

Fig. 2: Proposed framework for self-assessment of grasp affordance transfer. After predicting a grasp affordance region, the most suitable grasp is determined based on a library of prior task affordance experiences and a heuristic confidence metric.

II. RELATED WORK

Understanding grasp affordances for objects has been an active area of research for robotic manipulation tasks. Ideally, an autonomous agent should be able to identify all the tasks that an object can afford, and infer the grasp configuration that leads to a successful completion of each task. A common approach to tackle this challenge is via visual features, e.g. [4]–[7]. Methods based on visual grasp affordance detection identify candidate grasps either via deep learning architectures that detect grasp areas on an object [4]–[6], or via supervised learning techniques that obtain grasping configurations based on an object’s shape [7]. While these techniques offer robust grasp candidates, they uniquely seek grasp stability. Consequently, these methods cannot guarantee any level of performance when executing a task, and in fact, not even a successful task completion. In order to, move towards reliable task deployment on autonomous agents, there is the need to bridge the gap between grasp affordance detection and task-oriented grasping.

Grasp affordances: Work on grasp affordances aims at robust interactions between objects and the autonomous agent. However, it is typically limited to a single grasp affordance detection per object, thus reducing its deployment in real-world scenarios. Some works, such as [14], focus on relating abstractions of sensory-motor processes with object structures (e.g., object-action complexes (OACs)) to extract the best grasp candidate given an object affordance. Others use purely visual input to learn affordances using deep learning [6], [8] or supervised learning techniques to relate objects and actions [9], [15]–[17]. Although these works are successful in detecting grasp affordance regions, they hypothesise suitable grasp configurations based on visual features, rather than indicators that hint such proposals suitability to accomplish an affordance task.

Task affordances: The end goal of grasping is to manipulate an object to fulfil a goal-directed task. When the grasping problem is contextualised into tasks, solely satisfying the grasp stability constraints is no longer sufficient. Nonetheless, codifying grasp configurations with task success is still an open problem. Along this line, some works focus entirely on learning tasks where the object category does

not influence the outcome, such as pushing or pulling [15], [17]. Hence, reliable extraction of grasp configurations is neglected. Another approach is to learn grasp quality measures for task performance via trial and error [11]–[13]. Based on the experiences, these studies build semantic constraints to specify which object regions to hold or avoid. Nonetheless, their dependency on great amounts of prior experiences and the lack of generalisation between object instances remain to be the main hurdle of these methods.

Our work seeks to bridge the gap between grasp affordances and task performance existing in prior work. The proposed approach unifies grasp affordance reasoning and task deployment in a self-assessed system that, without the need for extensive prior experiences, is able to transfer grasp affordance configurations to novel object instances.

III. PROPOSED METHOD

An autonomous agent must be able to perform a task affordance in different scenarios. Given a particular object and task \mathcal{T} to perform, the robot must select a suitable grasp affordance configuration g_p^* that allows executing the task’s policy π_τ successfully. Only the correct choice of both g_p^* and π_τ leads to the robot being successful at addressing the task affordance problem. Despite the strong correlation between g_p^* and the π_τ execution performance, current approaches in the literature consider these elements to be independent. This results in grasping configurations that are not suitable for completing the task.

In this section, we introduce our approach to self-assess the selection of a suitable grasp affordance configuration according to an estimate of the task performance. Fig. 2 illustrates the proposed pipeline which (i) detects from visual information a set of grasping candidates lying in the object’s grasp affordance space (Section III-A), (ii) exploits a learnt library of task affordance policies to forward simulate the outcome of executing the task from the grasping candidates (Section III-B), and then (iii) evaluates the grasp configuration candidates subject to a heuristic confidence metric (Section III-C) which allows for one-shot transfer of the grasp proposal (Section III-D). Finally, in Section III-E, we detail how these components fit in the scheme of a robotic agent dealing with task affordance problems autonomously.

A. Prediction of Grasp Affordance Configurations

The overall goal of this work is, given an object’s grasp affordance region G^* , to find a grasp configuration g_p^* that allows the robot to successfully employ an object for a particular task. In the grasp affordance literature, it is common to visually detect and segment the grasp affordance region G^* using mapping to labels [6], [8], [9]. While these methods all predict g_p^* via visual detection hypotheses, none estimate the configuration proposals based on a task performance insight. This relational gap endangers a successful task execution. Instead, an autonomous agent should be capable of discerning the most suitable grasp that benefits the execution of a task.

To bridge this gap, in our method we consider a grasp affordance region G^* in a generic form such as the bounding box provided by [9] (see Fig. 3a). We are interested in pruning this region by finding multiple grasp proposal candidates. With this aim, we use the pre-trained DeepGrasp model [5], a deep CNN that computes reliable grasp configurations on objects. The output grasp proposals g_{p_i} from DeepGrasp, which do not account for affordance relation, are shown in Fig. 3b. The pruned region (see Fig. 3c), denoted as $g_{p_i} \in G^*$, provides a set of grasp configuration candidates that accounts for both reliability and affordability.

B. Library of Task Affordances

The success of an affordance task \mathcal{T} lies in executing the corresponding task policy π_τ from a suitable grasp configuration g_p^* . This is a difficult problem given that the π_τ and g_p^* are codefining [10]. Namely, the task’s requirements constrain the possibly suitable grasp configurations g_p^* , at the same time that the choice of g_p^* conditions the outcome of executing the task’s policy π_τ . Additionally, determining whether the execution of a task is successful requires a performance indicator. To cope with this challenge, we build on our previous work [18] to learn a library \mathcal{L} of task affordances from human demonstrations. The library aims at simultaneously guiding the robot on the search of a suitable task policy π_τ while informing about its expected outcome α_τ when successful. All these elements serve as the basis of the method described in Section III-C to determine g_p^* via self-assessment of the candidates $g_{p_i} \in G^*$.

In this work, we build the library of task affordances as:

$$\mathcal{L} = \{ \mathcal{T}_1 \rightarrow \{ \pi_{\tau_1}, A_{\tau_1} \}, \dots, \mathcal{T}_n \rightarrow \{ \pi_{\tau_n}, A_{\tau_n} \} \}, \quad (1)$$

where π_τ is a policy encoding the task in a generalisable form, and $\alpha_\tau \in A_\tau$ is a set of possible successful outcomes when executing π_τ . In our implementation, π_τ is based on dynamic movement primitives (DMPs) [19], [20]. DMPs are differential equations encoding behaviour towards a goal attractor. We initialise the policies via imitation learning, and use them to reproduce an observed motion while generalising to different start and goal locations, as well as task durations.

Regarding the set of possible successful outcomes $\alpha_\tau \in A_\tau$, we provide the robot with multiple experiences. We define the outcome α_τ as the state evolution of the object’s action region S_O through the execution of the task. We employ mask RCNN (M-RCNN) [21] to train a




Fig. 3: Prediction of grasp affordance configurations for the pouring task. (a) Patch affording the pouring task, (b) reliable grasp configurations from DeepGrasp, (c) pruned space for reliable grasp candidates that afford the task pouring.

model that detects objects subparts as action regions S_O . As exemplified in Fig. 4, the action region state provides a meaningful indicator of the task. This information is used as the basis for our confidence metric, which evaluates the level of success of an affordance task for a grasping proposal.

C. Search-Based Self-Assessment of Task Affordances

The task policies π_τ learnt in Section III-B allow a previously experienced task from any candidate grasp $g_{p_i} \in G^*$ to be performed. Nonetheless, executing π_τ from any grasp configuration may not always lead to suitable performance. For example, Fig. 4 depicts the case where grasping the mug from g_{p_1} prevents the robot from performing a pouring task as adequately as when grasping it from g_{p_2} .

We propose to self-assess the outcome of executing the task’s policy π_τ from $g_{p_i} \in G^*$ before deciding the most

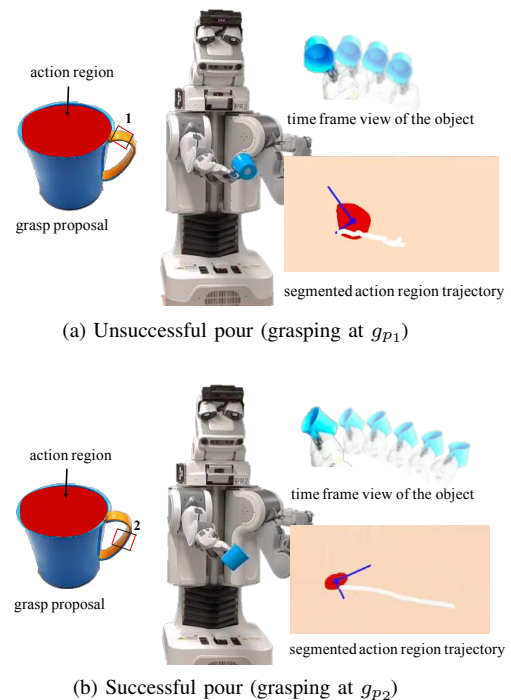


Fig. 4: Example of a pouring task from two different grasp configurations. Each situation illustrates the raw two-dimensional (2-D) camera input of the object and the segmented action region that affords the pouring task.

suitable grasp configuration g_p^* on a new object. This is efficiently done by forward simulation of the DMP-encoded π_τ . From each roll-out, we look at the object’s state action region $\bar{\alpha}_\tau$ as a suitable task performance indicator. To this aim, we consider the entropy between the demonstrated successful task outcomes α_τ and the simulated outcome $\bar{\alpha}_\tau$ in the form of Kullback-Leibler divergence [22]:

$$D(\alpha_\tau || \bar{\alpha}_\tau) = \sum_{i \in I} \alpha_\tau(i) \log \left(\frac{\alpha_\tau(i)}{\bar{\alpha}_\tau(i)} \right), \quad (2)$$

which results in a low penalisation when the forward simulated outcome $\bar{\alpha}_\tau$ is similar to a previously experienced outcome in A_τ , and a high penalisation otherwise. Then, we propose to rank the grasping candidates $g_{p_i} \in G^*$ according to a confidence metric which estimates the suitability of a candidate g_{p_i} for a given \mathcal{T} as:

$$C(g_{p_i}) = \max_{\alpha_\tau \in A_\tau} D^{-1}(\alpha_\tau || \bar{\alpha}_\tau). \quad (3)$$

Finally, we select the grasping configuration g_p^* among all grasping candidate $g_{p_i} \in G^*$ as:

$$g_p^* = \arg \max_{g_{p_i} \in G^*} C(g_{p_i}) \quad s.t. \quad C(g_{p_i}) > \delta, \quad (4)$$

which returns the grasp configuration with highest confidence of successfully completing the task. This assessment is subject to a minimum user-defined confidence level δ that rejects under-performing grasp configuration proposals. As explained in the experimental setup, such a threshold is adjusted from demonstration by a binary classifier.

D. One-Shot Self-Assessment of Task Affordances

The search-based strategy presented in Section III-C in the grasp affordance region can be time and resource consuming if performed for every single task affordance problem. Alternatively, we propose to augment the library in (1) with an approximate of the prior experienced outcomes α_τ per grasp configuration g_{p_i} , such that it allows for one-shot assessment. Namely, we extract the spatial transform of all experienced grasps with respect to the detected grasp affordance region G^* . The relevance of these transforms is ranked in a list R according to their confidence score computed following (3). Therefore, the augmented library is denoted as:

$$\mathcal{L} = \{ \mathcal{T}_1 \rightarrow \{ \pi_{\tau_1}, A_{\tau_1}, R_{\tau_1} \}, \dots, \mathcal{T}_n \rightarrow \{ \pi_{\tau_n}, A_{\tau_n}, R_{\tau_n} \} \}. \quad (5)$$

At deployment time, we look at the spatial transform from the new grasping candidates that resembles the most well-ranked transform in R . This allows us to hierarchically self-assess the candidates by order of prospective success.

E. Deployment on Autonomous Agent

Algorithm 1 presents the outline of SAGAT’s end-to-end deployment, which aims at improving the success of an autonomous agent when performing a task. Given visual perception of the environment, the desired affordance, the pre-trained model to extract the grasp affordance relation (see Section III-A), the model to detect the action region,

Algorithm 1: deployment of SAGAT

```

1 Input:
2   CVF: camera visual feed
3   affordance: affordance choice
4   graspAffordance: grasp affordance model
5   actionRegion: MRCNN learnt model
6   libTaskAffordances: task affordance library
7 Output:
8    $g_p^*$ : most suitable grasp affordance configuration
9 begin
10   $G^* \leftarrow$  graspAffordance(CVF, affordance)
11   $S_O \leftarrow$  actionRegion(CVF, affordance)
12   $g_p \leftarrow$  libTaskAffordances( $G^*$ , affordance)
13  while not isEmpty( $g_p$ ) do
14     $g_{p_i} \leftarrow$  popHighestConfidence( $g_p$ )
15     $\bar{\alpha}_\tau \leftarrow$  forwardSimulateTask( $g_{p_i}$ ,  $S_O$ )
16    if prospectiveTaskSuccess( $\bar{\alpha}_\tau$ ) then
17      return  $g_{p_i}$ 
18  return none

```

and the learnt library of task affordances (see Section III-B to Section III-D) (lines 2 to 6), the end-to-end execution is as follows. First, the visual data is processed to extract the grasp affordance region (line 10) and the object’s action region (line 11). The resulting grasp affordance region along with the desired affordance are used to estimate the grasp configuration proposals on the new object using the library of task affordances as prior experiences (line 12). The retrieved set of grasp configuration candidates is analysed in order of decreasing prospective success (line 13 to line 17) until either exhausting all candidates or finding a suitable grasp for the affordance task. Importantly, the hierarchy of the proposed self-assessment analysis allows for one-shot transfer of the grasp configuration proposals, i.e. to find, on the first trial, a suitable grasp affordance by analysing the top-ranked grasp candidate. Nonetheless, the method also considers the case that exhaustive exploration of all candidates might be required, thus ensuring algorithmic completeness.

Notably, the proposed method is not dependant on a particular grasp affordance or action region description. This modularity allows the usage of the proposed method in a wide range of setups. We demonstrate the generality of the proposed method by first, using multiple state-of-the-art approaches for grasp affordance detection, and then, determining the improvement on task performance and deployability when used altogether with our approach.

IV. EXPERIMENTAL EVALUATION AND DISCUSSION

The proposed methodology endows a robot with the ability to determine a suitable grasp configuration to succeed on an affordance task. Importantly, such a challenge is addressed without the need for extensive prior trials and errors. We demonstrate the potential of our method following the experimental setup described in Section IV-A and a thorough eval-

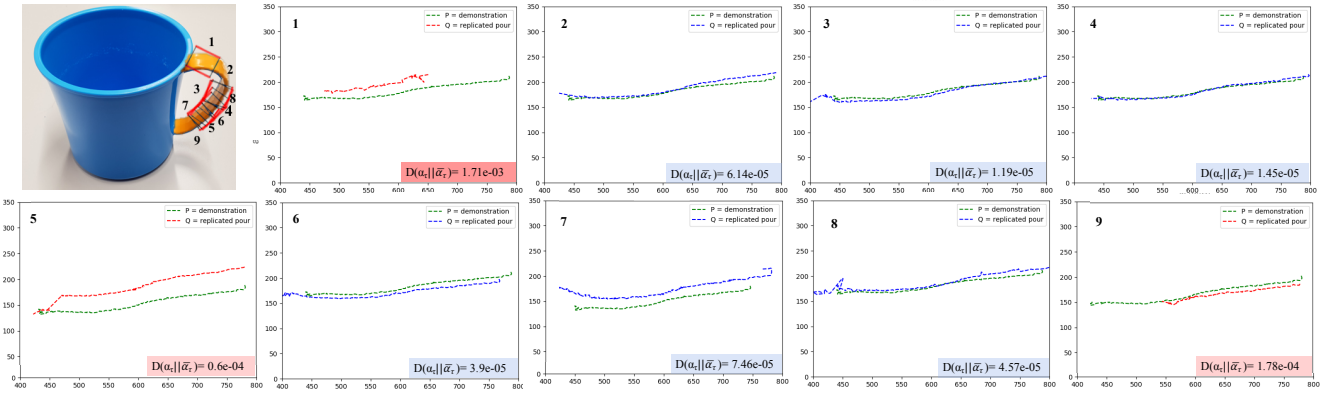


Fig. 5: Entropy measurements on the 2-D frame for the pouring task. We consider as reference a socially acceptable pouring demonstration (green) against successful (blue) and undesired (red) task repetitions from different grasp candidates. The candidates are numbered with the corresponding observed effect. Successful tasks present low entropy whereas undesired effects have higher entropy. Our proposal exploits this relation to discern among grasp candidates at deployment time.

uation based on the following tests: (i) the spatial similarity between learnt and computed configurations across objects (Section IV-B), (ii) the accuracy of the task affordance deployment when transferred to new objects (Section IV-C), and (iii) the performance of our proposal when compared to other methodologies (Section IV-D).

A. Experimental Setup

The end-to-end execution framework presented in Algorithm 1 is deployed on a PR2 robotic platform, in both simulated and real-world scenarios. We use a Kinect mounted on the PR2’s head as our visual sensor and the position sensors on the right arm joints to encode the end-effector state pose for learning the task policies in the library.

We evaluate the proposed approach with an experimental setup that considers objects with variate affordable actions and suitable grasping configurations. Particularly, the library of task affordances is built uniquely using the blue mug depicted in Fig. 5, but evaluated with the objects depicted in Fig. 6. As can be observed, the training and testing sets present a challenging and significant variability on the grasp affordance relation. Our experimental setup also considers multiple affordances, namely: pouring, handover and shaking. The choice of these affordances is determined by those being both common among the considered objects and socially acceptable according to [9].

The task policy and its expected effect corresponding to each affordance are taught to the robot via kinaesthetic demonstration. The end-effector state evolution is used to learn the task policy in form of a set of DMPs, and the state evolution of the container’s action region segmented on the 2-D camera frame to learn the expected effect. As depicted in Fig. 5 for the pouring task, the learnt policy is replicated 9 times from different grasping candidates, including suitable grasp affordances (blue) and undesired deployments (red).

The collected demonstrations are used to adjust the confidence threshold in (4) via a binary classifier, where the confidence level computed following (3) is the support, and

the label {“successful”, “undesired”} is the target. Only successful deployments are included in the library.

B. Spatial Similarity of Grasp Configurations

Our method allows the system for one-shot transfer of grasp configurations to new objects. As explained in Section III-D, we rank the grasp candidates on new objects as those that closely resemble the experiences stored in the library of task affordances. This approximation is based on the expectation that similar spatial configurations should offer similar performance when dealing with the same task. In this set of experiments, we demonstrate the validity of such a hypothesis by evaluating the spatial similarity between the proposals estimated on new objects and the ones previously identified as suitable and stored in the library.

For an object, we calculate the Euclidean distance between the segmented action region S_O and the obtained grasp configuration g_p^* . Fig. 7 shows the obtained distances denoted as $d_h(S_O, g_p^*)$. The blue horizontal line represents the mean distance obtained during the demonstrations. Overall, we observe similar distances from action regions to grasp configurations across objects. For dissimilar cases such as 4 and 5 (ashtray and bowl respectively), the difference is



Fig. 6: Novel objects to test the self-assessed grasp transfer.




Fig. 7: Visualisation of the dissimilarity metric between an object’s action region and the corresponding suitable grasp configuration, in comparison to the mean dissimilarity observed during the demonstrations (d_{hd} , blue horizontal line).

given by the fact that the obtained grasping region for most of the tasks lies on the edges of the object compartment. Even though these grasping configurations are relatively close to the action region, we will see on Table I that the average performance of the tasks is preserved.

To further evaluate similarity across obtained grasping configurations, we are also interested in how much the system prunes the grasping space based on the information stored in the library. As defined in (4), we use a confidence threshold for the pruning process of the grasping space. Thus, based on the prior of well-performing grasp configurations, highly dissimilar proposals are not considered on the self-assessed transfer process. Fig. 8 depicts the rejection rate of grasp configuration proposals per task affordance. From the plot, we see that the pouring task shows the highest rejection rate, especially for objects that have handles. This hints that for this task the grasping choice is more critical.

C. One-Shot Transfer of Task Affordances

The second experimental test analyses the performance of our method when addressing task affordances on new

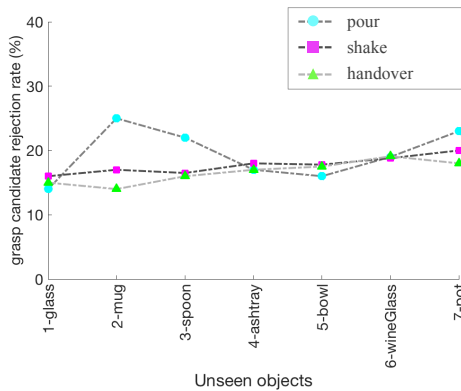


Fig. 8: Rejection rate of grasp candidates with prospective unsuccessful task deployment. Grasp configurations, as extracted with DeepGrasp [5], that do not relate to the prior on successful task deployment, as stored in the library, are rejected in the one-shot transfer scheme.

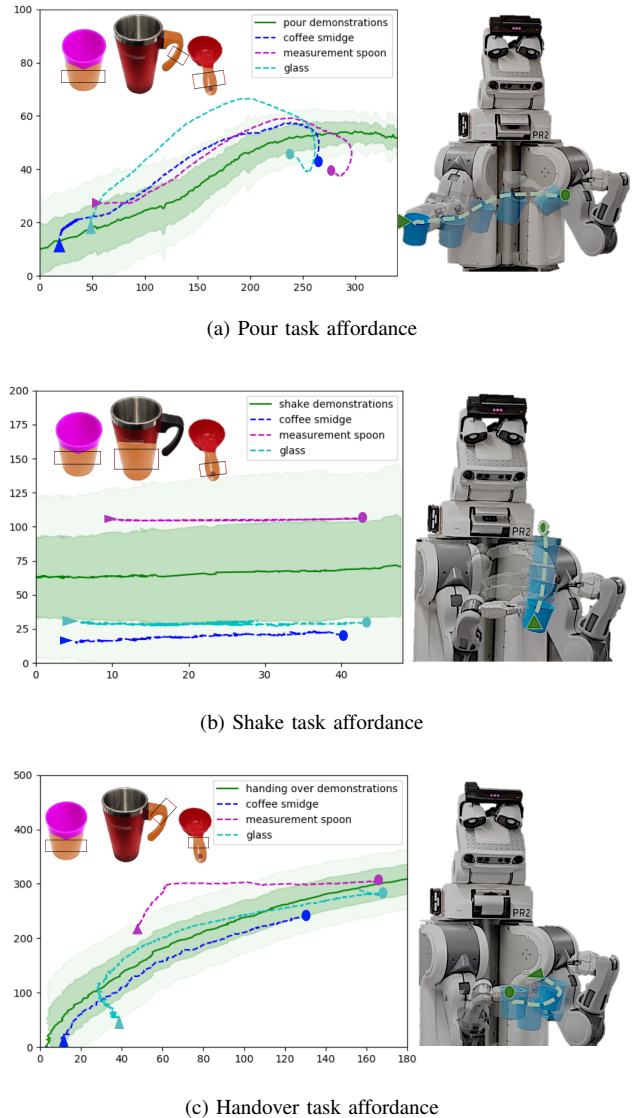


Fig. 9: Task affordance performance when deployed on novel objects (colour-coded lines) in comparison with the multiple successful demonstrations (green scale distribution).

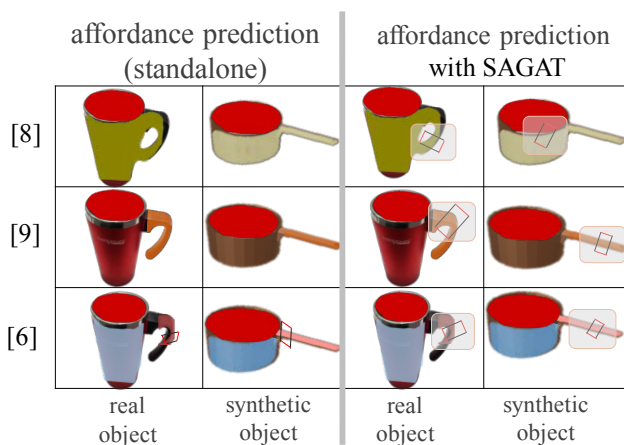


Fig. 10: Comparison of grasp affordance detection for the task of pouring with state-of-the-art methods and SAGAT. The resulting grasp configuration proposals obtained with SAGAT are highlighted for better visualisation.

objects. The goal of this evaluation is to determine if the chosen grasp configuration enables objects to perform the task affordance as successfully as the prior stored in the library. Fig. 9 depicts the mean and variance (green scale) of the prior experiences in the library for the tasks pour, shake and handover. Each task was performed with three real objects with notably different features: a travel mug (dark blue), measurement spoon (magenta) and a glass (blue). The resulting effect when performing the tasks from the computed grasping configuration is colour-coded on top of the prior experiences distribution.

Subject to the task affordance, the three objects show different grasp affordance regions. After the one-shot self-assessment procedure, the computed grasp configurations are the most spatially similar to the most successful grasp configuration in the experience dataset. Importantly, as shown in Fig. 9, this strategy is invariant to the initial and final states of the task. This is reflected in the obtained task affordance effect, which falls inside the variance of the demonstrations.

D. Comparison of Task Deployment Reliability

The last experimental test is to demonstrate at which level the proposed method enhances the task deployment reliability when used in conjunction with methods for grasp affordance detection [6], [8], [9]. To conduct this evaluation, we use the open-source implementations of [6], [8], [9] on all objects illustrated in Fig. 6, in the real and simulated robotic platform. The obtained grasp regions are used to execute the task in two different ways: (i) in stand-alone fashion, i.e. as originally proposed, and (ii) as input of our SAGAT approach to determine the most suitable grasp candidate. Fig. 10 shows some examples of the grasp affordance detected with the previously mentioned methods and our approach.

We use the policies in the learnt library of task affordances to replicate the pour, shake and handover tasks on each object, for each grasp affordance, and for each method when used as stand-alone and combined with SAGAT. This results

	[8]	[8]+SAGAT	[9]	[9]+SAGAT	[6]	[6]+SAGAT
Pour	70%	82%	72%	83%	73%	85%
Shake	84%	87%	85%	87%	86%	88%
Handover	80%	85%	81%	86%	82%	86%

TABLE I: Comparison of success rates on task affordance deployment when using state-of-the-art grasp affordance extractors as stand-alone and with our method.

in a total of 126 tasks deployments on the robotic platform¹. Table I summarises the obtained results. As can be observed, deploying a task using state-of-the-art methods on grasp affordance detection provides an average success rate of 79.2% across tasks. With our approach, the deployability success is enhanced for all the tasks, with an average rate of 85.4%. Interestingly, the 5.2% improvement is not equally distributed across tasks; more challenging tasks experience a higher success rate. This is the case of the pouring tasks where deployability success is increased by 11.67%.

V. CONCLUSIONS AND FUTURE WORK

In this paper, we presented a novel experience-based pipeline for self-assessment of grasp affordance transfer (SAGAT). Our approach enhances the deployment reliability of current state-of-the-art methods on grasp affordance detection, by extracting multiple grasp configuration candidates from a given grasp affordance region. The outcome of executing a task from different grasp candidates is estimated via forward simulation. These estimates are evaluated and ranked via a heuristic confidence function in relation to task performance and grasp configuration candidates. Such information is stored in a library of task affordances, which serves as a basis for one-shot transfer estimation to identify grasp affordance configurations similar to those previously experienced, with the insight that similar regions lead to similar deployments of the task. We evaluate the method's efficacy on novel task affordance problems by training on a single object and testing on multiple new ones. We observe a significant performance improvement up to approximately 11.7% in our experiments when using our proposal in comparison to state-of-the-art approaches on grasp affordance detection. Experimental evaluation on a PR2 robotic platform demonstrates highly reliable deployability of the proposed method to deal with real-world task affordance problems.

This work encourages multiple interesting directions for future work. Our follow-up work will study a unified probabilistic framework to infer the most suitable grasp affordance candidate. We envision that this will allow sets of actions and grasps to be predicted when dealing with multiple correlated objects in the scene. Consequently, including a task planning layer that connects actions with grasp affordances. Another interesting extension is the assessment of the end-state comfort-effect for grasping in human-robot collaboration tasks, such that the robot's grasp affordance considers the human's grasp capabilities.

¹A compilation of experiments can be found in: https://youtu.be/nCCc3_Rk8Ks

VI. ACKNOWLEDGEMENTS

This paper was written while Paola Ardón and Èric Pairet were on research visits to the Human-Centred Robotics lab at University of Washington (Seattle, WA, USA), and Kavraki's Lab at Rice University (Houston, TX, USA), respectively. The authors thank both groups for their support.

REFERENCES

- [1] L. Jamone, E. Ugur, A. Cangelosi, L. Fadiga, A. Bernardino, J. Piater, and J. Santos-Victor, "Affordances in psychology, neuroscience, and robotics: A survey," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 10, no. 1, pp. 4–25, 2018.
- [2] H. Min, C. Yi, R. Luo, J. Zhu, and S. Bi, "Affordance research in developmental robotics: a survey," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 8, no. 4, pp. 237–255, 2016.
- [3] J. Gibson, "The theory of affordances," in *Perceiving, Acting, and Knowing: Toward and Ecological Psychology* (R. Shaw and J. Bransford, eds.), pp. 62–82, Hillsdale, NJ: Erlbaum, 1977.
- [4] I. Lenz, H. Lee, and A. Saxena, "Deep learning for detecting robotic grasps," *International Journal of Robotics Research*, vol. 34, no. 4-5, pp. 705–724, 2015.
- [5] F.-J. Chu, R. Xu, and P. A. Vela, "Real-world multiobject, multigrasp detection," *IEEE Robotics and Automation Letters*, vol. 3, no. 4, pp. 3355–3362, 2018.
- [6] F.-J. Chu, R. Xu, and P. A. Vela, "Learning affordance segmentation for real-world robotic manipulation via synthetic images," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1140–1147, 2019.
- [7] J. Bohg and D. Kragic, "Learning grasping points with shape context," *Robotics and Autonomous Systems*, vol. 58, no. 4, pp. 362–377, 2010.
- [8] T.-T. Do, A. Nguyen, and I. Reid, "Affordancenet: An end-to-end deep learning approach for object affordance detection," in *International Conference on Robotics and Automation (ICRA)*, 2018.
- [9] P. Ardón, È. Pairet, R. P. Petrick, S. Ramamoorthy, and K. S. Lohan, "Learning grasp affordance reasoning through semantic relations," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 4571–4578, 2019.
- [10] L. Montesano, M. Lopes, A. Bernardino, and J. Santos-Victor, "Learning object affordances: From sensory-motor coordination to imitation," *IEEE Trans. Robotics*, vol. 24, pp. 15–26, 2008.
- [11] K. Fang, Y. Zhu, A. Garg, A. Kurenkov, V. Mehta, L. Fei-Fei, and S. Savarese, "Learning task-oriented grasping for tool manipulation from simulated self-supervision," *The International Journal of Robotics Research*, p. 0278364919872545, 2019.
- [12] A. Mandlekar, Y. Zhu, A. Garg, J. Booher, M. Spero, A. Tung, J. Gao, J. Emmons, A. Gupta, E. Orbay, *et al.*, "Roboturk: A crowdsourcing platform for robotic skill learning through imitation," in *Conference on Robot Learning*, pp. 879–893, 2018.
- [13] O. Kroemer, E. Ugur, E. Oztop, and J. Peters, "A kernel-based approach to direct action perception," in *2012 IEEE International Conference on Robotics and Automation*, pp. 2605–2610, IEEE, 2012.
- [14] N. Krüger, C. Geib, J. Piater, R. Petrick, M. Steedman, F. Wörgötter, A. Ude, T. Asfour, D. Kraft, D. Omrčen, *et al.*, "Object-action complexes: Grounded abstractions of sensory-motor processes," *Robotics and Autonomous Systems*, vol. 59, no. 10, pp. 740–757, 2011.
- [15] D. Song, K. Huebner, V. Kyrki, and D. Kragic, "Learning task constraints for robot grasping using graphical models," in *Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ International Conference on*, pp. 1579–1585, IEEE, 2010.
- [16] L. Montesano and M. Lopes, "Learning grasping affordances from local visual descriptors," in *Development and Learning, 2009. ICDL 2009. IEEE 8th International Conference on*, pp. 1–6, IEEE, 2009.
- [17] B. Moldovan, P. Moreno, M. van Otterlo, J. Santos-Victor, and L. De Raedt, "Learning relational affordance models for robots in multi-object manipulation tasks," in *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pp. 4373–4378, IEEE, 2012.
- [18] È. Pairet, P. Ardón, M. Mistry, and Y. Petillot, "Learning and composing primitive skills for dual-arm manipulation," in *Annual Conference Towards Autonomous Robotic Systems*, pp. 65–77, Springer, 2019.
- [19] A. J. Ijspeert, J. Nakanishi, H. Hoffmann, P. Pastor, and S. Schaal, "Dynamical movement primitives: learning attractor models for motor behaviors," *Neural computation*, vol. 25, no. 2, pp. 328–373, 2013.
- [20] È. Pairet, P. Ardón, M. Mistry, and Y. Petillot, "Learning generalizable coupling terms for obstacle avoidance via low-dimensional geometric descriptors," *IEEE Robotics and Automation Letters*, vol. 4, no. 4, pp. 3979–3986, 2019.
- [21] K. He, G. Gkioxari, P. Dollár, and R. Girshick, "Mask r-cnn," in *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.
- [22] F. Pérez-Cruz, "Kullback-leibler divergence estimation of continuous distributions," in *2008 IEEE international symposium on information theory*, pp. 1666–1670, IEEE, 2008.