

Towards Learning Abstract Representations for Locomotion Planning in High-dimensional State Spaces

Tobias Klamt and Sven Behnke

Abstract—Ground robots which are able to navigate a variety of terrains are needed in many domains. One of the key aspects is the capability to adapt to the ground structure, which can be realized through movable body parts coming along with additional degrees of freedom (DoF). However, planning respective locomotion is challenging since suitable representations result in large state spaces. Employing an additional abstract representation—which is coarser, lower-dimensional, and semantically enriched—can support the planning.

While a desired robot representation and action set of such an abstract representation can be easily defined, the cost function requires large tuning efforts. We propose a method to represent the cost function as a CNN. Training of the network is done on generated artificial data, while it generalizes well to the abstraction of real world scenes. We further apply our method to the problem of search-based planning of hybrid driving-stepping locomotion. The abstract representation is used as a powerful informed heuristic which accelerates planning by multiple orders of magnitude.

I. INTRODUCTION

Most robot locomotion planners feature search- or sampling-based methods which perform well in 2D and 3D planning problems, e.g., driving robot locomotion [1]–[3]. However, many areas of operation, e.g., in search and rescue missions, have challenging properties. Suitable robots need to adapt to those environments to provide fast, safe, and energy efficient locomotion which can be realized through, e.g., tracked flippers, legs, or wheels at adjustable limbs. Including these capabilities in the planning problem results in high-dimensional state spaces which may lead to extensive searches for search- and sampling-based approaches.

In recent years, intensive research has been performed on learning-based motion planning approaches [4]–[8] which learn to map a given situation to a reasonable action without performing extensive searches. However, the necessary amount of training data and the required network complexity strongly depend on the size of the considered maps and the number of state space dimensions. Thus, at the current state-of-the-art, learning-based planners are restricted to small maps or low-dimensional planning problems.

A well investigated idea to accelerate planning in large state spaces is abstraction. An abstract representation has a coarser resolution or fewer state space dimensions and compensates this information loss through additional features to increase semantics. Given such an abstract representation,

All authors are with University of Bonn, Computer Science Institute VI, Autonomous Intelligent Systems, Bonn, Germany klamt@ais.uni-bonn.de, behnke@cs.uni-bonn.de. This work was supported by the European Union’s Horizon 2020 Programme under Grant Agreement 644839 (CENTAURO).

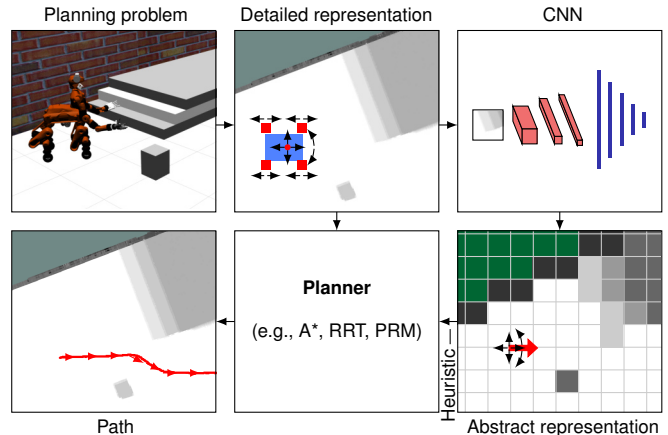


Fig. 1. A CNN is used to generate an abstract representation of a detailed planning problem which is employed to support the planner.

the necessity of the detailed representation during planning may be reduced to certain situations. However, a valuable property of an abstract representation to be used along with a detailed representation (e.g., as a heuristic or for coarse-to-fine planning) is that the same actions induce the same costs in both representations (*cost similarity*).

In [9], we have proposed a high-dimensional search-based planner for hybrid driving-stepping locomotion which we have extended in [10] to multiple levels of abstraction. This has shown promising results since planning has been significantly accelerated while the resulting path quality has stayed comparable. However, abstract representations were manually designed and parametrized to obtain *cost similarity* which is a challenging and exhausting task.

In this paper, we propose a method to support the implementation of abstract representations through convolutional neural networks (CNNs), which map a spatially small planning task to a corresponding cost assessment for the shortest path. Since these small planning tasks represent a coarse, low-dimensional set of actions, the CNN represents a *cost similar* abstract representation for the high-dimensional planning cost function (see Fig. 1).

We train the CNN on generated artificial data and evaluate it on simulated and real-world sensor data. Furthermore, the method is used to generate a powerful heuristic for hybrid driving-stepping locomotion planning, but it can be easily transferred to other domains, e.g., walking locomotion. The results indicate that the proposed method outperforms our manually tuned approach in terms of abstraction quality, while eliminating tuning efforts, and that the proposed heuristic accelerates path planning by multiple orders of magnitude, compared to popular heuristics.

II. RELATED WORK

Most robot motion planning approaches are either sampling-based, such as Rapidly-exploring Random Trees (RRT) [1] or Probabilistic Roadmaps (PRM) [2], search-based, such as A* [3] or a combination of those [11]. Low-dimensional motion planning in 2D or 3D state spaces, can be seen as solved with these approaches. However, it is still challenging to solve high-dimensional, large planning problems since the required computational power and memory significantly increase with an increasing state space size.

A solution to handle large environment sizes is multi-resolution planning [12]. To handle high-dimensional state spaces, a local adaptation of the robot representation is an option. In previous work [9], we have proposed a search-based approach to plan hybrid driving-stepping locomotion. Similarly, Dornbush et al. [13] have planned multi-modal paths for a humanoid with a search-based planner. Both approaches handle the occurring high-dimensional state spaces by separating the planning problem with respect to the locomotion mode and apply high-dimensional planning only if required. Nevertheless, both works suffer the problem of handling large scenarios in feasible time since the high-dimensional represented areas are still too large.

However, those approaches only neglect information in their coarse/low-dimensional representations which might result in wrong assessments, especially for complex terrain. This is addressed by abstraction: Representations are coarser but semantically enriched to compensate the information loss. A theoretical basis for abstraction for search-based planning has been given by Holte et al. [14]. In [10], we have extended hybrid driving-stepping locomotion planning to three levels of abstraction. With increasing abstraction, the environment is represented in a coarser resolution but with additional hand-crafted features such as height differences or terrain classes. In addition, the robot representation has a coarser resolution and less dimensions with increasing abstraction. The costs functions were manually tuned to obtain *cost similarity*. This was done by iteratively comparing costs on a small set of exemplary tasks and adjusting parameters. The abstract representations accelerate planning by multiple orders of magnitude while the path quality stays comparable. Especially the utilization of the most abstract representation as a heuristic leads to significant speedup. However, the design of descriptive features and tuning of cost functions require extensive manual parametrization and are very dependent on the used set of exemplary tasks.

In recent years, learning-based approaches for solving robot motion planning problems have been proposed. In [4] and [5], CNNs have been trained to map camera images directly to motor commands, e.g., for manipulation tasks or steering of a self-driving car. However, the long-term goal-directed behavior of such approaches is usually poor or the training would require unreasonable amounts of data and time. Tamar et al. [6] have proposed a differentiable approximation of the value iteration algorithm which can be represented as a CNN—the Value Iteration Networks. Their

performance has been evaluated on small 2D grid worlds. Similarly, Karkus et al. [7] have proposed QMDP-Net which is also capable of planning in 2D grid worlds. Srinivas et al. [8] have proposed Universal Planning Networks which map images of the initial and goal scene to actions. These three approaches point out the general problem of learning-based approaches at the current state-of-the-art: The required amount of training data and the required network complexity are not manageable for large, high-dimensional planning problems.

To summarize, learning-based planning approaches can handle local problems with limited state space sizes quickly without performing extensive searches. In contrast, traditional planning approaches show good goal-directed behavior but might get stuck in extensive searches for complex high-dimensional problems. Hence, it promising to combine these approaches and merge the advantages of both. Faust et al. [15] use a reinforcement learning agent to learn short-range, point-to-point navigation policies for 2D and 3D action spaces which capture the robot dynamic and task constraint without considering the large-scale topology. Sampling-based planning is used to plan waypoints which give the planning a long-range goal-directed behavior.

In contrast to that work, we combine learning- and search-based planning to handle 7-dimensional hybrid driving-stepping locomotion planning. A CNN represents the cost function of an abstract representation of the high-dimensional planning problem employed as a heuristic to accelerate planning.

III. PROBLEM STATEMENT

Given a planner which uses an environment representation (\mathcal{E}), a robot representation (\mathcal{R}_d), a corresponding action set (\mathcal{A}_d), and a cost function (\mathcal{C}_d). \mathcal{E} is a map with an arbitrary number of features describing each cell. \mathcal{R}_d represents all required DoF of the robot kinematics which are necessary to address the planning problem. \mathcal{A}_d contains all actions which can be executed by the robot such that $r_{d,i} + a_{d,j} = r_{d,i+1}$, an action $a_{d,j} \in \mathcal{A}_d$ connects two successive robot states $r_{d,i}, r_{d,i+1} \in \mathcal{R}_d$ while inducing the costs $\mathcal{C}_d(r_{d,i}, a_{d,j})$.

A second, abstract representation, consisting of \mathcal{E} , \mathcal{R}_a , \mathcal{A}_a , and \mathcal{C}_a , can be used to support the planning. \mathcal{R}_a describes the robot state in a low-dimensional state space although this might not suffice to describe the robot state in enough detail for execution. The correspondence between an abstract robot state $r_{a,i} \in \mathcal{R}_a$ and a detailed robot state $r_{d,i} \in \mathcal{R}_d$ is given through the transformation

$$r_{a,i} = \mathcal{T}_{d \rightarrow a}(r_{d,i}) \quad (1)$$

and vice versa with

$$r_{d,i} = \mathcal{T}_{a \rightarrow d}(r_{a,i}). \quad (2)$$

\mathcal{A}_a describes actions $a_{a,j}$ to move an abstract robot state $r_{a,i} \in \mathcal{R}_a$ to a successive state $r_{a,i+1} \in \mathcal{R}_a$. The resolution of \mathcal{A}_a is coarser compared to \mathcal{A}_d , such that an action sequence

$$\mathcal{T}_{a \rightarrow d}(r_{a,i}) + a_{d,j} + a_{d,j+1} + \dots + a_{d,j+k} = \mathcal{T}_{a \rightarrow d}(r_{a,i+1}), \quad (3)$$




Fig. 2. Architecture of the proposed CNN. Input are a height map patch and the goal state. Although it is not fed into the network, the start state is depicted as a red arrow for better understanding. Output are the *feasibility* and *costs* values. Convolutional layers are visualized as red cuboids; blue lines show fully connected layers. If not stated different, convolutions have a padding of 0 and a stride of 1.

is necessary in the detailed representation to perform the least cost transition between two successive robot states in the abstract representation, while the difference between the abstract costs $C_a(r_{a,i}, a_{a,j})$ and the detailed costs $C_d(\mathcal{T}_{a \rightarrow d}(r_{a,i}), a_{d,j}, \dots, a_{d,j+k})$ should be minimized to obtain *cost similarity*.

While \mathcal{R}_a and \mathcal{A}_a can be easily defined, C_a needs extensive tuning. We propose to represent C_a as a CNN to avoid these tuning efforts and improve abstraction quality.

IV. NETWORK DESIGN

We propose a regular CNN architecture—consisting of convolutional layers and successive fully connected layers—to learn the abstract cost function (see Fig. 2). Input are a height map patch with 72×72 pixels and the three-dimensional abstract goal state $r_{a,g}$. The start state $r_{a,s}$ is assumed to be always in the map patch center with a fixed orientation and is not fed into the network. For a given resolution of 2.5 cm, the map size is chosen such that for every abstract goal state $r_{a,g} = r_{a,s} + a_{a,j} \in \mathcal{A}_a$, the corresponding detailed goal state $r_{d,g} = \mathcal{T}_{a \rightarrow d}(r_{a,g})$ with any feasible leg configuration is completely inside this map patch. $r_{a,g}$ is defined in resolution steps relative to $r_{a,s}$.

Instead of only outputting costs which become infinite for infeasible queries, we output two values: The *feasibility* value describes whether there exists a path between $r_{a,s}$ and $r_{a,g}$, and, if so, the *costs* value describes the corresponding costs.

We discovered that key to a good abstraction performance are some convolutions with large kernel sizes which might be explained as follows: A first small convolution extracts descriptive map features from the input height map. Next, a second convolution possesses a kernel size which is similar to the size of a robot foot and thus can determine if foot placement is possible for each kernel position. The kernel size of the third convolution is chosen such that it covers the maximum action length for an individual foot. Hence, it can find connections between feasible foot positions in a certain distance which is valuable for, e.g., steps. The following convolutions and max pooling operations with small kernels




Fig. 3. Hybrid driving-stepping locomotion robots. a) Momaro, b) Centauro, c) corresponding detailed robot representation (blue = robot base, red squares = feet, arrows visualize the DoF).

do further processing on the actions and are followed by six fully connected layers.

The last fully connected layer is split: While *costs* are output directly, the *feasibility* output is processed by a sigmoid function since it is Boolean.

V. LEARNING ABSTRACTION OF HYBRID DRIVING-STEPPING LOCOMOTION PLANNING

We apply the proposed method to hybrid driving-stepping locomotion planning for our platforms e.g., Momaro [16] and Centauro [17] (see Fig. 3 a, b). Both are able to perform omnidirectional driving and stepping motions. A detailed robot representation which matches both robots is depicted in Fig. 3 c.

A. Detailed Representation

The environment representation \mathcal{E} is a height map which is generated from registered point clouds using the method by Droschel et al. [18] (Fig. 4 a). Regarding C_d , foot costs (see Fig. 4 b) and base costs, which describe the costs to place a single foot/the base at a given position/in a given state on the map, are computed from this height map. Foot costs and base costs are merged to state costs. The robot is represented in 7D states $r_d \in \mathcal{R}_d = (r_x, r_y, r_\theta, f_1, \dots, f_4)$ with the robot base state (r_x, r_y, r_θ) and the relative longitudinal position of each foot f_1, \dots, f_4 as shown in Fig. 3 c. Positions have a resolution of 2.5 cm while there are 64 discrete orientations. Lateral foot positions are fixed and foot heights are computed after a result path is found.




Fig. 4. Detailed environment representation: a) input height map, b) foot cost map (yellow = untraversable by driving, olive = unknown).




Fig. 5. Robot actions in \mathcal{A}_d : a) omnidirectional driving with fixed orientation, b) turning with fixed position, c) moving a foot relative to the base while keeping ground contact, d) longitudinal base shift, e) step. Grid and orientation resolution are enlarged to facilitate visualization.

Robot actions in \mathcal{A}_d are (see Fig. 5):

- omnidirectional driving within a 20-neighborhood with fixed orientation,
- turning to the next discrete orientation,
- moving an individual foot relative to the robot base while keeping ground contact,
- moving the base longitudinal relative to the feet, and
- performing a step with a single foot.

Steps are represented as the direct transition from a pre-stepping to a post-stepping state. Only those steps in the result path are refined to detailed motion sequences which consider robot stability and the detailed stepping motion. Each action carries costs with respect to the occurring foot and base costs that the individual robot elements experience. For driving locomotion, a large angular difference between the robot orientation and driving direction is punished with higher costs to prefer driving forward which brings advantages to the perception of the environment directly in front of the robot and when switching to stepping motions in the sagittal direction. An A*-based planner which uses the above presented representation is used to plan hybrid driving-stepping locomotion paths. More details can be found in [9].

B. Abstract Representation

\mathcal{R}_a contains 3D robot states $r_a = (r_x, r_y, r_\theta)$ which describe the robot position with a resolution of 10 cm and the orientation in 16 discrete steps. Individual foot configurations are neglected. \mathcal{A}_a contains

- moving the robot within a 20-neighborhood with fixed orientation (see Fig. 5 a) and
- turning to the next discrete orientation with fixed position (see Fig. 5 b).

Transforming a detailed robot state to an abstract robot state ($\mathcal{T}_{d \rightarrow a}$) is done by neglecting the foot positions and matching the position and orientation to the coarse resolution of the abstract state space. The transformation from an

abstract to a detailed robot state $\mathcal{T}_{a \rightarrow d}$ is more complicated: For all detailed robot base states that match the abstract state, we search the least cost foot configuration while preferring configurations which are close to the neutral robot configuration (Fig. 5 a). The detailed robot state with the minimum state costs is the transformation result.

C. Network Training

Training data is generated artificially. Hence, large datasets can be produced without considerable effort. A map generator produces height maps of the desired network input size. The following obstacles are placed randomly in those height maps:

- cuboid shaped obstacles of random size,
- walls of random length and height, and
- staircases of random width with a random number of stairs (with random height and length).

We produce 2,000 maps of each of the following categories:

- one/two/three cuboid obstacles,
- one/two walls,
- one cuboid obstacle and one wall,
- one staircase,
- one staircase and one wall, and
- one staircase whose orientation is in the interval $[-\frac{\pi}{16}, \frac{\pi}{16}]$ around the robot orientation. Those maps are used to set a learning focus on stair climbing.

For each map, we define 22 abstract goal states r_{a,g_i} with respect to \mathcal{A}_a . The start state $r_{a,s}$ is always in the map center with a fixed orientation. $r_{a,s}$ and r_{a,g_i} are transformed to \mathcal{R}_d using $\mathcal{T}_{a \rightarrow d}$. For some maps, a valid detailed start state $r_{d,s}$ cannot be found due to obstacles. Those maps are deleted. In total, we get a set of 11,327 maps with 249,194 tasks. Subsequently, we search for a shortest path from $r_{d,s}$ to r_{d,g_i} with our detailed A*-planner. For each task, we save the *feasibility* flag which describes if a path could be found. *Costs* are saved for all feasible tasks.

The network is trained using the SGD optimizer with a learning rate of 0.0001 and a momentum of 0.9. We use a BCE loss function for the *feasibility* and a L1 loss function for the *costs*. The *costs* loss is only considered in the backpropagation if the task is feasible. Losses are weighted with $\mathcal{W}_{feasible}$ and \mathcal{W}_{costs} , both starting at 1. If no improvement by means of a decreasing loss is achieved in three successive training epochs, the corresponding loss weight is divided by 5. This dynamic is applied to both losses individually. For evaluation, a threshold of 0.5 is used to make the *feasibility* output Boolean. A validation set which includes 100 maps of each mentioned category is generated and used to evaluate the training performance (Fig. 6). We train the network for 100 epochs and choose the state with the best results on the validation set for our experiments.

D. Abstract Representation as Heuristic

We utilize the learned abstract representation as a heuristic for planning in the detailed representation. For a given goal state $r_{d,g}$, a one-to-any 3D Dijkstra search is started from

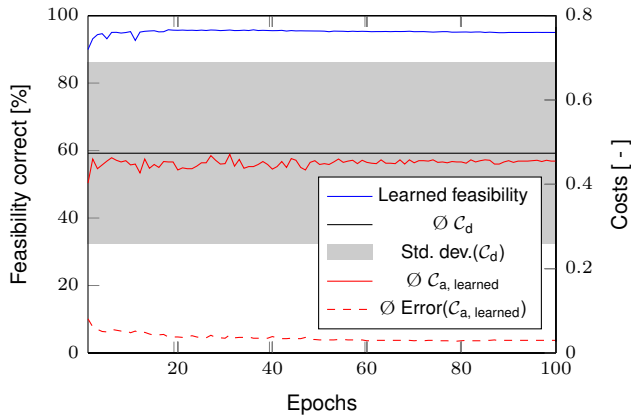


Fig. 6. CNN training performance. The detailed cost function C_d is shown as a base line. The stated error describes the cost difference between the detailed and the abstract cost functions.

$r_{a,g} = \mathcal{T}_d \mapsto_a(r_{d,g})$ and explores the whole map in the abstract representation. During that search, neighbor states r_{a,n_i} for a state $r_{a,m}$ are generated through abstract actions $a_{a,i}$ such that $r_{a,n_i} + a_{a,i} = r_{a,m}$ while respective costs are computed by the CNN which is fed with the respective height map patches. Start and goal of each action are exchanged since this search is running backwards, starting at the planner goal state. While actions which are assessed as infeasible are neglected, feasible actions are assigned the corresponding *costs* output.

Consequently, each abstract state in the map carries the estimated costs of the shortest path to $r_{d,g}$. When planning in the detailed representation, the planner uses these cost estimations as an informed heuristic.

Please note that we cannot prove that this heuristic always underestimates costs, and thus, we cannot prove admissibility for the generation of optimal paths. We rather focus on the generation of paths with a satisfying quality in feasible time and thus, accept sub-optimality to speedup planning.

The CNN is implemented using Python 2.7 and PyTorch 0.4.1. The planner is implemented in C++. Communication is realized via ROS. Code for the CNN, the training data generator and the framework to use the CNN as a heuristic is available online¹.

VI. EXPERIMENTS

We evaluate the proposed approach in two experiments which compare the abstraction quality to the manually tuned abstraction of our previous work [10] and show the performance of the proposed heuristic to plan hybrid driving-stepping locomotion. A video which shows additional footage of the experiments is available online².

A. Abstraction Quality

The abstraction quality is evaluated on three data sets:

- *random*: We generate 200 random maps of each category resulting in a set of 1,124 maps with 24,728 tasks.

¹https://github.com/AIS-Bonn/planning_abstraction_net

²https://www.ais.uni-bonn.de/videos/ICRA_2019_Klamt/




Fig. 7. Example tasks of the *random* test set (a), *simulated* test set (b), and *real* test set (c). Red arrows show start states (not fed into the CNN), green arrows show goal states.

TABLE I
ABSTRACTION QUALITY EVALUATION

	<i>random</i>	<i>simulated</i>	<i>real</i>
$\emptyset C_d$	0.476	0.466	0.509
Std. dev.(C_d)	0.222	0.202	0.236
<i>feasibility</i> correct, CNN	95.04%	96.69%	92.62%
$\emptyset C_{a,CNN}$	0.453	0.469	0.446
$\emptyset \text{Error}(C_{a,CNN})$	0.027	0.013	0.081
<i>feasibility</i> correct, man.tuned	79.27%	65.35%	69.77%
$\emptyset C_{a,man,tuned}$	0.435	0.402	0.429
$\emptyset \text{Error}(C_{a,man,tuned})$	0.057	0.021	0.103

- *simulated*: Height map patches of the desired size are cut out from height maps of simulated planning scenes. This set includes 77 maps with 1,694 tasks.
- *real*: Height map patches of the desired size are cut out from height maps that were generated from laser scanner measurements during real world experiments. This set includes 109 maps with 2,398 tasks.

We compared the performance to the manually tuned abstraction approach from our previous work. Finally, costs for the tasks in the detailed representation C_d are stated as a base line. We evaluate the *feasibility* and *costs* output. A correct *feasibility* assessment means that the abstract representation outputs the same *feasibility* value as the detailed representation. Only if both representations assess a situation as feasible, *costs* are considered and give an evaluation of the *costs similarity* of the two representations. Figure 7 shows some example tasks. The abstraction performance of the proposed CNN is shown in Tab. I.

The results indicate that the CNN *feasibility* output is significantly better compared to the manually tuned abstraction. While the latter has problems in simulated and real world robot environments, the CNN assesses a correct *feasibility* for > 92.62% of the tasks throughout all test sets.

Regarding the *costs* assessment, the average *costs* error of the CNN is smaller compared to the manually tuned abstraction on all test sets. The error is particular small when seen in relation to the large distribution of the base line costs. While the error of the proposed CNN is < 5.67% of the absolute *costs* on the *random* and *simulated* test sets, it is 15.9% on the *real* test set. This might be explained by noisier sensor measurements which result in noisier height maps.

B. Application to Planning

We designed a 10×10 m arena in Gazebo simulation which includes typical locomotion tasks for Centauro in search and rescue missions (Fig. 8). Environment perception is realized through a continuously rotating Velodyne Puck 3D laser scanner with spherical field-of-view at the

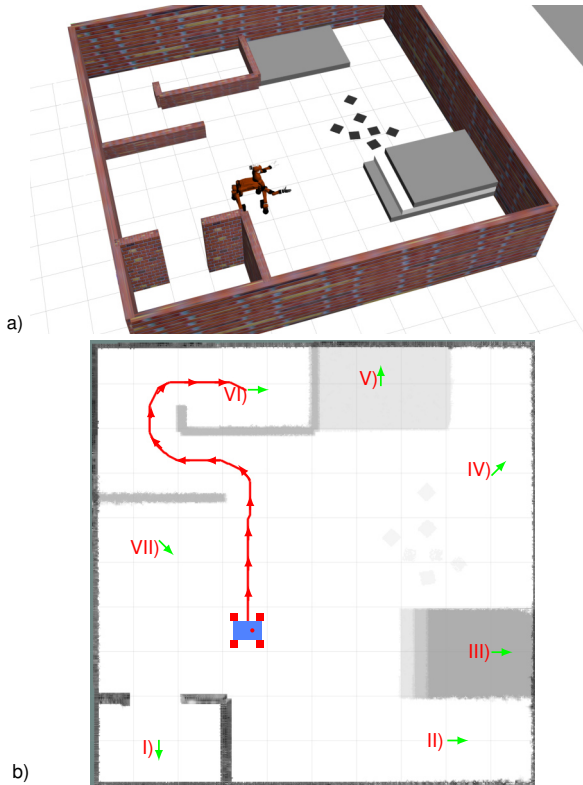


Fig. 8. Locomotion planning experiment. a) Gazebo arena with Centauro. b) Height map with start state (blue/red) and goals (arrows): I) Behind a narrow door, II) next to stairs, III) on top of stairs, IV) behind some clutter, V) on a platform, VI) inside a labyrinth, and VII) behind the robot. The red path is the resulting path to VI with the proposed heuristic and $\mathcal{W} = 1.25$.

robot head. Sensor measurements are processed to registered point clouds and used for localization using the method by Droschel et al. [18]. Height maps are generated from these point clouds. The used system is equipped with an Intel Core i7-8700K@3.70 GHz, 64 GB RAM and an NVidia GeForce GTX 1080Ti with 11 GB memory.

For all abstract states, map patches and neighbors with costs are precomputed by the CNN. This takes 239 s, is only required once per map, and can be incrementally updated if parts of the map change. The one-to-any Dijkstra search which starts from each goal state and generates the heuristic takes 0.049 s in average. We use the weighted-A* planner to plan paths to all goals while using the learned abstract representation as a heuristic. We compare the planning performance to a geometric heuristic. This combines Euclidean distances with rotational differences and is admissible. Hence, when used with a weight $\mathcal{W} = 1$, results are optimal. Both heuristics are evaluated with multiple $\mathcal{W} \geq 1$ to also obtain fast, sub-optimal solutions.

Figure 9 visualizes the planner performance for both heuristics and different \mathcal{W} . Table II summarizes the resulting speedup and cost increase compared to the optimal solution. The results indicate that the proposed abstraction-based heuristic accelerates planning by multiple orders of magnitude while, in particular for $\mathcal{W} = 1.25$, path costs stay comparable. This significantly outperforms the geometric heuristic. Especially for challenging tasks such as the stairs




Fig. 9. Planning times (including heuristic generation) and path costs for all goal states. The heuristic which is based on the learned abstract representation is compared to the geometric heuristic.

TABLE II
HEURISTIC PERFORMANCE

\mathcal{W}	Abstract representation			Geometric	
	1.0	1.25	2.0	1.25	2.0
speedup factor	27.80	708.5	10,860	12.00	27.88
costs increase	+4.77%	+10.5%	+33.1%	+6.07%	+33.9%

(III) and the labyrinth (VI), our heuristic was mandatory to obtain a solution in feasible time. This can be explained by the fact that the geometric heuristic has no information about the environment and thus the planner may expand many states before considering expensive actions. In contrast, the proposed abstraction-based heuristic uses its costs assessments to support the planner in its goal-directed behavior by including knowledge about the environment.

VII. CONCLUSION

In this paper, we propose a 3-dimensional abstract representation for a high-dimensional locomotion planning problem which employs a CNN to learn the cost function. The CNN maps a local planning task, consisting of a map patch and goal state, to a costs assessment for this task. We demonstrate how such an abstract representation can generate an informed heuristic for search-based high-dimensional planning. Experiments show that such a heuristic accelerates planning by multiple orders of magnitude, especially for challenging tasks. We further show that the learned representation outperforms a manually tuned abstract representation from previous work while eliminating tuning efforts.

REFERENCES

- [1] S. M. LaValle, "Rapidly-exploring random trees: A new tool for path planning," *Computer Science Dept., Iowa State University*, 1998.
- [2] L. Kavraki, P. Svestka, J.-C. Latombe, and M. H. Overmars, "Probabilistic roadmaps for path planning in high-dimensional configuration spaces," *IEEE Transactions on Robotics and Automation*, vol. 12, no. 4, pp. 566–580, 1996.
- [3] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Transactions on Systems Science and Cybernetics*, vol. 4, no. 2, pp. 100–107, 1968.
- [4] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research (JMLR)*, vol. 17, no. 1, pp. 1334–1373, 2016.
- [5] M. Bojarski, D. Del Testa, D. Dworakowski, B. Firner, B. Flepp, P. Goyal, *et al.*, "End to end learning for self-driving cars," *arXiv preprint arXiv:1604.07316*, 2016.
- [6] A. Tamar, Y. Wu, G. Thomas, S. Levine, and P. Abbeel, "Value iteration networks," in *Advances in Neural Information Processing Systems (NIPS)*, 2016.
- [7] P. Karkus, D. Hsu, and W. S. Lee, "QMDP-Net: Deep learning for planning under partial observability," in *Advances in Neural Information Processing Systems (NIPS)*, 2017.
- [8] A. Srinivas, A. Jabri, P. Abbeel, S. Levine, and C. Finn, "Universal planning networks," *arXiv preprint arXiv:1804.00645*, 2018.
- [9] T. Klamt and S. Behnke, "Anytime hybrid driving-stepping locomotion planning," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017.
- [10] T. Klamt and S. Behnke, "Planning hybrid driving-stepping locomotion on multiple levels of abstraction," in *IEEE/RSJ International Conference on Robotics and Automation (ICRA)*, 2018.
- [11] J. Li, S. Liu, B. Zhang, and X. Zhao, "RRT-A* motion planning algorithm for non-holonomic mobile robot," in *IEEE Society of Instrument and Control Engineers Annual Conference (SICE)*, 2014.
- [12] S. Behnke, "Local multiresolution path planning," in *RoboCup 2003: Robot Soccer World Cup VII*, ser. Lecture Notes in Computer Science, Springer, vol. 3020, 2003, pp. 332–343.
- [13] A. Dornbush, K. Vijayakumar, S. Bardapurkar, F. Islam, M. Ito, and M. Likhachev, "A single-planner approach to multi-modal humanoid mobility," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2018.
- [14] R. C. Holte, M. Perez, R. Zimmer, and A. MacDonald, "Hierarchical A*: searching abstraction hierarchies efficiently," in *Symposium on Abstraction, Reformulation, and Approximation*, 1995.
- [15] A. Faust, O. Ramirez, M. Fiser, K. Oslund, A. Francis, J. Davidson, and L. Tapia, "PRM-RL: long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning," in *IEEE/RSJ International Conference on Robotics and Automation (ICRA)*, 2018.
- [16] M. Schwarz, T. Rodehutsors, D. Droschel, M. Beul, M. Schreiber, N. Araslanov, I. Ivanov, C. Lenz, J. Razlaw, S. Schüller, D. Schwarz, A. Topalidou-Kyniazopoulou, and S. Behnke, "NimbRo Rescue: Solving disaster-response tasks with the mobile manipulation robot Momaro," *Journal of Field Robotics*, vol. 34, no. 2, pp. 400–425, 2017.
- [17] T. Klamt, D. Rodriguez, M. Schwarz, C. Lenz, D. Pavlichenko, D. Droschel, and S. Behnke, "Supervised autonomous locomotion and manipulation for disaster response with a centaur-like robot," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.
- [18] D. Droschel, M. Schwarz, and S. Behnke, "Continuous mapping and localization for autonomous navigation in rough terrain using a 3D laser scanner," *Robotics and Autonomous Systems*, vol. 88, pp. 104–115, 2017.