

Speeding up 6-DoF Grasp Sampling with Quality-Diversity

Johann Huber^{*1}, François Héli on^{*1}, Mathilde Kappel^{*1},
 Elie Chelly¹, Mahdi Khoramshahi¹, Faiz Ben Amar¹ and St ephane Doncieux¹

Abstract—Recent advances in AI have led to significant results in robotic learning, including natural language-conditioned planning and efficient optimization of controllers using generative models. However, the interaction data remains the bottleneck for generalization. Getting data for grasping is a critical challenge, as this skill is required to complete many manipulation tasks. Quality-Diversity (QD) algorithms optimize a set of solutions to get diverse, high-performing solutions to a given problem. This paper investigates how QD can be combined with priors to speed up the generation of diverse grasps poses in simulation compared to standard 6-DoF grasp sampling schemes. Experiments conducted on 4 grippers with 2-to-5 fingers on standard objects show that QD outperforms commonly used methods by a large margin. Further experiments show that QD optimization automatically finds some efficient priors that are usually hard coded. The deployment of generated grasps on a 2-finger gripper and an Allegro hand shows that the diversity produced maintains sim-to-real transferability. We believe these results to be a significant step toward the generation of large datasets that can lead to robust and generalizing robotic grasping policies.

I. INTRODUCTION

Grasping is a skill of great interest in robotics as it is a prerequisite for many manipulation tasks [1]. The previously prevalent analytical-based methods [2] are gradually giving way to data-driven strategies since the beginning of the 21st century [3]. However, the challenging exploration aspect of grasping hinders the bootstrapping of the learning process, as most grasps attempted by a policy initialized at random yield no reward [4]. Many studies have addressed this issue by employing imitation learning [5]–[7], parallel grippers [8], [9], and top-down movements [10]–[12]. However, these approaches restrict the operational space, constraining the policies’ adaptability.

The recent advances in robotic learning demonstrate the capabilities of data-based approaches for skill acquisition [13]–[17]. These results rely on modern artificial intelligence methods that require a tremendous amount of high-quality data to generalize to unknown scenes. It has led to the release of large datasets [18], many of them focusing on grasping [9], [19], [20]. Acquiring such high-quality datasets is becoming critical to allow high representational power architecture to achieve generalization [13] or to build efficient foundation models [14] to transfer skills between platforms and scenarios.

^{*} equal contribution and corresponding authors

¹Sorbonne Universit , CNRS, Institut des Syst mes Intelligents et de Robotique, ISIR, F-75005 Paris, France {huber, helenon, kappel, chelly, khoramshahi, benamar, doncieux}@isir.upmc.fr

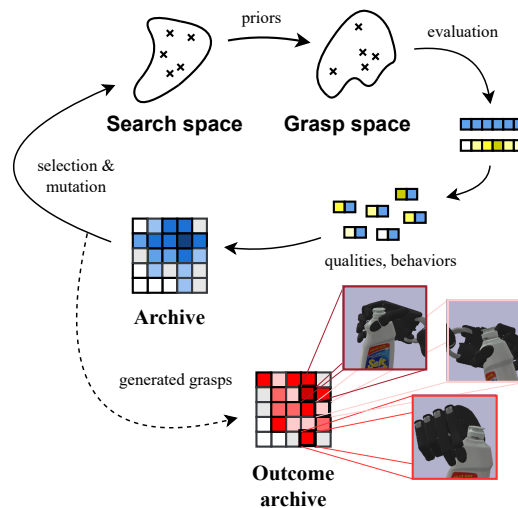


Fig. 1. **Overview of the proposed framework.** It consists of a population-based algorithm that allows the efficient exploration of the space of possible grasp positions from a genotype space. The optimization process is driven by searching for diverse and high-performing grasps using a structured archive of previously found solutions [28]. The genotype space is designed to leverage the robotic priors that are commonly used in 6-DoF grasp sampling [23]. All grasps ever produced are added to an outcome archive, which is the output of the algorithm.

Great efforts have been made to build large sets of real data [9], [18], but the acquisition of such data is very slow and expensive. Recent works leverage simulated scenes to speed up the data generation process [4], [19], [20] while assuring the sim-to-real transfer through dedicated quality criteria [21].

The advent of data-driven approaches has led to the formalization of grasping as a 6 Degrees-of-Freedoms (6-DoF) pose synthesis problem that can be predicted by a trained model [22]. Recently, Eppner et al. [23] conducted a systematic study to identify the best grasp sampling schemes among the most commonly used ones in the literature. In the Evolutionary Algorithms field, Huber et al. [4] studied how Quality-Diversity (QD) algorithms can be applied to generate diverse datasets of robust reach-and-grasp trajectories.

This paper shows that QD methods can significantly speed up the 6-DoF grasp sampling. In particular:

- We propose a framework that combines robotic priors with QD algorithms to generate large sets of diverse and robust grasps;
- We show that variants based on this framework outperform state-of-the-art 6-DoF grasp synthesis schemes on 4 grippers and a dozen of standard objects in simulation;

- Experiments conducted on a physical Franka Emika Panda gripper and an Allegro hand show that the proposed method generate grasps that successfully transfer into the real world while preserving diversity.

The code had been made publicly available¹. More details can be found on the project website².

II. RELATED WORKS

Learning to Grasp in Robotics. Many paradigms have been explored since the advent of data-driven approaches in the field, including reinforcement learning [24], [25], and learning from a few demonstrations [6], [7]. However, these methods require constraining the search space to make the problem tractable. Many data-greedy approaches were proposed to get generalization capabilities [9]–[11]. Recently were leveraged Diffusion models [15], Variational autoencoders [16], or neural representations [17] to sample grasps poses on unknown objects. Interestingly, most of them exploit automatically generated data in their learning process.

Automatic Generation of Grasping Datasets. Datasets for robotic learning can be directly collected in the real world [10], [18], but this process is very time and cost-expensive. To circumvent this problem, many works focus on simulation. Data were originally annotated with analytic criteria [11], [26], which are still used nowadays for top-down grasps [27]. On the contrary, more and more works rely on a physics engine to simulate the gripper-object interaction. These methods leverage robotic priors to speed up the data generation, resulting in separated studies for 2-finger [8], [19], [23] and multi-fingered [20] grippers.

However, there is no consensus on the best methods to generate diverse grasps for n-finger grippers efficiently. Eppner et al. [23] put a stepstone on that matter through a systematic comparison of standard sampling schemes for 6-DoF grasps poses on a parallel gripper. We extend this analysis to n-finger by including new QD-based sampling schemes.

Quality Diversity. Quality-Diversity methods are optimization algorithms that aim to generate a set of diverse and high-performing solutions to a given problem [28]. Recently, Huber et al. [4] demonstrated that QD could generate datasets of diverse and high-performing reach-and-grasp trajectories for n-finger grippers and proposed a Domain-Randomization-based method for estimating the probability of successfully completing the sim-to-real transfer [21]. We extend this work to 6DoF grasp sampling to investigate how QD can speed up the automatic generation of grasping data.

III. METHOD

A. 6DoF Grasp Sampling

Let Θ be the *parameter space*, and $\theta \in \Theta$ an *individual*. Let G be the space of possible grasps given a rigid object and a gripper, and $\mathcal{G} \subseteq G$ be the space of successful grasps.

¹https://gitlab.isir.upmc.fr/l2g/qd_grasp_6dof

²<https://qdgdrasp.github.io/>




Fig. 2. **Commonly used priors in 6DoF grasp sampling.** (Left) approach-based ; (Right) antipodal-based. R_g is the frame associated with the gripper. The cylinder is an illustration of a targeted object.

This section aims to define how robotic priors can be used to project elements from Θ to G .

Definition. Let $g \in SE(3) \times \mathbb{R}^n$ be a grasp, with n being the number of internal DoF. To reduce the search space on multi-fingered grippers, we close each hand with synergies. However, some grippers need flexibility in the initial joint pose to exploit their gripping capabilities. Thus, we have $g \in SE(3) \times S_y \times \mathbb{R}^k$ with $S_y \subseteq \mathbb{N}^{+*}$ the space of predefined synergies with $\text{Card}(S_y) = m$ the number of synergies for a given hand, and k the number of joints for which the initial state is included in the search space. For a parallel gripper, we have $m = 1$ and $k = 0$, and thus $g \in SE(3)$, similarly to [23]. Details on the synergies and controlled initial joints are provided in the supplementary materials.

Prior-based Sampling schemes. The present article relies on the most commonly used schemes for 2-finger grippers [23]. To our knowledge, the commonly used methods to generate diverse multi-fingered grasps rely on several optimization steps [20], [29], [30]. The present work aims to compare methods that produce grasps poses by directly sampling in a search space without further optimization steps. We let multi-steps methods for future work. The sampling schemes evaluated by Eppner et al. include the *antipodal-based* and *approach-based* ones (Fig. 2). To have a meaningful random baseline, we also introduce a *contact-based* sampling scheme. Let $R_g = (\vec{x}_g, \vec{y}_g, \vec{z}_g)$ be the frame associated with the end effector such that (\vec{z}_g, \vec{x}_g) generates the palm plane and \vec{y}_g drives the gripping direction.

Approach-based prior: Let \vec{n} be the normal on the object surface at the targeted reference point P_g . This method aims to align \vec{n} with \vec{y}_g such that the angle verifies:

$$(\vec{n}, -\vec{y}_g) = \nu \leq \nu_a \quad (1)$$

where ν_a is the maximum half-aperture of the approach cone. The parameter θ to explore G is thus:

$$\theta = (x_r, y_r, z_r, d, \nu, \xi, \omega) \quad (2)$$

where $P_g = (x_r, y_r, z_r)$ defines the targeted reference pose, d is the gripper distance to contact point, ν is defined as in eq. (3), ξ is the cone revolution angle around \vec{n} , and ω is the

gripper rotation around \vec{y}_g . The maximum half-aperture has been set to $\nu_a = \pi/4$, like in Eppner et al. [23].

Antipodal-based prior: Let $\vec{u} = -\vec{n}$, the opposite vector to the normal \vec{n} at the first contact point $P_1 = (x_{r_1}, y_{r_1}, z_{r_1})$. The last point found on the object surface along \vec{u} defines the second contact point P_2 . The normal of the object surface at P_2 is noted \vec{n}_2 . This method consists of finding gripper poses that apply opposite forces on P_2 and P_1 . Therefore, the following criterion must be verified:

$$(\vec{n}_1, -\vec{n}_2) \leq \nu_t \quad (3)$$

where ν_t is the maximal allowed difference between \vec{n}_1 and \vec{n}_2 to define an antipodal grasp. We set $\nu_t = \pi/6$. The parameter θ is thus:

$$\theta = (x_{r_1}, y_{r_1}, z_{r_1}, \Omega) \quad (4)$$

where Ω is the gripper orientation around \vec{u} .

The contact-based methods have been introduced to get baselines with minimal priors. It consists of the approach-based method with $\nu_a = \pi$. Consequently, sampling the parameter space can easily make the gripper overlap with the object. If the approach assumption is as efficient as reported in the literature [23], the contact-based variants should be significantly less sample efficient – while avoiding gripper poses that are too far from the object to grasp it.

B. Leveraging Quality-Diversity for 6DoF-Grasp Sampling

This section aims to detail how we can explore G from Θ to discover solutions in \mathcal{G} .

Definition. This study uses QD standard notations [28]. Let $\mathcal{B} \subseteq \mathbb{R}^{n_b}$ be the *behavior space*, and $\phi_{\mathcal{B}} : \Theta \rightarrow \mathcal{B}$ the *behavior function*, which assigns a *behavior descriptor* $b_{\theta} = \phi_{\mathcal{B}}(\theta)$ to each θ . The *fitness function* is $f : \Theta \rightarrow \mathbb{R}$, and $d_{\mathcal{B}} : \mathcal{B}^2 \rightarrow \mathbb{R}$ is a distance function within \mathcal{B} . The goal is to generate an *archive* A such that:

$$\begin{cases} \forall b \in \mathcal{B}_{reach}, \exists \theta \in A, d_{\mathcal{B}}(\phi_{\mathcal{B}}(\theta), b) < \epsilon \\ \forall \theta' \in A, \theta' = \operatorname{argmax}_{\theta \in N(b_{\theta'})} f(\theta) \end{cases} \quad (5)$$

where $\mathcal{B}_{reach} \subseteq \mathcal{B}$ is the *reachable behavior space*, $\epsilon \in \mathbb{R}^{+*}$ defines the density of \mathcal{B}_{reach} paving, and $N(b_{\theta'}) = \{\theta \mid neighbor_{d_{\mathcal{B}}}(b_{\theta}, b_{\theta'})\}$ is the set of solutions with close projections in \mathcal{B} . $\phi_{\mathcal{B}}$ is deterministic.

Efficiently exploring G from Θ . The overall algorithmic principle is presented in Fig. 1. It consists of a standard QD framework, where a mutation-selection process optimizes an archive A of solutions. Individuals are projected to the grasp space by exploiting the previously described priors. Depending on the gripper, the individual θ has additional values to search among the space of possible joint initial positions and synergies.

To search within the space of contact points on the object surface, the mesh is first preprocessed to get a uniform coverage of its surface. Let S_c be the resulting set of contact points for a given object. The first components of θ are in practice used to define the position $(x_{c_f}, y_{c_f}, z_{c_f})$ of a contact point finder P_f . The first targeted reference point (either P_g or P_1 depending on the prior) is defined as




Fig. 3. **Considered scenes.** Grasp sampling schemes are evaluated on YCB objects [32] in the Pybullet [31] simulators. Experiments involve a FE Panda gripper, a Barrett hand, an Allegro hand and a Shadow hand.

the closest point to P_f in S_c . This method has been used to allow a computationally efficient, physically meaningful local search. A direct search within a flattened list of contact points would limit the locality to the way the mesh has been built, such that two points far away in the list can, in practice, be very close on the object's surface.

The components of individuals are normalized to lie in $[-1, 1]$ by limiting the search space to the bounding box of the targeted object. Each component is projected in its own interval when projecting θ from Θ to G to eventually get a grasping $g \in G$ that consists of a 6DoF pose in $SE(3)$ and a constant closure strategy (defined by the initial joint states and the synergy).

Evaluation. After a population of genomes has been sampled, the computed corresponding grasps are evaluated in simulation. A 3d model of the end effector is set to the computed position in $SE(3)$. The position is already discarded if there is an overlapping between the end effector and the object. Otherwise, a constant closure is then applied to the gripper, with respect to the individual additional components if necessary (i.e. initial joint states and synergy). If the fingers are in contact with the object, a verification phase is conducted – similarly to Eppner et al. [23]. It consists of applying a fixed shaking pattern on the end effector. The quality of the grasp (called the *fitness* f) is computed as the number of shakes the grasp g has resisted.

Population-based optimization. Quality-Diversity methods rely on a behavioral characterization associated with each evaluated solution [28]. We extend [4] formulation of the reach-and-grasp problem to the 6-DoF grasp pose sampling problem by exploiting the same behavior space \mathcal{B} , which is the Cartesian position of the end effector when applying the forces on the object surface. Each QD method has its own mutation-selection pattern. An *outcome archive* A_o is distinguished from the optimization archive [4]. The output A_o is eventually filled with a large variety of robust grasps.

IV. EXPERIMENTS

This section provides the most essential information about the experiments. Details of hyperparameters and implementations are provided in Appendix I.

Grippers and objects. A comparative analysis has first been conducted in the simulator Pybullet [31] (Fig. 3). It involves 4 grippers: a Franka Emika Panda 2-finger gripper, a 3-finger Barrett hand 280, a 4-finger Allegro hand, and a 5-finger Shadow hand. The multi-fingered grippers are

controlled with synergies to remain in the 6-DoF grasp pose paradigm. A subset of 10 objects from the YCB object dataset [32] were used to evaluate the sampling.

Compared methods. The following methods are compared: *contact_rand*, *approach_rand* and *antipodal_rand*. They are the standard prior-based sampling schemes described in section III-A. Elements are randomly sampled from the search space, similarly to Eppner et al. [23]. Some variants of QD methods with robotic priors have also been compared: *ME_scs* [4], which is state-of-the-art for reach-and-grasp trajectory generation, and *CMA_MAE* [33], a general purpose state-of-the-art QD algorithm. *ME_rand* [34] has also been added as a baseline. Note that all prior-based methods are reported with a dedicated prefix (*approach_**, *antipodal_**, or *contact_**). QD methods without prefixes rely on a genome that directly encodes g with spherical coordinates and Euler angles w.r.t. the object frame.

The 3 kinds of priors are studied on the 2-finger gripper. However, the antipodal grasps as described in Eppner et al. are not applicable to multi-fingered grippers. The antipodal-based methods are thus discarded for the other hands. Note that variants of approach-based methods are applied in the literature to generate grasp poses for multi-fingered grippers [37]–[39].

Evaluation metrics. To identify the best 6-DoF grasp sampling methods for multiple grippers, we compute an estimation of the successful grasp space coverage $cvg(\mathcal{G})$, similarly to Eppner et al. [23]. To do so, 5 runs of 10M evaluated samples for *ME_scs* and *approach_me_scs* have been computed. All the generated grasps are then added to a single set of successful grasps S_G . To limit the size of S_G , we only keep the Cartesian position of the end effector while applying the forces on the object surface and round it to 0.01 m. S_G contains a single occurrence of each grasp pose. The *grasp pose coverage* $cvg(\mathcal{G})$ is computed by extracting the Cartesian poses of each successful grasp found in the outcome archive A_o . Positions are then rounded to 0.01 m. The coverage is calculated as the ratio of elements found in S_G divided by $Card(S_G)$. Therefore, $cvg(\mathcal{G}) \in [0, 1]$, with $cvg(\mathcal{G}) = 1$ being the optimal value.

V. RESULTS AND DISCUSSION

Simulated Grasps Synthesis. Fig. 4 shows the evolution of the grasp space coverage with respect to the number of 6-DoF grasps samples evaluated in the simulation. The QD state-of-the-art method with limited priors (*contact_ME_scs*) significantly outperforms the standard grasps sampling schemes (*contact_rand*, *antipodal_rand*, *approach_rand*) regardless of the platform. It shows that QD can speed up grasp sampling compared to standard methods. The usage of the contact prior is here critical, as it leads to a significant increase of performance compared to *ME_scs*.

Interestingly, the results obtained on the standard schemes for the Panda gripper are lower than the one reported by Eppner et al. [23]. This difference might come from the model of contacts in bullets and the simulation parameters that could make the grasps harder in Bullet than in FleX.

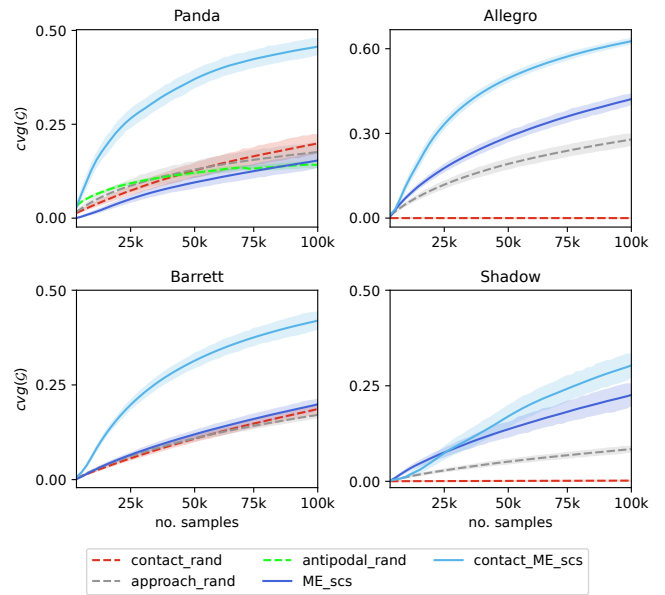


Fig. 4. **QD-based vs standard prior-based sampling schemes.** Ratio of diverse, successful grasp found w.r.t. the number of evaluated samples. The dashed lines are commonly used in the literature; the plain ones are QD-based. *contact_ME_scs* outperforms the standard schemes and the raw QD method by a large margin on the 4 considered grippers.

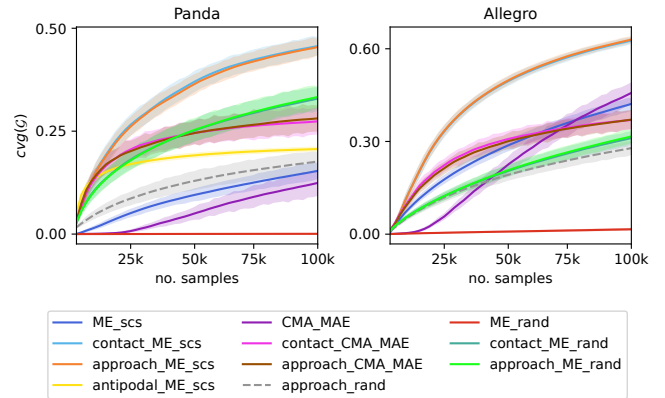


Fig. 5. **Leveraging priors with QD state-of-the-art.** Same principle as in Fig. 4. All the tested QD variants outperform standard prior-based sampling schemes. The pressure for quality optimization makes QD.contact variants behave similarly as QD.approach, such that the algorithm generates solutions that verify the usually hard-coded approach criterion (Fig. 6).

Another difference is the grasp coverage metrics: Eppner et al. include the orientation of the gripper pose for measuring diversity, while we only keep the Cartesian position. This might shift the plateau of strongly constrained methods like *antipodal_rand* to a higher grasping coverage. However, the relative performances are comparable.

Fig. 5 shows the evolution of the grasp space coverage for several combinations of state-of-the-art QD algorithms and robotic priors on the Panda gripper and the Allegro hand. QD variants with priors reached higher coverage than standard prior-based methods. Variants based on *ME_scs* are the best-performing ones, extending Huber et al. [4] results on reach-and-grasps trajectories to 6-DoF poses. The antipodal-based




Fig. 6. **Distribution of ν from successful grasps generated with contact-based variants.** Each bin describes the probability $p(\nu)$ of finding grasps poses with a ν angle between the normal to the hand palm and the normal at the nearest contact point on the object surface. By focusing on the most promising part of the search space, the QD method generates grasps that match the approach prior (i.e. $\nu \in [0, \pi/4]$).




Fig. 7. **Examples of successful grasps coverage per methods.** Each voxel is the end effector position of a simulated successful grasp. The heatmaps describe the robustness of each grasp – the hottest the voxel, the more robust the grasp. Leveraging QD results in a more sample-efficient search, eventually producing a set of grasp poses that cover the whole object surface.

variant of ME_scs quickly plateaus, showing that this prior is too conservative to explore \mathcal{G} . Moreover, it can be noticed that QD variants with either the approach or the contact priors do not reach significant differences in coverage. This is unexpected, as there is a clear difference between standard random-based variants for approach and contact (see Fig. 4). An explanation can be found in Fig. 6, which shows the distribution of ν variables for contact variants of QD methods. The distributions mostly lie between 0 and $\pi/4$, which is the constraint imposed by the approach prior. By focusing on the most promising part of the search space, QD variants generate grasp samples in which the palm normal is aligned with the normal on the object surface at the targeted contact point. In other words, QD contact-based variants automatically find priors that are usually hard-coded to speed up the generation of the grasps.

Fig. 7 provides examples of FR3 grasps poses for standard 6-DoF grasp sampling scheme, as well as one of the two best-performing QD method, after 100k evaluation samples. The QD-generated grasps successfully cover the whole object’s surfaces, including both fragile and robust grasp regions. It is well-known in the machine learning community that high-



Fig. 8. **Examples of transferred grasps.** The proposed method generates diverse grasp poses that are robust in the physical world.

QD Framework	Reach-and-grasp	6DoF	
Gripper	Panda	Panda	Allegro
$\eta^{sim2real}$	0.84	0.95	0.72

TABLE I
MEASURED SIM2REAL TRANSFER RATIOS.

quality datasets include good and bad examples and do not ignore some parts of the targeted distribution. Several works build diverse grasp poses datasets using standard sampling scheme [19], [23]. The obtained results suggest that similar datasets can be built significantly faster by leveraging QD-based optimization methods.

Real robot experiments. To evaluate the exploitability of the grasps generated by ME_scs , grasping repertoires have been generated with the same hyperparameters, except from the usage of a Domain-Randomization-based fitness that showed a significant alignment with sim2real transfer [21]. More details on the experimental setups can be found in Appendix II. Table I reports the obtained sim2real transfer ratios $\eta^{sim2real}$ on a couple of objects when deploying the grasps with higher fitness in A_s . The 6DoF framework leads to comparable results with the reach-and-grasp framework, reaching 95% of transfer on the Panda gripper. Deployments on the Allegro hand resulted in a lower ratio ($\eta^{sim2real} = 72\%$). Most of the failures are due to 2-finger grasp synergies, which do not apply enough force for lifting the power drill. All the attempts succeeded on the bleach cleanser, the bowl, and the mug. In Fig. 8 are provided some examples of successfully transferred grasps. Those experiments show that the generated grasps can successfully be exploited in the physical world.

Limitations. The grasp generation experiments have been conducted with CPU parallelization only. Related works leverage GPU parallelization for scaling the generation of diverse 6-DoF grasp poses [23]. Applying QD methods on GPU environments would imply a stochasticity inherent to parallelization. Recent works explore how to apply QD to stochastic environments [35]. Exploring how to extend the present work to GPU parallelization is a promising way of improving the generation of diverse grasping datasets.

Another limitation is the grippers’ control. There is no fine control of the force intensity here, and the dexterous hands are limited to synergy-based grasps. The use of simple sampling methods for dexterous hands is also restrictive. The proposed framework aims to speed up the generation of diverse and robust grasps for different grippers with minor adaptations. Exploiting multi-step optimization methods

[20], [29] within this framework might be a promising way to generate large datasets of diverse grasps for dexterous hands beyond synergies.

VI. CONCLUSIONS

This work demonstrates the potential of combining Quality-Diversity methods with priors to speed up the 6-DoF grasp sampling schemes and build large datasets of diverse and robust grasps. Our proposed method can easily be applied to several robotic platforms, from parallel 2-finger grippers to dexterous hands. Experiments conducted on real robots show that the diversity of found solutions maintains sim-to-real transferability. We believe such a method can help the latest robotic learning models to converge to generalizing grasping policies by leveraging large datasets of simulated grasps.

ACKNOWLEDGMENT

This work was supported by the Sorbonne Center for Artificial Intelligence, the German Ministry of Education and Research (BMBF) (01IS21080), the French Agence Nationale de la Recherche (ANR) (ANR-21-FA11-0004) (Learn2Grasp), the European Commission’s Horizon Europe Framework Programme under grant No 101070381 and from the European Union’s Horizon Europe Framework Programme under grant agreement No 101070596. This work used HPC resources from GENCI-IDRIS (Grant 20XX-AD011014320). Many thanks to Louis Annabi and Olivier Serris for our insightful discussions.

REFERENCES

[1] Hodson, R. (2018). A gripping problem: designing machines that can grasp and manipulate objects with anything approaching human levels of dexterity is first on the to-do list for robotics. *Nature: Spotlight Robotics*.

[2] Nguyen, V. D. (1988). Constructing force-closure grasps. *The International Journal of Robotics Research*, 7(3), 3-16.

[3] Zhang, H., Tang, J., Sun, S., Lan, X. (2022). Robotic grasping from classical to modern: A survey. *arXiv preprint*.

[4] Huber, J., H el on, F., Coninx, M., Ben Amar, F., Doncieux, S. (2023). Quality Diversity under Sparse Reward and Sparse Interaction: Application to Grasping in Robotics. *arXiv preprint*.

[5] Qin, Y., Su, H., Wang, X. (2022). From one hand to multiple hands: Imitation learning for dexterous manipulation from single-camera teleoperation. *IEEE RAL*, 7(4), 10873-10881.

[6] Wang, P., Manhardt, F., Minciullo, L., Garattoni, L., Meier, S., Navab, N., & Busam, B. (2021, September). DemoGrasp: Few-shot learning for robotic grasping with human demonstration. *IROS 2021*.

[7] Sefat, A. M., Angleraud, A., Rahtu, E., & Pieters, R. (2022, August). SingleDemoGrasp: Learning to Grasp From a Single Image Demonstration. *CASE 2022*.

[8] Depierre, A., Dellandrea, E., & Chen, L. (2018, October). Jacquard: A large scale dataset for robotic grasp detection. *IROS 2018*.

[9] Fang, H. S., Wang, C., Gou, M., Lu, C. (2020). Graspnet-1billion: A large-scale benchmark for general object grasping. *CVPR 2020*.

[10] Levine, S., Pastor, P., Krizhevsky, A., Ibarz, J., & Quillen, D. (2018). Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International journal of robotics research*, 37(4-5), 421-436.

[11] Mahler, J., Liang, J., Niyaz, S., Laskey, M., Doan, R., Liu, X., ... Goldberg, K. (2017). Dex-net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp metrics. *arXiv preprint*.

[12] Yang, J., Tan, W., Jin, C., Liu, B., Fu, J., Song, R., Wang, L. (2023). Pave the Way to Grasp Anything: Transferring Foundation Models for Universal Pick-Place Robots. *arXiv preprint*.

[13] Chi, C., Feng, S., Du, Y., Xu, Z., Cousineau, E., Burchfiel, B., & Song, S. (2023). Diffusion policy: Visuomotor policy learning via action diffusion. *arXiv preprint*.

[14] Team, O. M. (2023). Octo: An Open-Source Generalist Robot Policy.

[15] Urain, J., Funk, N., Peters, J., & Chalvatzaki, G. (2023, May). Se (3)-diffusionfields: Learning smooth cost functions for joint grasp and motion optimization through diffusion. *ICRA 2023*.

[16] Barad, K. R., Orsula, A., Richard, A., Dentler, J., Olivares-Mendez, M., & Martinez, C. (2023). GraspLDM: Generative 6-DoF Grasp Synthesis using Latent Diffusion Models. *arXiv preprint*.

[17] Chen, H., Xu, B., & Leutenegger, S. (2024). FuncGrasp: Learning Object-Centric Neural Grasp Functions from Single Annotated Example Object. *arXiv preprint*.

[18] Padalkar, A., Pooley, A., Jain, A., Bewley, A., Herzog, A., Irpan, A., ... & Jain, V. (2023). Open x-embodiment: Robotic learning datasets and rt-x models. *arXiv preprint*.

[19] Eppner, C., Mousavian, A., Fox, D. (2021, May). Acronym: A large-scale grasp dataset based on simulation. *ICRA 2021*.

[20] Turpin, D., Zhong, T., Zhang, S., Zhu, G., Liu, J., Singh, R., ... & Garg, A. (2023). Fast-Grasp’D: Dexterous Multi-finger Grasp Generation Through Differentiable Simulation. *arXiv preprint*.

[21] Huber, J., H el on, F., Watrelet, H., Amar, F. B., & Doncieux, S. (2024). Domain Randomization for Sim2real Transfer of Automatically Generated Grasping Datasets. *ICRA 2024*.

[22] Newbury, R., Gu, M., Chumbley, L., Mousavian, A., Eppner, C., Leitner, J., ... & Cosgun, A. (2023). Deep learning approaches to grasp synthesis: A review. *IEEE Transactions on Robotics*.

[23] Eppner, C., Mousavian, A., & Fox, D. (2019, October). A billion ways to grasp: An evaluation of grasp sampling schemes on a dense, physics-based grasp data set. *ISRR 2019*.

[24] Chen, Y. L., Cai, Y. R., & Cheng, M. Y. (2023). Vision-Based Robotic Object Grasping—A Deep Reinforcement Learning Approach. *Machines*, 11(2), 275.

[25] Zhou, W., & Held, D. (2023, March). Learning to grasp the ungraspable with emergent extrinsic dexterity. *CoRL 2023*.

[26] Goldfeder, C., Ciocarlie, M., Dang, H., Allen, P. K. (2009, May). The columbia grasp database. *ICRA 2009*.

[27] Vuong, A. D., Vu, M. N., Le, H., Huang, B., Huynh, B., Vo, T., ... & Nguyen, A. (2023). Grasp-anything: Large-scale grasp dataset from foundation models. *arXiv preprint*.

[28] Cully, A., Mouret, J. B., & Doncieux, S. (2022, July). Quality-diversity optimisation. *GECCO Companion 2022*.

[29] Yao, K., & Billard, A. (2023). Exploiting Kinematic Redundancy for Robotic Grasping of Multiple Objects. *IEEE Transactions on Robotics*.

[30] Turpin, D., Wang, L., Heiden, E., Chen, Y. C., Macklin, M., Tsogkas, S., ... & Garg, A. (2022, October). Grasp’d: Differentiable contact-rich grasp synthesis for multi-fingered hands. In *European Conference on Computer Vision* (pp. 201-221). Cham: Springer Nature Switzerland.

[31] Coumans, E., Bai, Y. (2016). Pybullet, a python module for physics simulation for games, robotics and machine learning.

[32] Calli, B., Walsman, A., Singh, A., Srinivasa, S., Abbeel, P., Dollar, A. M. (2015). Benchmarking in manipulation research: The ycb object and model set and benchmarking protocols. *arXiv preprint*.

[33] Fontaine, M., & Nikolaidis, S. (2023, July). Covariance matrix adaptation map-annealing. *GECCO 2023*.

[34] Mouret, J. B., & Clune, J. (2015). Illuminating search spaces by mapping elites. *arXiv preprint*.

[35] Chatzilygeroudis, K., Cully, A., Vassiliades, V., & Mouret, J. B. (2021). Quality-diversity optimization: a novel branch of stochastic optimization. In *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*.

[36] H el on, F., Huber, J., Amar, F. B., & Doncieux, S. (2024). Toward a Plug-and-Play Vision-based Grasping Module in Robotics. *arXiv preprint*.

[37] Roa, M. A., Argus, M. J., Leidner, D., Borst, C., & Hirzinger, G. (2012, May). Power grasp planning for anthropomorphic robot hands. *ICRA 2012*.

[38] R. Grimm, M. Grotz, S. Ottenhaus, and T. Asfour, “Vision-based robotic pushing and grasping for stone sample collection under computing resource constraints,” *ICRA 2021*.

[39] C. Pohl, K. Hitzler, R. Grimm, A. Zea, U. D. Hanebeck, and T. Asfour, “Affordance-based grasping and manipulation in real world applications,” *IROS 2020*.

[40] Miller, A. T., & Allen, P. K. (2004). Graspit! a versatile simulator for robotic grasping. *RAM 2004*.

Supplementary Materials

APPENDIX I

EXPERIMENTAL DETAILS: SIMULATION

Objects. Following YCB objects were used in the data generation experiments. For the Panda gripper: banana, bleach cleanser, bowl, cracker box, mug, power drill, rubiks cube, spatula, tennis ball; for the Barrett hand: banana, bleach cleanser, bowl, cracker box, power drill, rubiks cube, spatula, tennis ball; for the Allegro hand: banana, bleach cleanser, bowl, chips can, cracker box, mug, power drill, rubiks cube, spatula, tennis ball; for the Shadow hand: banana, bleach cleanser, bowl, mug, power drill, rubiks cube, spatula, tennis ball.

Algorithms. Let μ be the population size, λ the number of offspring, k the number of neighbors considered for novelty computation, and N_e the maximum number of evaluation. We set: $\mu = \lambda = 500$, $k = 15$, $N_e = 100k$. All offspring are mutated with a probability $ind_{pb} = 0.3$ to modify each gene. For a fair comparison, all ME-derived methods sample $\mu = \lambda$ individuals for offspring generation at each iteration. The mutation operator applied by default to all the methods is a Gaussian perturbation of 0 mean and 0.1 standard deviation. This value of σ has been obtained by doing grid search, trying to get the highest possible value of $avg(\mathcal{G})$ after 20k evaluated samples. The variance of optimal value of σ between methods was so small that we set it to 0.1 for all the mutation-based methods. For *CMA.MAE* variants, we used the same parameters as in [33]: The emitter batch size is set to 36, and the number of emitters to 15, $f_{min} = -1$ and $\alpha = 0.01$.

The fitness f is computed by applying a predefined shaking pattern, similarly to [23]. The fitness of a grasp pose g is equal to the number of shakes applied to the gripper after the closure without making the object fall. Let N_s the number of performed shakes. We set $N_s = 2$: the first perturbation is a translation shake, the second one is a rotation - both around a fixed axis in the simulated world. Therefore, the optimal fitness is $f^* = N_s = 2$.

Grippers. Following are the details on the way each gripper is controlled. For the Panda gripper: The fingers opened as wide as possible, and closed at grip time; for the Barrett hand: the initial position of the joints that orient the proximal part of each fingers with respect to the palm are defined by the parameter, such that the search space allows to get several initial configuration of those fingers. All the fingers are then uniformly closed; for both the Allegro and the Shadow hand: the hand is closed with respect to synergies. The synergies consist of a uniform closure of specific fingers, including thumb-index, thumb-mid, thumb-index-mid, and all the fingers. All fingers are initialized in wide open position. The thumb palm-to-proximal joint is always set such that the thumb is orthogonal to the palm. This kind of synergies have been used for their simplicity, their efficiency, and the ease with which they can be adapted to different grippers. Other kind of synergies can be found in the literature, like the eigen grasps [40]. Studying the optimal

way to generate grasps for n-fingers is beyond the scope of this paper.

APPENDIX II

EXPERIMENTAL DETAILS: PHYSICAL WORLD

Data generation. The fitness has been replaced with a quality criterion dedicated to sim2real transfer. We use a variant of the Mixture Domain Randomization criterion proposed in [21], removing the variance on the joint state - as it is redundant with the variance on the object state in the 6DoF grasp prediction context. Let N_{MDR} be the number of retrials, σ_{o_p} the object position variance along each axis, and σ_{o_o} its orientation variance. We set $N_{MDR} = 100$, $\sigma_{o_p} = 0.005m$, $\sigma_{o_o} = 30^\circ$. We kept the same variance of friction coefficient as in [21]. The shaking process is similar to the simulated experiments. Consequently, the best possible fitness is $f_{MDR}^* = N_s \times N_{MDR} = 200$. The probability $\eta^{sim2real}$ for a given grasp g to successfully transfer in the physical world can thus be estimated as:

$$\eta^{sim2real}(g) = \frac{f(g)}{f_{MDR}^*}$$

All the remaining hyperparameters are similar to those used in the simulated experiments. The grasps are then sorted with respect to their fitness f_{MDR} .

Real world setup. Real experiments involve a Franka Research 3 arm with a parallel gripper and an Allegro-Hand V4 with gravity compensation and some compliance. The same arm is used for both grippers. A 3D-printed mounting adaptor is used to fix the Allegro hand on the arm. The Allegro hand is controlled with a joint impedance controller with gravity compensation. No additional material that could increase adhesion is used to make grasping easier. The synergies are the same as the ones used in the simulated experiments. The gravity compensation is computed with respect to the fixed arm base. The motion planning of the end effector is conducted with RRT connect through *Moveit!*, assuming that the object and the table are collision bodies. The object pose detection is conducted by leveraging a hand-eye calibrated Realsense D435i, using the vision pipeline from H el enon et al. [36]. The angle between the camera point-of-view and the table plan is roughly 30° .

Experiments. The real experiment involves 4 YCB objects: the bowl, the mug, the bleach cleanser and the power drill. It is worth noting that among the sorted high-performing grasps, the ones that do collide with the table were discarded. Almost all the deployed trajectories have successfully been transferred on the physical gripper. The failures are primarily due to a misalignment between the applied forces in simulation and the real world. This also raises a limitation of the 6DoF setup widely used in the literature. A finer control of applied forces could overcome this limitation.