



UNIVERSITY OF LEEDS

This is a repository copy of *Human-like Planning for Reaching in Cluttered Environments*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/158051/>

Version: Accepted Version

Proceedings Paper:

Hasan, M orcid.org/0000-0002-7477-7133, Warburton, M orcid.org/0000-0001-5309-4424, Agboh, WC orcid.org/0000-0002-0242-0215 et al. (6 more authors) (2020) Human-like Planning for Reaching in Cluttered Environments. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). ICRA 2020, 31 May - 04 Jun 2020, Paris, France. IEEE , pp. 7784-7790. ISBN 978-1-7281-7396-2

<https://doi.org/10.1109/ICRA40945.2020.9196665>

© 2020, IEEE. This is an author produced version of a paper accepted for publication in 2020 International Conference on Robotics and Automation (ICRA). Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

See Attached

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Introducing a Human-like Planner for Reaching in Cluttered Environments

Mohamed Hasan¹, Matthew Warburton², Wisdom C. Agboh¹, Mehmet R. Dogar¹, Matteo Leonetti¹, He Wang¹, Faisal Mushtaq², Mark Mon-Williams² and Anthony G. Cohn¹

Abstract—Humans, in comparison to robots, are remarkably adept at reaching for objects in cluttered environments. The best existing robot planners are based on random sampling in configuration space- which becomes excessively high-dimensional with a large number of objects. Consequently, most of these planners suffer from limited object manipulation. We address this problem by learning high-level manipulation planning skills from humans and transfer these skills to robot planners. We used virtual reality to generate data from human participants whilst they reached for objects on a cluttered table top. From this, we devised a qualitative representation of the task space to abstract human decisions, irrespective of the number of objects in the way. Based on this representation, human demonstrations were segmented and used to train decision classifiers. Using these classifiers, our planner produced a list of waypoints in task space. These waypoints provide a high-level plan, which can be transferred to an arbitrary robot model and used to initialize a local trajectory optimiser. We evaluated this approach through testing on unseen human VR data, a physics-based robot simulation and real robot experiments. We find that this human-like planner outperforms a state-of-the-art standard trajectory optimisation algorithm and is able to generate effective strategies for rapid planning, irrespective of the number of objects in a cluttered environment. Our dataset and source code are publicly available¹.

I. INTRODUCTION

Imagine grasping a yoghurt tub from the back of a cluttered fridge shelf. For humans, this is a trivial task and one that even very young children are able to perform exquisitely. Yet, for a robot to be able to execute the same action with such ease, there are a number of non-trivial questions that it must resolve. Should the robot try to navigate through the available free space? Or might it be better to move obstructing items away first? In which case, which object should be moved first and where to?

Standard robot motion planning approaches focus on identifying a collision-free trajectory that satisfies a set of given constraints [1] and the majority of current planning techniques are based on random sampling [2], [3], [4], [5] of the configuration space. A defining feature of these sampling-based planners (SBPs) is the use of a set of probing samples drawn to uniformly cover the state space. To accelerate the planning process, it is thus desirable to devise non-uniform sampling strategies that favor sampling in regions where an optimal solution might lie [6]. Finding such regions is non-




Fig. 1: Overview of the *HLP* approach.

trivial, but, as we set out in the opening, humans can find near-optimal solutions very quickly.

Predicated on human expertise, imitation learning from demonstration (LfD) techniques are increasingly being adopted by researchers for robot motion planning [7], [8] [9], [10]. For example, researchers have demonstrated the use of neural networks for learning the dynamics of arm motion from human data [11], whilst others have shown the utility of combining planning and learning-based approaches to facilitate goal-directed behavior during human-robot interactions [12]. Alternative approaches to learning from human data include learning *qualitative* task representations. Evidence indicates that humans recognize and interact with spatio-temporal space in a more qualitative rather than quantitative manner [13], [14], [15]. As such, previous work integrated qualitative spatial representations (QSRs) with manipulation planning at different levels [16], [17]. Importantly, these approaches avoid the pitfalls of SPBs [18], [19], [20] - which only allow a small number of objects due to the curse of dimensionality,

We propose a novel approach to the problem of reaching through a cluttered environment based on geometric reasoning and workspace partitioning [1] (similar to cell decomposition) augmented by QSRs [14]. To achieve this, we collected data from human participants for our scenario using virtual reality (VR)[21]. Demonstrations were segmented based on a spatial model of the task space. This segmentation resulted in a set of state-action pairs used to train classifiers via standard machine learning techniques. These trained classifiers are used by a human-like planner (HLP) which, during testing, generates a high-level plan in task space. Finally, the generated high-level plan was forwarded to a robot-specific low-level controller for inverse kinematics (IK) and local optimisation if required. The resulting method is a human-like planning (HLP) algorithm for reaching objects in cluttered environments (see Fig. 1 for an overview of this

¹School of Computing, University of Leeds, UK.

²School of Psychology, University of Leeds, UK.

¹<https://github.com/m-hasan-n/human-like-planning.git>

framework).

Our novel contributions include:

- the development of a high-level planning algorithm that learns from human participants interacting in VR;
- a new qualitative space/action representation to mitigate the problem of high dimensionality; and
- empirical demonstrations of the utility of this high-level planner – showing that it is scalable and can work in conjunction with any existing low-level planner in a seamless manner.

II. AN OVERVIEW OF THE HLP

Given a cluttered environment $\xi = \{X^s, X^t, X_i^o\}$ represented by a start position X^s , a target position X^t and positions of N movable objects $X_i^o, i = 1, \dots, N$, we consider the problem of planning a high-level path ρ from X^s to X^t that is most likely to be planned by a human. A high-level plan $\rho = \{X_i^k, a_i^k\}$ is represented by a set of M keypoints X_i^k labeled by a set of associated actions $a_i^k, i = 1, \dots, M$.

Our framework (Fig. 1) starts by learning from human demonstrations collected through VR. Each demonstration is given in the form of an environment representation ξ in addition to an associated human trajectory τ^h . This framework runs in two phases, training and testing. In the training phase (green arrow path of Fig. 1): given pairs of $\{\xi_{train}, \tau^h\}$, demonstrations are segmented into abstracted state-action pairs $\{S, A\}$ that are used to train decision classifiers. In the testing phase, given ξ_{test} of a new scene, the HLP uses the trained decision models to generate ρ . The generated high-level plan can then be transferred to an arbitrary robot model. A robot-specific low-level controller solves for IK and performs local optimisation if required.

Demonstration segmentation is based on our modeling of the task space in order to extract required state-action pairs. Modeling of the task space is detailed in Sec. IV. Feature extraction and decision classifiers are explained in Sec. V whilst the planning algorithm used in the testing phase is given in Sec. VI. Demonstration of the transfer from high-level planning to robot control is provided in the Experiments section.

III. VR DATASET COLLECTION

A dataset ² of human demonstrations was collected by asking 24 participants to perform 90 reaching trials towards objects placed on a cluttered table top in a virtual environment (Fig. 2). The table top surface was surrounded by transparent screens from all but the front side and the work space dimensions were tailored to suit human arm movement. The position and rotation of the hand, elbow and upper arm (nearest point to shoulder) were tracked and sampled at 90Hz.

Participants were asked to initiate each trial by first moving towards a home position, which was indicated by a transparent pink ball. After onset, a target end-point a

red-coloured can appeared along with a total of 6 obstacles placed on two rows. Participants were instructed to interact with the scene to pick up the can and move it to the start point. Participants could achieve this by either navigating around the obstacles or picking up and moving the obstacles to a new position. Trials were failed if any part of the arm interacted with the obstacles, if any moved object hit the edge of the workspace, if the can touched any of the obstacles or if participants knocked over an obstacle with their hand.

IV. MODELING THE TASK SPACE

Devising a model of the task space that enables one to learn from demonstrations and generalize with different environment settings is of critical importance. To this end, we designed the row-based structure shown in Fig. 1 and Fig. 3, with two rows each containing three objects and aligned start-target positions. This structured environment helped model the task space and train decision learners. All VR demonstrations used for training were generated according to this structure and we show in the Experiments section how the HLP is able to generalize to different environment settings.

We model the environment task space ξ qualitatively as a union of three disjoint sets: ξ^o representing the occupied space while the free space is given by the union of gaps ξ^g and connecting space ξ^c sets. Occupied space represents all objects in the scene in addition to walls³ of the surface on which objects rest. A gap $g_i \in \xi^g$ is a qualitative representation of the free space existing between two objects o_j and $o_k, j \neq k$ at the same row. The connecting space models the remaining free space used to connect from start to first row, between rows and from second row to target.

Based on this representation of the task space, the action space A is discretized into two primitive actions: A^o moving an object and A^g going through a gap (Fig. 2). A primitive action occurs at a keypoint in the $\xi^o \cup \xi^g$ space, i.e. an action applies to an object or a gap at a specific row. Actions at different rows are connected through the connecting space ξ^c .

Keypoints at objects and gaps are conceptually analogous to the idea of keyframe-based learning from demonstration [22]. They are also comparable to nodes in RRTs and PRMs [23], [19]. Similar to interpolation between keyframes or connecting graph/tree nodes, keypoints are locally connected in ξ^c .

The high-level plan is generated hierarchically in three levels: path, segment and action. At the top level, a path ρ comprises a sequence of consecutive segments. Each segment is a qualitative description of a subset of ξ^c in which the planner (connects) interpolates between two consecutive actions. Finally, an action is one of the two primitive actions in A . For example, considering an environment structured in N_R rows, a path consists of N_R-1 segments. Each segment connects between a pair of consecutive rows. One action

³A set of four virtual objects are used to represent top, left, right and bottom walls and added to the occupied space to avoid slipping objects off the surface.

²<https://doi.org/10.5518/780>




Fig. 2: Sample keyframes in a VR trial. This structure of two rows with three obstacles in each row is fixed among all trials. However, our approach can generalize to different structures as seen in Sec. VII

takes place at a row and applies to a gap or an object. Start and target points are directly connected to the actions at first and last rows respectively.

V. DECISION CLASSIFIERS

Decision learners are used to locally capture rules underlying human planning skills. We design such learners as classifiers that map from state (feature) domain X to action (class) domain Y . Training examples are drawn from a distribution D over $X \times Y$ where the goal of each classifier [24], [25], [26] is finding a mapping $C : X \mapsto Y$ that minimizes prediction errors under D . The classifier output may be represented by the posterior probability $P(Y|X)$.

In this section, we introduce the features extraction required for decision classifiers. Features⁴ are extracted based on the spatial representations used by humans to recognize the task space. We heavily use relational spatial representations like distance, size, orientation and overlap. In our working scenario, there are four main decisions during a plan: (1) which gap to go through? (2) which object to move? (3) to which direction the object should be moved? (4) which segment to connect two consecutive actions? Except for the object-moving direction, we designed our learners as binary classifiers.

⁴Features are appropriately normalized to relevant factors like surface size, objects area, number of objects and arm dimensions.




Fig. 3: Modeling the task space. A 2D projection is shown for a structure defined by two rows with three objects (and hence four gaps) each, start location and target object. Space around an object is discretized into eight direction classes. Size of each direction block depends on the object size and a scaling factor α .

A. Gaps and Objects Classifiers (C_g, C_o)

These classifiers have a binary output defining the probability of either selecting or rejecting a gap or an object. Intuitively, humans prefer gaps close to their hand, having proper size, leading to the target and not requiring acute motions. Hence, the features vector, X_g , input to gap classifier C_g is defined by the distances from gap to initial d_{x^g, x^s} and target d_{x^g, x^t} positions, gap (size) diagonal l_g , orientation of the gap-to-start $\theta_{g,s}$ and gap-to-target $\theta_{g,t}$ lines.

$$X_g = [d_{x^g, x^s}, d_{x^g, x^t}, l_g, \theta_{g,s}, \theta_{g,t}]^T \quad (1)$$

Similarly, object features are given by distances d_{x^o, x^s} and d_{x^o, x^t} , object diagonal l_o , the orientation angles $\theta_{o,s}$, $\theta_{o,t}$ in addition to object's overlap with the target $l_{o,t}$ and a measure of the free space volume/area around object a_{ofs} .

$$X_o = [d_{x^o, x^s}, d_{x^o, x^t}, l_o, \theta_{o,s}, \theta_{o,t}, l_{o,t}, a_{ofs}]^T \quad (2)$$

Space around a given object is discretized into a set Δ of eight classes (forward FF , forward left FL , left LL , back left BL , back BB , back right BR , right RR and forward right FR) covering object's neighborhood (Fig. 3). The size of each direction block depends on the object size and an arbitrary scaling factor. The free space in each block is computed and a_{ofs} is given by the sum of free space in the eight blocks.

B. Object Direction Classifier (C_d)

If the action is to move an object, we learn appropriate moving direction from human data. To estimate the moving direction, an object is described by the direction of human hand $h_d \in \Delta$ when approaching the object, orientation of an object-to-target line $\theta_{o,t}$ and amount of free space a_{fs} in each surrounding direction around the object. This is a multi-class classifier whose output $Y_d = \Delta$.

$$X_d = [h_d, \theta_{o,t}, a_{fs}]^T \quad (3)$$

C. Connecting Segments Classifier (C_c)

Each segment connects two actions applied on gaps and/or objects. Let e_1 and e_2 denote the elements representing these gaps/objects at the first and second rows respectively. It is desirable to avoid actions that are considerably apart from each other, having no or small overlap with target and expected to result in large collision. Hence, a segment feature vector consists of the signed horizontal $d_{(e1, e2)_x}$ and vertical $d_{(e1, e2)_y}$ distances, overlap $l_{c,t}$ between $d_{(e1, e2)_x}$ and

the target object, segment orientation θ_c w.r.t a straight line connecting initial and target positions and the collision ξ_c expected to accompany motion through this segment.

Segment collision is computed as the overall overlapping volume (area) between human forearm (link) and ξ^o . The overlapping area is approximated by the intersection polygon area and found by scanning the surrounding area using sampling lines.⁵

$$X_c = [d_{(e1,e2)_x}, d_{(e1,e2)_y}, l_{c,t}, \theta_{cC_s}]^T \quad (4)$$

VI. HLP ALGORITHM

The *HLP* algorithm uses the trained classifier models explained in the previous section to generate the high-level plan in the testing phase. The algorithm starts by locating rows $R_i, i = 1, \dots, N_R$ and gaps ξ^g in the scene.

For each i -th row, objects and gaps classifiers are called (Lines 2-3) to identify the best N_{i_o} objects and N_{i_g} gaps respectively. Selected gaps and objects in (i)-th row are connected to their counterparts in the next ($i+1$)-th row through connecting segments. Selecting a total $N_g = \sum_{j=i}^{i+1} N_{j_g}$ gaps and $N_o = \sum_{j=i}^{i+1} N_{j_o}$ objects at a pair of consecutive rows results in $N_g N_o$ segment combinations. These combinations are constructed and passed to the segments classifier that selects the best N_c connecting segments (Lines 4-6).

Algorithm 1 The Human-Like Planner (HLP)

Input: Environment representation $\xi = \{X^s, X^t, X_i^o\}$

Output: High-level path ρ

Locate rows R and gaps ξ^g

- 1: **for all** R **do**
 - 2: Compute gaps feature vector X_g
 $G_{selected} \leftarrow C_g(X_g)$
 Compute objects feature vector X_o
 $O_{selected} \leftarrow C_o(X_i^o)$
 - 3: **end for**
 - 4: **for all** pairs of consecutive rows **do**
 - 5: $C \leftarrow$ Segment Constructor ($G_{selected}, O_{selected}$)
 Compute segments feature vector X_c
 $C_{selected} \leftarrow C_c(X_c)$
 - 6: **end for**
 - 7: **for all** $C_{selected}$ **do**
 - 8:
 - 9: **if** $a^o \in C_{selected}$ **then**
 - 10: Compute object-direction feature vector X_d
 Object direction = $C_d(X_d)$
 Augment $C_{selected}$ by expected object's location
 - 11: **end if**
 Compute arm configuration feature vector X_a
 Estimate arm configuration: $R_a(X_a)$
 Compute expected path collision ρ_ζ
 - 12: **end for**
 Select the path with minimum collision score
-

⁵Although there are standard graphics polygon clipping algorithms to solve this kind of problems, we preferred an approximate but faster line-sampling approach.

A segment connects two actions, one of which may belong to the moving-object class. Hence, the object moving direction is estimated by the C_d classifier using same convention of space discretization in Sec. V-B. The expected object location after moving is found and added to the segment's sequence of keypoints (Lines 7-11).

Next, the human-arm configuration is estimated (explained in next section) at each keypoint and estimated configurations are used to evaluate the overall expected collision of each segment (Lines 11-12). To this end, we get candidate segments between rows that are labeled by a measure of their overall expected collision. For a two-row structure, the best segment is picked as the one with least likely collision.

A. Arm Configuration and Collision Detection

In the 2-row structure, a candidate path has one segment that is connected to the start and target points. Having a number of candidate paths, we selected the one with minimum overall expected collision. Overall expected path collision is found by computing collision of full-arm motion with ξ^o between path keypoints. Therefore, arm configurations are firstly estimated at each keypoint and then expected arm collision is computed.

A human arm is modeled as a planar arm with four joints at neck, shoulder, elbow and hand. The arm configuration is represented by two angles: θ_{sh} between neck-shoulder and upper arm links and θ_{el} between upper arm and forearm links. The arm configuration at a given keypoint K_t is estimated by regression. Input features to the configuration regression model R_a are: hand direction $h_d \in \Delta$ when approaching K_t , arm configuration at the previous keypoint K_{t-1} and signed horizontal and vertical distances between the two keypoints.

$$X_a = [h_d, \theta_{sh_{t-1}}, \theta_{el_{t-1}}, d_{(k1,k2)_x}, d_{(k1,k2)_y}]^T \quad (5)$$

By estimating arm configurations at keypoints and hence joint positions, full arm collision is computed as sum of its links collision during motion along keypoints from start to target. Collision of each link is computed using the same approach as in Sec. V-C. Intersections between sampling lines and objects are found and the area of the resulting intersection polygon defines the collision amount. To this end, a number of candidate paths are generated, each labeled with a measure of its overall expected collision ρ_ζ . This step completes the *HLP* algorithm and generates a plan defined by the path having minimum expected overall collision.

VII. EXPERIMENTS AND RESULTS

The *HLP* approach was tested through three protocols: human *VR* data, robot simulation and real-robot experiments. The *VR* dataset was randomly split into two disjoint sets: approximately 80% (19 subjects and 1710 trials) for training and 20% (5 subjects and 450 trials) for testing. This cross-subject splitting, (i.e. no overlap between training and testing data in terms of subjects) allowed us to test the generalization of the proposed approach with subjects who had never seen during training. This splitting was repeated for five




Fig. 4: A human-like manipulation plan for reaching a target amongst 8 obstacles. Note that the testing environment is different from the training one in terms of the number and shape of obstacles.

folds of randomly selected subjects for both sets. Standard (MATLAB) support vector machines with Gaussian kernel and Gaussian process regression were used to implement classifiers and regression models respectively. The same trained models and planning algorithm were used for all experiment protocols.

A. Experiments on VR Data

This protocol tested the *HLP* on unseen, but similarly distributed data to the training set. The objective was to measure the similarity of the generated plans to the “ground truth” human plans. This similarity was computed for a scene having N_R rows by the following metric:

$$s_{HLP} = \frac{1}{2N_R} \sum_{n=1}^{N_R} I(D_n)(I(D_n) + I(E_n)) \quad (6)$$

where $I(\cdot)$ is the indicator function which is 1 when its argument is true and 0 otherwise, D_n is a Boolean variable that is true if *HLP* action is same as human action at the n -th row and E_n is a Boolean variable that is true if the same element (gap or object) is selected by both *HLP* and human action. To illustrate, if both the *HLP* and the human participant decided to move an object at the first row and then go through a gap at the second row, then this would be quantified as a 50% similarity rate. This could increase to 100% if both selected the same object and the same gap.

Mean and standard deviation of the 5-fold results are shown in Table I. Accuracy of the gaps and objects classifiers are 95% and 85% respectively. It is worth noting that we compared similarity of the *HLP* output to a specific participant’s plan at a time and then reported the mean similarity. This means that *HLP* similarity is counted if it exactly matches a specific testing subject who was never seen during training. On average, our planner was 70% similar to the test participant plans. More specifically, the *HLP* decided the same action as a human plan 79% of the time, while it selected the same element (specific gap or object) 67% of the time.

B. Robot Experiments

Through robot experiments, we compared the *HLP* with a standard physics-based stochastic trajectory optimisation (*STO*) approach [27], [28], [29]. These algorithms were implemented using the Mujoco [30] physics engine and the Deepmind control suite [31]. We assumed a world consisting of objects on a table with a 10-DOF robot as shown in Fig. 4.

As a baseline algorithm, we used a state-of-the-art physics-based stochastic trajectory optimiser [27], initialized with a straight line control sequence from the end-effector’s start pose to the target object.

IK solutions for the high-level plan keypoints were found and connected with straight lines in robot’s configuration space to generate a control sequence. This was passed to a trajectory optimiser as an initial candidate control sequence. Therefore, for the *HLP*, the number of actions in a given control sequence varied depending on the high level plan. In contrast, the baseline approach (*STO*) initialized the trajectory optimiser using a straight line control sequence to the goal.

We compared performance of the *HLP* and *STO* quantitatively through success rates and planning times of simulation experiments and qualitatively through visual inspection of manipulation plans in real robot experiments.

1) *Robot Simulation Experiments*: Simulation experiments evaluated the performance of our planner in scenes generated from a distribution considerably different to that of the training data. Generalization was investigated across different numbers, sizes and types of objects. A plan was classified as being successful if at the final state the target object was inside the gripper’s pre-grasp region and if no other object dropped off the table during the trial. We recorded the total planning time for successful trials of each planner. Planning time which comprised an initial time used to generate the high level plan (straight line for *STO*), and an optimisation time. Two simulation categories were considered:

Generalization I– Performance of the *HLP* was tested in simulation scenes with the same two-row structure used in the *VR* demonstrations. Here, the generalization element involved varying the dimensions of the table and objects. Table II summarizes results of 100 scenes from this category. The planner *HLP* substantially outperformed the *STO* by a large margin, indexed through success rates and a reduction

TABLE I: Results (mean and standard deviation) of the 5-fold *VR* test experiment.

Metric	Mean	STD
C_g accuracy	0.95	0.002
C_o accuracy	0.85	0.005
s_{HLP} (overall)	0.70	0.011
$s_{HLP}(I(D_n))$	0.79	0.016
$s_{HLP}(I(E_n))$	0.67	0.012



Fig. 5: Human-like Manipulation (HLP) plan (top), and Baseline (STO) plan (bottom). HLP algorithm can reason over the task space and navigate through the free space around the obstacles similar to what people may do. On the other hand, STO is biased towards its straight-line initialization and hence had to push objects around which made it 1.5 times slower than HLP.

TABLE II: Generalization I Simulation Scenes Results

	Success rate(%)	Init. time(s)	Opt. time(s)	Total time(s)
HLP	94	0.59	0.97	1.56
STO	84	0.04	17.84	17.88

TABLE III: Planning times (mean) for Generalization II. Note that the actual planning time of HLP (Init. time) is approx. fixed irrespective of the number of obstacles.

No. of Objects	HLP			STO
	Init. time(s)	Opt. time(s)	Total(s)	Total(s)
5	0.60	1.70	2.30	1.85
7	0.61	2.65	3.26	3.68
9	0.63	5.90	6.53	4.98

in planning time.

Generalization II– Our second test involved investigating how our approach would fare when generalizing to different numbers of objects, different types of objects and different environment settings. We considered conditions with 5, 7 and 9 objects with two shape classes: cylinders and boxes. Importantly, the start and target positions were no longer aligned in this experiment. 100 randomly sampled scenes were generated for each object-number condition. For each random scene, we selected a random⁶ initial robot configuration.

The success rates for 100 random scenes for each of the three object categories were computed. The rates for both planners were relatively similar, (93%, 93% and 95%) for *HLP* and (96%, 96% and 98%) for *STO* for 5, 7 and 9 objects respectively. Results of planning time comparison are given in Table III. The time required for generating a high-level plan by *HLP* (Init. Time) is fixed irrespective of number of objects.

⁶We uniformly sampled a start point along the front edge of the table, and a corresponding random end-effector orientation (in the plane) and found the initial robot configuration using inverse kinematics

2) *Real Robot Experiments*: In a set of real robot experiments, we used a Robotiq two finger gripper attached to a UR5 arm mounted on an omnidirectional robot (Ridgeback). We obtained object positions using a depth camera mounted above the table. We conducted a limited number of experiments- 4 sample scenes per 7, 9 and 11 objects. We then ran both the *HLP* and the *STO* producing a total of 24 robot runs⁷. Sample results are shown in Fig. 5 where the *HLP* favored going through the free space avoiding pushing objects while the *STO* was biased to its initialized straight-line path. Other scenes can be found in the attached video⁸.

VIII. DISCUSSION

This work addresses the challenge of human-like computing i.e. the endowment of systems with capabilities derived from modelling human perception, cognition and action. Humans are known to recognize the spatio-temporal world in a qualitative manner. Hence, instead of cloning the human behavior from demonstrations, we used *QSR* in order to segment demonstrations in the action space. Extracting human skills at such a high level helps to model a planner that can: (1) generalize to different number of objects without increasing the actual planning time and (2) seamlessly connect with an arbitrary robot model.

Many cluttered environments can be clustered into regions that geometrically approximate to our definition of rows. For an arbitrary number of rows, the *HLP* can be run recursively for row pairs by defining a set of sub-targets. Moreover, generalization may be improved by augmenting training data with more generalized scenes, using more powerful classifiers and running in a closed-loop manner. These topics will be addressed in our future work.

⁷Four scenes, three number-of-object categories, and two methods.

⁸https://youtu.be/aMIZP_SYa0I

IX. CONCLUSIONS

We used VR to collect data from human participants whilst they reached for objects on a cluttered table-top. From this data, we devised a qualitative representation of task space to segment demonstrations into keypoints in the action space. State-action pairs were used to train decision classifiers that constituted the building blocks of the HLP algorithm. Through testing in VR, robot simulation and real robot experiments against a state-of-the-art planner, we have shown that this HLP is able to generate effective strategies for planning, irrespective of the number of objects in a cluttered environment. We conclude that the optimisation of robot planning has much to gain by extracting features from understanding the processes underlying human action selection and execution.

ACKNOWLEDGMENT

This work is funded by the EPSRC (EP/R031193/1) in their Human Like Computing initiative. Support was also received from EPSRC studentship 1879668 to the third author, the European Union's Horizon 2020 research and innovation programme under grant agreement 825619 (AI4EU), and Turing Fellowship from the Alan Turing Institute to the ninth author.

REFERENCES

- [1] J.-C. Latombe, *Robot motion planning*. Springer Science & Business Media, 2012, vol. 124.
- [2] J. E. King, J. A. Hausteine, S. Srinivasa, and T. Asfour, "Nonprehensile whole arm rearrangement planning on physics manifolds," in *ICRA*, 2015.
- [3] M. u. Din, M. Moll, L. Kavraki, J. Rosell *et al.*, "Randomized physics-based motion planning for grasping in cluttered and uncertain environments," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 712–719, 2017.
- [4] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The international journal of robotics research*, vol. 30, no. 7, pp. 846–894, 2011.
- [5] M. Elbanhawi and M. Simic, "Sampling-based robot motion planning: A review," *Ieee access*, vol. 2, pp. 56–77, 2014.
- [6] B. Ichter, J. Harrison, and M. Pavone, "Learning sampling distributions for robot motion planning," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 7087–7094.
- [7] R. Fox, R. Berenstein, I. Stoica, and K. Goldberg, "Multi-task hierarchical imitation learning for home automation," in *2019 IEEE 15th International Conference on Automation Science and Engineering (CASE)*. IEEE, 2019, pp. 1–8.
- [8] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2018, pp. 1–8.
- [9] M. Rana, M. Mukadam, S. R. Ahmadzadeh, S. Chernova, and B. Boots, "Towards robust skill generalization: Unifying learning from demonstration and motion planning," in *Intelligent robots and systems*, 2018.
- [10] H. Tan and K. Kawamura, "A computational framework for integrating robotic exploration and human demonstration in imitation learning," in *2011 IEEE International Conference on Systems, Man, and Cybernetics*. IEEE, 2011, pp. 2501–2506.
- [11] H. Ravichandar and A. Dani, "Learning contracting nonlinear dynamics from human demonstration for robot motion planning," in *ASME 2015 Dynamic Systems and Control Conference*. American Society of Mechanical Engineers Digital Collection, 2016.
- [12] M. Lawitzky, J. R. Medina, D. Lee, and S. Hirche, "Feedback motion planning and learning from demonstration in physical robotic assistance: differences and synergies," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 3646–3652.
- [13] J. O. Wallgrün, F. Dylla, A. Klippel, and J. Yang, "Understanding human spatial conceptualizations to improve applications of qualitative spatial calculi," in *27th International Workshop on Qualitative Reasoning*, 2013, p. 131.
- [14] J. Chen, A. G. Cohn, D. Liu, S. Wang, J. Ouyang, and Q. Yu, "A survey of qualitative spatial representations," *The Knowledge Engineering Review*, vol. 30, no. 1, pp. 106–136, 2015.
- [15] A. G. Cohn, J. Renz, and M. Sridhar, "Thinking inside the box: A comprehensive spatial representation for video analysis," in *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2012.
- [16] M. Mansouri and F. Pecora, "More knowledge on the table: Planning with space, time and resources for robots," in *2014 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2014, pp. 647–654.
- [17] M. Westphal, C. Dornhege, S. Wöflf, M. Gissler, and B. Nebel, "Guiding the generation of manipulation plans by qualitative spatial reasoning," *Spatial Cognition & Computation*, vol. 11, no. 1, pp. 75–102, 2011.
- [18] I. A. Şucan and L. E. Kavraki, "A sampling-based tree planner for systems with complex dynamics," *IEEE Transactions on Robotics*, vol. 28, no. 1, pp. 116–131, 2011.
- [19] R. Geraerts and M. H. Overmars, "A comparative study of probabilistic roadmap planners," in *Algorithmic Foundations of Robotics V*. Springer, 2004, pp. 43–57.
- [20] S. Karaman, M. R. Walter, A. Perez, E. Frazzoli, and S. Teller, "Anytime motion planning using the rt," in *2011 IEEE International Conference on Robotics and Automation*. IEEE, 2011, pp. 1478–1483.
- [21] J. Brookes, M. Warburton, M. Alghadier, M. Mon-Williams, and F. Mushtaq, "Studying human behavior with virtual reality: The unity experiment framework," *Behavior research methods*, pp. 1–9, 2019.
- [22] B. Akgun, M. Cakmak, K. Jiang, and A. L. Thomaz, "Keyframe-based learning from demonstration," *International Journal of Social Robotics*, vol. 4, no. 4, pp. 343–355, 2012.
- [23] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," in *2011 IEEE international conference on robotics and automation*. IEEE, 2011, pp. 723–730.
- [24] B. Moldovan, P. Moreno, D. Nitti, J. Santos-Victor, and L. De Raedt, "Relational affordances for multiple-object manipulation," *Autonomous Robots*, vol. 42, no. 1, pp. 19–44, 2018.
- [25] E. Alpaydin, *Introduction to machine learning*. MIT press, 2009.
- [26] J. Langford and B. Zadrozny, "Reducing t-step reinforcement learning to classification," 2003.
- [27] W. C. Agboh and M. R. Dogar, "Pushing fast and slow: Task-adaptive planning for non-prehensile manipulation under uncertainty," in *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2018.
- [28] W. C. Agboh and M. R. Dogar, "Real-time online re-planning for grasping under clutter and uncertainty," in *IEEE-RAS 18th International Conference on Humanoid Robots (Humanoids)*, 2018.
- [29] W. C. Agboh, D. Ruprecht, and M. R. Dogar, "Combining coarse and fine physics for manipulation using parallel-in-time integration," *International Symposium on Robotics Research (ISRR)*, 2019.
- [30] E. Todorov, T. Erez, and Y. Tassa, "Mujoco: A physics engine for model-based control," in *IROS*. IEEE, 2012.
- [31] Y. Tassa, Y. Doron, A. Muldal, T. Erez, Y. Li, D. de Las Casas, D. Budden, A. Abdolmaleki, J. Merel, A. Lefrancq, T. P. Lillicrap, and M. A. Riedmiller, "Deepmind control suite," *CoRR*, vol. abs/1801.00690, 2018. [Online]. Available: <http://arxiv.org/abs/1801.00690>