Deep Visual Perception for Dynamic Walking on Discrete Terrain

Avinash Siravuru, Allan Wang, Quan Nguyen, and Koushil Sreenath

Abstract—Dynamic bipedal walking on discrete terrain, like stepping stones, is a challenging problem requiring feedback controllers to enforce safety-critical constraints. To enforce such constraints in real-world experiments, fast and accurate perception for foothold detection and estimation is needed. In this work, a deep visual perception model is designed to accurately estimate step length of the next step, which serves as input to the feedback controller to enable vision-in-theloop dynamic walking on discrete terrain. In particular, a custom convolutional neural network architecture is designed and trained to predict step length to the next foothold using a sampled image preview of the upcoming terrain at foot impact. The visual input is offered only at the beginning of each step and is shown to be sufficient for the job of dynamically stepping onto discrete footholds. Through extensive numerical studies, we show that the robot is able to successfully autonomously walk for over 100 steps without failure on a discrete terrain with footholds randomly positioned within a step length range of [45:85] centimeters.

I. INTRODUCTION

A. Problem Definition

The objective of this work is to build and systematically evaluate a Deep Visual Perception system to be used by a planar dynamic walking robot in order to autonomously walk on discrete terrain. We build a realistic visual simulator to generate the robot's first person view and combine it with an accurate physics simulator of the bipedal robot walking on discrete terrain. The physics simulator also contains an inner-loop safety-critical controller that can generate stable and safe limit cycle walking of a desired step length [19]. In this setup, we train a deep neural network to estimate the step length (distance to the next stepping location) using a single sampled image preview that is obtained at the beginning of each step. Detecting footholds and estimating distance is a classic object localization problem similar to object grasping in robotic manipulation, however in the case of locomotion there are additional challenges due to the time-critical and safety-critical nature of the problem.

Note that, we limit our attention to only autonomous planar walking, and accordingly, only predict step length information. This simplification allows us to keep the focus on the visual simulator development, custom CNN design (to bound worst-case estimate) and perception-control integration. However, the method itself can be extended to 3D



Fig. 1: At the beginning of each step, the convolutional neural network predicts the step length based on per-step visual preview from a torso-mounted camera. A suitable limit cycle is chosen from the gait library [19] and executed to achieve the predicted step length.

walking without loss of generality as evidenced by (a) our prior work on dynamic walking over terrain with varying step width and step height in addition to step length [20], [19], and (b) the perception system in this paper that takes an image rendered from a 3D scene as input without making any geometric simplifications due to the planar walking.

B. Motivation

Conventionally, robot perception involves parsing the entire scene, labeling objects of interest, and feeding this information for planning and control. As the number of tasks increase and the decision-making time drops, searching the entire scene for all tasks is expensive and trade-offs are inevitable. Reactive vision-in-the-loop control for sub-tasks, wherever possible, reduces overhead on higher-level planners and injects more dynamism into the system. The objective of this work is to serve this need for the sub-task of walking safely on discrete terrain.

In fact, experiments with humans and cats have shown that during over 50 - 60% time of the gait cycle, gaze is invested in target detection and the higher-level spatial navigation problem [21]. When required to walk on complex terrain requiring accurate foot placement, humans operate with intermittent visual samples of the foothold location and use the information in a feed-forward manner to adjust step length just 1 - 2 steps a priori [21], [17]. For walking on discrete terrain, the key finding here is that, instead of active modulation of gait through visual feedback during the entire stance phase, humans prefer to adjust their gait one-two step ahead using an intermittent visual preview and execute an energy-efficient ballistic motion [16]. If the foothold remains constant, continuous visual feedback may even be unnecessary. Our work is inspired by these

This work is supported in part by grants from the Pennsylvania Infrastructure Technology Alliance, the Google Faculty Research Award, and in part by NSF grant IIS-1526515.

A. Siravuru, A. Wang, and Q. Nguyen are with the Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA 15213, USA {avinashs, allanwan, quannguyen}@cmu.edu.

K. Sreenath is with the Dept. of Mechanical Engineering, University of California, Berkeley, CA 94720, koushils@berkeley.edu.

biological findings and we wish to elicit similar behavior from the dynamic bipedal robot ATRIAS. Our paper makes the following contributions:

- We present a custom deep convolutional neural network (CNN) architecture for the task of estimating next step location given a sampled visual preview of the terrain at the end of the current step. A synthetic outdoor dataset with dynamic real world textures, varying lighting conditions and camera positions is also developed. The network, trained on this dataset, returns an average prediction error of 1.62 cm and a worst-case error of 10.38 cm.
- 2) We integrate the CNN-based step length estimator with an inner-loop safety-critical controller to enable vision-in-the-loop dynamic walking on discrete terrain. The robot with the visual step-length estimator and the safety-critical controller is shown to successfully walk at least 100 steps without failure, with the step lengths randomly sampled from a uniform distribution of 45 - 85 cm.

II. RELATED WORK

Deep Learning for Robotic Vision: Recent advances of Deep Learning in fields like computer vision, natural language processing, speech recognition, etc. has lead to tremendous research interest in extending these gains to robotics. Large amounts of data available for network training and parallel computation helped accelerate this effort. From learning end-to-end visuomotor policies for object manipulation [15], to learning to fly UAVs in cluttered environments [23], or in self-driving cars [4], [6], deep learning is impacting all major robotics sub-domains. In humanoid robotics also, CNNs were used for innovative applications like surface friction estimation from images, to help in slip prediction [5]. However, it is very challenging to build end-to-end fully data-driven policies for stable and safe limit cycle walking, let alone on discrete terrain.

Perception in Legged Locomotion: Perception for bipedal locomotion primarily focused on foot-step planning for statically stable or linear dynamical model-based walkers. Usually, LIDAR-camera combination is preferred in this case. Accurate high resolution depth data obtained from Lidar is used for safe footstep detection and planning [7], [11]. Unlike these walkers, dynamics robots have point feet, move much faster and therefore need faster execution and the ability to pick footholds of any size. This makes the search problem harder on the full 3D map. Vision-in-the-loop walking with gait adjustment (comparable to our approach) was implemented on a Quadruped in [3] to operate on steep slopes and dense vegetation. However, the problem of discrete terrain is not addressed. We believe our solution is complimentary to their effort and a combined solution could pave the way for true rough terrain navigation.

III. ROBOT MODEL AND CONTROLLER SUMMARY

Having presented an overview of related work, we will now briefly develop the dynamical model and controller for TABLE I: Optimization constraints.

Motor Toque	$ u \leq 7 \mathrm{Nm}$
Impact Impulse	$F_e \leq 15 \text{ Ns}$
Friction Cone	$\mu \le 0.6$
Vertical Ground Reaction Force	$F_{st}^v \ge 200 \text{ N}$
Mid-step Swing Foot Clearance	$h_f _{s=0.5} \ge 0.1 \text{ m}$

achieving stable walking with precise foot placements.

A. Dynamical Model for Walking

We consider the ATRIAS bipedal robot with configuration q and state $x = (q, \dot{q})$ that denotes the generalized positions and velocities of the robot, with u denoting the joint torques. A hybrid model of walking can then be expressed as

$$\begin{cases} \dot{x} = f(x) + g(x)u & x \notin S \\ x^+ = \Delta(x^-) & x \in S, \end{cases}$$
(1)

where S is the switching guard, Δ is the reset or impact map, and the superscripts – and + denoting pre- and postimpact variables respectively. A detailed description of the robot and a derivation of its model can be found in [22].

B. Periodic Gait Design using Virtual Constraints

Virtual constraints are kinematic relations that synchronize the evolution of the robot's coordinates via continuous-time feedback control, thereby simplifying control of high degreeof-freedom systems, [25]. Virtual constraints are expressed as an output vector

$$y(q) = h_0(q) - h_d(s(q), \alpha),$$
 (2)

to be asymptotically zeroed by a feedback controller, with one virtual constraint typically imposed per each actuator. Here $h_0(q)$ specifies the variables to be controlled, and $h_d(s, \alpha)$ specifies the desired evolution of the controlled variables, parametrized by the the coefficient α and the gait phasing variable s which goes from zero at the start of the gait to one at the end of the gait.

A nonlinear constrained optimization is used to find the coefficient α so as to create a periodic orbit satisfying a desired step length, while respecting physical constraints on torque, motor velocity, and friction cone. The cost function is taken as the integral of squared torques normalized by step length,

$$J = \frac{1}{L_{step}} \int_0^T ||u(t)||_2^2 dt,$$
 (3)

and the constraints for the optimization are formulated as given in Table I, see [25, Sec. 6.6.2] for more details. The optimization is solved through a fast direct collocation framework from [13].

Since walking over stepping stones involves changes in the step length, one way to easily transition between different step lengths is to create a 2-step periodic gait that explicitly considers the step length of the current step as well as the subsequent step. The optimization process presented above can be used to design a 2-step periodic orbit with step lengths (l_0, l_1) with coefficient $\alpha(l_0, l_1)$ as shown in Fig. 2a.



Fig. 2: (a) 2-Step periodic walking with changing step lengths only. The walking gait is 2-step periodic therefore the step length of the second step and that of the initial condition are the same $(l_2 = l_0)$ and (b) Gait interpolation for the problem of changing step length.

With a collection of 2-step periodic orbits, We can then transition at each step between multiple 2-step periodic orbits in a MPC-like manner with small transients. However, to prevent an explosion of the number of periodic gaits that need to be optimized for, here we use only four 2-step periodic gaits corresponding to step lengths $(l_{min}, l_{min}), (l_{min}, l_{max}), (l_{max}, l_{min}), (l_{max}, l_{max})$ and use bilinear interpolation to find the coefficients α for a desired gait of step lengths (l_0^d, l_1^d) as shown in Fig. 2b. This work builds off recent work on periodic walking gait libraries in [19], [18], [9], [10].

C. Feedback Control

The presented optimization results in a desired walking gait encoded through $h_d(s(q), \alpha)$ in (2). The goal for the feedback control is to drive $y(q) \rightarrow 0$. In this work, we use a nonlinear feedback controller to enforce exponential stability for the hybrid system through input-output linearization, [2]. If y(q) has vector relative degree 2, then its second derivative can be expressed through Lie derivatives as,

$$\ddot{y} = L_f^2 y(q, \dot{q}) + L_g L_f y(q, \dot{q}) u.$$
 (4)

We can then apply the following pre-control law

$$u(q,\dot{q}) = u^*(q,\dot{q}) + (L_g L_f y(q,\dot{q}))^{-1} \ \mu, \tag{5}$$

where the nominal feedforward component $u^*(q, \dot{q}) := -(L_g L_f y(q, \dot{q}))^{-1} L_f^2 y(q, \dot{q})$ with feedback $\mu = -K_p y - K_d \dot{y}$ stabilizes the system. This combination of the 2-step periodic gait libary and the feedback controller results in precise footplacements with specified step lengths in a safety-critical manner. Formal guarantees on safety for dynamic walking on discrete terrain can also be achieved through control barrier functions [18].

IV. DIRECT PERCEPTION FOR ATRIAS

Having presented the robot's dynamical model and the safety-critical controller designed to walk on discrete terrain when provided with accurate step length information, we will now present a systematic way to build and train a deep visual perception model that estimates the step length from



Fig. 3: A collage of textures use for the synthetic outdoor dataset generation: (a) Background Textures, (b) Stone Textures.

a single monocular image. The system will take an input of the upcoming terrain through a front-facing camera at the beginning of each step. This image is to be fed to a CNN to estimate the step length for the next step. The step length estimate is then fed to the gait-library based controller to enable the robot to precisely land on the next foothold.

This CNN-based deep perception model has two critical components. Firstly, we need a large corpus of step-length annotated imagery of the robot's front person view while walking. This dataset is used to train the model for accurate step length estimation. The systematic methodology used to create this synthetic dataset is described in the next subsection. Secondly, we need a suitable deep neural network architecture that can best approximate the complex nonlinear mapping from image to step length estimate. The network needs to be tuned methodically to not only obtain the best test accuracy but also bound the worst-case prediction. These details are summarized after the next subsection.

A. Dataset Generation

To generate the image dataset, we use a popular opensource graphics software called Blender [12] to programatically generate realistic scenes. To create a discrete terrain scene, we need four key details: 1) Camera location and intrinsic parameters 2) Stepping stone location, 3) Lighting model and location, and 4) Color and texture information of both the stepping stone and background. These parameters will be randomized in ranges larger than what the robot may encounter in order to account for error accumulation over time.

The above four parameters are randomized for image generation in the following manner:

Camera Location: The camera location is measured with respect to the stance foot position. For each image, we randomly sampled from a range of $[-10 : 20] \times [-10 : 10] \times [80 : 120]$ (cm) to obtain the x, y, z offsets of the camera from the stance foot, respectively. The ZED Stereo Camera [14] model is used for rendering the images. However, since we only focus on the step length, we only generate monocular images for this study.

Stone Location: From [19], we note that, the robot was able to walk on a discrete terrain where the step lengths ranged from [20:90] (cm). True step length is the distance from the robot's stance foot to the next stone's center. For this work we uniformly sample step lengths from [15:95] (cm) range and arrange the stones in a single column accordingly.



Fig. 4: Sample images from the *Synthetic Outdoor Dataset* (SOD).

Based on the width of the robot, we spread out stones [30:35] (cm) away of the robot's center-line. Moreover, they are alternately positioned on either sides of the center-line, as shown in Fig. 4. Finally, the stone size itself is varied randomly between [10:20] (cm) in length and [60:90] (cm) in width, respectively. However, the stone height is kept fixed at 15 cm.

Lighting Properties: The light source is always facing the current stepping stone. However, its position is randomly chosen from a semi-spherical dome-like space above the robot with a 6 meters constant radius to the current stonecenter.

Texture: Here we select real world textures for both stones and background. We chose 12 unique textures comprising of grass, sand, water, pebbled terrain, etc. for the background and 14 unique textures for the stone including granite, cement, brick, wood, etc. Texture samples are shown in Fig. 3 Each scene is rendered by randomly sampling a texture pair from the available collection.

The images are rendered with a resolution of 223×223 pixels. Additionally, we crop the left and right 15% of the image to further reduce computational overhead. The final image resolution is 223×149 . We generate 40,000 images and call it the Synthetic Outdoor Dataset (SOD). Sample images are shown in Fig. 4. Having presented the dataset generation details, we will present the neural network architecture design next.

B. Network Architecture Summary

For training our object detector, we propose a custom neural network architecture. It consists of six convolutional layers of 32 filters each, followed by two dense layers, both with 256 neurons each. Unlike traditional designs where the number of features maps increase with depth, we found that, a constant number of feature maps throughout does better and needs fewer parameters. The kernel size of each filter is (4×4) . Batch Normalization and Max Pooling are applied after every two convolutional layers. Additionally, we apply Batch Normalization just before the final output layer as well. We use *relu* activations in all the layers except the last one, where we use a *linear* activation function instead. Fig. 5 summarizes the CNN architecture. The detector is trained using the *Mean Squared Error* loss and the *Adam* optimizer. We use a learning rate of 1e-4 along with a suitable learning rate decay policy. We use Keras API [8] with TensorFlow backend [1] to build and train our model. We trained for 40 epochs with a batch size of 50, on a Intel i7 machine with an NVIDIA Titan X GPU. The model has roughly 2.5 million trainable parameters. Finally we split the dataset into Training, Validation and Test sets, each comprising of 28900, 5100, 6000 images, respectively.

The step length is estimated as follows: Using the joint encoders on the robot, the location of the camera is first estimated. The deep neural network only learns to predict the distance from the camera to the next stone center. Therefore, the predicted step length is the sum of these two distances. Note that the camera position information is not used during training.

While deep networks have remarkable function approximation abilities, they have many tunable hyper-parameters whose fine-tuning critically impacts network performance and generalization ability. In the next section, we systematically outline our network design and customization process while examining the impact of each hyper-parameter on reducing worst-case test error.

C. Hyper-parameter Search

Deep neural networks have a very high dimensional hyperparameter space, where almost every single building block can be optimized. Most applied deep learning papers use existing architectures and leverage their transfer learning properties. Few papers explain in sufficient detail, the impact of each hyper-parameter on the learning outcomes. Unfortunately, hyper-parameter choices don't always generalize to all problems and it is therefore worthwhile to carefully tweak and specifically examine them for individual problems. Important insights drawn from this exploration for our problem are summarized below. Note that, all the results reported below are on the test set.

- Roughly 89% drop in error occurs within the first 20 epochs. Therefore, a wider hyper-parameter coarse search was carried out on models trained for 20 epochs while for the finer search, the models were trained for 40 epochs.
- As already identified in [24], Dropout with any probability or placement in the network worsens performance.
- Batch Normalization really boosts performance and results in around 33% drop in mean absolute error. More interestingly it leads to over 55% drop in the worst-case prediction error (or the maximum prediction error on the test set).
- Using L2 Regularization (default is 0.01) for only the fully connected layers helps further reduce the worst-case prediction error.
- We tested the Architecture with 5 kernel sizes, (2,2), (4,4), (6,6), (3,3), (2,4). We observed that rectangular kernels had the largest worst-case error, fol-



Fig. 5: Deep Network Architecture for predicting step length. Each convolutional block consists of two convolutional layers followed by a Max-Pool and Batch Normalization layers. Two identically dimensioned fully connected layers are used. Finally, the output activation function is linear.

lowed by the (3,3) kernel. Surprisingly, even numbered kernels did a better job, against conventional wisdom. The best kernel was (4,4).

• Adding an extra convolutional block (ie., two additional convolutional layers followed by a max pool and a batch norm) increased the mean absolute error. Unlike classification tasks where depth almost always helps, in regression tasks, localization accuracy is affected by max-pooling layers beyond some depth. Therefore, for regression tasks one must find the sweet spot between depth (complexity) and accuracy.

Finally, based on the above mentioned hyper-parameter search, an optimal network architecture is designed and it is trained with the synthetic outdoor dataset. The qualitative and quantitative results obtained are presented in the next section.

V. RESULTS AND DISCUSSION

In this section, we study the performance of the deep perception model and then integrate it with a physics-based simulator and gait-library based controller to numerically realize and analyze autonomous dynamic walking on randomly generated discrete terrain.

A. Step Length Prediction Performance

Once trained, we expect the step length predictor to accurately detect the next stepping stone and output it's distance in centimeters. Note that, each image will have anywhere between 1-5 stepping stones. Therefore, even though they have the same texture and geometry, our perception framework needs to overlook other stones and actively seek out the first one. In this situation, we believe perspective distortion helps in better distinguishing the stone of interest. Further, due to the safety-critical nature of this problem, in addition to describing network performance based on mean squared error, we will also report the variance and the worst-case prediction.

Test Avg. Loss	Std. Dev. of Loss	% Above 5 cm	Max Pred. Error
(in cm)			(in cm)
1.618	1.32	2.116	10.38







Fig. 6: Visualizing step length prediction performance through plots of (a) Predicted versus True Step Length values and (b) Prediction Error Histogram. Both plots are for the test data obtained from the Synthetic Outdoor Dataset.

Error is unavoidable in function approximation. However, in a safety-critical scenario like discrete terrain walking, the default approach of choosing the network that gives the least mean squared error could be detrimental as the worst-case estimate could still be off the safety limits. In order to avoid this issue, in this work, hyper-parameters were tuned with the objective to find the least possible worst-case estimate with the available dataset. The performance of the CNN was evaluated on the 6000 sample test data and is visualized in Fig. 6 and summarized in Table II.

1) Best and Worst predictions:: In addition to studying the qualitative learning outcomes like average loss, standard deviation of loss, etc., we also visualize images of the best-8 and the worst-8 predictions in Fig. 7 to visually interpret which parameters the model could and couldn't generalize over. From the figure, it is clear that the model is able to generalize over the various foreground and background textures, including bright and dim lighting conditions. However, shadows contributed to some of the higher estimation errors.





Fig. 7: Snapshots of the best-8 and worst-8 predictions of the neural network alongwith the corresponding prediction error in centimeters. The desired (red line) and predicted (blue line) step lengths are marked along with pixel coordinates of the resulting foot placement location (yellow dot).

B. Simulation Results

In this section, we integrate the CNN estimator with a physics-based robot simulator in III-A and the gait-library based controller in III-B to evaluate the closed-loop autonomous operation of the robot. The robot dynamics are simulated in Matlab. At the beginning of each step, robot's current position (specifically the camera position) is supplied to Blender to render a first-person-view synthetic outdoor image using the information from the terrain generator. This image is in turn supplied to the Convolution Neural Network (implemented in Keras and TensorFlow) to predict the step length for the next step. Provided with this predicted step length, the closed-loop dynamics of the robot and controller are simlated for one step to enable the robot to take the step forward. The process is repeated for subsequent steps. A schematic of the numerical simulation pipeline is shown in Fig. 8.

Numerical simulations are carried out with stones randomly placed with the inter-stone distance within the [45: 85] cm range. Recall that the camera position and stone location were uniformly sampled from dissimilar ranges in Section IV. The label used for training the CNN is the distance to the camera which is difference of stone location and camera position. Therefore, the distribution of labels used for training is no longer uniform. Accounting for this fact, the step length ranges have been adjusted to only test the CNN in the range where there was enough data to guarantee a good learning outcome.

In this study, we render images with light fixed in an overhead position and focus on carefully examining sensitivity to potential failure modes like camera position or step length going outside the range used for dataset generation during continuous simulation. (We have already noted earlier that shadows are a failure mode.) Our evaluation is based on two metrics, 1) Perception Error and 2) Foot Placement Error. Prediction Error is defined as the difference between Predicted Step Length and True Step Length and it is purely an artifact of the perception module. Similarly, Foot Placement Error is defined as the difference between footcontact-point and stone-center. This error is the cumulative error of both perception and control.

The robot was simulated for 100 steps and these two errors are plotted, as shown in Fig. 9. Note that the prediction error is bounded within a 5 cm range from the center for the most part and doesn't show significant accumulation of error over time. In Fig. 9(b), the foot placement error is also plotted for the case where perception was not used (called WOP - With Out Perception) and with perception (using SOD synthetic outdor dataset) for comparison. In an attempt to minimize effort, the conservative controller always understeps. In contrast, the prediction model over-estimates more frequently. Therefore, the errors cancel more often and lead to a desirable and mostly-bounded stepping behavior.

VI. CONCLUSIONS AND FUTURE WORK

In this work, we outlined a systematic way to design a CNN-based predictor that can estimate step lengths for a dynamic bipedal walker operating in discrete terrain. It is shown empirically that a feed-forward gait adjustment based on intermittent visual feedback is sufficient to walk on a discrete terrain where high-speed prediction and accurate foot placement are critical. Several visual factors that impact the predictor's performance are identified for further refinement. This paper integrates gait optimization, nonlinear control, vision, and deep learning.

As part of future work, the direct perception model will be extended to predict other related gait parameters like step width, step height and yaw angle, enabling autonomous rough terrain navigation. Moreover, since multiple steps are visible in an image, accuracy could be improved through Recurrent Neural Networks. Domain adaptation techniques that transfer the predictor from simulation to real world images will enable testing the algorithm on the real robot.

REFERENCES

[1] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, et al., "Tensorflow: A system for large-scale machine learning," in *Proceedings of the 12th USENIX* Symposium on Operating Systems Design and Implementation (OSDI). Savannah, Georgia, USA, 2016.



Fig. 8: Simulation pipeline for autonomous dynamic walking on discrete terrain. The pipeline integrates gait optimization, nonlinear control, vision, and deep learning. Simulation Video: https://youtu.be/ijJAPapU7qI.





Fig. 9: (a) Prediction Error, and (b) Foot Placement Error plots for 100 step walking simulation with step lengths varying within [45 : 75] cm. In (b) WOP indicates error *without perception* while SOD indicates errors when step length is estimated using the *synthetic outdoor dataset*-based image preview.

- [2] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle, "Rapidly exponentially stabilizing control lyapunov functions and hybrid zero dynamics," *IEEE Transactions on Automatic Control (TAC)*, vol. 59, no. 4, pp. 876–891, April 2014.
- [3] M. Bajracharya, J. Ma, M. Malchano, A. Perkins, A. A. Rizzi, and L. Matthies, "High fidelity day/night stereo mapping with vegetation and negative obstacle detection for vision-in-the-loop walking," in *IEEE/RSJ International Conference on Intelligent Robots and Systems* (*IROS*), 2013.
- [4] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, L. D. Jackel, M. Monfort, U. Muller, J. Zhang, *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [5] M. Brandao, Y. M. Shiguematsu, K. Hashimoto, and A. Takanishi, "Material recognition cnns and hierarchical planning for biped robot locomotion on slippery terrain," in *Humanoid Robots (Humanoids)*, 2016 IEEE-RAS 16th International Conference on. IEEE, 2016, pp. 81–88.
- [6] C. Chen, A. Seff, A. Kornhauser, and J. Xiao, "Deepdriving: Learning affordance for direct perception in autonomous driving," in *Proceedings of the IEEE International Conference on Computer Vision*, 2015, pp. 2722–2730.
- [7] J. Chestnutt, Y. Takaoka, K. Suga, K. Nishiwaki, J. Kuffner, and S. Kagami, "Biped navigation in rough environments using on-board

sensing," IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2009.

- [8] F. Chollet, "Keras," 2015.
- [9] X. Da, O. Harib, R. Hartley, B. Griffin, and J. Grizzle, "From 2d design of underactuated bipedal gaits to 3d implementation: Walking with speed tracking," *IEEE Access*, vol. PP, no. 99, pp. 1–1, 2016.
- [10] X. Da, R. Hartley, and J. W. Grizzle, "First steps toward supervised learning for underactuated bipedal robot locomotion, with outdoor experiments on the wave field," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2017.
- [11] M. F. Fallon, P. Marion, R. Deits, T. Whelan, M. Antone, J. McDonald, and R. Tedrake, "Continuous humanoid locomotion over uneven terrain using stereo fusion," *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 2015.
- [12] L. Flavell, "Beginning blender: open source 3d modeling, animation, and game design," *The expert's voice in open source*, 2010.
- [13] M. S. Jones, "Optimal control of an underactuated bipedal robot," Master's thesis, Oregon State University, ScholarsArchive@OSU, 2014.
- [14] S. Labs, "Zed stereo camera." [Online]. Available: https://www.stereolabs.com/
- [15] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *Journal of Machine Learning Research*, vol. 17, no. 39, pp. 1–40, 2016.
- [16] J. S. Matthis, S. L. Barton, and B. R. Fajen, "The critical phase for visual control of human walking over complex terrain," *Proceedings* of the National Academy of Sciences, p. 201611699, 2017.
- [17] J. S. Matthis and B. R. Fajen, "Humans exploit the biomechanics of bipedal gait during visually guided walking over complex terrain," *Proceedings of the Royal Society of London B: Biological Sciences*, vol. 280, no. 1762, p. 20130700, 2013.
- [18] Q. Nguyen, X. Da, J. W. Grizzle, and K. Sreenath, "Dynamic walking on stepping stones with gait library and control barrier," *Workshop on Algorithimic Foundations of Robotics*, 2016.
- [19] Q. Nguyen, X. Da, W. Martin, H. Geyer, J. W. Grizzle, and K. Sreenath, "Dynamic walking on randomly-varying discrete terrain with one-step preview," in *Robotics: Science and Systems (RSS)*, 2017.
- [20] Q. Nguyen, A. Hereid, J. W. Grizzle, A. D. Ames, and K. Sreenath, "3d dynamic walking on stepping stones with control barrier functions," in *Decision and Control (CDC), 2016 IEEE 55th Conference on.* IEEE, 2016, pp. 827–834.
- [21] A. E. Patla, "Understanding the roles of vision in the control of human locomotion," *Gait & Posture*, vol. 5, no. 1, pp. 54–69, 1997.
- [22] A. Ramezani, J. W. Hurst, K. Akbari Hamed, and J. W. Grizzle, "Performance Analysis and Feedback Control of ATRIAS, A Three-Dimensional Bipedal Robot," *Journal of Dynamic Systems, Measurement, and Control*, vol. 136, no. 2, 2014.
- [23] F. Sadeghi and S. Levine, "CAD2RL: Real singel-image flight without a singel real image," 2017.
- [24] J. Tobin, R. Fong, A. Ray, J. Schneider, W. Zaremba, and P. Abbeel, "Domain randomization for transferring deep neural networks from simulation to the real world," *arXiv preprint arXiv:1703.06907*, 2017.
- [25] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, J. Choi, and B. Morris, *Feedback Control of Dynamic Bipedal Robot Locomotion*, ser. Control and Automation. Boca Raton, FL: CRC, June 2007.