Automated Planning to Support Physical Rehabilitation

Alessandro Umbrico[1], Roberta Bevilacqua[2], Marco Benadduci[2], Amedeo Cesta[1], Francesca Fracasso[1], Alessandro Leone[3], Elvira Maranesi[2], Mauro Marzorati[4], Andrea Orlandini[1], Giovanna Rizzo[4] and Gabriella Cortellessa[1]

[1]National Research Council of Italy; Institute of Cognitive Sciences and Technologies
[2]National Institute of Science and Health for Aging (INRCA-IRCCS)
[3]National Research Council of Italy; Institute of Microelectronics and Microsystems
[4]National Research Council of Italy; Institute of Biomedical Technologies

Abstract
This work presents a contribution about the use of a temporal planning framework for the synthesis of a personalized rehabilitation exercise program in the shape of dancing sessions. A first part of the paper focuses on the modeling workflow to encode/represent all the technical and clinical information necessary to the planner. A second part describes the temporal planning framework and the multi-objective reasoning approach adopted to support the desired (clinical) objectives. An experimental evaluation then assesses the technical feasibility of planning framework and its capability of reasoning on different features to support (clinical) objectives.

Keywords
Temporal Planning, Timeline-baed Planning, Heuristic Search, Multi-objective Search, Planning and Scheduling

1. Introduction

According to the last EU Commission's Report on Aging, from 2019 to 2070, the share of the age cohorts above 65 years in the EU population is expected to rise markedly from 20% to 30%, with the share of those aged 80 and over doubling from 6% to 13%. By contrast, the share of the age group 20-64, namely the working-age population, would fall from 59% to 51% of the total population [1]. Under these circumstances, understanding healthy aging and age-related diseases represents an imminent challenge, and investing in more sustainable healthcare services becomes an urgency since the decrease of working-age

Italian Workshop on Artificial Intelligence for an Ageing Society (AIxAS 2021), November 29th, 2021
alessandro.umbrico@istc.cnr.it (A. Umbrico); r.bevilacqua@inrca.it (R. Bevilacqua); m.benadduci@inrca.it (M. Benadduci); amedeo.cesta@istc.cnr.it (A. Cesta); francesca.fracasso@istc.cnr.it (F. Fracasso); alessandro.leone@cnr.it (A. Leone); e.maranesi@inrca.it (E. Maranesi); mauro.marzorati@itb.cnr.it (M. Marzorati); andrea.orlandini@istc.cnr.it (A. Orlandini); giovanna.rizzo@itb.cnr.it (G. Rizzo); gabriella.cortellessa@istc.cnr.it (G. Cortellessa)
0000-0002-1184-5944 (A. Umbrico); 0000-0002-3851-3552 (R. Bevilacqua); 0000-0002-0703-9122 (A. Cesta); 0000-0001-9442-870X (F. Fracasso); 0000-0002-8970-3313 (A. Leone); 0000-0002-2414-3773 (E. Maranesi); 0000-0003-1093-2162 (M. Marzorati); 0000-0001-6458-5202 (A. Orlandini); 0000-0002-6341-1304 (G. Rizzo); 0000-0002-9835-1575 (G. Cortellessa)
© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).
CEUR Workshop Proceedings (CEUR-WS.org)
CEUR-WS.org/Vol-3108/paper5.pdf

population will lead to an increase in the demand of healthcare professionals in face of an increased need for a higher level of care for future assistance.

In this context, a relevant research trend is Socially Assistive Robotics which aims at realizing increasingly supportive, proactive and personalized assistive services for the aging or frail population, and current results in this field suggest the development of potential new and innovative processes in healthcare [2]. Many projects and research initiatives like e.g., [3, 4, 5] do exists which have pursued this objective by focusing on different aspects of social and/or health assistance. Another example is represented by an ongoing research project, SI-Robotics that proposes a holistic approach to assistive and monitoring interventions in pro-active assisted living applications, in different contexts such as the house of an older person, hospitals and residential facilities. The offered services are tailored according to the user's needs and the context of delivering, and include: i) physical activity monitoring (occupational, leisure time and household activity) and coaching; ii) physiological monitoring (chronic diseases, ageing features, psychological status, both emotional and cognitive, stress level); iii) habits recognition, lifestyle changes and promotion; iv) physical/cognitive decline assessment and support; v) lifestyle support and physical/cognitive stimulation (nutrition, cooking, physical activity) and; vi) advanced tele-operation with caregivers. In order to address all these challenges, SI-Robotics relies, among others, on the integration of a number of AI technologies ranging from Knowledge Representation and Reasoning, for user modeling and personalization to Machine Learning and Automated Planning for a continuous proactive and adaptive assistance.

The project pursues a mixed-initiative approach by putting clinicians and therapists into the AI-based loop for the synthesis and monitoring of personalized and adaptive assistance [6, 7]. Automated Planning plays a central role in supporting domain experts (i.e., clinicians/therapists). Planning technologies are, in this project, crucial to support therapists in making decisions and merge clinical objectives with health needs of different end-users and different assistive scenarios.

This work focuses on the use of timeline-based planning technologies to support physical rehabilitation programs. Following a mixed-initiative methodology, plans are personalized according to the specific clinical objective selected by a therapist. Therapists specify the (clinical) objective of a rehabilitation session and the planner synthesizes a rehabilitation exercise program as a set of suitable stimuli. The planning process reasons on a number of clinical and technical features of known stimuli in order to select those that meet clinical requirements but also can be carried out correctly by users. The mixed-initiative methodology supports a continuous refinement of planning knowledge according to therapists' feedback. The paper first presents the overall therapist-in-the-loop methodology pursued to elicit and continuously refine suitable planning knowledge. Then an experimental evaluation assesses the reasoning capabilities of the planner and its efficacy in addressing clinical objectives. The following section focuses on a specific target of old population and explains the role of the planning technology to support their rehabilitation programs.

2. Physical Therapy and Parkinson's Disease

Neurological disorders are the leading source of disability globally [8]. In the Global Burden of Disease, Injuries, and Risk Factors Study, Parkinson's Disease (PD) is described as the fastest growing in prevalence, disability, and deaths. The rising prevalence of PD is linked with the rising of the ageing population considering that age above 65–70 years is a well-established risk factor. Advancing age is indeed associated with a faster rate of disease motor progression, decreased levodopa responsiveness, more severe gait and postural impairment, and more severe cognitive decline, exiting in the development of dementia [9]. Within the unmet care needs of PD population, the identification of strategies to support the feeling of being active part of the society, is the most requested one by the patients and their families [10].

To this purpose, guidelines recommend physical therapy early at the onset of the disease, to support mobility by counteracting the insurgence of motor symptoms [11]. PD patients, in fact, are routinely treated through rehabilitative approaches aimed at improving static and dynamic balance, recovery of walking, falls and mobility [12]. Within non-pharmacological treatments, many studies [13] have recently demonstrated that regular physical exercise practice, predominantly aerobic, have a beneficial effect on balance and gait functional mobility [14]. In particular, novel interventions based on different types of dance (e.g., Tango or Irish dance) have been designed to recover the normal gait of patients with PD [15, 16].

One of the objectives of SI-Robotics is therefore to realize dance-based training/rehabilitation conducted by a social robotic coach and supervised by the physiotherapist. The service will be integrated in the daily activity of the IRCCS-INRCA rehabilitation Unit, located in Ancona (Italy), that routinely conducts group therapy with older PD patients at different stage of disease severity [1].

2.1. The Role of the Planner

Within the designed dance-based rehabilitation program, the planner is in charge of supporting the therapist in the construction of rehabilitation sessions (i.e., synthesis of suitable rehabilitation exercise programs). The planner is responsible for the synthesis of rehabilitation exercise programs and thus for the selection of the physical exercises that patients will perform during each rehabilitation session.

More specifically, the therapist provides the planner with information about the rehabilitation session (e.g., song time and song speed) and the health needs of participating users. These needs characterize the clinical objective of the session and thus determine the "shape" of the resulting rehabilitation exercise. Given this input, the planner builds a "choreography" of a dance session by selecting the combinations of dancing steps (i.e., stimuli) that best fit the desired clinical objective.

As we will show in the next section, the execution of each dancing step is seen as the administration of a particular physical stimulus addressing one or more health-

[1] https://www.inrca.it/inrca/home.asp?ling=en

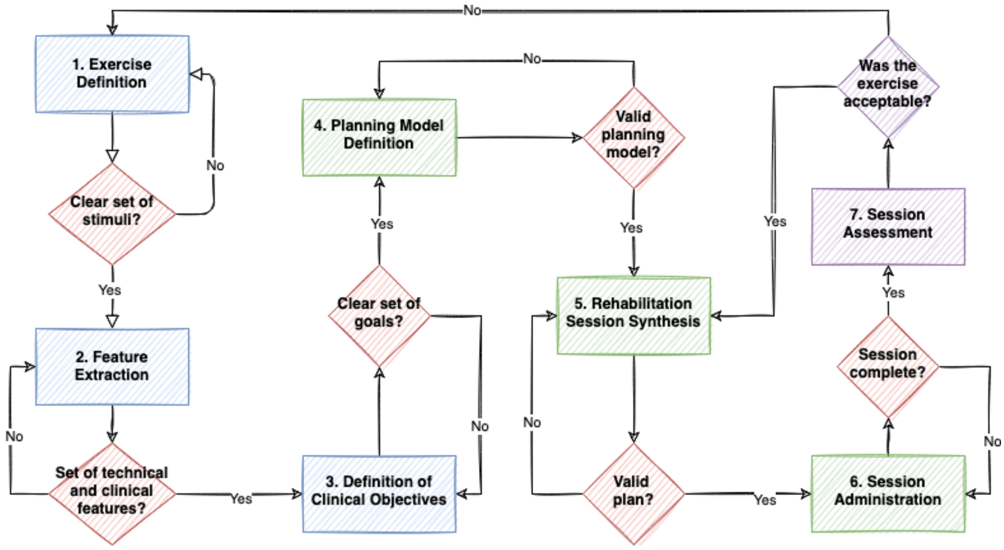

Figure 1: Planning model definition and validation process

related aspects of a user. The selection of such stimuli should therefore be contextualized according to the selected objective and thus to patients' health-related needs.

3. Materials and Methods

3.1. Knowledge Acquisition and Model Validation

To achieve the desired objective a planning system should be endowed with carefully designed domain knowledge. Broadly speaking, a planner should know a possibly rich set of stimuli. It should know the effects of such stimuli on the health-related needs of a user. It should be able to contextualize such stimuli with respect to different (clinical) objectives in order to synthesize effective plans. In addition, a planner should know technical constraints concerning the physical execution and concatenation of dancing steps like e.g., "spatial effects" of motions, in order to synthesize plans that are valid in the considered scenario. The acquisition of such knowledge is not trivial and entails continuous refinements and constant interactions with therapists. Figure 1 shows the mixed workflow implemented for the definition and refinement of the domain knowledge necessary to support the planning process.

As shown, the workflow is iterative and may entail (mixed) iterations at different levels of the knowledge definition and validation process. The flow starts with the Exercise Definition and thus with the definition of the dancing steps that will be considered in the synthesis process. These steps represent the primitive stimuli the planner should concatenate to build rehabilitation exercise programs. The definition of these steps is guided by the therapist in order to define a portfolio of physical stimuli that are suitable for the considered target of end-users.

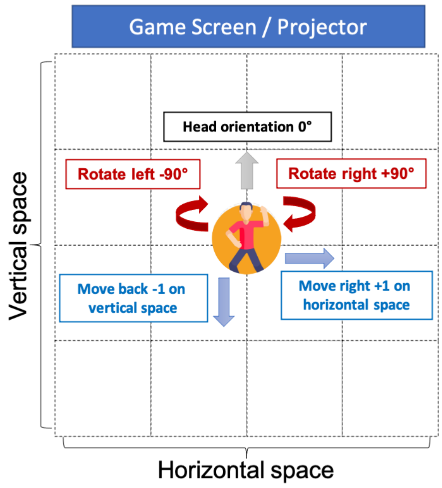

Figure 2: Layout of a rehabilitation session

The next two steps of the flow aims at: defining a number of features that characterize each dancing step from a technical and clinical perspective and; defining a set of clinical objectives that characterize how to concatenate known dancing steps. After these three steps a complete planning knowledge is ready for the definition of a suitable planning model. The next three steps are therefore mainly guided by the planning expert with the aim of synthesizing and executing contextualized rehabilitation exercises.

The execution of each session (i.e., rehabilitation exercise program) is monitored and assessed by a therapist who validates the outcomes of the exercise and the contribution of the planner. The last step of the flow ("Session Assessment") therefore can "approve" the executed exercise so that the developed planning framework can be used for next sessions. Otherwise, the therapist can "reject" the executed exercise and thus trigger a new iteration of the flow in order to refine the planning knowledge and "tune" the behavior of the planner to better fit the desired objectives.

3.1.1. Definition of Stimuli and Spatial Constraints

Putting aside clinical and rehabilitation aspects, the role of the planner is to decide a sequence of motions (i.e., dancing steps) patients should perform over time. To correctly plan such sequences the planner take into account spatial constraints in order to correctly combine/concatenate users' motions.

Figure 2 generally describes the physical layout of the scenario. Users are put in front of a wide screen where an avatar shows the dancing steps they are supposed to perform in a game-like way. The two-dimensional grid of the figure shows a schematic representation of the "dance floor" and possible motions of a user. For the sake of simplicity we here consider only one user but the layout can be easily extended to the case with more users.

A user is placed at the center of the layout with the head pointed to the game screen. Users' motions consist of steps along vertical and horizontal axes and body rotations. It is important to point out that rotations change body orientation and thus the reference system of a user. This means that the axis affected by a motion depends on the body orientation of the user. For examples, a step forward with orientation 0 degree (i.e., with the head pointed to the game screen) affects the vertical axis. A step forward with orientation 90° (i.e., after a body rotation to the right) instead affects the horizontal axis. Each motion of a user consumes one "space unit" along a particular axis. Taking into account the default position (i.e., the user placed at the center with the head pointed to the game screen) a user can perform a maximum number of N steps forward and backward along the vertical axis (for a total available space of $2N$ units) and can perform a maximum number of M steps rightward and leftward along the horizontal axis (for a total available space of $2M$ units).

Given this layout it is necessary to characterize the basic motions a user can perform. These steps constitute the stimuli a planner should concatenate to synthesize rehabilitation exercises (i.e., the plans). As a result a dataset of (primitive/basic) steps has been defined containing a total number of 134 steps. Each step is characterized with respect to the "effects" that related motions have on a user in the considered layout. Technical features characterize "spatial effects" of each step and other technical information necessary to the planner. Table 1 shows the structure of the dataset and describes the defined technical features.

3.1.2. Definition of Clinical Features and Objectives

In addition to the technical knowledge about the defined dancing steps, the planner needs clinical knowledge to build sessions that meet clinical requirements. To synthesize effective rehabilitation exercise programs indeed the planner should know how available stimuli (i.e., dancing steps) affect the health state of users. To this aim, the second step of the workflow in Figure 1 generates an additional dataset to characterize clinical parameters of dancing steps. A number of clinical features are defined to characterize how the execution of a particular step "stimulate" the health-state of a user. Table 2 shows the structure of the dataset with a description of the defined clinical features.

The third step of the workflow of Figure 1 then concerns the definition of clinical objectives (i.e., planning goals). This information is crucial to characterize the expected qualities of plans and thus synthesize rehabilitation exercises that are effective with respect to the health conditions of users. Assuming that the choice of the clinical objective is made by therapists according to their knowledge about the status of the participants, two main cases have been considered:

- Stimulation of physical equilibrium. The rehabilitation exercise should stimulate users' capability of keeping a good equilibrium and thus coordinate their motions. The "intensity" of the exercise is not central in such a case. Rather it is important to administrate a set of stimuli that properly train coordination and body balance of users.
- Stimulation of metabolic response. The rehabilitation exercise should stimulate users' energy expenditure and cardiovascular fitness. The "intensity" of the exercise

Table 1
Technical features

Name	Type	Description
Step Id	Unique String	A string uniquely identifying a basic step within the dataset.
Step Name	String	Textual description providing a simple description of the motion a user is supposed to perform.
Step duration	Integer $\in \mathbb{N}^+$	Information about the expected number of "beats" necessary to execute a step. Each dancing session has a basic bpm (beats per minute like e.g., 60 or 120 bpm) denoting the frequency of the song and the relative speed of the dancing steps. For example, in the case of a song with 60 bpm (i.e., 1 beat per second) a step with a duration of 2 beats is expected to be executed in 2 seconds. In the case of a song with 120 bpm (i.e., 2 beats per second) a step with a duration of 2 beats is expected to be executed in 1 second.
Body Orientation	Integer $\in \{0, 90, 180, 270\}$	Information about the orientation of the body during the execution of the step. Basic steps are distinguished with respect to the body orientation assumed during their execution. For example the dataset may contain an entry "step forward with right foot with orientation 0" and an entry "step forward with right foot with orientation 180". This helps the planner to contextualize the motion and correctly "propagate" spatial constraints.
Horizontal Translation Vertical Translation	Integer $\in [-1, 1]$	Information about spatial effects of a step on the horizontal/vertical axis of the layout. A positive number (1) denotes a movement of user's body to the right-side/top-side of the layout. A negative number (-1) denotes a movement of user's body to the left-side/bottom-side of the layout. Body translations are expressed with respect to the reference system of the layout since they depend on the body orientation of a user. For example, a rightward performed with 0° orientation moves user's body to the right-side of the layout (1). A rightward step performed with 180° orientation instead moves user's body to the left-side of the layout (-1).
Horizontal Requirement Vertical Requirement	Integer $\in [-1, 1]$	Information about the spatial requirements of a step on the horizontal/vertical axis of the layout. A positive or negative value denotes steps that do not change the current position of a user but require some space to be performed like e.g., "rise right arm". Similarly to the translations these requirements are expressed with respect to the layout reference system in order to contextualize spatial constraints according to body orientation of users.

is central in this case. It is important to identify and administrate a set of stimuli that achieve a proper level of cumulative energy in order to stimulate the physical resistance of users.

In general each clinical objective entails the capability of reasoning on a number of (heterogeneous) features. A planner should therefore be capable of pursuing a multi-objective perspective in order to synthesize effective plans.

Table 2
Clinical features

Name	Type	Description
Step Id	Unique String	A string uniquely identifying a basic step within the dataset.
Effort on Lower Body Parts	Integer $\in [0, 10]$	Estimation of the muscular effort on the legs required for the correct execution of the associated dancing step.
Effort on Upper Body Parts	Integer $\in [0, 10]$	Estimation of the muscular effort on the arms required for the correct execution of the associated dancing step.
Energy	Integer $\in [0, 10]$	Estimation of the amount of energy required for the execution of the associated dancing step. This information allows the planner to contextualize dancing steps with respect to the energy metabolism and the "training level" of a user.
Coordination	Integer $\in [0, 10]$	Characterize the difficulty of a dancing step with respect to both cognitive and motor aspects. This information is crucial to adapt the technical difficulty of a choreography to the actual motor skills and cognitive capabilities of users.
Body Balance	Integer $\in [0, 10]$	Characterize the difficulty of a dancing step with respect to the maintenance of body equilibrium.

3.1.3. Execution and Validation of Synthesized Plans

A planning instance can be defined after the third step of the workflow and thus the definition of a complete planning knowledge is completed. In this regard, the workflow of Figure 1 wants to stress the continuous validation of domain experts (i.e., therapists). The "Session assessment" step indeed assesses the quality of the executed exercise (i.e., plan). Feedback from therapists is crucial to incrementally refine the defined model of stimuli, the planning model and the implemented selection criteria. Next sections describe how the gathered knowledge is actually modeled into a (timeline-based) planning specification and how clinical features are used to adapt the search of a planner to achieve the clinical objective selected by the therapist.

3.2. Timeline-based Planning

Timeline-based Planning has been introduced in early 90s [17] by taking inspiration from the Control Theory. A key aspect of this paradigm is the integration of planning and scheduling and the synthesis of plans as envelopes of valid (and synchronized) temporal behaviors of some domain entities. This aspect have characterized the practical use of this paradigm in several real-world scenarios, especially from space like e.g., [18, 19].

In the current work, we apply the timeline-based approach to the synthesis of personalized rehabilitation exercises by extending the solving capabilities of an open-source timeline-based framework called PLATINUm [20]. Before entering into the details of the defined planning model and the experimental evaluation this section provides the reader with some backgrounds about the planning formalism and the general solving capabilities of the framework.

3.2.1. The Formalism in a Nutshell

A timeline-based specification describes the valid behaviors of a number of domain entities. Given this description, a timeline-based planning process synthesizes a set of (temporally) flexible behaviors (i.e., timelines) that describe how these entities should evolve to achieve some objectives. According to the formalization proposed in [21], domain features are modeled by means of state variables.

Definition 1. A State Variable is a tuple $SV = \langle V, T, D, \gamma \rangle$ describing valid behaviors of a domain entity:

- V is a set of values $v_i \in V$ representing states or actions an entity can perform or assume over time.
- $T : V \rightarrow 2^V$ is a state transition function describing for each value $v_i \in V$ its possible successors.
- $D : V \rightarrow \mathbb{T} \times \mathbb{T}$ is a duration function specifying for each value $v_i \in V$ its expected duration bounds, expressed in some temporal domain $\mathbb{T}$ (typically $\mathbb{N}^+$).
- $\gamma : V \rightarrow \{c, pc, u\}$ is a controllability tagging function specifying the controllability property of a value.

Controllability properties characterize the execution of SVs' values with respect to the dynamics of the environment.

Definition 2. A value $v_i \in V$ of a state variable $SV = \langle V, T, D, \gamma \rangle$ is:

- Controllable ($\gamma(v_i) = c$) if the system can decide both the start and end times of its execution.
- Partially-controllable ($\gamma(v_i) = pc$) if it can only be started by the system, while its end time can only be observed.
- Uncontrollable ($\gamma(v_i) = u$) if it can only be observed and thus the system can neither decide its start nor its end.

Information about controllability and temporal flexibility are crucial to deal with temporal uncertainty and support robust execution of temporal plans (see the controllability problem [22]).

A flexible timeline for a state variable SV_i is a sequence of temporal intervals called tokens that describes an envelope of valid temporal behaviors.

Definition 3. If $SV_i = (V, T, \gamma, D)$ is a state variable, a token x_j for the variable has the form:

$$x_j = (v_k, [e_j, e_j^{'}], [d_j, d_j^{'}], \gamma(v_k))$$

where $v_k \in V$ is the value assumed by the token x_j, $[e_j, e_j^{'}]$ is the end-time interval of x_j (with $e_j \leq e_j^{'}$) and $[d_j, d_j^{'}]$ is the minimum and maximum duration of x_j (with $d_j \leq d_J^{'}$).

Synchronization rules specify additional constraints that are necessary to synthesize timelines that achieve desired objectives (i.e., planning goals).

Definition 4. A synchronization rule has the form

$$a_0[SV_0 = v_0] \rightarrow \ a_1[SV_1 = v_1], ..., a_n[SV_n = v_n].C$$

where every $a_i[SV_i = v_i]$ is a token variable denoting a temporal interval in which a state variable SV_i assumes the value v_i. The left-hand part of the synchronization rule ($a_0[SV_0 = v_0]$) is called the trigger of the rule. The set C specifies temporal relations between token variables.

Synchronization rules with the same trigger are treated as disjunctions and represent alternative constraints that should hold between different sets of token variables.

These are the main concepts composing a timeline-based planning specification. The reader may refer to [21] for a complete description of the formalism.

3.2.2. Iterative Refinement of Temporal Behaviors

PLATINUm[2] is an open-source timeline-based framework compliant with the described formalism. The synthesis of a valid solution plan π is implemented as a general plan refinement procedure.

Given a set of state variables ($\mathcal{SV}$) and synchronization rules ($\mathcal{S}$) a partial plan is initialized from facts and goals of a planning problem. A partial plan represents a set of partially instantiated timelines whose behavior is neither complete nor valid. The solving process iteratively refines such timelines until a complete and valid set of SVs' behaviors is synthesized. In this regard, a flaw is a general condition affecting the completeness or validity of a timeline. Flaws may concern either tokens to be added to a timeline (i.e., planning flaws) or, overlapping tokens of a timeline (i.e., scheduling flaws). Each iteration of the planning process detects flaws on the current partial plan, selects the flaw to solve and refines the plan by applying possible solutions. A solution is found when the "current plan" does not contain any flaw i.e., the timelines of the plan are valid and complete. Algorithm 1 summarizes this refinement procedure.

Two decision points of Algorithm 1 are particularly crucial. One is the selection of the flaw to solve to refine the "current plan". Although this is not a backtracking decision, it strongly affects efficiency and determines the way planning and scheduling refinements are interleaved. Heuristics $\mathcal{H}_\phi$ encapsulates evaluation criteria that help Algorithm 1 to select "the most promising" flaw to solve among those found on the current (partial) plan.

Another decision point concerns the selection of the plan to extract from the fringe. Heuristics $\mathcal{H}_\pi$ in this case encapsulates criteria that help Algorithm 1 to evaluate the partial plans that compose the fringe and select the "most promising plan" to refine. This is a backtracking decision and is particularly relevant with respect to the quality of synthesized plans.

[2] https://github.com/pstlab/PLATINUm

Algorithm 1 Timeline-based plan synthesis

Input: $\mathcal{SV}$, $\mathcal{S}$, $\mathcal{H}_\pi$, $\mathcal{H}_\phi$
Output: $\pi = (FTL, R)$

$\Pi \leftarrow \emptyset$
$\pi \leftarrow$ initialize$(\mathcal{SV}, \mathcal{S})$
while $\neg$ isSolution$(\pi, \mathcal{SV}, \mathcal{S})$ do
  $\Phi \leftarrow$ flaws$(\pi, \mathcal{SV}, \mathcal{S})$
  $\Phi^* \leftarrow$ chooseFlaws $(\Phi, \mathcal{H}_\phi)$
  for $\phi \in \Phi^*$ do
    $\Pi \leftarrow$ refine $(\pi, \phi.$resolvers$)$
  $\pi \leftarrow$ choosePlan $(\Pi, \mathcal{H}_\pi)$
return π

3.2.3. Pareto-based Heuristic Search

To reason on the different clinical and technical aspects of the problem the search of the solving process follows a multi-objective approach. Specifically, we have extended PLATINUm by integrating a Pareto-based heuristic search to support heterogeneous metrics and objective functions. Algorithm 1 thus relies on a multi-objective heuristic $\mathcal{H}_\pi$ that evaluates and compares partial plans π_i by taking into account clinical feature. Similarly to other works [23, 24], for each metric j an evaluation function $f_j(\pi_i)$ is computed as the sum of two elements:

$$f_j(\pi_i) = g_j(\pi_i) + h_j(\pi_i) \tag{1}$$

A cost element $g_j(\pi_i)$ estimates a metric j on the consolidated part of the partial plan π_i. It takes into account the tokens that are part of the timelines of the plan π_i. A heuristic element $h_j(\pi_i)$ instead estimates a metric j on the possible refinements of the timelines of π_i. It analyzes pending goals of the partial plan π_i (i.e., the agenda) and evaluates the metric j on the possible projections of the timelines (i.e., tokens that could be added to the timelines in future refinements). Since projections represent alternative refinements and thus alternative behaviors, the heuristic value $h_j(\pi_i)$ compute the minimum and maximum estimated values of the projections. The minimum estimated value is generally considered in order to guarantee the admissibility of the heuristic and thus do not overestimate the actual cost of refinements.

Equation 1 is used by $\mathcal{H}$ to evaluate all the considered metrics j on all the partial plans π_i of the fringe. The concept of Pareto dominance is used to establish relationships between partial plans.

Definition 5. Given a set $\mathcal{F} = \{f_1, ..., f_n\}$ of objective functions that a (timeline-based) planning process aims at optimizing, a partial plan π_i is said to dominate a partial plan π_j (with $i \neq j$) if

$$\forall\ f_k \in \mathcal{F}, f_k(\pi_i) \bowtie f_k(\pi_j)$$

where $\bowtie = \{<\}$ in the case that function f_k should be minimized and $\bowtie = \{>\}$ in the case that function f_k should be maximized.

The concept of Pareto dominance guides the search towards the Pareto set of the fringe i.e., the subset of (not-explored) partial plans that do not dominate each other. Such partial plans represent suitable trade-offs with respect to the considered objective functions. The choice among plans belonging to the Pareto set and thus the selection of a specific solution is then made by "prioritizing" objective functions (solution polarization).

4. Results

4.1. A Planner for SI-Robotics

Given the described timeline-based framework, this section further describes the defined planning model and the objective functions that support the desired clinical objectives.

4.2. Modeling "Choreography" Constraints

The domain specification consists of a number of state variables characterizing the structure of a rehabilitation exercise and a number of synchronization rules characterizing possible combinations of dancing steps.

Three types of state variable are defined. A goal state variable describes high-level planning requests with a specified song duration and song rhythm expressed in bpm (beats per minute). A level state variable describes the logical steps composing rehabilitation exercises. The values of this state variable "sample" the whole exercise in a number of blocks or logic steps that are necessary to contextualize the structure of the exercise to the selected rhythm of a song. The number of logical steps that will compose a plan (i.e., a rehabilitation exercise) depends on the duration of the song and its bpm. Consider for example a song with a duration of 2 minutes and 60 bpm. In the case of blocks with a fixed duration of 10 bpm a total number of 12 blocks (or logical steps) will compose the exercise. Specifying 120 bpm with the same song duration and "block size" instead the exercise will be composed by 24 blocks. In general, the values of this state variable can be distinguished between "simple steps" that are directly mapped to known (primitive) steps, and "complex steps" that are mapped to patterns specifically designed by therapists. A step state variable then describes all the primitive steps known by the planner. Synchronization rules map each simple step to all known primitive steps and each complex step to the combinations of primitive steps defined by therapists. At this level each rule represents a disjunctive choice and allow a planner to reason about alternative rehabilitation exercises. Considering the dataset defined through the workflow of Figure 1, each simple step is characterized by a branching factor of 134 (alternative) choices of implementation/decomposition.

In addition to these state variables two consumable resources are defined to model and reason about spatial constraints of the layout. The vertical and horizontal axes of Figure 2 are modeled as two consumable resources with maximum capacity $2N$ and $2M$

Table 3
Aggregated results of the configured planning instances

	Clinical features			Solving time by layout (s)				Solving time by song duration (s)			Success Rate
	balance	coordination	energy	4×4	6×6	8×8	10×10	60	90	120	
OB1	46,63	30,98	22,35	131,76	87,93	7,40	7,13	5,79	38,64	174,62	97,92%
OB2	38,19	24,06	18,73	103,32	86,33	7,04	7,44	5,80	31,11	72,91	85,42%
DFS	37,53	22,33	16,67	54,04	38,28	6,78	7,35	6,30	31,37	64,96	91,67%

respectively. The levels of the two resources implicitly map the current position of a user and thus the amount of space available in the vertical and horizontal areas of the layout. Let us consider for example a horizontal axis resource with maximum capacity $M = 4$ and level $l = 3$. This means that the user has $l = 3$ units of space available on the left-side of the layout and $M - l = 1$ unit of space available on the right-side of the layout. The same applies to the vertical axis resource.

According to the body orientation of a user, each primitive step (i.e., each motion) consumes or produces a certain amount of resource according to the required body translation (see the technical features of Table 1). Let us consider a user with 0° orientation and current position encoded by resource levels $l_v = 1$ (vertical axis) and $l_h = 3$ (horizontal axis). A motion requiring the user to "make a leftward step" is modeled as an activity consuming the horizontal axis resource of 1 unit. The execution of such a step would decrease the level of the resource from $l_h = 3$ to $l_h = 2$.

Steps that do not require an actual body translation but "only" some spatial requirement (see again the technical features of Table 1) are treated in a similar way by combining productions and consumptions of associated resources. Consider a user with 0° orientation and current position encoded by resource levels $l_v = 1$ (vertical axis) and $l_h = 3$ (horizontal axis). A motion requiring the user to "raise (and lower) the right arm to the shoulder" is modeled as an activity that produces the horizontal axis resource of 1 unit during its execution but that consumes the same amount of resource at its end. Namely, the level of the horizontal axis resource is raised to $l_h = 4$ during the execution of the step and lowered to $l_h = 3$ at its end (i.e., restoring the starting state of the resource). Such resource constraints are modeled by means of dedicated synchronization rules (one for each dancing step).

4.3. Reasoning on Clinical Features

The rules specified into the planning model allows the planner to synthesize sequences of dancing steps that are correct with respect to the spatial constraints of the layout and the expected duration (and timing) of the selected song. The planner then should decide which sub-set of steps best fits the specific health-related needs of a particular rehabilitation session. To this aim, the search strategy of the planner should be able to combine and evaluate different metrics (i.e., different clinical features of Table 2) according to the clinical objective selected by the therapist.

From an optimization perspective the planning problem can be described by: a horizon

$\mathcal{H} \in \mathbb{N}^+$ defined according to the specified song duration; a set $\mathcal{L} = 1, ..., n$ of expected blocks' IDs denoting the logical steps (defined according to the specified duration and bpm of the song) and; a set $\mathcal{S} = 1, .., n$ of known primitive steps' IDs. For each block $i \in \mathcal{L}$ and step $j \in \mathcal{S}$ a binary decision variable is defined as follows:

$$s_i^j = \begin{cases} 1, & \text{if step } j \text{ is selected during block } i \\ 0, & \text{otherwise} \end{cases} \tag{2}$$

the assignment of values to decision variables s_i^j should be such that the duration of the resulting plan "covers" the entire duration of the song.

$$\sum_i \sum_j d^j s_i^j \geq \mathcal{H} \tag{3}$$

where $d^j \in \mathbb{N}$ is the duration of step $j \in \mathcal{S}$ as specified in the dataset (see the technical features of Table 1).

Taking into account the clinical features of Table 2, clinical objectives are encoded as two different multi-objective functions. The objective concerning the stimulation of physical equilibrium can be seen as the problem of maximizing the "effects" of the rehabilitation exercise on the overall cognitive/motor aspects of the users as well as maintenance of body equilibrium.

$$\underset{\pi}{\text{maximize}} \{f_b, f_c\} \tag{4}$$

where functions f_b and f_c denote respectively the cumulative impact of a plan π on body balance and body coordination of users.

$$f_b = \sum_i \sum_j s_i^j \; balance(j) \tag{5}$$

$$f_c = \sum_i \sum_j s_i^j \; coord(j) \tag{6}$$

The objective concerning the stimulation of metabolic response can be seen as the problem of maximizing the "effects" of the rehabilitation exercise on the metabolism of the users while keeping the difficulty of the coordination low.

$$\underset{\pi}{\text{maximize}} \{f_e\}, \underset{\pi}{\text{minimize}} \{f_c\} \tag{7}$$

where functions f_e denotes the cumulative energy demand of a plan π.

$$f_e = \sum_i \sum_j s_i^j \; energy(j) \tag{8}$$

These objective functions determine the way a planner should reason on the clinical features of plans. These objective functions are therefore encapsulated into the planning framework as different search strategies $\mathcal{H}_\pi$ in order to guide the search of Algorithm 1 towards the solutions that best fit the desired clinical objective.

5. Discussion

We have made an experimental evaluation of the developed planning framework with the objective of assessing the efficacy of synthesized plans. To this aim we have defined a number of planning problems by varying the following domain parameters: (i) song time with values 60 seconds, 90 seconds and 120 seconds; (ii) song bpm with values 60 bpm and 120 bpm; (iii) layout dimension with values $N \times M$ (i.e., the capacity of the vertical/horizontal axis resources) equal to $\{ 4 \times 4, 6 \times 6, 8 \times 8$ and $10 \times 10 \}$.

For each problem we have configured and run three planning instances: (i) a planning instance labeled with $OB1$ supports stimulation of physical equilibrium (i.e.,Equation 4); (ii) a planning instance labeled with $OB2$ supports stimulation of metabolic response (i.e., Equation 7); (iii) a last planning instance labeled with DFS does not support a specific objective in order to show the behavior of a planner without a specific heuristic guide. A number of 24 experiments for each planning configuration has been run (i.e., a total of 72 experiments). Table 3 shows average results out of 3 repetitions of each experiment and a timeout of 3 minutes (i.e., 180 seconds).

Results show initial but promising results concerning the capability of the planning framework to deal with both spatial constraints and clinical requirements. A first part of results show and compare the qualities of plans generated by the three planning instances. Specifically, Table 3 shows the average scores of the discussed clinical features. As expected, planner DFS without heuristics does not synthesize plans that are effective with respect to the desired clinical objectives. In this case indeed plans do not show neither the best value of balance nor the best value of energy. The planner $OB1$ instead effectively addresses the clinical objective concerning stimulation of physical equilibrium. This configuration synthesizes plans with the highest values of balance and coordination.

The performance of planner $OB2$ instead is not effective as expected. As can be seen from Table 3 indeed the planner does not achieve the best level of energy which is in particular lower than the value achieved by planner $OB1$. Also, planner $OB2$ is the configuration with the lowest success rate denoting a difficulty of the heuristic in effectively guiding the search to achieve the associated clinical objective (Equation 7). As a general observation, solving behavior and the efficacy of the heuristics strongly depend on the encoded domain knowledge and thus the scores assigned to the steps of the dataset. In this regard, the behavior of $OB2$ (compared to the behavior of $OB1$) is clear signal of the need for refining domain knowledge (see the workflow of Figure 1) in order to better characterize and differentiate the “impact” of known stimuli on users.

For what concern the technical requirements and planning time, Table 3 shows how layout and song duration affect solving performance. In particular, the parameters of the layout determine capacity constraints of the underlying reservoir resources that strongly affect the reasoning overhead due to consistency check of resource profiles. As can be seen indeed a lower resource capacity (i.e., smaller layout) entails a harder planning problem since a higher number of resource peaks (i.e., resource-related flaws) is detected. Similarly a higher song bpm entails a harder planning problem since the planner should consider a higher number of steps.

Although initial and with some room for improvements (especially concerning planner

*OB*2), the experimental assessment shows promising results. Experiments show the capability of “tailoring” the general-purpose search of the (timeline-based) planning framework to different clinical objectives. These results therefore represent a solid basis to further improve the reasoning capabilities of the framework and incrementally address a wider set of clinical objectives and features.

6. Conclusions

The paper has described an ongoing work about the use of timeline-based planning for the synthesis of personalized rehabilitation exercise programs. An experimental evaluation has shown the technical feasibility of the developed framework. Next steps will focus on improving the solving capabilities of the framework and performing experiments with real users. In this regard, this work lays the foundations for the integration of learning mechanisms suitable to “specialize” and improve solving performance. We plan to integrate data from real users to automatically synthesize “specialized heuristics” and pruning mechanism to reduce planning time.

Acknowledgments

This work is supported by “SI-Robotics: SocIal ROBOTICS for active and healthy ageing” project (Italian M.I.U.R., PON – Ricerca e Innovazione 2014-2020 – G.A. ARS01_01120).

References

[1] A. W. Group, The 2021 ageing report: Economic budgetary projections for the eu member states (2019-2070), 2021.

[2] O. Vermesan, A. Bröring, E. Tragos, M. Serrano, D. Bacciu, S. Chessa, C. Gallicchio, A. Micheli, M. Dragone, A. Saffiotti, P. Simoens, F. Cavallo, R. Bahr, Internet of robotic things : converging sensing/actuating, hypoconnectivity, artificial intelligence and IoT Platforms, in: Cognitive hyperconnected digital transformation : internet of things intelligence evolution, 2017, pp. 1–35.

[3] M. El Kamali, L. Angelini, M. Caon, G. Andreoni, O. A. Khaled, E. Mugellini, Towards the NESTORE e-Coach, Proceedings of the 2018 ACM International Joint Conference and 2018 International Symposium on Pervasive and Ubiquitous Computing and Wearable Computers - UbiComp ’18 (2018).

[4] D. Casey, H. Felzmann, G. Pegman, C. Kouroupetroglou, K. Murphy, A. Koumpis, S. Whelan, What People with Dementia Want: Designing MARIO an Acceptable Robot Companion, in: K. Miesenberger, C. Bühler, P. Penaz (Eds.), Computers Helping People with Special Needs, Springer, 2016, pp. 318–325.

[5] S. Coradeschi, A. Cesta, G. Cortellessa, L. Coraci, J. Gonzalez, L. Karlsson, F. Furfari, A. Loutfi, A. Orlandini, F. Palumbo, F. Pecora, S. von Rump, A. Štimec, J. Ullberg, B. Ötslund, GiraffPlus: Combining social interaction and long term monitoring

for promoting independent living, in: The 6th International Conference on Human System Interactions (HSI), 2013, pp. 578–585.

[6] A. Umbrico, A. Cesta, G. Cortellessa, A. Orlandini, A holistic approach to behavior adaptation for socially assistive robots, International Journal of Social Robotics (2020).

[7] R. De Benedictis, A. Umbrico, F. Fracasso, G. Cortellessa, A. Orlandini, A. Cesta, A Two-Layered Approach to Adaptive Dialogues for Robotic Assistance, in: RO-MAN. The 29th IEEE International Symposium on Robot and Human Interactive Communication, 2020.

[8] R. de la Torre, S. de Sola, G. Hernandez, M. Farré, J. Pujol, J. Rodriguez, J. M. Espadaler, K. Langohr, A. Cuenca-Royo, A. Principe, L. Xicota, N. Janel, S. Catuara-Solarz, G. Sanchez-Benavides, H. Bléhaut, I. Dueñas-Espín, L. del Hoyo, B. Benejam, L. Blanco-Hinojo, S. Videla, M. Fitó, J. M. Delabar, M. Dierssen, Safety and efficacy of cognitive training plus epigallocatechin-3-gallate in young adults with down's syndrome (tesdad): a double-blind, randomised, placebo-controlled, phase 2 trial, The Lancet Neurology 15 (2016) 801 – 810. doi:https://doi.org/10.1016/S1474-4422(16)30034-5.

[9] G. Levy, The Relationship of Parkinson Disease With Aging, Archives of Neurology 64 (2007) 1242–1246. doi:10.1001/archneur.64.9.1242.

[10] K. R. Chaudhuri, N. Titova, Societal Burden and Persisting Unmet Needs of Parkinson's Disease, European Neurological Review 14 (2019) 28–35.

[11] S. Keus, E. Hendriks, B. Bloem, KNGF guidelines for physical therapy in patients with Parkinson's disease, Dutch J Physiother 114 (2004).

[12] A. G. Robinson, A. M. Dennett, D. A. Snowdon, Treadmill training may be an effective form of task-specific training for improving mobility in people with Parkinson and disease and multiple sclerosis: a systematic review and meta-analysis, Physiotherapy 105 (2019) 174–186.

[13] R. Bevilacqua, E. Maranesi, M. Di Rosa, R. Luzi, E. Casoni, N. Rinaldi, R. Baldoni, F. Lattanzio, V. Di Donna, G. Pelliccioni, G. R. Riccardi, Rehabilitation of older people with Parkinson's disease: an innovative protocol for RCT study to evaluate the potential of robotic-based technologies, BMC neurology 20 (2020) 186–186. doi:10.1186/s12883-020-01759-4.

[14] M. dos Santos Delabary, I. G. Komeroski, E. P. Monteiro, R. R. Costa, A. N. Haas, Effects of dance practice on functional mobility, motor symptoms and quality of life in people with parkinson's disease: a systematic review with meta-analysis, Aging Clinical and Experimental Research 30 (2018) 727–735. doi:10.1007/s40520-017-0836-2.

[15] A. M. Carapellotti, R. Stevenson, M. Doumas, The efficacy of dance for improving motor impairments, non-motor symptoms, and quality of life in Parkinson's disease: A systematic review and meta-analysis, PLOS ONE 15 (2020) 1–28.

[16] D. Volpe, M. Signorini, A. Marchetto, T. Lynch, M. E. Morris, A comparison of Irish set dancing and exercises for people with Parkinson's disease: A phase II feasibility study, BMC Geriatrics 13 (2013) 54.

[17] N. Muscettola, HSTS: Integrating Planning and Scheduling, in: Zweben, M. and Fox, M.S. (Ed.), Intelligent Scheduling, Morgan Kauffmann, 1994.

[18] A. Cesta, G. Cortellessa, S. Fratini, A. Oddi, Developing an End-to-End Planning Application from a Timeline Representation Framework, in: IAAI-09. Proc. of the 21^{st} Innovative Application of Artificial Intelligence Conference, Pasadena, CA, USA, 2009.
[19] A. Jonsson, P. Morris, N. Muscettola, K. Rajan, B. Smith, Planning in Interplanetary Space: Theory and Practice, in: AIPS-00. Proceedings of the Fifth Int. Conf. on AI Planning and Scheduling, 2000.
[20] A. Umbrico, A. Cesta, M. Cialdea Mayer, A. Orlandini, Integrating resource management and timeline-based planning, in: The 28th International Conference on Automated Planning and Scheduling (ICAPS), 2018.
[21] M. Cialdea Mayer, A. Orlandini, A. Umbrico, Planning and execution with flexible timelines: a formal account, Acta Informatica 53 (2016) 649–680.
[22] P. H. Morris, N. Muscettola, T. Vidal, Dynamic Control of Plans With Temporal Uncertainty, in: International Joint Conference on Artificial Intelligence (IJCAI), 2001, pp. 494–502.
[23] M. Khouadjia, M. Schoenauer, V. Vidal, J. Dréo, P. Savéant, Pareto-based multiobjective ai planning, in: Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, AAAI Press, 2013, p. 2321–2327.
[24] M. Do, S. Kambhampati, Sapa: A multi-objective metric temporal planner, Journal of Artificial Intelligence Research 20 (2003) 155–194.