Human-guided gait learning for a bipedal walking robot

Lars Besselaar (4489136), Erik van den Broek (4466764),
Dennis Koens (4603591) and Marlies Popken (4568400)

***Abstract*— Machine learning approaches are largely applied to robots that are made for thorough, repetitive learning. While some applications may benefit from these techniques, more delicate robots, such as bipedal walkers, tend to benefit from gentler methods. In this research, the viability of an alternative teaching method has been evaluated. The method consists of two phases; the first phase teaches a basic, imperfect gait by means of kinesthetic teaching, while the second phase adjusts and improves the aforementioned gait by means of human corrective feedback. To improve the effectiveness of the feedback phase and to increase intuitiveness, interfaces have been designed. The interfaces have been evaluated by means of a small experiment, involving participants assessing the intuitiveness and performance of the designs. The interface coined kinesthetic feedback has proven to be the most intuitive, while the interface that uses keyboard feedback is more effective for incremental adjustments. Although the method as a whole produced rigid and inflexible gait, the taught gait was stable. It can therefore be concluded that the method is effective when applied to repetitive and consistent motions.**

I. Introduction

The absence of bipedal robots in today's modern society is indicative of how challenging it is to teach a robot how to walk. With robot intelligence increasing, their potential to move without human guidance has not been commercially used yet. One possible cause of this difficulty is the lack of alternative teaching methods besides machine learning approaches.

The goal of this research was to evaluate the viability of an unconventional teaching method, which includes demonstrative teaching and human corrective feedback. As part of this evaluation, two independent interfaces have been compared. The main criteria for this comparison were intuitiveness and effectiveness.

Reinforcement learning and other trial-and-error methods are consciously not considered in this research, mainly due to the mechanical challenges these methods present. Trial-and-error methods involve falling and failing over a large period of time. Many robots, including the one used in this research, are not made to fall hundreds of times, nor are they built to run continuously for hours or days [1]. Even if trial-and-error methods were used on robots that are equipped to handle these impacts and running times, they would still have to run for a long period, in the order of days, until they were functional. While some robots are made for thorough, repetitive learning, some are not. More delicate robots, like bipedal walkers, tend to benefit from gentler methods, such as human-guided learning. The absence of reliable and gentle methods supports the case for further research into kinesthetic methods.

Addressing this problem, this paper describes research into a relatively new, more hands-on approach to gait learning. First, a human teacher will demonstrate the desired motion to the robot by means of kinesthetic teaching [2]. Kinesthetic teaching is one of the branches of demonstrative teaching and holds the most potential in this method for its simplicity and natural intuitiveness. This procedure will result in a motion that closely resembles the desired motion, but is not yet independently stable. Second, an interface has been designed to allow the human teacher to provide corrective feedback to the robot [3]. Using the interface, the teacher is able to adjust the robot's motion until it performs the desired motion.

The bipedal robot used in this research is called LEO. LEO consists of a body and two legs, each leg containing three independent Dynamixel RX-28 motors [4]. The motors represent the hips, knees and ankles that humans have, as can be seen in Figure 1. All commands and scripts are executed on an external device, while the Dynamixel servos contain their own internal controller. Further information about LEO's wiring and control system can be found in Appendix A.

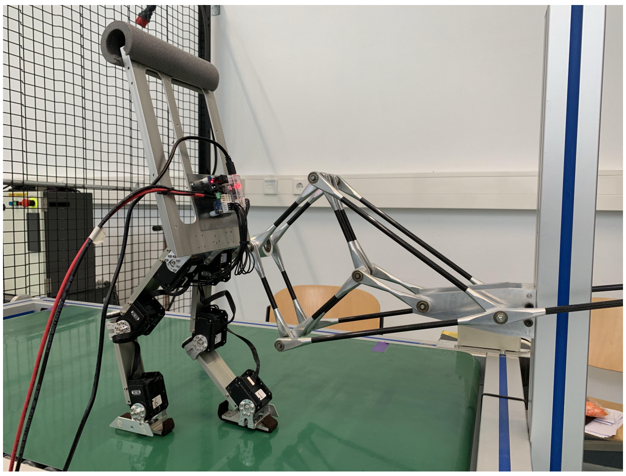

Fig. 1. Bipedal walking robot LEO. Visible are the six Dynamixel RX-28, the U2D2 power hub, the treadmill, and the supporting multibar linkage.

II. Methods

The method used to teach LEO how to walk is divided into two phases: kinesthetic teaching and human corrective feedback. Section A about kinesthetic teaching discusses the recording of the demonstrated motion, and the processing of the recorded data. Section B about human corrective feedback discusses the underlying algorithm used to adjust the robot's trajectory, and the interface used to give feedback.

A. Kinesthetic teaching

Kinesthetic teaching is used to teach LEO its most basic gait. During this first phase, a human teacher physically guides LEO's legs to demonstrate a walking motion, which is recorded by the robot in the form of angular servo positions. The recorded steps are then processed in order to create a smooth, symmetric gait.

1) Recording the demonstrated motion: The first requirement for successful recording, is stable and fast communication between the receiving device and the robot. Because the continuous motion shown by the user needs to be translated into a discrete dataset, sampling frequency plays an important role. When the communication speed is inadequate, the sampling frequency may be limited by the maximum rate of information exchange between the device and the robot. In this research, the communication speed was not high enough, causing the sampling frequency to reach its limit at $f = 8$ Hz.

To simulate a forward walking motion, LEO has been placed on a treadmill. Due to technical limitations, however, the treadmill is only able to run at one consistent speed. On top of the treadmill, a multibar linkage has been mounted. LEO can be attached to the multibar linkage, which serves two purposes with regard to LEO's stability. Firstly, it restricts LEO's frame to only move vertically or rotate around its point of attachment. This prevents LEO from falling sideways and provides additional support while walking. Secondly, the linkage is connected to a counterweight. The weight counterbalances LEO's own weight, making it easier for him to carry his own weight.

One recording is recommended to contain approximately 5 to 15 steps. A lower number of steps would decrease the reliability of the resulting motion, and a higher number of steps has no effect on the quality of the result. Once the motion is demonstrated, the collected data is stored in a matrix, where each row represents a sample at a certain time point and each column represents all recorded angles of one of the six servos. This data format is maintained throughout the entire method. To play back the recorded motion, rows of angle values are sent to the six motors at the same frequency as they are recorded. The angle values these rows contain, serve as setpoints for the motors to move towards.

2) Processing the recorded data: In order to translate the recorded motion into a motion that is as stable as possible, and able to be looped, processing is needed. In this method, the processing of the recorded data is threefold.

First, the steps in the recording need to be extracted and averaged. To extract individual steps from the recording, an average period of the steps has to be determined. Using the Fourier transform, the step frequency within the trajectories of each motor is found in the frequency domain. With this frequency, the recording is divided into individual steps. All extracted steps are then averaged for each motor, resulting in six single steps.

Once the single steps are extracted and averaged, they need to be synchronized in order to compare the left and the right steps. Because of the misalignment of motors between LEO's legs, angle values between the left and right joints are offset. Compensating for this offset allows for comparison, and thus for averaging between both legs. The result of this process is a single, unified step, produced by averaging the trajectories of the left and right step. This step can be projected onto the left and right leg, producing a completely symmetric gait.

Lastly, the motion of the universal step will be looped. The starting point of this motion will often not have the same angle value as its end point. This effect causes the motion to have a noticeable twitch at the looping point of its trajectory. To smoothen the motion at this looping point, an algorithm has been written that pulls the starting position, the end position, and a small number of values around those positions, closer together.

These three processes conclude the kinesthetic teaching. The result is a gait that closely resembles the originally shown motion, but serves as a better foundation for the feedback phase than the originally demonstrated motion.

B. Human corrective feedback

The feedback phase serves to improve the gait that has been taught during the kinesthetic teaching. To improve the effectiveness and to increase the intuitiveness of the feedback phase, two interfaces have been created. Experiments have been performed with these feedback interfaces, searching for the most intuitive implementation.

1) The underlying feedback algorithm: Although each interface is different, both are built upon the same feedback algorithm. The main premise of the feedback algorithm is the translation of one instance of feedback into a gradual adjustment of the original trajectory. This has been realized by creating a function that takes a timepoint and an adjustment value, and returns a spread of values. Two types of spreads are being considered: a normal distribution, and a parabolic distribution. When feedback is given by the user, the algorithm produces its feedback values. The spread of values is then added to or subtracted from the original trajectory, depending on the direction of the feedback, creating a new trajectory for each joint. The spread of values still possesses the same important properties as the original instance of feedback; its peak height is equal to the adjustment value, and its center is aligned with the time at which the feedback was given. The width of the spread is chosen to be constant, but adjustable.

2) Keyboard feedback: The first interface functions with the use of keyboard inputs. As LEO performs the gait that he has learned from the kinesthetic teaching, the user is able to provide him with live feedback. This means that the improved gait trajectory is being calculated based on the user's feedback, while the original gait is still being performed. The user provides feedback with two parameters: which joint to adjust, and which direction to adjust the joint towards. The value with which the joint trajectory is being adjusted, is set to be a small number, so larger adjustments can be realized by multiple instances of feedback. To give feedback, the user first selects the joint to give feedback to, by pressing either Q, W, A, S, Z or X. Each key represents a single joint in LEO's leg. The P-key will deselect a joint and allow the user to switch the focus of the feedback to another joint, but until the P-key is pressed, the selected joint will stay selected. The user is now able to adjust the trajectory of this single joint as many times as they want. By pressing the O-key, the servo of that joint will adjust its motion forwards by adding the spread of values to its trajectory at the time at which the key is pressed. By pressing the L-key, the trajectory is adjusted by the subtraction of the same spread of values. This interface requires the user to be precise, both in terms of timing, as well as in terms of realizing what feedback the robot needs.

3) Kinesthetic feedback: The second interface is an implementation that relies more on human hands-on interference with the robot's gait, which is why this implementation will now be referred to as kinesthetic feedback. With this interface, the user is able to pause LEO's walking motion at the position they would like to adjust, by pressing the spacebar. Once paused, the user selects either a single joint or one entire leg, depending on which they prefer to reposition. The single joints are selected using the same keys as the previous interface, and the entire left or right leg are selected using keys '1' or '2', respectively. The torques on the servo(s) in that selected joint or leg will then be disabled, allowing the user to freely move them around towards the desired position. When satisfied with the adjustment, the spacebar is to be pressed again, to enable the torque(s). LEO will not start moving until the spacebar is pressed once more, to allow for the user to safely remove their hands from the servos. Unlike the previous interface, the value with which the joint trajectory is being adjusted is not constant; instead, it is defined as the angular difference between the original servo position and the adjusted servo position. By defining the value this way, the modified trajectory will smoothly move through the demonstrated, desired position. The interface therefore allows for larger adjustments per instance of feedback, instead of the incremental approach that the first interface takes.

C. Experiments

To gauge the intuitiveness and effectiveness of the two previously mentioned interfaces, a small-scale experiment has been set up. Six participants are invited and divided into two groups. All participants will perform the method in its entirety twice; once using keyboard feedback, and once using kinesthetic feedback. To eliminate bias, one group will use keyboard feedback on their first execution of the method, and the other group will first use kinesthetic feedback. After experiencing both interfaces, they are asked which interface they think is more intuitive, which interface they prefer to use, and which produced gait they are more content with. For further information about the experiment, see Appendix B.

III. Results

For the sake of continuity, the results are divided into the same sections as the method. In section A about kinesthetic teaching, the results of recording and processing the motion will be discussed, supported by the opinions of the participants. In section B about human corrective feedback, the effects of the implemented algorithm, and the results of each implemented interface will be presented, again supported by the opinions of the participants.

A. Kinesthetic teaching

1) Recording the demonstrated motion: Playing back the demonstrated motion often did not result in a successful gait. The quality of the motion, produced by the kinesthetic teaching, largely depends on the skill of the demonstrator. Demonstrating the desired motion proved to be a harder task than expected, especially when done for the first time. However, after a maximum of 4 minutes of practice, most demonstrators were able to mimic a reasonable gait using LEO's legs. A comparison can be found in Figure 2.

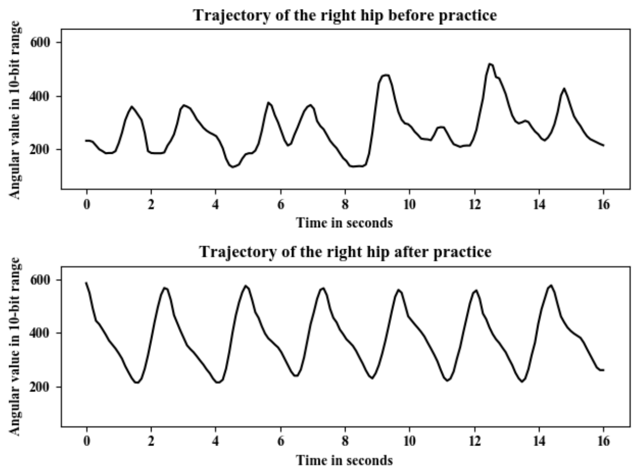

Fig. 2. The first graph shows the trajectory of the right hip during the demonstration of a participant, the second graph shows the trajectory of the right hip during the demonstration of the same participant, but after 4 minutes of practice.

2) Processing the recorded data: The fact that participants had four minutes to practice the demonstration of the motion proved to be influential; most participants did not need more than the given four minutes to teach themselves satisfactory demonstrative skills. The quality of the demonstration is important, because a clear and consistent step frequency has to be present. Without a consistent frequency, it is hard to average multiple steps into a single, unified step. However, even when the demonstrated motion was adequate, the result was not always a successful gait. LEO tended to tip forward or backward when the motion was processed and replayed, regardless of the quality of the demonstration. Even though the isolated motion that LEO performed seemed correct, the synchronization between his motion and the treadmill seemed to have been lost in the processing of the demonstration – he walked either too fast, or too slow.

Regardless of the desynchronization, however, the main functions within the processing phase worked as intended. In Figure 3, a comparison can be found between the three stages within the processing of the demonstrated gait.

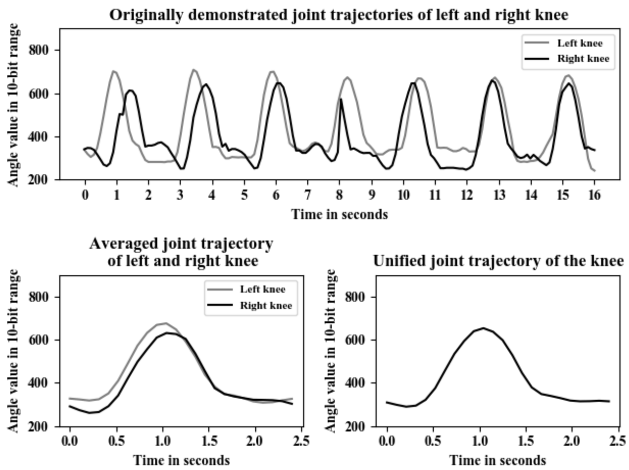

Fig. 3. A step-by-step comparison of the effects of each process that is executed onto the demonstrated motion. The top graph shows the original trajectories of the left and right knee, the bottom left graph shows the average step for each knee, and the bottom right graph shows the unified step.

B. Human corrective feedback

1) The underlying feedback algorithm: When the human teacher provided LEO with feedback, the feedback-induced adjustment closely resembled the intended adjustment, regardless of the effectiveness of the intended adjustment. The specific shape of the spread of values that are to be added or subtracted to the original trajectory proved to be influential. When first testing with a normally distributed spread and a parabolic spread, it was found that the boundaries of the parabolic spread, where the feedback values transitioned into the original trajectory, were too harsh. Because the values of a normal distribution approach zero at positive and negative infinity, this distribution functioned much better in this implementation. In Figure 4, a side-to-side comparison has been shown, to indicate the difference between the type of value spreads.

The width of the spread of values proved to be crucial as well. When the spread was chosen to be too narrow, the feedback would be too abrupt, causing the motion to twitch when the trajectory reached the point where feedback was given. When the spread was chosen to be too wide, the feedback would be too general, influencing almost every position in its trajectory instead of the ones that needed the feedback. This parameter has shown to be constant; regardless of the interface used, the optimal width of the

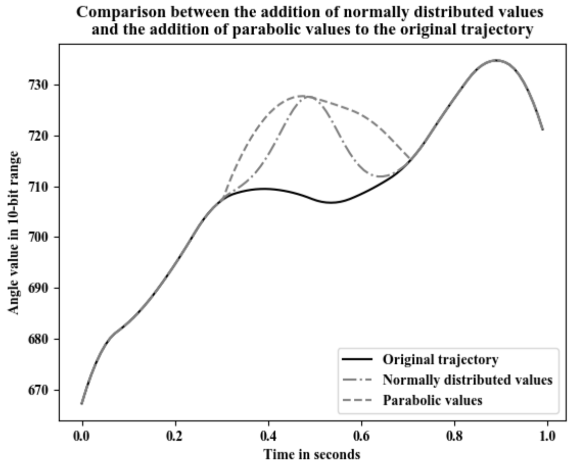

Fig. 4. A comparison between two considered spreads of values; a normally distributed spread, and a parabolic spread. The width of the spread in this figure is not representative of the actual width, but chosen to be smaller to clearly illustrate the difference between the two.

feedback amounted to about 0.4 times the total number of values in a single step. This width resulted in the smoothest motion, and the most noticeable improvement.

2) Keyboard feedback: The first interface, where the user provided LEO with feedback using keyboard inputs, initially proved to be difficult to operate. Because feedback was given while LEO was walking, participants initially found it difficult to give feedback at the exact position they intended to adjust. This observation, combined with the results of the post-experiment questionnaire, indicates that this interface is not intuitive. With time, however, participants learned to operate the interface more easily, which allowed them to provide the robot with more precise feedback. Eventually, participants were able to create stable gait that could walk for at least 10 steps without falling forward or backward.

One problem users encountered while using this interface was that they were able to provide one timepoint of LEO's motion with only one instance of feedback per step. Larger adjustments proved to be a challenge, because the value with which the trajectory is adjusted per instance of feedback is constant and predetermined. If they wanted to adjust a specific instance of the step by a larger amount, they would have to time their multiple instances of feedback consistently over multiple steps. Furthermore, users expressed that the angle adjustments by means of up-and-down feedback occasionally felt counter-intuitive, since the direction in which a joint needs to be moved depends on the orientation of the legs.

3) Kinesthetic feedback: The second interface, coined kinesthetic feedback, quickly proved to be more intuitive for users, as the answers to the post-experiment questionnaire show. The questionnaire showed that users also preferred to use the kinesthetic feedback in general. Remarkably, regardless of their experience with the specific interface, all participants preferred the gait produced with kinesthetic feedback over the gait produced with keyboard feedback. Participants also provided more instances of feedback when using the kinesthetic interface.

However, participants had difficulties applying small, precise adjustments using this interface. When participants selected a joint or leg to adjust, and the torques of the servos were disabled, much of the initial position was instantly lost. Occasionally, participants did not remember how the joint(s) were originally positioned, causing their feedback to have no frame of reference.

IV. Discussion

In this section, the performance of the method will be discussed, as well as recommendations for further research.

A. Performance

The results show that while the kinesthetic teaching did not perform perfectly, it fulfilled its role well. Because the purpose of kinesthetic teaching is only to provide a base for the human teacher to give feedback to, nothing more than an imperfect walking motion should be expected. The most prevalent issue within the process was the desynchronization of speed between the walking motion and the treadmill. The issue most likely roots in the processing of the recorded motion. Although inconvenient, the issue is not severe enough to cause any further problems throughout the method.

Users found that the kinesthetic interface was most useful for applying large adjustments to LEO's gait. On the other hand, users found that the keyboard interface was most useful for applying small adjustments. These findings indicate that both interfaces are viable, but each with their own main purpose. In general, however, most users found the kinesthetic interface to be more intuitive, and expressed that their preference of use lies with the kinesthetic feedback.

An important problem in this method lies in the inadequate communication speed between the controlling device and LEO. Because previous research with LEO has not encountered this problem while working in a Matlab-environment, and this research has been mainly performed in Python, it is suspected that the Python library provided by Dynamixel is the origin of the inadequate communication speed.

B. Recommendations for future research

1) Applications in other fields: In this research, the developed method is exclusively used to teach a bipedal robot how to walk. The method, however, has potential for purposes other than teaching a gait. Introducing a way of teaching robots a simple and repetitive motion, not by machine learning, but by demonstration, creates an opportunity for easier introduction of robots into many fields. For example, farmers will be able to use a single robot model to perform multiple tasks, without the need of engineering knowledge. Another application for this method is assembly in factories [5]. In assembly lines, automation will be significantly less expensive due to the non necessity of custom programming by engineers. Instead, workers themselves will be able to teach the robot how to move.

In other fields of use, however, the method described in this paper would not suffice, due to its limited versatility. The taught motion is inflexible and rigid, whereas most realistic scenarios require some form of adaptability. Because adaptability can be obtained with the use of machine learning, integration between the two methods could produce better results [6]. One idea would be to integrate aspects of reinforcement learning into kinesthetic feedback. For instance, multiple types of gait (longer steps, shorter steps, lifting the feet higher) could be taught to the robot using kinesthetic feedback. The robot would then undergo reinforcement learning to decide which gait to perform under which conditions. This way, simplicity would be conserved while improving versatility.

2) Improvements to the method: One possibility for improvement would be the implementation of a controller that responds to variable torques in the joints. This would be possible, because the servos in LEO's legs are equipped to read currents, which are proportional to its torque values. For example, the controls could be softened if the servos sense a higher torque value when putting a foot on the ground, and vice versa. This would result in lower impacts and softer movements, which subsequently improves performance on rough terrain. Using the current controller implementation, however, this method would not be feasible, because the communication between the computer and the motors is simply too slow. The aforementioned improvement requires the motors to send their torque values to the computer while also sending goal positions, which would be too much information for the amount of time they have to communicate. This problem has also caused haptic interfaces to lose their potential. A haptic interface, where the user would provide LEO with feedback by tapping his legs, would run into the same communication problems regarding torques. More information on this implementation can be found in Appendix C.

Another possibility for improvement is to combine the two designed interfaces. Because the kinesthetic interface is most useful for large and obvious adjustments, and the keyboard interface excels when incrementally adjusting, the combination of the two could have potential when used sequentially. This idea is more extensively discussed in Appendix D.

Lastly, the implementation of an undo-function in the feedback interface is recommended. When users applied feedback that did not improve LEO's gait, it was frustrating for them not to be able to undo their feedback. Their options were either to counteract the feedback by giving more feedback, or to start the process over; neither of which are desirable options. This improvement, along with others, has been realized after completion of the research. More information on these improvements can be found in Appendix E.

V. Conclusion

During the research this paper describes, the viability of a kinesthetic approach to gait learning has been evaluated.

Kinesthetic teaching has been applied to teach the robot a basic walking motion, and two interfaces for human corrective feedback have been designed to further adjust and improve its gait. Both interfaces have proven to be effective in different situations. Keyboard feedback is recommended when applying feedback incrementally, and when small adjustments are more likely needed than large adjustments. Due to its intuitiveness, kinesthetic feedback is recommended when the method is to be used by inexperienced users. It is also preferred when the necessary adjustments are more likely to be large than small.

The results of the research have indicated that the designed method of human guided gait learning is viable, because the participants were able to intuitively teach a robot how to walk without the need for extensive reinforcement learning. The main disadvantage of the method, however, is the lack of adaptability in the resulting motion. It can therefore be concluded that the method is effective when applied to repetitive and consistent motions.

ACKNOWLEDGEMENTS

The authors would like thank prof. dr. Robert Babuska, prof. dr. ir. Heike Vallery and dr. Carlos Celemin Paez, for their fruitful discussions and helpful feedback. The authors would also like to thank the Department of Cognitive Robotics at the TU Delft for allowing LEO and his treadmill to be available for research.

REFERENCES

[1] I. Koryakovskiy, H. Vallery, R. Babuška, and W. Caarls, "Evaluation of physical damage associated with action selection strategies in reinforcement learning," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 6928–6933, 2017.

[2] M. Hersch, F. Guenter, S. Calinon, and A. Billard, "Dynamical system modulation for robot learning via kinesthetic demonstrations," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1463–1467, 2008.

[3] S. Chernova and A. L. Thomaz, "Robot learning from human teachers," *Synthesis Lectures on Artificial Intelligence and Machine Learning*, vol. 8, no. 3, pp. 53–64, 2014.

[4] E. Schuitema, M. Wisse, T. Ramakers, and P. Jonker, "The design of leo: a 2d bipedal walking robot for online autonomous reinforcement learning," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 3238–3243.

[5] A. Muxfeldt, J.-H. Kluth, and D. Kubus, "Kinesthetic teaching in assembly operations - a user study," vol. 8810, 10 2014.

[6] M. E. Taylor, H. B. Suay, and S. Chernova, "Integrating reinforcement learning with human demonstrations of varying ability," in *The 10th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*. International Foundation for Autonomous Agents and Multiagent Systems, 2011, pp. 617–624.

APPENDIX

A. LEO's wiring

The Dynamixel motors on LEO are connected via 4-pin RS-485 cables. Therefore, a dongle is needed to be able to connect the motors to a laptop. To realize this, the U2D2-dongle from Dynamixel is used to convert the serial signal from the RS-485 cables to a USB-signal. The dongle is placed on the U2D2 power hub, which splits the RS-485 cable from the U2D2-dongle into two cables, one for the right leg and one for the left leg. Furthermore, the power hub is powered by an external power supply, allowing it to simultaneously provide power to the Dynamixel motors. An overview of this setup and the corresponding wires can be seen in Figure 5.

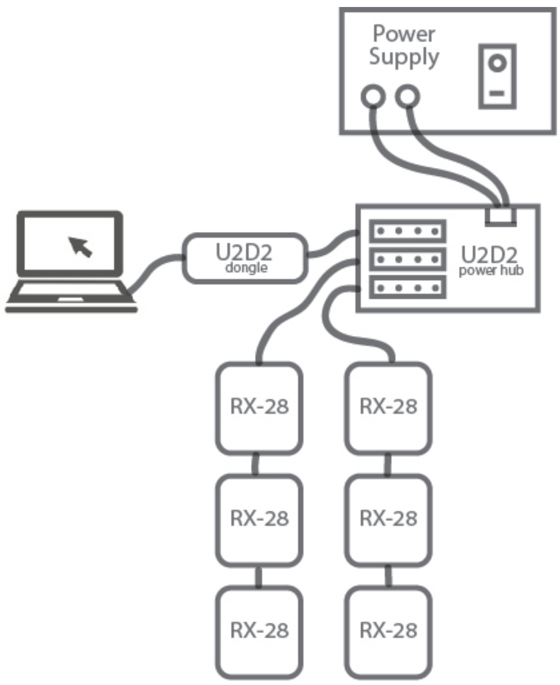

Fig. 5. Wiring diagram of LEO.

B. Experiment

During the experiment, keyboard and kinesthetic feedback interfaces have been tested. The goal of the experiment was to discover which interface is more preferable and more intuitive for users.

1) The setup: The participants started the experiment by practicing the kinesthetic teaching. All participants were given a maximum of 4 minutes to practice. The next part consisted of two full cycles of kinesthetic gait teaching. To prevent bias, they redid their demonstration for each cycle. If participants were to provide feedback to the same initial demonstrated motion twice, once for each interface, the subject could already know what feedback the motion needs during the second iteration.

One cycle of kinesthetic gait teaching consisted of two steps. The first step was kinesthetic teaching, where the participant showed LEO the desired walking motion. Second, the participant gave feedback to the demonstrated gait.

The participants had been divided into two groups. To eliminate bias, one group used keyboard feedback on their first execution of the method, and the other group first used kinesthetic feedback.

During the feedback phase, an observer kept track of the given feedback. Every time a hip, knee or ankle joint was selected, this was noted. The only interest was the feedback frequency – the size of the feedback was not taken into account. Because feedback size can vary when using the kinesthetic feedback, but can not vary when using the keyboard feedback, multiple instances of feedback on the same joint during the keyboard feedback was counted as a single instance. This way, both interfaces could be fairly compared.

For example, if during keyboard feedback, the right knee received three instances of feedback, only one instance would be noted. Once another joint was selected, a second instance of feedback would be noted. If the right knee was once again selected after this second instance of feedback, a third instance of feedback would be noted. Similarly, when the kinesthetic interface was used, feedback on an entire leg would be counted as three instances of feedback.

After performing the two full cycles, the participants answered the following questions:

- Which feedback interface, keyboard or kinesthetic, did you find more intuitive?
- Which feedback interface, keyboard or kinesthetic, do you prefer to use?
- Which produced gait has your preference?

2) The results: In this research, six participants were able to participate in the experiment. The results are shown in Figure 6 and Figure 7

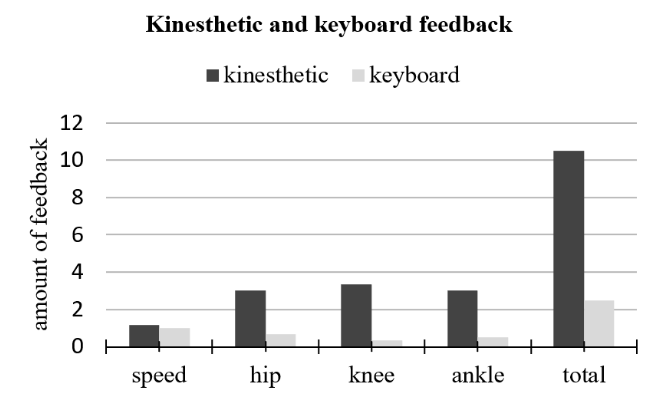

Fig. 6. Average amount of feedback on LEO needed to produce a robust gait.

C. Haptic feedback

Interesting to note is the role that haptic interfaces played in the development of the method. When the first iteration of feedback had been realized using keyboard inputs, the first

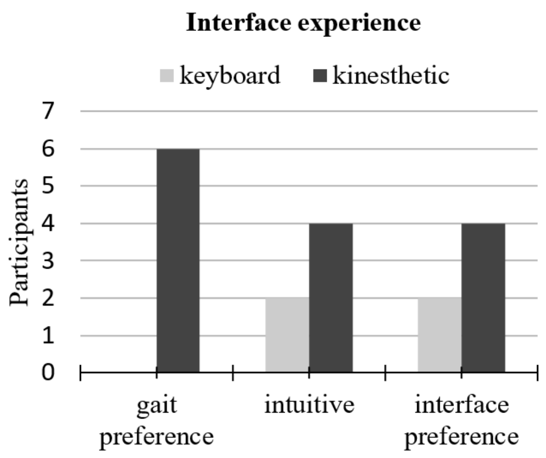

Fig. 7. Participant interface experience, keyboard feedback versus kinesthetic feedback.

idea for improvement was to implement a haptic element into the feedback phase. The user would have to tap one of the joints in the legs while it was walking, and the load spike in that specific joint would be interpreted as feedback, triggering the same underlying algorithm as the other interfaces use. This method, however, proved to be plagued with complications. Firstly, the idea of hands and fingers near powerful motors created a safety hazard that was hard to solve. Secondly, load spikes would be hard to detect, because the placement of a foot or the locking of a joint would be similar to a correcting tap on a motor. Lastly, while the interface would definitely be more intuitive, the ease of use would barely improve. Having to tap a motor at exactly the right time, while it is executing its motion, would undermine its intuitiveness.

D. Combined Interface

Because the kinesthetic feedback and keyboard feedback interfaces complement each other well, they have been combined into a single interface. Using this interface, users can, at any point, decide which way they want to give feedback; either kinesthetically or using keyboard inputs. Using this combination, users can apply large adjustments using kinesthetic feedback, and perfect the gait using incremental feedback with the keyboard interface. Or, they can decide which interface they find most pleasant to use, and use that interface for the entire feedback phase. In further research the user experience of this combined interface should be more elaborately tested.

E. Improvements

During the experiments participants expressed their needs for some improvements. These improvements concern the keyboard feedback and kinesthetic feedback and are listed below:

- Kinesthetic feedback
  - It is possible to subsequently adjust joints or legs in one instance of feedback
  - If the wrong joints are selected and disabled, one can now cancel this selection
  - A hotkey is added to undo the last instance of feedback
  - It is now possible to erase all given feedback and to restart the feedback phase on the original motion
- Keyboard feedback
  - The deselection of the joints using the p-key turned out to be confusing, so the user can now select another joint without needing to press the p-key
  - It is now possible to erase all given feedback and to restart the feedback phase on the original motion
  - Adjusting the speed can now be done using the same controls in both the kinesthetic feedback and the keyboard feedback, for better ease of use

Furthermore, a calibration mode has been added. Now, in case a motor would be replaced, the setup with Leo and its treadmill can easily be fixed by performing the calibration. Lastly, it is now possible to run the keyboard feedback and the kinesthetic feedback simultaneously, as previously discussed in Appendix D.