First Steps Towards Planning for Targeted Medicine

Lee-or Alon, Hana Weitman, Gal A. Kaminka

MAVERICK Group, Computer Science Department, Bar-Ilan University, Israel
alonlee1@biu.ac.il, weitman@cs.biu.ac.il , galk@cs.biu.ac.il

Abstract

Targeted nanomedicine involves the use of nanometer-scale devices that target specific tissues, while reducing side effects. It creates opportunities for ground-breaking applications in medical imaging, diagnosis, and treatment. These applications require careful planning of the type, dosage, and timing of administrating nano-devices, which cause cascading reactions, ending with a medically-desired result. Such planning is currently only carried out manually. We present a novel representation of targeted nanomedical treatment problems, related—but distinct—from familiar temporal planning approaches. We discuss preliminary steps towards planning and scheduling approaches for solving such problems. The formulation opens novel directions for the use of planning in targeted medicine.

Introduction

Targeted nanomedicine involves the use of nanometric-scale devices that target specific tissues, while significantly reducing side effects. It creates opportunities for ground-breaking, highly innovative applications in medical imaging, diagnosis, and treatment. Nanoparticles (NPs) are nanometric-scale material structures with a diameter of 1 to 100 nm. Each nanoparticle has unique properties and is designed to react with a predefined group of materials.

The usage of nanoparticles for medical applications has rapidly grown in recent years. The properties of different nanoparticle types are exploited for better detection and treatment of damaged organs. Using a single nanoparticle type might not be sufficient to achieve the desired medical goal, and so novel treatment plans provide schedules for nanoparticle injections. An example of such a requirement is combination therapy, where multiple treatments are used to mitigate a single disease. In addition, using various nanoparticle types can increase efficacy in comparison to the usage of each type alone (Turan et al. 2019).

Timing the administration of each nanoparticle type is necessary to maximize the benefits drawn from the nanoparticles. Many factors must be considered in order to increase the interactions between different nanoparticles, such as nanoparticle's circulation time, attaching probability, and half-life property. These factors depend on the physiochemical properties of the nanoparticles; their size, shape, charge, and surface property. Timing different nanoparticle types requires careful planning of administration times.

We propose a novel temporal planning representation that supports continuous non-linear actions for medical treatments and takes into account drug interactions. In this planning domain, states represent the amount of each nanoparticle type and material (e.g., drug) in different organs, and actions are the administration of nanoparticles to the patient. The domain's constraints are the safety constraints of a patient. We then discuss problem formulations using this representation. The formulation opens novel directions for the use of planning in medicine.

There are current planners and planning languages that support continuous actions and may be a basis for representing this domain (Fox and Long 2006; Scala et al. 2016; Piotrowski et al. 2016). For this domain, we require the ability to model parallel actions with concurrent non-linear effects that may affect the same state variable. We are currently examining the usage of existing planners and planning languages to support this domain.

Targeted Nanomedicine: An Opportunity for AI Planning

Targeted nanomedicine involves the use of nanometer-scale devices, often referred to as nanoparticles[1] which are designed such that by themselves, or in combination with other nanodevices, cause desired reactions in specific target tissues at the therapeutic site in the body (a biological site of interest, for example, a tumor, or a specific organ), while reducing side effects. This is the distinction from non-targeted medicine, where the administered compounds spread in the body without restricting the effect to specific tissues. As an example, most older forms of chemotherapy cause damage to healthy and cancerous cells alike, while newer *targeted* chemotherapy destroys cancer cells preferentially, and thus can be more aggressive in terms of the toxicity of the compounds administered to the cancerous tissues.

We believe that the design of targeted nanomedicine treatments presents high-impact opportunities for AI planning.

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

[1] and varieties such as nanospheres, nanorods, nanorobots/nanobots, etc.

The current process is completely manual, and is very slow to produce meaningful outcomes. In this section, we explain the fundamental concepts and terminology of targeted nanomedicine, and then argue for its suitability for automation by AI planning. The next section will formally define targeted nanomedicine planning processes.

Biodistributions of Medical Nanoparticles

Nanoparticles have many medical applications, such as cancer treatment, diabetes, infections, and cosmetics (von Maltzahn et al. 2011; Souto et al. 2019; Kirtane et al. 2021; Raj et al. 2012). Some nanoparticles are designed so that their properties (geometric size and shape, electric charge, and their chemical composition) can allow targeting different tissues in the body to achieve various medical goals.

When nanoparticles are introduced into the body (e.g., by an injection or oral), they are generally carried by the bloodstream, completing a circulation cycle approximately three times a minute. The circulation time of nanoparticles in the bloodstream is prolonged compared to conventional drugs, since they have good solubility in the plasma and due to their compatibility with the immune system. As a result, their clearance is decreased and accumulation in different tissues (including the therapeutic site) is improved, thereby increasing therapeutic efficacy and reducing side effects (Cho et al. 2008). They may penetrate to a specific organ that the blood vessels reach, as a function of their biochemical affinity to the specific organ. Thus, over time, nanoparticles can accumulate in different tissues in the body, and may also be cleared from the organs and the body. In general, all nanoparticles disintegrate at some point and are cleared from the body.

We refer to all possible biological sites (tissues, organs, tumor) as *bio-sites*. One or more of these—the therapeutic sites—are the target of the treatment, while others are exposed—and may accumulate—nanoparticles as a side effect. The time-changing distribution of the nanoparticles across different bio-sites is referred to as their **biodistribution**. It can be represented by a set of trajectories: given a time since the injection, each trajectory is defined by a function that gives the level of accumulated nanoparticles per a standard unit of mass, for a specific bio-site (including but not limited to a therapeutic site).

Figure 1 presents the biodistribution of the nanoparticle trajectories in mice, resulting from the injection of two types of nanoparticles. The horizontal axis shows the organs (bio-sites): *heart*, *liver*, *spleen*, *lungs*, and *kidney* in which the biodistribution of nanoparticles was sampled *in-vivo* (mice). *Bloodstream* is treated as a bio-site in itself—a common practice in nanoparticle literature—despite its system-wide in connecting all bio-sites. The vertical axis marks the percentage of the injected dose (ID%) per gram of tissue. Each organ has four bars, showing the levels (in ID% units) after 1, 6, 16, and 24 hours from the administration from left to right. The first point displayed for each bio-site is at one hour from administration. The nanoparticle biodistribution is not uniform across the body: the accumulation of nanoparticles varies between organs and changes over time. The biodistribution was shown to be highest after the first hour and then decreased gradually. This process is similar for all the substances administered to the body, but the rate (timing) of the accumulation and the clearance process is different for each nanoparticle type. Indeed, the figure shows two biodistributions for two differently-shaped nanoparticles (Akhtar et al. 2019). The biodistributions are different, as the shape affects the affinity of the nanoparticles to the different bio-sites, i.e., the nanoparticle shape affects the biodistribution. In particular, concentration (ID%/gr) is found to be higher for one-type of nanoparticle compared to the second-type. Furthermore, one nanoparticle type mainly accumulates in the spleen, while the other preferentially attaches to the liver.

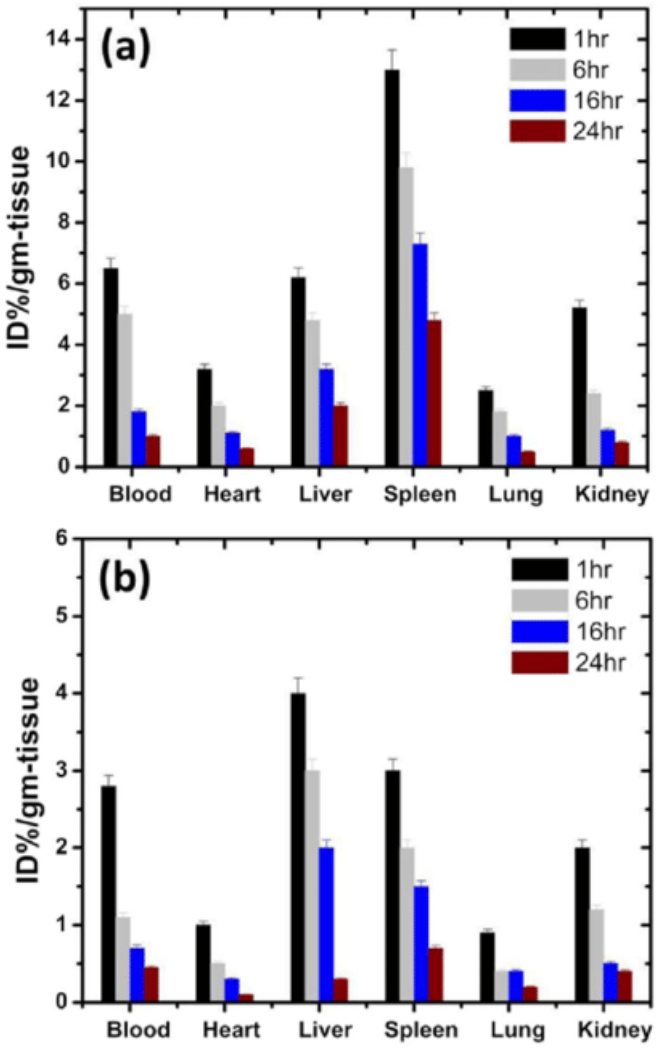

Figure 1: Biodistribution of nanoparticles of different shapes, in the bloodstream, heart, liver, spleen, lungs, and kidney, at different times from the injection. Figure (a) shows actual in-vivo results from injecting mice with one type of nanoparticle, which we term N_1. Figure (b) shows the same, for a second type, which we term N_2. The graphs are adapted from (Akhtar et al. 2019).

Using Biodistributions for Targeted Medicine

Some nanoparticles have biodistribution trajectories that are directly targeted. That is, they accumulate faster and higher at therapeutic sites (Sykes et al. 2014). Modifying the surfaces of nanoparticles by conjugation of a specific protein or other biomarkers, it is possible to use them as a basis for a direct targeted medical treatment (Pearce and OReilly 2019).

However, the real promise of targeted nanomedicine comes from combining different nanoparticles. A recent tumor-detecting approach proposed by von Maltzahn et al. 2011 serves as an example. The approach divides the treatment into two steps, each using a different nanoparticle. The first step begins by injecting "signaling" nanoparticles which attach to the tumor and induce coagulation. After 72 hours, the signaling nanoparticles are entirely cleared from the body, but the elevated coagulation level is high in the tumor microenvironment. Then, a second type of nanoparticle called "receivers" are injected. Over 24 hours, these accumulate in therapeutic sites with elevated levels of coagulation. They conducted two experiments; Firstly, the receiver nanoparticles were tested for treatment by near-infrared irradiation, causing them to heat up and destroy cancer cells. In the other experiment, they were served for diagnosis, by imaging, for detection and localization of tumors.

This approach utilizes two different nanoparticles, which can provide targeted treatment only in combination. The first nanoparticle can target tumors, but cannot be used for imaging or for heating. The second nanoparticle cannot target tumors.

We posit that the process of choosing nanoparticles for use, determining their dosage, and timing and method of administration, is essentially a temporal planning process. The planner is given a description of all the nanoparticles it can administrate, and their bio-diversity trajectory effects. It is given a goal that is specified in terms of biological properties and their levels. It is given safety constraints that prohibit dosage and accumulation from reaching toxic levels, or interacting with each other in harmful ways.

A Representation for Nanomedicine Treatment Planning Problems

We show a concrete representation of the nanomedicine treatment planning problem in terms that are familiar to the AI planning community: bio-diversity *states*, injection *actions* that transition the system between states, and targeted medical *goals*. Along the way, the domain requirements will raise discussion of additional concepts such as derived predicates (to represent interactions between bio-chemical properties), and safety constraints (to prevent the generation of feasible—but harmful—plans). Using this representation the next section will address preliminary solution approaches.

States

We begin by discussing *fluents*, as the building block from which states can be represented. A bio-chemical **property** of a bio-site (therapeutic or otherwise) is represented as a *fluent*, which has a numeric value, representing the concentration level of the property in the given site, in standard units. For example, accumulation of nanoparticles of type N_1 in the liver can be represented as the fluent $liver_{n_1} = x$. It means x nanoparticles are present per gram of liver tissue.

Properties are not limited to measuring nanoparticle concentrations. The levels of coagulation factor in bio-sites, the levels of insulin or glucose in the blood: these are all measurements which are familiar to anyone who has done a blood test, and follow the same pattern, also with established statistical values on normal and abnormal ranges. Biochemical tests estimate these values from measurements, and medically, it is convenient to use concentration per mass, as this allows normalization for different organ sizes and masses.

Nanoparticles affect properties in several ways. First, the accumulation of nanoparticles in a bio-site directly affects their concentration. Second, compounds of which the nanoparticles are composed, can directly add to the levels of the same, pre-existing compounds in the bio-site. For example, if nanoparticles are made of gold (a common theme), the amount of gold in the bio-site increases as nanoparticles accumulate. Third, the presence of nanoparticles can trigger biological reactions that change (reduce or elevate) other properties. Indeed, in the von Maltzahn example we discussed earlier, the signaling nanoparticles cause coagulation, which lingers even after they are cleared.

A set of fluents can therefore be used to describe the bio-chemical status of different bio-sites, and together, of a body. However, as properties affect each other locally (as discussed above), it makes sense to collect together the properties of the same bio-site.

We, therefore, represent a body state s as a set of bio-site property sets,

$$s := \{O_1, \ldots, O_n\} \tag{1}$$

where each **bio-site** (or **organ**) O_j is itself a set of fluents representing properties $\{O_j[1] = v_1^j, \ldots, O_j[m] = v_m^j\}$. $O_j[i]$ is used to refer to the fluent representing property i in bio-site O_j, and v_i^j is its value.

We sometimes simply use v_i to represent the value of property i, as a notational shorthand for the fluent, when the bio-site is understood from the context. The **initial state** of a patient's body may be represented by setting the values of properties, in each organ, to current (normal or abnormal) values. If no nanoparticles were previously injected, then their associated property values will be 0.

An example of a state is presented in Table 1. The data of nanoparticle type N_1 is estimated by the leftmost column of each organ in Figure 1a, when injecting a dosage of 100 type-N_1 nanoparticles. In this example, each organ has only two properties: nanoparticle type N_1, and a hypothetical property p_1. This state represents the domain one hour from the nanoparticle administration, assuming the initial state is the null vector.

Property \ Organ	Blood	Heart	Liver	Spleen	Lung	Kidney
N_1	6.5	3.1	6.2	13	2.5	5
p_1	3.2	430.02	0.001	3.99	32.3	4.3

Table 1: The state representation of the patient's body one hour after injecting a dosage of 100 type-N_1 nanoparticles. Columns represent organs. Rows represent property values. Here, one row shows the accumulation levels type N_1 nanoparticles, and the second row shows the values for a hypothetical property p_1, both after one hour from administration.

Actions

Actions involve the administration of l nanoparticles of type n at time t, represented by the action template $n(l,t)$[2]. This allows multiple administrations (e.g., by injection) of the same nanoparticle type, at different times, regardless of dosage. However, parallel injections of the same type, at the same time, are not allowed by this representation.

The effect of an action is given by the biodistribution trajectories of the administrated nanoparticles, as they change the associated properties (accumulated nanoparticles of type n in different bio-sites) over time. Each action (an administration of nanoparticles) has a multidimensional continuous non-linear effect. Its multidimensionality is in both time and place as it affects several bio-site simultaneously and changes over time. As the effects of an action depend on the parameter l (the dosage), they are specified relative to l.

Figure 2 serves to illustrate. It shows the temporal effect of nanoparticle type N_1 administration in mice. Consider this to be the non-grounded effects of action $n_1(l,0)$. The horizontal axis shows the time passed from administration (0). The amount of injected nanoparticles l is unknown, as this is not a fully-grounded action. Thus the vertical axis measures the concentration per gram of tissue as *percentage of the injection dosage* l. Each line shows the effect of the nanoparticle type N_1 administration on a different bio-site.

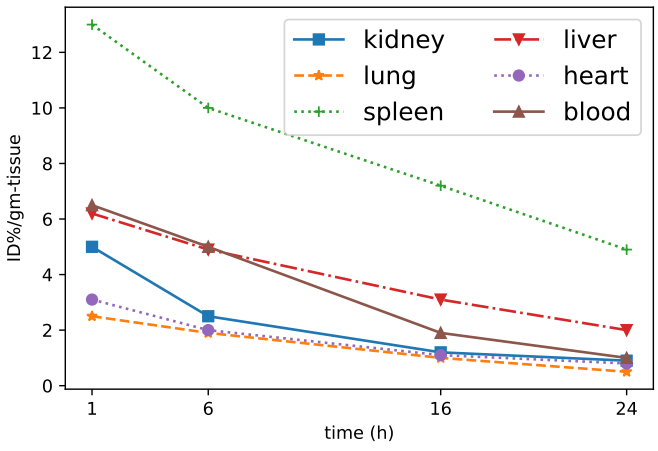

Figure 2: N_1 type nanoparticles administration effect on the paitent's blood, heart, liver, spleen, lung, and kidney over time. The data is estimated from Figure 1a.

[2]For notation simplicity, we distinguish between an action $n()$ and a nanoparticle type n by adding parentheses to the action label.

The data presented in this Figure is estimated from (Akhtar et al. 2019). We simply extracted the data presented in Figure 1a and created an accumulation graph for each organ. We assume that prior to the injection, the subject had no particles in its body. As one can see, the effect of an action is multidimensional, as the administration of nanoparticles affects multiple organs simultaneously. Moreover, it has a continuous effect that changes over time.

To ease the discussion of how action effects are represented and computed, we begin first by discussing simpler special cases of actions, before the more elaborate general case.

Baseline Effects of a Single Action. The simplest case of an action is one that affects a single property. The biodistribution of the administrated nanoparticles only affects each organ directly, modifying a single property on each organ: the accumulated number of nanoparticles that are present in the organ. We refer to this effect as the baseline trajectory of the nanoparticle type accumulation behavior (**baseline** in short).

We follow bio-chemical literature common practices, and assume that the biodistribution trajectories of injection actions are given in percentages of the initial dosage, for the standard mass unit. Computing the grounded effects of an action $n(l,t_0)$ is carried out by computing a function of the initial dosage l and the time since injection $t-t_0$, which yields a percentage in the associate biodistribution trajectory. For simplicity, we assume this function is linear, i.e. it is a simple multiplication of the percentage by l.

Formally, the baseline effect of an action $n(l,t_0)$ determines the properties $O_i[n] = b_n^i$ (property n in every organ O_i). It is determined for time $t \geq t_0$ by:

$$b_n^k[t] = f_n(t-t_0,k)\cdot l \qquad (2)$$

where $f_n(t-t_0,k)$ is the value of the biodistribution trajectory for organ O_k of nanoparticle type n at time absolute time t, when the injection took place in time t_0. In Figure 2, we see the plots of f where k is one of *blood, liver, kidney, heart, lung, spleen*, for an injection at time $t_0 = 0$.

Derived Effects. Some properties are affected by others, as demonstrated by (von Maltzahn et al. 2011). The accumulating signaling nanoparticles (whose concentration levels in any organ k can be computed by Eq. 2) creates a coagulation cascade in the organ k. The level of coagulation in this case is not a simple function of the number of signaling nanoparticles, as coagulation lingers after the signaling nanoparticles clear the organ.

We therefore model the effects of interactions between properties as follows. Suppose a property d depends on other properties $p_1, \ldots, p_z$, i.e, the value v_d of property d is dependent on the values $v_1, \ldots, v_z$ (all in the same bio-site, for simplicity of the illustration). We assume this dependency is captured by a function D_d which determines the value of d at time $t \geq t_0$) as follows:

$$v_d[t] = D_d(t, v_d[], \{v_i[]\})$$

where $v_d[]$ is the past trajectory (historical values, up to time t) of d in the same bio-site, and likewise $v_i[]$ stands for the historical values of property p_i. We do not allow same-time cycles in this computation; if d at time t depends on p at time t, then p cannot depend on d at the same time.

The reader should keep in mind that although the processes are continuous, the currently available bio-data is not. To the best of our knowledge, biologists' smallest measurement scale is minutes. Thus, we allow finite discretization of time. We will also assume a Markov property of the calculation, and simplify the above to:

$$v_d[t] = D_d(v_d[t-1], \{v_i[t]\}) \tag{3}$$

Going back to the example, the coagulation level v_c at time t can be computed—*derived*—as a function of the baseline level of signaling nanoparticles (s), given by $b_s[t]$:

$$v_c[t] = D_c(v_c[t-1], b_s[t])$$

Conditional Effects of a Single Action. The baseline effects of an action, described above, assume that the accumulation and clearance of administered nanoparticles are a function of the biodistribution trajectories *alone*, and do not depend on the bio-chemical conditions of the bio-site, as represented by other properties. However, in reality, the biodistribution trajectories change depending on such conditions.

Indeed, some nanoparticle types accumulate differently under different conditions, i.e., their accumulation and clearance depend on the value of other properties. An example of such behavior can be found in the von Maltzahn example discussed earlier (von Maltzahn et al. 2011): The baseline behavior of the receiver nanoparticle changes due to the presence of coagulation; it targets tissues with increased coagulation.

Note that this case is a special case of the derived effects discussed earlier. Here, the property describing the accumulation of nanoparticles is not directly described by the baseline (Eq. 2), but rather, the effects are a function of a different property, which itself is derived from others (Eq. 3).

Combining both equations, we describe the value of the accumulation of a nanoparticle type n at time t in organ O_k is by $S_n^k[t]$, given a function that describes the accumulation based on both the baseline and other properties:

$$S_n^k[t] = C_n(S_n^k[t-1], b_n^k[t], \{S_j^k[t]\}, \{v_p^k[t]\}) \tag{4}$$

where $S_n^k[t-1]$ is the concentration level of n at time $t-1$, $b_n^k[t]$ is the baseline value for particles if type n at time t (Eq. 2), $\{S_j^k[t]\}$ is the set of levels of other particle types ($j \neq n$) on which n's accumulation may be dependent, and similarly $\{v_p^k[t]\}$ is a set of levels of other properties (Eq. 3).

This is the general form for computing the effects of an action. $b_n^k[t]$, $S_j^k[t]$,$v_p[t]$ are all values of properties on which $S_n^k[t]$ is dependent. So in fact, $S_n^k[t]$ is computed as a derived effect.

For the von Maltzahn example, the accumulation of the first, signaling nanoparticles (α) in any biosite k is given by:

$$S_\alpha^k[t] = b_\alpha^k[t] = f_\alpha(t, k) \cdot l$$

which leads to increased coagulation level:

$$v_c^k[t] = D_c^k(v_c^k[t-1], S_\alpha^k[t])$$

and both of these factors affect the accumulation of the second, receiving nanoparticles (β) as follows:

$$S_\beta^k[t] = C_\beta(S_\beta^k[t-1], b_\beta^k[t], v_c^k[t])$$

The State-Transition Function $\delta()$. Let us denote $S_i^k[t]$ as the value of property i at time t in organ O_k. The effect of an action $n()$ on property i of organ k at time $t+1$ depends on the number of particles from $n()$ on organ k at time $t+1$ and their effect on property i.

We now generalize the effect compositions to consider multiple actions affecting the same property, at the same time, in parallel. Let us denote the set of actions taken between time 0 and t by N_t. The value of any nanoparticle accumulation property, $S_i^k[t]$, is the sum of effects of all actions in N_t:

$$S_i^k[t] = \sum_{a \in N_t} S_i^k[t]_a \tag{5}$$

where $S_i^k[t]_a$ is the value computed by Eq. 4 for any action $a = n(l, t_0) \in N_t$. Any such $S_i^k[t]_a$ includes the computation of any derived properties on which S_i^k is dependent. The start time of the action is its administration time. For simplicity, we assume a single administration point.

The state-transition function $\delta()$, analogous to the classical planning definition, will compute Eq. 5 for time t, for all properties, in all bio-sites (Eq. 1).

Goal States and Safety Constraints

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). This is a relatively straightforward extension of how we define goal states in classical planning and factored state representations. The general preferences of reaching a goal-state quicker would be familiar to planning researchers. The result of the planning process is a timed sequence of injection actions, which will cause the levels of specific properties to reach their target values.

However, the medical safety of the plan is a key concern, and it affects states that are reachable (and may be reached) at a time *after* the goal state has been reached. In particular,

any plan returned by the planner must conform to safety constraints imposed by the medical professional, and is given to the planner as part of the planning problem.

Safety constraints 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 tailored for patients based on their specific health conditions. 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 organ k can be expressed as follows: $v_j^k \leq h$, or $v_j^k \geq h$.

A safety constraint can also constrain an interaction between two properties, e.g., when drugs interact with each other. For example, we may represent a constraint that if a property value i in an organ k is greater than a given threshold h_i, then the value of property j in the same organ must be less than a threshold h_j: Namely, $v_i^k > h_i \Rightarrow v_j^k < h_j$.

As the effects of action have a duration, the first state in which the goal state is reached (i.e., certain properties meet goal conditions), is not the last state in which actions affect change in the body. Given a time t_g in which the goal conditions are first met, safety constraints must be satisfied not only in the interval $[0, t_g]$ but also in the interval $[t_g, \infty)$ (though generally, the effects of actions do have finite durations).

Planning and Scheduling Treatments

Using the definitions above, we can now define both planning and scheduling problems.

Planning Problem Definition. Targeted nanomedicine planning problems involve deciding both on the type of the nanotypes used, as well as on their dosage and their timing. It is represented as a tuple $\langle S, A, s_0, \delta, S_G, C \rangle$ where S is the set of states, $s^0 \in S$ is the initial state, A is a set of actions, S_G is the set of goal states, described as a set of conditions on the goal properties, and C is a set of safety constraints. δ is the state transition function, which is composed of all the functions that define S_i^k and any derived properties. Given a treatment problem, a valid solution plan is an ordered set of actions (parallel actions allowed) that reaches a goal state and does not violate any of the safety constraints both before and after the goal state is reached. If there are multiple valid plans, the plan with the shortest overall time (from initial injection until the nanoparticles are cleared from the body) will be chosen. Reducing the treatment time in medical procedures is important as it might reduce possible side effects and accelerate healing (Sykes et al. 2014).

Scheduling Problem Definition and Preliminary Approach. A special case of the planning is a scheduling problem, where the type and number of nanoparticle injections is given in the problem description, and the only open question is their timing, such that all safety constraints are satisfied.

We are currently investigating an approach towards solving the scheduling problem, using a hill-climbing approach[3]. Hill-climbing is an iterative heuristic search algorithm that aims to find a local optimal solution regarding a given objective function. It starts at a random state and tries to improve the current solution at each iteration. For the scheduling problem, the algorithm attempts to reschedule the injection times until it reaches a local optima. It then repeats this process for different possible order of injections, and random initial times. Hill-climbing allows searching in a continuous domain, thus it is a strong candidate for solving the described above scheduling problem.

To apply this technique, one needs to define an objective function that the algorithm tries to minimize (or maximize). Our objective function examines several criteria: The first criterion to consider is the total number of unsatisfied constraints in the current solution, as a valid solution must satisfy all safety constraints. Another criterion is the overall treatment time. Treatment begins at the first injection and ends once all injected particles were cleared out of the patient's body (after the last injection in the treatment plan). A final criterion is the number of injections, which generally should be minimized.

Discussion and Related Work

Each action (an administration of nanoparticles) has a multidimensional, continuous, non-linear, durative effect. Its multidimensionality is in both time and place as it affects several organs simultaneously and changes over time. In particular, multiple actions can take place in parallel and simultaneously affect the same state variable.

Penberthy and Weld proposed a temporal planner with continuous effects (Penberthy and Weld 1994). Their planner, however, assumes that there is no continuous effect without a direct explicit action causing it. This generally does not hold in the targeted nanomedicine domain, as interactions between different properties or natural processes in the body have continuous effects, but are not modeled as actions, but instead as derived properties.

PDDL (Planning Domain Definition Language) is one of the most popular planning language families. Fox and Long suggest an extension of PDDL (PDDL+) to model events and processes (Fox and Long 2006), allowing for continuous effects. Several planners extend the original PDDL+ to support non-linear effects (Scala et al. 2016; Piotrowski et al. 2016). We currently examine the possibility of using such planners to represent this domain.

Boutilier and Brafman proposed a planner that models simultaneous actions (Boutilier and Brafman 2001). The actions in their planner are single points in time, whereas our problem requires durative actions.

Michalowski et al. (Michalowski et al. 2021) proposed a planner for multi-morbid patients. Their planner receives a patient's information, optimization function, planning horizon length, and medical goals and outputs a conflict-free plan. Their planner uses existing medical procedures and chooses the combination of plans that minimizes (or maximizes) the objective function, replacing adverse actions between chosen sub-plans. They use *knowledge repositories* as a reference for medical actions that can serve to substitute one another. In comparison, the planner we suggest

[3]We thank Sven Koenig for suggesting this approach.

aims to generate a plan from scratch. We do not use templates from existing medical procedures. Furthermore, the knowledge of adverse actions and their equivalent actions is not inferred from the actions' conditions and effects, but is given as facts. We do not specify which actions cannot be taken at the same plan and which actions have similar results. Instead, we specify a health plan's constraints and actions' effects and preconditions, and the planner infers the rest.

Conclusion and Future Work

We propose that targeted nanomedicine treatments are a novel domain for AI Planning. To represent such problems, we propose a novel representation of temporal planning problems. Planning in this domain requires planning for parallel actions, each of which can have multi-dimensional, continuous, non-linear effects. Additionally, we discussed a scheduling variant of the treatment planning problem and suggested solving it with a hill-climbing algorithm. We are continuing to develop planning and scheduling algorithms and will evaluate them in experiments using real-world and synthetic data.

References

Akhtar, S.; Khan, Q.; Anwar, S.; Ali, G.; Maqbool, M.; Khan, M.; Karim, S.; and Gao, L. 2019. A comparative study of the toxicity of polyethylene glycol-coated cobalt ferrite nanospheres and nanoparticles. Nanoscale Res Lett 14: 386.

Boutilier, C.; and Brafman, R. I. 2001. Partial-order planning with concurrent interacting actions. *Journal of Artificial Intelligence Research*, 14: 105–136.

Cho, K.; Wang, X.; Nie, S.; Chen, Z.; and Shin, D. M. 2008. Therapeutic nanoparticles for drug delivery in cancer. *Clinical cancer research*, 14(5): 1310–1316.

Fox, M.; and Long, D. 2006. Modelling mixed discrete-continuous domains for planning. *Journal of Artificial Intelligence Research*, 27: 235–297.

Kirtane, A. R.; Verma, M.; Karandikar, P.; Furin, J.; Langer, R.; and Traverso, G. 2021. Nanotechnology approaches for global infectious diseases. *Nature Nanotechnology*, 16(4): 369–384.

Michalowski, M.; Rao, M.; Wilk, S.; Michalowski, W.; and Carrier, M. 2021. MitPlan 2.0: Enhanced Support for Multimorbid Patient Management Using Planning. In *International Conference on Artificial Intelligence in Medicine*, 276–286. Springer.

Pearce, A. K.; and OReilly, R. K. 2019. Insights into active targeting of nanoparticles in drug delivery: Advances in clinical studies and design considerations for cancer nanomedicine. *Bioconjugate chemistry*, 30(9): 2300–2311.

Penberthy, J. S.; and Weld, D. S. 1994. Temporal planning with continuous change. In *Aaai*, volume 94, 1010.

Piotrowski, W. M.; Fox, M.; Long, D.; Magazzeni, D.; and Mercorio, F. 2016. Heuristic Planning for PDDL+ Domains. In *AAAI Workshop: Planning for Hybrid Systems*, volume 16, 12.

Raj, S.; Jose, S.; Sumod, U.; and Sabitha, M. 2012. Nanotechnology in cosmetics: Opportunities and challenges. *Journal of pharmacy & bioallied sciences*, 4(3): 186.

Scala, E.; Haslum, P.; Thiébaux, S.; and Ramirez, M. 2016. Interval-based relaxation for general numeric planning. In *ECAI 2016*, 655–663. IOS Press.

Souto, E. B.; Souto, S. B.; Campos, J. R.; Severino, P.; Pashirova, T. N.; Zakharova, L. Y.; Silva, A. M.; Durazzo, A.; Lucarini, M.; Izzo, A. A.; et al. 2019. Nanoparticle delivery systems in the treatment of diabetes complications. *Molecules*, 24(23): 4209.

Sykes, E. A.; Chen, J.; Zheng, G.; and Chan, W. C. 2014. Investigating the impact of nanoparticle size on active and passive tumor targeting efficiency. *ACS nano*, 8(6): 5696–5706.

Turan, O.; Bielecki, P.; Tong, K.; Covarrubias, G.; Moon, T.; Rahmy, A.; Cooley, S.; Park, Y.; Peiris, P. M.; Ghaghada, K. B.; et al. 2019. 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): 4352–4360.

von Maltzahn, G.; Park, J.-H.; Lin, K. Y.; Singh, N.; Schwöppe, C.; Mesters, R.; Berdel, W. E.; Ruoslahti, E.; Sailor, M. J.; and Bhatia, S. N. 2011. Nanoparticles that communicate in vivo to amplify tumour targeting. *Nature materials*, 10(7): 545–552.