

# ShapeGrasp: Zero-Shot Object Manipulation with Large Language Models through Geometric Decomposition

Samuel Li, Sarthak Bhagat, Joseph Campbell, Yaqi Xie, Woojun Kim, Katia Sycara, Simon Stepputtis

**Abstract**—Task-oriented grasping of unfamiliar objects is a necessary skill for robots in dynamic in-home environments. Inspired by the human capability to grasp such objects through intuition about their shape and structure, we present a novel zero-shot grasping method leveraging a geometric decomposition of the target object into simple, convex shapes that we represent in a graph structure, including geometric attributes and spatial relationships. Our approach employs minimal essential information – the object’s name and the intended task – to facilitate zero-shot task-oriented grasping. We utilize the commonsense reasoning capabilities of large language models to dynamically assign semantic meaning to each decomposed part and subsequently reason over the utility of each part for the intended task. Through extensive experiments on a real-world robotics platform, we demonstrate that our grasping approach’s decomposition and reasoning pipeline is capable of selecting the correct part in 92% of the cases and successfully grasping the object in 82% of the tasks we evaluate. Additional videos, experiments, code, and data are available on our project website: <https://shapegrasp.github.io/>.

## I. INTRODUCTION

In-home environments present a significant challenge for real-world robotics, primarily due to their highly unstructured nature. As a result, these environments frequently contain novel objects not present in the robot’s training environment; however, interacting with such objects, particularly grasping them in a task-dependent manner, is a necessary skill. Grasping an object in a way that facilitates a certain task, namely *task-oriented grasping*, requires a system to not only detect an object but also to reason over the utility of its parts. For example, when picking up a hammer with the goal to “hand it over” (see Fig. 2, left), it should be grasped by the head to facilitate a safe handover. Large Language Models (LLMs) provide the capability of such commonsense reasoning and can be utilized for task-oriented grasping with only a minimal set of contextual information, namely the object’s name and desired task [1], [2]. However, zero-shot task-oriented grasping remains challenging, particularly since current approaches are computationally expensive and may require additional information beyond the object’s name and intended task for high performance, such as desired object part names [1], [2], thus limiting their zero-shot performance.

In this work, we propose *ShapeGrasp*, a **robust grasping framework based on a geometric decomposition of the target object and shape-based semantic part reasoning leveraging LLMs, capable of conducting task-oriented grasping of novel objects**. This approach is inspired by the

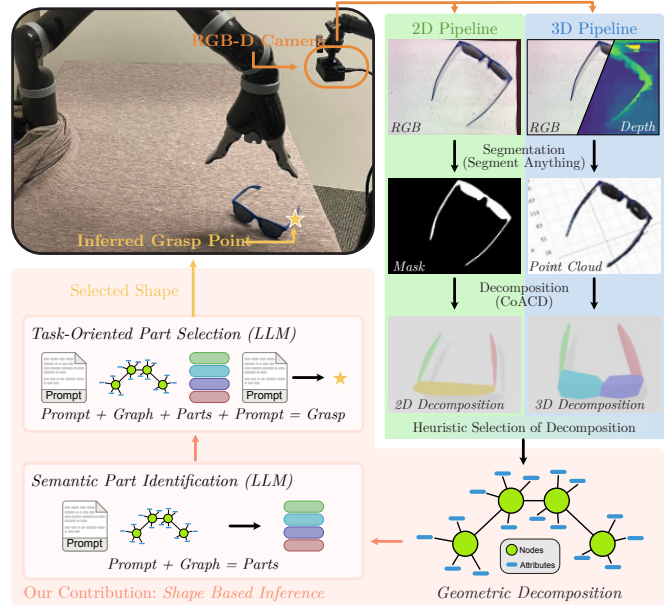


Fig. 1: **The ShapeGrasp Pipeline:** Given a target object, our RGB+D-based approach decomposes the object into basic convex parts. We propose a heuristic approach to decide which decomposition to use before converting it into a shape graph, allowing an LLM to utilize its commonsense reasoning to identify part semantics and task suitability.

human ability to interact with a novel object by analyzing its geometric composition, relating it to prior knowledge, and inferring each part’s utility [3] before utilizing this knowledge to identify a suitable part for an intended task. Unlike prior approaches, we introduce a multi-step reasoning approach that leverages a graph of basic shapes that compose the object and a multi-stage LLM prompt that a) assigns semantic meaning to each part of the object given its name and b) reasons over the task utility of each part to select the most suitable part for the specific task.

Figure 1 provides a high-level overview of the functionality proposed in ShapeGrasp. Starting with a passive monocular RGB+D image, we segment the target object using a pre-trained Segment Anything Model (SAM) [4], receiving a 2D mask image and the object’s relevant point cloud by applying the mask to the depth image. We then calculate two approximate convex decompositions using a pre-trained CoACD [5]: 1) a 3D decomposition using the masked point cloud and 2) a 2D decomposition using the image’s SAM mask. ShapeGrasp heuristically identifies

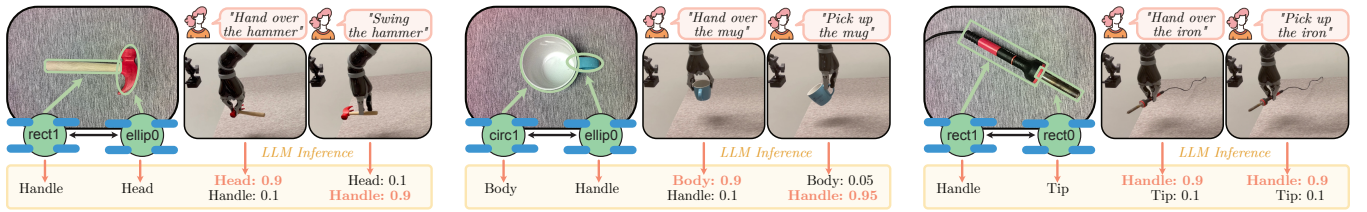


Fig. 2: Different resulting grasps given our shape-based inference pipeline. Parts in orange-boldface are ultimately grasped. Green circles and blue lines represent the part-graph decomposition with each entity’s associated attributes.

a suitable object decomposition, approximates each convex hull with a basic geometric shape, including a combination of rectangles, circles, triangles, and ellipses, and represents these parts and their spatial relationships as a graph. We then utilize an LLM to determine the semantic meaning of each shape given the name of the object and identify a suitable part to grasp using the assigned fine-grained semantic part labels and desired task. Through this process, inspired by the Chain of Thought [6] approach, we improve the reasoning capabilities of our method. Finally, the object is grasped at the identified part by calculating the centroid and principle components of its masked point cloud to determine the grasping location and orientation. Our work provides the following contributions:

- We introduce *ShapeGrasp* – a novel approach for zero-shot task-oriented grasping leveraging an object decomposition by creating a graph of basic geometric shapes, allowing for improved reasoning over each part’s task-oriented utility.
- Our proposed method is a lightweight approach in both the graph construction stage and reasoning stage, compared to other zero-shot grasping methods, utilizing only basic information and a single RGB+D image.
- Through extensive experiments on a robot platform, we demonstrate that our proposed approach outperforms current state-of-the-art methods.

## II. RELATED WORKS

Robotic grasping [7] aims to determine the best way to grasp objects and conventionally includes two approaches: the analytical approach [8], [9] and the data-driven approach [10], [11], [12]. The **analytical approach** aims to identify an appropriate grasp pose through analytic models that consider geometric conditions. For instance, [9] analyzes the geometry of the point cloud and identifies the appropriate grasping points based on a set of geometric conditions. Prior works have shown that understanding geometric information dramatically benefits robotic grasping [8], [9], as well as tasks including 3D geometry reasoning [13] and building synthetic tools [14]; however, these approaches require the availability of 3D geometry, which is very computationally expensive, and are sensitive to noise, thereby, restricting their applicability in real-world environments. On the other hand, the **data-driven approach** [10], [11], [12], [15] train models to predict grasp points. Prior works propose end-to-end grasp detection networks for partial, noisy point clouds [11], [12]

and for leveraging semantic information [10] in a supervised manner. In addition, techniques based on leveraging a semantic knowledge graph [16], shape segmentation techniques [17], and a physics simulator [18] have been considered. These approaches are limited in their ability to generalize to unseen objects due to requiring computationally expensive training.

Besides the approaches mentioned above, LLMs and VLMs have recently been successfully utilized in robotics grasping tasks due to their reasoning ability [19], [2], [20], [21]. Closest to our work is LERF-TOGO [2], which utilizes a vision-language model in a zero-shot fashion to output a distribution of potential grasping points given a natural language query indicating the object as well as the part-name that should be grasped. Specifically, LERF-TOGO reconstructs a 3D object mask based on DINO embeddings [22] and then uses it to generate grasping points via language models with conditional LERF [23]. While LERF-TOGO requires detailed natural language queries containing the object and part name, as well as several views of the scenes to reconstruct the 3D mask, our approach, *ShapeGrasp*, is able to generate grasping points with minimal semantic information about the object, based solely on a single RGBD top-down view of the scene. This is achieved by decomposing the object into a graph containing geometric shapes, their spatial relationships, and attributes, before using the graph as a prompt for an LLM. Thus, our approach leverages the advantages of the analytic approach—*understanding geometric information* [24]—and the large models-based approach—*the zero-shot reasoning ability* [25], which have been shown to be effective.

LLMs and VLMs have successfully been applied to the robotics domain [26], [27], [28], including task-oriented grasping [2], [29]. In such context, LLMs have been integrated into planning procedures in various ways, such as providing semantic knowledge [26] and performing complex reasoning in the form of an inner monologue [28]. In addition, [27] leverages the zero-shot reasoning capabilities of VLMs for the pick-and-drop task requiring object detection, navigation, and grasping. Despite the benefits of utilizing LLMs and VLMs, which include no need for additional training and providing common-sense reasoning capability, naive utilization of LLMs and VLMs still have limitations stemming from their inherent shortcomings, such as indecisiveness, lack of domain knowledge, hallucination, and the black-box problem [30]. To address these limitations and en-

hance reasoning, we infuse LLMs with structured knowledge in the form of symbolic graphs that capture the geometry, such as the shape of the object, and the relationships between different parts of the object. This infusion has been shown to be effective in preventing LLMs from deviating into the realm of fictitious information, thereby ensuring a connection to factual data [24], [31], [30], [32], thus enhanced reasoning capabilities for task-oriented grasping.

### III. SHAPE-BASED GRASPING

In this section, we introduce ShapeGrasp, our approach to zero-shot task-oriented grasping of novel objects by leveraging a graph of basic geometric shapes that compose the object. Given an RGB+D input image  $I \in \mathbb{R}^{H \times W \times C}$ , our approach  $g, \theta = f_{SG}(I)$  estimates a grasp location  $g \in \mathbb{R}^3$  and rotation  $\theta \in \mathbb{R}^1$  for the robot to pick up the object.

The function  $f_{SG}(\dots)$  represents our modular grasping pipeline, *ShapeGrasp*, composed of the following modules: First, we introduce our approach to segmenting and retrieving convex decompositions of the target object (see Sec. III-A) by using a pipeline of pre-trained models. Then, we discuss our first contribution, selecting a suitable decomposition through an automatic heuristic (see Sec. III-B). Finally, we introduce our novel approach to reasoning over the geometric composition of the target object by using a graph-based representation of the object decomposition and utilizing an LLM for the reasoning process (see Sec. III-C).

#### A. Segmenting the Object

In this section, we introduce our image-processing pipeline utilizing SAM [4] to retrieve an object mask and CoACD [5] to generate a convex decomposition of the object.

1) *Retrieving the Object Mask and Point Cloud*: To obtain the segmentation mask for input image  $I$ , we utilize user input defining a set of “in-points” on the target object, used by *Segment Anything Model* (SAM) to retrieve the object’s full 2D segmentation mask  $M$ . Following this step, in-points are discarded, and a point cloud  $\mathcal{P}$  of the object is retrieved utilizing the mask and depth information in our input image  $I$ . We posit that the particular method for retrieving such object masks is not central to our overall approach and can be facilitated by various alternative approaches.

2) *Convex Decomposition with CoACD*: Given an object mask  $M$  and point-cloud  $\mathcal{P}$ , we independently retrieve a convex decomposition for each of the two inputs. To this end, we utilize CoACD, which is a recent approach for the convex decomposition of 3D meshes, specifically designed to retain fine-grained object features, which are important to preserving the original functionality, particularly in interactive settings. To utilize this approach, we convert the 2D mask  $M$  into a 3D mesh by interpreting it as a plane and retrieving a mesh from point cloud  $\mathcal{P}$  through voxelization. The depth data utilized in the 3D pipeline allows for more intricate decomposition with CoACD, uncovering features at various elevations within the object that might be overlooked by the 2D method. However, simultaneously, 2D decompositions provide a fast approximation that can be beneficial,

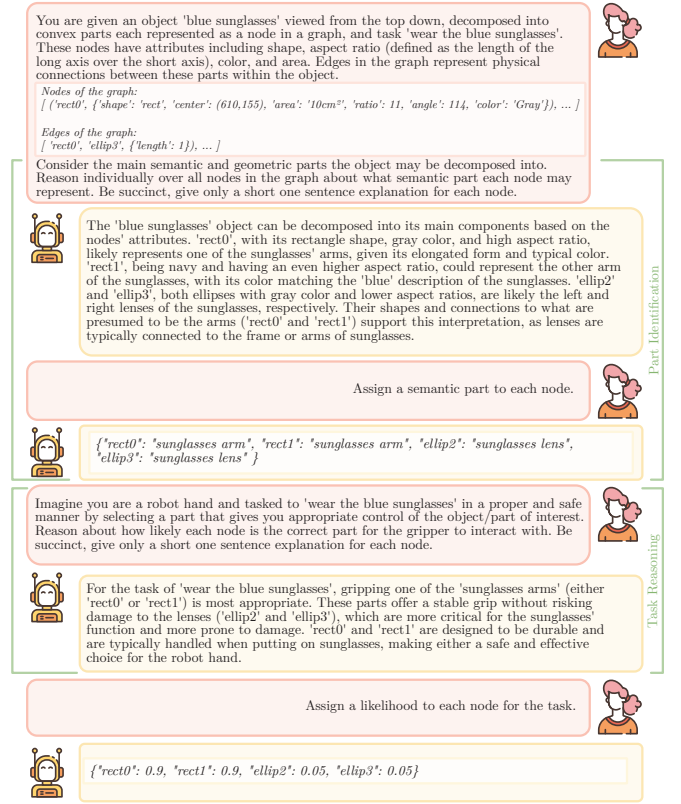


Fig. 3: ShapeGrasp Prompting: Given a geometric decomposition graph, we infer a suitable grasping point through a chain of four consecutive prompts, two for task-agnostic part identification, and two for task reasoning. Ablations of these prompts can be found in Tab. III.

particularly for concave objects like mugs. We compute two separate convex decompositions, one from the 2D mask and one from the 3D point cloud, termed  $\mathcal{C}_{2D}$  and  $\mathcal{C}_{3D}$ , respectively, as each decomposition exhibits a set of different desirable properties. Selecting the appropriate decomposition depends on the object the system interacts with. In the next section, Sec. III-B, we discuss our automated heuristic utilized for this purpose.

#### B. Geometric Decomposition

In this section, we discuss the following contributions: a) our heuristic approach to select which decomposition to use for a particular object, and b) how we create an object-graph  $\mathcal{G}$  that relates shapes with each other and stores the attributes of each part.

1) *Heuristic Selection of Decomposition*: With our heuristic, we address the problem that certain objects (e.g., concave mugs) are inherently unsuitable for convex decomposition, while other objects may have little visual features or reflective surfaces that result in low-confidence depth maps. To address this problem, we introduce a heuristic  $\mathcal{C}^* = h(\mathcal{C}_{2D}, \mathcal{C}_{3D})$  to choose between the two decompositions that ultimately provides a singular decomposition  $\mathcal{C}^*$ . Fundamentally, we lower the decomposition threshold until both the 2D and 3D pipelines result in more than a single part. After



that, our heuristic  $h(\dots)$  chooses a preferred decomposition given the following criteria: If the 3D decomposition  $\mathcal{C}_{3D}$  results in more than a set threshold of parts,  $\omega$ , or if the percentage of depth points,  $\alpha$ , with high confidence required to retain the 3D decomposition is too low, our heuristic  $h(\dots)$  chooses the 2D decomposition  $\mathcal{C}_{2D}$ . Formally, our heuristic which selects between 2D and 3D decompositions at valid thresholds can be defined as follows:

$$h(\mathcal{C}_{2D}, \mathcal{C}_{3D}) = \begin{cases} \mathcal{C}_{3D}, & \text{if } |\mathcal{C}_{3D}| \leq \omega \wedge \text{conf.} \geq \alpha \\ \mathcal{C}_{2D}, & \text{otherwise.} \end{cases} \quad (1)$$

where  $|\mathcal{C}|$  is the number of convex parts found in the respective decomposition of the object. Section IV-A.1 provides empirical evidence for our default  $\gamma$  value. With the selected heuristic  $\mathcal{C}^*$ , we create an object-graph  $\mathcal{G}$  to represent the object as a composition of multiple basic geometries.

2) *Structured Object-Graph Creation*: After selecting an appropriate decomposition  $\mathcal{C}^*$ , we project it back onto the original input image  $I$  and create an object-graph  $\mathcal{G}$  describing the composition of the target object. Each segmented part is represented as a node in the graph, accompanied by attributes for each shape derived from the segmented image. The primary attribute is an approximating shape primitive, chosen from an isosceles triangle, rectangle, circle, or ellipse. To select an appropriate shape primitive, we approximate each convex hull with a simplified polygon given a pre-defined error threshold  $\varepsilon$ . The resulting points from this simplification dictate the fitted shape as follows:

- **Isosceles triangles** are formed by modifying any three points to equalize the leg lengths and adjust the base.
- **Rectangles** are formed by finding the rotated rectangle of the minimum area enclosing part.
- **Circles** are identified by measuring the shape factor, or the ratio of the part’s area to the area of the bounding circle, with a threshold of 0.9 for circle classification.
- **Ellipses** are formed by fitting the part inside a rectangle (see above) if the fit reduces errors further.

These geometric shapes allow for the determination and inclusion of additional attributes, such as aspect ratio, calculated based on the long and short sides of each shape (or major/minor axis) and the angle of the long side, as well as the centroid. Object color is derived by bucketing the RGB color spectrum based on the standard 16 web colors and selecting the most prevalent color as the node attribute. Edges within graph  $\mathcal{G}$  are drawn to connect nodes whose convex parts share boundaries or intersect in the segmentation, with the shared boundary length as an edge attribute.

### C. Grasp Inference through Shape Reasoning

To select an appropriate grasping point for the object, we propose to leverage the common-sense knowledge encoded in the LLM across two interaction stages, inspired by the Chain of Thought approach [6], further improving our approach’s reasoning capabilities: 1) semantic reasoning over each shape described in the graph  $\mathcal{G}$  and 2) selecting the most appropriate shape that facilitates the task-oriented grasp.



Fig. 4: Overview of the 38 objects used in our study, inspired by the objects introduced in LERF-TOGO [2].

We leverage a prompt template, described in Fig. 3, that depends only on the target object and desired task. To ensure the intended output structure, we utilize TypeChat [33] in addition to the presented prompts.

1) *Semantic Part Identification*: In the first stage, given the target object name, task, and constructed graph, the LLM is instructed to reason about the subparts of the object described in the graph and what semantic part (e.g., “handle” and “blade”, for a knife) each represents. To conduct this reasoning, we first ask for an unstructured, free-form answer in which the LLM explains its thought. As a follow-up to this response, the LLM is tasked to assign a single semantic part to each graph node in a structured manner.

2) *Task-Oriented Part Selection*: After the semantic reasoning is complete, the LLM is instructed to reason about the task utility of each part, given its graph representation and assigned semantics from the first stage. Similar to the first stage, this is accomplished in two steps: a free-form reasoning and explanation stage, which the LLM then uses to assign a final task-conditioned suitability score to each node, which determines the selected grasp.

### D. Selecting a Grasp Pose

In the final step, we select the shape with the highest predicted score. To facilitate the grasp, we calculate the centroid of the chosen part and derive its corresponding 3D coordinates from the depth information of the input image  $I$ . To consider rotations, we calculated the principle components of the masked subpart in the point cloud and grasp along the largest component.

## IV. EXPERIMENTS

We evaluate our approach in real-world experiments and demonstrate its effectiveness in grasping a diverse range of household objects across a variety of tasks. Figure 4 presents the 38 objects from 12 different categories included in our experiments. Section IV-A discusses our heuristic decomposition pipeline with regards to both, the dynamic selection of the decomposition threshold (Sec. IV-A.1), as well as the 2D vs 3D pipelines (Sec. IV-A.2). Section IV-B demonstrates our approach on a real-world robotic platform

and compares it against state-of-the-art baselines, as well as additional qualitative experiments in section IV-B.1 and an LLM ablation study in section IV-B.2. All of our experiments are conducted with a Kinova Jaco robotic arm equipped with a three-finger gripper, coupled with a fixed Oak-D SR passive stereo-depth camera for RGB and depth perception.

**Pipeline Configuration:** In order to retrieve convex decompositions of sufficient quality, we empirically set the threshold  $\omega$ , deciding when to select the 2D decomposition over the 3D decomposition to  $\omega = 10$ . Additionally, we set parameter  $\alpha$ , defining the maximal percentage of depth points that can exhibit low confidence, to  $\alpha = 15\%$ . While these values have been set empirically, we note that  $\alpha$  is not very sensitive and does not require intricate tuning.

**Dataset:** We created a dataset of 38 objects covering 12 general categories of items and 49 tasks, as shown in Fig. 4, inspired by the objects and tasks used in LERF-TOGO [2].

**Metrics:** To evaluate the effectiveness of our approach, we employ three metrics: “Part Identification”, “Part Selection”, and “Successfully Lifted”. “Part Identification” measures the accuracy of the semantic part assignment, across all parts globally and across only the ground truth parts. “Part Selection” quantifies the proficiency of our model to select the correct part for the task-oriented grasp. For this metric, the selection is deemed correct if it aligns with ground-truth parts based on LERF-TOGO [2]’s part queries and common-sense human judgment with respect to the task, safety, and stability. “Lift Success” indicates the percentage of successfully lifted objects at the ShapeGrasp selected part for the given task. The failure number is the number of parts incorrectly identified during part selection, plus failed grasps at parts that were correctly selected. For this metric, we count a lift as failed if the “Part Selection” was incorrect in the first stage.

#### A. Automatic Geometric Decomposition

1) *Dynamic Threshold Selection:* In both the 2D and 3D pipeline, the convex decomposition error threshold  $\gamma$  is an important element that facilitates grasping success by controlling the number of segmented object parts. This threshold plays a pivotal role in balancing the accuracy of the decomposition with the manageability of the resulting parts. A high error threshold  $\gamma > 0.2$  that does not over-segment a geometrically complex object may fail to decompose a simple object – the entire object may be approximated as a single convex part. Conversely, a low error threshold  $\gamma < 0.1$  may decompose complex objects into overly many parts, complicating the resulting graph and making reasoning challenging. As shown in Figure 5, we run the pipeline at thresholds between 0.01 and 0.35 and evaluate the number of decomposed parts and the lifting success rate on sunglasses and screwdriver, with the former being a more complex object as evidenced by the consistent larger number of resulting parts at each threshold. This experiment reveals that any single threshold may not be sufficient to accommodate all objects. At thresholds conducive to successful lifting for the sunglasses, the screwdriver often fails to decompose;

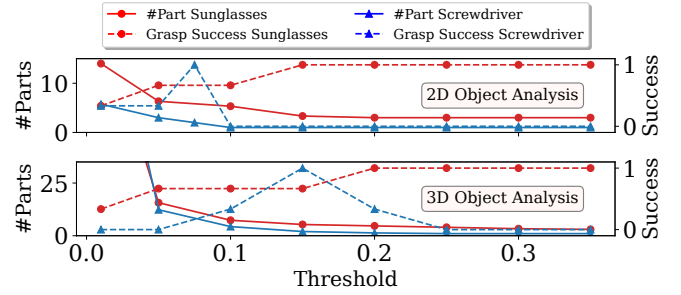


Fig. 5: Threshold analysis for 2D (top)/3D (bottom) decompositions on sunglasses (complex) and screwdriver (simple).

naively taking the centroid of the entire object as the grasp point is inherently unsafe and does not allow room for reasoning, which we consider a failure case. While the error threshold may easily be set and tuned in our pipeline by a user, our heuristic fully automates the inference process.

Our pipeline provides a generalizable zero-shot approach for grasping by balancing object complexities and geometries. Therefore, we propose the selection algorithm described in Section III-B that considers the number of decomposed parts to set an object-specific threshold  $\gamma$ . An initial threshold is set at  $\gamma = 0.2$  for 3D meshes and  $\gamma = 0.15$  for 2D meshes. The error allowance for 3D objects should naturally be higher to accommodate potential noise in the depth perception. These thresholds are set conservatively and tend to under-decompose objects. To address this, we iteratively decrease the threshold by 0.025 for each object until a valid decomposition is achieved (an object is decomposed into two or more parts). For example, a valid decomposition and successful lift are achieved within 2 – 3 iterations for the screwdriver, while the initial threshold is satisfactory for sunglasses. We further find that in our experiments, high thresholds that do not decompose an object execute in less than a second, making this search both efficient and effective.

2) *2D vs 3D Decompositions:* After setting decomposition thresholds, the choice between the 2D and 3D pipelines for the final decomposition from which to build the graph is a critical step. As discussed in Section III-B, the heuristic algorithm we employ uses the depth confidence and number of decomposed parts (too many parts may indicate noisy or concave surfaces) to facilitate this choice. Table I shows the “Part Selection” performance across all object-task pairs between the 2D, 3D, and heuristic that selects between them, each of which uses the optimal threshold search. An interesting observation is that while our heuristic selection does indeed result in the best performance (92%), the 2D pipeline exhibits stronger performance (86%) than the 3D pipeline (83%). While this may be initially counterintuitive, this result demonstrates and is consistent with the conclusion that the flexibility of our pipeline allows us to dynamically adapt to settings where depth information may be inconsistent, low quality, or incompatible with the convex decomposition method. In our real-world experiments, we utilize an Oak-D SR passive depth camera, which provides high-quality

Model	(using GPT-4)	Part Selection
1 ShapeGrasp (2D only)		0.86
2 ShapeGrasp (3D only)		0.73
3 ShapeGrasp (Heuristic, no obj)		0.51
4 ShapeGrasp (Heuristic)		<b>0.92</b>

TABLE I: Part selection accuracy on different variants of ShapeGrasp. We evaluate 2D-only, 3D-only, and heuristic decompositions. For the heuristic, we also explore a “no obj” variant that omits the object name from the prompt and task.

Model	Part ID	Part Sel.	Success	Time
1 GraspGPT [20]	N/A	0.37	0.31	150
2 GPT4-Vision [34]	N/A	0.82	0.73	20
3 ShapeGrasp (Starling)	0.54 (0.63)	0.65	0.57	25
4 ShapeGrasp (GPT-4)	<b>0.84 (0.90)</b>	<b>0.92</b>	<b>0.82</b>	30

TABLE II: Results on ShapeGrasp compared to GraspGPT and GPT4-V baselines. Part ID is the semantic “Part Identification” accuracy across all parts (and across only the ground truth parts), Part Sel. is the “Part Selection” accuracy, “Success” is the robot’s “Lift Success” for the task-oriented grasp, and “Time” is the typical inference time in seconds of each method.

depth estimates in well-lit environments on matte objects with noticeable elevated features—the same features we may care to segment for grasping—however, results deteriorate with concave, reflective, and transparent surfaces, where the 2D pipeline excels due to the sole reliance on RGB data. The prevalent noise in real-world settings often necessitates a reliance on 2D decomposition for accuracy and robustness via the depth confidence threshold, while still leveraging depth data in the 3D pipeline when it is confident and useful.

The threshold search, coupled with the automatic selection between the 2D and 3D pipelines, forms an algorithm that determines decomposition outputs that are reasonable and suitable for grasping.

### B. Zero-Shot Task-Oriented Grasping

For real-world task-oriented grasping experiments, we pair the ShapeGrasp pipeline with a Kinova Jaco arm designed to grasp and lift objects at the selected part (see Table II). This setup employs a straightforward method for selecting grasp poses, drawing directly from the attributes defined in our symbolic object part graph and masked point cloud (see III-C), with the goal of executing a top-down grasp.

We employ two baselines to evaluate the efficacy of our proposed approach, ShapeGrasp on the “Part Selection” and “Lift Success” metrics: **GraspGPT [20]**, which is a current state-of-the-art approach for zero-shot task-oriented grasping and **GPT4-V [34]**, a foundation model trained with internet-scale data with visual input modality, prompted with language instructions to select the correct task-oriented part.

Our empirical findings indicate a significant performance advantage of our method over GraspGPT [20], underscoring the efficacy of our structured, symbolic object part graph

in conjunction with LLM reasoning. The performance gain using our pipeline over GraspGPT [20] is 55% and 51% (see rows 1 and 4 in Table II) for the “Part Selection” and “Lift Success” metrics, respectively. We further analyze the discrepancy between ShapeGrasp and GraspGPT results. With respect to GraspGPT results, we note that it depends on GraspNet [29], and certain objects may be inaccurate or fail in the grasp sampling stage when the depth quality is noisy or poor, which may occur due to our static monocular depth camera. GraspGPT is also limited to tasks and objects related to previously known concepts from their training set. ShapeGrasp demonstrates robustness to the same noisy static depth inputs, while featuring zero-shot performance that is more lightweight than GraspGPT (see runtime numbers in Table II).

An important hypothesis that motivates our vision pipeline that constructs the symbolic object graph is that directly processing object part features and spatial relationships, and providing this information in a structured way for LLM reasoning, is more robust and performant than relying on VLMs for end-to-end reasoning. Though VLMs are considerably larger and more expensive models, performance on low-level features and relationships within parts of an object image may be unreliable and subject to hallucinations [35].

To test this hypothesis and directly compare our vision and reasoning pipeline to a VLM, we established a privileged GPT4-Vision baseline that benefits from the same object segmentations generated by the heuristic decomposition pipeline and skipping the graph-based reasoning stages (contributing to the faster runtime). This baseline is grounded by coloring each part and assigning an integer index label for clarity. We confirm GPT4-Vision’s capability to interpret segmented and grounded input images through a series of questions and human-verified responses. Grasp poses for the GPT4-Vision selected parts are determined in the same manner as a selected part to ensure comparability. Our method shows significant success rate gains over GPT4-Vision, by 10% and 9% on the evaluation metrics (see Tab. II).

We further test the modularity of ShapeGrasp by evaluating the full pipeline using Starling [36], a much smaller and more efficient open-source LLM, as the inference backend instead of GPT-4. Performance across all metrics, while lower than the larger and more powerful GPT-4, remains meaningful and higher than the GraspGPT [20] baseline, and also demonstrate the coupling between the “Part Identification” and “Part Selection” results.

1) *Qualitative Results:* The flexibility and generalizability of ShapeGrasp and LLM shape-inference allow us to explore more complex interactions that may be further enhanced by the incorporation of additional object and robot attributes. For example, Fig. 2 shows how LLM semantic reasoning over our shape graph enables effective execution of the “hand over” task. As the LLM understands which node in the graph corresponds to the handles in objects, it can prioritize that part for the human in the “handover” interaction. Additional attributes can also be incorporated and reasoned over in an object-specific way; while both the

You are a three-fingered gripper with max width X cm. Pick up the wine bottle.




Fig. 6: ShapeGrasp results incorporating width attributes and variable gripper width constraints.

“mug” and the “soldering iron” are given the attribute of being “hot”, the task-oriented reasoning stage is able to make the common-sense inference that the level of heat and risk exhibited by these two objects differ dramatically. For the handover task, while the mug is grasped by the hot body, which “minimizes the risk of spilling hot liquid and ensures a comfortable handover”, the soldering iron is still grasped by the handle, which “positions the hot tip away from both the robot and the human” (see Fig. 2).

Fig. 6 demonstrates how the task-oriented selection stage is also able to consider robot attributes. Here, a “width” attribute is added to the shape graph, along with a variable maximum gripper width. While the robot prefers to “pick up the wine bottle” by the bottle body, when the maximum gripper width becomes less than the body width, it switches to the narrower neck, maintaining a successful grasp.

2) *LLM Ablations*: To validate the efficacy of our model and the comprehensiveness of the LLM interaction stages, we conducted an ablation study with a focus on the “Part Selection” evaluation metric, as detailed in Table III. Applying the previously described decomposition selection algorithm utilizing both the 2D and 3D pipelines, we examine the impact of each LLM reasoning stage: the *semantic parts reasoning* stage and *task-oriented parts reasoning* stage that comes before the final task score assignments. We evaluate the results using neither reasoning stage (LLM directly assigns only task scores), only one reasoning stage, and the full reasoning procedure.

The ablation study indicates that the most effective performance is achieved when both part and task reasoning stages are employed in the LLM interaction (92%). The part identification stage is empirically the more important of the two stages, leading to a performance of (87%). This stage is likely important because when employed, the LLM’s final selection is forced to depend on its own semantic part identification. Intuitively, the semantics of a part are essential in determining its affordances and suitability for a task. This argument is further supported by the accuracy in the part identification stage (see Table II), 84% globally, and 90% across the ground truth parts. These numbers lead to an important observation: not every part in an object, or even the ground-truth part, needs to be correctly identified; it may be sufficient if even a single critical part to either grasp or

	Task Reasoning	Part Identification	Part Selection
1			0.43
2	✓		0.57
3		✓	0.84
4	✓	✓	<b>0.92</b>

TABLE III: Ablation study on the LLM reasoning within ShapeGrasp. While a part to grasp is always produced, reasoning can be done without identifying parts and/or without reasoning over the task. See Fig. 3 for detailed prompt information.

avoid is correctly identified. For example, in a pair of blue sunglasses, the ground truth part “arm” was misidentified as the “frame”, but the “lens”—which is critical to avoid—was correctly identified. As such, the task-oriented selection stage correctly avoided the “lens” segment. Put in another way, the quality of the automatic geometric decomposition and the accuracy of the parts reasoning stage effectively provides another novelty: *semantic part segmentation*, which naturally should be coupled with task-oriented part selection.

The task reasoning demonstrates a notable performance gain as well (57%) over the no-reasoning result (43%); this may indicate that even once semantics are assigned, explicit task-conditioned reasoning is still beneficial, especially as a single object may have numerous appropriate uses and tasks.

In Table I, we also explore the performance when the target object name is withheld. This leads to a significant performance drop (51%), likely due to the challenge of assigning semantic parts to an unknown object. In this setting, the assigned semantics may not be meaningful, with the final results largely determined by geometric attributes; however, reasoning while not having any semantic object knowledge may be an interesting avenue for future work.

## V. CONCLUSION, LIMITATIONS, AND FUTURE WORK

In this work, we present ShapeGrasp, a novel approach that reasons over an object’s geometric composition of basic shapes to find a suitable semantic part and associated grasp pose that facilitates a particular task. Our novelty, namely the representation of the object as a graph of basic geometric shapes, allows an LLM to effectively reason over the semantic utility of object parts and ultimately select the most suitable part given a desired task. Through extensive experiments on real-world hardware, we demonstrated that our approach can efficiently utilize a single, static RGB+D camera image for zero-shot grasping of novel objects and outperform current state-of-the-art approaches.

We recognize that the synergy of convex part decomposition and the LLM’s ability to assign semantic meaning to each segment is a powerful feature of our method while, at the same time, introducing a dependency on a “reasonable” decomposition. While we mitigate this dependency through the introduction of our heuristic by automatically tuning the decomposition pipeline, further integration into alternative decomposition methods may further improve our



pipeline. In future work, we will investigate the utility of our approach with even less information, particularly when no semantic information about the object is provided (initial results in Tab. I, “no obj”), which may be enough to infer an object’s utility given only its decomposed parts. Such improved reasoning could be facilitated by elevating our shape-based decompositions into the third dimension and utilizing 3D-geometries. Further, we will investigate different LLM prompting approaches to further improve the model’s reasoning capabilities.

#### ACKNOWLEDGEMENTS

We want to acknowledge the support from DARPA under grant FA8750-23-2-1015, AFOSR under grants FA9550-18-1-0251 and FA9550-18-1-0097, and ARL under grant W911NF-19-2-0146 and W911NF-2320007.

#### REFERENCES

- [1] A. Murali, W. Liu, K. Marino, S. Chernova, and A. Gupta, “Same object, different grasps: Data and semantic knowledge for task-oriented grasping,” in *Conference on Robot Learning*, 2020.
- [2] A. Rashid, S. Sharma, C. M. Kim, J. Kerr, L. Y. Chen, A. Kanazawa, and K. Goldberg, “Language embedded radiance fields for zero-shot task-oriented grasping,” in *7th Annual Conference on Robot Learning*, 2023.
- [3] H. P. O. de Beeck, K. Torfs, and J. Wagemans, “Perceived shape similarity among unfamiliar objects and the organization of the human object vision pathway,” *Journal of Neuroscience*, vol. 28, no. 40, pp. 10 111–10 123, 2008.
- [4] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” *arXiv:2304.02643*, 2023.
- [5] X. Wei, M. Liu, Z. Ling, and H. Su, “Approximate convex decomposition for 3d meshes with collision-aware concavity and tree search,” *ACM Transactions on Graphics*, vol. 41, no. 4, p. 1–18, Jul. 2022. [Online]. Available: <http://dx.doi.org/10.1145/3528223.3530103>
- [6] J. Wei, X. Wang, D. Schuurmans, M. Bosma, F. Xia, E. Chi, Q. V. Le, D. Zhou *et al.*, “Chain-of-thought prompting elicits reasoning in large language models,” *Advances in neural information processing systems*, vol. 35, pp. 24 824–24 837, 2022.
- [7] H. Zhang, J. Tang, S. Sun, and X. Lan, “Robotic grasping from classical to modern: A survey,” *ArXiv*, vol. abs/2202.03631, 2022.
- [8] A. Ten Pas, M. Gualtieri, K. Saenko, and R. Platt, “Grasp pose detection in point clouds,” *The International Journal of Robotics Research*, vol. 36, no. 13–14, pp. 1455–1473, 2017.
- [9] B. S. Zapata-Impata, P. Gil, J. Pomares, and F. Torres, “Fast geometry-based computation of grasping points on three-dimensional point clouds,” *International Journal of Advanced Robotic Systems*, 2019.
- [10] E. Jang, S. Vijayanarasimhan, P. Pastor, J. Ibarz, and S. Levine, “End-to-end learning of semantic grasping,” *arXiv preprint arXiv:1707.01932*, 2017.
- [11] B. Zhao, H. Zhang, X. Lan, H. Wang, Z. Tian, and N. Zheng, “Regnet: Region-based grasp network for end-to-end grasp detection in point clouds,” in *2021 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2021, pp. 13 474–13 480.
- [12] A. Alliegro, M. Rudorfer, F. Frattin, A. Leonardis, and T. Tommasi, “End-to-end learning to grasp via sampling from object point clouds,” *IEEE Robotics and Automation Letters*, vol. 7, no. 4, 2022.
- [13] D. Paschalidou, L. V. Gool, and A. Geiger, “Learning unsupervised hierarchical part decomposition of 3d objects from a single rgb image,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1060–1070.
- [14] Z. Liu, W. T. Freeman, J. B. Tenenbaum, and J. Wu, “Physical primitive decomposition,” in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 3–19.
- [15] A. Agarwal, S. Uppal, K. Shaw, and D. Pathak, “Dexterous functional grasping,” in *Conference on Robot Learning*, 2023.
- [16] A. Murali, W. Liu, K. Marino, S. Chernova, and A. Gupta, “Same object, different grasps: Data and semantic knowledge for task-oriented grasping,” in *Conference on robot learning*. PMLR, 2021.
- [17] Y. Lin, C. Tang, F.-J. Chu, and P. A. Vela, “Using synthetic data and deep networks to recognize primitive shapes for object grasping,” in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2020, pp. 10 494–10 501.
- [18] K. Fang, Y. Zhu, A. Garg, A. Kurenkov, V. Mehta, L. Fei-Fei, and S. Savarese, “Learning task-orientative grasping for tool manipulation from simulated self-supervision,” *The International Journal of Robotics Research*, vol. 39, pp. 202 – 216, 2018.
- [19] R. Mirjalili, M. Krawez, S. Silenzi, Y. Blei, and W. Burgard, “Langrasp: Using large language models for semantic object grasping,” *arXiv preprint arXiv:2310.05239*, 2023.
- [20] C. Tang, D. Huang, W. Ge, W. Liu, and H. Zhang, “Graspgpt: Leveraging semantic knowledge from a large language model for task-oriented grasping,” *IEEE Robotics and Automation Letters*, 2023.
- [21] C. Tang, D. Huang, L. Meng, W. Liu, and H. Zhang, “Task-oriented grasp prediction with visual-language inputs,” *arXiv preprint arXiv:2302.14355*, 2023.
- [22] M. Caron, H. Touvron, I. Misra, H. Jégou, J. Mairal, P. Bojanowski, and A. Joulin, “Emerging properties in self-supervised vision transformers,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2021, pp. 9650–9660.
- [23] J. Kerr, C. M. Kim, K. Goldberg, A. Kanazawa, and M. Tancik, “Lerf: Language embedded radiance fields,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023.
- [24] W. Ding, S. Feng, Y. Liu, Z. Tan, V. Balachandran, T. He, and Y. Tsvetkov, “Knowledge crosswords: Geometric reasoning over structured knowledge with large language models,” *ArXiv*, vol. abs/2310.01290, 2023.
- [25] J. Huang and K. C.-C. Chang, “Towards reasoning in large language models: A survey,” *ArXiv*, vol. abs/2212.10403, 2022.
- [26] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman *et al.*, “Do as i can, not as i say: Grounding language in robotic affordances,” *arXiv preprint arXiv:2204.01691*, 2022.
- [27] P. Liu, Y. Orru, C. Paxton, N. M. M. Shafuallah, and L. Pinto, “Ok-robot: What really matters in integrating open-knowledge models for robotics,” *arXiv preprint arXiv:2401.12202*, 2024.
- [28] W. Huang, F. Xia, T. Xiao, H. Chan, J. Liang, P. Florence, A. Zeng, J. Tompson, I. Mordatch, Y. Chebotar *et al.*, “Inner monologue: Embodied reasoning through planning with language models,” in *Conference on Robot Learning*. PMLR, 2023, pp. 1769–1782.
- [29] A. Mousavian, C. Eppner, and D. Fox, “6-dof grasnet: Variational grasp generation for object manipulation,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019.
- [30] S. Pan, L. Luo, Y. Wang, C. Chen, J. Wang, and X. Wu, “Unifying large language models and knowledge graphs: A roadmap,” *ArXiv*, vol. abs/2306.08302, 2023.
- [31] F. Moiseev, Z. Dong, E. Alfonseca, and M. Jaggi, “SKILL: Structured knowledge infusion for large language models,” in *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2022.
- [32] L. F. Yang, H. Chen, Z. Li, X. Ding, and X. Wu, “Chatgpt is not enough: Enhancing large language models with knowledge graphs for fact-aware language modeling,” *ArXiv*, vol. abs/2306.11489, 2023.
- [33] Microsoft. [Online]. Available: <https://microsoft.github.io/TypeChat/>
- [34] OpenAI, “Gpt-4 technical report,” 2023.
- [35] M. Yuksekgonul, F. Bianchi, P. Kalluri, D. Jurafsky, and J. Zou, “When and why vision-language models behave like bags-of-words, and what to do about it?” in *The Eleventh International Conference on Learning Representations*, 2023.
- [36] B. Zhu, E. Frick, T. Wu, H. Zhu, and J. Jiao, “Starling-7b: Improving llm helpfulness harmfulness with rlai,” November 2023.