Analysis of Open-Loop Grasping From Piles

Elődó Páll Oliver Brock

Abstract— **This paper offers an explanation of why humans can effortlessly grasp objects from a pile. We identified a regularity in objects' motion when pushed, namely, an object separates and stabilizes in front of the pusher. We devise an open-loop grasping strategy leveraging this regularity in piles of nearly identical objects. Our real robot robustly grasps round objects beside a wall with success rates between 95% and 100% without visual or tactile feedback. We analyze our grasping strategy extensively both in real-world and simulated experiments. We observe that object roundness improves grasping and the motion pattern also manifests in small piles beside a wall. Our qualitative simulation can approximate the real robot's grasping behavior, and we apply open-loop grasping in an warehouse pick-and-place application.**

I. Introduction

Humans grasp objects from piles with incredible ease. As an example of this, consider picking up a single nut from a container of nuts or eating popcorn from a box while watching a movie. In these contexts, the robustness of human grasping may seem surprising at first sight. After all, complex contact dynamics are playing out during grasping, leading to substantial motion within the pile. It is inconceivable that humans possess precise models and perform accurate simulations of the pile's motion. But then how do they do it? One explanation could be the incredible sense of touch in the human hand. But in this paper, we show that there is another explanation.

Grasping from piles of almost identical objects is, in fact, *enabled* by the complex interaction forces and resulting motion patterns among the individual objects, rather than rendered complicated. This is because the interactions exhibit a strong regularity that *enables* robust grasping. When grasping from such a pile, visual or tactile feedback are no longer necessary. Grasping can be performed robustly and open-loop (just like eating popcorn). This paper aims to investigate this phenomenon to enable robotic pile grasping (or bin picking) with very simple, open-loop strategies.

This paper analyzes and explains the motion patterns that occur during grasping in piles of nearly identical objects. Based on the resulting insights, we devise a very simple, open-loop (no visual or tactile feedback) grasping strategy that—for a variety of object geometries—achieves highly robust grasping from piles (up to 100% success). We validate these results with different types of end-effectors, both in real-world and in simulated experiments. We also examine the dependency of pile cardinality (number of objects in a pile) and grasp success, finding that already small piles exhibit the regularity required for robust, open-loop grasping. Our experiments reveal that features of the environment that constrain the pile's motion, such as walls or corners (two walls), substantially lower the required pile cardinality for the desired motion patterns to occur. We also demonstrate that it is possible to assess the suitability of the open-loop pile grasping strategy from simulation experiments. Finally, we test the open-loop grasping in a bin-picking application.

Both authors are with the Robotics and Biology Laboratory, Technische Universität Berlin, Germany. We gratefully acknowledge the funding provided by the European Commission (SOMA, H2020-ICT-645599) and by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy - EXC 2002/1 "Science of Intelligence" - project number 390523135.

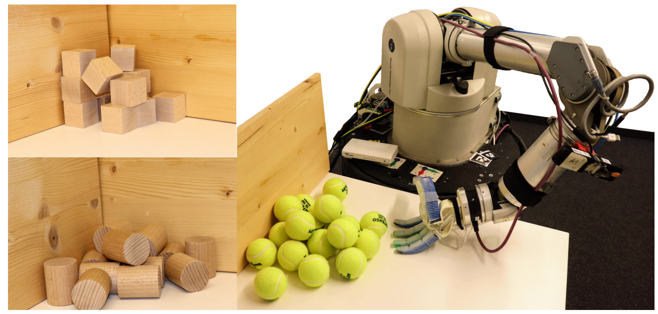

Fig. 1: We analyze open-loop grasping in random piles of wooden cubes (top-left), cylinders (bottom-left), or tennis balls (right) supported by a static wall or corner, using the Barrett WAM 7DOF arm with the RBO Hand 2.

II. Related Work

Common grasping methods rely on accurate visual perception. They detect individual objects and their geometrical properties [1, 2, 3, 4] to compute or learn feasible grasp poses for simple end-effectors like suction-cups or parallel grippers. Some methods increase perception accuracy with interactive perception [5, 6] or pile-separation [7, 8, 9]. Our approach is entirely different because we will grasp without object detection but exploiting motion patterns in piles. Moreover, we want to use an anthropomorphic hand due to its versatile use after and beyond grasping.

Since we postulate that motion patterns enable open-loop grasping from piles, we look at related work to see how they exploited motion patterns. Motion patterns arise from some form of restriction.

Inertial properties and gravity constrain an object's motion. The resulting motion pattern can be exploited to increase a robot's dexterity [10, 11] with simple motion primitives. An object's motion is also constrained through contact with the environment. A static contact state reduces an object's degrees of freedom, and it is leveraged to stabilize the object for grasping [12, 13], or re-orienting and re-grasping [14].

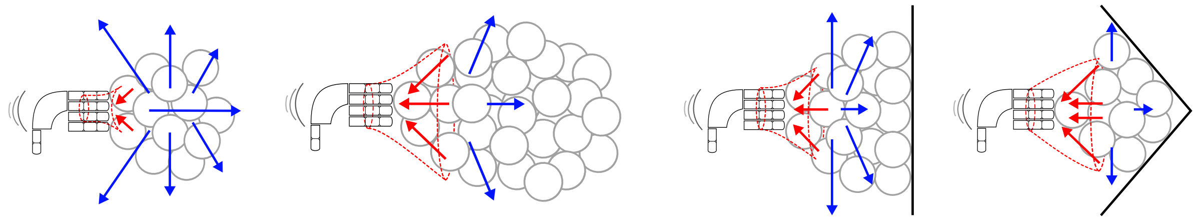

Fig. 2: All four sketches show a hand pushing a pile from left to right, where the red arrows indicate the relative motion of objects inside a virtual (red) funnel. The blue arrows show the absolute motion of the objects outside this funnel. For large enough piles, the funnel separates and pushes an object into the hand while the rest of the pile expands less. A wall or corner restricts the pile's expansion. Hence smaller piles supported by a wall or corner also exhibit the same behavior.

Quasi-static interactions produce predictable sliding motions enabling dexterous manipulation [15] or rearranging clutter around a goal object [16]. Contact restricts a robot's motion too. This was leveraged to decrease the robot's configuration-state uncertainty [17, 18, 19] and therefore increase execution robustness. Perception becomes simpler because we only need to detect environmental affordances [20] and not a precise model of the environment.

A motion pattern arises when a compliant hand interacts with rigid objects. The shape of the hand deforms due to contact forces adapting to the object. This deformation regularity can be approximated [13] or not even modeled [12] but leveraged for grasp planning. This way, a compliant hand can robustly grasp with a simple open-loop controller [21, 22].

The presented methods interact with a single movable object and leverage motion patterns to simplify planning, perception, and control. We expect to gain similar benefits if we identify a motion pattern in a collection of movable objects and leverage it for grasping.

III. Analysis of Motion Patterns in Piles of Nearly Identical Objects

We analyze and explain the motion patterns in piles of nearly identical objects to motivate open-loop grasping.

We observed regularity in objects' motion when a hand applies a force on a pile. Objects in front of the hand are actively pushed toward the hand by the rest of the pile, while other objects spread radially, as visualized in Figure 2. The pile spreads radially because interaction forces between the objects also spread radially, observed experimentally in a collection of rigid objects [23]. This regularity effectively separates an object in front of the hand. When the pile is large enough, the separated object is stabilized by the rest of the pile, and the hand can slide under it. The object is stabilized because some portion of the interaction forces is absorbed by friction between objects and between objects and the environment [24]. The described motion pattern, object separation and stabilization, arises even in smaller piles when a static wall supports the pile.

Object separation and stabilization arise due to a restriction in objects' motion emerging from the interaction between the robot, objects, and the environment. The literature refers to such contact-based motion constraints as *environmental constraints* (EC) [25]. We will use EC interchangeably with motion patterns. Yet, material sciences define a collection of interacting objects as a granular material [26]. Thus, we will refer to motion patterns in a pile of objects a *granular EC*.

The granular EC is radically different from existing ECs because the constrained motion manifests in a collection of movable objects and accomplishes two goals for grasping: object separation and stabilization. Interestingly, even though it is difficult to predict the individual objects' motion accurately when pushed, the higher-level motion pattern arises consistently. Hence, we can exploit it with open-loop grasping and without modeling individual objects' motion.

This is one instance of such EC where we can abstract physical interactions away from perception and control, but the first one that enables open-loop grasping from piles. Our analysis contributes an experimental procedure to identify and characterize ECs, which opens up the possibility of finding similar ECs that will simplify perception, control, and planning for other manipulation tasks.

IV. Exploitation of Granular Environmental Constraints

We define the open-loop grasping strategy that exploits the granular EC. The grasp strategy is composed of three motion primitives. Each primitive moves the hand on a straight line, and it is executed with a simple operational-space controller. First, the hand approaches the pile, then pushes it from the side, and finally, it grasps a separated and stabilized object. The transitions between the primitives are distinct contact events. The contact events occur when the hand interacts with the pile or the environment. The three primitives are sketched in Figure 3, shown in Figure 4, and detailed below:

1) Approach Primitive: the hand moves into proximity to the edge of the pile by executing a free-space motion and lowers until it detects the surface's normal force. When a static wall or corner supports the pile, we consider the wall's or corner's location, as shown in Figure 3 where the relative hand pose is represented by α, β, and θ angles. α is the relative hand orientation to the pile, β is the offset angle between the fingers and the direction of motion, and θ is the slope on the fingers. We need to detect the pile's center, spread, and static walls to devise this primitive.

2) Push Primitive: the hand slides on the support surface and pushes the pile, shown on the third sketch in Figure 3. The primitive is parameterized by the sliding speed $|v_{\text{hand}}|$

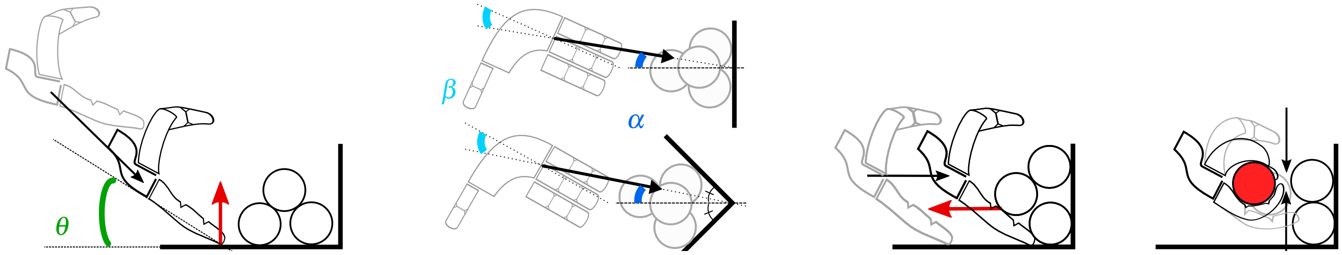

Fig. 3: 2D sketch of the grasp strategy in the order of the motion primitive execution: *approaching* the pile from free space, where θ (green) is the finger's slope, α (blue) and β (cyan) are the relative hand orientation to the pile and a supporting wall or corner. The normal force (vertical red arrow) triggers the *pushing* primitive, which terminates when a contact force (horizontal red arrow) is detected, and finally, the hand *grasps* an object that rolled on it.

Fig. 4: Leveraging granular ECs for grasping from piles next to a corner (top row) or wall (bottom row) when bin picking.

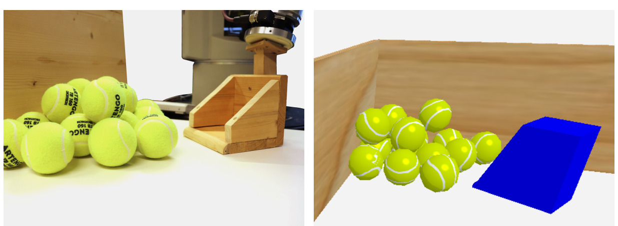

Fig. 5: The real shovel-like end-effector approaches a pile of tennis balls next to a wall (left) and similarly in a multi-body dynamics simulation where the pile is in a corner (right).

and the force threshold F_{grasp}. The next primitive is triggered when the force (in the pushing direction) reaches F_{grasp}.

3) Grasp Primitive: the hand simply closes its fingers, assuming that an object was separated and rolled into the hand when F_{grasp} triggered.

V. Experimental Analysis of Open-Loop Grasping

We analyze the grasping strategy with a real robot and in simulation. We conducted extensive real-world experiments to examine grasping with granular ECs. First, we will prove our initial hypothesis in Sec. V-B that the granular EC enables open-loop grasping from piles of nearly identical objects. We further analyze grasping robustness concerning objects' shape in Sec. V-C and the end-effector's morphology in Sec. V-D. We also examine the dependency between the pile's cardinality (number of objects in a pile) and the environment in Sec. V-E. We investigate the influence of the

Objects	**mass [g]**	**size [mm]**		**R**
	real	real	sim	
light spheres	60	$r = 27$	$r = 56$	4.3
heavy spheres	180	$r = 27$	$r = 56$	5
cylinder	45	$h = 50,\ r = 20$	$h = 50,\ r = 45$	9
cube	65	$l = 45$	$l = 50$	–
wooden lime	29	$h = 68,\ r = 49$	–	5.6
real apple	168	$h = 64,\ r = 70$	–	7.6

TABLE I: We used similar objects' mass and size in real-world and simulated experiments, where l, h, r are length, height, and radius respectively. In our industrial application in Sec. V-H, R is the average grasp attempts to pick four objects successfully, where the perfect performance is $R = 4$.

object's mass on grasping in Sec. V-F. We will show that qualitative simulation can verify the granular EC's existence even for a significant sim-to-real gap in Sec. V-G. Finally, we will apply open-loop grasping in an industrial use case in Sec. V-H.

A. Experimental Setup

In the real-world experiments, we use a Barrett WAM 7DOF arm with the compliant and anthropomorphic RBO Hand 2 [27]. We execute grasping strategies using ATI FTN-Gamma force-torque sensors mounted between the hand and the wrist. We assume to know the pile's and vertical walls' location. In a multi-body dynamics simulation, we analyze the *push primitive's* behavior qualitatively across different problem properties, but we report only on the relevant ones due to the page limit. We use a simple, rigid, shovel-like end-effector shown in Figure 5. We also built the shovel to compare it with the RBO Hand 2. We compute the grasp success rate from 20 real-world and 50 simulated grasp attempts. For the real-world bin-picking application, we average the grasp attempts from three picking sessions per object type.

We build random piles by dropping objects in a cube-like frame with size 25cm for pile cardinality $|\text{Pile}| = 18$, 18cm for $|\text{Pile}| = 14$, and 15cm for $|\text{Pile}| = 10$. We list the real and simulated object's properties in Table I. All simulated objects have the same friction properties using the tennis ball's properties [28]. The restitution coefficient (COR) of all objects is a linear function of the mass, where COR(60g)$= 0.72$ [28], and COR(180g)$= 0.37$ using our experimental observation on our sand-filled heavy tennis balls. The simulated end-effector has lateral friction 0.5, spinning-, and rolling-friction 0.001, and COR=0.8. For bin picking, we also used real apples and net-bags of three wooden limes. A grasp attempt is successful if one or more objects are in the hand or shovel.

B. Granular EC Exploitation Leads to Robust Grasping

We want to demonstrate that granular EC exploitation with our hand does not require the definition nor detection of conditions for grasp success. The conditions for grasp success are complicated, but the EC helps by making grasp success unconditional. Our results (Figure 6) show that three hand-tailored grasp strategy, which uses object-based conditions, never outperforms the pile-based strategy, which relies only on the granular EC.

Through five iterations, we hand-tailored the conditions for successful grasping from piles of light and heavy tennis balls and cylinders with the RBO Hand 2. We achieved the best performance for tennis balls by centering the gap between the middle and ring fingers with an object and choosing an object that moves toward the pile's center when pushed straight toward a wall or corner ($\alpha = \beta = 0°$). Like tennis balls, we choose a cylinder with its curved surface toward the hand or standing on its base. We compare the grasp success rate of these two object-based strategies to the pile-based strategy.

The pile-based strategy has only pile- and environment-based conditions (granular EC) for grasping. We explored different conditions, namely, α and β, and the best is $\alpha = \beta = 0°$ for light and heavy tennis balls, and $\alpha = 45°, \beta = -20°$ for cylinders. The robot grasps from random piles of 18 light or heavy tennis balls or 14 cylinders next to a wall or corner to compare the object- and pile-based strategies.

In all six cases, The results show that the hand-tailored strategy never outperforms the pile-based strategy when grasping with our hand limited to the considered objects. Therefore, we do not need to devise object-based grasping conditions because the granular EC fulfills the conditions implicitly. Moreover, it is increasingly difficult to characterize grasp success for heavy tennis balls that are difficult to grasp. This is why our hand-tailored strategy counteracts the granular EC, which results in a significant performance gap when a wall supports a pile of heavy tennis balls.

We conclude that granular ECs can fulfill grasping conditions, and therefore, we don't need to compute the strategy for robust grasping.

C. Object Roundness Improves Grasping

We analyze grasp robustness concerning the object's shape. Hence, we execute open-loop grasping from piles of spheres, cylinders, and cubes with various strategy parametrizations $(\alpha,\ \beta) \in \{(0°, 0°), (30°, \text{-}10°), (30°, 20°), (45°, \text{-}20°)\}$. Objects' mass is m =60g, the pile was next to a wall, and its cardinality is $|\text{Pile}| \in \{14, 15, 18, 20\}$. We expect that open-loop grasping succeeds more likely for round objects than cubes because cubes can compact, inhibiting object separation and rolling into the hand.

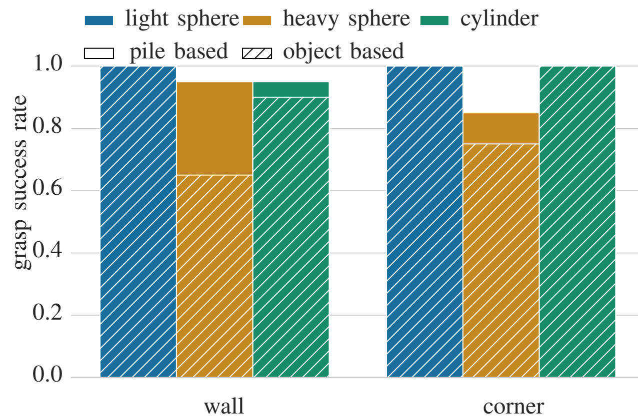

Fig. 6: Grasping from a pile does not require object-based strategy due to the granular EC's help. The real robot robustly grasp from random piles of light (blue) and heavy tennis balls (yellow) and cylinder (green) constrained by a wall or corner irrespective of the strategy.

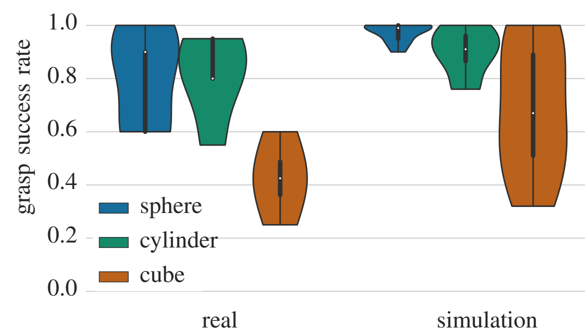

Fig. 7: We can exploit the granular EC for round objects (blue and green) for a wide range of actuation parameters, unlike cubes (red). We can use simulation to identify the EC because the behavior is qualitatively similar to the real-world.

We sample a grasp success probability for each instantiation of the above-described problem space by averaging multiple executions (see Sec. V-A). To compare the success rate distributions between different object shapes, we show in Figure 7 a kernel density estimation of the sampled grasp success rates. The violin plots are limited only to the observed data. For now, we analyze the real-world results, and we will discuss the simulation results in Sec. V-G.

The results confirm that open-loop grasping is more robust for round objects. The mean grasp success rates are $\geq 80\%$ for both tennis balls and wooden cylinders. For the wooden cubes, it is $< 50\%$. We can devise open-loop grasping strategies for round objects in piles, like nuts in a can.

D. The End-Effector's Compliance Improves Grasping

We analyze the influence of the end-effector's morphology (shape and structure) on grasp robustness. Therefore, we execute the grasp strategy with the anthropomorphic, complaint RBO Hand 2, and shovel-like, rigid end-effector. We grasp from piles of 18 light and heavy tennis balls and 14 cylinders or cubes next to a wall using $\alpha = \beta = 0°$

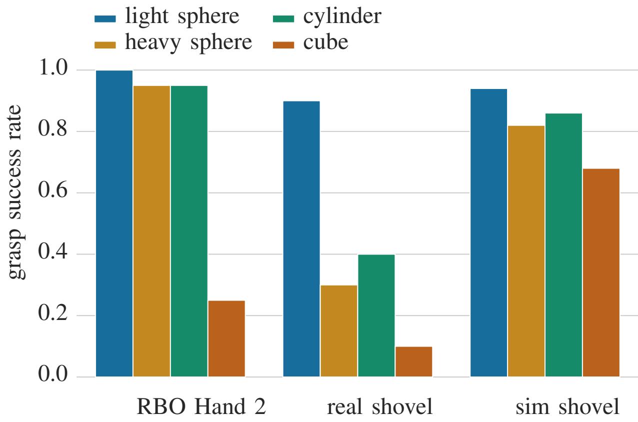

Fig. 8: The real robot performs best with the RBO Hand 2 for round objects, while the real shovel can only exploit the EC for light tennis balls. All real and simulated experiments show qualitatively similar grasp behavior. Simulation alone cannot indicate an end-effector's limitation concerning the object's shape due to the quantitative differences.

as discussed in Sec. V-B. The implementation of the *push primitive* (see Sec. IV) differs for the hand and the shovel. We use an operational space controller for the hand due to its compliance, but the rigid shovel requires an impedance controller to achieve safe sliding on the table and interaction with the pile. We expect that the hand performs better due to its shape and structure compared to the shovel.

Figure 8 shows the real and simulated mean grasp success rates, as described in Sec. V-A. The real robot performs best with the hand. The hand achieved between 95% and 100% grasp success for round objects, while the shove could scoop up only light tennis balls with 85%. Both end-effectors performed poorly for cubes, which indicates a limitation of the granular EC.

Though the shovel can exploit the granular EC, only the hand allows open-loop grasping. This is because we need closed-loop force control (impedance controller) to slide the shovel on the table. The *grasp* primitive also benefits from the hand's compliance because the fingers can adapt to the object's shape while closing. The hand's shape also helped object separation because the space between fingers creates a guiding rail for separating and rolling objects into the palm.

We executed over 900 grasp attempts and observed grasps of multiple objects within one attempt. We grasped more than one object 15% for 600 successful grasps with the RBO Hand 2 and 8.75% for 160 with the shovel. The number of grasped objects and grasp success depended on the object's and end-effector's size, as shown in Figure 9. In case single object grasping is desired, only the hand could solve it with in-hand-manipulation.

In conclusion, a human-like end-effector enables robust open-loop grasping of round objects supported by a wall, like eating popcorn from a box.

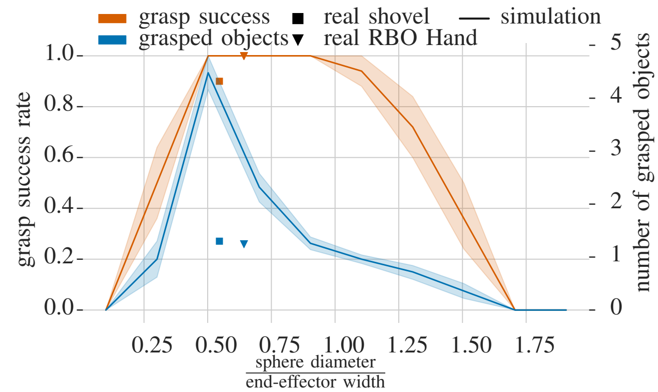

Fig. 9: The grasp success and number of grasped objects depend on the the ration between object's and end-effector's size shown for light spheres in a corner.

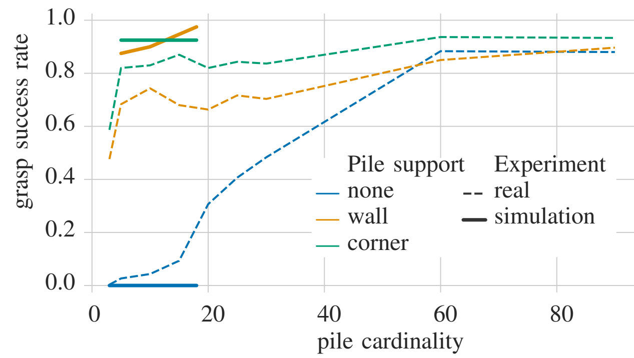

Fig. 10: Our grasping strategy can exploit the granular EC even for small piles when supported by a static wall or corner with the real robot (solid lines) and in simulation (dashed lines). The simulation shows that a pile alone supports granular EC exploitation from a certain cardinality.

E. The EC Manifests in Small Piles Supported by a Wall

We analyze grasp robustness concerning the pile's cardinality and static environmental constraints, like vertical walls. Hence, we grasp from piles with various cardinality, in the real-world $|\text{Pile}| \in \{5, 18\}$, and in simulation $|\text{Pile}| \in \{3, 5, 10, 15, 20, 25, 30, 60, 90\}$, with and without a static wall or corner. The simulated sphere's mass is $m_{\text{object}} \in \{60, 120, 180\}$ and the real tennis balls have $m_{\text{object}} \in \{60, 180\}$. We ran only ten grasp attempts for piles of 18 heavy tennis balls without any vertical support on the real robot because the grasp primitive never triggered. Since we could not grasp from a pile of 18 heavy tennis balls, we did not run other experiments with fewer or lighter objects, and we assume that those attempts would also fail. We expect that grasp succeeds for large enough piles or if a static wall or corner is beside the pile.

In Figure 10, we visualize the mean grasp success rate with segment-wise linear interpolation for both real-world (solid lines) and simulation (dashed lines) using the best strategy parametrization for spherical objects (see Sec. V-B). Even though the real robot fails to grasp from piles without a supporting wall, we observed that 40% of the attempts could have succeeded with a sensorized hand [29].

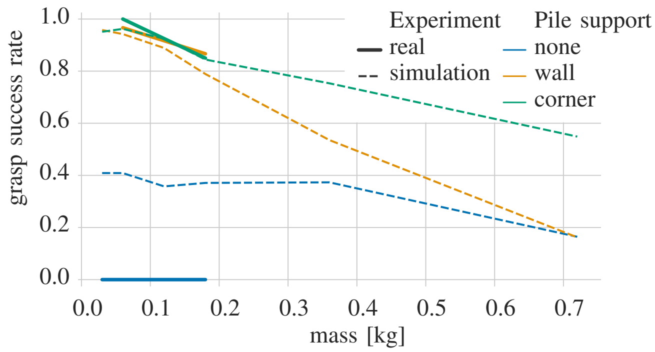

Fig. 11: Lighter objects are easier to grasp because their inertia is less, so objects can move within the pile, and the EC can arise over time. Moreover, static supports can further constrain the pile's motion similar to the object's mass and hinder granular EC exploitation for heavy objects.

The results prove open-loop grasp robustness, more than 85%, even for smaller piles supported by a wall or corner.

F. Lighter Objects Are Easier to Grasp

We analyze the influence of object's mass on our approach. We grasp from piles of spheres with various mass $m_{\text{object}} \in \{60, 180\}$ in the real world, and in simulation $m_{\text{object}} \in \{30, 60, 120, 180, 360, 720\}$. The other problem properties are the same as presented in Sec. V-E. We expect an increase in grasp success for lighter objects.

Figure 11 shows the grasp success rates with segment-wise linear interpolation. In the case of piles without any support, grasping fails, as discussed in the previous section. However, the results confirm that lighter objects are easier to grasp.

We conduct a second experiment to show that we can improve grasp success for heavy objects. We grasp 18 heavy tennis balls in a corner because the corner further constrains objects' motions. We expect that if we increase the hand's velocity and grasping force threshold, grasp robustness increases because heavier objects require more force to induce the motion pattern. We successfully increased grasp success from 60% to 85% by increasing the velocity $|v_{hand}|$ from 0.1 to 0.25m/s and force threshold F_{grasp} from 17 to 25N.

G. Real and Simulation Results Are Qualitatively Similar

We want to show that simulation is qualitatively similar to real-world open-loop grasping. Hence, we simulate open-loop grasping likewise to the four previous sections. We described both real-world and simulation experiments in those sections, and now, we only interpret the results.

We showed that object roundness improves open-loop grasping in the real world. Figure 7 also confirms a qualitative similarity between real-world and simulation results. For both cases, distributions' mean increases with the object roundness. We can observe qualitative similarity concerning object roundness in Figure 8, where we compared different end-effectors. However, this experiment shows a significant sim-to-real gap between real- and simulated-shovel, similarly between the hand and simulated shovel for cubes. Therefore, simulation cannot indicate an end-effector's limitation.

The real robot robustly grasps from smaller piles if those are next to a wall or corner. This is also true in simulation, as shown in Figure 10. The simulation indicates that grasp robustness increases with the pile's size, but the real-world executions only support it when the pile is next to a wall. The simulation also indicates that a larger pile ($|\text{Pile}| \geq 60$) could enable grasping even without static support, but we did not test it on the real robot due to practical reasons.

Finally, the real-world and simulation results indicate (see Figure 11) that lighter objects are easier to grasp when supported. In the simulation, we also showed the relationship between grasp success and relative object size to the end-effector in Figure 9.

We conclude that simulation is qualitatively similar to real open-loop grasping and should be used to identify and motivate real robot experiments for future ECs.

H. Bin Picking Application With Granular EC Exploitation

Finally, we show the applicability of open-loop grasping in an industrial grocery logistics use case of Ocado [30], where a robot has to picks and place N objects of the same type to fulfill an order, and the bin's location is known.

In our adaption of the task, $N = 4$ and the random piles are 18 apples, tennis balls, cylinders, or five net-bags of limes. The key difference to Ocado's setup is that we tilt the bin 5°, so objects remain beside the same wall after each grasp attempt. This way, we do not need to detect the pile. An operator decides to grasp from the pile constrained either by the wall or a corner. If the robot grasps more objects than required (e.g., it already picked three, and next, picks two objects), the robot drops them back to the pile.

We present the average grasp attempts to pick four objects in Table I and in this video[1], which shows that open-loop grasping can be applied in an industrial use-case. The results indicate that granular EC generalizes for irregular object shapes too. In the future, we want to learn a heuristic for alternating between the wall or corner support, which might require perceiving geometrical properties of the pile.

VI. Conclusion

Environmental constraints simplify control and perception and increase grasping robustness. This paper presents the novel granular EC based on motion patterns in piles of almost identical objects. This EC effectively separates and stabilizes an object from the pile. We demonstrated the existence of the EC. Our grasp strategy requires no grasp success characterization for robustness. Without object detection and a simple operational space controller, we achieved success rates between 95% to 100% for round objects besides a wall. We analyzed the grasp robustness in real-world experiments and showed that a static wall enables robust grasping from smaller piles. We showed generalization for different objects and environments. We validated that qualitative simulation can approximate grasping behavior. Finally, we successfully applied open-loop grasping in an industrial use case, demonstrating the granular EC's relevance.

[1] https://youtu.be/9bjqHQ8NUNo

References

[1] K. Rahardja and A. Kosaka, “Vision-based bin-picking: recognition and localization of multiple complex objects using simple visual cues,” in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems. (IROS)*, vol. 3, Nov. 1996, pp. 1448–1457 vol.3.

[2] M. Nieuwenhuisen, D. Droeschel, D. Holz, J. Stckler, A. Berner, J. Li, R. Klein, and S. Behnke, “Mobile bin picking with an anthropomorphic service robot,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2013, pp. 2327–2334.

[3] J. Mahler, M. Matl, V. Satish, M. Danielczuk, B. DeRose, S. McKinley, and K. Goldberg, “Learning ambidextrous robot grasping policies,” *Science Robotics*, vol. 4, no. 26, Jan. 2019.

[4] M. Pozzi, G. Salvietti, J. Bimbo, M. Malvezzi, and D. Prattichizzo, “The closure signature: A functional approach to model underactuated compliant robotic hands,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 2206–2213, July 2018.

[5] D. Katz, M. Kazemi, J. A. Bagnell, and A. Stentz, “Clearing a pile of unknown objects using interactive perception,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2013, pp. 154–161.

[6] H. v. Hoof, O. Kroemer, and J. Peters, “Probabilistic Segmentation and Targeted Exploration of Objects in Cluttered Environments,” *IEEE Transactions on Robotics*, vol. 30, no. 5, pp. 1198–1209, Oct. 2014.

[7] M. Gupta and G. S. Sukhatme, “Using manipulation primitives for brick sorting in clutter,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2012, pp. 3883–3889.

[8] L. Chang, J. R. Smith, and D. Fox, “Interactive singulation of objects from a pile,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, May 2012, pp. 3875–3882.

[9] M. Dogar, K. Hsiao, M. Ciocarlie, and S. Srinivasa, “Physics-based grasp planning through clutter,” *Robotics: Science and Systems (RSS)*, 2012.

[10] M. T. Mason, “Progress in Nonprehensile Manipulation,” *The International Journal of Robotics Research*, vol. 18, no. 11, pp. 1129–1141, Nov. 1999.

[11] N. Chavan Dafle, A. Rodriguez, R. Paolini, B. Tang, S. S. Srinivasa, M. Erdmann, M. T. Mason, I. Lundberg, H. Staab, and T. Fuhlbrigge, “Extrinsic dexterity: In-hand manipulation with external forces,” in *Proc. of IEEE International Conference on Robotics and Automation (ICRA)*, May 2014, pp. 1578–1585.

[12] C. Eppner and O. Brock, “Planning Grasp Strategies That Exploit Environmental Constraints,” in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 4947–4952.

[13] K. Hang, A. S. Morgan, and A. M. Dollar, “Pre-Grasp Sliding Manipulation of Thin Objects Using Soft, Compliant, or Underactuated Hands,” *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 662–669, Apr. 2019.

[14] H. Marino, M. Ferrati, A. Settimi, C. Rosales, and M. Gabiccini, “On the Problem of Moving Objects With Autonomous Robots: A Unifying High-Level Planning Approach,” *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 469–476, Jan. 2016.

[15] N. Chavan-Dafle and A. Rodriguez, “Prehensile pushing: In-hand manipulation with push-primitives,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 6215–6222.

[16] M. Dogar and S. Srinivasa, “A framework for push-grasping in clutter,” *Robotics: Science and Systems VII*, 2011.

[17] M. C. Koval, N. S. Pollard, and S. S. Srinivasa, “Pre- and post-contact policy decomposition for planar contact manipulation under uncertainty,” *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 244–264, Jan. 2016.

[18] C. Phillips-Grafflin and D. Berenson, “Planning and Resilient Execution of Policies For Manipulation in Contact with Actuation Uncertainty,” *Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2016.

[19] A. Sieverling, C. Eppner, F. Wolff, and O. Brock, “Interleaving Motion in Contact and in Free Space for Planning Under Uncertainty,” in *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Vancouver, 2017, pp. 4011–4017.

[20] C. Eppner and O. Brock, “Visual Detection of Opportunities to Exploit Contact in Grasping Using Contextual Multi-Armed Bandits,” in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Sep 2017, pp. 273–278.

[21] M. Catalano, G. Grioli, E. Farnioli, A. Serio, C. Piazza, and A. Bicchi, “Adaptive synergies for the design and control of the pisa/iit softhand,” *The International Journal of Robotics Research*, vol. 33, no. 5, pp. 768–782, 2014.

[22] R. Deimel and O. Brock, “Soft Hands for Reliable Grasping Strategies,” in *Soft Robotics, Soft Hands, Grasping*. Berlin Heidelberg: Springer-Verlag, 2015, pp. 211–221.

[23] E. I. Corwin, H. M. Jaeger, and S. R. Nagel, “Structural signature of jamming in granular media,” *Nature*, vol. 435, no. 7045, pp. 1075–1078, June 2005.

[24] K.-H. Park, Y.-H. Jung, and T.-Y. Kwak, “Effect of Initial Granular Structure on the Evolution of Contact Force Chains,” *Applied Sciences*, vol. 9, no. 22, p. 4735, Jan. 2019.

[25] C. Eppner, R. Deimel, J. lvarez Ruiz, M. Maertens, and O. Brock, “Exploitation of environmental constraints in human and robotic grasping,” *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 1021–1038, June 2015.

[26] H. M. Jaeger, S. R. Nagel, and R. P. Behringer, “Granular solids, liquids, and gases,” *Reviews of modern physics*, vol. 68, no. 4, p. 1259, 1996.

[27] R. Deimel and O. Brock, “A novel type of compliant and underactuated robotic hand for dexterous grasping,” *The International Journal of Robotics Research*, vol. 35, no. 1-3, pp. 161–185, 2016.

[28] R. Cross, “Measurements of the horizontal coefficient of restitution for a superball and a tennis ball,” *American Journal of Physics*, vol. 70, no. 5, pp. 482–489, 2002.

[29] G. Zller, V. Wall, and O. Brock, “Active Acoustic Contact Sensing for Soft Pneumatic Actuators,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. Paris, France: IEEE, May 2020, pp. 7966–7972.

[30] H. Mnyusiwalla, P. Triantafyllou, P. Sotiropoulos, M. A. Roa, W. Friedl, A. M. Sundaram, D. Russell, and G. Deacon, “A bin-picking benchmark for systematic evaluation of robotic pick-and-place systems,” *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1389–1396, 2020.