

NL2Contact: Natural Language Guided 3D Hand-Object Contact Modeling with Diffusion Model

Zhongqun Zhang¹, Hengfei Wang¹, Ziwei Yu², Yihua Cheng^{1*},
Angela Yao², and Hyung Jin Chang¹

¹ University of Birmingham, UK

² National University of Singapore, Singapore

Abstract. Modeling the physical contacts between the hand and object is standard for refining inaccurate hand poses and generating novel human grasp in 3D hand-object reconstruction. However, existing methods rely on geometric constraints that cannot be specified or controlled. This paper introduces a novel task of controllable 3D hand-object contact modeling with natural language descriptions. Challenges include i) the complexity of cross-modal modeling from language to contact, and ii) a lack of descriptive text for contact patterns. To address these issues, we propose NL2Contact, a model that generates controllable contacts by leveraging staged diffusion models. Given a language description of the hand and contact, NL2Contact generates realistic and faithful 3D hand-object contacts. To train the model, we build *ContactDescribe*, the first dataset with hand-centered contact descriptions. It contains multi-level and diverse descriptions generated by large language models based on carefully designed prompts (e.g., grasp action, grasp type, contact location, free finger status). We show applications of our model to grasp pose optimization and novel human grasp generation, both based on a textual contact description.

1 Introduction

Understanding and modeling physical contact between hands and objects [3, 4, 37, 47, 49] can contribute to applications in animation [23], virtual reality [44], augmented reality [39, 46], and robotics [32, 52]. The contact could also improve hand-object 3D pose and shape estimation, either reconstructed from images [10, 42, 53] or generated from given object models [18, 24, 27].

Conventional contact modeling methods typically apply PointNet [30] on hand-object point clouds to infer the contact status of each point. These methods leverage geometric constraints for modeling physically realistic contact. However, their results do not align well with human grasps [27]. As shown in Figure 1, the method ContactOpt [10], despite being state-of-the-art, produces a result in which all the fingers make contact with scissors. While mostly feasible physically, this is unrealistic and does not correspond to human usage of scissors.

* Corresponding author.

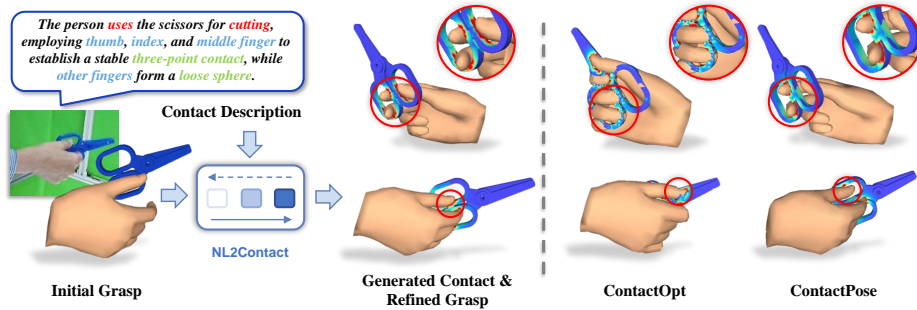


Fig. 1: Illustration of the NL2Contact setting. Our method learns to model 3D hand-object contact from language description. The modeling contact can be used to grasp pose optimization. By utilizing the predicted contact to refine the initial grasp, our performance surpasses that of ContactOpt [10] and ContactPose [4]. It’s worth noting that ContactPose is the annotation from the motion capture system.

Recently, there has been growing interest in synthesizing natural hand-object interactions based on object semantics and affordance, as well as human behavior and intent. Existing methods typically use intent labels, such as interaction verbs [28, 48] and object part-level affordance [17], to achieve controllable hand-object interaction modeling. Interaction verbs, like *use* and *handle* [48], describe the behavior and intent during the interaction but are typically high-level and vague to guide precise contact modeling. Affordance, labeled on 3D point clouds of objects, indicates feasible regions where interactions occur. However, the regions are coarse and require prior knowledge of object categories. These methods face significant challenges in precisely controlling hand-object contact modeling.

This paper introduces a novel task, natural language-guided 3D hand-object contact modeling. Inputting a natural language description of hand-object interaction, our task outputs physically realistic contacts aligned with the description, as shown in Figure 1. Compared to verbs and affordances, natural language descriptions provide freer and more precise expressions of contact patterns. To this end, we propose the *ContactDescribe* dataset, which provides 3D hand-object contact and fine-grained natural language descriptions. To our knowledge, this is the first dataset providing fine-grained descriptions for hand-object contact.

We create *ContactDescribe* using a novel text annotation scheme. Initially, we designed hand-centered prompts. Subsequently, we aim to annotate detailed natural language descriptions based on these prompts. However, manual annotation of text descriptions is time-consuming, and annotator expertise significantly influences the results. In this work, we propose leveraging large-scale language models (LLMs) [21] for annotation. We carefully craft an input containing our prompts and let the power of LLM provide descriptions; with well-selected prompts, the resulting descriptions are diverse and fine-grained.

We also propose a new cross-modal latent diffusion model for our task. Directly modeling hand-object contacts from text descriptions is non-trivial, since natural

language descriptions do not capture the precise physical 3D interaction between hand and object. As such, we design a two-stage approach. Initially, we generate a hand pose from a noisy hand input and text description, where the hand pose aligns with the text description but may not establish optimal contact with the object. We then further model the hand-object contact from the text description and the object’s point clouds, with the hand pose serving as a prompt. To our knowledge, we are the first to model 3D hand-object contacts using natural language descriptions.

In summary, our contributions are four-fold:

- We explore a new task of modeling 3D hand-object contact from natural descriptive languages.
- We create a new dataset, *ContactDescribe*, consisting of contact maps and corresponding descriptive languages generated by LLMs.
- We design a new staged diffusion model to generate 3D hand-object contacts from the descriptive language.
- We propose a text-to-hand-object fusion network that efficiently extracts contact information from point clouds and natural language to guide the reverse diffusion process.

We evaluate the efficacy of our approach on two datasets: the newly introduced *ContactDescribe* dataset and the HO3D dataset [13]. Additionally, we validate the capability of our methods to effectively capture controllable 3D hand-object contact through our natural language annotations. Our dataset and proposed method provide a new research direction for the community to model interactions with natural language.

2 Related Work

Hand-object contact modeling. Modeling contact between hands and objects is important for grasp reconstruction and synthesis [14, 26, 36, 40, 41]. Existing studies [10, 18, 24, 42, 47, 49, 51, 53, 54] inferred contact patterns from hand-object point clouds with geometry constraints such as penetrations. ContactOpt [10], GraspTTA [18], and Contact2Grasp [24] generated the contact regions based on PointNet [30]. ContactGen [27] further generated contact location with the hand part and touch direction. Additional constraints from RGB images and sequences also benefit the contact modeling. S²Contact [42] proposed estimated the contact with visual and geometry constraints. TOCH [53] introduces the temporal constraint to model contact sequences. However, research focusing on text-to-contact modeling remains relatively limited.

Text-guided diffusion models. Diffusion models [16, 25] are probabilistic generative models that use the Markov chain to convert a simple Gaussian distribution to a target data distribution by gradually adding and removing noise. Recently, Rombach *et al.* [33] proposed the latent diffusion model (LDM) and has shown the real powers of LDM in text-to-image synthesis [45]. Inspired by LDM, Ye *et al.* [50] also adopt a text-conditioned diffusion model for guiding the 3D

generation of hand-object interactions. Furthermore, text is often used in human motion generations [2, 11, 22]. Karunratanakul *et al.* [19] proposed the Guided Motion Diffusion, which can generate controllable and diverse human motions given a text prompt. Text inputs also show potential improvements in robotic grasping [7, 38]. Ha *et al.* [12] presented a framework for robotic skill learning, which uses natural language (*e.g.* put the toy into the left bin) to guide high-level grasp planning. This work uses the diffusion model to explore the controllable 3D hand-object contact generation from text description.

3 *ContactDescribe* Dataset

Dataset collection. Learning a controllable 3D hand-object contact modeling from natural language requires 3D contact data with corresponding text descriptions. No dataset to date includes detailed text for describing hand-object contact patterns (see comparison in Table 1). To bridge this gap, we created the *ContactDescribe* dataset, which contains paired point cloud-text data. The hand-object point clouds and essential annotations (such as MANO [34] parameters, object meshes and 6D poses, and hand-object contact maps.) are taken from ContactPose [4]. *ContactDescribe* extends the ContactPose dataset by adding annotations of contact pattern descriptions for each hand-object point cloud. To enhance the naturalness and diversity of these descriptions, we leverage Large Language Models (LLMs), such as ChatGPT [1], to generate the final sentences given hand-centered contact prompts, as depicted in Figure 2.

Multi-level language description. As illustrated in Figure 2, we newly designed a coarse-to-fine multi-level text description to describe the hand-object contact systematically. The high-level text describes the grasp action, the middle-level text details the grasp type and finger status, while the low-level text specifies the contact location on which hand joints.

Prompt annotation. We first manually annotate each hand-object point cloud to obtain the contact prompts. We aim to depict contact from a hand-centric perspective and prioritize descriptions that enhance adaptability across diverse object categories and attributes.

As illustrated in Figure 2, our prompt annotations of the contact parts contain several single-choice questions covering grasp action, basic grasp type, contact location on each finger, and free finger status. The basic grasp types assume the involvement of all fingers in contact, including ‘*wrap*’, ‘*pinch*’, ‘*lift*’, and ‘*twist*’. However, grasp type alone cannot precisely describe contact patterns and lacks human-intuitive understanding. Moreover, the diverse vocabulary used to describe grasps presents challenges in achieving uniform descriptions. Consequently, we ask annotators to specify the contacts of each finger to refine the depiction. Three descriptive options are provided for fingers with free contact: ‘*sphere*’, ‘*full extension*’, and ‘*partial extension*’. To make our language more natural and be used by untrained users in daily chat, we avoid using highly specialized vocabulary to describe hand joints, such as the Distal Interphalangeal Joint (DIP Joint).

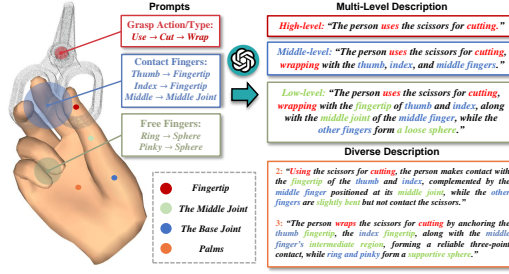


Fig. 2: An example of the dataset with multi-level language descriptions. Given the proposed hand-centered contact prompts, our dataset leverages ChatGPT3 [1] to generate multi-level and diverse free-form text descriptions of contact patterns.

Datasets	Label		Grasp Finger Contact		
	Verb	Sentence	type	status	location
HO3D [13]	✗	✗	✗	✗	✗
YCBAfford [8]	✗	✗	✓	✗	✗
AffordPose [17]	✗	✗	✓	✗	✗
HOI4D [28]	✓	✗	✓	✗	✗
ContactPose [4]	✓	✗	✓	✗	✗
OakInk [48]	✓	✗	✓	✗	✗
Ours	✓	✓	✓	✓	✓

Table 1: Intent labels comparison between different datasets of hand-object interaction. Most existing datasets collect verb-level intent labels. They are insufficient to prompt controllable contact modeling. Our dataset is the first to provide detailed sentence-level contact descriptions.

Sentence generation. To increase description diversity, we leverage the ChatGPT3 model [1] for sentence generation and provide the following prompts:

1. I will give you some prompts describing a hand gesture when humans use the scissors; please help me reword a sentence. But remember not to describe the shape or attribute of the object.
2. Prompts are [grasp action], [grasp type], [contact location], [free fingers status].
3. To tell ChatGPT some basic knowledge about the hand, e.g. using the Proximal/DIP joint, what the sphere finger is, etc.
4. Help me reword four sentences to a different format but keep their meaning. You can incorporate the descriptions I just provided.

We further leverage ChatGPT to check the sentence and give feedback to annotators automatically, and annotators rephrase inappropriate sentences manually.

Detailed Statistics. *ContactDescribe* contains 2,300 unique grasps of 25 household objects (e.g. scissors, bottles, eyeglasses, etc.) from 50 participants. Our dataset consists of 11,500 language descriptions, with each grasp corresponding to five different sentences.

4 Methodology

4.1 Preliminary: Contact Modeling

The contact map on hand point cloud $C_H \in \mathbb{R}^{N \times 1}$ and object point cloud $C_O \in \mathbb{R}^{N \times 1}$, where N is the number of points, denotes the contact probability of each point, with values constrained within the $[0, 1]$ range. As shown in Figure 1 and 3, the redder the color, the higher the probability of the point being in contact. The contact modeling network is applied to the point cloud and outputs its contact map. Our goal is to model the contact map from the input text description. The contact map can be easily applied for the following two tasks:

Grasp pose optimization leverages contact map to refine inaccurate 3D hand pose, e.g. the estimation from a single RGB image. It first predicts the contact map

on both hand and object mesh. To optimize the hand pose, the hand MANO [34] parameters are iteratively updated to minimize the difference between the current contact maps and the predicted contact maps.

Human grasp generation generates hand parameters as human grasps from a given object mesh based on the contact map. It first generates hand MANO mesh and then predicts the contact map between the object and the generated hand. Finally, the hand-object meshes are fitted using the contact map.

4.2 Method Overview

Given inputs $\{\tilde{\mathcal{H}}, \mathcal{O}, \mathbb{T}\}$, where $\mathcal{O} \in \mathbb{R}^{2048 \times 3}$ and \mathbb{T} are the object point cloud and the contact description respectively. The initial hand MANO pose $\tilde{\mathcal{H}} = (\theta, \beta, t^H, R^H)$ consisting of pose, shape, translation, and rotation respectively is an optional input. In this work, we consider initial hand poses derived from an image-based algorithm [14], though our method is not limited specifically and can accept alternatively derived hand poses. We aim to generate a 3D contact map \hat{C}_H and \hat{C}_O on both the hand and object mesh. To achieve this, a three-module generative network, NL2Contact, is proposed as shown in Figure 3. The text-to-hand-object fusion module (Section 4.3) efficiently extracts contact information from both point clouds and natural language to guide the reverse diffusion process. A staged diffusion module (Section 4.4) generates hand MANO pose and subsequently predicts 3D hand-object contact maps. The generated hand pose is then refined through optimization using the generated contact map (Section 4.5). We elaborate on these submodules in the following sections.

4.3 Text to Hand-Object Fusion

Encoders. Consider a set of inputs $\{\tilde{\mathcal{H}}, \mathcal{O}, \mathbb{T}\}$, the NL2Contact extracts global feature $f_\theta^g \in \mathbb{R}^{64}$ and tokens embedding $f_{text} \in \mathbb{R}^{768 \times n}$ from the hand pose $\tilde{\mathcal{H}}$ and text \mathbb{T} by two feature encoders VPoser [29] and BERT [9], and extracts global and local features $f_{obj}^g \in \mathbb{R}^{64 \times N}$ and $f_{obj}^l \in \mathbb{R}^{1024}$ from object point clouds \mathcal{O} using the PointNet++ [31], where N is the number of points and n is the number of words in text.

Text-to-hand fusion. We apply MLP layers to the tokens embedding f_{text} to get a global text feature $f_{text}^g \in \mathbb{R}^{64}$. Then we fuse f_{text}^g with the hand pose feature f_θ^g and object global feature f_{obj}^g via two multi-head attention modules [6] in a cascaded manner to get the latent condition for the hand pose denoising process. The first multi-head attention module takes the text feature f_{text}^g as a query to extract the object shape information from the object global feature f_{obj}^g . The second attention takes the hand pose feature f_θ^g as a query to extract the semantic hand gesture information from the text feature f_{text}^g . Such design allows a switch between grasp pose optimization and human grasp generation.

Text-to-object fusion. This module aims to fuse text information into hand-object point clouds to get the latent condition for guiding the hand-object contact denoising process. In particular, we first transfer the generated hand pose \mathcal{H}' ,

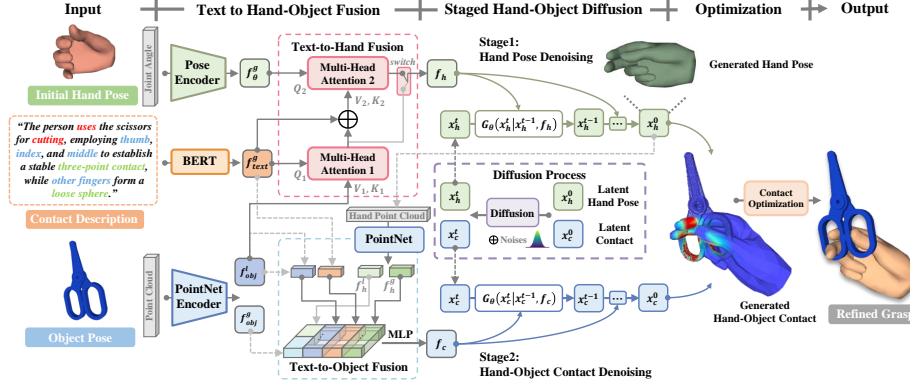


Fig. 3: NL2Contact pipeline. We propose a novel method to model 3D hand-object contact using language description. Our framework is composed of 1) a Text-to-Hand-Object Fusion module, efficiently fusing the geometry information from both point clouds and semantic information from natural language, 2) a staged latent diffusion model to first generate the hand pose fitting the hand description and then generate the 3D hand-object contact map conditioned on the descriptions and the generated hand pose. 3) contact optimization, which iteratively optimizes the hand pose using the generated contact. We also added a switch to toggle between its application for grasp pose optimization and human grasp generation.

which outputs from the first diffusion stage to the hand point by a differentiable MANO Layer [34]. Then, we extract global and local features $f_h^g \in \mathbb{R}^{64 \times 778}$ and $f_h^l \in \mathbb{R}^{1024}$ from hand point clouds using the PointNet++ [31]. To obtain the per-point hand-object feature, we combine the hand local feature f_h^l and object local feature f_{obj}^l together, including an additional binary per-point feature indicating whether the point belongs to the hand or the object. The global features of hand f_h^g , object f_{obj}^g and text f_{text}^g are then duplicated N times, N is $(2048+778)$, for matching the shape of hand-object per-point feature. All these features are then concatenated and fused by MLP layers to produce the conditional embeddings f_c for the contact generator below.

4.4 Staged Hand-Object Diffusion Models.

Conditional latent diffusion models. Latent diffusion models (LDMs) [33] learns a data distribution $q(x^0)$ by reversing a progressive noise diffusion process from a prior Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$. Given a sample $x^0 \sim q(x^0)$ from the latent space, a fixed diffusion process $q(x^{1:t} | x^0) = \prod_{t=1}^t q(x^t | x^{t-1})$ is used to add Gaussian noise $q(x^t | x^{t-1}) := \mathcal{N}(\sqrt{1 - \beta_t}x^{t-1}, \beta_t \mathbf{I})$, producing a chain of noisy latent codes $\{x^1, x^2, \dots, x^t\}$, where β_t is the variance defined in the forward process. The latent code x^0 is extracted using a pre-trained Encoder [43] from the raw data, which reduces the computational requirements by learning this distribution in a latent space. Then, the denoiser $G_\theta(x^{t-1} | x^t) =$

$\mathcal{N}(\mu_\theta(\mathbf{x}^t, t), \sigma_t^2 \mathbf{I})^3$, typically a U-Net [35], is trained to estimate the noise from those samples. The training objective can be simplified to:

$$\begin{aligned} \mathcal{L}_{LDM} &= \mathbb{E}_{\mathbf{x}^0, \varepsilon \sim \mathcal{N}(0,1), t} \left[\|\varepsilon - G_\theta(\mathbf{x}^t, t, f)\|_2^2 \right], \\ G_\theta(\mathbf{x}^{0:t}) &= G(\mathbf{x}^t) \prod_{t=1}^t G_\theta(\mathbf{x}^{t-1} | \mathbf{x}^t, f), \end{aligned} \quad (1)$$

where f is the latent codes of conditions, such as text and images.

Hand pose diffusion. Figure 3 shows that we use LDMs conditioned on Text-to-Hand embeddings f_h to model the hand pose generation process. Following the LDMs, we apply encoder E_p to encode the hand pose ground truth \mathcal{H}_{gt} into a latent code and reshape it to dimensions $x_h^0 \in \mathbb{R}^{d \times d}$. As E_p , we use VPoser [29], a variational autoencoder for human pose representations. Then, a forward diffusion process adds Gaussian noise to the input \mathbf{x}^0 to transfer it to \mathbf{x}^t . This latent code is subsequently fed into a denoiser G_θ^1 , U-Net [35], along with condition embedding f_h and t to predict the $\hat{\mathbf{x}}_h^0 = G_\theta^1([\mathbf{x}_h^t, f_h, t])$, where $\mathbf{x}^0 = E_p(\mathcal{H}_{gt})$. Finally, the decoder from VPoser generates the hand pose \mathcal{H}' from the reconstructed $\hat{\mathbf{x}}_h^0$. For the denoising process at the timestep t , the training objective minimizes the following:

$$\mathcal{L}_{\text{pose}} = \mathbb{E}_{\mathbf{x}_h^0, \varepsilon \sim \mathcal{N}(0,1), t} \left[\|\varepsilon - G_\theta^1(\mathbf{x}_h^t | t, f_h)\|_2^2 \right]. \quad (2)$$

Hand-object contact diffusion. We freeze the hand pose diffusion’s parameters and pre-train a PointNet [31] as the contact encoder to extract the latent code from the hand-object contact ground-truth \mathcal{C}_{gt} and then reshape it to $x_c^0 \in \mathbb{R}^{32 \times 32}$. We get the condition embedding f_c by sending the generated hand pose into the Text-to-object fusion module. Similar to the hand pose denoising, this embedding is then fed into the U-Net denoiser G_θ^2 along with x_c^t and t to predict the $\hat{\mathbf{x}}_c^0$. The loss function is:

$$\mathcal{L}_{\text{contact}} = \mathbb{E}_{\mathbf{x}_c^0, \varepsilon \sim \mathcal{N}(0,1), t} \left[\|\varepsilon - G_\theta^2(\mathbf{x}_c^t | t, f_c)\|_2^2 \right]. \quad (3)$$

The PointNet contact decoder produces the contact map $\hat{\mathcal{C}}$ from the predicted $\hat{\mathbf{x}}_h^0$. At test time, a latent vector is randomly sampled from the Gaussian distribution and progressively denoised by U-Net to generate contact.

4.5 Contact Optimizations.

Given the generated hand pose \mathcal{H}' , our target is to optimize it to find the best parameter θ and β that best complies with the generated hand-object contact map $\hat{\mathcal{C}}$. To achieve this, we iteratively minimize the difference between the current contact \mathcal{C}' and the generated contact $\hat{\mathcal{C}}$ and update the hand pose parameters. We applied the contact map on both hand and object point clouds, where values are

³ We slightly abuse the notion θ to represent learning of model parameters.

bounded within the $[0, 1]$ range to indicate the contact status of the points. The current hand-object contact map \mathcal{C}' at each step is computed by DiffContact [10]. The loss for the contact optimization is formulated as follows:

$$\mathcal{L}_{opt} = \left\| \mathcal{C}'_H - \hat{\mathcal{C}}_H \right\|_2^2 + \lambda_o \left\| \mathcal{C}'_O - \hat{\mathcal{C}}_O \right\|_2^2 + \lambda_{pen} \mathcal{L}_{pen}, \quad (4)$$

where \mathcal{L}_{pen} is the penetration loss to penalize the penetration between the hand and the object, similar to ContactOpt [10].

5 Experiments

5.1 Implementation Details

Our model is trained in an end-to-end manner. A PointNet++ [31] is pre-trained on the ContactPose [4] dataset as the contact encoder. We sample 2048 points from the object mesh as the input. We adopt the Adam optimizer [20], a batch size of 64, and a learning rate of $1.5e^{-3}$ with 50 epochs for training hand pose denoising and another 100 epochs for hand-object contact generation. The training takes about 11 hours on a single V100 GPU with 9.5GB memory. Regularization terms, λ_O and λ_{pen} , are assigned values of 5 and 3, respectively. During the contact optimization, we apply the random restart to avoid a local optimum following [10]. NL2Contact takes around 3 seconds to generate per contact map.

5.2 Datasets

ContactDescribe dataset. The *ContactDescribe* dataset is an enhanced dataset based on the ContactPose. ContactPose is the first dataset [4] that combines hand-object contact information with hand pose, object pose, and RGB-D images. It encompasses 2,306 distinct grasps involving 25 household objects, executed with two functional intents by 50 participants, and includes over 2.9 million RGB-D grasp images. To ensure fair comparisons with the ContactOpt [10] and S²Contact [42], we adhere to the Perturbed ContactPose dataset, wherein hand meshes undergo modification through the introduction of additional noise to MANO parameters. This dataset comprises 22,624 training grasps and 1,416 testing grasps. We annotated five contact descriptions for each grasp instance. Finally, we train our NL2Contact on the *ContactDescribe* dataset.

Described HO3D dataset. HO3D [13] datasets are used for evaluating 3D hand-object interactions, which are captured in a real-world setting and involve 10 different subjects performing various fine-grained manipulations on one of the 10 objects from the YCB models [5]. We employ the HO3D dataset to assess the generalization capability of our proposed method when faced with previously unseen objects. We choose a subset from HO3D following ContactOpt [10] (around 14k grasps) and manually annotate the textual descriptions.

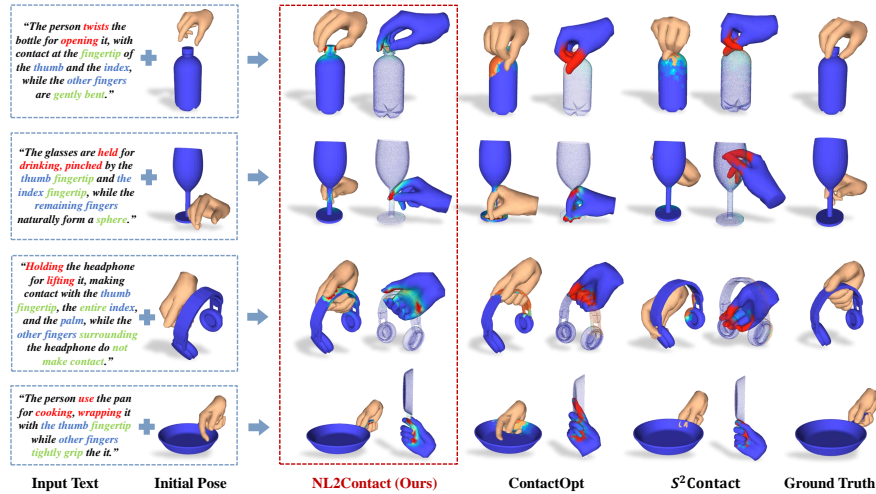


Fig. 4: We conduct qualitative comparisons with state-of-the-art methods. Our method generates accurate and controllable contact, matching the input text description. ContactOpt [10] and S^2 Contact [42] are uncontrollable contact generation methods. They tend to generate a grasp with all five fingers engaged. Interestingly, we provide a great initial pose in the first row. However, ContactOpt and S^2 Contact still generate a five-finger grasp which is different from the initial pose.

5.3 Evaluation Metrics

Hand error. We use the Mean Per-Joint Position Error (MPJPE) (mm) over 21 joints to evaluate the grasp refinement. We calculate the average L2 per-joint kinematic error with the ground truth hand pose.

Contact quality. We adopt the intersection volume (cm^3), contact coverage (%) and Contact Precision/Recall (%). The intersection volume (denoted as Inter. in tables) is acquired by voxelizing both the hand and object with a voxel size of 0.5cm. Contact coverage (denoted as Cover. in tables) calculates the proportion of grasps in contact with objects, defined as the percentage of hand points within $\pm 2\text{mm}$ of the object’s surface. Contact precision/recall (denoted as Pr. and Re. in tables) calculates the similarity between the generated contact and ground-truth contact with a threshold of 0.4.

Diversity of grasp generation. We cluster generated grasps into 20 clusters using K-means and then evaluated the average Cluster Size (denoted as Diversity).

Physical stability. We placed the object and its corresponding grasp in a simulator. We evaluated the simulation displacement (SD) by calculating the displacement (cm) of the object’s center relative to its center of mass over a specified duration under the influence of gravity. We randomly generated 20 grasps and reported the average SD.

Table 2: Quantitative results of grasp pose optimization on *ContactDescribe* and Described HO3D dataset. Due to the natural language guidance, we can model accurate hand-object contact. The optimization from our predicted contact achieves the lowest MPJPE. We also show that contact descriptions improve the generalization of our method to out-of-domain objects on HO3D dataset.

Method	<i>ContactDescribe</i> Dataset					Described HO3D Dataset		
	MPJPE↓	Contact Quality				Method	MPJPE (mm)↓	Inter. (mm ³)↓
		Inter. ↓	Cover. ↑	Pr. ↑	Re. ↑			
Perturbed Pose	79.9	8.4	2.3%	9.9%	11.5%	Initial Pose [14]	11.4	9.26
ContactNet [18]	45.2	15.6	18.4%	31.6%	47.6%	ContactOpt [10]	9.5	5.71
ContactOpt [10]	25.1	12.8	19.7%	38.7%	54.8%	TOCH [53]	9.3	4.67
S ² Contact [42]	29.4	12.2	22.2%	42.5%	56.1%	S ² Contact [42]	8.7	3.52
NL2Contact	21.7	7.1	30.5%	49.2%	59.9%	NL2Contact	8.4	4.39

5.4 Evaluating Grasp Pose Optimization

We first evaluate the effectiveness of the proposed NL2Contact application on grasp pose optimization from the hand-object point cloud.

Baseline methods. We compare our method with three strong geometry-guided grasp optimization baselines. 1) ContactOpt [10] estimates contact maps on both hand and object with PointNet-based DeepContact. 2) S²Contact [42] proposed a graph-based GCN-Contact network to predict hand-object contacts and a semi-supervised method to improve generalization on out-of-domain datasets. 3) Jiang *et al.* [18] also provides a ContactNet to estimate the contact map from the hand-object point cloud, but it only outputs the object contact.

Results on *ContactDescribe*. We evaluate the NL2Contact on the ContactDescribe dataset, in which we perturb the annotated MANO parameters. The hand error after perturbation is around 80 mm. This tests the ability to improve hand poses using the predicted contact maps. Table 2 shows a quantitative comparison between different methods. Our method significantly outperforms the other methods by a large margin. We achieve the lowest MPJPE score, outperforming ContactOpt [10] by 4.4 mm. Besides, our method significantly reduced the intersection volumes while other methods increased it. As shown in Figure 4, the ContactOpt and S²Contact always have penetration cases. For the contact quality, the results illustrate that our method can generate more similar contact with the ground-truth contact (gathering by thermal camera). We can observe that the NL2Contact is the most effective one for hand pose refinement. Other methods only boost the contact quality, and the hand pose still exhibits large errors after contact optimization. However, our method leverages the language description to provide semantic information for generating the contact map and effectively reduces hand error while improving the contact quality.

Results on Described HO3D. We demonstrate the generalization ability of our method on the HO3D dataset. Since the hand model (MANO) is consistent across different datasets, most contact modeling methods can directly infer the contact without being retrained. Nevertheless, the domain gap still exists, occasionally resulting in insufficient contact or significant penetration. As indicated in Table 2, our method’s performance is close to the best method, S²Contact [42], in the

Table 3: Quantitative results of grasp generation. Our method achieves the lowest penetration, best contact, and highest diversity compared to ContactGen [27].

Method	Inter.↓	Conver.↑	Diversity↑	SD↓
GrabNet [37]	15.50	99%	2.06	2.34
GraspTTA [18]	7.37	76%	1.43	5.34
ContactGen [27]	9.96	97%	5.04	2.70
Ours	5.89	99%	5.91	2.31




Fig. 5: Visualization of diverse grasp generation. We observe ContactGen [27] always generates grasps with all five fingers engaged, while our method generates grasps consistent with the text.

metric of Intersection Volume. The S^2 Contact used a complex semi-supervised pipeline and carefully selected pseudo-label to model contact for an unseen object. This illustrates that our designed hand-centered contact description is generic and adapts to the large domain among different datasets. Besides, compared with S^2 Contact, our method achieves notably lower hand error.

Qualitative results. We present our main experimental results in Figure 4. We observe that our proposed method can generate accurate contacts that exactly match the input concrete descriptions, and then the initial hand pose can be improved by contact optimization. Other methods, e.g. ContactOpt [10] and S^2 Contact [42] always generate full-five finger grasps. Besides, their method is not controllable; results after refinement may exhibit larger hand pose errors, while our method, guided by textual prompts, ensures that the optimized results align with the target hand pose.

5.5 Evaluating Human Grasp Generation

Our method can be extended to generate human grasps from a given input object and contact description. We train the model on the *ContactDescribe* dataset and test on unseen objects from the HO3D dataset.

Baseline methods. We compare our method with the state-of-the-art human grasp generation approach, ContactGen [27], which generates contact for the input object and then adopts model-based optimization to predict grasps.

Results. We first quantitatively evaluate the grasp generation for the unseen object from the HO3D dataset. GraspTTA [18] are trained on the ObMan dataset [15], and GrabNet [37] and ContactGen [27] are trained on the GRAB dataset [37]. We randomly generate 20 grasps. As indicated in Table 3, our method achieves the best performance in all metrics. GraspTTA and GrabNet are struggling with poor diversity. Comparison results of grasp generation for binoculars are shown in Figure 5. ContactGen always generates a grasp with all five fingers touched, while our method can control the generation using the text.

Table 4: Ablation study of multi-level text input. We demonstrate that the more low-level, the more precise the grasp and location of the contact points.

Method	MPJPE↓	Contact Quality		
		Inter. ↓	Pr. ↑	Re. ↑
ContactOpt [10]	25.1	12.8	38.7%	54.8%
High-level	23.9	8.9	42.7%	55.6%
Middle-level	22.5	8.0	46.8%	57.3%
Low-level	21.7	7.1	49.2%	59.9%

Table 5: Ablation study of our method. We demonstrate that the full model of NL2Contact outperforms all other variants.

Method	MPJPE↓	Contact Quality		
		Inter. ↓	Pr. ↑	Re. ↑
w/o text-guided	24.3	11.7	39.9%	56.2%
w/o text2hand fusion	23.5	10.3	41.6%	57.9%
w/o text2ho fusion	22.9	9.7	43.9%	58.4%
w/o staged diffusion	27.3	13.6	38.5%	52.1%
with VAE	21.9	7.4	47.0%	56.3%
full model	21.7	7.1	49.2%	59.9%

5.6 Ablation Study

Effect of multi-level text input. To validate the effectiveness of the different level text descriptions, we conducted the experiments by varying the specificity of the text and showing the grasp optimization results in Table 4. The initial grasp pose is also from Hasson *et al.* [14]. We find that the more *low-level* text, the more precise the grasp and location of the contact points. Notably, the high-level text is removed from all contact-related descriptions and still significantly outperforms ContactOpt.

Effect of language description. To validate the effectiveness of the language description-guided contact modeling, we conduct the experiments with the following settings: (a) Without text input: the contact description is not used. At the Text-to-Hand-Object fusion stage, we use the multi-head attention 2 module to fuse the object global feature with the hand pose feature to produce conditional embedding. At the text-ho stage, we do not concatenate the text global feature. (b) Without Text-to-Hand fusion module: the text feature does not fuse to the hand pose feature using cascaded multi-head attention, instead by directly concatenating. (c) Without the Text-to-HO fusion module, the text feature does not fuse with the hand-object point cloud feature while fusing the global hand-object feature into the local hand-object feature. The results of the ablation study are shown in Table 5. Comparing (a) with the full model, we find the hand error significantly increases while the contact quality decreases. This result verifies that introducing contact description supervision improves the effectiveness of our model. However, our results are still better than those of ContactOpt; we consider the reason that our diffusion model is more effective than PointNet for modeling contact. Comparing (b) and (c), we find that Text-to-Hand fusion is more effective because this can correct the initial input to fit the gesture in the description and reduce the complexity of text-to-contact generation.

Effect of the Staged Diffusion model. To validate the effectiveness of the staged diffusion, we conduct the experiments with the following settings: (d) Without the staged diffusion model: we do not generate hand-object contact with two stages. Instead, we use a single diffusion model to generate the contact map directly. To this end, we use PointNet to extract the hand feature from the initial hand pose and the Text-to-HO fusion module to fuse all features. (e) With VAE:

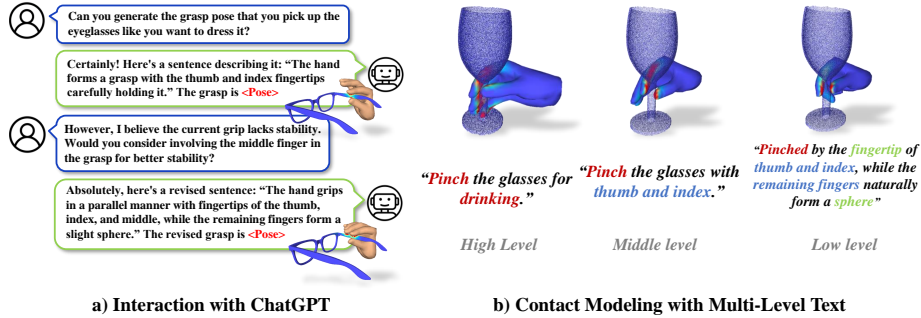


Fig. 6: a) We show the seamless integration of our method with ChatGPT, allowing for interactive control to generate an initial grasp and refine it. Users offer high-level prompts, and ChatGPT generates detailed contact descriptions, then feeds descriptions into our network to generate the grasp. b) Our model generates more accurate grasps as text descriptions become more detailed.

we implement our pipeline using two VAE generators instead of staged diffusion. As illustrated in Table 5, the performance declines significantly under setting (d) without generated hand pose. This result demonstrates that directly cross-modal mapping text to hand-object contact is non-trivial. Setting (e) also shows that other generative models, such as VAEs, apply to our task. However, the diffusion model performs better due to its iterative denoising mechanism, which is more effective in learning a data distribution.

Applications with ChatGPT. Our method can provide more potential downstream applications, such as the teleoperation of robotic hands and virtual avatar editing. In this experiment, we integrate NL2Contact with ChatGPT [1] to facilitate interactive control. As illustrated in the left part of Figure 6, the user initially generates a grasp of the scissors using a high-level description. The large language model comprehends how to use the scissors and provides a low-level contact description. Subsequently, our method generates the initial grasp by low-level description. If the generated contact is unsatisfactory, our method can optimize the grasp pose based on different levels of user feedback, as shown in the right part of Figure 6.

6 Conclusion

This paper proposes a novel task - natural language-guided 3D hand-object contact modeling. We introduce NL2Contact, a novel cross-modal latent diffusion model for contact modeling, and *ContactDescribe* dataset, which provides 3D hand-object contact and fine-grained natural language descriptions. Our method can be applied to grasp pose optimization and novel human grasp generation. In the future, we would like to explore the modeling of dynamic contacts with text.

Acknowledgements

This research was supported by the MSIT(Ministry of Science and ICT), Korea, under the ITRC(Information Technology Research Center) support program(IITP-2024-2020-0-01789) supervised by the IITP(Institute for Information & Communications Technology Planning & Evaluation), the National Research Foundation Singapore and DSO National Laboratories under the AI Singapore Programme (Award Number: AISG2-RP-2020-016), China Scholarship Council (CSC) Grant No. 202208060266 and No.202006210057.

References

1. <https://openai.com/blog/chatgpt/>
2. Ahn, H., Ha, T., Choi, Y., Yoo, H., Oh, S.: Text2action: Generative adversarial synthesis from language to action. In: ICRA (2018)
3. Brahmbhatt, S., Ham, C., Kemp, C.C., Hays, J.: ContactDB: Analyzing and predicting grasp contact via thermal imaging. In: CVPR (2019)
4. Brahmbhatt, S., Tang, C., Twigg, C.D., Kemp, C.C., Hays, J.: ContactPose: A dataset of grasps with object contact and hand pose. In: ECCV (2020)
5. Calli, B., Singh, A., Walsman, A., Srinivasa, S., Abbeel, P., Dollar, A.M.: The ycb object and model set: Towards common benchmarks for manipulation research. In: ICAR (2015)
6. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: ECCV (2020)
7. Cheang, C., Lin, H., Fu, Y., Xue, X.: Learning 6-dof object poses to grasp category-level objects by language instructions. In: ICRA (2022)
8. Corona, E., Pumarola, A., Alenya, G., Moreno-Noguer, F., Rogez, G.: Ganhand: Predicting human grasp affordances in multi-object scenes. In: CVPR (2020)
9. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805 (2018)
10. Grady, P., Tang, C., Twigg, C.D., Vo, M., Brahmbhatt, S., Kemp, C.C.: ContactOpt: Optimizing contact to improve grasps. In: CVPR (2021)
11. Guo, C., Zou, S., Zuo, X., Wang, S., Ji, W., Li, X., Cheng, L.: Generating diverse and natural 3d human motions from text. In: CVPR (2022)
12. Ha, H., Florence, P., Song, S.: Scaling up and distilling down: Language-guided robot skill acquisition. CoRL (2023)
13. Hampali, S., Rad, M., Oberweger, M., Lepetit, V.: Honnotate: A method for 3D annotation of hand and object poses. In: CVPR (2020)
14. Hasson, Y., Varol, G., Laptev, I., Schmid, C.: Towards unconstrained joint hand-object reconstruction from RGB videos. In: 3DV (2021)
15. Hasson, Y., Varol, G., Tzionas, D., Kalevtykh, I., Black, M.J., Laptev, I., Schmid, C.: Learning joint reconstruction of hands and manipulated objects. In: CVPR (2019)
16. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. NeurIPS (2020)
17. Jian, J., Liu, X., Li, M., Hu, R., Liu, J.: Affordpose: A large-scale dataset of hand-object interactions with affordance-driven hand pose. In: ICCV (2023)

18. Jiang, H., Liu, S., Wang, J., Wang, X.: Hand-object contact consistency reasoning for human grasps generation. In: ICCV (2021)
19. Karunratanakul, K., Preechakul, K., Suwajanakorn, S., Tang, S.: Guided motion diffusion for controllable human motion synthesis. In: ICCV (2023)
20. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
21. Kojima, T., Gu, S.S., Reid, M., Matsuo, Y., Iwasawa, Y.: Large language models are zero-shot reasoners. *Advances in neural information processing systems* **35**, 22199–22213 (2022)
22. Kong, H., Gong, K., Lian, D., Mi, M.B., Wang, X.: Priority-centric human motion generation in discrete latent space. In: ICCV (2023)
23. Lakshmipathy, A.S., Feng, N., Lee, Y.X., Mahler, M., Pollard, N.: Contact edit: Artist tools for intuitive modeling of hand-object interactions. *ACM Transactions on Graphics (TOG)* (2023)
24. Li, H., Lin, X., Zhou, Y., Li, X., Huo, Y., Chen, J., Ye, Q.: Contact2grasp: 3d grasp synthesis via hand-object contact constraint. *IJCAI* (2022)
25. Liu, N., Li, S., Du, Y., Torralba, A., Tenenbaum, J.B.: Compositional visual generation with composable diffusion models. In: ECCV (2022)
26. Liu, S., Jiang, H., Xu, J., Liu, S., Wang, X.: Semi-supervised 3D hand-object poses estimation with interactions in time. In: CVPR (2021)
27. Liu, S., Zhou, Y., Yang, J., Gupta, S., Wang, S.: Contactgen: Generative contact modeling for grasp generation. In: CVPR (2023)
28. Liu, Y., Liu, Y., Jiang, C., Lyu, K., Wan, W., Shen, H., Liang, B., Fu, Z., Wang, H., Yi, L.: Hoi4d: A 4d egocentric dataset for category-level human-object interaction. In: CVPR (2022)
29. Pavlakos, G., Choutas, V., Ghorbani, N., Bolkart, T., Osman, A.A., Tzionas, D., Black, M.J.: Expressive body capture: 3d hands, face, and body from a single image. In: CVPR (2019)
30. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: PointNet: Deep learning on point sets for 3D classification and segmentation. In: CVPR (2017)
31. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: PointNet++: Deep hierarchical feature learning on point sets in a metric space. In: NeurIPS (2017)
32. Qin, Y., Wu, Y.H., Liu, S., Jiang, H., Yang, R., Fu, Y., Wang, X.: Dexmv: Imitation learning for dexterous manipulation from human videos. In: ECCV. pp. 570–587 (2022)
33. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: CVPR (2022)
34. Romero, J., Tzionas, D., Black, M.J.: Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics (ToG)* (2017)
35. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: MICCAI (2015)
36. Sener, F., Chatterjee, D., Shelepov, D., He, K., Singhania, D., Wang, R., Yao, A.: Assembly101: A large-scale multi-view video dataset for understanding procedural activities. In: CVPR (2022)
37. Taheri, O., Ghorbani, N., Black, M.J., Tzionas, D.: GRAB: A dataset of whole-body human grasping of objects. In: ECCV (2020)
38. Tang, C., Huang, D., Ge, W., Liu, W., Zhang, H.: GraspGPT: Leveraging semantic knowledge from a large language model for task-oriented grasping. *IEEE Robotics and Automation Letters* (2023)
39. Tendulkar, P., Surís, D., Vondrick, C.: Flex: Full-body grasping without full-body grasps. In: CVPR (2023)

40. Tse, T.H.E., Kim, K.I., Leonardis, A., Chang, H.J.: Collaborative learning for hand and object reconstruction with attention-guided graph convolution. In: CVPR (2022)
41. Tse, T.H.E., Mueller, F., Shen, Z., Tang, D., Beeler, T., Dou, M., Zhang, Y., Petrovic, S., Chang, H.J., Taylor, J., et al.: Spectral graphormer: Spectral graph-based transformer for egocentric two-hand reconstruction using multi-view color images. In: ICCV (2023)
42. Tse, T.H.E., Zhang, Z., Kim, K.I., Leonardis, A., Zheng, F., Chang, H.J.: S2Contact: Graph-based network for 3d hand-object contact estimation with semi-supervised learning. In: ECCV (2022)
43. Van Den Oord, A., Vinyals, O., et al.: Neural discrete representation learning. NeurIPS (2017)
44. Wang, H., Zhang, Z., Cheng, Y., Chang, H.J.: High-fidelity eye animatable neural radiance fields for human face. BMVC (2023)
45. Wang, H., Zhang, Z., Cheng, Y., Chang, H.J.: Textgaze: Gaze-controllable face generation with natural language. MM (2024)
46. Wu, Y., Wang, J., Zhang, Y., Zhang, S., Hilliges, O., Yu, F., Tang, S.: Saga: Stochastic whole-body grasping with contact. In: ECCV (2022)
47. Xie, W., Zhao, Z., Li, S., Zuo, B., Wang, Y.: Nonrigid object contact estimation with regional unwrapping transformer. In: ICCV (2023)
48. Yang, L., Li, K., Zhan, X., Wu, F., Xu, A., Liu, L., Lu, C.: Oakink: A large-scale knowledge repository for understanding hand-object interaction. In: CVPR (2022)
49. Yang, L., Zhan, X., Li, K., Xu, W., Li, J., Lu, C.: CPF: Learning a contact potential field to model the hand-object interaction. In: ICCV (2021)
50. Ye, Y., Hebbar, P., Gupta, A., Tulsiani, S.: Diffusion-guided reconstruction of everyday hand-object interaction clips. In: ICCV (2023)
51. Yu, Z., Yang, L., Xie, Y., Chen, P., Yao, A.: Uv-based 3d hand-object reconstruction with grasp optimization. BMVC (2022)
52. Zhang, H., Ye, Y., Shiratori, T., Komura, T.: Manipnet: neural manipulation synthesis with a hand-object spatial representation. ACM Transactions on Graphics (ToG) (2021)
53. Zhou, K., Bhatnagar, B.L., Lenssen, J.E., Pons-Moll, G.: TOCH: Spatio-temporal object-to-hand correspondence for motion refinement. In: ECCV (2022)
54. Zhu, Z., Wang, J., Qin, Y., Sun, D., Jampani, V., Wang, X.: Contactart: Learning 3d interaction priors for category-level articulated object and hand poses estimation. 3DV (2024)