

Towards Explainable General Medication Planning*

Lee-or Alon*, Hana Weitman, Alexander Shleyfman and Gal A. Kaminka

Computer Science Department, Bar-Ilan University, Israel

Abstract

The ability to produce explanations for automated systems in healthcare domains is crucial for establishing trust between users and the system. Despite the growing demand for explainable artificial intelligence in medical domains, to the best of our knowledge, there are no existing works on explainability for medication planning. In this work, we propose a visualization method for medication planning domains to make the automatic planning process transparent to users, thereby fostering the desired trust.

Keywords

Medication Planning, Explainable AI Planning, Personalized Medicine

1. Introduction

Personalized medication planning is the process of generating a plan of drug administrations that meets a given set of medical goals that are specific to the individual patient. The planning process must take into account general health safety constraints, helpful or harmful interactions between drugs, and individual physiological differences in responses to medications. The resulting personalized medication plan defines *what drugs* are administered, *when*, and at *what dosage*: too little is ineffective; too much is toxic.

Medication planning is a complex process, manually carried out by healthcare professionals. Its complexity is often encountered in mitigating harmful drug interactions in patients with multiple diseases [1], or in *combination therapy*, where multiple medications are used to synergistically improve therapeutic effects while minimizing side effects [2, 3]. Indeed, a combination of drugs can result in effects no drug can achieve alone [4].

To demonstrate the difficulty of solving medication planning problems, let us examine a relatively simple instance. Suppose the system is to consider only two types of medicine, both affecting a specific property of interest p . Assume that for both types, the time it takes the medicine to clear the body is 24 hours. Let us also assume there are five different available dosages for each medicine. The medical objective is to have property p reach a level of at least 51 in the spleen, but not exceed 53.8. While this problem seems small, solving it could be challenging due to the many possible treatment plans. Assuming each medicine is allowed to be administered at most once, there are $24 \cdot 24$ administration times and $5 \cdot 5$ dosage options, resulting in 14,400 combinations.

Not only are there numerous treatment combinations, but each administration's effects are

EXPLIMED-24: ECAI Workshop on Explainable Artificial Intelligence for the Medical Domain

*Corresponding author.

✉ alonlee1@biu.ac.il (L. Alon); weitman@cs.biu.ac.il (H. Weitman); alexash@biu.ac.il (A. Shleyfman); galk@cs.biu.ac.il (G. A. Kaminka)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

also non-linear and multidimensional, as they simultaneously impact several organs and change over time. Furthermore, because each patient may require a different treatment plan tailored to their specific needs and conditions, a new plan must be created for each individual. All of these characteristics make the problem even more challenging to compute manually.

To address the complexities of personalized medication planning, researchers have turned to artificial intelligence (AI) planning techniques. AI planning is a form of sequential decision-making over time. It aims to find an ordered set of actions that leads from an initial state to a goal state. In medication planning, the AI planning process aims to automatically find an ordered set of drug administrations that progresses from the current health condition of a patient to the desired medical condition. It must not violate any medical safety constraints along the way.

The wide variety of treatment combinations, as well as the non-linearity and multidimensionality of the treatment effects, make it challenging for users to comprehend the implications of a given plan. Moreover, as medication plans vary between patients, explanations or clarifications of each plan must be generated anew.

The complexity and sensitivity of medical treatment problems result in a high demand for explainability. Since mistakes in the treatment plan can have major irreversible consequences, including death, users are not willing to follow automatically generated treatment plans without understanding the reasoning and potential implications. Indeed, the use of explainable artificial intelligence planning (XAIP) [5] is important here. Its goal is to make the planners transparent to users, thereby establishing users trust.

We propose a framework for explainable general medication planning (XGMP). The approach we suggest visualizes the personalized treatment plan's effect on the patient's body, thus making the plan transparent and easy to understand, for healthcare professionals and even for non-professional users. We evaluated this approach using the general medication planning (GMP) representation by Alon et al., with medical data as reported in [6, 7]. To the best of our knowledge, no steps were taken towards explainable medication planning.

2. Background

Explainability is an emerging and crucial field in artificial intelligence (AI). Its primary objective is to make the automated decision-making processes of AI systems transparent and interpretable to users by providing explanations that elucidate the underlying reasoning behind the system's outputs. These explanations aim to foster trust and confidence in AI systems, enabling users to comprehend the rationale behind the decisions made by the system.

The demand for explainability becomes paramount in domains where human lives are at stake, such as healthcare and the legal system. In the healthcare domains, the consequences of erroneous decisions can be severe, e.g., misdiagnoses, and improper treatments. Consequently, users in this field are more likely to demand comprehensive explanations for the AI system's outputs, as the implications of incorrect decisions can be far-reaching and potentially life-altering.

Personalized medicine further increases the necessity for explainability. In personalized medicine, there is no single protocol to treat a certain condition. The treatment is tailored to the

individual patient based on their background medical conditions, preferences, and factors such as age. The variety of possible treatments under various conditions complicates the problem and requires different explanations, as the treatment for one patient may not be relevant for another, even if the same medical goal is required.

We begin with a short description of the medication planning process and basic terminology. We then discuss the need for explainability in this domain.

2.1. Medication Planning

The *general medication planning* problem (GMP) is concerned with selecting drugs to be administered, as well as determining the dosage and schedule of the chosen medications [8, 9, 10]. It is therefore, in some aspects at least, a generalization of dosing regimen planning [11], which assumes drugs have already been selected, and deals with administration dosages and schedules.

Medication planning is also closely related to the process of planning treatments for patients with multiple diseases, by merging available multiple single-disease clinical guidelines. The latter process includes substituting drugs when adverse or redundant interactions occur, adjusting and scheduling tests to monitor for such interactions, and other related tasks [12, 13, 14, 15, 16, 17, 18]. This process produces plans that span weeks or months, and involve selecting drugs from the set recommended by guidelines. It may also recommend medical tests, and actions to take per their results. In contrast, GMP is carried out from first principles, personalizing the dosage and hourly medication schedules, using models of how medicines (drugs) spread through the body and interact with it, and with each other (pharmacokinetic and pharmacodynamic models—see below). However, it does not address testing, or chronic conditions.

Pharmacokinetics and Pharmacodynamics Medication planning is a model-based planning approach. It uses models that predict how drugs spread in the body (pharmacokinetic models), and how they interact with in (pharmacodynamic models). These are explained below.

Once a drug is introduced into the body, it is generally absorbed, carried and circulating by the bloodstream throughout the body. The drug reaches various biological sites (*bio-sites*) and may accumulate for some time, before it is eventually cleared out of the body. The concentration of a drug in various bio-sites, known as its *biodistribution*, undergoes changes over time, which can be described by *pharmacokinetic* (PK) models of varying complexity. These models range from simple 1-3 compartment exponential decay models [19, 20] to more advanced models that account separately for multiple kinetic processes (see [21]). Alternatively, biodistribution trajectories can also be represented by explicit curves [22, 6], obtained from clinical trials.

For example, in Figure 1, we see the biodistribution trajectories of a specific drug administered to a mouse (nanoparticle #11, in [6]). Drug concentrations (percentage of initial dosage per gram of tissue) were measured in four bio-sites (*kidney, lung, spleen, liver*), at several time points (measured in hours since the administration at time $t_0 = 0$). Such trajectories change between medicines, but may also change between patients. The horizontal axis shows the time since administration. The vertical axis shows the concentration per gram of tissue as a percentage of the injection dosage. Each line shows the PK trajectory at a different bio-site.

When the drug reaches a target bio-site, it may affect the properties of that bio-site. These effects can be characterized using *pharmacodynamic* (PD) models [23]. PD models describe the




Figure 1: Biodistribution trajectories of nanoparticle #11 in mice (from [6]).




Figure 2: Illustration of the connection between pharmacokinetic models and pharmacodynamic models. Scheme was taken from [23].

relationship between the drug concentration at any given bio-site, or the body as a whole in simple models, and the resulting therapeutic effect.

PK and PD models are combined to form *PKPD* models [23], which predict the expected magnitude of drug therapeutic effect over time. Figure 2 illustrates the connection between PK and PD models. The PK model yields the concentration of the administered drug in the body at a specific time post-administration. Subsequently, the PD model utilizes this data to calculate the biochemical effect of that drug on the patient's body. PKPD models (and their component models) have been and continue to be an active area of research in medicine and pharmacology, with entire journals devoted to their investigation.

Goals of Medication Planning Medication planning involves medical goals that are specified in terms of properties of different bio-sites (or the body taken as a whole), taking into account temporal pharmaceutical dynamics and kinematics. It combines information about the rate of

accumulation and clearance of drugs in different bio-sites (from PKPD models) with information about toxicity and personal health constraints and patient activities to meet target levels of the drug or its biological effects. The process then *selects* drugs, *determines their dosage*, and the *schedule* of their administrations to a patient.

2.2. Explainable Medication Planning

Alon et al. [9, 10] presented the general personalized medication planning, and used a planning representation using PDDL+ [24] to plan using multiple drugs, affecting multiple bio-sites over time (e.g., as in Fig. 1). This representation allows for an arbitrary number of medicines, each may be administered repeatedly if needed. The interactions of the drugs are modeled, so that the planner can avoid harmful interactions, and replace one drug with another (or with a combination of drugs). This extends the work of Alaboud et al. [8], which introduced the use of automated (AI) planning to address medication planning for maintenance goals, of a single drug and its associated PK model. Unfortunately, these investigations has not addressed the need for explaining the resulting plans.

There are various aspects of explainable artificial intelligence planning (XAIP), all of which are missing in medication planning. One common XAIP technique is to intervene in the problem representation, forcing the planner to execute the user’s suggestions [25, 26]. The user iteratively asks questions, where each question yields a new plan. By comparing the original plan with the new plan derived from the user questions, users infer the reasons that led to the original plan. An XAIP investigation related to a medical domain is presented by Korikov et al. [27], which considers the appointment scheduling problem. They use a counterfactual explanation technique that explains to the user what should have been different in order to achieve the user’s suggested outcome.

In this paper, we focus on visualization, as the basis for interaction with a user. In general, visualization methods can also be utilized to explain the planner choices. Chakraborti et al. [28] introduce a visualization of the top- k plans as a graph where nodes represent actions and edges represent the transitions between actions. This visualization does not allow visualizing durative actions (whose effects change over time), or actions taken simultaneously. Similarly, Kumar et al. [29] present a visualization system for classical planning domains, i.e., all variables have binary values. They allow for both changes in the domain (in action structures) and in the problem (in initial state and goal states). As medication planning is not carried out in classical planning domains (durative actions, simultaneous actions, constraints), it is not compatible with with the proposed visualization method.

3. The General Medication Planning Representation

We will briefly describe the GMP PDDL+ representation as was proposed by Alon et al. [9, 10].

A PDDL+ planning problem [24] can be described by the following tuple: $\langle \mathcal{V}, \mathcal{S}, s_0, \mathcal{C}, \mathcal{G}, \mathcal{A}, \hat{\mathcal{E}}, \hat{\mathcal{P}} \rangle$ where \mathcal{V} is a set of state variables either propositional or numeric, \mathcal{S} is a set of states, where each state is a complete assignment of values to all variables $v \in \mathcal{V}$, $s_0 \in \mathcal{S}$ is an initial state, \mathcal{C} is a set of constraints on possible assignments of values, and \mathcal{G} is a goal description (a set of conditions over variables). \mathcal{A} is a set of instantaneous actions

that change the values of variables when selected by the agent, and $\hat{\mathcal{E}}, \hat{\mathcal{P}}$ sets of events and processes (resp.) that change the values of variables instantaneously or overtime, outside of the control of the agent. A plan is a timed sequence of (parallel) actions, which starts from the initial state and reaches a goal state while not violating any constraint in \mathcal{C} .

From a medical perspective, a patient’s body can be viewed as a set B of bio-sites, such as organs and blood. Basic pharmacological models often depict the entire body as a single bio-site ($|B| = 1$), but in more complex models multiple bio-sites are represented.

Their work allows for the representation of B bio-sites, where each bio-site is represented as a set of P biochemical properties. Their values are numeric fluent indicating the concentration levels or other measures of interest and generally vary between bio-sites.

Table 1 shows an example of 12 property variables, used in the experiments. The property m_{11} (first row) measures the concentration of nanoparticle #11 (Fig. 1) in six different bio-sites, at a specific time (e.g., $liver[m_{11}] = 2.97$, and $kidney[m_{11}] = 9.2$). Data was taken from [6]. A second property, measuring the mu-opioid receptor (MOR) activity, is shown in the second row. Its values in this case are derived from PKPD model parameters reported elsewhere [7]. The *initial state* s_0 of a patient’s body may be represented by setting the values of properties, in each bio-site, to current values. For properties measuring drug concentration, initial values in all bio-sites are zero.

Properties P \ Organs B	Blood	Heart	Liver	Spleen	Lung	Kidney
m_{11}	1.6	0.73	2.97	2.34	1.81	9.2
MOR activity	26.4	20.28	30.003	28.79	27.27	33.55

Table 1

Illustration of the state of a patient’s body 24 hours post-administration. Columns represent bio-sites. Rows represent property values.

The administration of dosage d of a drug type m at time t is represented as a PDDL action. This representation allows for repetitive administrations of the same drug type as long as they are not administered at the same time. Drugs from different types may be taken in parallel.

Each drug administration may affect several bio-sites and several properties simultaneously. Note that several drug administrations may affect the same bio-site property simultaneously, even if these administrations were not at the same time, since administrations have durative effects.

The medicine effects over time are represented by PDDL+ events. For every medicine m , there is *at least* a single property m in every bio-site b , i.e., $b[m] \in \mathcal{V}$, which represents the concentration level of medicine m in bio-site b . In this representation, the medicine levels are estimated directly from biodistribution trajectories (e.g., Fig. 1) in the problem description in PDDL+.

As a drug is accumulated in a bio-site (measured by its concentration level), it causes changes in other biochemical properties within the same bio-site. These changes can be predicted using PD models. The combination of the PK and PD model types, known as a PKPD model, allows for the estimation of how the accumulation and clearance of a drug change in biochemical properties influence various bio-sites over time [19, 23, 30].

The PKPD drug effects are also represented using PDDL+ events. This representation utilizes the direct action (direct effect) model, a common PKPD model in medical literature [23, 30]. This model describes the relationship between the time-dependent concentration, and the effects of the drug, measured in relevant units varying between drugs.

Different drugs may affect the same property simultaneously. As these will be handled by different events, their effects will increase the value of the property according to the PKPD effect of the associated medicine type. This naturally follows the Loewe *additive* drug interaction model [31, 23], whereby drugs can affect the same property, but at different “strength”. Contra-indicated drugs (may not be taken together) are handled by constraints (see below).

Goals \mathcal{G} and Safety Constraints \mathcal{C} Given the definitions of states and actions above, it seems a simple matter to define goal states in terms of target levels for properties of interest, at a specific set of bio-sites (therapeutic sites). However, medically, the planner must also ensure that the levels of all properties are maintained at safe levels, *before* the target levels are reached, as well as *after*.

They use events to impose limits on the maximal and/or minimal values of a property at any moment. These limits can come from medical defaults, or they may be personalized for specific health conditions of a patient. For example, if a patient has diabetes, the glucose level must stay below a given threshold h at all times. Such a constraint on the property j of bio-site b can be expressed as $b[j] \bowtie h$, where $\bowtie \in \{>, \geq, =, \leq, <\}$. Constraints can be also placed to prevent interactions between drugs.

The goal description \mathcal{G} has two components in the PDDL+ representation of GMP. The first involves specifying target levels for properties in the set of therapeutic sites. These target levels can be personalized and differ between patients. The second component ensures that constraints are maintained after these target levels are achieved.

Once the goal conditions are first satisfied at time t_g , safety constraints should be upheld not only in the interval $[0, t_g]$ but also in the extended interval $[t_g, \infty)$, bearing in mind that action effects have finite durations. Thus, a second subgoal introduced using PDDL+ checks that all administered medication had been eliminated from the patient’s body *after* the first component has been achieved.

Personalization Patients, who seek treatment for the same goal, vary in their medical history (e.g., background medical conditions leading to differences in safety constraints) and treatment preferences (e.g., due to age, sex, levels of activity). Two patients with the same medical goal may still require different treatment plans due to differences in background medical conditions, such as diabetes, pregnancy, etc. Patient diversity may cause differences in their PKPD responses, both in biodistribution trajectories, as well as E_{max} and EC_{50} parameters.

4. Explainable General Medication Planning

Using clinical data from mice and rats, we implemented medication planning problems in PDDL+ as described. The PKPD models are taken from databases of possible nanoparticle-based drug

carriers [6] and pain relief drugs [7]. We use the ENHSP-20 numeric planner [32] to solve medication problems (see [10] for more details).

4.1. Visualizing the Plan PKPD Effects

We propose visualizing the PK and PD effects to help users understand the treatment's impact on the patient body and easily compare between plans.

Consider a medication planning problem where the goal is to achieve a minimal value of 44 for mu-opioid receptor (MOR) activity in the heart and a maximal value of 46 (i.e., there is a safety constraint of 46 for this property in the heart)¹. The planner suggested the following plan (denoted as plan A; dosage is measured in units of nM (nano-molar)):

```
0: (administer_med m9 a1 d100)
0: --waiting-- [2.0]
2.0: (administer_med m10 a1 d100)
2.0: --waiting-- [28.0]
```

This plan suggests administering a dosage 100nM of medicine *m9*, waiting two hours and then administering a dosage 100nM of medicine *m10*. Lastly, the planner waits another 26 hours until clearance occurs at time 28 from the beginning of the plan. While this very common plan format is detailed enough to follow, the exact effect of the plan on the body remains unclear. Understanding the pharmacokinetic (PK) and pharmacodynamic (PD) effects among users will increase their trust in the system and the suggested treatment and making the process transparent.

Figure 3 presents the visualization of the plan A (described above). The horizontal axis shows the time (in hours) since the start of the plan (first administered drug). The vertical axis shows MOR activity levels. The red line shows the PKPD effect of administering 100 nM of *m9* alone at time 0, while the blue line shows the PKPD effect of administering 100 nM of *m10* alone two hours after the plan start (i.e., had it been injected without *m9* being present). The horizontal line (light blue, a level of MOR activity of 46) describes the safety constraint. The total PKPD effect of the two drugs is shown in purple. All of these are presented for the heart: the curves would be different for other bio-sites.

The visualization also highlights limitations of the clinical data. The sudden drops to zero towards the end of the plan span (i.e., 23–25 hours) signify that the different drugs have cleared the bio-site and thus their effect is reduced. In reality, we expect such clearance to be more gradual. However, the planning process is restricted to using the actual clinical data given in [6], which provides measurements at a resolution of a few hours (i.e., approximately 3–10 data points, depending on the drug and bio-site). The planner uses linear interpolation for points in-between measured data. As the actual clearance time is not given, the planner extrapolates it from the last point given in the data, towards zero. The result is that the clearance time is arbitrarily set as an hour following the last measured data point available from the database.

¹MOR activity levels are associated with pain relief.

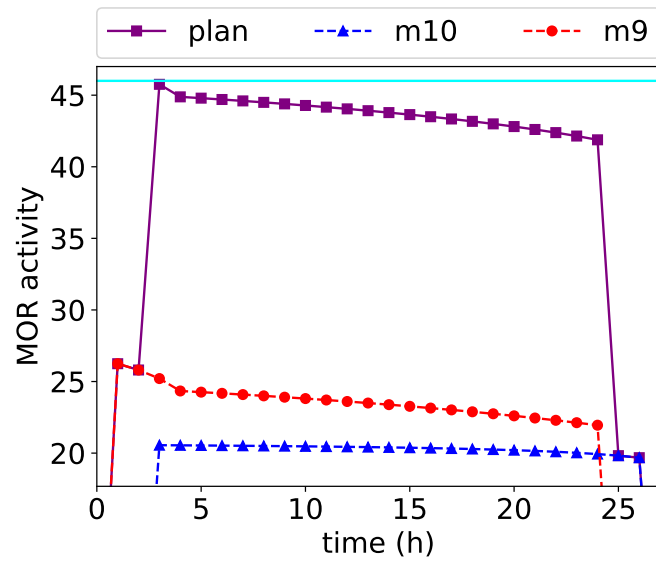


Figure 3: Plan A. The PKPD effect (in the heart) of administering medicine *m9* at time 0 with dosage 100nM (red) and medicine *m10* with dosage 100nM at time 2 (blue) and in combination (purple).

4.2. Visualize Alternative Plans or Changes in Medical Conditions

The PK and PD effects of an alternative plan can be visualized side-by-side with the original plan. This approach highlights the differences between the plans, allowing for a clear and easy comparison of their respective effects. Such differences may be due to variations in drug administration timing or dosages, or due to differences in medical conditions or considerations.

Comparing plans with alternative schedules for the same drugs The visualization can demonstrate differences in effects due to variations in administration timing and dosages. Consider the plan effect described in Figure 3. One might ask: "What would be the effect of the plan if both medicines were administered simultaneously?". The plan (denoted plan B) produced by the planner in this case is two hours shorter than plan A:

```
0: (administer_med m9 a1 d100)
0: (administer_med m10 a1 d100)
0: --waiting-- [26.0]
```

Figure 4 presents the PKPD effect of executing plan B. Contrasting it with plan A (Fig. 3), we indeed see that while plan B is two hours shorter than plan A, it violates the medical safety constraint.

Comparing Plans with Alternative Medical Considerations Personalization of the medical safety constraints, or other medical considerations, may also alter the plan. It may help the user visualize and contrast alternative treatment plans, resulting from such considerations.




Figure 4: The effect of executing plan B (purple).




Figure 5: The PKPD effect (in the heart) of administering 20nM of medicine $m5$ (orange) vs. the PKPD effect in the heart of administering 40nM of medicine $m11$ (green). The horizontal cyan line shows the safety constraint.

Consider the medication planning problem of achieving MOR activity of 29 in the heart, maintaining it lower than 37 in the same bio-site. The planner initially suggested administering 20nM of $m5$, which in our experiments is an opiod (specifically, *Morphine*). Due to its severe potential side-effects (one of which is long-term formation of dependence), a medical professional

may ask for an alternative that excludes Morphine. When we pose this to the planner, it indeed finds a plan using a different medicine (*m11*) at a dosage of 40nM.

Figure 5 clearly shows the differences between the two suggested treatment plans. While both plans achieve the medical goal without violating the safety constraint, the plan involving medicine *m5* has a shorter duration of 26 hours, compared to the plan using medicine *m11* (50 hours). Additionally, the *m5* plan has a lower maximal MOR activity in the heart compared to the *m11* plan. However, the *m11* plan achieves the desired medical goal (MOR activity of 27) more quickly.

5. Conclusion

We introduced a visualization method for the general medication planning (GMP), a relatively new and underexplored area in personalized medical treatment planning. The experiments conducted using real medical data showcased the effectiveness of this method in making unclear treatment plans easily understandable, even for non-professional users.

While the visualization method offers simplicity and clarity, it currently is not able to present other reason for using/not using certain medications, such as potential nausea, dizziness, weakness, and resistance to the administered medicines. Modeling such effects remains a future work.

Acknowledgements

Alexander Shleyfman's work was partially supported by ISF grant #2443/23. G.K. thanks K. Ushi, as always.

References

- [1] M. Dawes, Co-morbidity: we need a guideline for each patient not a guideline for each disease, *Family Practice* 27 (1).
- [2] S. K. Singh, A. Mohammed, O. A. Alghamdi, S. M. Husain, New approaches for targeting drug resistance through drug combination, in: *Combination Therapy Against Multidrug Resistance*, Academic Press, 221–246, 2020.
- [3] O. Turan, P. Bielecki, K. Tong, G. Covarrubias, T. Moon, A. Rahmy, S. Cooley, Y. Park, P. M. Peiris, K. B. Ghaghada, et al., Effect of dose and selection of two different ligands on the deposition and antitumor efficacy of targeted nanoparticles in brain tumors, *Molecular pharmaceutics* 16 (10) (2019) 4352–4360.
- [4] G. von Maltzahn, J.-H. Park, K. Y. Lin, N. Singh, C. Schwöppe, R. Mesters, W. E. Berdel, E. Ruoslahti, M. J. Sailor, S. N. Bhatia, Nanoparticles that communicate in vivo to amplify tumour targeting, *Nature materials* 10 (7) (2011) 545–552.
- [5] M. Fox, D. Long, D. Magazzeni, Explainable planning, arXiv preprint arXiv:1709.10256 .
- [6] M. Kumar, P. Kulkarni, S. Liu, N. Chemuturi, D. K. Shah, Nanoparticle biodistribution coefficients: A quantitative approach for understanding the tissue distribution of nanoparticles, *Advanced Drug Delivery Reviews* 194 (2023) 114708.

- [7] E. J. Santos, N. Nassehi, E. W. Bow, D. R. Chambers, E. S. Gutman, A. E. Jacobson, J. A. Lutz, S. A. Marsh, K. C. Rice, A. Sulima, et al., Role of efficacy as a determinant of locomotor activation by mu-opioid receptor (MOR) ligands in female and male mice. II. Effects of novel MOR-selective phenylmorphans with high-to-low MOR efficacy, *Pharmacology Research & Perspectives* 11 (4) (2023) e01111.
- [8] F. K. Alaboud, A. Coles, Personalized medication and activity planning in PDDL+, in: ICAPS, 492–500, 2019.
- [9] L. Alon, H. Weitman, G. A. Kaminka, First Steps Towards Planning for Targeted Medicine, in: KEPS, 2023.
- [10] L. Alon, H. Weitman, A. Shleyfman, G. A. Kaminka, Towards Personalized Medication Planning, in: KEPS, 2024.
- [11] F. Sime, M. Roberts, J. Roberts, Optimization of dosing regimens and dosing in special populations, *Clinical Microbiology and Infection* 21 (10) (2015) 886–893.
- [12] S. Wilk, W. Michalowski, M. Michalowski, K. Farion, M. M. Hing, S. Mohapatra, Mitigation of adverse interactions in pairs of clinical practice guidelines using constraint logic programming, *JBI* 46 (2) (2013) 341–353.
- [13] L. Piovesan, P. Terenziani, A Constraint-Based Approach for the Conciliation of Clinical Guidelines, in: IBERAMIA, vol. 10022 of *LNCS*, 77–88, 2016.
- [14] D. Riaño, A. Collado, Model-Based Combination of Treatments for the Management of Chronic Comorbid Patients, in: AIME, vol. 7885 of *LNCS*, 11–16, 2013.
- [15] L. Piovesan, P. Terenziani, A Mixed-Initiative Approach to the Conciliation of Clinical Guidelines for Comorbid Patients, in: AIME: Knowledge Representation for Health Care, vol. 9485 of *LNCS*, 95–108, 2015.
- [16] I. Sánchez-Garzón, J. Fdez-Olivares, E. Onaindía, G. Milla, J. Jordán, P. Castejón, A Multi-agent Planning Approach for the Generation of Personalized Treatment Plans of Comorbid Patients, in: AIME, 23–27, 2013.
- [17] J. Fdez-Olivares, E. Onaindia, L. Castillo, J. Jordán, J. Cózar, Personalized conciliation of clinical guidelines for comorbid patients through multi-agent planning, *AIME* 96 (2019) 167–186.
- [18] M. Michalowski, M. Rao, S. Wilk, W. Michalowski, M. Carrier, MitPlan 2.0: Enhanced Support for Multi-morbid Patient Management Using Planning, in: AIME, 276–286, 2021.
- [19] C. Hull, Pharmacokinetics and pharmacodynamics, *British Journal of Anaesthesia* 51 (7) (1979) 579–594.
- [20] P.-L. Toutain, A. Bousquet-mélou, Plasma terminal half-life, *Journal of veterinary pharmacology and therapeutics* 27 (6) (2004) 427–439.
- [21] L. E. Gerlowski, R. K. Jain, Physiologically-based pharmacokinetic modeling: principles and applications, *Journal of pharmaceutical sciences* 72 (10) (1983) 1103–1127.
- [22] S. Akhtar, Q. Khan, S. Anwar, G. Ali, M. Maqbool, M. Khan, S. Karim, L. Gao, A comparative study of the toxicity of polyethylene glycol-coated cobalt ferrite nanospheres and nanoparticles. *Nanoscale Res Lett* 14: 386, 2019.
- [23] D. F. Wright, H. R. Winter, S. B. Duffull, Understanding the time course of pharmacological effect: a PKPD approach, *British Journal of Clinical Pharmacology* 71 (6) (2011) 815–823.
- [24] M. Fox, D. Long, Modelling mixed discrete-continuous domains for planning, *JAIR* 27 (2006) 235–297.

- [25] B. Krarup, M. Cashmore, D. Magazzeni, T. Miller, Model-based contrastive explanations for explainable planning, in: 29th International Conference on Automated Planning and Scheduling-ICAPS 2019, 2019.
- [26] M. Cashmore, A. Collins, B. Krarup, S. Krivic, D. Magazzeni, D. Smith, Towards explainable AI planning as a service, arXiv preprint arXiv:1908.05059 .
- [27] A. Korikov, A. Shleyfman, C. Beck, Counterfactual explanations for optimization-based decisions in the context of the GDPR, in: ICAPS 2021 workshop on explainable AI planning, 2021.
- [28] T. Chakraborti, K. P. Fadnis, K. Talamadupula, M. Dholakia, B. Srivastava, J. O. Kephart, R. K. Bellamy, Visualizations for an explainable planning agent, arXiv preprint arXiv:1709.04517 .
- [29] A. Kumar, S. L. Vasileiou, M. Bancelhon, A. Ottley, W. Yeoh, Vizxp: A visualization framework for conveying explanations to users in model reconciliation problems, in: Proceedings of the International Conference on Automated Planning and Scheduling, vol. 32, 701–709, 2022.
- [30] M. A. Felmler, M. E. Morris, D. E. Mager, Mechanism-Based Pharmacodynamic Modeling, Humana Press, Totowa, NJ, ISBN 978-1-62703-050-2, 583–600, 2012.
- [31] M. C. Berenbaum, Synergy, additivism and antagonism in immunosuppression: a critical review, *Clinical and Experimental Immunology* 28 (1) (1977) 1–18.
- [32] E. Scala, A. Saetti, I. Serina, A. E. Gerevini, Search-Guidance Mechanisms for Numeric Planning Through Subgoal Relaxation, in: ICAPS, 226–234, 2020.