUNICATION SYSTEMS ENGINEERING

Presenter: Panagiotis Symeonidis

Associate Professor

http://panagiotissymeonidis.com

psymeon@aegean.gr

Outline

① Deep Reinforcement Learning for Drug Recommendations

② Graph-based Methods for Drug Recommendations along with Explanations

③ Predicting Optimal dosing of a Drug (DEMO)

Reinforcement Learning Methods for Personalized Health

Published work

1. Deep Reinforcement learning for Medicine Recommendation
(Symeonidis et al., IEEE BIBE 2022)

2. Mortality Prediction and Safe Drug Recommendation for Critically-ill Patients
(Symeonidis et al., IEEE BIBE 2022)

Problem Definition



- Can we use lab tests, health interventions, and the medicines of a patient to make predictions about her health?



- Can we use Reinforcement learning to provide drug combinations to Medical Doctors for their patients?

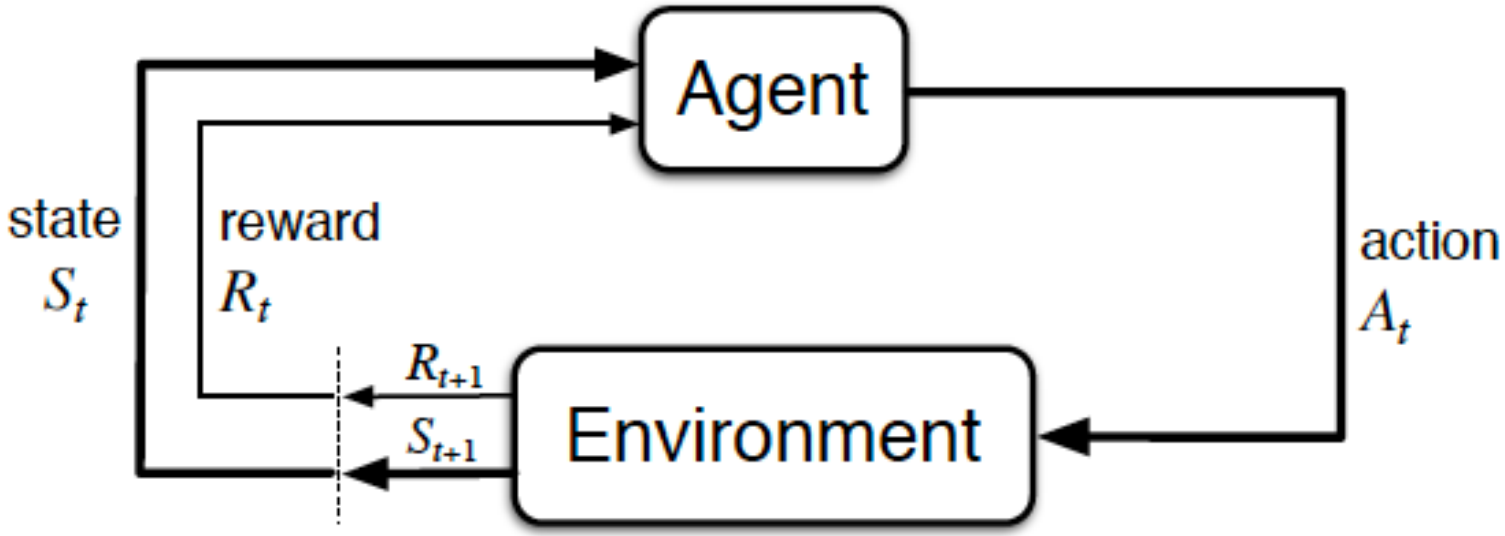


- Can we minimize the unwanted side effects among the drugs of a recommended treatment to have less toxicity?



- Can we maximize the synergistic effect of drugs for faster or more effective patient's recovery?

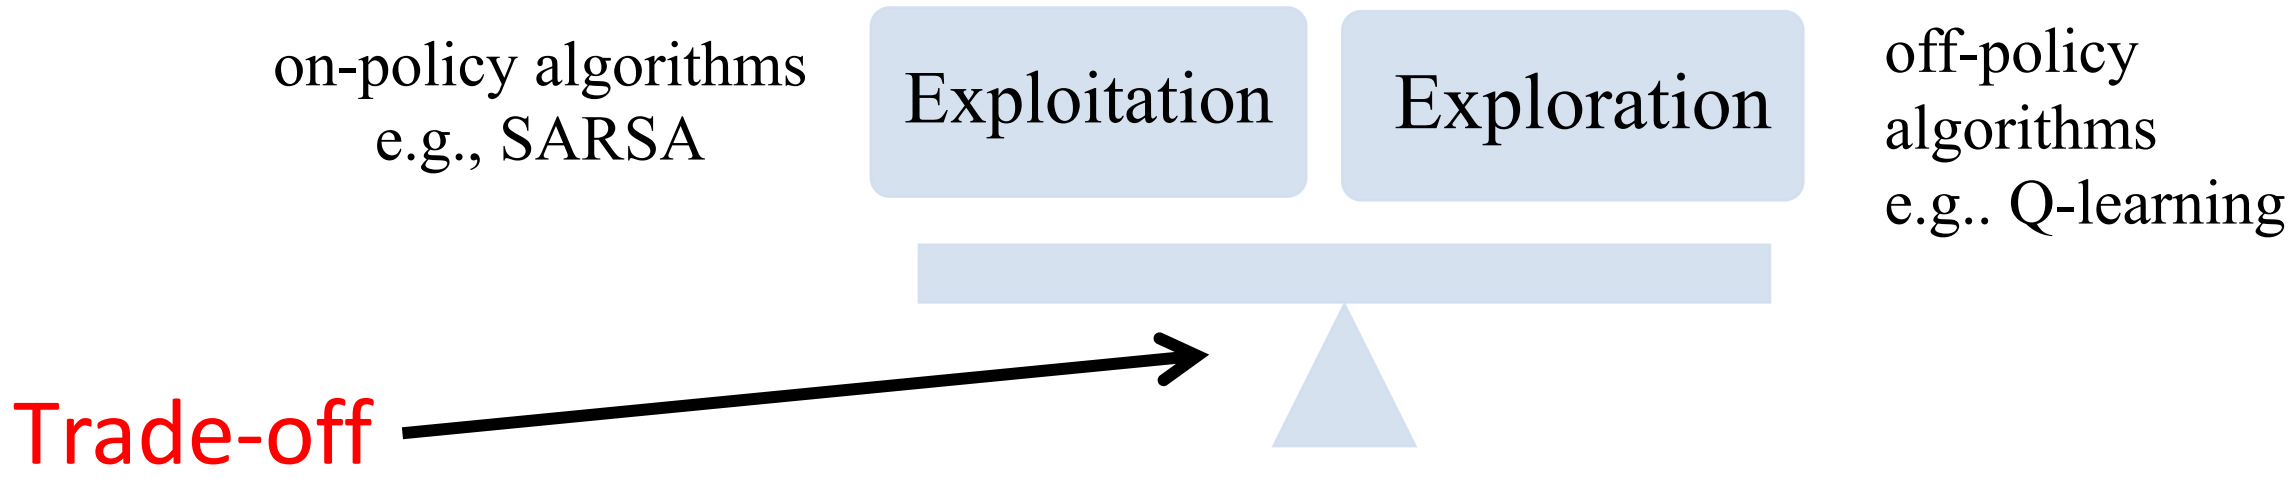
Main Components of Reinforcement Learning



- S: set of states
- A: set of actions
- T: S×A → S′ transition probability function from state S to a new state S′ after taking an action.
- R: S × A → ℝ reward function for any possible{state, action} pair.
- π: S → A: policy π function which depicts the action to take when you are in a particular state

Exploration vs. Exploitation

- **Reinforcement learning is in the middle of supervised (e.g. classification) and unsupervised (e.g. clustering) learning.**
- **After a repetitive process of "trial and error", the Agent learns the environment it interacts with.**
- **The selection of the optimal action in a state is known as Exploitation.**



Related work in optimal choice of medication

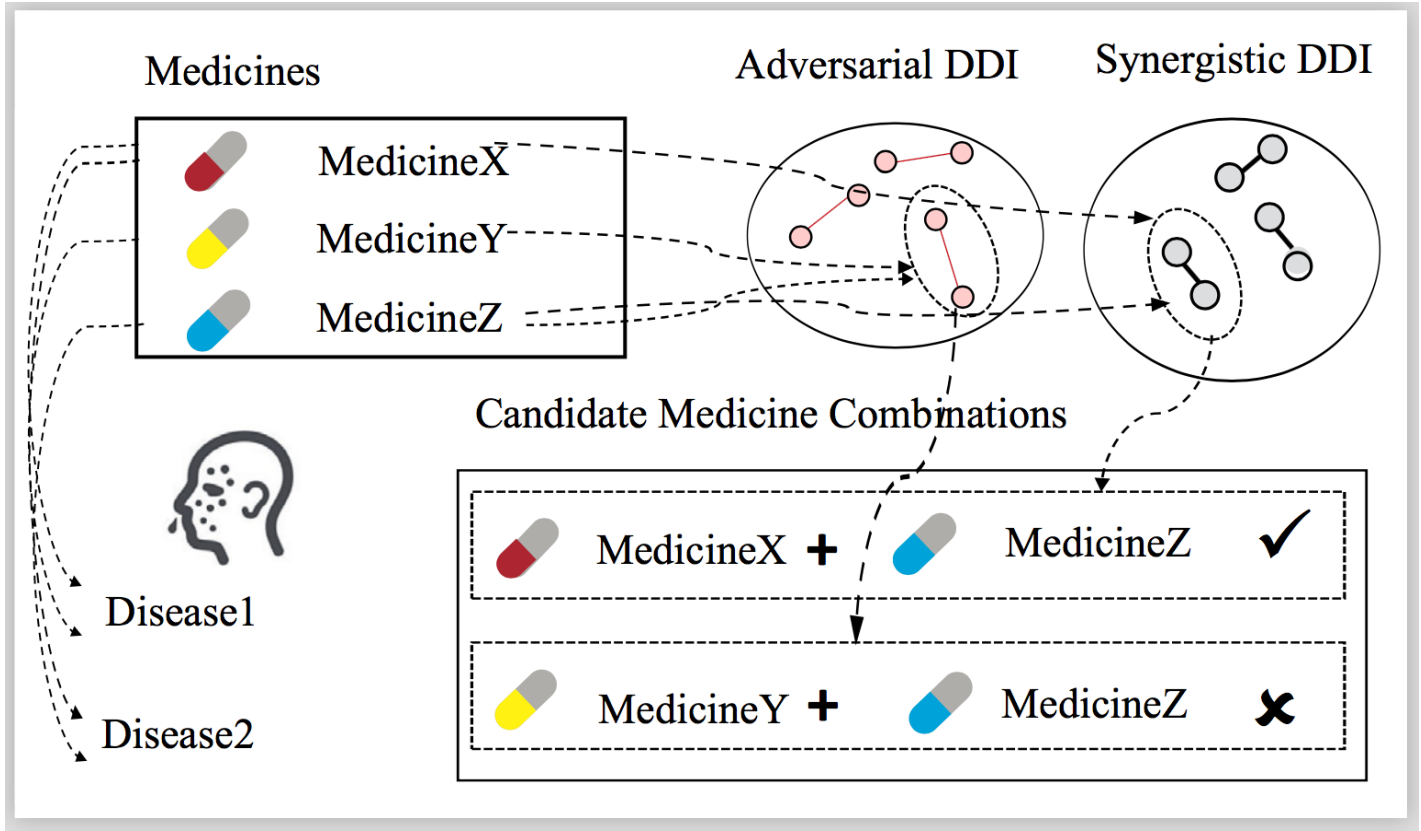
GAMENet: Graph augmented memory networks for recommending medication combination (Shang et al., AAAI 2019)

- Combines Drug Drug Interaction (DDI) knowledge with Electronic Health Records (EHR) graph using a GCN

- COMPNet: Order free Medicine Prediction with Graph Convolutional Reinforcement Learning (Wang et al., CIKM 2019).
  - Combines CNN (Convolutional), RGCN (Relational Graph) and DQN (Deep Q-Learning) Neural Networks.

Our Motivation

- Effectiveness: Provide Accurate Medicine Combinations
- Safety: Reduce the unwanted drug side effects for less toxicity
- Fast recovery: increase the synergistic action of drugs

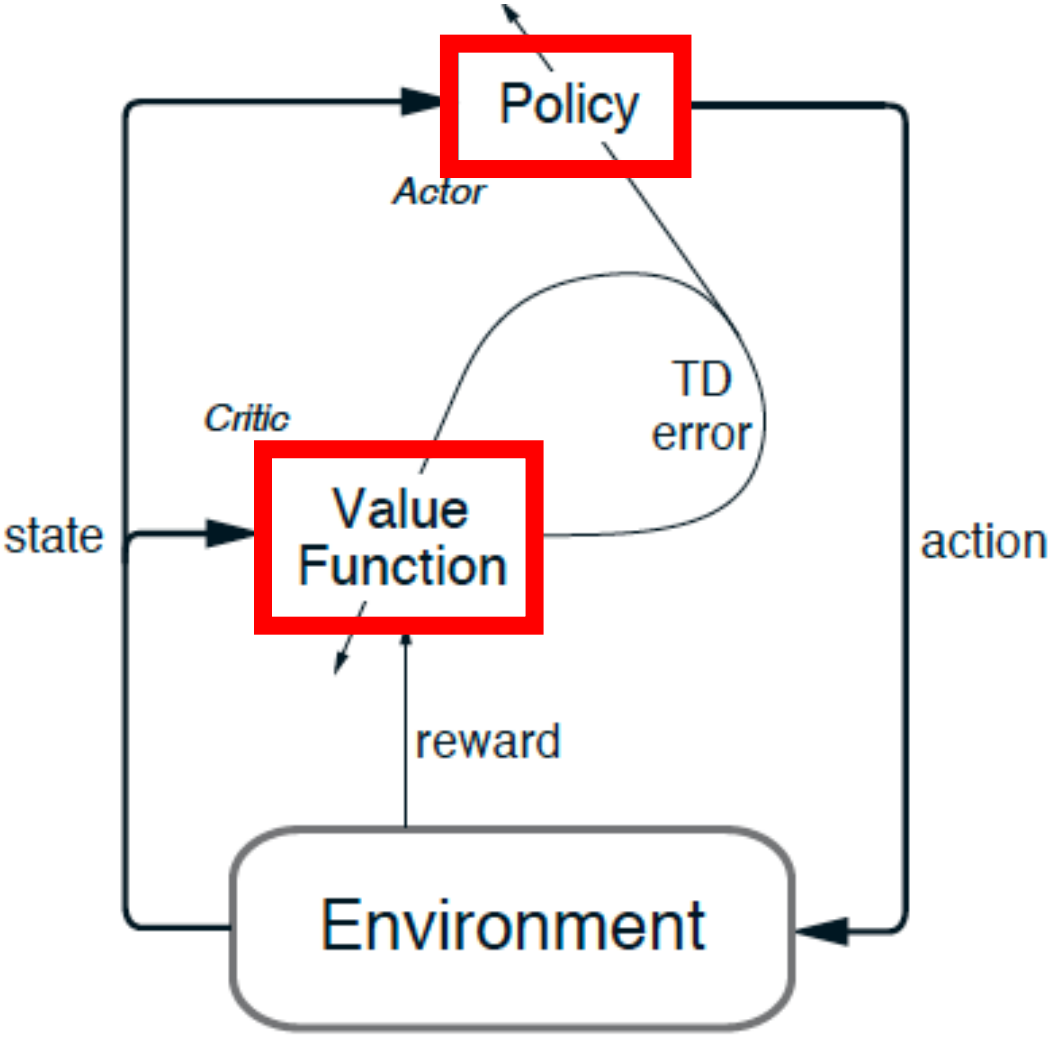


Our proposed Method
Advantage Actor Critic (A2C)

We adopt A2C, which consists of 2 neural nets:

1. gradient policy π function (actor),
2. the state value function (critic)

A2C combines the advantages of both off-policy RL methods, which are value-based, and on-policy methods, which directly optimize the current policy.



The Advantage Function and the State Value

- The Advantage function computes the benefit between the rewards we get

$$A(s,a) = Q(s,a) - V(s)$$

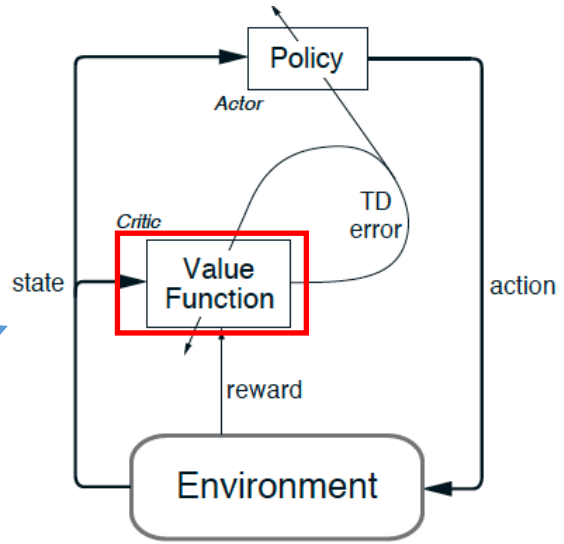
where V(s) is the value state function.

- Alternatively, we may use R instead Q(s, a), which holds the accumulated future rewards as follows:

$$A(s,a) = R - V(s)$$

where $R = max \sum_{t=0}^{\infty} \gamma_t \cdot r_t$

State Value

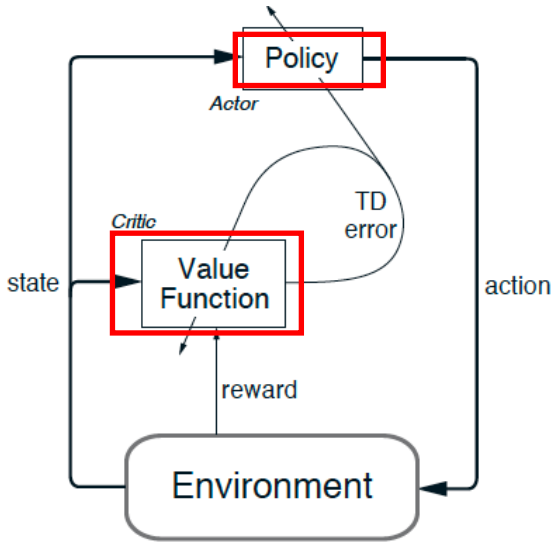


Gradient Policy function of A2C Algorithm

- The "gradient policy function" tries to directly train the intelligent agent's policy π (the action that the agent chooses at any given time):

Gradient Policy Loss $$L(\theta) = \mathbb{E}_t[\pi_\theta(a \mid s) \cdot A(s,a)]$$

where $\pi_\theta(a \mid s)$ is the probability distribution of policy π based on event θ for selecting action α when we are on state s, and $A(s, a)$ is the Advantage function



Advantage Actor Critic (A2C) algorithm

- **A2C computes for a state s**
  - **the loss function of the state value V(s) and**

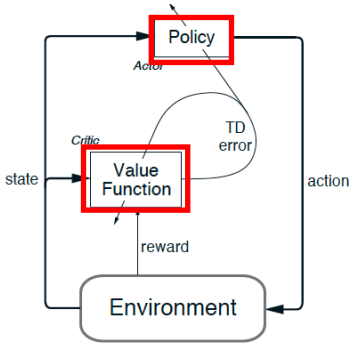
(critic) State Value Loss: $L_v = \sum(R_t - V(s_t; \theta_v))^2$

  - **the loss function of the gradient policy π,**

(actor) Gradient Policy Loss: $L_p = -\log(\pi(a_t \mid s_t; \theta)) \cdot A(s_t, a_t; \theta, \theta_v) - \beta \cdot H(\pi(s_t; \theta))$

and propagates them to the 2 neural nets.

$H(\pi(s_t; \theta)$ is the entropy for policy π and controls how much our model will explore or exploit!



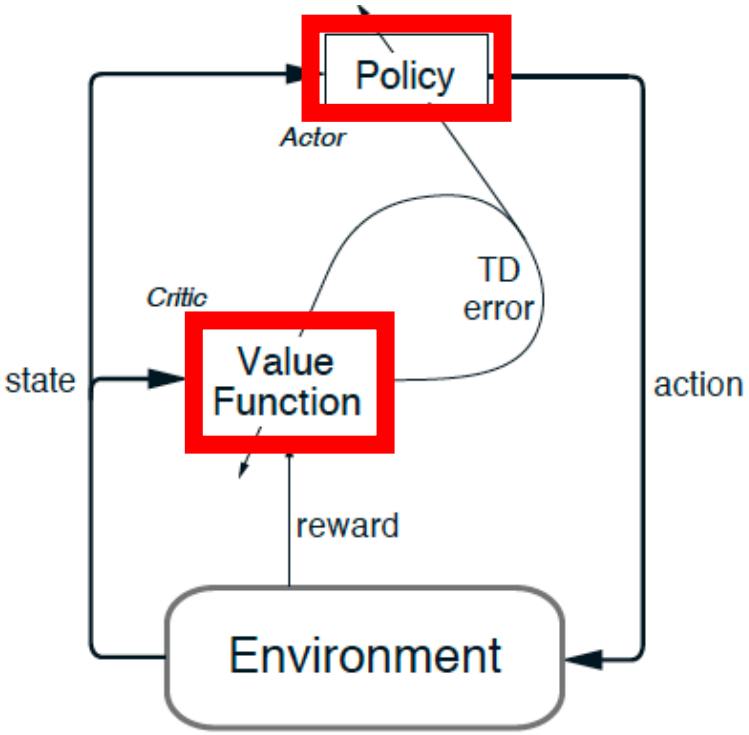
The Update Rules of A2C

- The update rules for parameters θ and θ_v of actor and critic respectively are as follows:

(critic) Value of state loss: $$\theta_v \leftarrow \theta_v + \nabla_{\theta_v} \cdot L_v$$

and

(actor) Gradient Policy's loss: $$\theta \leftarrow \theta + \nabla_{\theta} \cdot L_p$$



Data sets

Input data : (i) patient clinical status (ii) lab tests (iii) health interventions
(iv) Medicines (v) unwanted drug-drug interactions (vi) synergistic drug effects

MIMIC III sub set

- # patients 5,847
- # treatments 13,732
- # diagnoses 1,954
- # medicines 138
- # patient-drug interactions 122,482
- # Drug-Drug side effects 460

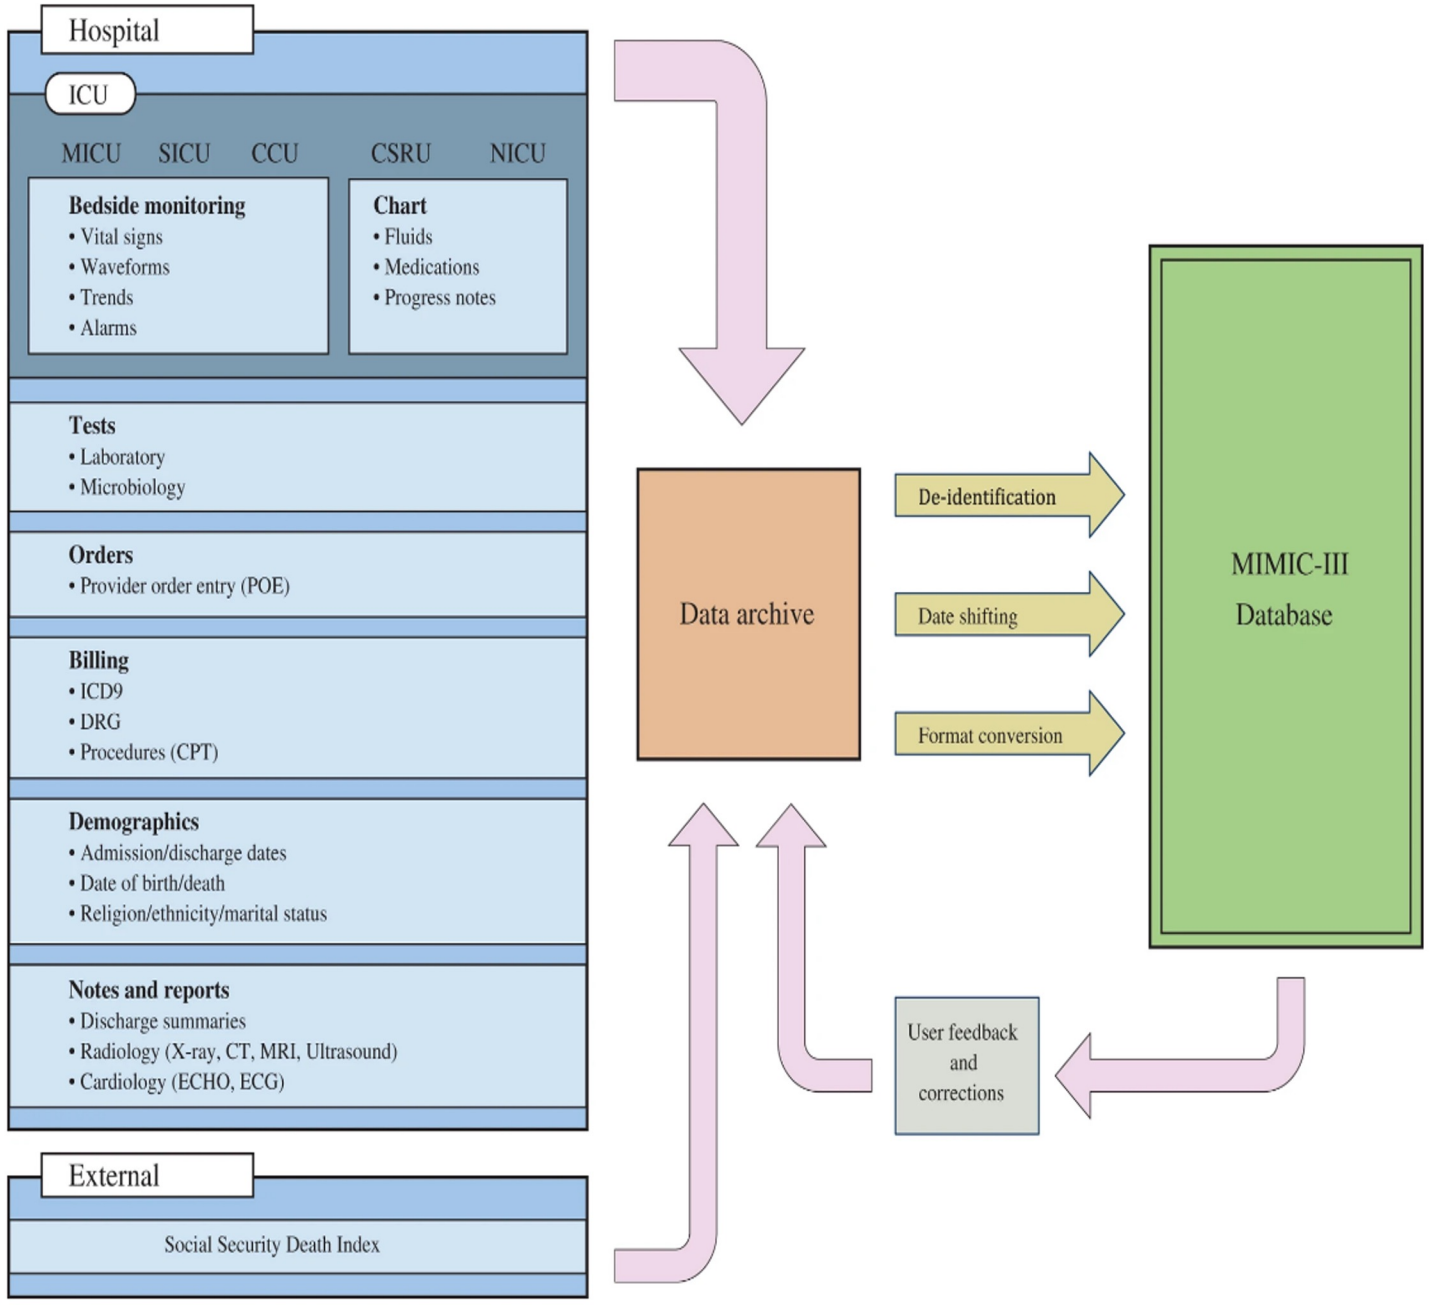
Cancer data set

- # patients 3,287
- # treatments 55,871
- # diagnoses 56
- #medicines 218
- # patient-drug interactions 350,026
- # Drug-Drug synergistic effect. 570

- Data De-identification and Anonymization
- Critical data are removed and shifted chronologically

MIMIC III Data set

- Beth Israel Deaconess Medical Center in Boston, Massachusetts
- 38,597 distinct adult patients
- 4,579 charted observations (’chartevents’)
- 380 laboratory measurements (’labevents & inputevents’)
- Demographic, Procedures/Interventions, Medications
- Patient follow-up after discharge
- 44 GB of medical data!



Johnson, Alistair EW, et al. "MIMIC-III, a freely accessible critical care database." *Scientific data* 3.1 (2016): 1-9.

Evalutation Protocol: Precision vs. Recall

- **Precision is** the fraction of relevant drugs out of all predicted drugs

$$Precision = \frac{tp}{tp + fp}$$

False positive are drugs which were predicted as relevant but the patient does not need them

- **Recall** is the fraction of relevant drugs out of all prescribed drugs from the doctors

$$Recall = \frac{tp}{tp + fn}$$

False negatives are the drugs which were relevant but were not predicted as being so

Since we should correctly identify all those drugs that were prescribed by the doctor, recall should be considered as more important metric than precision.

Evalutation Protocol: Drug-Drug unwanted side effects Rate

$$DDI\ rate = \frac{1}{m}\sum_{i}^{m}\frac{\left|\left(c_j, c_k\right) \in Y_i \& \left(c_j, c_k\right) \in \mathcal{A}_{ddi}\right|}{\sum_{j,k} 1}$$

where $\sum_{j,k} 1$ is the total number of possible pairwise drug combinations, $\mathcal{A}_{ddi}$ represents all DDI pairs in the DDI knowledge base, i is a sample index in the test set representing a patient and m is the size of the test set (i.e., the number

$$\text{DDI change \%} = \frac{\text{DDI Rate (EHR+DDI)} - \text{DDI Rate (EHR)}}{\text{DDI Rate (EHR)}} \times 100.$$

Top-10 most common unwanted side effects

('arterial_pressure_decreased', 28568)
('anaemia', 27006)
('Difficulty_breathing', 26037)
('nausea', 25190)
('diarrhea', 23848)
('edema_extremities', 21981)
('body_temperat_increased', 21806)
('Hypoventilation', 21322)
('chest_pain', 21013)
('dizziness', 20204)

Reinforcement learning (RL) algorithms used in our experiments

Advantage-Actor Critic Reinforcement Learning (A2C)

[Mnih et al., 2016]

Double Deep Q-learning Network algorithm (DDQN)

[Hasselt, 2010]

DDQN is an improvement of DQN. DQN's problem is the overestimation of Q-values. DDQN uses a second Q-table which is a copy of the original Q-table but only with the last values computed per action-state pair. So old main trends cannot overburden new trends.

Comparison with other state-of-the-art methods

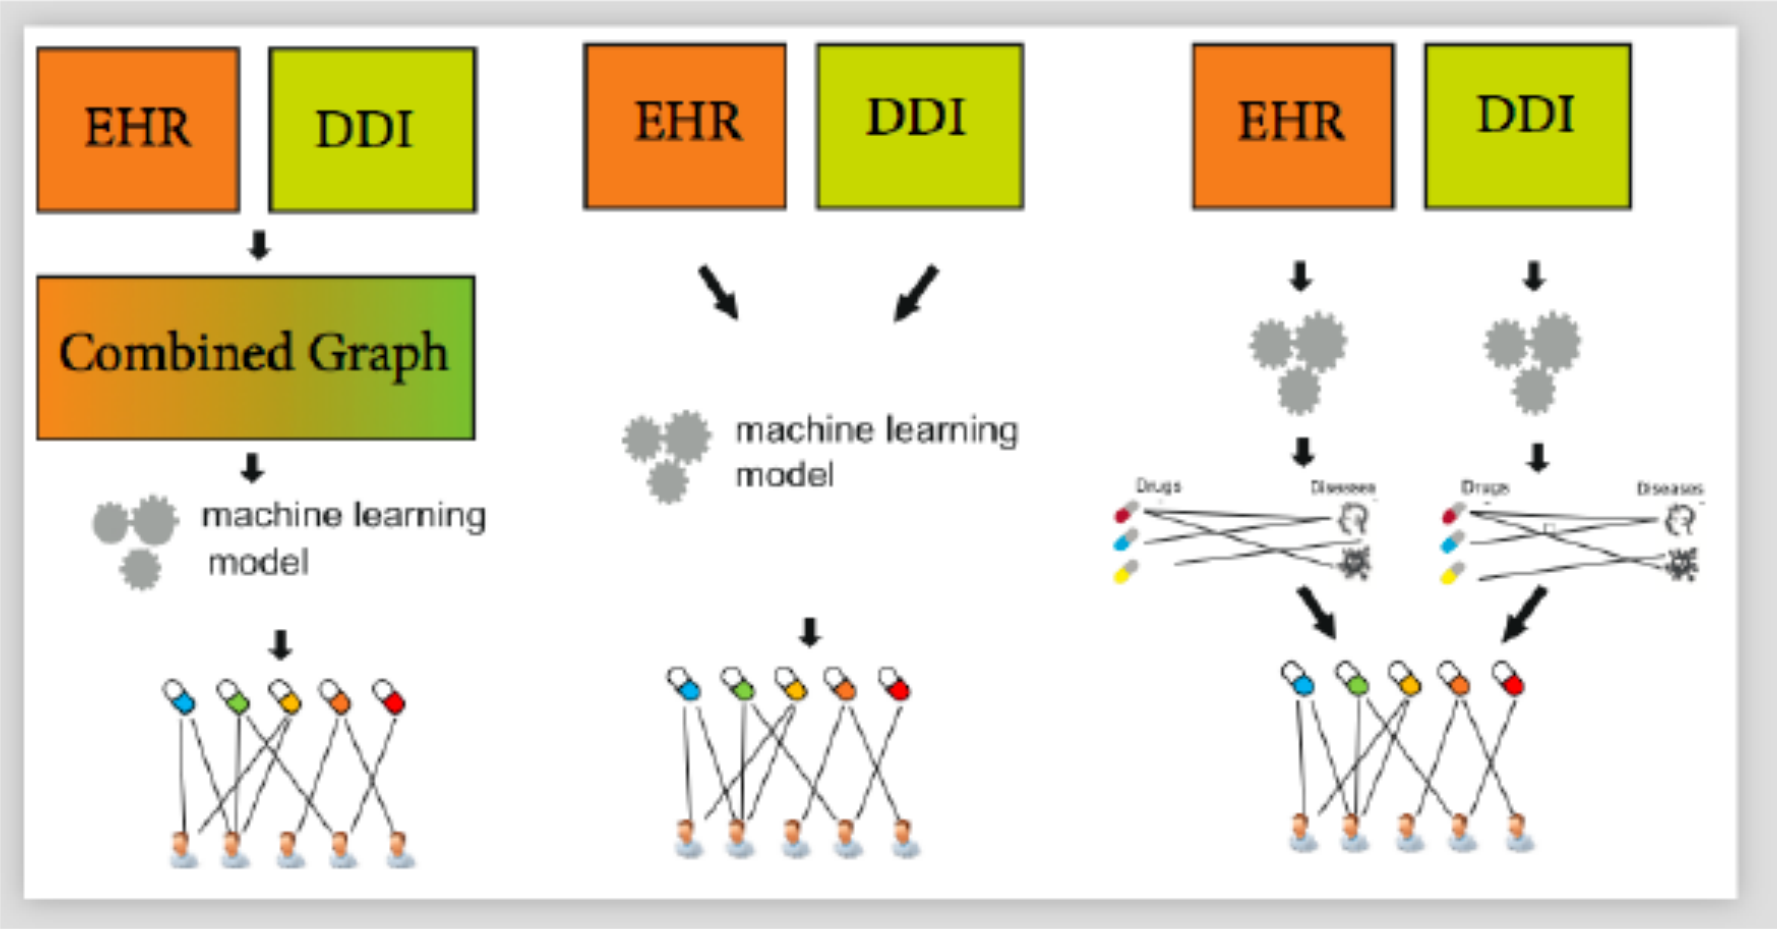
Collective Matrix Factorization (CMF)

[Singh et al., 2008]

Random Walk with Restart (RWR)

[Tong et al., 2006]

Electronic Health Records (EHR) & Drug Drug Interactions (DDI)



(RWR) (CMF, A2C, DDQN) (HYBRID)

Q-learning algorithm

Algorithm 1: Q-learning

Input:

States $S = \{s_1, s_2, ..., s_n\}$
Actions $A = \{a_1, a_2, ..., a_n\}$
Reward Function $R : S \times A \to \mathbb{R}$
Transition Probability Function $T : S \times A \to S$
Policy π Function, $\pi : S \to A$
Learning rate $\lambda \in [0, 1]$, usually $\lambda = 0.1$
Reward attenuation coefficient $\gamma \in [0, 1]$

Output:

A table Q which holds the value of each action a in state s

procedure Q-LEARNING($S, A, R, T, \lambda, \gamma, \beta$)
  Initialization of the values of the Q-table
  while Q-values have not converged **or** we have not reached a specified number of iterations **Do**
    Choose randomly a state $s \in S$
    while state s is not terminal **Do**
      Choose based on *policy* π with β prob. a random action a_t, else
      take with prob. $1 - \beta$ the best action $\pi(s_t) \leftarrow \max_a Q(s_{t+1}, a)$
      $a_t \leftarrow \pi(s_t)$ ▷ $\pi : S \to A$
      $r \leftarrow R(s_t, a_t)$ ▷ reward reception
      $s_{t+1} \leftarrow T(s_t, a_t)$ ▷ transition to the new state
      $Q(s_{t+1}, a) \leftarrow (1 - \lambda)\ Q(s_t, a_t) + \lambda \times (r + \gamma\ max_a Q(s_{t+1}, a))$
      $s \leftarrow s_{t+1}$
    end while
  end while
  return Q-table
end procedure

Collective Matrix Factorization (CMF)

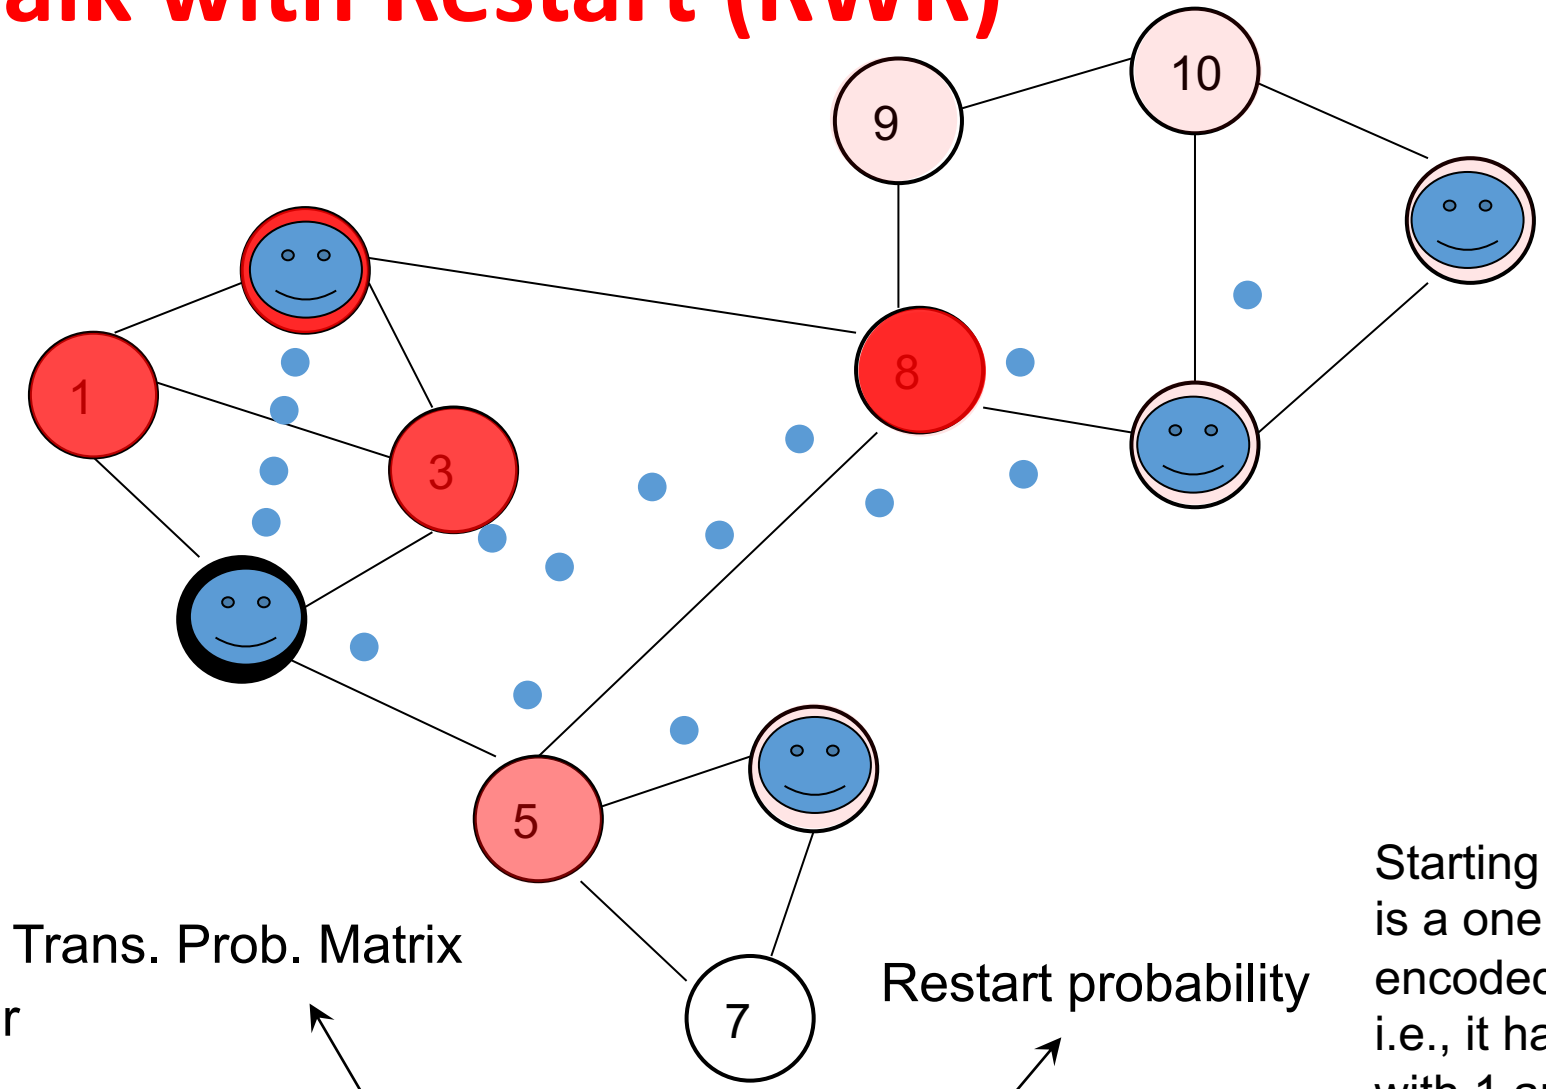
Let us suppose that matrices $X \in R^{I \times J}$ and $Y \in R^{I \times K}$ have the first mode in common. The objective function for their joint factorization can be formulated as:

$$f(\mathbf{U}, \mathbf{V}, \mathbf{W}) = \left\| \mathbf{X} - \mathbf{U}\mathbf{V}^{\mathsf{T}} \right\|^2 + \left\| \mathbf{Y} - \mathbf{U}\mathbf{W}^{\mathsf{T}} \right\|^2$$

where ||.|| denotes the Frobenius norm. The goal is to find the matrices $U \in R^{I \times R}$, $V \in R^{J \times R}$ and $W \in R^{K \times R}$ that minimize the objective function.

Note that U, i.e., the factor matrix extracted from the shared mode, is common in factorization of both X and Y.

Random Walk with Restart (RWR)



Ranking vector

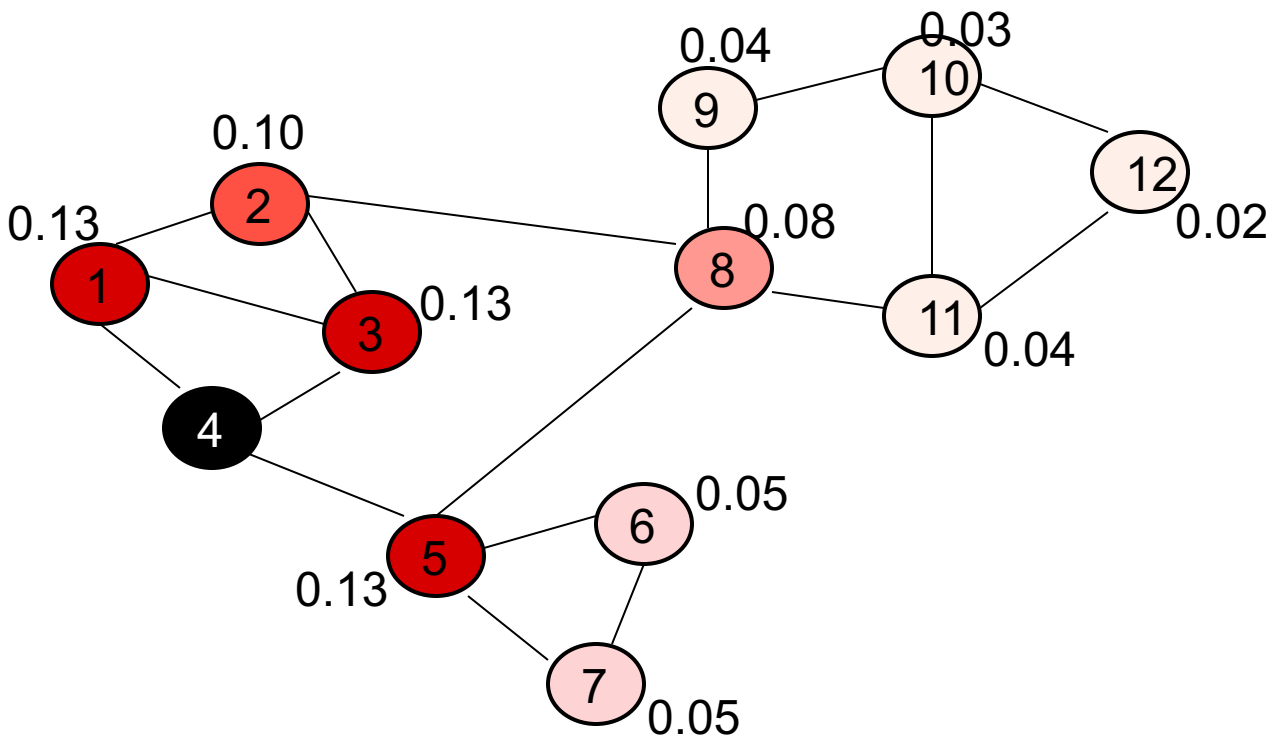
Trans. Prob. Matrix

Restart probability

Starting vector is a one hot encoded vector. i.e., it has one entry with 1 and the rest are zeros.

$$\vec{r}_i = c\tilde{W}\vec{r}_i + (1-c)\vec{e}_i$$

Random Walk with Restart (RWR)



	Node 4
Node 1	0.13
Node 2	0.10
Node 3	0.13
Node 4	0.22
Node 5	0.13
Node 6	0.05
Node 7	0.05
Node 8	0.08
Node 9	0.04
Node 10	0.03
Node 11	0.04
Node 12	0.02

Nearby nodes, higher scores

More red, more relevant

Ranking vector

$\vec{r}_4$

Example of DDQN's Recommended Medication combination

Patient	Method	Recommended Medication Combination
Diagnosis: Ischemic heart disease Pulm. embolism, Hypothyroidism Depression, Morbid obesity, Valvular disease	DDQN	sodium chloride, potassium chloride, linezolide, amikacin, labetalol, digoxin, norepinephrine, aztreonam
	Groundtruth	Sodium chloride, D-glucose,Insulin, Calcium Gluconate, Acetaminophen, Potassium chloride, Meperidine, Morphine Sulfate, Metoclopramide Magnesium Sulfate', Aspirin, Docusate sodium, oxycodon-acetaminophen, ranitidine

DDQN predicts successfully labetalol and digoxin for heart disease, but fails to predict any medication for hypothyroidism and depression.

For this specific patient, DDQN recommends drugs by attaining 25% recall, when we compare it with the ground truth medicines

Sensitivity analysis of RL methods

z parameter controls the magnitude of the reward

	Cancer Data Set			MIMIC III data set		
$(z_t, -z_t)$-values	(2, -2)	(5, -5)	(3, -1)	(2, -2)	(5, -5)	(3, -1)
DDQN	0.0485	0.0127	**0.1532**	0.0111	0.0373	**0.0628**
A2C	0.1067	0.1317	**0.1626**	0.0864	0.0821	**0.0991**

β parameter controls the exploitation vs. exploration of A2C

	Cancer Data Set			MIMIC III data set		
β-values	0.01	0.1	1	0.01	0.1	1
A2C - precision	**0.1626**	0.1353	0.1411	**0.0991**	0.0927	0.0899
A2C - recall	**0.2575**	0.2141	0.2234	**0.2364**	0.2208	0.2140

Comparison with other state-of-the-art methods

Model	Cancer Data Set				MIMIC III data set			
	Prec.	Recall	DDI rate	DDI Change %	Prec	Recall	DDI rate	DDI Change %
CMF(EHR)	0.1221	0.1008	0.0001	0 %	0.0861	0.2319	0.5656	0 %
CMF(EHR+DDI)	0.0276	0.0344	0.0034	+3,300.0 %	0.0746	0.2248	0.1154	-390.0 %
RWR(EHR)	0.0423	0.0674	0.0016	0 %	0.0848	0.2259	0.0150	0 %
RWR(EHR + DDI)	0.0341	0.0564	**0.0170**	966 %	0.0848	0.2259	**0.0150**	0 %
DDQN(EHR)	0.1532	0.2425	0.0001	0 %	0.075	0.179	0.53334	0%
DDQN(EHR + DDI)	0.0717	0.1124	0.0002	+200.0%	0.0628	0.1495	0.3555	-50.0%
A2C(EHR)	**0.1626**	**0.2575**	0.0001	0 %	**0.0991**	**0.2364**	0.64445	0%
A2C(EHR + DDI)	0.0625	0.0978	0.0002	+200.0 %	0.0532	0.1261	0.6222	-3.6 %

- Effectiveness: Provide Accurate Medicine Combinations
- Faster Recovery: increase the synergistic action of drugs to get a higher share of inhibited cancer cells
- Effectiveness: Provide Accurate Medicine Combinations
- Safety: Reduce the unwanted drug side effects for less toxicity

Contribution

- Reinforcement learning algorithms can provide accurate drug recommendations.
- We used synergistic DDIs for faster healing.
- We used adversarial DDIs to minimize the toxicity in the patients' treatment and suggest safer drug combinations.
- As future work, we want to define different levels of unwanted drug side effects and re-train our algorithms.

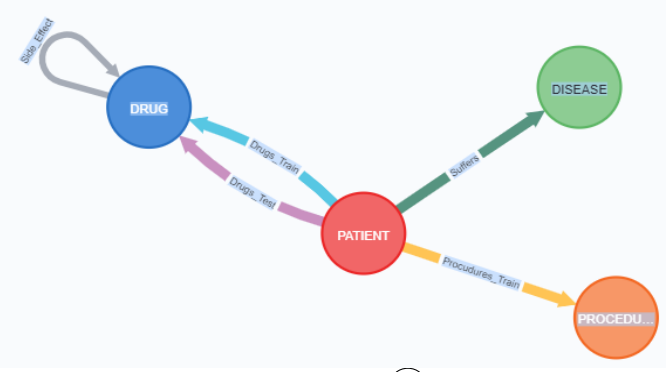
Graph-based Methods for Personalized Health

Published work

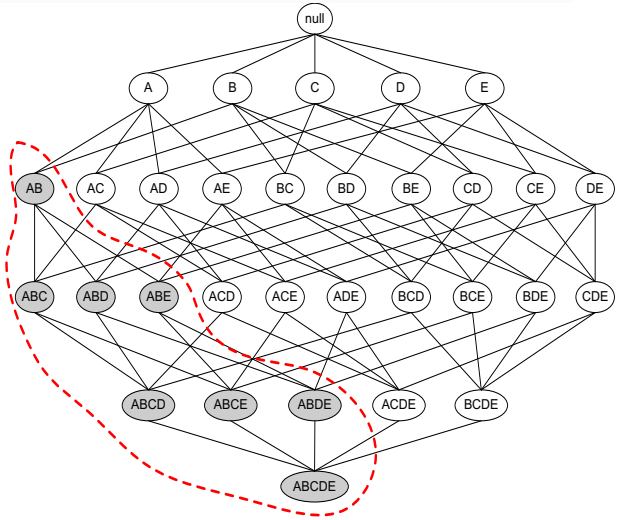
1. Safe, effective and explainable drug recommendation based on medical data integration (Symeonidis et al., UMUAI journal, 2022)
2. Treatment Recommendations for COVID-19 Patients along with Robust Explanations (Symeonidis et al., IEEE CBMS 2021)
3. Recommending what drug to prescribe next for accurate and explainable medical decisions (Symeonidis et al., IEEE CBMS 2021)

Problem Definition

- Can we integrate heterogeneous medical data and represent the patient's clinical status with a Knowledge Graph?



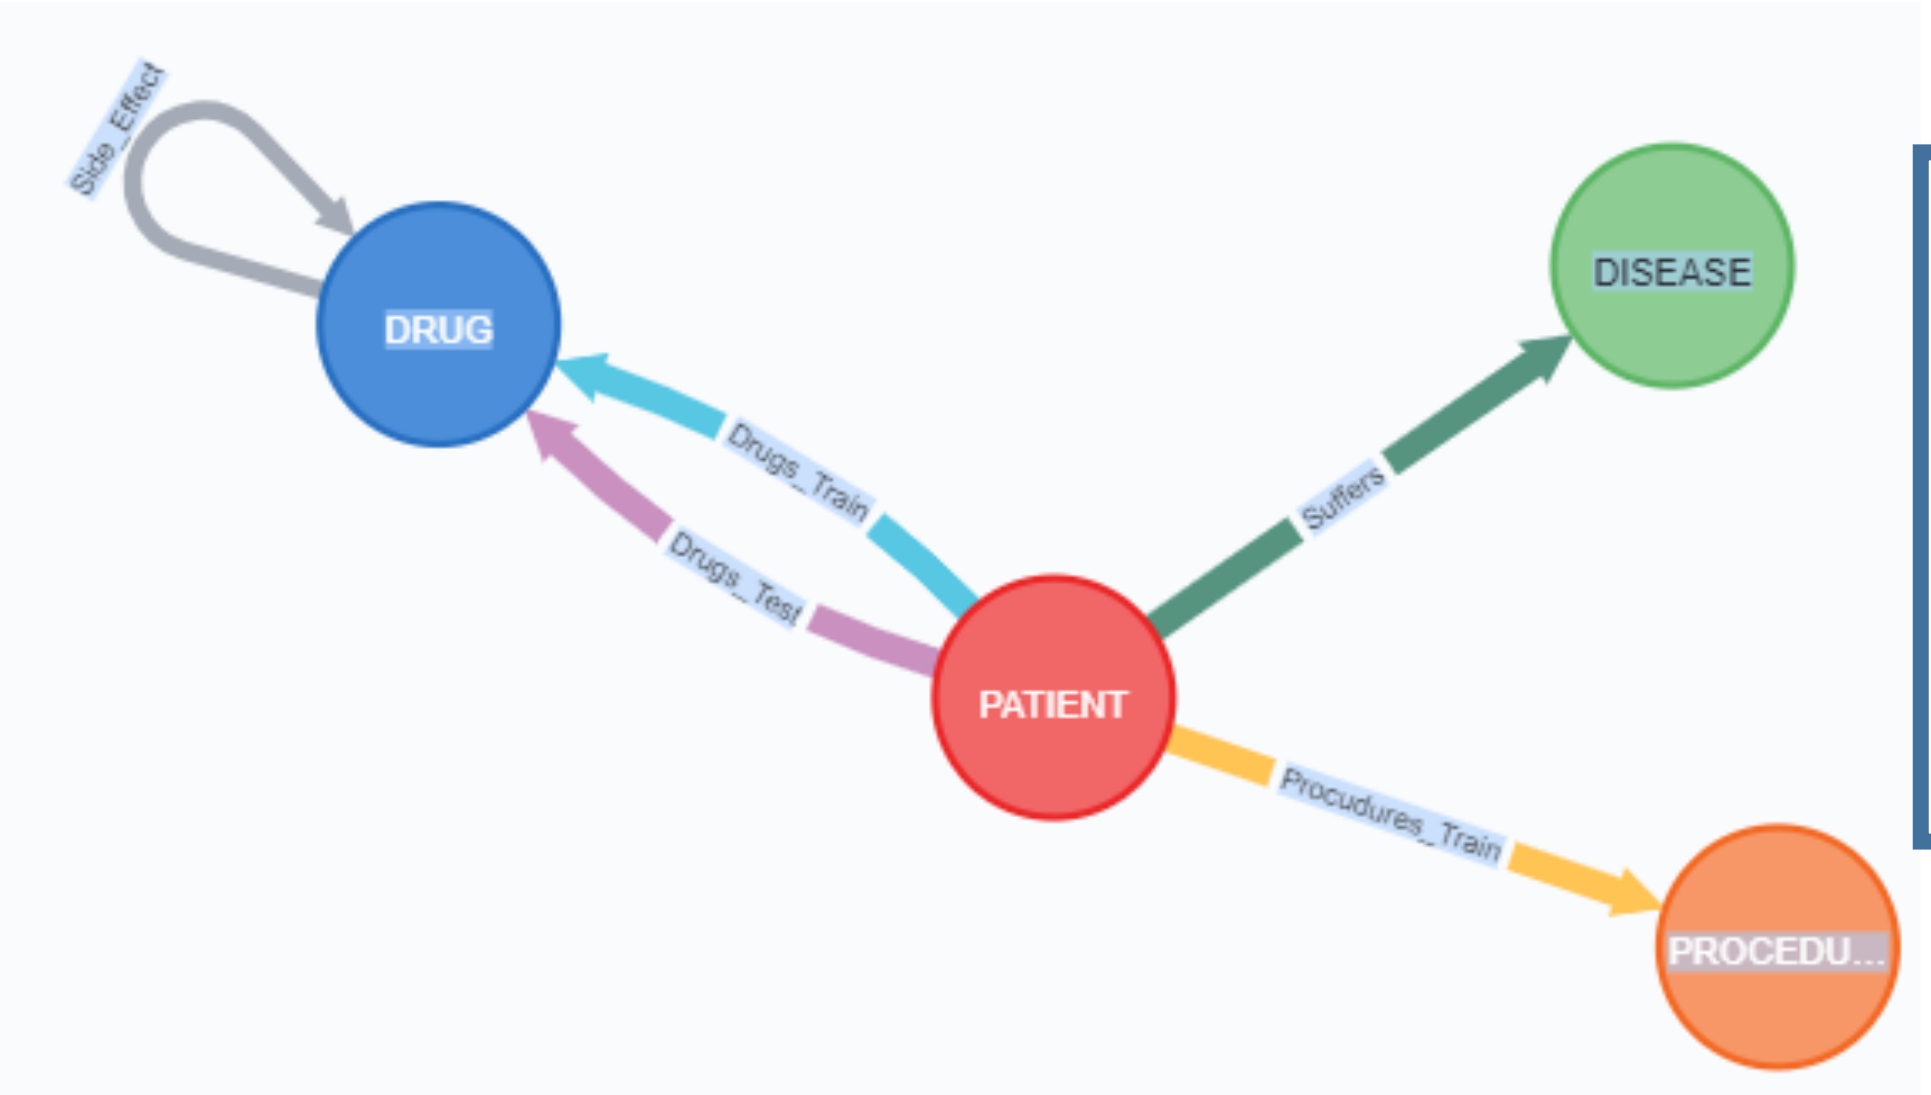
- How can we extract frequent patterns and infer relevance among the participating entities of a knowledge graph in Health?



- How can we support Medical Doctors with recommendation of drugs to be prescribed to patients and explain our recommendations?



Knowledge Graph for Medical Data (RDF triples)



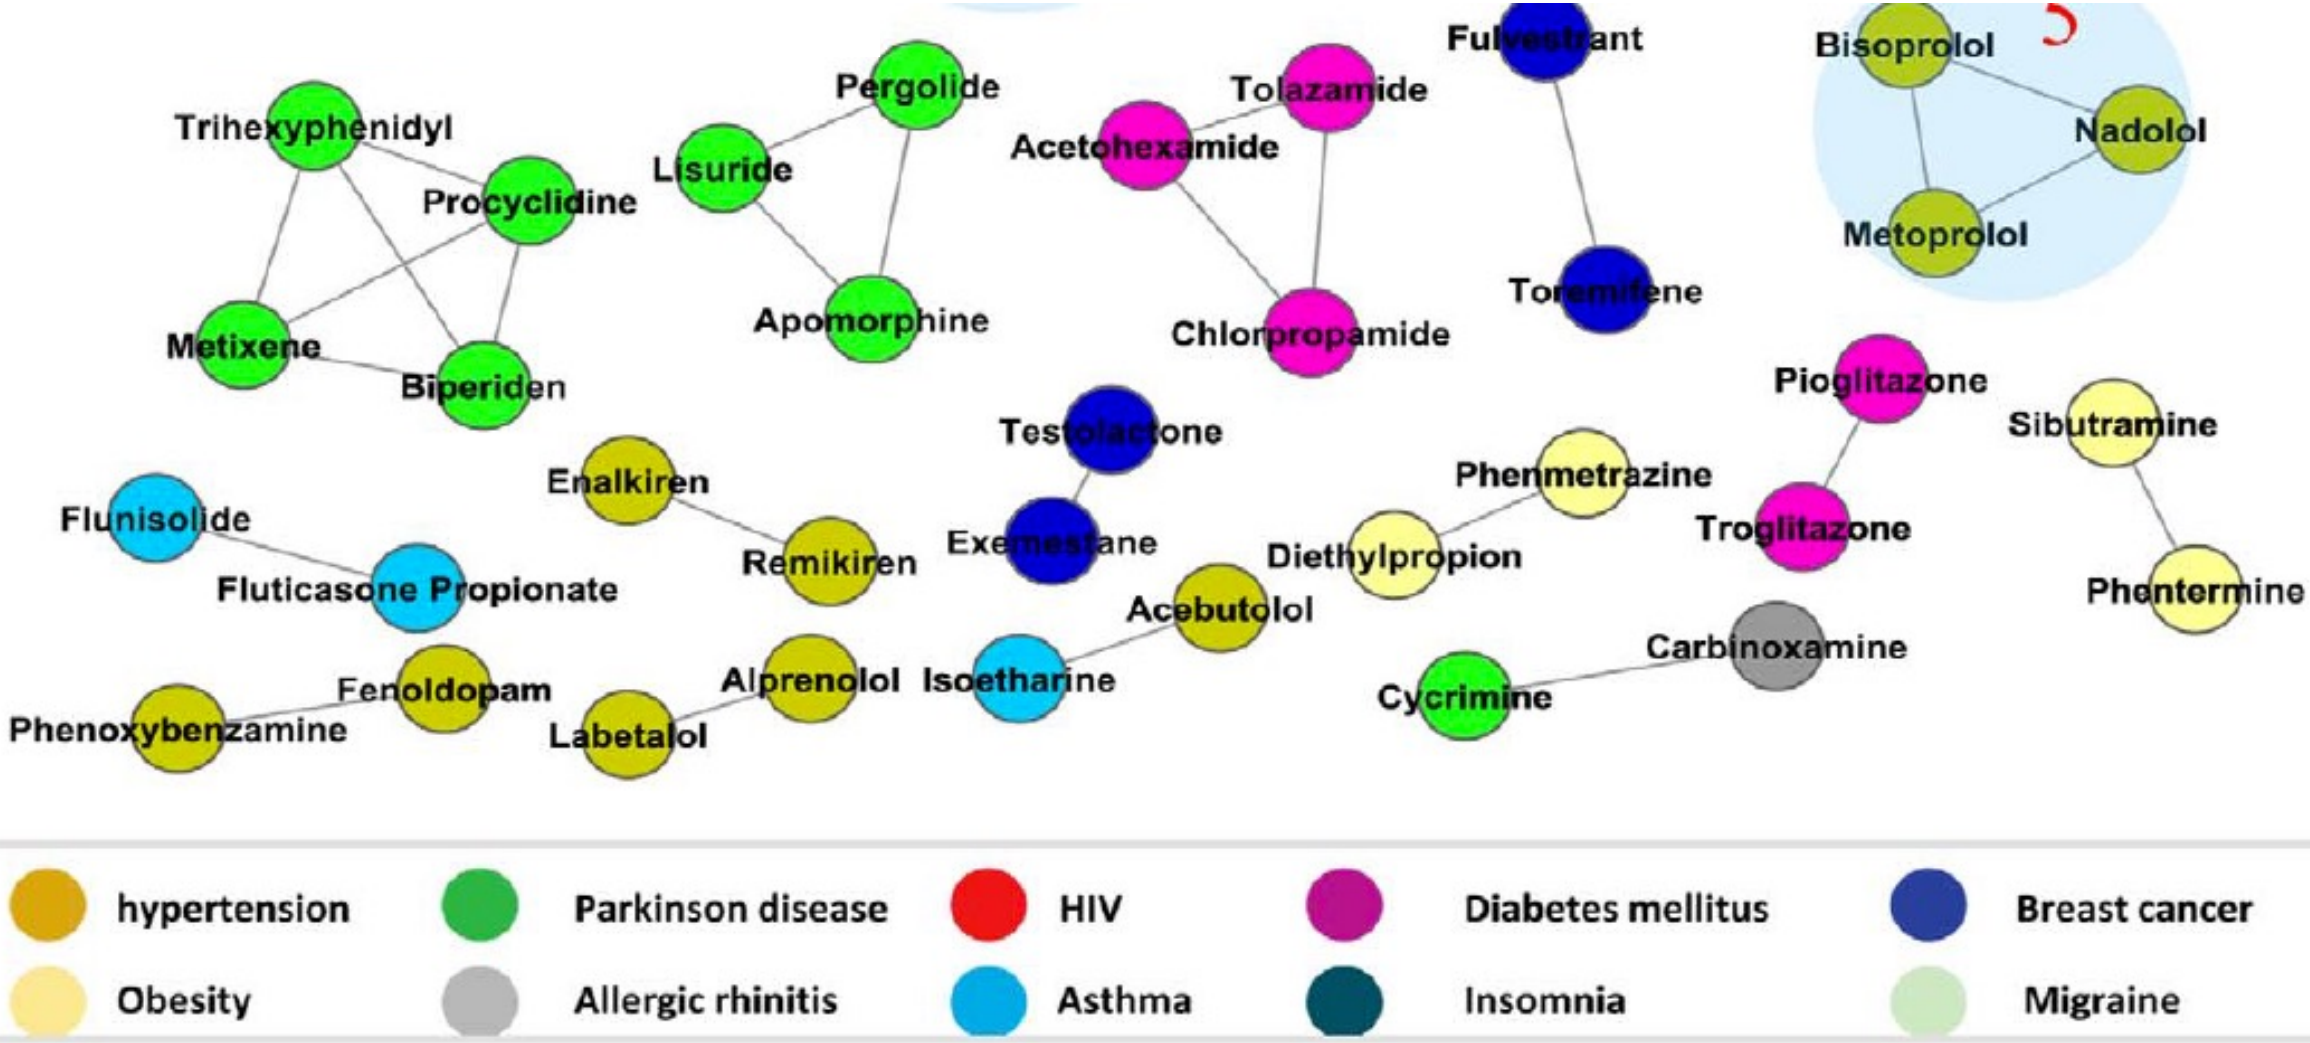
Participating entities

1. Patient
2. Drug
3. Disease
4. Procedure
5. Side Effect

Patients take Drugs to cure Diseases, but may also have side Effects.

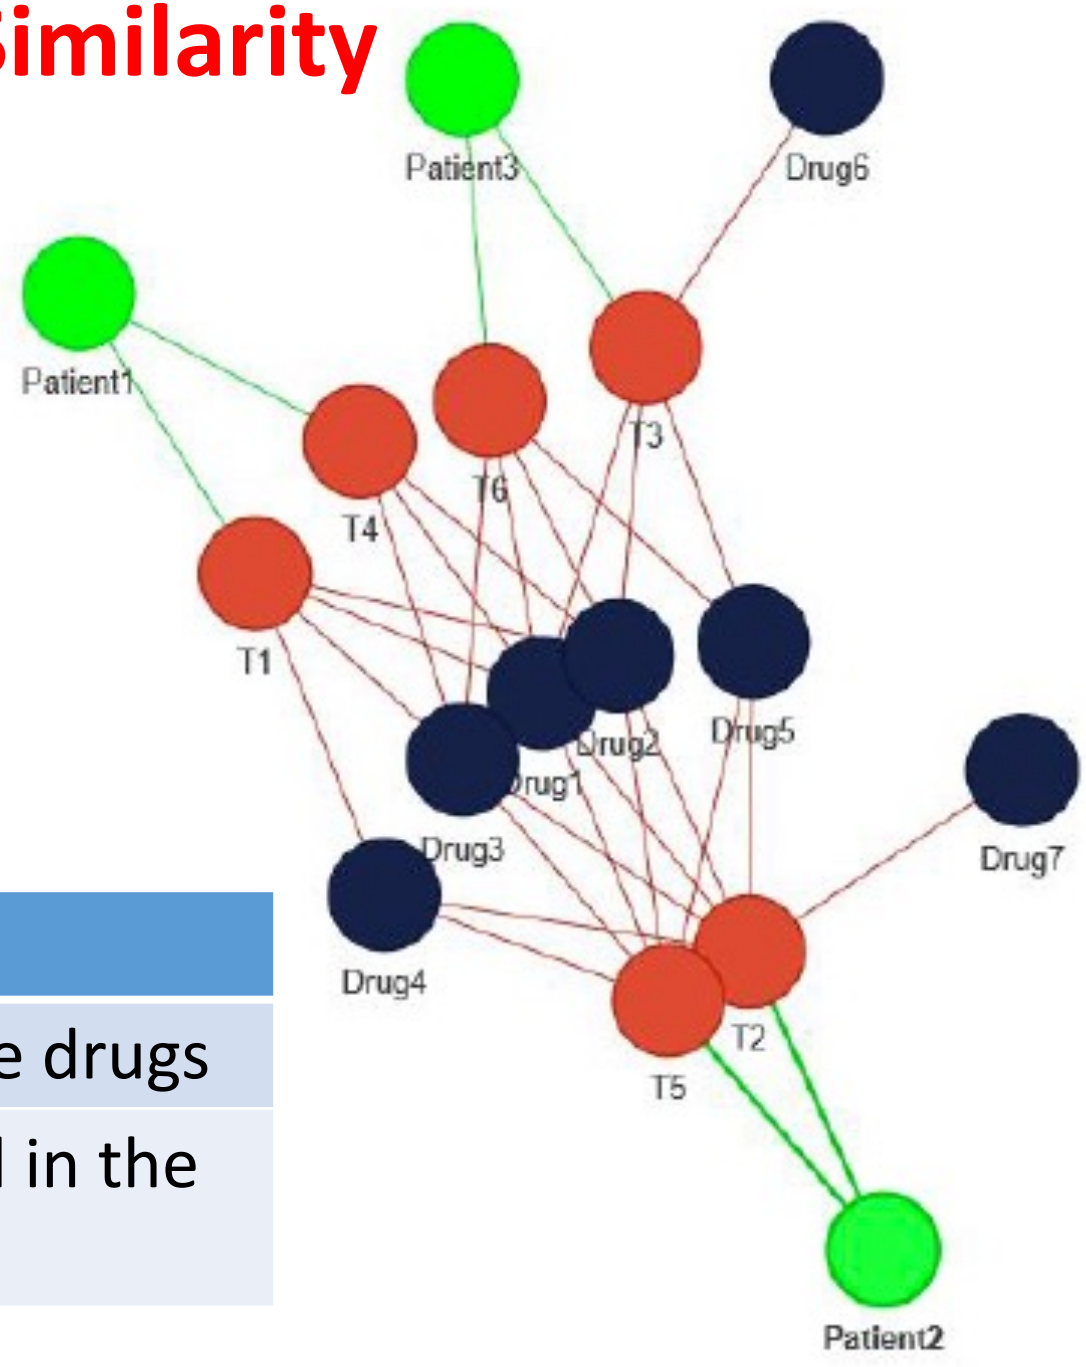
Similarity Search and Relevance in Graphs

- Different graph-based algorithms (SimRank, RWR, etc.) can be used to infer Drug-Drug Similarity Network.



Meta Path-based Similarity

- A meta path is a sequence of relations between different node types in a graph, which carries a specific semantic meaning



Meta Path	Semantic Meaning
PTDTP	Patients who took the same drugs
DTD	Drugs which are prescribed in the same treatment

Related Work

- Sun, Y., et al. **Pathsim: Meta path-based top-k similarity search in heterogeneous information networks**. (*VLDB* 2011)
  - Contribution: Defined meta paths on unweighted graphs

- Shi, Chuan, et al. ***Semantic path based personalized recommendation on weighted heterogeneous information networks***. *(CIKM 2015)*
  - Contribution: Defined meta paths on weighted graphs

- Shi, Chuan, et al. **SemRec: a personalized semantic recommendation method based on weighted heterogeneous information networks** *(WWW journal 2019)*.
  - Contribution: Combines multiple meta paths to infer similarity

The PathSim Similarity Measure

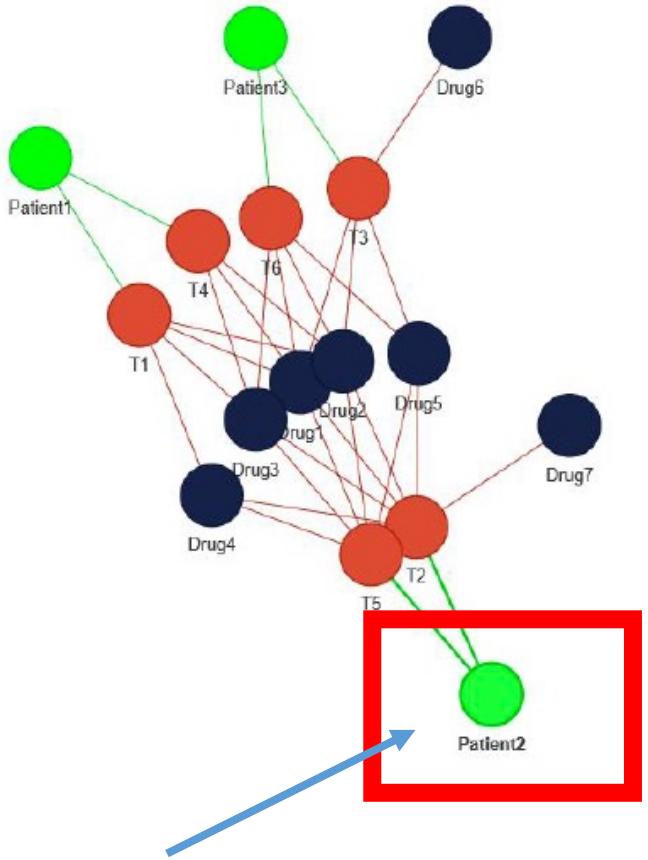
Definition 3. ***PathSim: A Single Meta path-based similarity measure*** *[22]. Given a symmetric meta path P, PathSim between two objects of the same type x and y is*

$$s(x,y) = \frac{2 * |p_{x \rightsquigarrow y} : p_{x \rightsquigarrow y} \in P|}{|p_{x \rightsquigarrow x} : p_{x \rightsquigarrow x} \in P| + |p_{y \rightsquigarrow y} : p_{y \rightsquigarrow y} \in P|}, \quad (1)$$

where $p_{x \rightsquigarrow y}$ is a path instance between x and y, $p_{x \rightsquigarrow x}$ is that between x and x, and $p_{y \rightsquigarrow y}$ is that between y and y.

- **The *intuition* behind it is that two similar peer objects should not only be strongly connected, but also share comparable visibility.**

Our proposed meta path-based Explanations



Which drug should we recommend for patient 2?

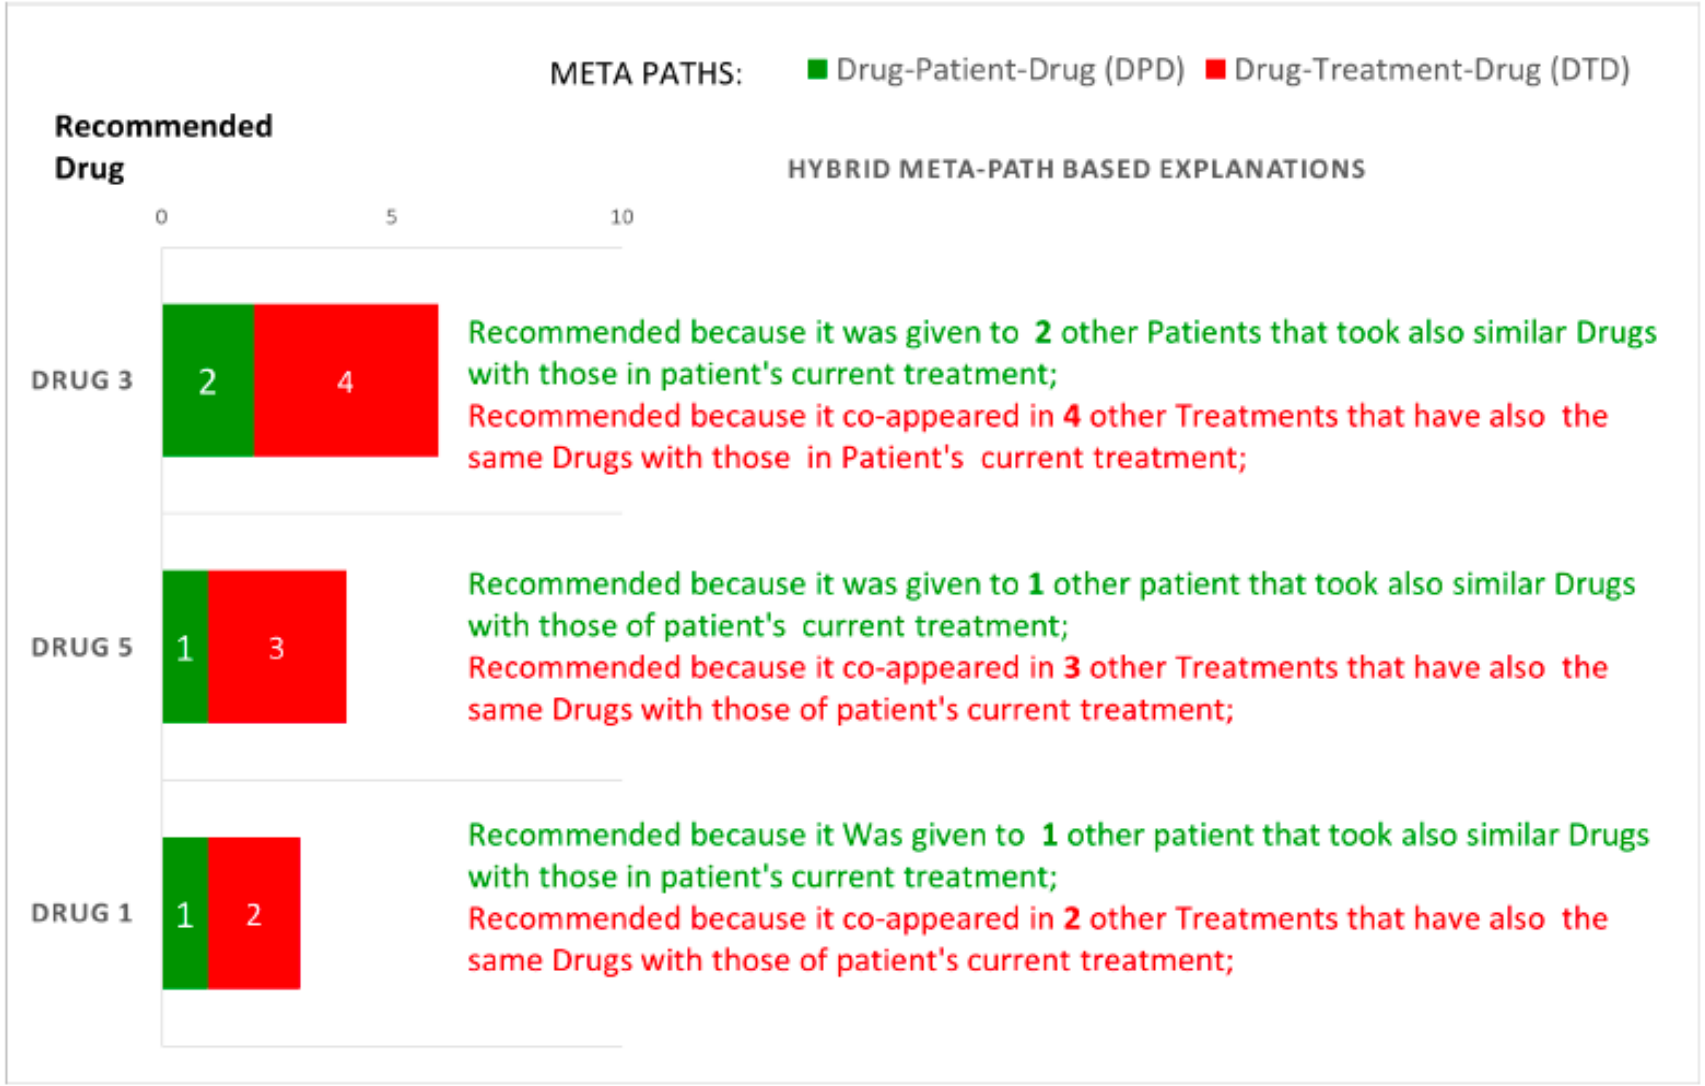


Fig. 4: Drug Recommendations along with explanation for patient 2. Drug 3 is recommended because it has the most frequent meta paths supporting it.

Experimental results in terms of Precision and DDI rate

Model	Cancer Data Set			
	Prec.	Recall	DDI rate	DDI Change %
PathSim(EHR)	0.4062	0.6414	0.0014	0 %
PathSim(EHR + DDI)	0.3070	0.5610	0.0054	+215.714 %
CMF(EHR)	0.1221	0.1008	0.0001	0 %
CMF(EHR+DDI)	0.0276	0.0344	0.0034	+3,300.000 %

Model	MIMIC III Data Set			
	Prec.	Recall	DDI rate	DDI Change %
PathSim(EHR)	0.3879	0.4457	0.4592	0 %
PathSim(EHR + DDI)	0.2976	0.3547	0.1734	-62.2386 %
CMF(EHR)	0.1261	0.2719	0.5656	0 %
CMF(EHR+DDI)	0.1146	0.2648	0.1154	-79.5968 %

Experimental results in terms of Explainability

$$Explain\ \ coverage(p, J) = \frac{\sum_{\forall (d, f_s^d) \in J} \min\{f_s^d, E_s(p, d)\}}{\sum_{\forall s \in S} E_s(p, d)}, \qquad (2)$$

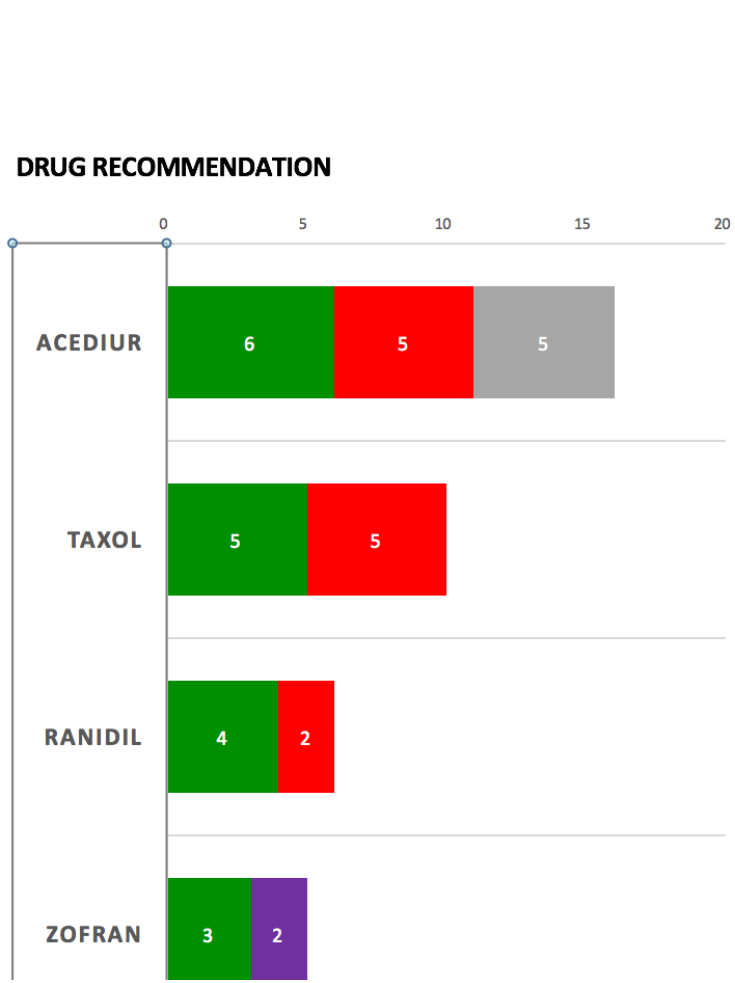
where the pair (d, f_s^d) denotes the overall frequency f of data source s with respect to drug d inside the hybrid explanation list J. Moreover, $E_s(p, d)$ is the frequency of d in the explainability matrix $E_s(p, d)$ of p over data source

Algorithms	Cancer Data Set	MIMIC III Data Set
	Expl. Cov.	Expl. Cov.
PathSim	**0.3300**	**0.4747**
CMF	0.0381	0.1307

Contributions

- We extract frequent patterns from a knowledge graph (KG) for Health
- We use PathSim to infer similarities among entities of the KG
- We explain adequately our drug recommendations based on meta paths

Meta path-based Drug Recommendations along with explanations



HYBRID EXPLANATIONS

Was given to **6** other Patients that took also similar Medicines with those in patient's current treatment;
Co-appeared in **5** other Treatments that have also the same Medicines with those in Patient's current treatment;
It cures similar diagnosed cancers with **5** medicines of patient's current treatment.

Was given to **5** other patients that took also similar medicines with those of patient's current treatment;
Co-appeared in **5** other Treatments that have also the same medicines with those of patient's current treatment;

Was given to **4** other patients that took also similar medicines with those in patient's current treatment;
Co-appeared in **2** other Treatments that have also the same medicines with those of patient's current treatment;

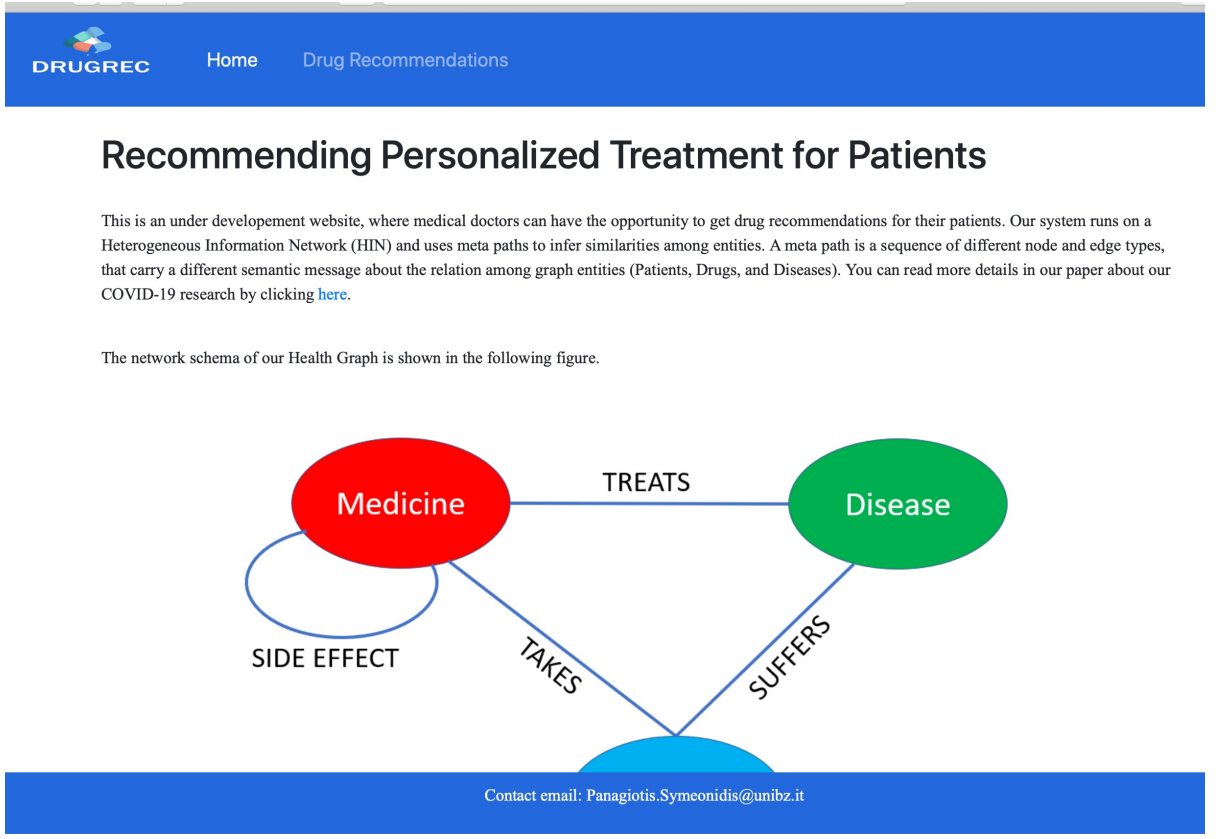
Was given to **3** other patients that took also similar medicines with those in patient's current treatment;

Download our demo notebook jupyter from https://github.com/hericonejito/health_recommendations

Check our DrugRec service for COVID patients

Our online system provides drug and treatment recommendations along with explanations for Covid Patients can be found:

http://drugrec.inf.unibz.it



Ongoing research: Alerting medical doctors over critical lab test measurements of patients

Parameters' Selection Lists of our Application

Groups of Subject IDs
- All Subject IDs
- Subject IDs Diagnosed with Diabetes
- Rest Subject IDs

Subject IDs
- 10017
- 10264
- 10424
- 10629
- 10633
- 10731
- 10806
- 10967
- 10976
- 11003

Chart Events Labels
- Glucose
- Glucose (70-105)
- Art.pH
- Arterial BP Mean
- Arterial BP [Diastolic]
- Arterial BP [Systolic]
- Arterial Base Excess
- Arterial CO2(Calc)
- Arterial PaCO2
- Arterial PaO2

Location
- HOSPITAL
- ICU_24H
- ICU_AFTER_24H

Age Range 5.0

Patient's Info

PATIENT INFORMATION

```
Gender   Female
Age      74 (hosp. entry)
Diabet.  YES
Death    AFTER ICU, 1199 days after ICU discharge
```

Diseases

DIAGNOSES

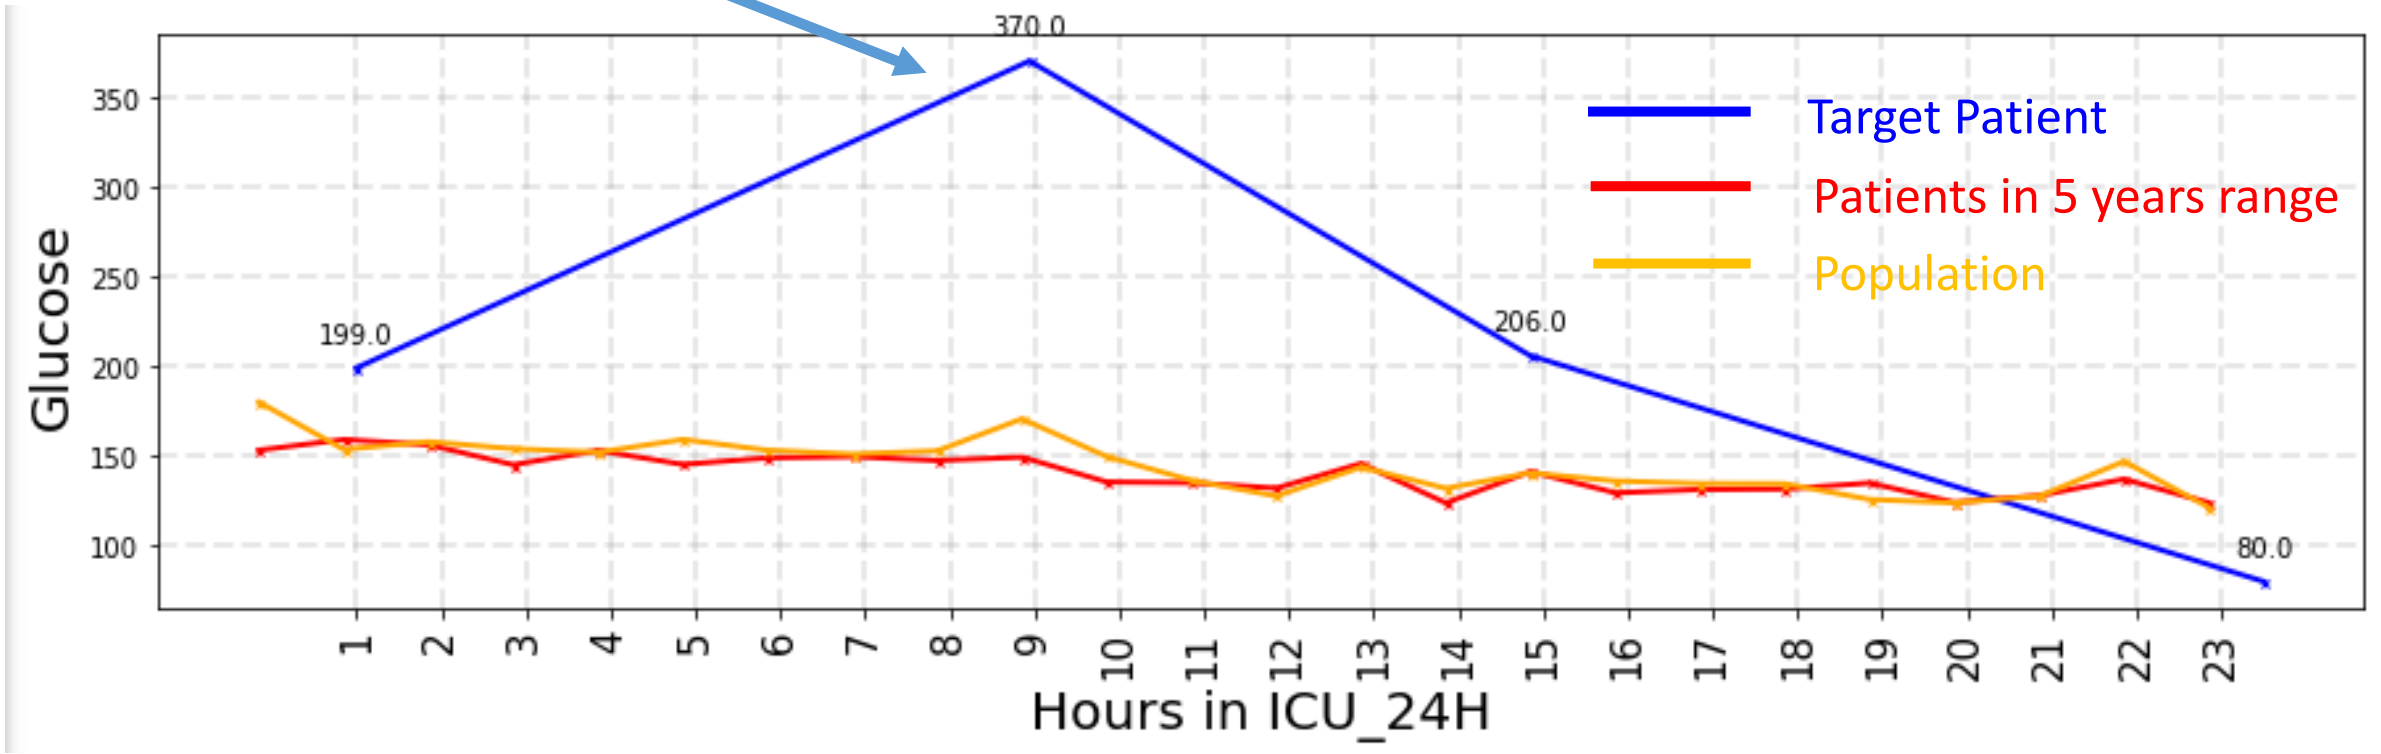
ICD9	DESCRIPTION
0380	Streptococcal septicemia
78552	Septic shock
5070	Pneumonitis due to inhalation of food or vomitus
5849	Acute kidney failure, unspecified
42731	Atrial fibrillation
1977	Malignant neoplasm of liver, secondary
1578	Malignant neoplasm of other specified sites of pancreas
2866	Defibrination syndrome
29590	Unspecified schizophrenia, unspecified
2764	Mixed acid-base balance disorder
99592	Severe sepsis
25070	Diabetes with peripheral circulatory disorders, type II or unspecified type, not stated as uncontrolled

PATIENT INFORMATION

```
Gender   Female
Age      74 (hosp. entry)
Diabet.  YES
Death    AFTER ICU, 1199 days after ICU discharge
```

Alerting doctors

Glucose measurements



PATIENT INFORMATION

```
Gender   Female
Age      74 (hosp. entry)
Diabet.  YES
Death    AFTER ICU, 1199 days after ICU discharge
```

Insulin doses

INPUTEVENTS

TIME	ITEM	AMOUNT	AMOUNT UNITS	RATE	RATE UNITS
9.3 Hours after ICU entry	Insulin			6	Uhr
9.3 Hours after ICU entry	Insulin	0	U		
9.8 Hours after ICU entry	Insulin	3	U		
9.8 Hours after ICU entry	Insulin			10	Uhr
10.8 Hours after ICU entry	Insulin	10	U		
10.8 Hours after ICU entry	Insulin			12	Uhr
11.8 Hours after ICU entry	Insulin			10	Uhr
11.8 Hours after ICU entry	Insulin	12	U		
12.8 Hours after ICU entry	Insulin			12	Uhr
12.8 Hours after ICU entry	Insulin	10	U		
13.8 Hours after ICU entry	Insulin	12	U		

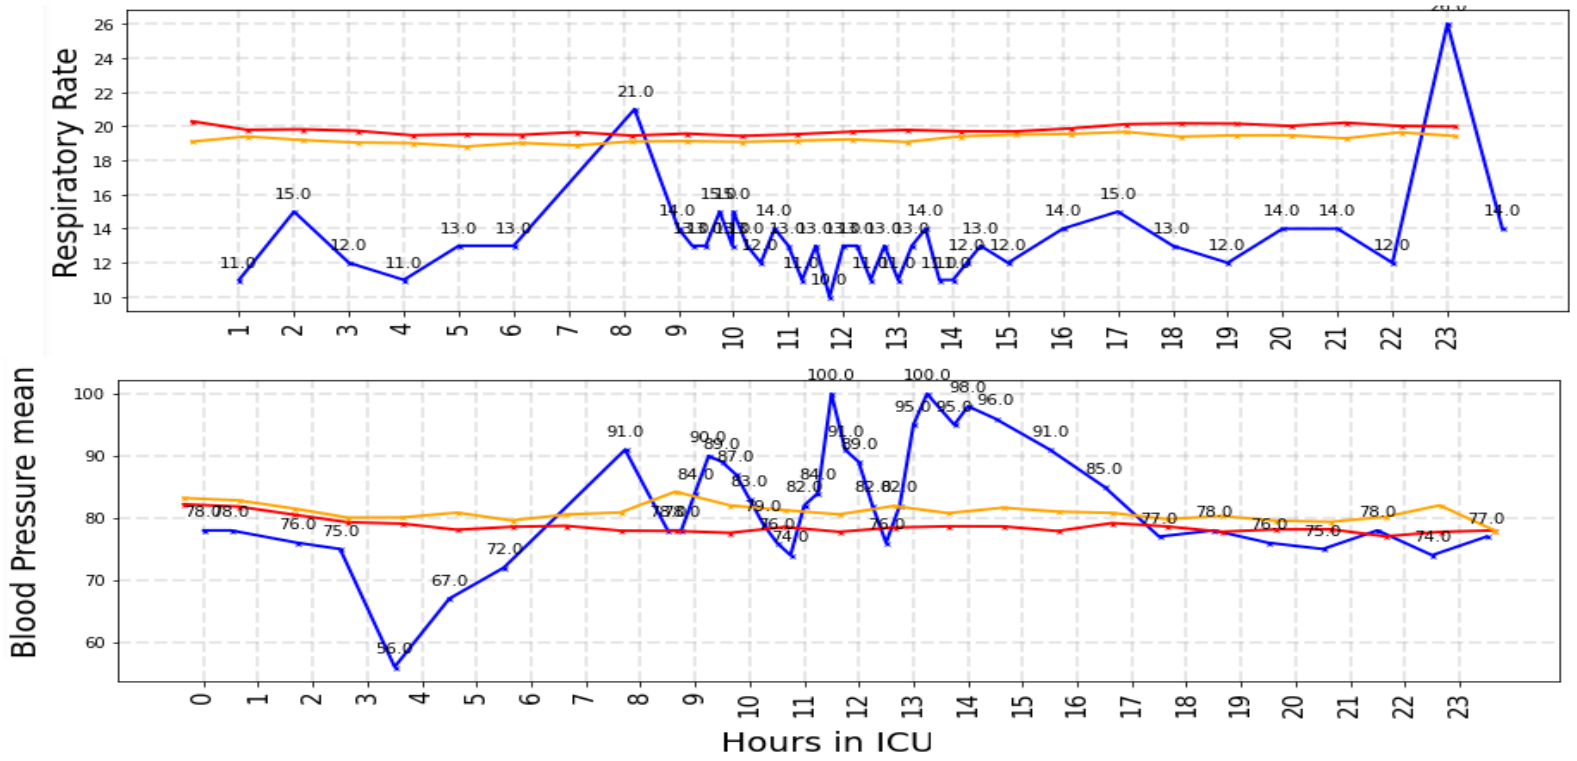
Another Example : Patient 83607

PATIENT INFORMATION

Gender Female
Age 49 (hosp. entry)
Diabet. YES
Death AFTER ICU, 92 days after ICU discharge

Diagnoses

40391	Hypertensive chronic kidney disease, unspecified, with chronic kidney disease stage V or end stage renal disease
7907	Bacteremia
5531	Umbilical hernia without mention of obstruction or gangrene
25042	Diabetes with renal manifestations, type II or unspecified type, uncontrolled
V4511	Renal dialysis status
0417	Pseudomonas infection in conditions classified elsewhere and of unspecified site

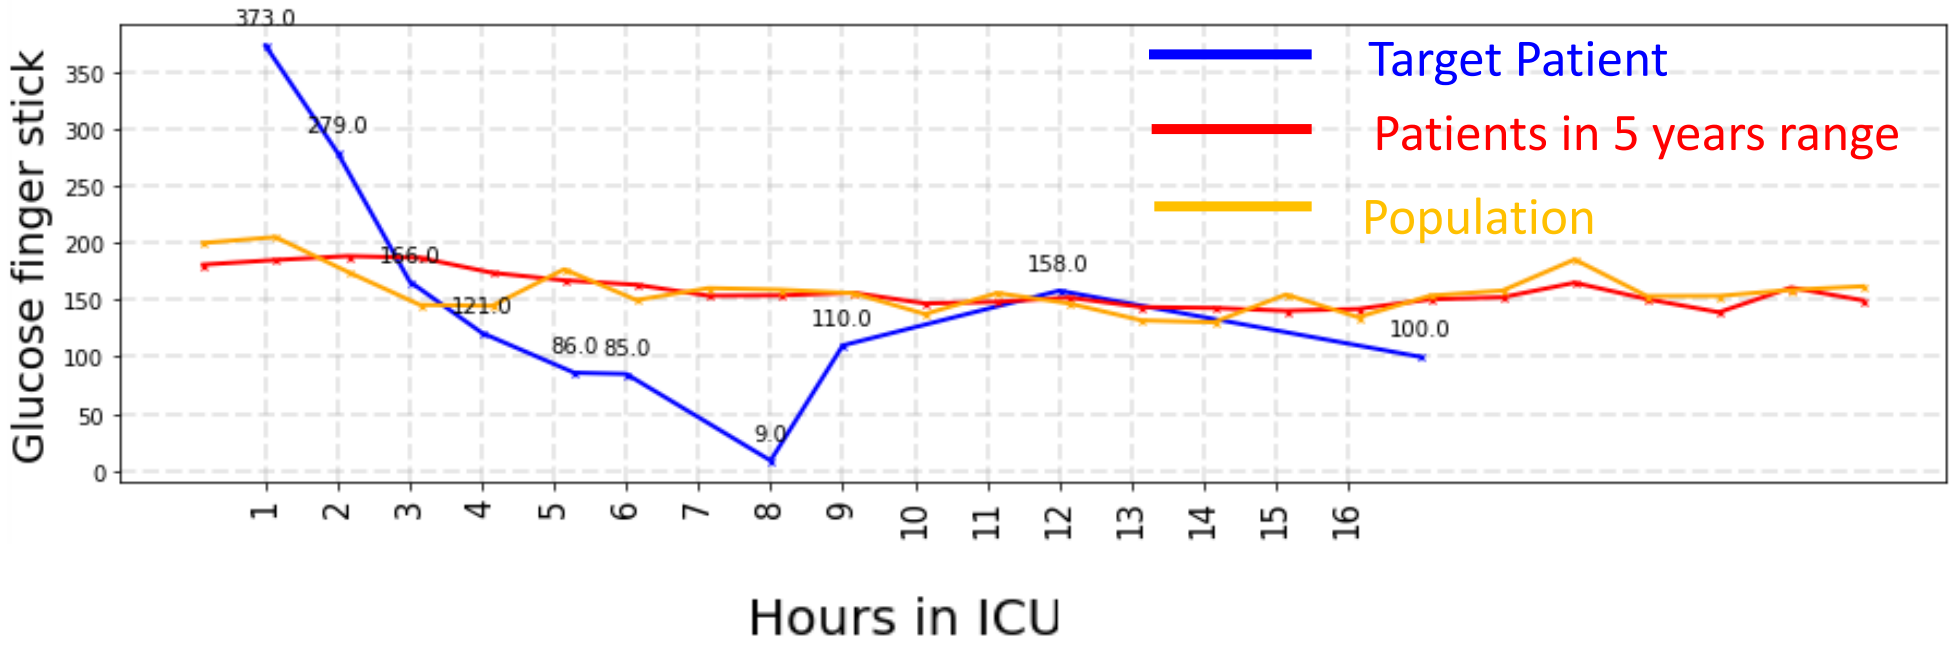


Target Patient
Patients in 5 years range
Population

Example Patient: 83607

INPUTEVENTS

TIME	ITEM	AMOUNT	AMOUNT UNITS	RATE	RATE UNITS	TABLE
0.0 Hours after ICU entry	Insulin - Regular	2.8	units	6	units/hour	MV
0.5 Hours after ICU entry	Insulin - Regular	7.66929756	units	10.0034316	units/hour	MV
1.3 Hours after ICU entry	Insulin - Regular	12.82340675	units	13.989171	units/hour	MV
2.2 Hours after ICU entry	Insulin - Regular	11.9855142	units	15.9806856	units/hour	MV
2.9 Hours after ICU entry	Insulin - Regular	25.28974176	units	9.9827928	units/hour	MV
5.5 Hours after ICU entry	Insulin - Regular	1.996558416	units	4.99139604	units/hour	MV
5.9 Hours after ICU entry	Insulin - Regular	7.73666349	units	2.99483748	units/hour	MV
11.9 Hours after ICU entry	Insulin - Regular	4	units			MV
12.1 Hours after ICU entry	Insulin - Glargine	7	units			MV



Ongoing work: Predicting Optimal Dosing of a Drug

Challenge for patients with diabetes

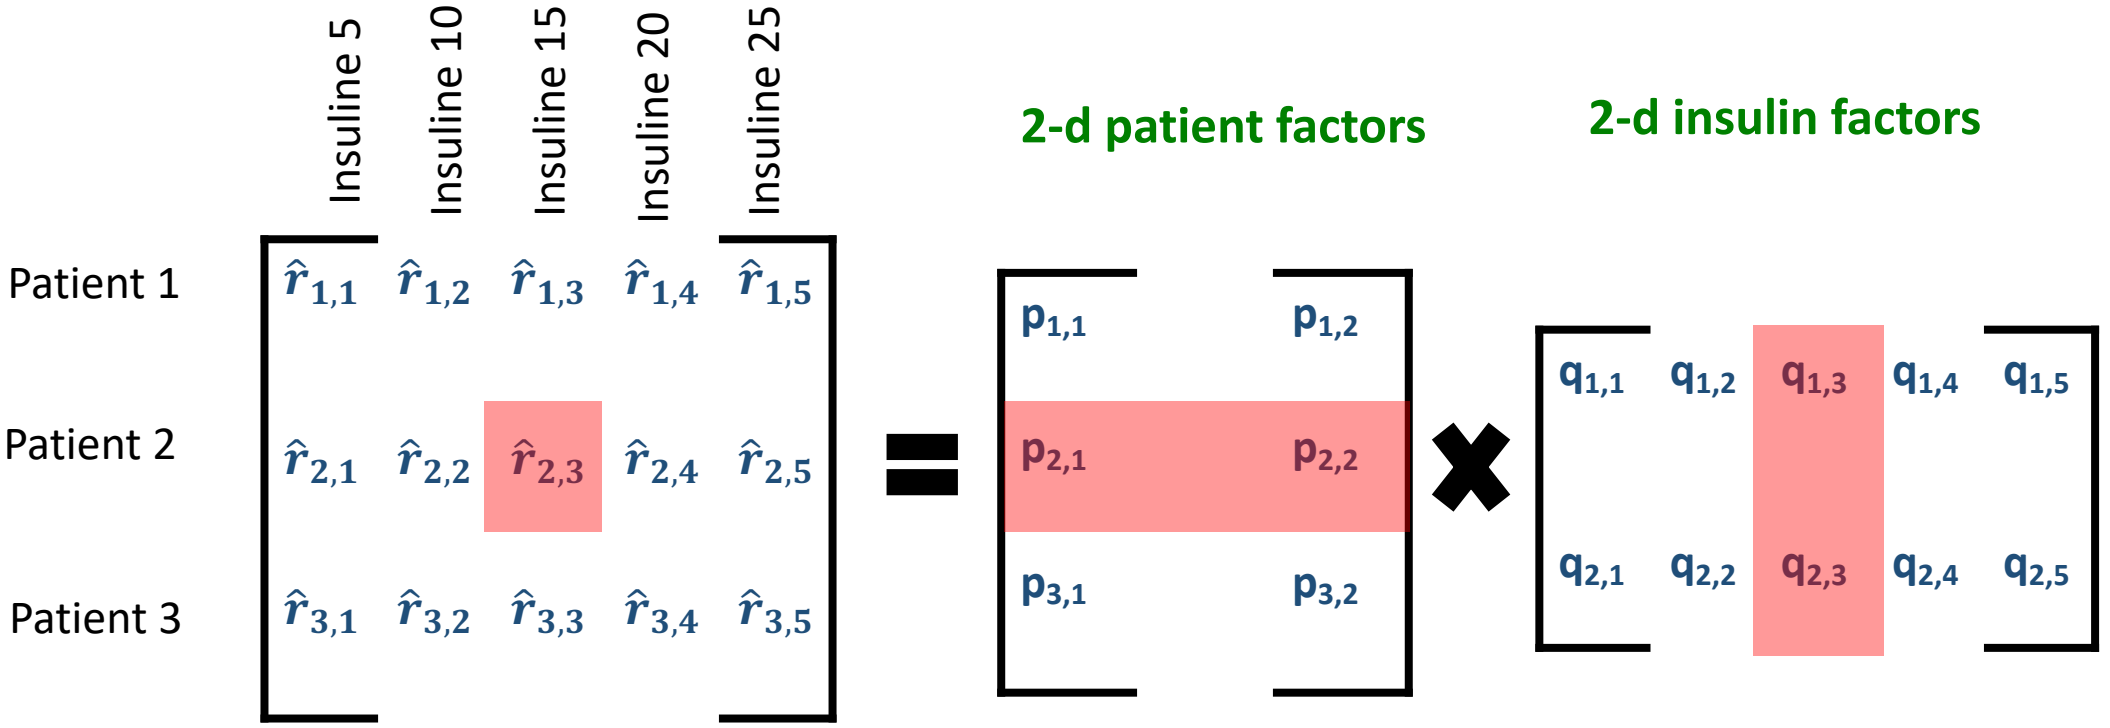
What is the optimal dosing of Insulin so that a diabetes patient does not have high glucose?

	Insulin 5 units	Insulin 10 units	Insulin 15 units	Insulin 20 units	Insulin 25 units
Patient 1	350 glucose	180 glucose	91 glucose (True)	60 glucose	0
Patient 2	150 glucose	100 glucose (True)	80 glucose	0	0
Patient 3	450 glucose	0	200 glucose	120 glucose	95 glucose (True)
Patient 4	0	120 glucose	97 glucose (true)	80 glucose	0
Patient 5	250 glucose	160 glucose	0	92 glucose (True)	0

This is a matrix completion problem!!!!

Matrix Factorization: Toy Example 1/2

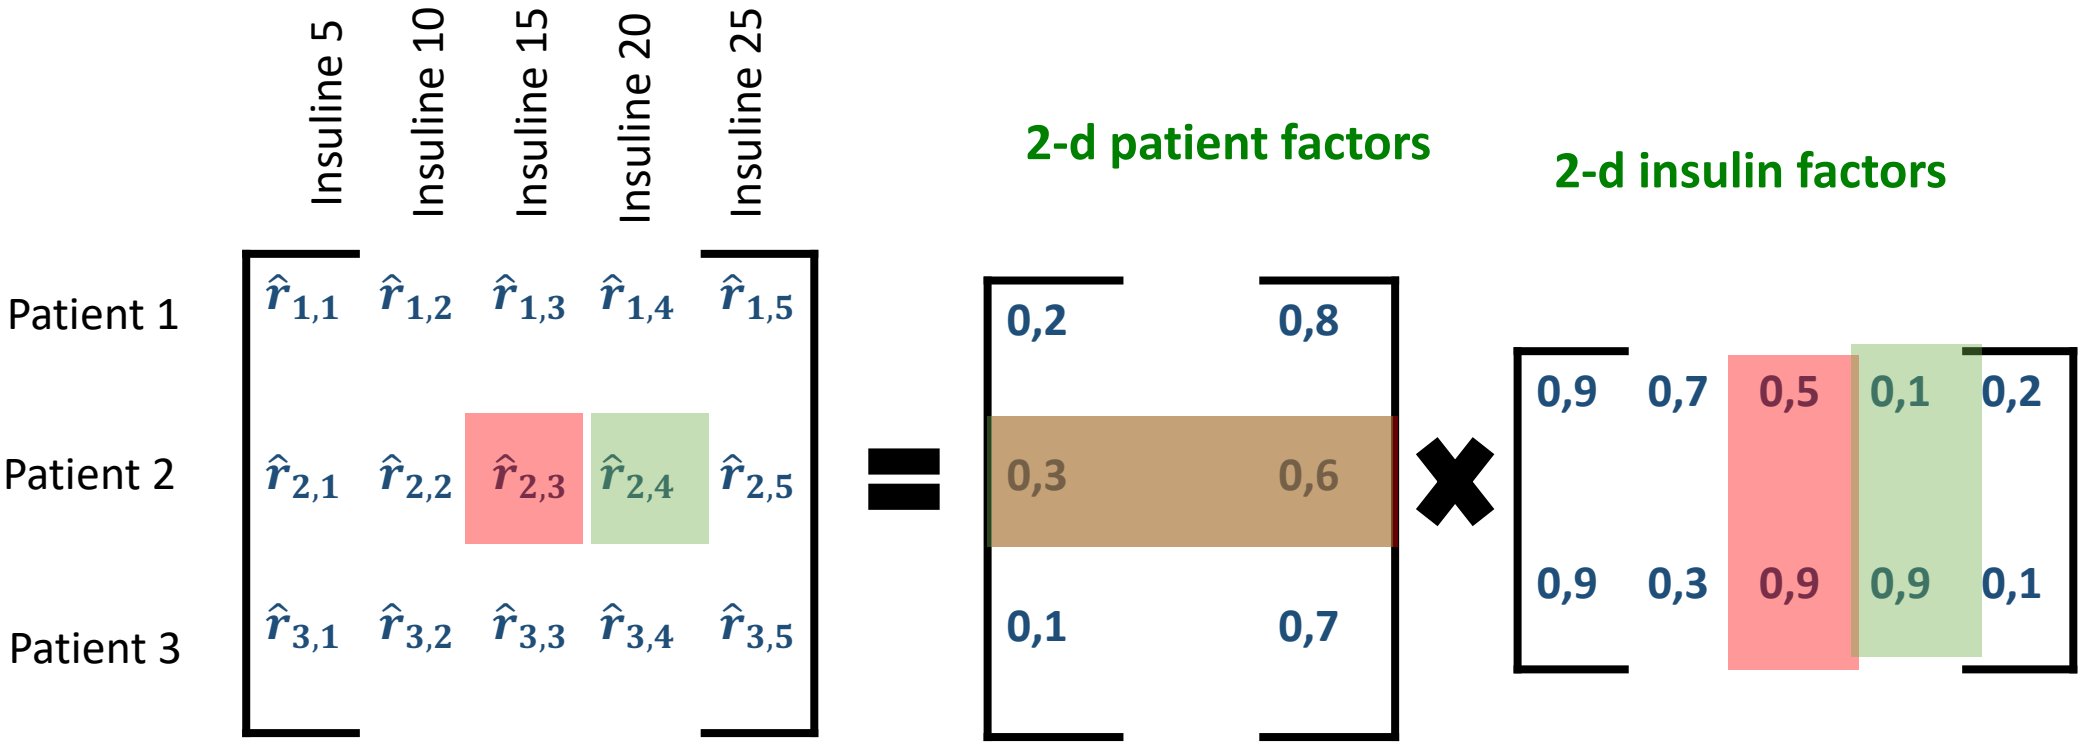
Patient-Insulin matrix with glucose values



$\hat{r}_{2,3} = p_{2,1} * q_{1,3} + p_{2,2} * q_{2,3}$ (Glucose prediction for Patient 2 on Insulin dose 15)

Toy Example 2/2

Patient-Insulin matrix with glucose values



$\hat{r}_{2,3}$ = 0,3 x 5 + 0,6 x 9 = 0,69

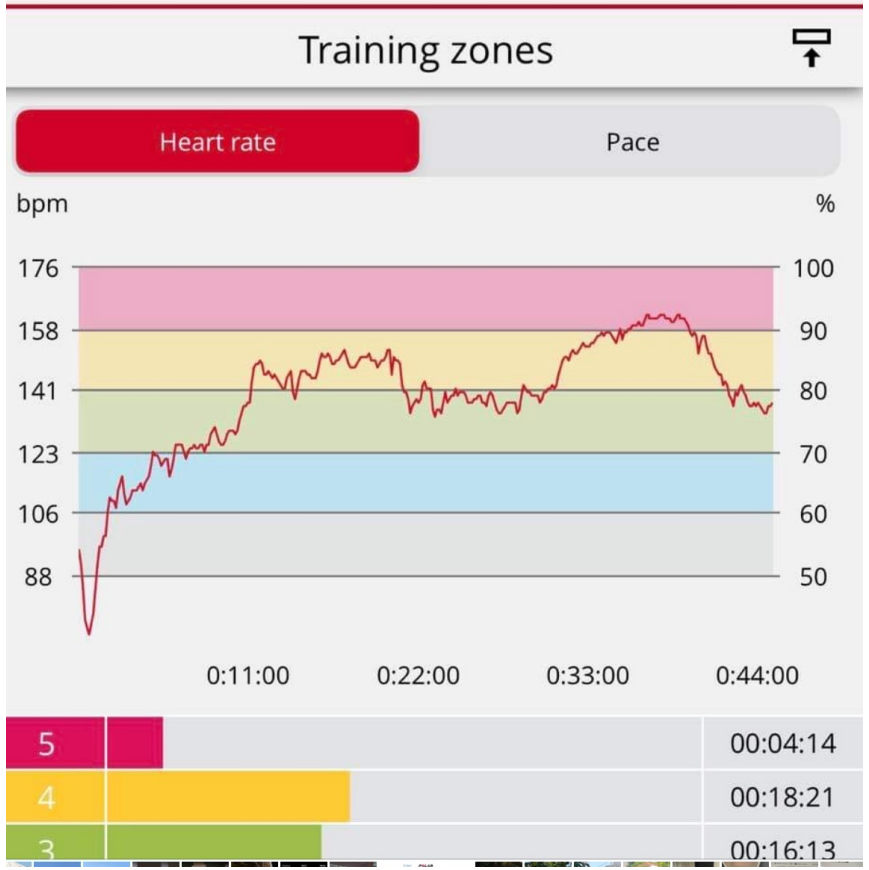
$\hat{r}_{2,4}$ = 0,3 x 0,1 + 0,6 x 0,9 = 0,57

Patient 2 should take insulin dose 15 instead of insulin dose 20.

Other possible research directions (1/2)

- Recommendations for the Physical Activity of patients through wearables and biosensors





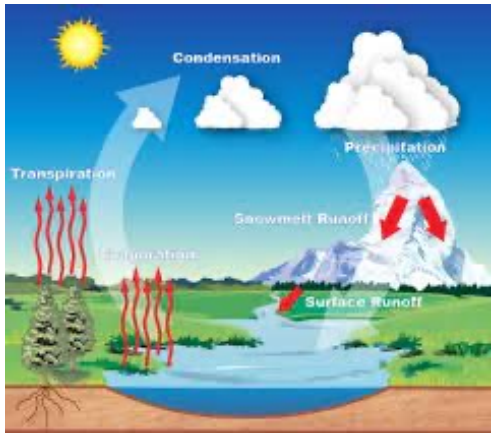
Zone	Effort	Target Heart Rate*	Training Benefit
ZONE 1	50% – 60%	-	Warmup / Recovery
ZONE 2	60% – 70%	-	Base Fitness
ZONE 3	70% – 80%	-	Aerobic Endurance
ZONE 4	80% – 90%	-	Anaerobic Capacity
ZONE 5	90% – 100%	-	Speed Training

Other possible research directions (2/2)

- Recommendations for the Food Choices of patients
  - (i.e., calories per meal and diet pattern)



- Incorporate environmental exposure (sun, humidity, etc.) into the training model of the recommender agent



Demonstration of an app to support medical doctors in alerting them over critical lab test measurements (Demo 1)

Proof of concept for optimal dosing of a drug (demo 2)

Explaining Drug Recommendations with Deep Learning

Wednesday 5/4/2023, University of Magdeburg

Thank you for your Attention! Any Questions?

UNIVERSITY OF THE AEGEAN



SCHOOL OF ENGINEERING

DEPARTMENT OF INFORMATION AND COMMUNICATION SYSTEMS ENGINEERING

Presenter: Panagiotis Symeonidis

Associate Professor

http://panagiotissymeonidis.com

psymeon@aegean.gr