Generating Realistic Sound with Prosthetic Hand: A Reinforcement Learning Approach

Taemoon Jeong[1], Sankalp Yamsani[2], Jooyoung Hong[2], Kyungseo Park[3], Joohyung Kim[2], Sungjoon Choi[1*]

Abstract— **In this study, we tackle the complex task of enabling prosthetic hands to accurately reproduce sounds, a crucial aspect for distinguishing between different materials through auditory feedback. Sound identification, such as discerning a drywall tap from that on a brick wall, significantly enhances the functionality and user experience of prosthetic devices. However, achieving this level of auditory feedback in prosthetic hands poses considerable challenges. We utilize reinforcement learning (RL) techniques to train prosthetic hands in emulating human-like sound characteristics, focusing on key auditory signals like amplitude and onset timing. Our approach integrates a detailed analysis of these sound attributes to direct the prosthetic hand's movements for the sound generation that mimics natural human actions. We developed a tailored reward function incorporating amplitude, onset strength, and timing criteria to ensure the prosthetic hand's movements align closely with the intended human-like sound output.**

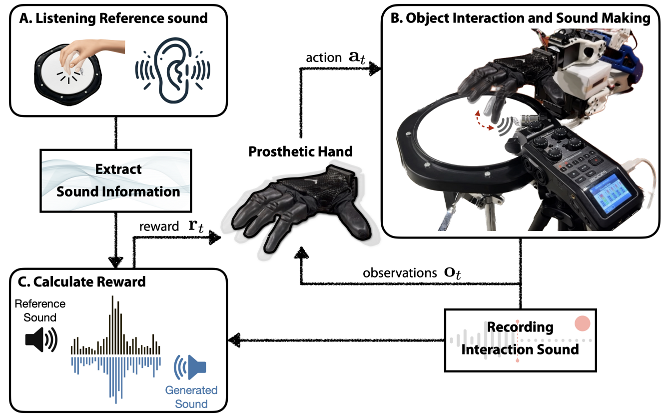

Fig. 1: RL framework overview. **A.** Listening to human-produced sound and extracting sound features such as amplitude and onset. **B.** Interaction with the object and sound generation according to the action of the prosthetic hand. **C.** Calculate the reward, including amplitude, onset, timing, and hit reward from the extracted sound information and recorded interactive sound.

I. Introduction

When inferring properties and characteristics about the object, non-visual cues often play a crucial role. For example, when we want to nail something heavy to the wall, we often determine the suitability of the wall material by tapping it with our fingers to listen to the resulting sound. A hollow sound might indicate drywall, while a denser, higher-pitched tone suggests plaster or brick. This seemingly simple act becomes surprisingly challenging for individuals using prosthetic hands. The absence of auditory feedback in prosthetic hands can limit the user's ability to interact with and understand their environment effectively. By incorporating realistic sound generation, prosthetic hands can provide users with valuable information about the objects they are touching, enhancing their overall experience and functionality.

While much of the research on prosthetic hands has focused on visual and tactile feedback [1]–[6], the use of electromyography (EMG) in controlling prosthetic hands has significantly improved responsiveness to muscle movements [7]–[9]. These approaches enable prosthetic hands to perform tasks effectively, such as creating finger gestures or grasping objects. However, these methodologies, despite their strong foundation, lack the crucial ability to generate sounds during object interaction, limiting performance in certain tasks. For instance, in situations where visual feedback is limited, such as in low-light conditions or when handling objects outside the user's field of view, auditory cues become increasingly important. By providing information about the material properties and the hand's interaction with the object, sound feedback can help users maintain a sense of control and awareness even when visual information is unavailable.

This paper explores methods for effective sound generation in prosthetic hands, aiming to bridge this gap and unlock their full potential. We focus on the problem of generating tapping sounds using a prosthetic hand that mimics human-produced sound. The tapping sound generation function can be integrated into the prosthetic hand control system, allowing users to activate it when needed, such as when exploring surface properties or interacting with objects in low-light conditions. Moreover, while prosthetic hand users can employ wrist or lower arm tapping to generate sounds, finger tapping offers a more natural and intuitive method for exploring surface properties, as it closely mimics the actions of a human hand. To this end, we utilize reinforcement learning (RL) to create effective tapping motions of the prosthetic hand that can generate realistic sounds. RL stands out for its ability to optimize decision-making processes in complex environments. This characteristic has enabled it to be applied extensively in fields like music and interactive audio applications [10]–[12].

[1]Taemoon Jeong and Sungjoon Choi are with the Department of Artificial Intelligence, Korea University, Seoul, Korea (email: taemoon-jeong, sungjoon-choi@korea.ac.kr)

[2]Sankalp Yamsani, Jooyoung Hong and Joohyung Kim are with the Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Champaign, IL, USA (email: yamsani2, jh97, joohyung@illinois.edu)

[3]Kyungseo Park is with the Department of Robotics and Mechatronics Engineering, Daegu Gyeongbuk Institute of Science and Technology (DGIST), Daegu, South Korea (e-mail: kspark@dgist.ac.kr).

[*]Corresponding author

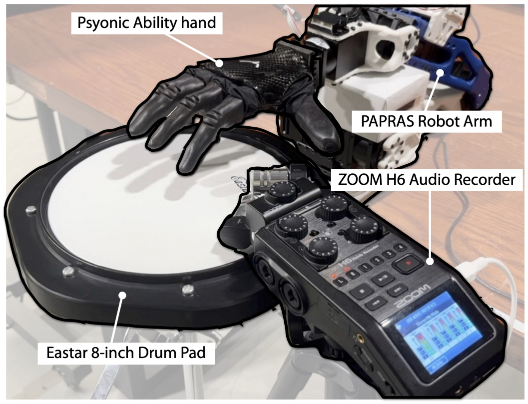

Fig. 2: The hardware setup includes PSYONIC Ability Hand, mounted on a 6-DOF PAPRAS robot arm. Sound recording is performed using a ZOOM H6 recorder. An 8-inch Eastar Drum Practice Pad is utilized as the tapping object.

Our system consists of three main steps. First, it records the human-produced sound and extracts features such as amplitude and onset. These features of the reference sound are used to calculate the reward. Second, the prosthetic hand moves its fingers based on the current policy and generates sound through interaction with the object. Third, the reward is calculated by inputting the reference sound and generated sound features, and the policy is updated. It is important to design a correct reward function in reinforcement learning. We design a reward function with four elements: amplitude, onset strength, onset timing, and hit so that the prosthetic hand can produce realistic sounds. Fig. 1 illustrates our system's framework.

Our contribution to this work can be summarized as follows:

- We formulate the task of effectively generating desired sounds through prosthetic control as reinforcement learning.
- We design a reward function to evaluate whether the sound is appropriately generated.

We would like to note that while we only focus on generating rather simple tapping sounds, this work is a preliminary study towards generating diverse interactive sounds using prosthetic hands. To the best of our knowledge, this is the first work to use reinforcement learning to train a policy for the prosthetic hands to make a sound that resembles the given ground truth sound.

II. Proposed Method

In this section, we present our method to generate tapping sounds, including single and double-beat sounds, by learning the motion of a prosthetic hand interacting with a drum pad, leveraging reinforcement learning (RL). Traditional RL approaches often rely on simulation environments due to their low sample efficiency. However, implementing a hand-drum pad interaction simulation to achieve realistic sound is extremely difficult because it involves multiple contact dynamics with high frequencies, resulting in a significant sim-to-real gap. To overcome this limitation, we propose a real-world learning approach. We train a policy by directly controlling the prosthetic hand and generating sound in a real environment using a high-fidelity audio recorder.

A. Formulation

We leverage a Markov decision process (MDP) framework [13] to model the sound generation process. This allows the prosthetic hand to learn and optimize its actions based on observations and rewards, progressively refining the sound quality. At each time step, t, the hand receives observations, $\mathbf{o}_t \in \mathbb{R}^6$, which consist of joint angle position, velocity, acceleration, and mean amplitude of sounds. Based on these observations, it selects an action, $a_t \in \mathbb{R}$, that adjusts its joint angle. The reward, r_t, is determined by four factors: the difference in mean amplitude from the reference sound, the similarity of sound onset sequences, the accuracy of sound onset timing, and the number of distinct sound events.

1) Observations: Our system generates tapping sounds using only the index finger among the fingers of a prosthetic hand. To guide the finger's movement, the system utilizes observations, consisting of the index finger's joint angle position at the current time step, q_t, joint angle position at the previous time step, q_{t-1}, joint angle velocity at the current time step, $\dot{q}_t$, joint angle velocity at the previous time step, $\dot{q}_{t-1}$, joint angle acceleration, $\ddot{q}_t$, and the mean amplitude of generated sound, $\tilde{m}_t$.

$$\mathbf{o}_t = \{q_{t-1}, q_t, \dot{q}_{t-1}, \dot{q}_t, \ddot{q}_t, \tilde{m}_t\}$$

The joint angle position, $q_t \in \mathbb{R}$, is limited to the range of [0.087, 1.047] radians (rad). $\dot{q}_t \in \mathbb{R}$ stands for joint angle velocity in radians per second (rad/s), and $\ddot{q}_t \in \mathbb{R}$ denotes joint angle acceleration in radians per second squared (rad/s^2). $\dot{q}_t$ and $\dot{q}_{t-1}$ are numerically calculated from q_t. $m_t \in \mathbb{R}$ is the mean amplitude of sound, typically measured in decibels (dB). This comprehensive set of observations, which includes both current and previous step values, forms the basis for our system's control and understanding of the tapping process.

2) Action: Given observations, the prosthetic hand learns a parameterized policy, π, to generate specific joint angle movements as actions. The policy, π, is a function that maps from observations to action, a target joint angle of an index finger. The action is bound within the limits of the prosthetic hand's joint, $a_t \in [0.087, 1.047]$.

3) Reward Structure: We aim to train the prosthetic hand to generate sounds closely matching a given reference audio. To achieve this, the agent receives a composite reward based on four critical aspects of the generated sound.

Amplitude Reward: The amplitude reward is based on the difference between the mean amplitude of the audio produced by the robot and a reference audio amplitude. The smaller the gap between the two amplitudes, the higher the reward. The amplitude reward is calculated as follows:

$$r_t^{\text{amp}} = e^{-|\tilde{m}_t - m_t|}, \quad 0 \le t \le T \tag{1}$$

where $\tilde{m}_t$ is the mean amplitude of the generated audio at time t, and m_t is the mean amplitude of the reference audio at time t. The exponential function accentuates the reward as the difference between the two amplitudes decreases, aligning the robot's audio output with the reference audio's amplitude. The amplitude reward encourages the robot to generate audio with a similar amplitude to the reference audio to have an overall sound quality similar to the reference audio.

Onset Strength Reward: The onset strength reward r_t^{onset} is derived from a dynamic time warping (DTW) [14] comparison between the onset strength sequence of the produced audio and the reference audio. The onset strength sequence captures the beats and strength of sounds. The closer the two sequences match, the higher the reward. The onset strength reward is as follows:

$$r_t^{\text{onset}} = \begin{cases} -DTW(\tilde{s}_{0:T}, s_{0:T}) & \text{if} \quad t = T \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

where $\tilde{s}_{0:T}$ is the onset strength sequence of the produced audio, and $s_{0:T}$ is the onset strength sequence of the reference audio. $DTW(\cdot,\cdot)$ is the normalized DTW distance between two sequences. The onset strength reward encourages the agent to generate sounds with attacks, as measured by the energy.

Onset Timing Reward: The onset timing reward r_t^{timing} is a measure of how accurately the agent matches the onset timing of the reference audio. The onset timing is the time at which the sound's energy suddenly increases, marking its beginning. While onset timing might seem like just a parameter for aligning signals in time, it plays a crucial role in maintaining a consistent rhythm in tapping. In real-world tapping, keeping a steady rhythm is important, and the relative time intervals between taps matter. The onset timing reward encourages the agent to learn this consistent rhythm and maintain the relative time intervals between taps. If the onset timings of the generated and reference audio are mismatched, it can lead to a perceived delay or unnatural rhythm in the generated tapping sound, which can degrade the quality and realism of the audio. Therefore, the closer the onset timings of the generated and reference audios, the higher the reward. The onset timing reward is as follows:

$$r_t^{\text{timing}} = \begin{cases} e^{-|\tilde{t}_{\text{onset}} - t_{\text{onset}}|} & \text{if} \quad t = T \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where $\tilde{t}_{\text{onset}}$ is the onset timing of the generated audio, and t_{onset} is the onset timing of the reference audio. The exponential function is used to accentuate the reward as the difference between the two onset timings decreases. This helps to ensure that the agent learns to generate sounds with onset timings that closely match the reference audio.

Hit Reward: The hit reward r_t^{hit} is a measure of whether the agent successfully strikes the drum pad. The hit reward is awarded when the maximum amplitude of the generated audio exceeds a specific threshold and when the number of beats detected in the generated audio corresponds to the number of beats in the reference audio. While a single tap may be sufficient in certain situations, learning to generate various tapping patterns, including multiple taps, is crucial for creating natural and realistic tapping sounds. Generating the correct number of taps is essential for maintaining the rhythmic pattern and natural feel of the tapping sound. If the agent generates too few or too many taps compared to the reference audio, it can result in an unnatural or inconsistent tapping pattern, which can degrade the realism and quality of the generated audio. By including the number of beats in the reward function, we ensure that the agent learns to generate the appropriate number of taps, whether it's a single tap or multiple taps, depending on the reference audio. The hit reward is as follows:

$$r_t^{\text{hit}} = \begin{cases} 1 & \text{if} \quad t = T \textit{ and } \tilde{h} = h \textit{ and } \hat{a}_t \geq \varepsilon \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

where $\tilde{h}$ is the number of beats detected in the generated audio. $\tilde{h}$, h is the number of beats in the generated and reference audio. $\hat{a}_t$ is the maximum amplitude of the generated audio at time t, and ε is the amplitude threshold. The hit reward is important for ensuring that the agent learns to generate loud sounds to strike the drum pad. It is also important to ensure that the agent learns to generate sounds with the correct number of beats.

Algorithm 1 Tapping Motion Generation

Initialize policy parameters θ, initial observations $\mathbf{o}_0$
Load reference sound
for $t = 0$ **to** T **do**
  Generate action a_t from policy $\pi_\theta(\mathbf{o}_t)$
  Execute action a_t on prosthetic hand
  Record generated sound for predefined duration
  Calculate rewards based on sound (Equation. 1-5)
  Observe new observations $\mathbf{o}_{t+1}$
  Store transition $(\mathbf{o}_t, a_t, r_t, \mathbf{o}_{t+1})$ in replay buffer
  if enough data in buffer **then**
    Calculate advantage estimate using GAE
    for $k = 1$ **to** number of epochs **do**
      Sample mini-batch from replay buffer
      Update π_θ using PPO clipped objective
    end for
  end if
end for
Return trained policy parameters θ

Total Reward: The total reward r_t is a weighted sum of the four individual reward components:

$$r_t = w_a r_t^{\text{amp}} + w_{os} r_t^{\text{onset}} + w_{ot} r_t^{\text{timing}} + w_h r_t^{\text{hit}} \tag{5}$$

where w_a, w_{os}, w_{ot}, and w_h are the weights for each reward component. The weights can be used to control the relative importance of each reward component. For example, if the agent has difficulty generating sounds with the correct number of beats, the weight for the hit reward can be increased. The total reward updates the agent's policy in the reinforcement learning algorithm.

B. Tapping Motion Generation

1) Architecture: We adopt the Proximal Policy Optimization (PPO) [15] reinforcement learning algorithm to train the prosthetic hand to generate sounds that closely match a given reference audio. PPO is a policy gradient algorithm that has been shown to be effective for a variety of tasks, including robotics. Our implementation of PPO uses an Actor-Critic architecture, which consists of two neural networks: an actor network and a critic network. The actor network is responsible for determining the action to take in a given state, while the critic network estimates the value of a state-action pair. Algorithm 1 represents our PPO-based sound generation procedure.

TABLE I: Training Settings

Parameter	Value
Random Seed	111
Optimizer	Adam
Learning Rate	3×10^{-4}
Number of Hidden Layers	2
Neurons per Hidden Layer	128
Action Velocity Boundaries	-8.0 to 8.0 rad/s
Joint Position Boundaries	0.087 to 1.047 radians
Audio Sampling Rate	44.1 kHz
Clipping Ratio	0.2
Discount Factor (γ)	0.99
GAE Parameter (λ)	0.95
Value Coefficient	0.5
Entropy Coefficient	0.01
Maximum Gradient	0.5

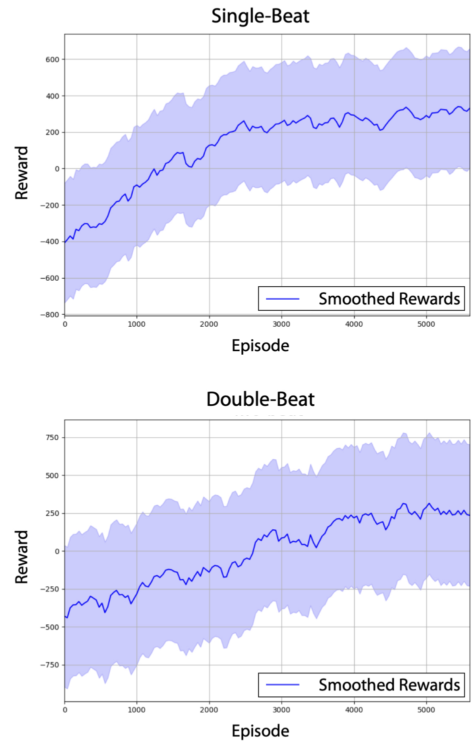

Fig. 3: The learning curve of a prosthetic hand to make tapping sounds, single-beat and double-beats. The blue line is the smoothed reward, and the light blue area indicates the variance.

2) Training Settings: We initialized our experiments with a random seed set to 111 to ensure optimal and reproducible results. Our training environment is built around the Proximal Policy Optimization (PPO) framework, driven by the Adam optimizer with a learning rate of 3×10^{-4}. We utilize feedforward neural networks as our actor and critic, characterized by two hidden layers, each consisting of 128 sizes. To guarantee safety and effectiveness in operations, we have defined the action velocity boundaries as a minimum of -8.0 and a maximum of 8.0 rad/s. This ensures that the prosthetic hand does not move too quickly or forcefully, which could lead to injury or damage. Correspondingly, the joint positions are bound between a minimum of 0.087 and a maximum of 1.047 radians, ensuring the feasibility and safety of movements. Understanding the intricacies of the audio environment, we sample our audio at a sample rate of 44.1kHz. This allows us to capture minute details and nuances from the agent's interactions, which is important for generating realistic and expressive sounds. Critical to PPO, our policy update is guided by a clipping ratio of 0.2. This ensures that the policy updates are not too large, which can lead to instability. The discount factor and GAE parameter are set to γ of 0.99 and λ of 0.95, respectively. These parameters control the importance of future rewards and the weighting of advantages, which are important for training the policy to generate long-term rewards. Regularization is addressed by setting a value and entropy coefficients to 0.5 and 0.01, respectively. This helps to prevent overfitting and encourage exploration. The maximum gradient is bounded at 0.5, ensuring stability in our optimization steps. Our training settings are summarized in TABLE I.

III. Experiments

In this section, we demonstrate the training of a prosthetic hand using a reward function defined by sound information and evaluate its ability to produce sounds similar to a reference sound. Emulating the interaction sounds between the hand and a drum pad is challenging, leading to a significant gap between simulation and reality. Therefore, we have implemented our system in a real-world environment rather than a simulated one. These experiments confirm that the prosthetic arm can produce a sound that matches the reference sound within the reinforcement learning system.

A. Experimental Setup

We used PSYONIC Ability Hand, a prosthetic hand with six degrees of freedom [16]. The hand was mounted on a 6-DOF PAPRAS robot arm [17], and one finger was controlled while the arm was held in a fixed pose. The sound was recorded using a ZOOM H6 recorder for 1-second intervals. Only mono sound information was used, despite the device's stereo capabilities. An 8-inch Eastar drum practice pad served as the tapping object, struck directly by the Ability hand to produce sound. The prosthetic hand performs

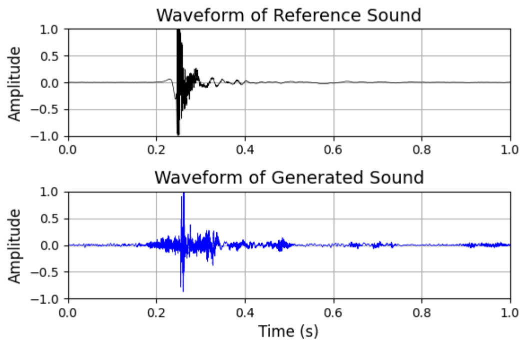

(a) Single-beat

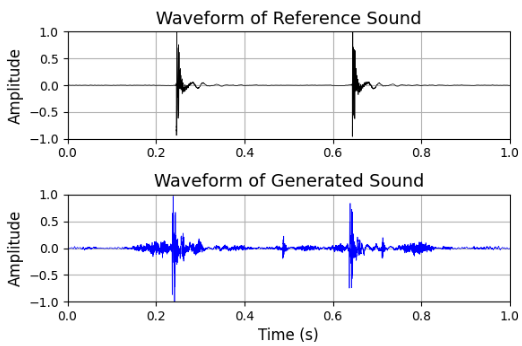

(b) Double-beat

Fig. 4: The waveforms of the reference sound and the generated sound for both single beats (a) and double beats (b).

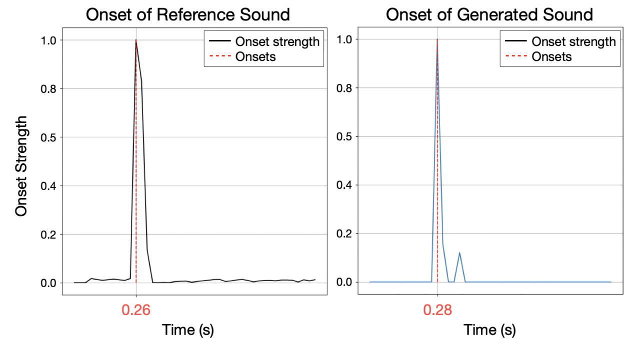

(a) Single-beat

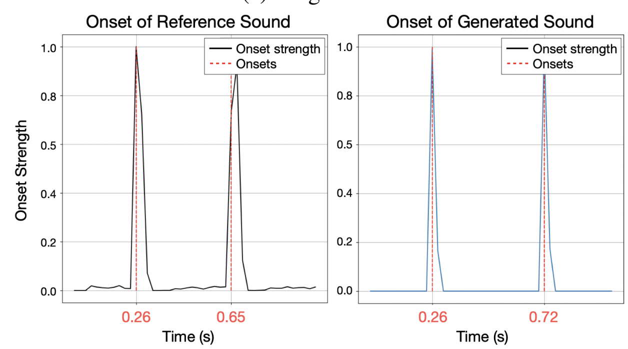

(b) Double-beat

Fig. 5: The onset strength and timing for both the reference sound and the generated sound in single beats (a) and double beats (b)

a tapping motion while the position of its wrist and the height of the drum pad are fixed and generate sound. Fig.2 represents our experimental hardware setup.

B. Experimental Results

Fig. 3 illustrates the learning curve of our prosthetic hand in generating tapping sounds, both single and double beats. We find that the agent achieves reasonable results after training 3000 episodes. Training an agent with a reward defined only by sound information is challenging. For example, motor noise, sudden external noise, and audio recorder operation error can cause the agent to receive low rewards even if it takes the correct action. To address this, our reward function reflects a variety of sound characteristics, not just a single characteristic, such as amplitude, onset strength, and timing. As a result, we observed that as the defined sound reward increased, the sound generated by the robot clearly became more similar to the shape of the reference sound, even in situations with external noise. This demonstrates the effectiveness of our reward function in guiding the agent towards accurate sound generation for prosthetic hands.

Fig. 4 presents the waveforms of the generated sound and the reference sound for both single beats (4a) and double beats (4b). While the generated sound exhibits some level of motor noise not present in the reference, the two waveforms are remarkably well aligned along the time axis. Notably, peak and low point demonstrate close resemblance, suggesting the generated sound shares similar frequency and amplitude characteristics with the reference.

Finally, we calculated onset strength and timing from the noise-filtered sound and evaluated the similarity. Fig. 5 represents plots of the onset strength and timing for both the reference and the generated sounds. Fig. 5a is results corresponding to a single-beat, where one can observe that the shape and timing of the onset strength plot from the robot hand's sound resemble the reference. Fig. 5b presents results for a double-beat, illustrating that the timing difference between the two beats also closely aligns with the timing difference in the reference. Fig. 6 represents the snapshot of a prosthetic hand movements for sound generation.

IV. Discussion

In this study, we utilize actual prosthetic hand hardware (PSYONIC Ability Hand) and an audio recorder to generate interaction sounds with a drum pad. This method can improve the realism of the prosthetic hand by incorporating auditory information other than visual and tactile into the prosthetic hand, providing an effect similar to that of a human hand. By more closely replicating the actions of a natural hand, this development can boost the confidence and comfort of users, facilitating smoother interactions in their daily lives. For instance, when users attempt to tap on different surfaces to distinguish materials, the generated tapping sounds can provide valuable auditory cues, enhancing their understanding of the environment and objects they interact with.

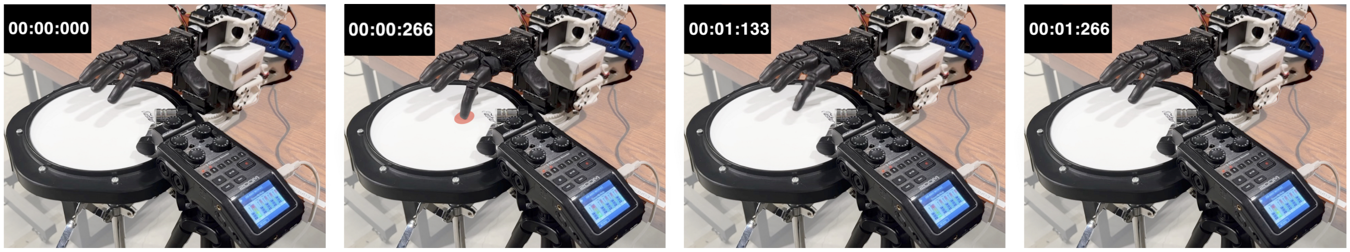

(a) Single-beat tapping motion

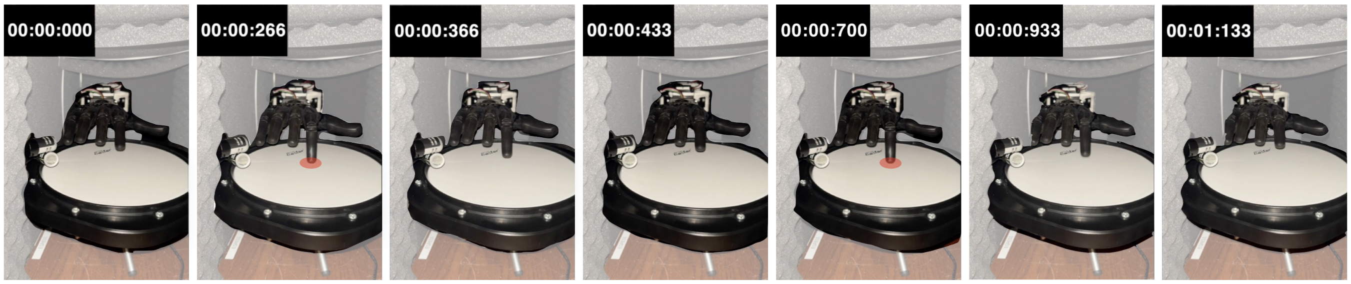

(b) Double-beat tapping motion

Fig. 6: Snapshot of prosthetic hand movements for sound generation. (a) The movement of the trained prosthetic hand when given a single beat sound of one-second duration, (b) The movement of the trained prosthetic hand when given a double beat sound of one-second duration.

The proposed methodology has the potential to impact the field of prosthetic hand research by introducing a novel sensory feedback modality. By demonstrating the feasibility of generating realistic sounds using reinforcement learning, this study opens up new avenues for enhancing the user experience and functionality of prosthetic devices. Compared to alternative methods, such as recording real finger motion and programming the prosthetic controller, our reinforcement learning approach offers several advantages. First, it allows for greater adaptability and flexibility, as the prosthetic hand can learn to generate appropriate sounds for a variety of surfaces and objects. Second, the learned policy can be fine-tuned and optimized over time, enabling the prosthetic hand to improve its performance based on user feedback and experience. Finally, our approach has the potential to be extended to other types of sounds and interactions beyond the tapping task described in this study, making it a versatile tool for enhancing the functionality of prosthetic devices.

Moreover, the technology presents the potential for increased sensory feedback. By producing sound corresponding to different movements, users can gain better spatial and situational awareness of their prosthetic limb. This could be particularly beneficial in scenarios where visual feedback is limited. For example, in low-light conditions or when handling objects outside the user's field of view, the auditory feedback can provide crucial information about the prosthetic hand's interactions with the environment.

However, it is important to acknowledge certain limitations encountered during the experiment. Creating a completely soundproof environment in real-world settings proved challenging. Ambient noise and actuator sounds from the prosthetic hand disrupted the recording of tapping sounds, impacting the reward mechanism as the prosthetic hand, despite performing correctly, often received unsatisfactory rewards. Furthermore, occasional malfunctions in the audio recording equipment led to sound recording errors, affecting the accuracy of reward calculations. To address these limitations, future studies could explore the use of soundproof booths or noise filtering technologies for the actuators to capture more accurate sound states and improve the learning process.

Based on the results of this study, we aim to expand beyond tapping sounds and explore training policies for prosthetic hands that use multiple fingers or tools like drumsticks to create sounds. This could lead to the development of more advanced and versatile prosthetic devices capable of producing a wide range of realistic sounds, further enhancing the user experience and functionality. Additionally, the proposed approach could be extended to other domains, such as robotic musicianship or industrial applications, where precise auditory feedback is crucial.

Acknowledgment

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program (Korea University), No. 2022-0-00871, Development of AI Autonomy and Knowledge Enhancement for AI Agent Collaboration, No. 2022-0-00612, Geometric and Physical Commonsense Reasoning based Behavior Intelligence for Embodied AI, and No. 2022-0-00480, Development of Training and Inference Methods for Goal-Oriented Artificial Intelligence Agents).

REFERENCES

[1] P. Fang, L. Tian, Y. Zheng, J. Huang and G. Li, ”Using thin-film piezoelectret to detect tactile and slip signals for restoring sensation of prosthetic hands,” 2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society, pp. 2565-2568, 2014.

[2] X. X. Liu, G. H. Chai, H. E. Qu and N. Lan, ”A sensory feedback system for prosthetic hand based on evoked tactile sensation,” 2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 2493-2496, 2015.

[3] J. DeGol, A. Akhtar, B. Manja and T. Bretl, “Automatic grasp selection using a camera in a hand prosthesis,” 2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 431-434, 2016.

[4] M. Sakuma, Y. Kito, Y. Tanaka, M. Yoshikawa and N. Kawashima, ”An Electric Cosmetic Prosthetic Hand with Vibrotactile Sense,” 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 4963-4966, 2020.

[5] L. Jabban, B. Metcalfe, D. Zhang and P. Iravani, ”Pressure Sensitive Skin for Prosthetic Hands: 2D Contact Location Determination Using Output Connections from a Single Side,” 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 4341-4344, 2020.

[6] J. Zhang, C. Chou, X. Wu, W. Pei and N. Lan, “Non-Invasive Stable Sensory Feedback for Closed-Loop Control of Hand Prosthesis,” 2022 44th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 2344-2347, 2022.

[7] A. Atasoy, E. Kaya, E. Toptas, S. Kuchimov, E. Kaplanoglu and M. Ozkan, ”24 DOF EMG controlled hybrid actuated prosthetic hand,” 2016 38th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC), pp. 5059-5062, 2016.

[8] L. Vargas, H. Huang, Y. Zhu and X. Hu, ”Stiffness Perception using Transcutaneous Electrical Stimulation during Active and Passive Prosthetic Control,” 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 3909-3912, 2020.

[9] J. A. George, T. N. Tully, P. C. Colgan and G. A. Clark, ”Bilaterally Mirrored Movements Improve the Accuracy and Precision of Training Data for Supervised Learning of Neural or Myoelectric Prosthetic Control,” 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), pp. 3297-3301, 2020.

[10] T. Fryen, M. Eppe, P. D.H. Nguyen, T. Gerkmann, and S. Wermter, “Reinforcement learning with time-dependent goals for robotic musicians,” arXiv preprint, 2020.

[11] H. Xu, Y. Luo, S. Wang, T. Darrell, and R. Calandra, “Towards learning to play piano with dexterous hands and touch,” In 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pp. 10410-10416, 2022.

[12] K. Zakka, L. Smith, N. Gileadi, T. Howell, X. B. Peng, S. Singh and P. Abbeel, “RoboPianist: A Benchmark for High-Dimensional Robot Control,” Conference on Robot Learning (CoRL), 2023.

[13] Puterman, L. Martin, “Markov decision processes: discrete stochastic dynamic programming,” John Wiley and Sons, 2014.

[14] M. Müller, “Dynamic time warping,” Information retrieval for music and motion, pp. 69-84, 2007.

[15] J. Schulman, F. Wolski, P. Dhariwal, A. Radford and O. Klimov, “Proximal policy optimization algorithms,” arXiv preprint, 2017.

[16] A. Akhtar, J. Cornman, J. Austin and D. Bala, “Touch feedback and contact reflexes using the PSYONIC Ability Hand,” MEC20 Symposium, 2020.

[17] J. Kim, D. C. Mathur, K. Shin, and S. Taylor, “PAPRAS: Plug-And-Play Robotic Arm System,” arXiv preprint, 2023.