

Unified Learning of Temporal Task Structure and Action Timing for Bimanual Robot Manipulation

Christian Dreher, Patrick Dormanns, Andre Meixner, and Tamim Asfour

Abstract—Bimanual manipulation requires both temporal task structure – which actions precede or overlap others – and concrete timing – when each action starts and how long it takes. Symbolic relations enable high-level reasoning about task structure and alternative execution sequences, while concrete timing parameters are essential for coordinating two hands at the execution level. Existing approaches address these two levels in isolation, leaving a gap between task planning and movement synchronization. This work presents an approach for learning symbolic and subsymbolic temporal task constraints from human demonstrations and deriving executable, temporally parametrized plans for bimanual manipulation. Specifically, our contributions are (i) a 3-D representation of timings with methods based on multivariate Gaussian Mixture Models to represent temporal relationships between actions on a subsymbolic level, (ii) a method based on the Davis–Putnam–Logemann–Loveland algorithm that finds and ranks all contradiction-free Allen relation assignments of a task, representing its modes, and (iii) an optimization-based planning system combining the identified symbolic and subsymbolic temporal task constraints to derive parametrized plans for robot execution. We quantitatively evaluate our approach and show that the inferred symbolic task constraints are more accurate than those of prior heuristic approaches and that the generated plans are temporally closer to human demonstrations than the most characteristic demonstration. Additionally, we qualitatively demonstrate its applicability through orchestrated task executions in simulation and on real robots.

I. INTRODUCTION

Temporal task structure and action timing are fundamental to bimanual manipulation. When unscrewing a component, it must be held *during* unscrewing, not *before* or *after*. When pouring liquid into a container, the pouring motion should *start* as the container is in position. These temporal relationships define the logical structure of bimanual tasks. Yet knowing that one action occurs *during* another is insufficient for execution – we also need concrete timings: *how long* does each action take, and *exactly when* should each action start? *Symbolic temporal task constraints* between two actions can be modeled as qualitative relations, specifically as Allen relations [1], such as *before*, *overlaps*, or *during*. They allow the robot to reason about alternative execution sequences, generate plans, and generalize to new situations. For instance, knowing that ‘transfer object’ must occur *after* ‘grasp object’ but can occur *before* or *during* ‘open container’ enables flexible and context-aware scheduling. *Subsymbolic temporal*

The research leading to these results has received funding from the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – SFB 1574 and the German Federal Ministry of Research, Technology and Space (BMFTR) under the Robotics Institute Germany (RIG). The authors are with the Institute for Anthropomatics and Robotics, Karlsruhe Institute of Technology, Karlsruhe, Germany. {c.dreher, asfour}@kit.edu




Fig. 1: Overview of our approach and contributions: *Temporal Relationship Assessment* evaluates all Allen relations observed in human demonstrations, *Temporal Task Constraint Inference* finds contradiction-free Allen relation assignments and most likely action timings, and *Temporal Planning* generates executable, temporally parametrized bimanual plans.

task constraints provide concrete durations, delays, and offsets that determine the execution quality – parameters to which we refer to as *timing*. They specify whether the robot waits 200 ms for stabilization before pouring, pours for 3 s or 5 s, and how precisely the hands synchronize.

This distinction between symbolic and subsymbolic temporal information reflects two fundamentally different levels of a robot’s temporal reasoning. Yet despite their complementary roles, existing approaches largely address them in isolation. Methods for learning task structure from demonstration focus on qualitative ordering and precedence [2]–[4], while methods for movement synchronization treat timing as a low-level control problem decoupled from task-level reasoning [5]–[10]. Our previous work [11] addressed bridging this gap by learning both symbolic and subsymbolic temporal task constraints from demonstration. However, it models the temporal relationship between each pair of action keypoints independently using univariate Gaussian Mixture Models (GMMs), capturing only marginal distributions and missing

the joint structure of action timings. Additionally, it is only capable of finding one assignment of Allen relations to each pair of actions in a task and does not include a planning system to generate executable plans for the whole task.

In this work, we present a unified approach that learns both symbolic and subsymbolic temporal task constraints from human demonstrations and uses them to generate temporally parametrized plans for robot execution. Our approach consists of three parts, illustrated in Fig. 1. First, we introduce a 3-D embedding of timings between two actions, the *timing space*, and propose multivariate GMM-based methods that jointly model the lengths of two actions and their relative offset to capture their temporal relationship. Second, we propose an approach based on the Davis–Putnam–Logemann–Loveland (DPLL) algorithm [12] that finds and ranks all contradiction-free Allen relation assignments for a task, enabling the identification of multiple task modes. Third, we develop an optimization-based planning system that takes the inferred symbolic and subsymbolic task constraints as input and produces a fully temporally parametrized plan that best reflects the human demonstrations. We quantitatively evaluate our approach using the KIT Bimanual Actions Dataset (Bimacs dataset) [13] and show that the inferred symbolic task constraints are more accurate than those of prior heuristic approaches and that the generated plans are temporally closer to human demonstrations than the most characteristic demonstration. Additionally, we qualitatively demonstrate its applicability through orchestrated task executions in simulation and on real robots using the Bimacs dataset and KIT Bimanual Manipulation Dataset (BiManip dataset) [14].

II. RELATED WORK

Temporal structure and timing are central to bimanual manipulation, yet the literature largely addresses the symbolic and subsymbolic levels in isolation. Symbolic temporal task constraints are important for task planning and are discussed in Section II-A. Subsymbolic temporal task constraints are needed for synchronized execution of bimanual tasks and are discussed in Section II-B. Apart from the work in [11], none of these works addresses the simultaneous learning of symbolic and subsymbolic temporal task constraints.

A. High-level: Task Modes and Temporal Task Structure

Task modes are different sequences of key actions that lead to the same goal through varying orderings. A *key action* is any action that changes the state of the primarily manipulated objects, such as a disassembly action separating two parts. This is referred to as *temporal task structure* [15]. Key actions stand in contrast to supportive actions like approaching objects for grasping or lifting objects for key actions. In the following, we discuss works on precedence relations and those that need more complex relations to handle bimanuality or simultaneous actions.

1) *Precedence Relations*: Several early works proposed learning different task modes for unimanual tasks, or tasks with one execution sequence [2]–[4]. Carpio et al. [16] propose a method based on Bayesian Networks to learn

the sequential temporal task structure in the context of human robot interaction. While effective for unimanual tasks, precedence-based representations do not scale naturally to bimanual tasks: two concurrent execution sequences create a combinatorial explosion of orderings, and slight qualitative changes in supportive actions do not constitute a fundamentally different task mode.

2) *Complex Relations*: Allen [1] systematically defines all 13 possible relations in which two intervals can be and provides multiple deductive methods e.g., to derive transitive relations. Ye et al. [17] consider parallel execution of independent actions using precedence graphs extended with concurrency, but this does not account for the complex temporal relationships between hands that bimanual manipulation requires. Dreher and Asfour [13] propose an inductive graph-based heuristic search method to learn symbolic temporal task constraints in bimanual manipulation tasks that are needed to understand task structure from partially contradictory demonstrations. Our DPLL-based algorithm builds on and extends this work by performing an exhaustive rather than heuristic search, finding and ranking all contradiction-free Allen relation assignments rather than a single most likely one. More recently, large language models and foundation models have been applied to high-level task planning [18]. However, these approaches operate at the level of symbolic action sequences and do not address subsymbolic timing – they cannot specify how long an action should take or how two concurrent actions should be temporally offset.

B. Low-level: MP Synchronization and Coupling

On a lower level, bimanual coordination or synchronization can be achieved by *coupling* Movement Primitives (MPs). This term implies that the execution of one MP is intertwined with the execution of another MP through some mechanism, and thus the synchronization is *emerging* from it. Another approach is to treat MPs as black boxes and parametrize the offsets between the two MPs of a bimanual action, as well as their duration with a targeted execution length. In this case, the synchronization is *assigned* at a higher level.

1) *Emerging Synchronization (Bottom-Up)*: Thota et al. [5] propose control laws to couple Dynamic Movement Primitives (DMPs) by feeding back the tracking error of one arm to the other arm in contrast to feeding it back to the arm itself to track the desired trajectory. For this work, the formulation of DMPs had to be adapted. Kulvicius et al. [6] propose learning weights of virtual spring coupling between two end-effectors controlled by DMPs, allowing flexible coordination to emerge from learned interaction forces. Gams et al. [7] propose to couple DMPs in bimanual tasks with force/torque feedback learned during the robot’s execution. This was later extended to discrete-periodic motion on a humanoid robot, where phase coupling between primitives drives synchronization [8]. Karlsson et al. [9] introduce temporal coupling between DMPs through a shared phase variable, ensuring convergence of both primitives regardless of execution speed variations. Although these bottom-up approaches yield robust synchronization, they require task-

specific coupling designs and do not generalize the temporal structure of a task across different execution modes or new situations. A different line of work bypasses explicit coupling altogether by learning end-to-end bimanual policies from demonstrations, e. g., transformer-based action-chunking [19]. Such approaches are reactive, but do not expose interpretable task constraints and require many demonstrations.

2) *Assigned Synchronization (Top-Down)*: An alternative is to treat MPs as black boxes and assign synchronization parameters (offsets and durations) from a higher level. The works in [20] and [10] achieve this in the context of bimanual manipulation using pre-defined Petri nets that encode the temporal structure of bimanual manipulation categories. These approaches are expressive, but require the temporal structure to be hand-designed rather than learned. The work in [11] is closest to our work. The authors propose a method for extracting symbolic and subsymbolic temporal task constraints from human demonstrations, where univariate GMMs are used to learn temporal differences between semantic action keypoints (starts and ends of actions). This approach is data-efficient but considers each pair of keypoints independently, missing the joint structure of action timings. In particular, it cannot capture correlations between action lengths and offsets, which are critical for generating consistent parametrized plans. In this work, we propose a multivariate GMM formulation to fully capture the temporal relationship between two actions.

To our knowledge, no existing approach simultaneously (i) learns symbolic temporal task structure at the level of Allen relations from demonstrations that may exhibit multiple task modes, (ii) models the full joint distribution of action timings using a principled representation, and (iii) integrates both levels into a single planning system that produces temporally parametrized plans for execution.

III. PROBLEM STATEMENT

In this section, we first introduce basic formalisms and definitions required in the scope of this work.

A. Definition of Task, Action, Demonstration

The following definitions are largely based on our previous work [11] with minor modifications and additions. A bimanual manipulation *task* consists of a set of actions $\mathcal{A}$, with $\mathcal{A} = \{\alpha_a, \alpha_b, \dots\}$ and where an *action* α is a tuple $\alpha = (v, o)$, with v denoting a verb and o an object. For actions enriched with temporal information – relevant when observing or reproducing an action with concrete start and end – the notation $a = (v, o, i)$ is used, where $i = [i^-, i^+]$ is a time interval with start i^- and end $i^+ > i^-$. We write $a^- = i^-$ and $a^+ = i^+$ to refer to the two semantic temporal keypoints of the action and define its *length* as $|a| = a^+ - a^-$. The *offset* ω_{ab} between two actions a_a and a_b is defined as the distance between their midpoints: $\omega_{ab} = \frac{1}{2}(a_b^- + a_b^+) - \frac{1}{2}(a_a^- + a_a^+)$. Let $\mathcal{D}$ be a set of demonstrations of a bimanual manipulation task. A *demonstration* $D \in \mathcal{D}$ is a tuple $D = (S_L, S_R)$ consisting of two action sequences S_L for the left hand, and S_R for the right hand, respectively. An *action sequence* S is an ordered collection of actions $S = (a_a, a_b, a_c, \dots)$.

Actions $a_a, a_b, a_c, \dots$ in an action sequence $S = (a_a, a_b, a_c, \dots)$ have a strict order and cannot overlap, thus it follows that $a_a^- < a_a^+ \leq a_b^- < a_b^+ \leq a_c^- \dots$ etc. A task consists of several subtasks, i. e., subsets of actions showing high temporal cohesion (symbolically or subsymbolically) [13]. A set of demonstrations $\mathcal{D}$ may exhibit multiple *task modes*, which are qualitatively different but equally valid orderings of subtasks. Additionally, each subtask may show different modes that manifest as slight variations of concurrent temporal relations, but rarely different sequences. Both types of modes reflect the natural variation in how humans perform the same task.

B. Definition of Task Assignments

A single set of demonstrations may contain various modes. This results in apparent contradictions when temporal relations between action pairs are unconditionally assessed. A robot needs to find one arrangement of all the actions in the task on a symbolic level so that they are consistent as a whole, and most likely given the demonstrations. To this end, we define a *task assignment* $t = (B, s_t)$, where B is a contradiction-free complete assignment of one single Allen relation $R \in R_A$ to each action pair $(\alpha_a, \alpha_b) \in \mathcal{A} \times \mathcal{A} : \alpha_a \neq \alpha_b$, R_A is the set of all 13 Allen relations [1], and a score s_t quantifying how well the assignment reflects the observed demonstrations. Formally, B is a set of 4-tuples $b = (\alpha_a, \alpha_b, R, s)$, with $(\alpha_a, \alpha_b) \in \mathcal{A} \times \mathcal{A} : \alpha_a \neq \alpha_b$ being an action pair where R is the Allen relation assigned to the action pair (α_a, α_b) , and $s \in [0, 1]$ being a value that expresses the degree to which the relation $\alpha_a R \alpha_b$ is supported by the demonstrations. We also consider a *partial task assignment* $\tilde{t}$ as a tuple $\tilde{t} = (B^?, B)$ as a means to obtain a complete task assignment t in the proposed algorithms. Here, B is the set of concrete assignments of Allen relations to action pairs and $B^?$ is the set of all action pairs that have not yet been assigned. An unassigned action pair $b^? \in B^?$ is a tuple $b^? = (\alpha_a, \alpha_b)$ with $(\alpha_a, \alpha_b) \in \mathcal{A} \times \mathcal{A} : \alpha_a \neq \alpha_b$. A task assignment makes a temporal relation a symbolic temporal task constraint because it guarantees that a single relation will not cause contradictions. In this work, our first aim is to identify all contradiction-free task assignments and rank them by their score s_t , i. e., likelihood given the demonstrations.

C. Definition of Timing

A task assignment captures the qualitative temporal structure of a task, but is insufficient for robot execution, which requires concrete timing parameters. We define the *timing* τ between two actions a_a and a_b as a compact representation of their concrete temporal relationship. Throughout this work, two equivalent representations will be used. Naively, the timing between the two actions can be represented as 4-D vector $\tau^{(4)} = (a_a^-, a_a^+, a_b^-, a_b^+)$, $\tau^{(4)} \in \mathbb{T}^4$, with the components being the starts and ends of both actions. However, a 3-D vector representation $\tau^{(3)} = (\lambda_a, \lambda_b, \omega_{ab})$, $\tau^{(3)} \in \mathbb{T}^3$ can also be used, where λ_a and λ_b correspond to the lengths of actions a_a and a_b respectively, and ω_{ab} to the offset between both actions. In the context of this work, these representations $\tau^{(4)}$ and $\tau^{(3)}$ can be used interchangeably, as they can be converted into each other. The $\tau^{(3)}$ form can be preferable for learning since it is

invariant to uniform time shifts. If the concrete representation is not important, we use the notation τ for brevity. To learn from several demonstrations, we need a model that represents the timing between two actions. Thus, we define a *timing model* as a multivariate GMM in timing space. A subsymbolic temporal task constraint is the most likely timing between two actions given a specific symbolic temporal task constraint, i.e., given the assigned Allen relation. In this work, our second aim is to infer subsymbolic temporal task constraints from timing models.

D. Definition of Plan and Temporally Parametrized Plan

A *plan* in the context of this work is a pair of action sequences (S_L, S_R) for the left and right hand, in which actions are qualitatively arranged to satisfy symbolic temporal task constraints. Since it is the output of a symbolic planner, it exhibits discrete and artificial timings. A *temporally parametrized plan* extends this by assigning concrete lengths and offsets to all actions, consistent with the symbolic constraints and as close as possible to the inferred subsymbolic temporal task constraints. Both are denoted by $P = (S_L, S_R)$, are equivalent to a demonstration, and differ only in integer or real timing values. In this work, our third aim is to derive a temporally parametrized plan given identified symbolic and subsymbolic temporal task constraints for robot execution.

IV. APPROACH

This section describes our approach with its three parts also illustrated in Fig. 1.

A. Temporal Relationship Assessment

Assessing temporal relationships between actions symbolically and subsymbolically is the cornerstone of our approach.

1) *Assessing Temporal Relations*: Temporal relations provide a qualitative source of temporal information needed in our approach. To assess which temporal relations between two actions are most likely to hold given all observed demonstrations, we adopt the method described in [11], which quantifies the degree to which each Allen relation holds between each action pair using univariate GMMs and fuzzy logic. The result is a soft, graded assessment of temporal relations for every action pair in the task.

2) *Assessing Timing*: Beyond qualitative relations, execution requires concrete timing information. We represent the timing between two actions a_a and a_b of a task as a 3-D vector $(\lambda_a, \lambda_b, \omega_{ab})$, where λ_a and λ_b correspond to the lengths of the actions and ω_{ab} to their midpoints' offset. All observed timings between two actions are expressed as such vectors and embedded in timing space $\mathbb{T}^3$ to train multivariate Bayesian GMMs, whose variational formulation automatically determines the required number of mixture components. Fig. 2 shows a visualization of the timing space with several demonstrations of 'pour milk' and 'hold bowl' together with the corresponding multivariate GMM that models the timing distribution. Allen relations can conveniently be represented in timing space by lines, areas, or volumes, as shown in Fig. 3.

The 3-D space $\mathbb{T}^3$ is better suited for representing observations, since it reduces the problem space for learning by

ignoring shifts of both intervals by the same amount as is the case in $\mathbb{T}^4$, i.e., shifts along $(1,1,1,1)$. Consider the two timings $(1,2,3,4) \neq (2,3,4,5)$: In $\mathbb{T}^4$, they are two different points even though they describe the exact same relative temporal arrangement between actions. In $\mathbb{T}^3$, both will be represented as the same point, accurately reflecting only their relationship with each other. This is also a benefit when considering norms and distances in $\mathbb{T}^3$, which is important for our approach. However, by naively using the lengths of two actions, i.e., $(\lambda_a, \lambda_b, \omega_{ab}) = (|a_a|, |a_b|, \omega_{ab})$, we would render the Euclidean norm $\|\tau\|_2$ for $\tau \in \mathbb{T}^3$ meaningless. Instead, since we are only interested in relative temporal arrangements as just discussed, we want the Euclidean norm $\|\tau\|_2$ to be equal to the minimum Euclidean norm $\|\tau^{(4)}\|_2$ considering shifts by the same amount:

$$\|\tau^{(3)}\|_2 = \min_{\varsigma} \left\| \tau^{(4)} + \varsigma \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \right\|_2, \text{ with } \tau^{(3)} = \begin{pmatrix} \lambda_a \\ \lambda_b \\ \omega_{ab} \end{pmatrix}, \tau^{(4)} = \begin{pmatrix} a_a^- \\ a_a^+ \\ a_b^- \\ a_b^+ \end{pmatrix}. \quad (1)$$

More precisely, with a scalar s and $\lambda_a = s|a_a|$ and $\lambda_b = s|a_b|$, we require that

$$\left\| \begin{pmatrix} s|a_a| \\ s|a_b| \\ \omega_{ab} \end{pmatrix} \right\|_2 = \min_{\varsigma} \left\| \begin{pmatrix} a_a^- \\ a_a^+ \\ a_b^- \\ a_b^+ \end{pmatrix} + \varsigma \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \right\|_2. \quad (2)$$

Solving for s yields $s = \frac{1}{\sqrt{2}}$. It follows that we embed a timing into our 3-D space $\mathbb{T}^3$ as follows:

$$\tau^{(3)} = \left(\frac{1}{\sqrt{2}}|a_a|, \frac{1}{\sqrt{2}}|a_b|, \omega_{ab} \right)^T. \quad (3)$$

Thus, using this embedding, we obtain a meaningful Euclidean norm in $\mathbb{T}^3$ that replicates the Euclidean norm in $\mathbb{T}^4$ if shifts along $(1,1,1,1)$ are ignored. The distance $d(\tau_1, \tau_2)$ between two timings τ_1 and τ_2 is then defined as

$$d(\tau_1, \tau_2) = \|\tau_2^{(3)} - \tau_1^{(3)}\|_2 = \min_{\varsigma} \left\| \begin{pmatrix} a_{a_2}^- \\ a_{a_2}^+ \\ a_{b_2}^- \\ a_{b_2}^+ \end{pmatrix} - \begin{pmatrix} a_{a_1}^- \\ a_{a_1}^+ \\ a_{b_1}^- \\ a_{b_1}^+ \end{pmatrix} + \varsigma \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix} \right\|_2. \quad (4)$$

To conclude, our proposed 3-D timing space $\mathbb{T}^3$ is isomorphic to the quotient space $\mathbb{T}^4 / \langle (1,1,1,1) \rangle$. In this space, we train GMMs which form the basis of the subsymbolic temporal task constraints that we obtain later.

B. Temporal Task Constraint Inference

After assessing the relationships between actions, we need to infer symbolic and subsymbolic temporal task constraints that can be used for planning task executions on a robot.

1) *Infer Symbolic Temporal Task Constraints*: In the previous part (Section IV-A.1), temporal relations have been assessed on an action pair basis using our previous approach [11]. However, this information is a superposition of all demonstrations and is not very useful yet for the execution on a robot, as it contains many contradictions from differently observed timings or even different task modes.




Fig. 2: Timing model in T^3 , modeling the relationship between two actions. Black dots are observed timings from several human demonstrations. Ellipsoids denote modes of the multivariate GMM.




Fig. 3: Visualization of selected Allen relations in T^3 as line (*equals*), areas (e.g., *finished by*) or volumes (e.g., *before*). Points in these regions can be attributed to the respective Allen relation.




Fig. 4: Conditioning a timing model to the Allen relation *finished by* (blue area). The orange point is the maximum of the GMM's probability density function restricted to that area.

Thus, we propose a DPLL-based approach that is able to find all contradiction-free assignments of Allen relations to the task and rank them by likeliness. This is a hard problem, since it is a 13-ary satisfiability problem, thus being NP-complete. Algorithm 1 shows the general outline. An initial task assignment $\tilde{t}_i$ is created with $\tilde{t}_i = (B_i^?, \emptyset)$ on line 3. Here, $B_i^?$ is the exhaustive set of all action pairs in the task, i.e., the Cartesian product of the set of all actions in the task without identical pairs. With $\tilde{t}_i$ as the starting point, the set of all partial task assignments is initialized on line 4. The main loop on line 6 continuously pops a partial task assignment $\tilde{t}$ from $\tilde{T}$, from which next possible assignments will be derived if the task assignment is not complete yet. A task assignment is not complete if there is at least one action pair to which no Allen relation has been assigned yet (line 8). Specifically, on line 9, one specific action pair will be chosen and all feasible Allen relations will be assigned to it using Algorithm 2. This creates new partial task assignments for each possible Allen relation for the chosen action pair, which are then pushed to $\tilde{T}$. If a popped task assignment

Algorithm 1 Find all feasible contradiction-free (contr.-free) task assignments sorted by highest score first.

```

1: function FINDASSIGNMENTS
2:    $B_i^? \leftarrow \{(\alpha_a, \alpha_b) \in \mathcal{A} \times \mathcal{A} \mid \alpha_a \neq \alpha_b\}$ 
3:    $\tilde{t}_i \leftarrow (B_i^?, \emptyset)$   $\triangleright$  initial partial task assignment
4:    $\tilde{T} \leftarrow \{\tilde{t}_i\}$   $\triangleright$  partial task assignments (stack)
5:    $T \leftarrow \emptyset$   $\triangleright$  contr.-free complete task assignments
6:   while  $\tilde{T} \neq \emptyset$  do
7:     pop  $\tilde{t} = (B^?, B)$  from  $\tilde{T}$ 
8:     if  $B^? \neq \emptyset$  then  $\triangleright$  if any unassigned
9:        $\tilde{T} \leftarrow \text{ASSIGNNEXT}(\tilde{t}) \cup \tilde{T}$ 
10:    else  $\triangleright$  if all assigned
11:       $s_t \leftarrow \text{SCOREASSIGNMENT}(B)$   $\triangleright$  score
12:      push  $t = (B, s_t)$  to  $T$ 
13:   Sort( $T$ )  $\triangleright$  sort descending with score
14:   return  $T$ 
```

turns out to be complete on line 8, it will be scored (line 11, using Algorithm 4) and pushed to the set of complete task assignments T (line 12). The sorted set (line 13) of complete and contradiction-free task assignments will be returned (line 14) after $\tilde{T}$ no longer contains any task assignment with any action pair that has not been assigned to an Allen relation.

Algorithm 2 shows how several new partial or complete task assignments are created, given a partial task assignment, by choosing one action pair and trying to assign each Allen relation to it. Each feasible assignment (identified on line

7 using Algorithm 3) of an Allen relation to an action pair will result in a new task assignment returned by this function. Note that this function may return an empty set if no Allen relation can be assigned to the selected action pair without contradictions, causing Algorithm 1 to backtrack.

Algorithm 2 Create next contradiction-free partial or complete task assignments from the given current partial task assignment $\tilde{t}$ by picking one unassigned action pair and assigning each feasible Allen relation.

```

1: function ASSIGNNEXT( $\tilde{t} = (B^?, B)$ )
2:    $T' \leftarrow \emptyset$   $\triangleright$  next task assignments
3:   pop  $b^? = (\alpha_a, \alpha_b)$  from  $B^?$ 
4:   for all  $R \in R_A$  do  $\triangleright$  for each Allen relation
5:      $s \leftarrow \text{SCORE}(\alpha_a, \alpha_b, R)$   $\triangleright$  score for  $\alpha_a R \alpha_b$ 
6:      $b' \leftarrow (\alpha_a, \alpha_b, R, s)$ 
7:     if ISFEASIBLE( $B, b'$ ) then
8:        $B' \leftarrow B \cup \{b'\}$   $\triangleright$  next assigned
9:       push  $t' = (B', B')$  to  $T'$ 
10:  return  $T'$ 
```

Algorithm 3 outlines the procedure to ensure that the given partial assignment B would still be feasible and free of contradictions if the next assignment b' is added to B . The next assignment b' is a 4-tuple with $b' = (\alpha_a, \alpha_c, R, s)$, i.e., an assignment of the Allen relation R with a score s between action α_a and α_c . In this work, we use the degree of membership of an Allen relation between two actions given the demonstrations [11]. Let $R_1 \circ R_2 \subseteq R_A$ be the composition of Allen relations, i.e., the set of relations that may hold between α_a and α_c given $\alpha_a R_1 \alpha_b$ and $\alpha_b R_2 \alpha_c$, as defined by the transitivity table of Allen [1, Fig. 4]. Then, to evaluate whether b' can be added without causing contradictions in B , the algorithm verifies for each action $\alpha_b \in \mathcal{A}$: $\alpha_a R_1 \alpha_b \wedge \alpha_b R_2 \alpha_c$ that $R \in R_1 \circ R_2$ holds, i.e., that $\alpha_a R \alpha_c$ is consistent with all already established relations involving a shared action α_b .

Algorithm 3 Ensure that the given partial assignment B would still be consistent with the next assignment b' according to the transitivity table of Allen [1, Fig. 4].

```

1: function ISFEASIBLE( $B, b' = (\alpha_a, \alpha_c, R, s)$ )
2:    $I \leftarrow \{(\alpha_b, R_1, R_2) \mid \exists \alpha_b \in \mathcal{A} : (\alpha_a, \alpha_b, R_1, s_1) \in B, (\alpha_b, \alpha_c, R_2, s_2) \in B\}$ 
3:   for all  $(\alpha_b, R_1, R_2) \in I$  do
4:     if  $R \notin R_1 \circ R_2$  then  $\triangleright$  with  $\alpha_a R \alpha_c, \alpha_a R_1 \alpha_b, \alpha_b R_2 \alpha_c$ 
5:       return false
6:   return true
```

Algorithm 4 shows the scoring function for a complete task assignment. It evaluates the sum of all individual assignment scores s with $b^i = (\alpha_a, \alpha_c, R, s)$ whose corresponding assigned relation R is an overlapping Allen relation, i.e., any Allen relation except for *before* or *after*. We exclude them from scoring because the assignment is solved per subtask, where precedence constraints do not contribute and concurrent relations carry the important information on action cohesion.

Algorithm 4 Score a complete task assignment B .

```

1: function SCOREASSIGNMENT( $B$ )
2:    $S \leftarrow \{s \mid (\alpha_a, \alpha_c, R, s) \in B \wedge R \text{ is an overlapping Allen relation} \}$ 
3:   return  $\sum S$ 

```

The heuristic rule-based approach from our previous work [13] is used to decompose the task into subtasks. Our proposed DPLL-based algorithm combined with this heuristic enables us to identify task modes in the sense of variations in a subtask, and task modes in the sense of macroscopic variations of subtask sequences in a task. As in [13], we assume that subtasks are constrained by precedence constraints only, or not at all. We argue this assumption is weak, as it merely leverages the demonstrator’s bimanuality: a human cannot perform two subtasks requiring tight temporal coordination at once, and actions that do require such coordination form a single subtask by construction.

The proposed algorithm has a time complexity of $O(13^n)$, with 13 Allen relations that can be assigned to each action pair, n being the total number of action pairs with $n = \frac{|\mathcal{A}|(|\mathcal{A}|-1)}{2}$, and $|\mathcal{A}|$ being the number of actions in that subtask. Apart from the divide-and-conquer strategy of decomposing the task into subtasks, another measure based on [13] can be taken to make the proposed algorithm tractable in practice. The relations *meets* and *before* (and their inverses) can often already be pre-assigned due to unambiguously high scores for a specific relation, effectively reducing the dimensionality of the search space by one for each pre-assigned relation.

2) *Infer Subsymbolic Temporal Task Constraints*: With our approach proposed in Section IV-B.1, we obtain a temporal task assignment, i.e., one concrete Allen relation per action pair in the task. A timing model as proposed in Section IV-A.2 contains the collective timing knowledge of an action pair for all Allen relations, but we can retrieve an optimal timing for a given Allen relation by conditioning. To this end, for Allen relations that impose any time point equality (i.e., lines or areas in timing space), we sample the GMM’s probability density, orthogonally project the samples to the respective Allen relation’s region in $\mathbb{T}^3$ and select the sample of maximum likelihood. In the case of an Allen relation without time point equalities (i.e., volumes in timing space), we just directly sample from the GMM in that region. An example of this process is illustrated in Fig. 4.

C. Temporal Planning

After the inference process, symbolic and subsymbolic temporal task constraints are organized in a graph structure. However, a graph cannot directly be used for the task execution on a robot, hence a plan must be drawn up that follows

the symbolic temporal task constraints tightly (Section IV-C.1) and the subsymbolic temporal task constraints as well as possible (Section IV-C.2).

1) *Symbolic Planning*: In [13] we developed a bimanual temporal planner to generate a synthetic dataset of bimanual manipulation demonstrations of a task by varying the sequences of actions with predefined lengths until a set of symbolic temporal task constraints is met. We extended this planner to additionally vary the lengths of actions in multiples of unit-lengths until qualitative solutions are found that satisfy the symbolic temporal task constraints at hand.

2) *Temporal Plan Parametrization*: The symbolic plan will be parametrized to meet the inferred subsymbolic temporal task constraints as well as possible. For this, we formulate an optimization problem where the symbolic temporal plan serves as hard constraints and the subsymbolic temporal task constraints as soft constraints. Figuratively, the optimizer takes the symbolic plan as a blueprint and shifts the action’s starts and ends so that the timings in the plan are as close to the inferred subsymbolic temporal task constraints as possible. Specifically, let τ_i be the timing and R_{τ_i} the respective Allen relation of the i -th action pair inferred from our approach, and let x_i be the corresponding timing of the same actions in the plan to optimize, with $i \in \{1, \dots, n\}$, where n is the total number of action pairs considered. Additionally, let R_{x_i} be the Allen relation the timing x_i resembles. Then the objective is to minimize the Euclidean norm of the vector of all timing differences, s.t. $x \in \Omega$, with decision variable $x = (x_1, x_2, \dots, x_n)$ and constraint set $\Omega = \{x \mid \forall i \in \{1, \dots, n\}: R_{x_i} = R_{\tau_i}\}$ as follows:

$$\min_x \left\| \left(d(\tau_1, x_1), \dots, d(\tau_n, x_n) \right)^T \right\|_2 \text{ s.t. } x \in \Omega. \quad (5)$$

We use a standard convex optimization problem solver (CVXPY) to find an optimal solution. The result is a fully temporally parametrized plan that can be used to parametrize the concrete durations and offsets of MPs for robot execution.

V. EXPERIMENTS AND EVALUATION

In this section, we will describe our evaluation and experiments and report the results.

A. Task Assignment Assessment

Experimental Setup: To quantitatively evaluate our DPLL-based approach (“Fuzzy + DPLL-based”), we compare its inferred symbolic temporal task constraints to those of the rule-based heuristic from [13] in two variants: using relative frequencies of observed Allen relations as input (“Rel. Freq. + Heuristic”), and using the fuzzy-logic-based formulation of Allen relations from [11] (“Fuzzy + Heuristic”). Since our approach is a full search that maximizes individual scores while ensuring contradiction-free assignments, we assume its result given all demonstrations to be a valid ground truth. This evaluation thus shows how quickly each approach converges to this desired assignment. Using the “prepare muesli” task of the Bimacs dataset, the model receives one demonstration at a time, and after each the constraints inferred by all methods are compared to this ground truth in a binary classification.




Fig. 5: Full temporally parametrized plans of two complex tasks of the Bimacs dataset generated by our approach. Task 5, “prepare muesli” on the left, and Task 6, “disassemble component” on the right. Boxes depict identified subtasks, with ‘hold bowl’ being an exception, showing that actions can span subtasks.




Fig. 6: Mean F_1 scores $\pm$ one standard deviation for an increasing number of demonstrations across 100 learning scenarios of two baselines compared to our approach (DPLL-based).




Fig. 7: Mean distance $\pm$ one standard deviation of a temporally parametrized plan generated with our approach to all demonstrations. The baseline is the distance of the most characteristic demonstration to all other demonstrations.

We repeat this process in 100 learning scenarios, each varying the order in which the model receives demonstrations, and report the mean F_1 score and standard deviation.

Results: The results in Fig. 6 show that our DPLL-based approach consistently outperforms both methods, especially for a small number of demonstrations. This can be attributed mainly to the fact that the DPLL-based approach does not rely on rule-based biases and seeks the most likely assignment given the demonstrations. Already after ca. 37 demonstrations, almost all models of the DPLL-based approach outperform the heuristic baselines. The mean inference time including temporal relationship assessment $\pm$ standard deviation of symbolic temporal task constraints over all demonstrations and learning scenarios is 0.2 ± 0.1 s for Rel. Freq. + Heuristic, 127.0 ± 25.7 s for Fuzzy + Heuristic, and 293.8 ± 303.5 s for our Fuzzy + DPLL-based approach. Note that the inference times of our DPLL-based approach include the time needed to infer the fuzzy-logic-based scores of Allen relations, since these are required as input.

B. Action Timing Quality Assessment

Experimental Setup: Quantifying how good a generated temporally parametrized plan is can be challenging, as there

is no objective *best* timing. Thus, in this evaluation, we compare the generated temporally parametrized plan with all demonstrations. Both our temporally parametrized plan and a demonstration are essentially the same data structure (cf. Section III-A). Let $P_a = (a_{p_1}, a_{p_2}, \dots, a_{p_n})$ be the set of actions in our temporally parametrized plan and $D_a = (a_{d_1}, a_{d_2}, \dots, a_{d_n})$ the set of corresponding actions in the demonstration. Then we can define the distance between P_a and D_a as the Euclidean norm of the vector of semantic action differences as a generalization of Equation 4:

$$d(P_a, D_a) = \min_{\varsigma} \left\| \begin{pmatrix} - \\ a_{d_1} \\ a_{d_1}^+ \\ a_{d_2}^- \\ \vdots \\ a_{d_n}^+ \end{pmatrix} - \begin{pmatrix} - \\ a_{p_1} \\ a_{p_1}^+ \\ a_{p_2}^- \\ \vdots \\ a_{p_n}^+ \end{pmatrix} + \varsigma \begin{pmatrix} 1 \\ 1 \\ 1 \\ \vdots \\ 1 \end{pmatrix} \right\|_2. \quad (6)$$

The model receives one demonstration at a time, after which the whole approach presented in Fig. 1 is executed again to obtain an updated parametrized plan with the exception of the task assignment being chosen manually. This isolates the timing parametrization from the symbolic search in the sense of an ablation to ensure comparable results. We repeat this process in 50 learning scenarios and report the mean distance $d(P_a, D_a)$ and variance that the respective temporally parametrized plan P_a has to all demonstrations compared to a baseline of the most characteristic demonstration D_a . The most characteristic demonstration is defined as the demonstration with the smallest distance to all other demonstrations known to the model at that time. A direct comparison to the approach based on univariate GMMs [11] is not possible, since it cannot produce plans. We use the “prepare muesli” task of the Bimacs dataset. Since the demonstrations are complex and feature many different task modes (cf. [13, Table V]), we compare the quality of the parametrization for the “add banana” subtask.

Results: Fig. 7 shows that our approach reliably yields a smaller distance to each individual demonstration compared to the baseline, especially for a small number of demonstrations. This means that our approach can derive a parametrization for the plan closer to all demonstrations than the most characteristic demonstration. Additionally, our formulation also allows to further condition the GMM to specific temporal parameters or parameter ranges, like limiting the duration of an action to a given interval.

C. Full Task Plan Assessment

Experimental Setup: In this evaluation, we show the result after running our full approach as depicted in Fig. 1 using two

example tasks of the Bimacs dataset, namely the “prepare muesli” task and the “disassemble component” task. All demonstrations are provided at once.

Results: Fig. 5 shows the resulting temporally parametrized plans generated by our approach. For the “prepare muesli”, 12 task modes (subtask sequences) were identified in total, while “disassemble component” exhibited only one; within individual subtasks, between 25 and 20 117 modes of slight variations were found, most of which are unlikely. Gaps of 1 s have been added between subtasks to demonstrate that from the inferred task constraints, we know that there is no tight temporal relationship that must be met, leaving the robot time to approach objects for grasping or similar.

D. Showcase: Orchestrated Execution

In this showcase, we demonstrate how entire executions of tasks are orchestrated using temporally parametrized plans both in simulation and on real robots. For this, demonstrations of tasks from the BiManip dataset [14] and recordings of an angle grinder disassembly were used to derive (i) plans as shown in Fig. 1 from the segmentation data, and (ii) a library of full-pose Via-Point Movement Primitives (VMPs) [21] associated to each action from motion capture and kinesthetic recordings, respectively. We arrange and parameterize the VMPs as dictated by the plan, and play back the generated executions in simulation using trajectories from an IK solver considering human-likeness criteria [22] or on a real robot. The attached video (youtu.be/B9oH8GLKSYg) shows several successful executions of synchronized bimanual manipulation tasks using a library of existing unimanual MPs to give a qualitative impression.

VI. CONCLUSION AND FUTURE WORK

We presented a unified approach to simultaneously learn temporal task structure and action timing. We contribute a novel representation of timings between two actions in a 3-D space, provide new methods to find assignments of Allen relations to each pair of actions of a task, and propose a novel planning system that can generate temporally parametrized plans for execution considering symbolic and subsymbolic temporal task constraints. A limitation is that it relies on the heuristic task decomposition method from our previous work [13] to identify subtasks prior to symbolic assignment. Since the runtime of our DPLL-based search grows exponentially with the number of action pairs considered, an imperfect decomposition that yields overly large subtasks can cause runtime spikes. In future work, we will investigate task decomposition methods to address this limitation. In addition, we believe that a combination of emerging and assigned synchronization approaches is necessary for a dynamic and goal-oriented orchestration of bimanual actions. Ultimately, more research on general task models that incorporate many task constraint modalities is needed.

REFERENCES

- [1] J. F. Allen, “Maintaining Knowledge About Temporal Intervals,” *Communications of the ACM*, vol. 26, no. 11, pp. 832–843, 1983.
- [2] S. Ekvall and D. Kragić, “Learning Task Models from Multiple Human Demonstrations,” in *International Symposium on Robot and Human Interactive Communication (RO-MAN)*, 2006, pp. 358–363.
- [3] M. Pardowitz, S. Knoop, R. Dillmann, and R. D. Zöllner, “Incremental Learning of Tasks From User Demonstrations, Past Experiences, and Vocal Comments,” *Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics)*, vol. 37, no. 2, pp. 322–332, 2007.
- [4] M. N. Nicolescu and M. J. Matarić, “Natural Methods for Robot Task Learning: Instructive Demonstrations, Generalization and Practice,” in *International Joint Conference on Autonomous Agents and Multiagent Systems*, 2003, pp. 241–248.
- [5] P. K. Thota, H. Ravichandar, and A. P. Dani, “Learning and Synchronization of Movement Primitives for Bimanual Manipulation Tasks,” in *Conference on Decision and Control (CDC)*, 2016, pp. 945–950.
- [6] T. Kulvicius, M. Biehl, M. J. Aein, M. Tamosiunaite, and F. Wörgötter, “Interaction Learning for Dynamic Movement Primitives Used in Cooperative Robotic Tasks,” *Robotics and Autonomous Systems (RAS)*, vol. 61, no. 12, pp. 1450–1459, 2013.
- [7] A. Gams, B. Nemec, A. J. Ijspeert, and A. Ude, “Coupling Movement Primitives: Interaction With the Environment and Bimanual Tasks,” *Transactions on Robotics (T-RO)*, vol. 30, no. 4, pp. 816–830, 2014.
- [8] A. Gams, A. Ude, and J. Morimoto, “Accelerating Synchronization of Movement Primitives: Dual-Arm Discrete-Periodic Motion of a Humanoid Robot,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 2754–2760.
- [9] M. Karlsson, A. Robertsson, and R. Johansson, “Convergence of Dynamical Movement Primitives with Temporal Coupling,” in *European Control Conference (ECC)*, 2018, pp. 32–39.
- [10] F. Krebs and T. Asfour, “Formalization of Temporal and Spatial Constraints of Bimanual Manipulation Categories,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 1302–1309.
- [11] C. Dreher and T. Asfour, “Learning Symbolic and Subsymbolic Temporal Task Constraints from Bimanual Human Demonstrations,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 5160–5167.
- [12] M. Davis, G. Logemann, and D. Loveland, “A Machine Program for Theorem-Proving,” *Communications of the ACM*, vol. 5, no. 7, pp. 394–397, 1962.
- [13] C. Dreher and T. Asfour, “Learning Temporal Task Models from Human Bimanual Demonstrations,” in *International Conference on Intelligent Robots and Systems (IROS)*, 2022, pp. 7664–7671.
- [14] F. Krebs, A. Meixner, I. Patzer, and T. Asfour, “The KIT Bimanual Manipulation Dataset,” in *International Conference on Humanoid Robots (Humanoids)*, 2021, pp. 499–506.
- [15] A. Billard, S. Calinon, R. Dillmann, and S. Schaal, “Robot Programming by Demonstration,” in *Handbook of Robotics*, 1st ed., B. Siciliano and O. Khatib, Eds., 2008, pp. 1371–1394.
- [16] E. Carpio, M. Clark-Turner, and M. Begum, “Learning Sequential Human-Robot Interaction Tasks from Demonstrations: The Role of Temporal Reasoning,” in *International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2019, pp. 1–8.
- [17] X. Ye, Z. Lin, and Y. Yang, “Robot Learning of Manipulation Activities with Overall Planning Through Precedence Graph,” *Robotics and Autonomous Systems (RAS)*, vol. 116, pp. 126–135, 2019.
- [18] T. Birr, C. Pohl, A. Younes, and T. Asfour, “AutoGPT+P: Affordance-based Task Planning with Large Language Models,” in *Robotics: Science and Systems (RSS)*, 2024.
- [19] T. Motoda, R. Hanai, R. Nakajo, M. Murooka, F. Erich, and Y. Domae, “Learning Bimanual Manipulation Via Action Chunking and Inter-Arm Coordination with Transformers,” in *International Conference on Automation Science and Engineering (CASE)*, 2025, pp. 1796–1801.
- [20] R. D. Zöllner, T. Asfour, and R. Dillmann, “Programming by Demonstration: Dual-Arm Manipulation Tasks for Humanoid Robots,” in *International Conference on Intelligent Robots and Systems (IROS)*, vol. 1, 2004, pp. 479–484.
- [21] T. Daab, N. Jaquier, C. Dreher, A. Meixner, F. Krebs, and T. Asfour, “Incremental Learning of Full-Pose Via-Point Movement Primitives on Riemannian Manifolds,” in *International Conference on Robotics and Automation (ICRA)*, 2024, pp. 2317–2323.
- [22] A. Meixner, M. Carl, F. Krebs, N. Jaquier, and T. Asfour, “Towards Unifying Human Likeness: Evaluating Metrics for Human-Like Motion Retargeting on Bimanual Manipulation Tasks,” in *International Conference on Robotics and Automation (ICRA)*, 2024, pp. 13 015–13 022.