2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

Accurate power consumption estimation method makes walking robots energy efficient and quiet

Giorgio Valsecchi[1], Andrea Vicari[2,3], Fabian Tischhauser[1], Manolo Garabini[3], and Marco Hutter[1]

***Abstract*— Power consumption is a frequently overlooked aspect in robotics, especially in the context of legged robots. Nevertheless, improving the efficiency of walking robots is crucial to overcome the current limitations in runtime. This work proposes a novel method for precisely estimating actuator power consumption based on LSTM neural networks. The performance of this approach is benchmarked against currently employed models and validated on real hardware using certified instruments. The proposed method is integrated into the Isaac Gym framework and utilized to train a power-efficient policy. Instead of optimizing for handcrafted cost functions, such as the often used torque-square minimization, our approach for the first time trains RL policies that minimize the effective energy consumption. Hardware results demonstrate a reduction of approximately 25% in the robot's total power consumption, with a notable 50% decrease observed for the knee actuator. Additionally, the newly developed policy generates significantly smoother and quieter motions.**

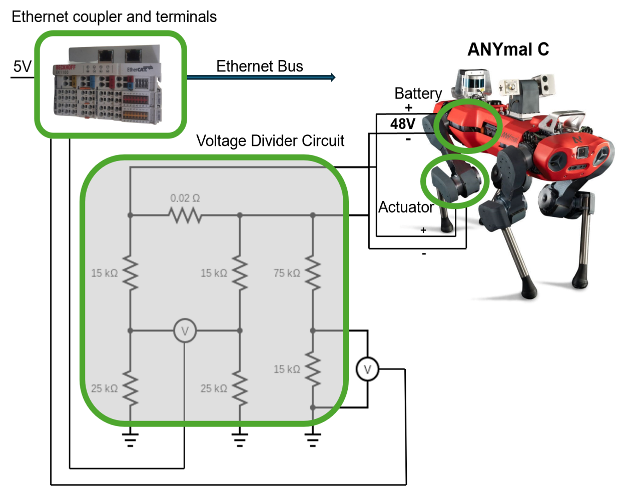

Fig. 1. Schematic representation of the power measurement setup.

I. INTRODUCTION

In recent years, the field of walking robots has undergone significant advancements. Originally conceived as prototypes designed to comprehend the fundamentals of locomotion, these machines have matured into fully developed products. Walking robots have found utility in various domains, including search and rescue operations, exploration, and automated surveillance. With the positive trend in their development, expectations regarding the performance of robots continue to rise.

Despite their versatility and growing capabilities, one of the significant challenges facing quadruped robots is their energy consumption. The autonomy of commercial quadrupeds ranges from 1.5 to 4 hours, a considerable distance from the requirements of most applications that demand extended shifts and continuous operation. As these robots gain autonomy and are deployed in more complex and prolonged missions, the demand for longer battery duration becomes increasingly critical and can only be met through improved efficiency.

This work was supported by ESA Contract Number 4000131516/20/NL/MH/ic and the European Union's Horizon 2020 research and innovation program under grant agreement No 101016970

[1] Author with the Robotic Systems Laboratory, ETH Zurich, Zurich 8092, Switzerland vgiorgio@ethz.ch
[2] Author with Scuola Superiore Sant'Anna, Italy
[3] Author with University of Pisa, Italy

The pursuit of energy-efficient quadruped robots is a necessity for sustainable and practical robotics deployment. Enhanced energy efficiency not only extends operational duration but also reduces the required battery size and capacity, thereby significantly lowering the overall energy costs associated with robotic operations. Furthermore, in remote or hazardous environments where human intervention is limited or impossible, the energy autonomy of quadruped robots becomes essential for mission success.

A. Related work

Despite not being a dominant focus, the topic of energy efficiency has been addressed by several works. More specifically, researchers have attempted in various ways to build energy-efficient robots. Designs aiming at transfering the advantages of passive walkers to actuated robots have been proposed [1], [2]. The MIT Cheetah project had a strong focus on energetic efficiency, particularly concerning the design of actuators [3]. StarlETH utilized springs in series with the gearbox (series elastic actuation) to store and reuse energy in the drivetrain [4], and the same approach is used by ANYmal [5]. The use of multiple actuators acting on the same joint (antagonistic actuation) has also been proposed as a possible way to increase efficiency [6], [7]. Several works have proposed using springs in parallel to the knee actuator to compensate for the torque induced by gravity [8]. Another recently explored approach involves the use of bio-inspired clutches and biarticular actuation (actuators coupling the motion of multiple joints) [9]. Optimized prototipe have been proposed as outcome of iterative processes leveraging simulations in the design phase [10]. The possibility of using actively variable transmission ratios has also been explored [11]. Another line of research attempted to solve the problem from a control point of view. Leveraging natually emerging passive limit cycles has been explored [12], as well as using virtual constraints and hybrid zero dynamics to leverage spring-loaded underactuated joints [13].

While all these works have achieved interesting results, they share two common shortcomings: they tend to introduce some level of mechanical complexity (with the exception of the MIT Cheetah), and they approach the efficiency challenge mostly from a design perspective, lacking a similar focus on the control side.

In terms of control, reinforcement learning is becoming the most utilized technique for quadrupedal legged robots, given its robustness and capability to handle nonlinearities and complex behavior. Platforms like Isaac Gym have democratized access to sophisticated simulation environments, enabling rapid and efficient training of control policies that perform with remarkable robustness [14]. Even in this context, the question of energy efficiency has been considered. Specifically, certain reward function terms are introduced for that purpose, typically penalizing the square of the joint torque, the joint mechanical power output, and acceleration [15], [16], [17]. These penalties are designed to discourage excessive use of force and rapid movements, supposedly contributing to more energy-conservative operations.

However, while these penalization methods are widely adopted for their intuitive explanation in reducing power consumption, they are based on strong simplifications of the full energy behavior of actuators. These simplistic penalty mechanisms, though seemingly logical, tend to overlook the deeper, more nuanced aspects of actuator behavior, which can be crucial for truly optimizing energy usage.

In the field of industrial manipulators the vast scale at which robots are deployed makes the topic of energy efficiency extremely relevant. In particular themes such as the selection of robots, motion planning, and task scheduling have been considered as potential ways to minimize energy consumption. Although also in this case most of the guidelines are qualitative in nature and not thoroughly validated, some works approached the problem in more data driven fashion, using neural networks to model the energy consumption of industrial arms [18], [19]

B. Contribution

This paper proposes a novel approach to minimize the energy consumption of robots by accurately characterizing actuators and employing model-free algorithms. The work focuses particularly on walking robots and utilizes ANYmal as a testing platform. Initially, we acquire a high-fidelity dataset that measures power consumption from robot actuators in real-world conditions. Subsequently, we use this dataset to train a neural network for estimating the power consumption of the robot actuators, comparing its performance against other metrics commonly used as proxies to penalize power consumption. The newly developed neural network is integrated into the IsaacGym framework and used to introduce a new penalty to the reward function, aiming to minimize energy consumption. Finally, we deploy the developed policy on ANYmal, measuring the actual power consumption for both individual joints and the entire robot. Additionally, we explore variations of the aforementioned reward function and benchmark the performance of the new policies against available baselines.

II. METHODS

A. Data acquisition

The adequate characterization of actuator power consumption behavior is at the core of our methodology. The initial step in characterizing a system involves measuring its properties and acquiring a relevant dataset. This task poses specific challenges, particularly for complex systems such as robotic actuators.

- The measurements must be sufficiently precise and accurate to minimize bias and noise in the experiments.
- Measurements of different quantities must be properly synchronized to capture the relationships between them accurately.
- The frequency at which data are acquired must be sufficiently high to capture the relevant processes of the actuators.
- The experiments must excite the actuators in modes that are representative of actual operation.

To meet these requirements, we constructed a custom measurement system setup (see Fig. 1), composed of a Beckhoff EtherCAT analog terminal (EL3102) and coupler (EK1100) interfacing with a voltage and current measurement circuit on the EtherCAT communication bus of ANYmal. This setup allows the acquisition of high-frequency measurements of a single actuator's power consumption while the robot is walking, synchronized with readings from other actuators. Additionally, we calibrated the setup against a certified power analyzer (Vitrek XT2640) to further enhance its accuracy and validate its performance.

We sampled actuator readings, including joint position and torque, at a rate of 400 Hz, and electrical parameters, voltage, and current, at 1000 Hz. The electrical readings were then downsampled to 400 Hz to ensure a coherent dataset.

We collected the dataset by measuring the individual power consumption of three different leg actuators: hip adduction/abduction (HAA), hip flexion/extension (HFE), and knee flexion/extension (KFE), for around ten minutes per actuator. The robot was manually commanded to perform a variety of motions using the standard controller Trekker, which sends position commands at the joint level, similar to other reinforcement learning controllers. Additionally, we manually disturbed the robot to further stress the actuators, aiming to collect a varied and representative dataset.

To enhance clarity and interpretability in the noisier datasets, a filtering process was necessary. We applied a Savitzky-Golay filter at 12.5 Hz. This filter is particularly adept at smoothing mid-range frequencies while preserving vital signal features like peaks and troughs [20].

Particular attention was given to how the dataset was utilized in the training process to enhance generalization while preventing the overlap of sequences. To distribute the dataset among training, validation, and test subsets, we initially segmented it into six parts, each of which was then divided in an 80%, 10%, 10% ratio. We then used the six subsets to perform cross-validation.

In addition to the training dataset, we conducted other experiments to test the robustness of the Power Prediction Network (PPN) performance. In these tests, we acquired data using different control policies and on another ANYmal C robot. Each experiment lasted between 2 or 3 minutes, employing various locomotion policies.

To avoid the introduction of bias, systematic errors, and model inaccuracies, we employ a model-free approach to estimate actuator power consumption. We selected Long Short-Term Memory (LSTM) networks as the ideal architecture for this task due to their ability to process sequences of data and contextual dependencies [21]. Our decision to use LSTMs is further supported by the fact that we work with relatively small datasets, typically spanning a maximum of 30 minutes of robot performance. LSTMs have demonstrated superior performance in such cases, making them a natural choice for our predictive model [22].

The primary goal of the PPN is to be integrated into the Isaac Gym framework, specifically for use in the reward function. For this reason, we utilize only labels available in the simulation environment and not all those available in the collected dataset. The features utilized by the network are joint position error, applied torque, and joint speed. Similarly, the frequency at which the network operates must be considered as high frequencies would slow down the simulation. In the Isaac Gym framework, physics are simulated at a frequency of 200 Hz, and the evaluation of the reward function occurs at 50 Hz, defining a minimum acceptable frequency for the PPN. Despite the possibility of adjusting these parameters, our observations indicated that higher frequencies did not enhance performance; therefore, we opted to maintain the default settings to facilitate a more effective comparison of future results. We chose to adopt a 200 Hz frequency for sampling both input and output sequences. We then take the average of the last four values estimated by the network to downsample it to 50 Hz.

In addition to the frequency of the PPN, the number of hidden states and the length of the time sequence fed into the network also have a significant impact on its performance and computational cost. We conduct several tests to evaluate these effects and select the network that allowed for the best trade-off.

We utilize the coefficient of determination, R^2, as a metric to assess how well the power consumption prediction method performs. R^2 is defined as $R^2 = 1 - \frac{\sum_i (y_i - f_i)^2}{\sum_i (y_i - \bar{y})^2} = 1 - \frac{SS_{res}}{SS_{tot}}$, were y_i, f_i, $\bar{y}$, SS_{res}, SS_{tot}, stand respectively for: the ground truth measurements, the value predicted by the model, the average ground truth value, the sum of squares of residuals, and the total sum of squares.

R^2 provides a measure of how well observed outcomes are replicated by the model, based on the proportion of the total variation of outcomes explained by the model. In case of a perfect prediction R^2 is equal to 1, for a model constantly predicting $\bar{y}$ is 0, and it can result in negative values for model that are less effective that predicting the average. We removed scale effects by normalizing the output of models that do not directly predict the power consumption so that their mean equals $\bar{y}$. This approach allows us to compare the performance of vastly different methods, such as those employed in previous studies [15], [16], [17].

We employed the Isaac Gym framework to train locomotion policies and assess the impact of the power estimation network. The framework's speed and efficiency allow for quick iterations. We used a baseline policy, trained with the standard reward function illustrated in [14], as a benchmark for the performance of other policies derived from the baseline.

Given our focus on the energetic aspects of walking gaits, we did not include a training curriculum with obstacles and stairs. Instead, we trained the policies to walk omnidirectionally on flat ground. Nonetheless, to enhance policy stability and robustness, we incorporated

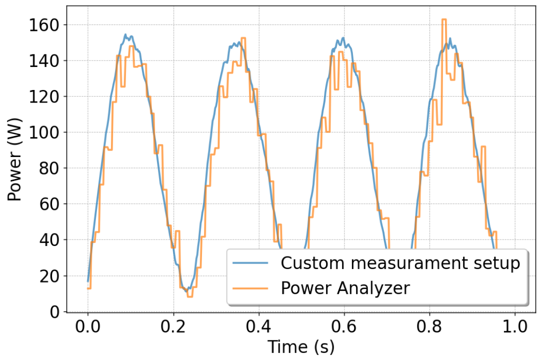

Fig. 2. Comparison between the total power consumption of ANYdrive, measured by our custom setup, and the certified power analyzer. In this experiment, the actuator is controlled in position without a load attached to it, tracking a sinusoidal reference signal with a frequency of 4 Hz and an amplitude of 1.5 rad.

rough terrain in the training. We applied uniform random noise ranging from -2 cm to $+2$ cm to the terrain, which included light slopes up to a 2% incline. The primary goal of this approach was to ensure the robot's stability when navigating environments that are both safe and nearly flat, effectively simulating the slight irregularities encountered in real-world conditions. We set up the training environment by deploying 4096 robots to concurrently learn for 2500 epochs.

We integrated the PPN into the Isaac Gym framework, creating one instance of the network for each actuator of every simulated robot, allowing us to estimate the individual actuator power consumption (and consequently the total for each robot). We then defined a new reward function term by multiplying the weigh $c_P = 4 \cdot 10^{-4}$ to the output of the PPN. We obtained the Power Minimization Policy (PMP) by adding the newly defined cost term to the baseline policy and removing the torque minimization term.

B. Testing

We evaluated the performance of our policies by deploying them on ANYmal C and conducting simple locomotion tasks, such as walking forward at a steady velocity and turning on the spot. We once again measured individual drive power consumption using the same setup as illustrated earlier. Additionally, we measured the entire robot's power consumption by replacing the battery with a calibrated power supply (PSB 9000 3U) and simultaneously logging the voltage and current at a frequency of 100 Hz.

III. RESULTS

A. Setup validation and dataset acquisition

We validated our measurement setup by comparing its readings with those of a certified power analyzer.

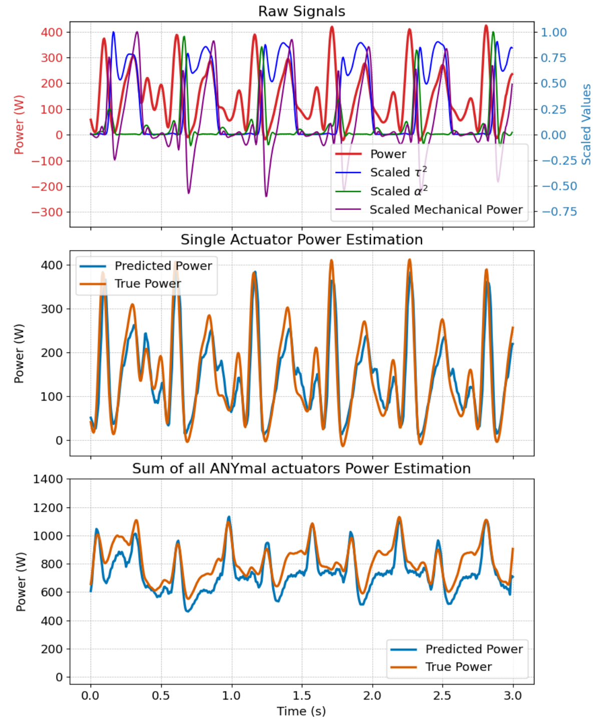

Fig. 3. Comparison between the ground truth power consumption and signals used to estimate it. In the top graph, the individual signals are scaled and compared with the power consumption of a single actuator. In the middle plot the PPN's power consumption estimation is compared with the measured ground truth, for a single actuator. In the bottom plot the estimation for the entire robot power consumption, obtained by summing the predictions for the twelve individual actuators, is compared to the one measured with a power supply.

Both measurement devices were connected to a single actuator which was controlled in various modes (position, velocity, torque), tracking reference signals such as sinusoids, square waves, and chirps. We manually aligned the readings and compared them, as shown in Fig. 2. The plot is representative of all the experiments we conducted, where the custom-made measurement setup performed very similarly to the certified one, albeit without capturing some of the oscillations occurring at higher frequencies. However, these oscillations have little impact on the overall power consumption.

B. Power Prediction Network

We trained and tested various neural networks with LSTM layers, differing only in the number of hidden states in the LSTM layers and for the length of the input sequence. We observed that the performance of the networks, measured by the coefficient of determination, reaches a plateau and does not improve with more complexity added to the network. For this reason we selected a network with 100 hidden states and an input sequence length of 0.3 s.

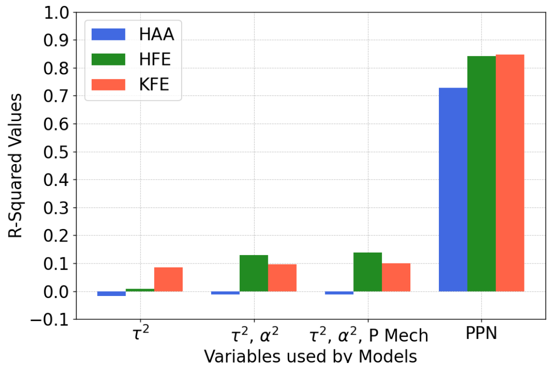

Fig. 4. Coefficient of determination comparison for different power consumption models. Different actuators are represented with various colors to highlight the differences.

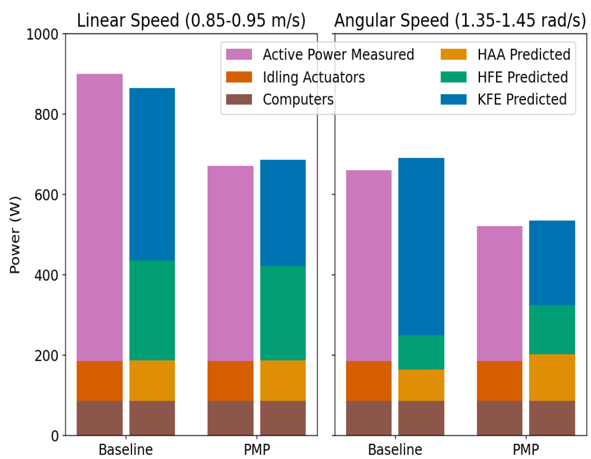

Fig. 5. Histogram showing the robot's total power consumption while performing the same locomotion task using different policies. Both the power consumption measured with the power supply and that estimated from the PPN are displayed. Additionally, the power consumption of the computation unit, as well as actuator idling power, is shown, along with the breakdown of the individual actuators' consumption.

C. Locomotion results

To provide a more qualitative understanding of the PPN performance, Fig. 3 compares the network's power estimation with the ground truth measured on the robot in an experiment not used for the training dataset, for both a single actuator and the entire robot. It can be seen that there is a strong coherence between the estimation and the ground truth, with peaks being captured accurately by the network. The raw signals used by other methods are also plotted (with a scaling factor to facilitate the understanding), to show the poor correlation with the power ground truth measurement.

Finally, we benchmark the performance of the PPN against that of other methods used for training policies. Specifically, we consider the methods employed to estimate power consumption and penalize in other relevant works, as mentioned in the introduction. Fig. 4 compares the R^2 values of the different models; it is clearly visible how the PPN exhibits significantly better performance. Two other observations can be made. The first is that it is very challenging to estimate power consumption using only the actuator torque, despite this being a common practice in the field. The second is that the HAA actuator's power consumption is the most challenging to predict. This can be explained by the average power consumption of the HAA actuator, which is significantly smaller, resulting in a lower signal-to-noise ratio. Consequently, all methods perform less well for HAA. Nevertheless, the PPN manages to achieve a reasonable performance even for HAA.

The introduction of the power minimization penalty did not affect the quality and stability of the training process, allowing policies to reach convergence. However, there is a five-fold increase in training time due to the additional computation of the network.

To assess the impact of the methodology, we simulated robots in Isaac Gym. The power estimation network was employed to evaluate the power consumption of each actuator for each policy. Standardized motions were used for comparison, specifically moving forward and rotating on the spot at a uniform velocity. Furthermore, we performed the same test with real hardware, trimming the data to use only sequences where the velocities were sufficiently constant to allow a meaningful interpretation.

Fig. 5 illustrates the power consumption for the two policies, both measured from the power supply and estimated from the PPN. Moreover, the power is broken down into multiple components, highlighting the impact of the computation units and the actuators' idling consumption. It is evident how the proposed methodology manages to reduce the overall power consumption by around 25%. It can also be observed that most of the power consumption reduction can be attributed to the KFE actuator, whose consumption is reduced by almost 50%. Interestingly, the consumption of the other actuators actually increases, this is particularly evident in turning motions. Nevertheless, a reduction in overall power consumption is achieved, with the benefit of more evenly loaded actuators.

We also experimented with higher weighs on the power minimization penalties, for both the baseline and the PMP. This approach shows only a marginal impact on the power consumption, indicating that using the wrong penalty cannot be correct by higher weights.

Fig. 6 presents time plots illustrating the dynamics of several relevant parameters throughout the gait cycle for both the baseline policy and the PMP. This provides a

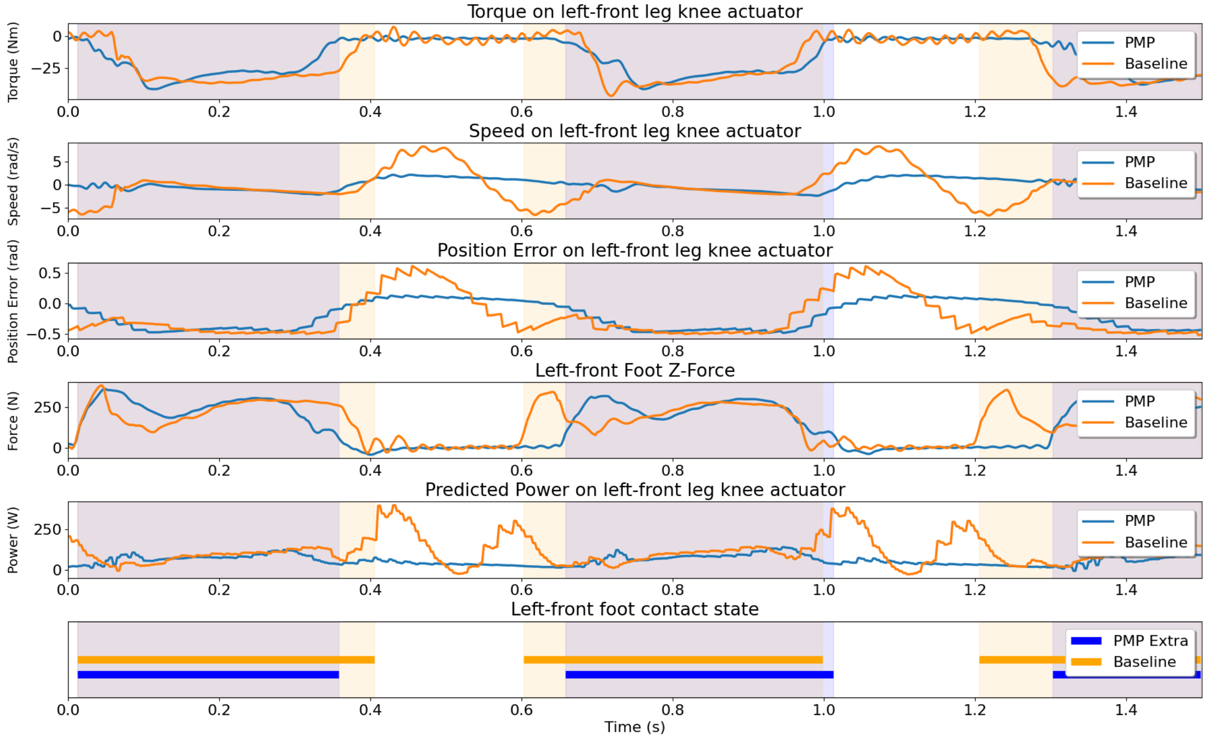

Fig. 6. Comparison between PMP and Baseline at the single actuator and leg level. The shaded areas represent the portion of the gait cycle during which the foot is in contact with the ground.

more detailed perspective on how the policies behave, with a particular focus on the KFE actuator, given its role as the primary source of power consumption and reduction. Several observations can be made:

- The PMP exhibits longer swing phases and lower gait cycle frequencies.
- The KFE torque in the PMP rises more slowly when the foot makes contact with the ground.
- The baseline KFE torque experiences high-frequency oscillations around zero during the swing phase.
- Although not visible in the plot, the PMP requires torques on average slightly higher that the baseline, contrary to the common belief that higher RMS torques imply higher power consumption.
- The KFE velocity is significantly lower for the PMP, with smaller and smoother changes over time.
- During the swing phase, the baseline policy shows two peaks of position error for the KFE.
- The estimated foot contact force dynamics are very similar for both policies, although the PMP has faster contact phases.
- For the baseline policy, the KFE has two peaks of power consumption, significantly higher than the average, which are completely absent for the PMP.

Overall, these observations align well with the visual impression given by the PMP policy, which results in smoother and less jerky robot motions. The foot impact with the ground is less harsh, making the robot less noisy. This is illustrated in Fig. 7, which shows the signal strength from a smartphone microphone located approximately 2 m away from the robot while it is controlled to turn on the spot at a constant velocity. The PMP consistently achieves lower levels than both Trekker and the baseline policy.

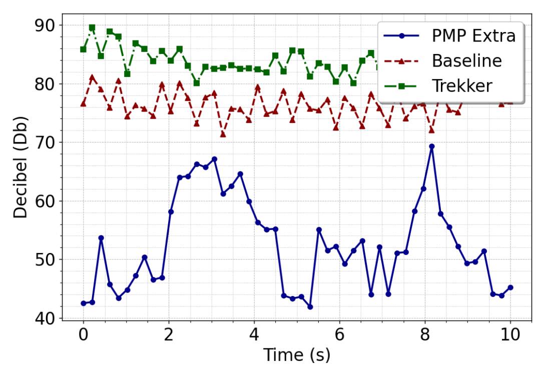

Fig. 7. Sound signal strength recorded with a smartphone's microphone from a distance of 2 m.

IV. CONCLUSION

In this work, we presented a novel approach for minimizing the power consumption of legged robots. We acquired a high-quality dataset by measuring the power consumption from the actuators of an ANYmal C during relevant tasks. Subsequently, we utilized this dataset to train an LSTM based neural network tasked with estimating instantaneous power consumption. We optimized the network to be computationally light without compromising its performance and benchmarked it against other power consumption methods used in similar applications, demonstrating its significantly better precision and accuracy. Further, we integrated the Power Prediction Network (PPN) into Isaac Gym and employed it to penalize power consumption during the training process. Finally, we measured the actual power consumption of the policies obtained with the new methodology and compared them against a baseline. We then analyzed the qualitative differences between the policies, highlighting factors contributing to power savings.

Despite the positive results, this contribution is preliminary in nature. In particular the increase in training time could prevent this technique from being used in a practical way. However, no time was spent yet to address this aspect. More accurate model selection or techniques like model compression, or even merging the PPN with the network used to model the actuator dynamics [15], could reduce the computational load introduced by the power estimation step. Two fundamental questions should be addressed in future works. The first one is: to which extent this technique affects the locomotion performance and how can this approach be improved, allowing operations at the efficiency-performance Pareto front? Moreover the approach should be fully integrated into more complex control architecture, merging the dynamically balancing efficiency and navigation over complex terrains. The second question is: to what extent can such an approach become more general? Tests with different types of actuators and robots should be performed to understand how applicable a similar approach could be in other applications, such as manipulators and exoskeletons.

REFERENCES

[1] S. H. Collins and A. Ruina, "A bipedal walking robot with efficient and human-like gait," in *Proceedings of the 2005 IEEE international conference on robotics and automation*. IEEE, 2005, pp. 1983–1988.

[2] P. A. Bhounsule *et al.*, "Design and control of ranger: an energy-efficient, dynamic walking robot," in *Adaptive Mobile Robotics*. World Scientific, 2012, pp. 441–448.

[3] S. Seok *et al.*, "Design principles for highly efficient quadrupeds and implementation on the mit cheetah robot," *2013 IEEE International Conference on Robotics and Automation*, pp. 3307–3312, 2013.

[4] M. Hutter *et al.*, "Starleth: A compliant quadrupedal robot for fast, efficient, and versatile locomotion," in *Adaptive mobile robotics*. World Scientific, 2012, pp. 483–490.

[5] ——, "Anymal-a highly mobile and dynamic quadrupedal robot," in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 38–44.

[6] N. G. Tsagarakis *et al.*, "An asymmetric compliant antagonistic joint design for high performance mobility," in *2013 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2013, pp. 5512–5517.

[7] W. Roozing *et al.*, "Development and control of a compliant asymmetric antagonistic actuator for energy efficient mobility," *IEEE/ASME transactions on mechatronics*, vol. 21, no. 2, pp. 1080–1091, 2015.

[8] H. Kolvenbach *et al.*, "Towards jumping locomotion for quadruped robots on the moon," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019, pp. 5459–5466.

[9] A. Badri-Spröwitz *et al.*, "Birdbot achieves energy-efficient gait with minimal control using avian-inspired leg clutching," *Science Robotics*, vol. 7, no. 64, p. eabg4055, 2022.

[10] J. Reher, E. A. Cousineau, A. Hereid, C. M. Hubicki, and A. D. Ames, "Realizing dynamic and efficient bipedal locomotion on the humanoid robot durus," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2016, pp. 1794–1801.

[11] G. Valsecchi *et al.*, "Actively variable transmission robotic leg," in *Climbing and Walking Robots Conference*. Springer, 2023, pp. 40–51.

[12] M. Spong *et al.*, "Passivity-based control of bipedal locomotion," *IEEE Robotics & Automation Magazine*, vol. 14, no. 2, pp. 30–40, 2007.

[13] K. Sreenath, H.-W. Park, I. Poulakakis, and J. W. Grizzle, "A compliant hybrid zero dynamics controller for stable, efficient and fast bipedal walking on mabel," *The International Journal of Robotics Research*, vol. 30, no. 9, pp. 1170–1193, 2011.

[14] N. Rudin *et al.*, "Learning to walk in minutes using massively parallel deep reinforcement learning," *PMLR 164:91-100*, 2022.

[15] J. Hwangbo *et al.*, "Learning agile and dynamic motor skills for legged robots," *Tech. Rep.*, 2019.

[16] J. Lee *et al.*, "Learning quadrupedal locomotion over challenging terrain," *Science Robotics*, vol. 5, no. 47, p. eabc5986, 2020. [Online]. Available: https://www.science.org/doi/abs/10.1126/scirobotics.abc5986

[17] T. Miki *et al.*, "Learning robust perceptive locomotion for quadrupedal robots in the wild," *Tech. Rep.*, 2022.

[18] J. Yan and M. Zhang, "A transfer-learning based energy consumption modeling method for industrial robots," *Journal of Cleaner Production*, vol. 325, p. 129299, 2021.

[19] M. Zhang and J. Yan, "A data-driven method for optimizing the energy consumption of industrial robots," *Journal of cleaner production*, vol. 285, p. 124862, 2021.

[20] S. T. WH Press, "Savitzky-golay smoothing filters," *Computers in Physics*, 1990.

[21] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, 1997.

[22] Trujillo-Guerrero *et al.*, "Accuracy comparison of cnn, lstm, and transformer for activity recognition using imu and visual markers," 2023.