International Journal of Social Robotics (2021) 13:981–998
https://doi.org/10.1007/s12369-020-00685-2

# High-Level Motor Planning Assessment During Performance of Complex Action Sequences in Humans and a Humanoid Robot

Theresa C. Hauge[1] · Garrett E. Katz[2] · Gregory P. Davis[3] · Di-Wei Huang[3] · James A. Reggia[3,4,5,6] · Rodolphe J. Gentili[1,5,6]

Accepted: 23 July 2020 / Published online: 17 August 2020


**Abstract**

Examining complex cognitive-motor performance in humanoid robots and humans can inform their interactions in a social context of team dynamics. Namely, the understanding of human cognitive-motor control and learning mechanisms can inform human motor behavior and also the development of intelligent controllers for robots when interacting with people. While prior humans and humanoid robot studies mainly examined motion planning, only a few have investigated high-level motor planning underlying action sequences for complex task execution. This sparse work has largely considered well-constrained problems using fairly simple performance assessment methods without detailed action sequence analyses. Here we qualitatively and quantitatively assess action sequences generated by humans and a humanoid robot during execution of two tasks providing various challenge levels and learning paradigms while offering flexible success criteria. The Levenshtein distance and its operators are adapted to the motor domain to provide a detailed performance assessment of action sequences by comparing them to a reference sequence (perfect sequence having a minimal number of actions). The results reveal that (i) humans produced a large variety of action sequences combining perfect and imperfect sequences while still reaching the task goal, whereas the robot generated perfect/near-perfect successful action sequences; (ii) the Levenshtein distance and the number of insertions provide reliable performance markers capable of differentiating perfect and imperfect sequences; (iii) the deletion operator is the most sensitive marker of action sequence failure. This work complements prior efforts for complex task performance assessment in humans and humanoid robots and has the potential to inform human–machine interactions.

**Keywords** Cognitive-motor control and learning · High-level motor planning · Humanoid robot · Human · Imitation · Action sequence

## 1 Introduction

Contrary to low-level sensorimotor planning, a limited effort has examined high-level cognitive-motor planning functions, which involve multi-step motor action sequences for achieving complex tasks. New measures of performance on complex cognitive-motor tasks executed by humans or robots can provide new means for robots to detect and adapt to a

Theresa C. Hauge and Garrett E. Katz have contributed equally to this work and are co-first authors.

**Electronic supplementary material** The online version of this article (https://doi.org/10.1007/s12369-020-00685-2) contains supplementary material, which is available to authorized users.

✉ Rodolphe J. Gentili
rodolphe@umd.edu

1 Cognitive Motor Neuroscience Laboratory, Department of Kinesiology, School of Public Health, University of Maryland, 4200 Valley Drive, College Park, MD 20742, USA

2 Department of Electrical Engineering and Computer Science, Syracuse University, New York 13244, USA

3 Department of Computer Science,, University of Maryland, College Park, MD 20742, USA

4 Institute for Advanced Computer Studies, University of Maryland, College Park, MD 20742, USA

5 Neuroscience and Cognitive Science Program, University of Maryland, College Park, MD, USA

6 Maryland Robotics Center, University of Maryland, College Park, MD, USA

human teammate's skill level promoting thus cooperation and trust. In particular, distance-based measures facilitate this goal by quantifying how similarly a human and a robot execute a task [1]. Such distance-based assessment can also be used by communities of robots to relate and share knowledge about their respective human collaborators to become socially knowledgeable partners [2, 3].

Most prior study of human high-level motor planning has mainly considered constrained tasks where individuals plan action sequences following specific rules and without varying conditions of challenge and/or practice paradigms (e.g., Towers of Hanoi). Also, behavioral analyses were largely based on basic metrics (e.g. response time, total number of moves, [4–6]). This prior work provides limited information about the action sequences since there is no indication of which specific actions were added, removed or switched, at which specific positions, in the sequences—information that is critical to assess high-level planning performance [4, 6]. This limited examination inadequately informs human neurocognitive processes, the development and assessment of smart robotic controllers, and human–robot interactions during complex task performance (e.g., equipment maintenance and cleaning [7, 8]).

Similarly, many robotic studies have assessed performance with fairly simple subjective (e.g., surveys; [9–15]), objective (e.g., time; [10, 13, 16, 17]), and coarse-grained measures (e.g., total step count; [13, 18–21]). These metrics cannot inform a fine-grained performance comparison between humans and humanoid robots. To our knowledge only one prior study has conducted a more refined analysis of high-level motor plans using alignment algorithms to compare action sequences executed in humanoid robots [1]. However, this effort did not examine the human or humanoid performance in a context where both agents learn to execute action sequences. Such a detailed analysis is critical since it can inform human and robot training needs and predict their interactions in situations where flexible solutions can be identified for learning to complete a complex task under various demands.

Thus, there is a need to examine the performance in both humans and humanoid robots, during complex action sequences. Here we focus on action sequences that can be performed with few constraints and flexible success criteria (the sequence can be executed in multiple ways while still reaching its goal) under various conditions of challenge and learning paradigms. Our primary interest is to improve characterization and understanding of high-level motor planning in both humans and humanoid robots when they must learn a complex task. Indeed, in real-world settings many tasks can be executed with action sequences which can fail to reach the task goal but also be completed in many ways while still reaching the same goal. There may also exist one (or more) way(s) to efficiently execute the task with a minimum number of actions. As such, learning could result in performance with various levels of success and efficiency depending on whether the sequence generated has a minimum number of actions or not. Thus, the performer could produce (i) a successful action sequence reaching the prescribed task goal with a minimal number of actions, (ii) an action sequence with additional steps (e.g., adding extraneous actions) but still reaching successfully the goal, or (iii) a failed attempt to reach the goal.

Learning action sequences through trial-and-error or imitation is highly relevant to fundamental human cognitive-motor mechanisms and allows humanoid robots to learn via a social exchange, avoiding burdensome manual programming [3, 7]. Here we take some first steps in examining in detail complex action sequence performance in both a humanoid robot and humans who practice a complex task where humans provide task goals and/or features to be replicated.

This work primarily aims to deploy a computational method that measures action sequence similarity, to examine the high-level structure of complex action sequences for assessing the performance under various conditions of challenge and learning in humans and a humanoid robot.[1] Specifically, we used action sequence similarity to determine the functional reasons for: (i) performance failure of complex action sequences (which action operations drove this failure) and (ii) performance differences between successful sequences with and without a minimum number of actions. Interestingly, other scientific fields, such as DNA sequence analysis in biology or phoneme sequence analysis in linguistics, encounter similar issues. These fields proposed solutions for sequence comparisons based on edit distance (e.g., Levenshtein distance (LD)) which computes the number of additions, deletions and substitutions of the bases/phonemes needed to transform one sequence into another that serves as a reference (a larger number of alterations indicates that both sequences further differ; [22]). The LD has been previously employed for comparing different gaze signals during human–robot interactions without however focusing on high level motor plan generation for upper extremity performance [23]. Among multiple potential applications for humans and/or robots, such a LD-based methodology could quantify how efficiently (here expressed as a minimal number of actions generated) humans and robots solve the same problem as well as how their solutions differ. This can inform their respective training and inform their human–robot team mental models. Indeed, if a human and a robot solve a task using a very different strategy, it is likely difficult for both to

[1] While the direct comparison between humans and the robot is conducted here, this is secondary and rather of an exploratory nature in this work.

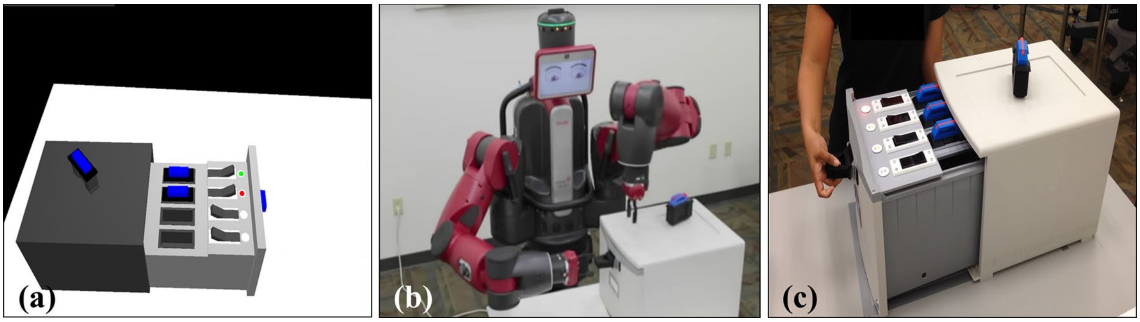


**Fig. 1** Top row: The disk drive dock task. **a** A screen shot from the interactive virtual world supported by SMILE where the objects on a table top (here a simulation of the disk-drive dock) can be manipulated by a demonstrator and those manipulations can be subsequently viewed by a human or a robot imitator. **b** The bimanual humanoid cognitive robot used in this work. **c** The disk drive dock (physical mock-up) used here was manipulated by a human. The drawer is partially opened showing three of the four slots occupied by drives, each with a corresponding black fasten-release toggle switch and a corresponding red-green indicator LED (red: malfunctioning; green: functioning correctly). A spare disk sits on the top of the cabinet. Pressing a toggle switch can change the state of the corresponding LED (red, green, off) controlled by a built-in Arduino microprocessor

predict or infer efficiently the future states and goal understanding of each other's behavior, leading ultimately to poor team dynamics [24–26].

Thus, taking inspiration from biology and linguistic fields, here we deployed a LD-based computational method to examine how adding, omitting or replacing actions in an action sequence can affect the learning task performance outcome (i.e., failure vs. success). We consider tasks where there exist (i) only one unique sequence that reaches the goal with a minimum number of actions and serves as a reference and (ii) an infinite number of action sequences that can be generated to successfully reach the goal while differing from the minimal sequence. Such a quantitative analysis is qualitatively complemented with graphical representations (sequence alignment display) to further illustrate any alterations of produced action sequences by humans and the humanoid robot compared to the reference sequence. Namely, it may be possible that the addition of extraneous actions even between the critical elements of the reference sequence is not a major determinant of failure, despite increased action sequence length, since the order of those critical components is preserved (the task may still be correctly performed and its goal reached). However, omitting or replacing an action of the reference sequence, unless compensated, may represent a greater threat to successful task completion. If these assumptions are correct, we predict that the occurrence of action omission and/or replacement should be more prominent during failure, and should be minimal or null when successfully executing (with a minimum of steps or not) complex action sequences. The results should be consistently observed both via qualitative (e.g., sequence alignment display) and quantitative (e.g., frequency of action omissions or additions) analyses. If the computational method proposed here is somewhat generalizable, it should provide consistent qualitative and quantitative results in different tasks involving various conditions of challenge and learning paradigms, whether the performer is a humanoid robot or human.

## 2 Materials and Methods

To examine whether our computational method was general enough to successfully assess action sequence performance in various situations facilitating or not mistakes and/or strategy diversifications, two tasks involving various conditions of challenge and learning were considered. The first task involved a fairly short motor sequence composed of eight actions to maintain a hard drive docking station (HDDS) along with limited rules and flexible success criteria and on which the robot and humans were trained in a similar way using imitation learning. Compared to the HDDS task, the second task aimed to substantially increase the conditions of challenge and allowed for different learning paradigms to solve a modified Tower of Hanoi (TOH) task. This task was manipulated to ensure that both humans and the robot were challenged as much as possible while considering their respective learning capabilities to still provide them a chance to perform the task. Both tasks are explained in detail below.

### 2.1 Hard Drive Docking Station Task

#### 2.1.1 Experimental Set-Up

Twelve human participants (see section S2.1 of the supplementary material for details) and a humanoid robot (Baxter; Rethink Robotics™) had to perform a complex task that involved a mock-up hard drive docking station. This station (Fig. 1a) was composed of a drawer that, when opened, allowed to manipulate disks in four hard drive slots, each being associated with a LED indicator and a toggle switch, which had to be switched off before inserting or removing the corresponding drive. A red LED associated with each drive indicated a faulty disk that had to be replaced with a new one. If a red LED was initially on and its corresponding switch was pressed, then this LED turned off, indicating that this drive was now disengaged and thus could be replaced. When a disk is successfully replaced and its switch is pressed, then its green LED turns on indicating that the corresponding disk is now engaged and functioning properly. The custom-made mock-up was controlled by an Arduino processor.

#### 2.1.2 Task Description

The HDDS task had flexible successful completion criteria, and limited constraints allowing for performing a real-world task in a very versatile manner. It was used in prior cognitive robotic research to examine high-level plan generation in humanoid robots [7, 8], and is well-suited for examining action sequences deployed by the various performers (humans and a humanoid robot) during action sequence imitation. It uses a given demonstration as an action sequence of *reference* having a minimal number of actions.

A humanoid robot (Baxter; Rethink Robotics™) and the human participants had to learn[2] to perform a complex task which consisted of operations on the HDDS task. Specifically, the participants had to learn to discard a faulty hard disk which was indicated by a red LED being on and replace it with a new one. First the action sequence needed to perform the task was demonstrated to the performer (humans or robot) via computer software used to record demonstrations in an artificial environment (SMILE; [28, 29]; see section S2.1 for details). After observing the demonstration, the performers (humans or robot) have to imitate the task immediately while employing the physical mock-up hard drive station placed in front of them (Fig. 1b, c).

Specifically, via the SMILE virtual environment operated by a human-controlled interface [28, 29], the video provides a demonstration of this complex action sequence composed of eight atomic actions executed in the following order: (1) the drawer was pulled to open the dock to make available the four hard drive slots, their corresponding toggle switches and LEDs (abbreviated dro, for drawer opened); (2) the faulty drive (here the first drive), which could be identified by its red LED, had to be turned off by pressing its toggle switch (i.e., t1, first toggle); (3) the hard drive had to be picked up to be removed from its slot (ud1), and (4) discarded in a bin (db); (5) a new spare drive was picked up (usp) and then (6) inserted in the empty first slot (ds1); (7) the same toggle switch was pressed again (t1) to turn on the green LED and (8) the drawer was closed (drc) (see supplementary material Table S1 for the atomic actions of this sequence). For human-readability we have chosen strings like 'dro' or 't1' that contain more than one letter, but each atomic action is thought of as a *single symbol*, and thus the smallest unit of activity in this task. This particular demonstrated sequence is called the *demonstrated sequence* or *reference sequence*.

While the human participants in this study practiced this task by watching a SMILE video, the humanoid robot did not process raw video footage. Specifically, while for humans multiple brain regions (e.g., dorsal and ventral stream) would be engaged to visually process and encode the actions composing the sequences shown in the video, such a process was not modeled in the neurocognitive architecture controlling the robot since the focus was on implementing high-level planning mechanisms. However, for the robot, the SMILE event record, which is auto-generated, provided a text-based transcription of the video to its neurocognitive architecture. This event record amounts to a transcribed demonstrated low-level action sequence, and the robot draws on its prior knowledge to infer the demonstrator's goals based on their actions. It then imitates by planning its own (potentially different) actions to achieve the same goals.

These robotic reasoning processes use rich knowledge representations including first-order logic with continuous variables, preconditions and temporal ordering constraints, and hierarchical action and task relationships. However, once a plan is generated, each constituent action is grounded (its variables are all bound to constants) and it can be treated as an atomic element of an alphabet for the purposes of LD computation (for more details for this inference and planning process see [8]). After learning, the robot produces a motor sequence that successfully imitates the complex action sequence and that is stored for further analysis. Thus, although compared to human participants, the neurocognitive architecture has enhanced pure memorization capabilities, it does not simply memorize the demonstrated action sequences but instead learns to understand their goals using cause-effect reasoning and decomposes them into actions that are then executed.

The human participants practiced the same imitation task. First, before watching any demonstration, an acclimation stage allowed the participants to become familiar with the physical disk drive dock station (see section S2.1 for details). A prompt provided the instructions including the meaning of the red and green LEDs and, importantly, also indicating that the LED should be turned off when adding or removing the hard drives. Once this familiarization phase was complete, the video was shown, demonstrating the action sequence

[2] For consistency between humans and the humanoid robot, here the term *learning* is employed in a general manner and reflects performance during the practice throughout the trials [27].

to the human participants who had to learn to imitate it. Five trials were performed by each human participant, each having the following structure and using the same video in order to avoid any possible confounding factors due to a potential transfer of learning from one task to another. At the start of each trial, the instructions were provided on a computer monitor via the prompt (and also verbally). Once the demonstration was over, the participants had to imitate the action sequence previously observed.

As the participants performed the task, a video recording using a digital camera was employed to collect the action sequences they performed, and these actions were subsequently manually transcribed into a sequence of single symbols as previously indicated (see supplementary material Table S2 for details). At the end of each trial, the hard drive dock station was reset to its initial state and the demonstration in the video ready to be played again in preparation for the next trial.

Success in this task required both: (i) the completion of the goal, which was to replace the faulty hard drive with a spare one placed on the top of the docking station, as well as (ii) following the basic principle that the LED should be turned off when hard drives are added or removed. If these two criteria were not met, a motor sequence was classified as a failure. It is critical to note that the combination of providing only two success criteria combined with a minimum number of instructions was chosen to minimize the constraints on action performance. Thus, it was possible for the participants to successfully perform the demonstrated action sequence without having to exactly imitate the demonstrated motor sequence (which had the minimum number of eight atomic actions executed in the specific order <dro, t1, ud1, db, usp, ds1, t1, drc> (see Table S2 for the complete set of abbreviations used) but also by performing the task in various manners while still reaching the goal (i.e., replacing the drive next to an initially red LED), and respecting the basic rule that LEDs should be switched off while inserting or removing drives. Although the reference sequence was fairly short, it still promoted real-world flexibility in its execution allowing to capture various high-level plans (i.e., motor sequences) deployed by the humans and the humanoid robot to complete the proposed cognitive-motor task.

## 2.2 Tower of Hanoi Task

### 2.2.1 Experimental Set-Up

Twenty participants faced the Tower of Hanoi (TOH) setup which included four or five disks stacked atop each other of gradually increasing diameters from top to bottom (denoted from 1 to 4 or from 1 to 5; 1 being the smallest disk atop) and three pegs (denoted A to C from the left to the right of the performer) (see section S2.2 for details on participants).

### 2.2.2 Task Description

The TOH task consisted of moving the entire tower from the leftmost (first peg, denoted A) to the rightmost (third peg denoted C) peg by performing an action sequence with a minimum number of moves. As previously mentioned the conditions of execution of the TOH were chosen to challenge as much as possible both the robot and humans while accounting for their own capability so that they could still be able to perform the task. Typically, to perform the TOH task, the three following rules are provided: (i) only one disk can be moved at a time; (ii) a disk may not be placed on the table or held in the hand while another disk is manipulated in space and (iii) a larger disk may not be stacked on top of a smaller disk. However to elevate the challenge relative to the HDDS task, the humanoid robot had to complete the TOH task with five disks and learning it through imitation without knowing the three rules. As such its success depended only on its capability to successfully imitate the sequence. In addition, it was required to perform the task with five disks resulting in a *reference sequence* having thirty-one atomic actions which had to be executed in a specific order to reach the task goal. Thus the reference sequence to be imitated was almost four times longer than that to complete the HDDS task, increasing substantially the computational cost for the cognitive control system of the robot while still promoting strategy diversification ($2^i-1$ where $i$ is the number of disks; here $i=5$; [30]) (see Table S1 for further details). No information about the recursive solution was provided. Under these conditions, the robot had to learn the TOH task via the SMILE virtual environment. After learning, the robot was able to successfully replicate the demonstrated complex action sequence and was stored for further analysis.

The human participants had also to perform the TOH task but with four disks since preliminary results revealed that, although being challenging, this amount was appropriate (five disks being too difficult). As a result the *reference sequence* included fifteen atomic actions which had to be executed in the particular order <1B, 2C, 1C, 3B, 1A, 2B, 1B, 4C, 1C, 2A, 1A, 3C, 1B, 2C, 1C> to reach the task goal ($2^i-1$; here $i=4$; [30]; see Table S1). As such, compared to the HDDS task, the action sequence with this TOH task was almost twice as long offering thus a greater challenge while promoting strategy diversification (see Table S1). To further increase the task difficulty relative to the HDDS task, while the three rules were explained to the human participants, they were trained via a trial-and-error learning paradigm where only the task goal was provided (i.e., stacking up the disks in increasing size from top to bottom) but without demonstrating the reference sequence. As for the robot no information regarding the recursive solution was provided.

Before practice human participants became accustomed to the set-up during a familiarization period (see section

S2.2 for details). As previously done for the HDDS task the same video data collection and analysis was conducted. After each trial, the participants were informed if they had violated one or more execution rules, then the TOH setup was reset to its initial state and the prompt was shown again in preparation for the next trial. Success in this task required the completion of the task goal without violating the three execution rules. If these criteria were not met, the motor sequence was considered as a failure.

### 2.3 The Levenshtein Distance Applied to Complex Action Sequences

The LD can be applied to the domain of complex motor action sequences by appropriately defining the corresponding alphabet of symbols, sequences, and operations. Thus each atomic "symbol" is here defined as an elementary action that is part of an action sequence. In general, the performance of a complex action sequence can be associated with an alphabet that can be defined as a finite set $\{A_1, A_2, \ldots A_i, \ldots, A_{n-1}, A_n\}$ where $A_j$ is the *jth* atomic action among all *N* possible actions. A *motor sequence* would be defined as a finite, ordered list of zero or more actions from the alphabet, potentially with repeats (see section S2.3 for a simple example of motor action alphabet).

In a motor context, the computation of the LD can be associated with three possible operators: (i) *insertion* of one action, (ii) *deletion* of one action, and (iii) *substitution* of one action for another one (i.e., this can be seen as a replacement). These three operations can be employed to transform any motor sequence into another one, by modifying only one action at a time. The *insertion* operator inserts a new action at any position in the action sequence, increasing the length of the motor sequence by one. The *deletion* operator removes an action at any position in the action sequence, decreasing the length of the sequence by one. The *substitution* operator replaces an existing action with a new action at the same position in the sequence, leaving thus the sequence length unchanged (see Fig. S1 of the supplementary material for an example of these operators).[3] In a motor context, the LD measures the overall distance between two different motor sequences, which is defined as the minimum number of atomic operations required to transform one sequence into the other [31, 32]. The computation of the LD can be efficiently performed by well-established dynamic programming methods such as the Wagner-Fischer algorithm [33]. The same computational analysis was employed for both the HDDS and the TOH tasks. Insertions, deletions, and substitutions are not part of the robot's planning process, but rather, conceptual edits that could be made to the "result" of the robot's planning process (i.e., the resulting planned action sequence), which would transform it into a demonstrated reference sequence. The robot's planning process involves prior knowledge about the preconditions and postconditions of each action, encoded in a hierarchical task network planning formalism. This prior knowledge includes symbolic conditions (e.g., a gripper must be empty to pick something up) and sub-symbolic computations (e.g., a motion planner must compute a successful reach motion to an object in order for it to be picked up) (for details of this planning process see [8]).

### 2.4 Action Sequence Performance Analysis

To capture the various ways of successfully completing the HDDS and the TOH tasks, the set of all potential actions that can be executed (i.e., the alphabet) had to be defined (see Table S2). For the HDDS task as well as the TOH task with four and five disks, an alphabet consisting of 17, 12 and 15 possible actions was identified (see section S2.4 and Table S2 of the supplementary material for details). After data collection, for the HDDS and TOH tasks, a total of 62 (60 from humans, 1 from the robot and 1 reference sequence from SMILE) and 77 (75 from humans, 1 from the robot and 1 reference sequence from SMILE) action sequences were examined, respectively. In each task, sequences included both successes and failures, based on the criteria defined above.

#### 2.4.1 Graphical Representation of Complex Action Sequences

Action sequence alignments were generated to provide a combined visualization of the various action sequences generated by human participants and the humanoid robot with respect to the reference sequences. The action sequence alignments were generated for both the failed and successful (perfect and imperfect) action sequences. This representation allows one to identify where divergences between the motor sequences occur, by aligning all sequences side by side. The spacing between successive actions in each sequence is systematically varied so that the common sub-sequences across all sequences are aligned as well as possible. The alignment allows both a visualization of the generated action sequences at a global level as well as a focus on details in a particular sequence.

#### 2.4.2 Levenshtein Distance and Operator Occurrence Analyses

In order to assess the performance for the humans and the humanoid robot for both tasks, each action sequence

[3] For consistency with prior work the standard LD considering only these three operators was employed.

performed by the humans or the robot was compared to the reference sequence by computing both the LD as well as the number of occurrences of each of the three operators (i.e., insertions (I), deletions (D) and substitutions (S); where LD = I + D + S). Then, three types of action sequence performance were identified and considered for statistical analyses: (i) failed action sequences where either the goal was not reached or it was reached by violating at least one of the execution rules, (ii) successful but imperfect action sequences where the goal was successfully reached by employing a different sequence than the reference sequence, and (iii) successful perfect action sequences that exactly matched the reference sequence (i.e., those have a LD = 0).

A first analysis examined the extent to which insertions, deletions and substitutions are markers of performance success or failure in the humanoid robot and humans. The average (computed across all trials and participants) LD and the occurrence of its three operators (insertion, deletion and substitution) for the failed action sequences were compared to (i) all the successful sequences; (ii) the successful imperfect sequences only and (iii) the successful perfect sequences only using a reference value of zero (representing a perfect performance having a zero LD with I = D = S = 0) via a series of *Wilcoxon signed rank* tests. Then the successful perfect and imperfect action sequences were further examined. To do this, the average LD and the occurrence of its three operators for the imperfect action sequences were compared to the value zero (LD = I = D = S = 0) via a series of one sample *Wilcoxon signed rank* tests. Finally, to assess the contribution of the three operators (insertion, deletion, substitution) to the LD for all action sequences (failed, imperfect/perfect successful) the LD and its operator occurrences were compared via a series of *Wilcoxon signed rank* tests. Also, a correlational analysis allowed to assess the relationships between each of the four metrics and the success or failure of action sequence.

A second analysis investigated the change in performance during the five practice trial. The average values of each of the four metrics plus the number of failed sequences were computed across participants for each single trial and then examined by employing a series of one way Friedman tests (with 5 repetitions for the factor trial).

A third analysis applied to the successful and failed action sequences assessed the ability of the performers to recall more items correctly at the beginning (i.e., primacy) and/or at the end (i.e., recency) of the action sequences [34, 35]. The occurrence of the LD operators at each of the positions of the atomic actions (eight and fifteen for the HDDS and the TOH, respectively) in the reference sequence to transform it into any given human sequence were computed. Then, to compare if some operators were employed more frequently than others at the beginning, middle or end[4] of the failed and imperfect successful action sequences, the number of insertions, deletions and substitutions were compared for these three periods via *Wilcoxon signed rank* tests.

Finally, an additional exploratory analysis directly compared the humans and robot action sequences to assess how the human motor plans differed from that employed by the humanoid robot and under which conditions (i.e., failed, successful imperfect/perfect outcome). While such comparison is appropriate for the HDDS task since humans and the robot were trained in the same way, for the TOH task the same comparison was more delicate since the humans and robots were trained differently and thus the results of this specific analysis has to be taken with caution. The average (computed across trials and participants) LD and the occurrence of its four operators obtained between the action sequence generated by the humanoid robot and those successfully produced by the humans were compared via a series of Wilcoxon signed rank tests due to the low power to detect any departure from a normal distribution with a limited sample size.

For all the statistical analyses mentioned above, the false discovery rate (FDR; [36]) was employed to control for the multiple statistical tests conducted on all the metrics that indexed the humans and humanoid robot performance. The same statistical analysis was employed for both the HDDS and the TOH task.

## 2.5 Robotic Imitation Learning System

Our robotic imitation learning system, CERIL, is an AI system that performs non-trivial causal inference. It has been described in detail elsewhere [7, 8, 28, 29], so we briefly summarize it here. It uses built-in cause-effect knowledge to infer the high-level goals of a demonstrator and imitate the goals rather than the actions. As opposed to rote mimicry of a demonstrator, imitating the goals allows the robot to generalize to new situations (e.g., different number of drives in different slots with different LED colors, and different number and positions for spare drives). A consequence of the robot's generalization ability is that, when presented with a new situation different from the demonstration, the robot may potentially choose a sequence of actions quite different from the sequence that was demonstrated. These differences can include manipulating different objects in different orders, and can also include changes that reflect the differences in the robot's and the demonstrator's embodiment and

[4] The beginning, middle and end of the sequence were defined as 25% of the total number of actions forming the reference sequence (e.g., the beginning included the actions #1–2 and #1–4 for the HDDS and TOH task, respectively).

constrained ranges of motion. For example, the robot may hand off a drive between grippers when only one gripper can reach a target location, while this hand off might have been unnecessary for a person. For similar reasons, the action sequence employed by a robotic imitator might be quite different than that of human learners imitating the same demonstration (see section S2.5 of the supplementary material for details on the robot's imitation processes).

## 3 Results

Analysis of the action sequences revealed that for both tasks, trials performed included a large variety of sequences among which the successful sequences do not include any deletion and no or few substitution operators contrary to the failed sequences (Fig. 2; compare the sequence 10, 12 vs. 13, 17 for the HDDS task as well as the sequences 4, 13 vs. 39, 40 for the TOH task; see Figs. S2, S3 and Tables S3, S4 and section S3 of the supplementary material for further information). Based on the criteria previously established, 65% and 96% of the trials were successful (either completely matched the demonstration reference sequence or differed but were still successful) whereas 35% and 4% of them were not successful in completing the HDDS and the TOH task, respectively (Fig. 3, see pie chart and Table 1 for details). For both tasks the humanoid robot successfully performed the action sequence in one single trial. While for the TOH task the humanoid robot perfectly replicates the sequence (i.e., LD = I = D = S = 0), for the HDDS task the robot did not strictly follow the demonstration with a LD and a number of insertions equal to one and no deletions or substitutions (i.e., LD = I = 1; D = S = 0; see Fig. 2a action sequence #12). For both tasks the number of insertions was significantly greater than the number of deletions and substitutions for the successful imperfect sequences indicating that increased LD was mainly driven by insertions. For the failed action sequences the contributions to the LD to both tasks tended to be more distributed among the three operators while the number of insertions was significantly larger than the number of substitutions for both the HDDS and TOH tasks (see section S3 of the supplementary material for further details).

### 3.1 Assessment of Failed and Successful action Sequences in Humans and the Humanoid Robot

#### 3.1.1 Hard Drive Docking Station

Statistical analysis revealed that the LD and the number of insertions for the failed action sequences generated by human individuals differ from those obtained with the perfect successful sequence ($z \geq 2.668$, $p < 0.023$ for both

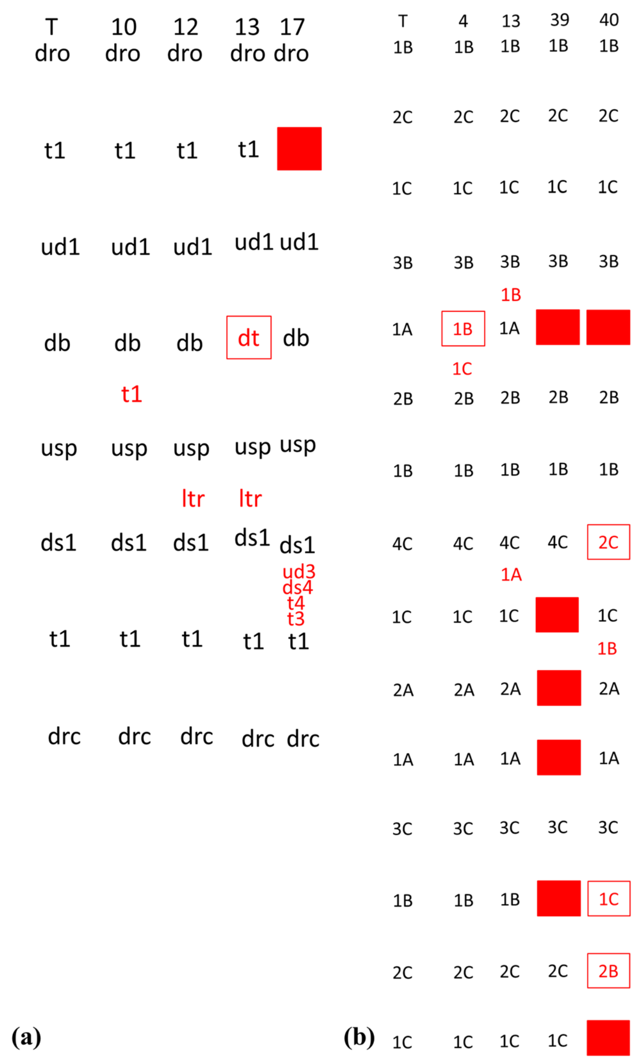


**Fig. 2** The full sequence alignment for the reference sequence T, two instances of successful and failed sequences for the HDDS (panel **a**) and the TOH (panel **b**) task. Each column shows an example sequence; each row shows different actions performed within the sequences. In the first column of each panel, "T" indicates that this is the "target" sequence (i.e., the reference sequence). The black actions indicate actions used in the reference sequence; these were the actions used for alignment. The red actions indicate extraneous actions that were not part of the reference sequence and instead were inserted during the human or robot trial. The red full boxes indicate a deletion for the corresponding action. The action denoted in red and boxed in red illustrate a substitution for the corresponding action by another one. For the HDDS task (panel **a**), the sequences 10 and 12 are imperfect whereas the sequences 13 and 17 are failed. For the TOH task (panel **b**), the sequences 4 and 13 are imperfect whereas the sequences 39 and 40 are failures. See Fig. S2, S3 and Table S3, S4 in the supplementary material for the alignment of all the sequences generated in both tasks

metrics; compare F and SP in Fig. 3a, b). Also, the number of deletions was significantly greater for the failed action sequences relative to those obtained for the successful imperfect or perfect ($z = 2.410$; $p = 0.030$ for both

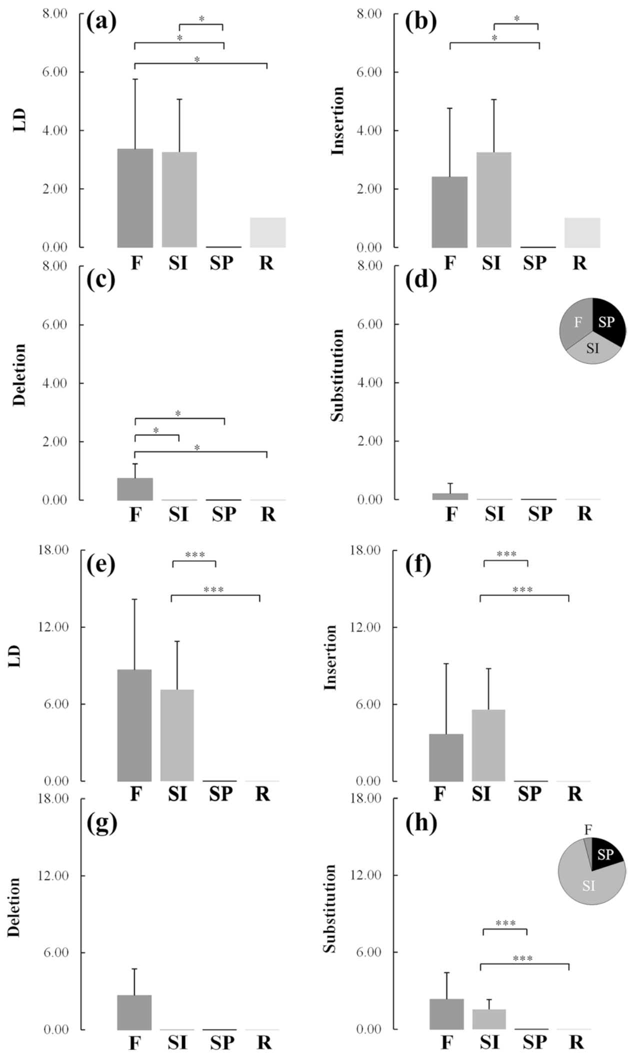


**Fig. 3** Average LD and numbers of the three operators for human participants for the failed (dark gray bars), imperfect (gray bars) and perfect (black bars; perfect replica of the demonstration) successful action sequences. The LD and its three operator for the humanoid robot action sequences (light gray bars). The panels **a**, **b**, **c** and **d**, represent the LD, the number of occurrences of insertion, deletion and substitution operators for the HDDS task, respectively. The panels e, f, g and h, represent the LD, the number of occurrences of insertion, deletion and substitution operators for the TOH task, respectively. F: failed action sequence; SI: successful imperfect action sequence; SP: successful perfect action sequence (including the reference sequence by definition) and R: action sequence generated by the humanoid robot. ***$p<0.001$; **$p<0.01$; *$p<0.05$

comparisons; compare F vs. SI and F vs. SP in Fig. 3c) action sequences. The same comparison did not reach the significance level for the number of substitutions ($p>0.303$ all comparisons considered; Fig. 3d).

Such prominence of deletions in the failed action sequences was confirmed by the correlational analysis that revealed that the number of deletions ($r=0.681$, $p<0.001$) was positively and significantly correlated to action

**Table 1** Performance outcome, average LD, I, D and S for the perfect, imperfect and failed action sequences for human execution

| | Perfect | Imperfect | Failed |
|---|---|---|---|
| *HDDS* | | | |
| Trial distribution | 33.33% | 31.67% | 35% |
| Average LD | 0 | $3.262 \pm 1.805$ | $3.359 \pm 2.396$ |
| Average I | 0 | $3.262 \pm 1.805$ | $2.415 \pm 2.353$ |
| Average D | 0 | 0 | $0.741 \pm 0.501$ |
| Average S | 0 | 0 | $0.204 \pm 0.351$ |
| *TOH* | | | |
| Trial distribution | 20% | 76% | 4% |
| Average LD | 0 | $7.128 \pm 3.675$ | $8.667 \pm 5.508$ |
| Average I | 0 | $5.582 \pm 3.206$ | $3.667 \pm 5.508$ |
| Average D | 0 | 0 | $2.667 \pm 1.202$ |
| Average S | 0 | $1.545 \pm 0.763$ | $2.333 \pm 2.082$ |

sequence failure. This was also observed for the number of substitutions ($r=0.495$, $p<0.001$) but not that of insertions ($r=0.219$, $p=0.090$). The LD was also positively and significantly correlated with action sequence failure, likely due to the contribution of deletion and substitution ($r=0.519$, $p<0.001$) (Fig. 4a; first row). The pronounced presence of deletion during failed action sequences is also consistent with the observation that this operator was absent whereas the humanoid robot successfully complete, although imperfectly (use of an extraneous action) this task (Fig. 3).

### 3.1.2 Towers of Hanoi

Considering performance of human individuals, the same statistical analysis performed for the TOH task revealed that although the LD and its operators were larger for the failed compared to the perfect successful action sequences, none of these differences reached the significance level ($p>0.05$; compare F and SP in Fig. 3e–g). Despite the lack of significance for the number of deletions, which was likely due to a small sample of failed trials leading to a lack of statistical power ($n=3$ representing 4% of all trials executed), it is the only metric which was found as non-zero for the failed trials compared to the perfect and imperfect sequences (i.e., see F vs. SI; F vs. SP; Fig. 3g). The specific presence of deletions in the failed action sequences was confirmed by the correlational analysis that revealed that the number of deletions ($r=0.990$, $p<0.001$) was the only metric which was positively and significantly correlated with action sequence failure (LD: $r=0.102$, $p=0.461$; I: $r=-0.024$, $p=0.912$; S: $r=0.006$, $p=0.961$) (Fig. 4a; second row). As for the HDDS task, the specific presence of deletions in action sequence failure is also in agreement with the result that this operator was absent while the humanoid robot generated a successful action sequence to execute the TOH task (see Fig. 3).

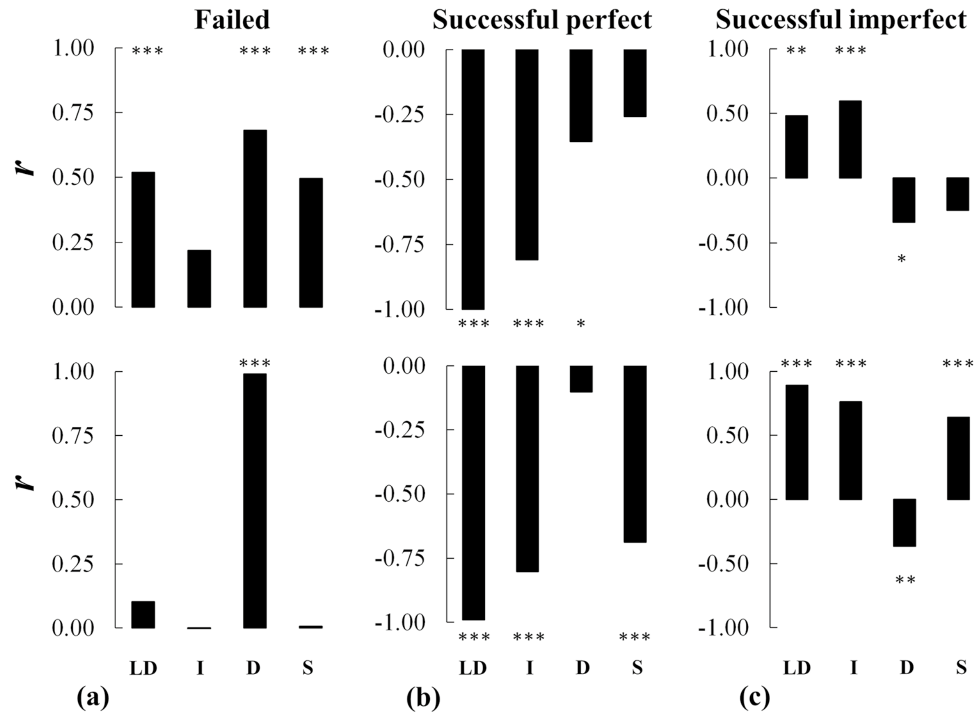


**Fig. 4** Correlations between the LD, number of insertions, deletions and substitutions and the failed (**a**), perfect (**b**) and successful but imperfect (**c**) execution of action sequences. The first and second rows represent the correlations obtained for the HDDS and the TOH tasks, respectively. I: Insertion; D: Deletion; S: Substitution. ***$p<0.001$; **$p<0.01$; *$p<0.05$

## 3.2 Successful Perfect and Imperfect Action Sequences in Humans and the Humanoid Robot

### 3.2.1 Hard Drive Docking Station

Statistical analysis conducted on the human data also revealed that the LD and the number of insertions (which was the sole contributor to the LD for imperfect sequences) were both greater for the imperfect compared to the perfect action sequences ($z=2.371$, $p=0.030$; for both metrics; compare SI and SP in Fig. 3a, b). This prominence of insertion as the main contributor of the LD was also observed with the humanoid robot performance since its LD exactly reflected the number of insertions which were generated to successful complete the HDDS task (Fig. 3a, b).

As expected, the four metrics (LD, I, D, S) were all negatively correlated with the successfully perfect action sequences. Significant correlations were obtained for the LD ($r=-1.00$, $p<0.001$), the number of insertions ($r=-0.809$, $p<0.001$), and of deletions ($r=-0.354$, $p=0.013$), while only a tendency was detected for the substitutions ($r=-0.257$, $p=0.056$) (Fig. 4b, first row). The LD ($r=0.481$, $p<0.001$) and the number of insertions ($r=0.595$, $p<0.001$) were positively and significantly correlated with the successful imperfect action sequences. However, the number of deletions and substitutions were negatively correlated with successful imperfect action sequences, while significant correlations were obtained for the number of deletions ($r=-0.340$, $p=0.013$) and a tendency was detected for substitutions ($r=-0.247$, $p=0.062$) (Fig. 4c, first row).

### 3.2.2 Towers of Hanoi

The same statistical analysis conducted for the TOH task revealed that the LD, the number of insertions (which was the main contributor to the LD for imperfect sequences) and the number of substitutions were larger for the imperfect relative to the perfect sequences for humans ($z\geq 3.625$, $p<0.001$; for these three metrics; compare SI vs. SP in Fig. 3e, f, h). For this task, the robot replicated the reference sequence (Fig. 3).

As anticipated, the four metrics were all negatively correlated with the successful perfect action sequences. Significant correlations were detected for the LD ($r=-1.00$, $p<0.001$), the number of insertions ($r=-0.802$, $p<0.001$), and substitutions ($r=-0.686$, $p<0.001$) but not deletions ($r=-0.102$, $p=0.461$) (Fig. 4b, second row). The same examination performed for the successful imperfect action sequences revealed that the LD ($r=0.890$, $p<0.001$), the number of insertions ($r=0.762$, $p<0.001$) and substitutions ($r=0.640$, $p<0.001$) were positively and significantly correlated with performance outcome. The number of deletions was negatively and significantly correlated with successful imperfect action ($r=-0.363$, $p=0.002$) (Fig. 4c, second row).

## 3.3 Practice, Primacy and Recency Effects

### 3.3.1 Hard Drive Docking Station

Statistical analysis conducted for the human individuals during the practice period revealed that the number of failed trials, the LD, and the occurrence of the three operators remained relatively unchanged throughout practice ($p > 0.05$; all metrics considered).

The humanoid robot did not reveal any change over the five trials since it was able to successfully perform the task after one trial and maintain this performance for the four remaining trials.

Finally, analysis of human data to explore any recency and primacy effects by examining whether the occurrence of the operators was more pronounced at the beginning, middle or end of the eight actions composing the reference sequence did not reveal any significant differences for the failed or imperfect successful action sequences ($p > 0.05$; all conditions considered). However, for the imperfect successful action sequences, the visual examination suggested that the insertions tended to further occur at the beginning and end of the sequences (Fig. 5a). The same visual examination of the results suggests that for the failed action sequences the deletions tend to occur more at the end and beginning of the sequence while the substitutions tend to rather occur in the middle of the sequence (see Fig. S4 in the supplementary material for further details). No primacy or recency effect was observed for the humanoid robot.

### 3.3.2 Towers of Hanoi

The same statistical analysis performed to assess changes during the practice period for the TOH task did not reveal any difference in the number of failed trials, the LD, and the occurrence of its three operators ($p > 0.05$; all metrics considered). As for the HDDS task, the humanoid robot did not reveal any change over the five trials since it was able to successfully perform the TOH task the first time and maintained the performance for the four other trials. However the analysis conducted for the imperfect successful action sequences revealed that the LD and the occurrence of the substitution operator were significantly larger in the middle relative to the beginning of the reference sequence ($z \geq 2.202$, $p < 0.026$; for all comparisons; see Fig. 5b). Similarly, the LD, as well as the number of insertion and substitutions were significantly greater in the middle compared to the end of the reference sequence ($z \geq 2.381$, $p < 0.015$; for all comparisons; see Fig. 5b). The same recency and primacy analysis did not detect any changes in the number of occurrence of the operators between the beginning, middle or end of the reference sequence for the failed action sequences ($p > 0.05$; all conditions considered; see Fig. S4 for further details). No primacy or recency effect was detected for the humanoid robot.

## 3.4 Direct Performance Comparison Between Humans and the Humanoid Robot

This work primarily aimed to deploy a computational tool to assess action sequence performance under various conditions of challenge and learning in humans and a humanoid robot. Thus the direct human versus humanoid robot comparison discussed in this section was secondary in this work and exploratory. When considering the HDDS task, the LD and number of deletions between the reference sequence and the human action sequences were significantly larger than those obtained between the reference sequence and the humanoid robot action sequence during task failure (LD: $z \geq 2.415$, $p < 0.030$ for both metrics; see F vs. R in Fig. 3a, c) while a similar number of insertions and substitutions was detected ($p > 0.281$ for all comparisons; see Fig. 3b, d). The same comparison for the imperfect successful sequences did not reveal any significant difference ($p > 0.05$; see SI vs. R in Fig. 3a–d). The same analysis for the TOH task did not reveal any human–robot difference for the failed sequences ($p > 0.05$ for all comparisons; see F vs. R in Fig. 3e–h) whereas for the imperfect successful human action sequences the LD, number of insertions and substitution significantly differed from those generated by the humanoid robot ($z \geq 3.625$; $p < 0.001$; see SI vs. R in Fig. 3e, f, h).

# 4 Discussion

Overall, when examining the action sequences generated by the human individuals and the humanoid robot in both tasks, which involved different amounts of challenge and learning, the main findings were that: (i) human participants generated a wide variety of action sequences in both tasks not only in terms of success versus failure, but also in terms of the different strategies used in imperfect successful trials, whereas the humanoid robot typically produced perfect or near-perfect sequences after just one trial; (ii) the LD and number of insertions reliably differentiate perfect and imperfect successful attempts and, (iii) deletion was the most sensitive marker of failure, representing the main contributor to performance failure independently of the level of challenge, learning paradigm and types of performer considered here.

## 4.1 Assessing Complex Action Sequence Performance in Human and Humanoid Robots

By assessing in detail the structure of each action sequence produced by the humans and the humanoid robot in both

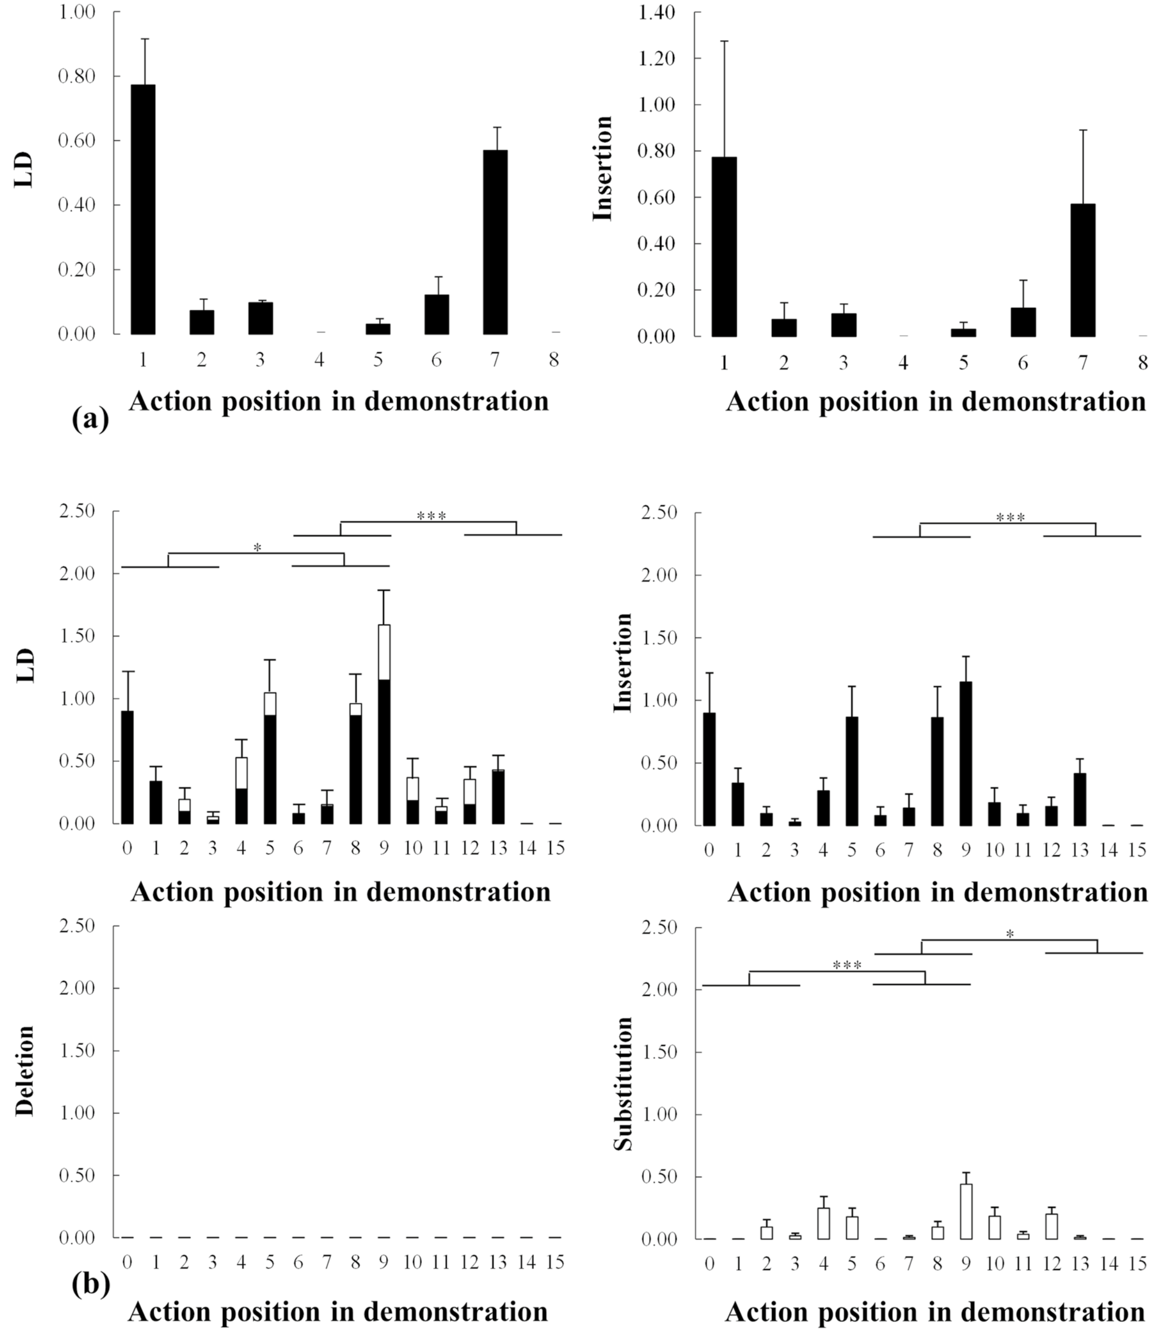


**Fig. 5** Average LD and number of occurrences of the three operators in humans at each of the eight and fifteen positions in the reference (demonstrated) action sequence, for the successful imperfect action sequence executed for the HDDS (panel **a**) and the TOH (panel **b**) task, respectively. For the HDDS task the number of deletions and substitutions were zero for the successful imperfect sequences, for clarity they are not represented here. For the same reason the successful perfect action sequences are not represented for both tasks since their LD and operator occurrences are all zero. Bar heights indicate number of operations at a specific position in the case of deletions and substitutions, and number of operations immediately after a specific position in the case of insertion. For example, no actions were ever inserted before the very first action (opening the drawer) or after the very last (closing the drawer) for the HDDS task

tasks it was possible to not only detect if the sequence was a perfect, imperfect (but still reaching the task goal) or failed action sequence, but also to determine specifically what actions were responsible for the performance outcome. For the HDDS task the findings revealed that a third of the human action sequences were not correctly executed since participants either did not press the release switch to turn off the disk before changing it, or simply failed to reach the task goal (i.e., incorrect order and/or incomplete sequence). Interestingly for the TOH task only 4% of the trials failed. Critically, while no statistical significance was detected likely due to a lack of statistical power because of the very few trials, as predicted, deletions only occurred in and caused failure trials; they may thus serve as a robust marker of failure. Possibly one cannot recover from deletion of a critical operation of the reference sequence, without reinserting the missing action at the same or a nearby position. The deletion operator as a marker of failed action sequence may have some degree of generality since consistent results were obtained for the two tasks considered here which involved different levels of challenge and learning paradigms. However, deletion as a marker of performance failure is also probably related to the specific type of tasks and learning processes considered in this work where there was a specific order of the actions followed to execute the task, one minimal reference sequence and few execution rules. While this result illustrates how, at least in this study, the proposed method can uncover performance markers which can complement more basic metrics, it is likely not valid for all tasks and/or learning types. Contrary to deletions, the substitutions revealed a less clear pattern of results. For the TOH task the use of the substitution operator was more frequent for the imperfect compared to the perfect sequences. This difference was not observed for the HDDS task.

When considering the distribution of a particular operator among the action positions that formed the reference sequence, no statistical differences were detected for the failed action sequences in both the HDDS and TOH tasks (Fig. S4). However, regarding the imperfect successful sequences, while no change in primacy and regency was identified for the HDDS task, some differences in LD, insertion and substitution were observed for the TOH task (Fig. 5a, b). For this task, the LD and number of substitutions were larger in the middle relative to the beginning and end of the sequence while the number of insertions was greater in the middle compared to the beginning of the sequence which is consistent with prior primacy/recency results for working memory (see actions 0–3, 6–9, 12–15 in Fig. 5b). Possibly while human individuals explore the high-level plan for the TOH task, its early and late components are better memorized compared to the middle component before execution. Although this phenomenon should be further examined, these results may potentially extend these notions initially identified in pure task memorization to certain complex motor tasks [37]. In addition, noticing that the trailing portion of a 5-disk solution is a 4-disk sub-problem and that the trailing portion of a 4-disk solution is a 3-disk sub-problem, a complementary possible explanation would be that performance improves near the end of the sequence since smaller sub-problems are easier to recognize and solve.

Our approach applied to the HDDS and TOH task revealed that among the human successful sequences, 51.27% and 20.83% of those perfectly replicated the demonstrated sequence (demonstration-human $LD = 0$) while 48.72% and 79.17% did not (demonstration-human $LD > 0$), respectively. However, it is difficult to identify which task had the best performance, since while the HDDS task had more flexible execution rules than the TOH task, the latter required more steps. Overall, individuals failed less when completing the TOH task (TOH: 4% and HDDS: 35% of failed trials) although 25% of the TOH trials could not be completed on time. However within successful trials, the TOH task led to inferior performance (TOH: 20% and HDDS: 33.33% perfect trials). Future work could further examine this by assessing the level of mental workload deployed by the individuals during performance [38–41].

For both tasks these successful but imperfect action sequences were of particular interest since they represent different ways to complete the task, something that is often observed in real-world settings. Our method revealed that for both tasks the lack of perfect match of these successful sequences was mainly due to the insertion of additional unnecessary actions during sequence imitation (no deletions and few substitutions were identified). For instance, instead of picking up the disk with the right hand and directly placing it in the available slot, the disk could be grasped with the left and then transferred to the right hand to be finally placed in the slot (see sequence #12, Fig. 2a). Such change of hands was observed in humans and in the humanoid robot. Similarly, for the TOH task two actions were added (disk 1 moved to the peg B; disk 1 to the peg A; sequence #13, Fig. 2b). Adding such extraneous actions (which could be done multiple times) increases the length of the sequence without however compromising the success of the task completion. Thus, while the presence of insertions may prevent a perfect imitation, they may have a low probability of impeding successful task completion, contrary to deletions. Possibly, inserting actions before/after each of the core atomic actions of the reference sequence is not problematic as long as these actions are produced in the correct order (and while following the task completion rules).

A possible explanation for production of extraneous actions for both the HDDS and TOH tasks by humans could be that (i) the multiple actions composing the reference sequence may be difficult to memorize (HDDS) or identify (TOH); (ii) since the task can be executed in various ways while still successfully reaching its goal, humans potentially learn to solve the task by generating multiple plans. Along those lines, while the analysis of the memory processes used here was not the primary aim of this work, it seems reasonable to ask if pure or logical memory processes would have been preferentially engaged during performance of failed, perfect and imperfect action sequences. When considering that pure (or rote) memory reflects encoding of information without its full understanding and/or context whereas logical memory is based on an intelligent understanding and logical thinking, it could be expected that a short sequence that can be solved without using a particular logical procedure may likely further promote the engagement of pure memory. Specifically, here, compared to the TOH, the completion of the HDDS task did not involve any particular solution to be identified such as a recursive procedure while also providing a much shorter sequence length (eight vs. fifteen actions) which was somewhat closer to the working memory capacity [42]. Thus, the execution of the HDDS task may have rather relied on pure memory whereas the performance of the TOH task may have further involved logical memory mechanisms. Keeping this in mind, and the fact that compared to HDDS, the TOH task revealed a smaller number of perfect/failed sequences but also a much greater number of successful imperfect sequences, it could be suggested that the pure and logical memory may have a more prevalent role in generating the perfect/failed and imperfect sequences, respectively. Although possible, the differential engagement of such memory mechanisms needs to be further investigated in future work.

Since here there was only one optimal sequence which served as a reference, the action sequences with a zero LD were all identical to the reference sequence and thus between themselves. However, those having the same, but non-zero, LD were not necessarily identical to each other since it is unknown to which actions the operators were applied. In particular, action sequences that have the same non-zero LD do not necessarily use the same high-level motor plan since the action sequence could have the same length but include different actions. Thus, having the same LD is a necessary but not a sufficient condition to have two sequences using the same high-level motor plan. To further study high-level plans deployed by the human and the humanoid robot for sequences with the same LD, the visual representation of actions sequences allows us to examine if an exact same path captured two or more sequences and which actions drove any discrepancies between those sequences. For instance for the HDDS task, the sequence generated by the humanoid robot (sequence #12, Fig. 2a), and another one produced by a human participant (sequence #10, Fig. 2a), both have a LD = 1 from the demonstration (both had one single different action inserted at a different position). Similarly, for the TOH task the sequences # 4 and 13 have the same LD (LD = 2) but the former has a substitution followed by an insertion whereas for the latter two actions were added. Without any visualization the same LD obtained for the two sequences would not reveal their different structure.

## 4.2 Comparison to Existing Approaches to Assess Complex Action Sequences

Our results complement past work by extending prior findings from human motor behavior and robotic studies that evaluated human and/or robot performance during execution of complex tasks (e.g., Towers of Hanoi, equipment maintenance and cleaning) with relatively basic task-specific metrics (e.g., number of moves, errors). Namely, here we have deployed a computational approach able to examine in detail the structure of high-level motor plans underlying action sequences via qualitative (visualization) and quantitative (LD and its operators) approaches which can potentially be applied to various tasks [1, 4–8]. More specifically, our current study complements past robotic work that has assessed complex motor tasks via fairly coarse-grained metrics (e.g., execution times, success rates) or survey-based subjective measures (e.g., trust, mental workload) to assess human–robot teaming [9–15, 43]. Our work also complements more refined metrics based on time-based measurements (e.g., reaction times of the robot/human teammate [10, 13, 16, 17]) or the segmentation of task execution into action sequences (e.g., number/cost of actions performed [9, 13, 18–21]). Nikolaidis and colleagues [10] model action sequences as a mixed-observability Markov decision process and compute distances between their transition probability matrices. Their approach works with a closed sensorimotor loop, handles both planning and analysis of robotic motion, and emphasizes sub-symbolic sensorimotor tasks, such as positioning and orienting a single box. It can be viewed as complementary to our work, which focuses on cognitive-level action sequences with many object positioning tasks (e.g., pickups, hand-offs, put-downs), and without any Markov requirement. The planning system we used with our robot, which uses less frequent sensory feedback control, but higher-level cognitive reasoning, is detailed in [8]. Complementary to [10], by indicating specific LD operators and specific steps in the action sequence where they occurred, our approach provides specific descriptions of exactly how two action sequences differ. Our approach also includes the visual alignment of sequences for qualitative analysis. Our distance metric complements this clustering-based approach on higher-level cognitive tasks, and since it does not rely on a Markov assumption, it can capture longer-term dependencies in action sequences. Only a few prior studies employed LD to study human or robot motor control, however the tasks considered were simpler compared to those employed here, being more sensorimotor in nature and with a limited or no cognitive component [44, 45].

The application of the LD to the motor domain can account for both the level of abstraction and the physical constraints. In the motor domain, the “alphabet” can include not only the action to be selected depending on the cognitive-motor stage using high-level abstract representations (e.g., opening the drawer), but also on sensorimotor coordination for action implementations (e.g., use of the left effectors) which rely on a lower level of abstraction. While adding performance constraints (e.g., requirements to exactly replicate the demonstration with the right arm only) can mitigate or even solve this problem, the task becomes less realistic. Also, in the motor domain, humans and/or robots interact with physical systems that provide task constraints that are naturally enforced (e.g., in the HDDS task, the dock must be opened first before any disk manipulation; imitators and demonstrators have different physical constraints such as the number of degrees of freedom). However, such specificities may not be encountered for motor tasks that are less complex and more constrained, limiting the use of different effectors and implicit rules. The studies mentioned above that employed LD in humans or robots did not face these two problems due to the use of fairly simple tasks [44, 45]. Thus, our approach can assess complex action sequences while still complementing simpler metrics currently available to comprehensively assess complex task performance by humans and humanoid robots.

### 4.3 Applications to Human, Humanoid Robots and Human–Robot Interactions

Our approach to analyzing human motor sequences could be employed to assess performance in both humans and humanoid robots. Specifically, the understanding of how human high-level motor plans are generated and adapted could inform the design of neurocognitive architectures for cognitive robots, enhancing their learning and performance capabilities. Our approach could also assess how human and humanoid robot high-level motor plans differ, allowing testing and prediction of the quality of human–robot team dynamics. Namely, it is possible that a humanoid robot learns to perform a complex maintenance task in a certain way whereas humans execute it in a very different manner. Thus, our approach could accurately quantify and predict how much the humanoid robot and humans differ when completing the task. As such, we could predict that as the level of discrepancy between the humanoid robot and a human increases (e.g., elevation of LD and of insertions occurrences) the human–robot team dynamics would translate from an adaptive to a maladaptive state [41, 46, 47]. Thus, such a quantification of strategy differences in the humans and humanoid robot could inform their respective training and also predict or at least inform how their shared mental models may differ or not. If a human and a robot solve a problem with very different strategies it may be difficult for both to predict or infer efficiently the future states and goal understanding of each other's behavior leading ultimately to poor team dynamics [24, 26, 48–50]. Such an approach focusing on high-level motor planning can complement prior work that has examined shared control at lower sensorimotor levels for enhancing human–robot interactions [41, 46, 51, 52]. Also, our approach could reveal whether some specific changes in the sequence (e.g., insertion of a specific action) may be the main drivers promoting adaptive or maladaptive team environments. This could be done in a context where the humanoid robot and humans perform or are trained together. The fact that here the humanoid robot could perform consistently the tasks learned after one single attempt whereas this was not observed for humans (likely needing more practice) suggests that the difference in robot and human learning capabilities would likely influence the human–robot collaborative learning dynamics. Although many factors can lead to a difference between human and robot learning capabilities, here a potential factor could be the process with which the actions of the sequence shown in the video were visually processed and encoded by each performer since the former recruit well-established neural mechanisms (e.g., dorsal and ventral stream) whereas a text-based transcription was provided to the robot via SMILE. Another source of difference could have been that the neurocognitive architecture of the humanoid robot had enhanced pure memorization capabilities relative to humans. However, it is critical to note that the neurocognitive architecture which controls the robot did not simply perform a rote memorization of the demonstrated action sequences, but instead inferred the goals, decomposed them into actions, and then executed those actions. As suggested earlier (see Sect. 4.1), although it is difficult to say if pure or logical memory mechanisms would have been preferentially engaged to perform the HDDS and TOH tasks in humans, for the neurocognitive architecture, the actions and the goals could potentially reflect an analog of pure and logical memory, respectively. Thus, the ratio computed between the goals and the actions for both HDDS and TOH tasks resulted in about 0.5 (1:2). This suggests that for both tasks a combination of pure and logical memory processes was employed by the robot. Although further examination is needed, both the humanoid robot and the human participants likely engaged a combination of pure and logical memory processes to complete the tasks considered here, however the differences in memory mechanisms of the humanoid robots and human individuals likely contributed to differences in performance.

Also, although our work involved a humanoid robot and humans, our approach could be deployed exclusively with robots or humans. Regarding the latter, this work could serve to compare high-level motor plans in both healthy and impaired individuals informing thus the compromised underlying cognitive-motor processes. Our approach could inform how the motor plan generation and more generally the human cognitive-motor processes are adapted during performance and learning of new complex tasks.

## 5 Conclusion and Future Work

This work is just a first step in examining qualitatively and quantitatively the cognitive-motor performance outcome in humans and robots during practice of complex action sequences. Our results provide different but complementary metrics to those already existing to assess complex task performance by quantifying the differences between action sequences. The LD in the cognitive-motor domain can be used with varying levels of abstraction in the context of physical constraints inherent to real-world interactions. While the LD is informative, the occurrence of its specific operators may, to some degree, provide a more refined assessment of cognitive-motor performance. Also, although this work was conducted in both human and robot, it can also be exclusively applied to one or the other in a context of motor performance and learning. Our approach not only allows us to examine cognitive-motor performance by assessing similarity between two given sequences, but also to identify the sources of these differences (e.g., if some specific portions of two sequences are similar/different, at which point these sequences differ (where branching occurs), recency/primacy effects). Thus, although the proposed approach, which is fairly simple to implement, is not a general measure of performance and further work is needed, it provides rich information which extends and complements other more basic metrics currently available.

As with any methodology, our approach also has limitations. One possible weakness of this method is the situation where a task includes many trials with many steps. In this case, the visual representation, and in particular the action sequence alignment used here, may not be able to represent clearly all of the trials at once but only under specific configurations. Also, although our computational tool was tested in two different tasks which taken together involved various levels of challenge and learning paradigms, only two tasks were employed here. Another limitation is related to the definition of optimal sequence. Here as a first step only one optimal sequence, defined as having a minimum number of actions, was considered. While our approach does not account for tasks having multiple optimal sequences as can be the case in real world settings, it can still compare two sequences, regardless of whether they are optimal or not. Under such conditions, our method could not determine the optimality of the sequence, however it could still assess how sequences (optimal or not) differ and ultimately inform how similarly two agents solve a given problem. In addition, if a performed action sequence is optimal, it will have a null LD to one of the optimal sequences that serve as references. Moreover, its distance *per-se* to other distinct optimal references, or the distance from sub-optimal sequences to one of these references would be non-zero and thus would vary depending on which distance metric is employed. For instance, the same limitations would likely arise with a more sophisticated LD which would include other operators (e.g., transposition) although the values of the distance computed would differ between the traditional LD as employed here and those having additional operators. Since for now it is uncertain how our approach could handle multiple optimal reference sequences, it is unclear if similar results would be obtained under such a condition.

Also, the minimum number of actions as optimality criteria was used since our approach aims to assess high-level planning which essentially relates to the number and order of actions in a sequence, which are the usual parameters considered in human planning studies [4, 6, 53, 54]. However, various optimality criteria (e.g., energy, kinematics-based) could also assess performance with some being possibly better suited than others for certain types of tasks [55]. In this case, a similar approach could be considered by employing various criteria (e.g., energy consumption, motion smoothness) for each action to compare two sequences while also computing the LD for some metrics resulting in multiple optimality criteria to assess the performance (similarly to multi-objective optimization). In this regard, for a given optimality criteria if multiple optimal sequences exist, the combination of all optimality criteria may result, at least in some cases, in only one unique optimal sequence and thus may solve the limitation mentioned above that our approach currently cannot handle multiple optimal sequences when using only one optimality criteria which is the minimum number of actions. Moreover, in some situations the current approach may provide limited information as could be the case for some tasks that are simple by nature or because of the way they are modeled. For the former, there is likely no need to use such an approach which however was initially developed to study performance in tasks which cannot be executed with basic action sequences. For the latter, if an atomic action is defined at a more general level the task can appear as simplified and our method may be less interesting. However, in this case we would lose the granularity of the action sequence analysis. Also, for certain contexts such as in the presence of perturbations or in a situation of complex interactions between two agents this approach may be more limited. For instance, in a collaborative context where two motor sequences are modified online and interleaved as two agents perform a task, our approach cannot currently capture these dynamics. Although, this approach, as a first step, did not aim to deal with such interacting contexts, it could still compare the action sequences performed by the human and the robot independently; this can inform their subsequent interactions in executing a given task. Although the usefulness of this approach may be limited for some situations, it could also compare action sequences executed in two different contexts to examine how their alterations are context-dependent.

Depending on the nature of the context this approach may need to be complemented by other metrics (e.g., energy) that altogether can collectively inform the performance.

Finally, we note that only a few trials were considered here and as such the same computational method could be employed to a larger number of trials to capture the learning dynamics. Future work could further examine performance in both humans and humanoid robots while considering (i) additional tasks to further assess the performance outcome markers of this LD-based method, (ii) more trials to assess learning dynamics throughout practice which should be captured by the LD and the occurrence of its operators, (iii) by considering more sophisticated LD computation (e.g., transposition operator, normalization [22, 32, 50]) or alternative string distance metrics other than LD, and (iv) in a social context of team dynamics.

**Acknowledgements** This work was supported by The Office of Naval Research (N00014-19-1-2044).

### Compliance with Ethical Standards

**Conflict of interest** The authors declare that they have no conflict of interest.

## References

1. Tenorth M, Ziegltrum J, Beetz M (2013) Automated alignment of specifications of everyday manipulation tasks. In: International conference on intelligent robots and systems (IROS), 2013 IEEE/RSJ, pp 5923–5928
2. Tenorth M, Perzylo AC, Lafrenz R, Beetz M (2013) Representation and exchange of knowledge about actions, objects, and environments in the roboearth framework. IEEE Trans Autom Sci Eng 10(3):643–651
3. Breazeal C, Buchsbaum D, Gray J, Gatenby D, Blumberg B (2005) Learning from and about others: towards using imitation to bootstrap the social understanding of others by robots. Artif Life 11(1–2):31–62
4. Goel V, Grafman J (1995) Are the frontal lobes implicated in "planning" functions? Interpreting data from the Tower of Hanoi. Neuropsychologia 33(5):623–642
5. Saint-Cyr JA, Taylor AE, Lang AE (1988) Procedural learning and neostriatal dysfunction in man. Brain 111(4):941–959
6. Welsh MC, Huizinga M (2005) Tower of Hanoi disk-transfer task: influences of strategy knowledge and learning on performance. Learn Individ Differ 15(4):283–298
7. Katz GE, Huang DW, Gentili RJ, Reggia JA (2016) Imitation learning as cause-effect reasoning. In: Steunebrink B, Wang P, Goertzel B (eds) Artificial general intelligence. AGI 2016, AGI 2016. Lecture notes in computer science, vol 9782. Springer, Cham
8. Katz G, Huang DW, Hauge T, Gentili R, Reggia J (2017) A novel parsimonious cause-effect reasoning algorithm for robot imitation and plan recognition. IEEE Trans Cognit Dev Syst PP(99):1–17
9. Nathan K, Takayama L, Matarić M (2010) Communication and knowledge sharing in human–robot interaction and learning from demonstration. Neural Netw 23(8):1104–1112
10. Nikolaidis S, Ramakrishnan, Gu K, Shah J (2015) Efficient model learning from joint-action demonstrations for human-robot collaborative tasks. In: Proceedings of the tenth annual ACM/IEEE international conference on human-robot interaction. ACM
11. Nikolaidis S, Hsu D, Srinivasa S (2017) Human-robot mutual adaptation in collaborative tasks: models and experiments. Int J Robot Res 36:618–634
12. Paxton C, Jonathan F, Hundt A, Mutlu B, Hager GD (2017) User experience of the CoSTAR system for instruction of collaborative robots. arXiv preprint arXiv:1703.07890
13. Steinfeld, Aaron, Fong T, Kaber D, Lewis M, Scholtz J, Schultz A, Goodrich M (2006) Common metrics for human–robot interaction. In: Proceedings of the 1st ACM SIGCHI/SIGART conference on human-robot interaction. ACM
14. Wang N, Pynadath D, Hill S (2015) Building trust in a human robot team with automatically generated explanations. ITSEC paper, vol 15315 (2015), pp 1–12
15. Zhang Y, Narayanan V, Chakraborti T, Kambhampati S (2015) A human factors analysis of proactive support in human–robot teaming. In: IEEE/RSJ international conference on. intelligent robots and systems (IROS)
16. Levine SJ, Williams BC (2014) Concurrent plan recognition and execution for human-robot teams. In: ICAPS'14 proceedings of the twenty-fourth international conference on international conference on automated planning and scheduling, pp 490–498
17. Shah J, Wiken J, Williams B, Breazeal C (2011) Improved human–robot team performance using chaski, a human-inspired plan execution system. In: Proceedings of the 6th international conference on Human–robot interaction. ACM
18. Butchibabu A, Sparano-Huiban C, Sonenberg L, Shah J (2016) Implicit coordination strategies for effective team communication. Hum Factors 58(4):595–610
19. Freedman RG, Shlomo Z (2017) Integration of planning with recognition for responsive interaction using classical planners. In: Thirty-first AAAI conference on artificial intelligence, San Francisco, CA, USA, pp 4581–4588
20. Salter T, Dautenhahn K, Boekhorst R (2006) Learning about natural human–robot interaction styles. Robot Auton Syst 54(2):127–134
21. Shah J, Breazeal C (2010) An empirical analysis of team coordination behaviors and action planning with application to human–robot teaming. Hum Factors 52(2):234–245
22. Ho T, Oh SR, Kim H (2017) A parallel approximate string matching under Levenshtein distance on graphics processing units using warp-shuffle operations. PLoS ONE 12(10):e0186251
23. Mohammad Y, Nishida T (2010) Controlling gaze with an embodied interactive control architecture. Appl Intell 32(2):148–163
24. Mathieu JE, Hener TS, Goodwin GF, Salas E, Cannon-Bowers JA (2000) The influence of shared mental models on team process and performance. J Appl Psychol 85(2):273–283
25. Scheutz M (2013) Computational mechanisms for mental models in human–robot interaction. In: Shumaker R (ed) Virtual augmented and mixed reality. Designing and developing augmented and virtual environments. VAMR 2013. Lecture notes in computer science, vol 8021. Springer, Berlin
26. Sheridan TB (2016) Human–robot interaction: status and challenges. Hum Factors 58(4):525–532
27. Magill RA, Anderson DI (2017) Motor learning and control: Concepts and applications, 11th edn. McGraw-Hill, Boston
28. Huang DW, Katz GE, Langsfeld JD, Gentili RJ, Reggia JA (2015) A virtual demonstrator environment for robot imitation learning. In: IEEE international conference on technologies for practical robot applications (TePRA), Woburn, MA, USA, pp 1–6
29. Huang DW, Katz GE, Langsfeld JD, Oh H, Gentili RJ, Reggia JA (2015) An object-centric paradigm for robot programming by demonstration. In: International conference on augmented cognition, Springer, Los Angeles, CA, USA, pp 745–756

30. Palmer DW (1996) Exploring recursion with variations on the Towers of Hanoi. J Comput Sci Coll 12(96):59–66
31. Levenshtein VI (1966) Binary codes capable of correcting deletions, insertions, and reversals. Sov Phys Dokl 10(8):707–710
32. Yujian L, Bo L (2007) A normalized Levenshtein distance metric. IEEE Trans Pattern Anal Mach Intell 29(6):1091–1095
33. Wagner RA, Fischer MJ (1974) The string-to-string correction problem. J ACM (JACM) 21(1):168–173
34. Deese J, Kaufman RA (1957) Serial effects in recall of unorganized and sequentially organized verbal material. J Exp Psychol 54:180–187
35. Murdoch BB (1962) The serial position effect of free recall. J Exp Psychol 64(5):482–488
36. Benjamini Y, Hochberg Y (1995) Controlling the false discovery rate: a practical and powerful approach to multiple testing. J R Stat Soc Ser B (Methodol) 57(1):289–300
37. Morrison AB, Conway AR, Chein JM (2014) Primacy and recency effects as indices of the focus of attention. Front Hum Neurosci 24(8):6 **(eCollection 2014)**
38. Gentili RJ, Jaquess KJ, Shuggi IM, Oh H, Lo LC, Tan YY, Domingues CA, Blanco JA, Rietschel JC, Miller MW, Hatfield BD (2018) Combined assessment of attentional reserve and cognitive effort under various levels of challenge with a dry EEG system. Psychophysiology 55(6):e13059
39. Shaw EP, Rietschel JC, Hendershot BD, Pruziner AL, Miller MW, Hatfield BD, Gentili RJ (2018) Measurement of attentional reserve and mental effort for cognitive workload assessment under various task demands during dual-task walking. Biol Psychol 134:39–51
40. Shuggi IM, Oh H, Shewokis PA, Gentili RJ (2017) Mental workload and motor performance dynamics during practice of reaching movements under various levels of task difficulty. Neuroscience 360:166–179
41. Shuggi IM, Shewokis PA, Herrmann JW, Gentili RJ (2017) Changes in motor performance and mental workload during learning of reaching movements: a team dynamics perspective. Exp Brain Res 236(2):433–451
42. Cowan N (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. Behav Brain Sci 24(1):87–185
43. Omote K, Feussner H, Ungeheuer A, Arbter K, Wei GQ, Siewert JR, Hirzinger G (1999) Self-guided robotic camera control for laparoscopic surgery compared with human camera control. Am J Surg 177(4):321–324
44. Gorbenko A, Popov V (2015) On the adjustment of the weights of the Levenshtein distance for the description of sequences of rhythmic motor primitives. Contemp Eng Sci 8(18):835–840
45. Holm L, Karampela O, Ullén F, Madison G (2017) Executive control and working memory are involved in sub-second repetitive motor timing. Exp Brain Res 235:787–798
46. Losey DP, McDonald CG, Battaglia E, O'Malley MK (2018) A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction. Appl Mech Rev 70(1):1–19
47. Miller MW, Presacco A, Groman LJ, Bur S, Rietschel JC, Gentili RJ, McDonald CG, Iso-Ahola SE, Hatfield BD (2014) The effects of team environment on cerebral cortical processes and attentional reserve. Sport Exerc Perform Psychol 3(1):61–74
48. Ahrndt S, Albayrak S (2016) Joint human-agent activities: challenges and definition, vol 9872. Springer, Cham, pp 105–112
49. Hanna N, Richards D (2018) The impact of multimodal communication on a shared mental model, trust, and commitment in human–intelligent virtual agent teams. Multimodal Technologies Interact 2:48
50. Honig S, Oron-Gilad T (2018) Understanding and resolving failures in human-robot interaction: literature review and model development. Front Psychol 9:861
51. Gillespie RB (2001) A general framework for cobot control. IEEE Trans Robot Autom 17(4):391–400
52. Boehm- P, Ghasemi AH, O'Modhrain S, Jayakumar O, Gillespie RB (2016) Architectures for shared control of vehicle steering. IFAC-PapersOnLine 49–19:639–644
53. Fitter NT, Kuchenbecker KJ (2016) Qualitative user reactions to a hand-clapping humanoid robot. In: Agah A, Cabibihan JJ, Howard A, Salichs M, He H (eds) Social robotics. ICSR 2016. Lecture notes in computer science, vol 9979. Springer, Cham
54. Vakil E, Lev-Ran Galon C (2014) Baseline performance and learning rate of conceptual and perceptual skill-learning tasks: the effect of moderate to severe traumatic brain injury. J Clin Exp Neuropsychol 36(5):447–454
55. Todorov E (2004) Optimality principles in sensorimotor control. Nat Neurosci 7(9):907–915


**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Theresa C. Hauge** received her M.A. in Kinesiology from the University of Maryland-College Park in 2018 where her work focused on the assessment of cognitive-motor performance during the execution of complex action sequences. Currently she is in the Ph.D. program in the Department of Applied Physiology & Kinesiology at the University of Florida.

**Garrett E. Katz, PhD,** is an assistant professor at Syracuse University. He received a Ph.D. in computer science from University of Maryland in 2017, an M.S. in mathematics from City College of New York, and a B.A. in philosophy from Cornell University. His current research focuses on artificial intelligence, neural computation, and robotics.

**Gregory P. Davis** is a PhD student at the University of Maryland, Department of Computer Science. His research interests lie at the intersection of artificial intelligence and the neural and cognitive sciences, and he aims to develop biologically-inspired neurocomputational models compositional memory and hypothetico-deductive causal reasoning.

**Di-Wei Huang, PhD,** received his PhD degree from the Department of Computer Science at the University of Maryland, College Park. His main interests include neural machine learning and neurocognitive architectures, as well as their applications such as robotics and recommendation systems.

**James A. Reggia, MD, PhD,** is a Professor of Computer Science at the University of Maryland, with joint appointments in the Institute for Advanced Computer Studies and in the Neuroscience & Cognitive Science Program. His research interests include neural computation, machine learning, robotics, and cognitive aspects of motor control; he has authored over 200 publications in these areas. His recent research has focused on imitation learning of complex sequential motor tasks in cognitive robotic systems and how such robotic learning relates to human learning.

**Rodolphe J. Gentili, PhD** is currently an Associate Professor in the Department of Kinesiology and a faculty member in the Neuroscience and Cognitive Science Program as well as the Maryland Robotics Center at the University of Maryland-College Park. His research examines the underlying cognitive-motor control and learning processes in humans by combining experimental, computational and robotics approaches. This integrative perspective can inform human motor behavior, the development of neuromimetic control systems and human–machine interactions.