

Open-World Amodal Appearance Completion

Jiayang Ao¹ Yanbei Jiang¹ QiuHong Ke² Krista A. Ehinger¹
¹The University of Melbourne ²Monash University




Figure 1. Examples of our open-world amodal completion using both specific (e.g., “polar bear”) and abstract (e.g., “What is the mammal in this image”) text queries. Our approach supports various applications, including image editing, novel view synthesis and 3D reconstruction.

Abstract

Understanding and reconstructing occluded objects is a challenging problem, especially in open-world scenarios where categories and contexts are diverse and unpredictable. Traditional methods, however, are typically restricted to closed sets of object categories, limiting their use in complex, open-world scenes. We introduce Open-World Amodal Appearance Completion, a training-free framework that expands amodal completion capabilities by accepting flexible text queries as input. Our approach generalizes to arbitrary objects specified by both direct terms and abstract queries. We term this capability reasoning amodal completion, where the system reconstructs the full appearance of the queried object based on the provided image and language query. Our framework unifies segmentation, occlusion analysis, and inpainting to handle complex occlusions and generates completed objects as RGBA elements, enabling seamless integration into applications such as 3D reconstruction and image editing. Extensive evaluations demonstrate the effectiveness of our approach in generalizing to novel objects and occlusions, establishing a new benchmark for amodal completion in open-world settings. Code and datasets available: <https://github.com/saraao/amodal>.

1. Introduction

Imagine taking a photograph where part of a landmark is hidden behind trees, or watching a video where an interesting object is partly occluded by other objects in the frame. What if you could automatically reconstruct the hidden parts of these objects to use them in photo editing, augmented reality (AR), or even 3D modelling? This capability, known as amodal completion [12], allows systems to infer and generate the occluded parts of objects, providing users with full representations of partially visible items. Amodal appearance completion is critical in applications such as AR [8, 19, 25], 3D reconstruction [17, 37], and content creation [2, 31, 39], where intuitive, language-based interaction allows users to specify objects directly [3, 9, 22]. Traditional amodal appearance completion methods, however, are typically constrained by fixed sets of object categories [35], or necessitating extensive training and limiting their applicability in diverse and changing environments [23]. The real world presents a more complex challenge: completing occluded objects that can belong to any categories and exhibit diverse, unseen visual features.

To address the gap, we propose **Open-World Amodal Appearance Completion**, a framework that generalizes amodal completion to arbitrary objects specified by flexible natural language queries—a task we term **reasoning**

amodal completion. Our framework supports text-based queries, allowing users to specify target object through both concrete terms and more abstract, context-dependent descriptions, as shown in Fig. 1. This flexibility introduces new possibilities for reasoning-driven amodal completion, where the model can infer and complete occluded objects based on the query alone, without predefined object categories or additional training.

Our approach is designed to be **training-free**, utilizing the extensive knowledge embedded within large, pre-trained models to deliver realistic amodal completion without requiring further data or retraining. By leveraging complementary capabilities across multiple models, our framework integrates segmentation, occlusion analysis, and progressive inpainting into a unified pipeline. Specifically, segmentation isolates the target object based on the text query, occlusion analysis identifies and resolves occluders, and inpainting reconstructs the occluded portions of the specified object with high fidelity.

The framework produces completed objects as RGBA elements, which are immediately ready for integration into downstream applications such as 3D scene reconstruction, AR, and image editing. By enabling intuitive, free-form text input and bypassing the need for retraining, we expand the application scope of amodal completion to complex, real-world scenarios where conventional closed-set methods fail.

Our primary contributions are as follows:

- **Open-World Framework for Amodal Completion:** We propose a training-free framework capable of completing arbitrary objects specified by natural language queries, advancing amodal appearance completion beyond traditional, closed-set approaches.
- **Reasoning Amodal Completion:** We introduce a reasoning-driven method that allows users to specify objects through both specific terms and abstract queries, enhancing amodal completion for intuitive human-computer interaction in real-world scenarios.
- **Adaptable Output for Enhanced Integration:** Our framework outputs completed objects as RGBA elements, facilitating seamless use in AR, novel view synthesis, 3D reconstruction, and image editing.

Through extensive evaluations, including a newly collected evaluation dataset and human preference studies, we demonstrate that our framework effectively generalizes to diverse objects and complex occlusions in real-world scenes. Our approach surpasses existing methods in both quantitative and user-centered metrics, establishing a new benchmark for open-world amodal appearance completion.

2. Related Work

Amodal Appearance Completion and Open-World Adaptability. Amodal completion aims to reconstruct oc-

cluded regions of objects to improve scene understanding [1]. Traditional approaches, however, are often limited to specific, predefined categories and require extensive training data. Category-specific models for humans [40, 43], vehicles [19, 36], food items [24] and indoor scene models [4, 5, 42] can effectively reconstruct occlusions within narrow classes but lack generalization to open-world settings.

Efforts to broaden amodal appearance completion include methods like Pix2gestalt [23], which synthesizes occluded appearances across various objects by training on large synthetic datasets. Other synthetic data-driven methods like CSDNet [42] and PACO [21] have been proposed to handle complex scenes but may face challenges in adapting from synthetic to real-world environments. Self-supervised methods, such as PCNet [38], reduce data dependency but still require visible masks for all objects in the image, making them vulnerable to segmentation errors. PD-MC (Progressive Diffusion with Mixed Context diffusion sampling) [35] partially leverages pre-trained models for broader generalization but remains constrained by its reliance on predefined categories.

Therefore, we need a robust framework that can handle arbitrary objects in the open world without relying on specific categories or extensive data annotations.

Reasoning-Driven and Adaptable Amodal Completion. Recent vision-language models (VLM) allow users to specify objects contextually, enriching human-computer interaction in vision tasks. General-purpose models like CogVLM [33] and VisionLLM [32] perform well in object recognition and language-guided tasks but struggle with reasoning about occlusions. LISA [15] offers improved reasoning for identifying visible objects through language input, yet it falls short when dealing with occluded regions. However, amodal completion requires reasoning beyond the visible to reconstruct these occlusions, providing a full representation of objects that benefits various applications. Increasingly, RGBA output formats are seen as vital for downstream applications, as they support flexible blending and realistic compositing [39]. Datasets like MULAN [31] underscore the importance of RGBA instances in amodal tasks. Altogether, there is a clear need for reasoning-driven amodal completion that adapts dynamically across open-world objects while producing adaptable RGBA outputs.

To address the limitations of existing approaches, our framework introduces a training-free method for open-world amodal appearance completion, accommodating flexible, language-guided queries to identify and complete arbitrary objects. By integrating RGBA outputs, our method ensures adaptable representations that enhance its applicability across diverse downstream tasks.

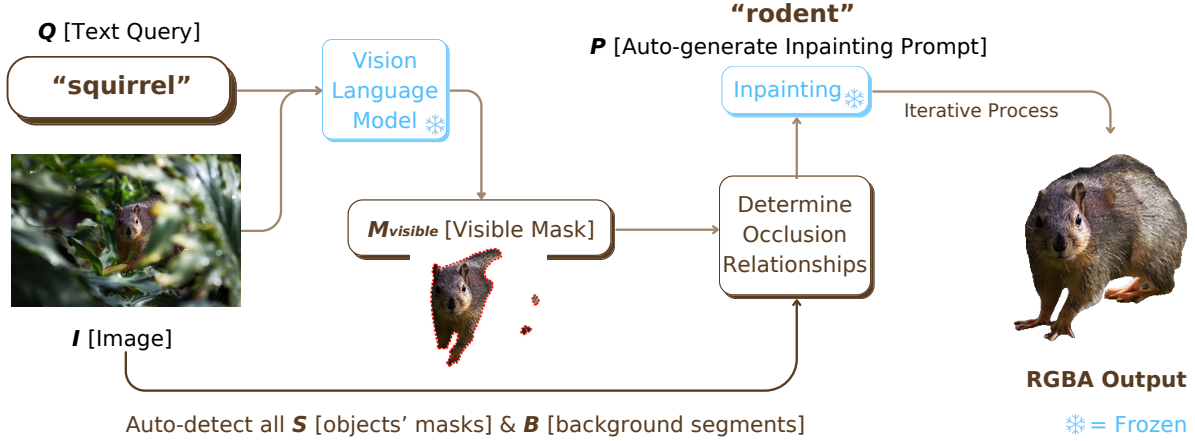


Figure 2. Overview of our framework. Starting with a text query, a VLM generates a visible mask to locate the target object in the input image. The framework then identifies all objects and background segments for occlusion analysis. An auto-generated prompt guides the inpainting model, which iteratively reconstructs the occluded object to produce a transparent RGBA amodal completion output.

3. Methodology

We introduce a novel reasoning-driven approach for reconstructing occluded objects across diverse, unrestricted categories. By leveraging the flexibility of natural language, users can specify target object through both concrete terms and abstract descriptions, allowing amodal completion that generalizes beyond the constraints of closed-set methods. This adaptability is achieved without additional training, allowing it to effectively handle open-world scenes and unknown objects. Fig. 2 presents an overview of our end-to-end inference pipeline, illustrating how flexible text-based input is processed to yield high-quality, transparent RGBA outputs. Our approach isolates the completed object as a layerable element, enabling straightforward integration in compositional applications.

In Section 3.1, we describe how the natural language query is processed to generate precise visible segmentation object mask. Section 3.2 details our occlusion analysis strategy, which entails identifying and masking occluders that occlude the target object. In Section 3.3, we present a heuristic method for selecting the inpainting text prompt, enabling targeted reconstruction of occluded regions based on the visible part of object. Section 3.4 outlines the iterative inpainting approach that progressively completes the occluded object and produces an RGBA output (an RGB image with an alpha mask).

Definitions. Let $I \in \mathbb{R}^{H \times W \times 3}$ represent the input image of height H and width W , with Q as a natural language query that specifies the target object in the scene. The query Q may explicitly name the object (e.g., “polar bear”) or describe it in context (e.g., “the animal in this image”). The output of the framework is an amodal completion $C \in \mathbb{R}^{H \times W \times 4}$, an RGBA representation where the first three channels represent RGB color values, and the

fourth channel (alpha) encodes transparency for adaptable layering in downstream applications. The amodal completion process includes generating an initial visible mask $M_{visible} \in \{0, 1\}^{H \times W}$, identifying the occluder mask M_{occ} , and iteratively updating these masks to achieve a fully reconstructed representation of the target object with background transparency.

3.1. Text Query Interpretation and Image Segmentation

One primary challenge in amodal completion is accurately isolating the target object in diverse, complex scenes, especially when the object is described in natural language. However, robust amodal completion also requires a comprehensive segmentation of all visible elements within the image, as any other object could potentially act as an occluder.

Given an input image $I \in \mathbb{R}^{H \times W \times 3}$ and a natural language query Q describing the target object, we first generate an initial mask $M_{visible} \in \{0, 1\}^{H \times W}$ that identifies the visible regions of I corresponding to the target object. This mask is derived from a pre-trained language-grounded vision model [15] that aligns Q with spatial areas relevant to the query, allowing the system to handle both specific and abstract descriptions. This initial mask $M_{visible}$ provides a foundation for understanding occlusion relationships with other elements in the scene.

Context-Aware Object Segmentation. While $M_{visible}$ offers an initial visible boundary for the target object, it is crucial to identify other objects and background areas within the image, as these may act as occluders. We employ the automatic image annotation system proposed in [27] to automatically identify all nameable objects in the image. The system first extracts a set of contextual tags $T = \{t_1, t_2, \dots, t_n\}$ from the image with a pre-trained

open-set image tagging model [11], where each tag t_i denotes a class label for a recognizable visual feature in I (e.g., “cat,” “plate,” or “bus”). Using T and I as inputs, the system leverage a pre-trained open-set object detector [20] combine with the Segment Anything model [13] to generate a set of segmentation masks $S = \{S_1, S_2, \dots, S_m\}$, where each mask S_i represents a visible objects within I . This context-aware segmentation framework enables the pipeline to delineate each visible object, including the target, with well-defined boundaries, ensuring all potential occluders are identified and isolated.

Handling Background Regions and Unknown Objects. In addition to identifiable objects, many scenes contain visually ambiguous or hard-to-describe areas, such as background elements, blurry or unidentifiable objects, clutter, and textures. Traditional segmentation methods may overlook these regions because they aren’t associated with object category labels, but these unlabelled (or “background”-labelled) regions can be occluders of a target object. Our approach handles these areas through an segmentation process which divides these “missing” areas into distinct segments without relying solely on object-category-driven segmentation.

After generating S , we identify unsegmented areas within I , defined as $B \subset I$, which consists of ambiguous background elements or hard-to-describe areas that may still interact with the target object through occlusion. These unsegmented regions pose a problem for occlusion analysis because they may be large and contain multiple objects or surfaces, only some of which may occlude the target object. To process these areas, we apply a sequence of morphological operations that refine and partition B into a collection of background segments. Specifically, we first generate a binary mask from B , where unsegmented pixels are assigned a value of 1 and all other pixels a value of 0. This binary mask is then eroded using a structured element to sharpen boundaries and separate loosely connected areas, preventing overlap with object segments in S . Following erosion, we apply dilation to the refined areas, which helps to re-expand and consolidate adjacent regions into distinct connected components, resulting in a set of background segments $B = \{B_1, B_2, \dots, B_k\}$, where:

$$B_j = \text{Morph}(I - \bigcup_{i=1}^m S_i), \quad \forall j \in \{1, 2, \dots, k\} \quad (1)$$

Here, Morph denotes the combined effect of the morphological operations applied to partition ambiguous regions into isolated segments.

This process allows our method of consider a wider range of occlusion conditions, such as target objects occluded by background (e.g., buildings partially hidden by undergrowth) and cases where the occluding objects are blurry, ambiguous, or difficult to recognize.

3.2. Occlusion Analysis

To reconstruct the missing parts of the target object, it is essential to identify where the object is occluded – these occluded areas are where completion is required. Our method generates an occluder mask for the target object that highlights the occluded areas. This mask is used to guide the inpainting process.

Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, the set of object masks $S = \{S_1, S_2, \dots, S_m\}$ and the set of background segments $B = \{B_1, B_2, \dots, B_k\}$, we define an occluder mask $M_{\text{occ}} \in \{0, 1\}^{H \times W}$ that aggregates all segments which occlude the target object.

Occluder Identification through Spatial Analysis. The occluder mask M_{occ} is derived by first evaluating the spatial relationships between the target’s visible mask M_{visible} and the set of segmented regions in $S \cup B$, where each segment potentially overlaps or interacts with the target. Our framework employs InstaOrderNet [16], a pre-trained model for occlusion orderings that takes pairwise segmentation masks and an image patch as input to determine the occlusion order between segment pairs, without requiring class labels. This approach enables identification of which segments within $S \cup B$ occlude the target object.

For each segment S_i or B_j in the scene, InstaOrderNet computes a binary occlusion indicator, assigning a value of $\text{occ} = 1$ if the segment occludes the target and 0 otherwise. The occluder mask M_{occ} is thus initialized as the union of all occluding segments in S and B identified by InstaOrderNet, formally expressed as:

$$M_{\text{occ}} = \bigcup_{\substack{i=1 \\ \text{occ}_i=1}}^m S_i \cup \bigcup_{\substack{j=1 \\ \text{occ}_j=1}}^k B_j \quad (2)$$

where \cup denotes pixel-wise union.

Boundary-Aware Occlusion. Scenes often present cases where the target object touches the image boundary, in which case amodal completion requires image expansion [35]. Thus we apply an adaptive refinement strategy that expands M_{occ} along the target’s boundary regions, guided by iterative boundary adjustments.

The process begins by checking for boundary contacts in M_{visible} , where the mask contacts any edge of I . For boundary-touching regions, we apply a dilation operation to the occluder mask, increasing its coverage along the contacted edges. Let $E \subset \{\text{top, bottom, left, right}\}$ denote the set of edges contacted by M_{visible} , and let d represent a structured dilation operation. We update M_{occ} as follows:

$$M_{\text{occ}} \leftarrow M_{\text{occ}} \cup \left(d(M_{\text{visible}}) \cap \bigcup_{e \in E} \text{edge}_e \right) \quad (3)$$

where edge_e refers to the pixels along the boundary edge e of the image, and \cap denotes pixel-wise intersection.

This expansion is applied iteratively, refining the occluder mask’s coverage until the mask stabilizes or the set of edges E is fully dilated to ensure coverage for the amodal completion process.

3.3. Prompt and Image Refinement for Inpainting

Because our system is designed to handle a wide range of natural language queries (e.g., “the animal in this image”) we cannot assume that the query alone will be sufficient to guide the inpainting step. We therefore combine the query with the image context to refine the inpainting prompt.

Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, visible mask M_{visible} , and a natural language query Q , we aim to produce an optimal inpainting prompt P that effectively guides the reconstruction process. Using a CLIP-based similarity comparison [26], we match the visible target object in $I_{\text{target}} = I \odot M_{\text{visible}}$ (where \odot denotes element-wise multiplication) with candidate descriptors from the image tags $T \cup Q$, selecting the descriptor that best matches the visible attributes of the target object. This similarity $S(t_i)$ for each candidate $t_i \in T \cup Q$ is computed as:

$$P = \arg \max_{t_i \in T \cup Q} \text{CLIP}(I_{\text{target}}, t_i) \quad (4)$$

where CLIP extracts features for visual-text alignment.

To further isolate the target, regions outside M_{visible} in I_{target} are swapped with a clean background I_{bkgd} , inspired by the background isolation approach in [35]. This prevents irrelevant areas from influencing the prompt. Our background swapping step ensures the prompt P focuses solely on the target object, improving inpainting precision. The prompt P is then used as a conditioning input for the inpainting model.

3.4. Iterative Inpainting and Amodal Completion

To achieve realistic amodal completion, our framework performs iterative masked inpainting to reconstruct occluded parts of the target object, terminating based on the occluder mask stability or a set iteration limit. After inpainting completes, a final seamless blending step merges the original visible regions with the reconstructed occluded parts, ensuring alignment with the original image and producing a transparent RGBA output that excludes unrelated background or extraneous elements.

Given an input image $I \in \mathbb{R}^{H \times W \times 3}$, visible mask M_{visible} , occluder mask M_{occ} , and inpainting prompt P , the reconstruction process begins by isolating the target object from the surrounding scene. We initialize an inpainting target I_{target} as follows:

$$I_{\text{target}} = I \odot M_{\text{visible}} + (1 - M_{\text{visible}}) \odot I_{\text{bkgd}} \quad (5)$$

where I_{bkgd} represents a clean background (similar to the natural, context-free background in [35]), ensuring that only the target object remains visible for inpainting.

The iterative process begins by setting $I_{\text{inpaint}}^{(0)} = I_{\text{target}}$, where $I_{\text{inpaint}}^{(t)}$ denotes the inpainting result at iteration t . At each step t , a pre-trained inpainting model [28] ϕ fills occluded regions in $I_{\text{inpaint}}^{(t)}$ based on P and the updated occluder mask $M_{\text{occ}}^{(t)}$:

$$I_{\text{inpaint}}^{(t+1)} = \phi(I_{\text{inpaint}}^{(t)}, M_{\text{occ}}^{(t)}, P) \quad (6)$$

In each iteration, the occluder mask $M_{\text{occ}}^{(t+1)}$ is refined to include newly inpainted regions, while an updated amodal segmentation mask M_{amodal} is computed to accurately bound the target object. M_{amodal} is computed iteratively as:

$$M_{\text{amodal}}^{(t+1)} = M_{\text{amodal}}^{(t)} \cup M_{\text{reconstructed}}^{(t+1)},$$

where $M_{\text{amodal}}^{(0)}$ is initialized as M_{visible} . $M_{\text{reconstructed}}^{(t+1)}$ is the mask for newly reconstructed regions, extracted from the inpainting output $I_{\text{inpaint}}^{(t+1)}$. This iterative refinement continues until reaching the final inpainted output $I_{\text{inpaint}}^{(T)}$, following the approach in [35].

Adaptive Termination for Efficiency. To ensure computational efficiency, we apply an adaptive termination criterion based on occluder mask stability. If the pixel-wise difference between consecutive occluder masks,

$$\Delta M_{\text{occ}}^{(t)} = \|M_{\text{occ}}^{(t+1)} - M_{\text{occ}}^{(t)}\|_{L_1} \quad (7)$$

falls below a threshold ϵ or the inpainting process reaches a maximum number of iterations T , the iteration stops, yielding the final inpainted result $I_{\text{inpaint}}^{(T)}$.

Blending for Seamless Integration. Once inpainting is complete, a final alpha blending step is applied to achieve a smooth transition between newly reconstructed and original visible region of the object. This blending operates from the interior toward a “transition region” of a predefined width along the boundary of the visible region. Specifically, the center of the original visible region maintains an alpha value of 1, and areas outside the visible region are set to 0. Within the transition region, alpha values gradually decrease from 1 to 0. The blended image I_{blend} is defined as:

$$I_{\text{blend}} = M_{\alpha} \cdot I_{\text{inpaint}}^{(T)} + (1 - M_{\alpha}) \cdot I_{\text{orig}}. \quad (8)$$

RGBA Output for Versatile Use. The final output is an RGBA image $C \in \mathbb{R}^{H \times W \times 4}$, with the first three channels corresponding to the RGB values of I_{blend} and the fourth channel M_{amodal} representing the binary amodal segmentation mask.

4. Experiments

4.1. Datasets Collection

Existing datasets for amodal appearance completion often have limitations, either due to constraints on a limited




Figure 3. Distribution of the top 50 most frequent categories in the our evaluation dataset.

Source	VG	COCO-A	Free Image	Laion	Total
# Img	1234	751	228	166	2379
# Instances	1294	845	249	177	2565
# Classes	269	256	147	97	553
# Img w/ 1 occldee	1184	663	211	156	2214
# Img w/ 1+ occldees	50	88	17	10	165

Table 1. Composition of the evaluation dataset. The broad variety of classes across different sources presents a challenge for open-world amodal appearance completion.

set of object categories or reliance on synthetic occlusion data [5, 6, 10, 18, 31, 44], which fall short of capturing the complexity of real-world scenarios. To provide a comprehensive evaluation for the open-world amodal appearance completion challenge, we gathered a diverse dataset with natural occlusions and a broad range of object categories across varied scenes.

Our evaluation dataset draws from four sources: COCO-A [44], Visual Genome (VG) [14], LAION [29], and copyright-free images from publicly accessible websites. Each source brings unique strengths to our evaluation dataset. COCO-A and VG provide natural scenes containing occluded objects real-world environments. Given that COCO-A was initially designed for semantic segmentation, not all images feature object-specific occlusions. Thus, we applied a filtering process to COCO-A, retaining only those images in which objects are occluded by other objects.¹ The addition of LAION and free-image sources introduces the diversity of internet-sourced images, covering a wide range of settings, lighting, and complex occlusion types characteristic of unconstrained environments.

The resulting dataset consists of 2379 images spanning 553 distinct target object classes (see Tab. 1). Three human annotators collected images containing occluded objects and provided a class label for each occluded object. Fig. 3 presents the top 50 most frequent occluded object cat-

¹Since the COCO-A dataset was originally developed for semantic segmentation, it focuses on semantic-level occlusions—such as occluded background elements—rather than solely object-level occlusions. Consequently, not all images contain relevant object occlusions; we removed these from our COCO-A selection. Details of the selection process are provided in the appendix.

egories, illustrating the distribution across classes. This diversity ensures that methods can be evaluated across a wide range of object categories and occlusion scenarios, resulting in a challenging testbed for open-world amodal appearance completion.

4.2. Implementation and Evaluation Metrics

Implementation Details. Our framework leverages publicly available models without additional training or fine-tuning to deliver high-quality amodal completions. The main models utilized include the Stable Diffusion v2 inpainting model [28] and the LISA-13B-llama2-v1 model [15], both publicly available.

All experiments were conducted on an NVIDIA A100 GPU. To ensure a fair comparison, we provide all methods with the same visible part of the occluded object as input, as other methods do not support text-based queries. For comparison methods, we ran inference using their publicly available pre-trained models with default settings. Our framework’s design maintains computational efficiency by limiting the inpainting process to a maximum of three iterations per object; further details on our configurations are provided in the appendix.

Evaluation metrics. Evaluating amodal completion on natural images with real-world occlusions presents unique challenges, as the ground-truth appearance of occluded regions is inherently unavailable. Thus, we use a combination of human evaluation and quantitative metrics. We center our evaluation on human assessment, which is critical for judging the subjective quality of amodal completion outputs.

We use CLIP score [26] to measure how well the amodal completions align with the class labels. Specifically, we compare the amodal completion of each object with the ground-truth class label associated with the object.

We also include appearance quantitative metrics. As only the visible parts of each object are available, these metrics offer a rough reference of completion quality. We compute LPIPS [41] for visual consistency, VGG16 feature similarity [7] for semantic consistency, SSIM and PSNR [34] for structural consistency—all by comparing the visible part of the object with the completed version.



Figure 4. Visual comparisons of amodal completions across different methods: Ours consistently outperforms others in terms of realism, handling complex occlusions, and producing plausible completions. Top to bottom: examples from VG, COCO-A, Free Image, LAION.

Dataset/ Method	VG	COCO-A	Free Image	LAION	Overall
PD w/o MC [†]	16.62%	15.38%	15.80%	11.49%	15.78%
PD-MC [†]	15.51%	14.20%	12.99%	9.42%	14.41%
Pix2gestalt	26.56%	31.24%	21.29%	31.83%	27.95%
Ours [†]	41.32%	39.17%	49.93%	47.27%	41.86%

Table 2. Human preference study across four different sources. The table shows the percentage of participants who preferred each method’s amodal completion results. [†] Methods that explicitly employ outpainting.

4.3. Comparisons with Other Methods

Since amodal completion is inherently subjective, human evaluation serves as a reliable measure of realism and completion quality. We conducted a preference study on Prolific crowdsourcing platform² with 180 participants, comparing our method to Pix2gestalt [23], PD-MC [35], and a baseline PD w/o MC (PD-MC without Mixed Context Diffusion). Participants were presented with the original image and four completed versions, each randomized to avoid positional bias. They were instructed to choose the completion that appeared most realistic, considering both visible and occluded regions. To ensure data reliability, “gold standard” questions with straightforward choices were included, and only participants who passed 75% of these checks were retained in the final analysis.

Our method achieved the highest overall preference in the evaluation dataset, as shown in Tab. 2. Across all four




Figure 5. Model preference of human evaluators by agreement levels. X-axis shows the number of images. 3/3 denotes full agreement among three evaluators per image, 1/3 indicates no consensus. “Ours” shows the strongest consensus on completion quality.

sources, our approach consistently ranked first, demonstrating robustness in handling diverse occlusions across open-world scenarios. The agreement levels among participants further highlight our performance. As illustrated in Fig. 5, in cases where all three participants agreed on a single preferred completion (3/3), our method was selected 57% of the time, followed by Pix2gestalt at 28%, showing the perceptual quality and realism of our results. Even in cases of partial consensus (2/3) or no majority (1/3), our approach still retained a lead, indicating a consistent preference even when there was divergence in evaluator opinion.

Fig. 4 shows visual comparisons of our method with others on challenging mutual occlusions across open-world scenarios. In the dog and horse images, which involve extensive outpainting, only our approach reconstructs the full body accurately. Similarly, in the book and Oscar statuette examples, ours achieves coherent and realistic completions, preserving structural details that other methods miss. This highlights our method’s strength in handling complex ob-

²<https://www.prolific.com/>

Dataset	Method	Class Relevance	Visual Consistency	Semantic Consistency	Structural Consistency	
		↑ CLIP	↓ LPIPS	↑ Feature Similarity	↑ SSIM	↑ PSNR
VG	PD w/o MC	28.254	0.586	0.406	0.459	7.855
	PD-MC	28.367	0.578	0.413	0.463	7.986
	Pix2gestalt	27.672	0.429	0.554	0.726	11.318
	Ours	28.470	0.310	0.658	0.732	13.075
COCO-A	PD w/o MC	27.426	0.673	0.318	0.381	6.246
	PD-MC	27.383	0.664	0.328	0.382	6.394
	Pix2gestalt	26.998	0.471	0.524	0.695	10.601
	Ours	27.612	0.351	0.609	0.718	12.184
Free Images	PD w/o MC	28.190	0.730	0.268	0.305	5.128
	PD-MC	28.333	0.720	0.279	0.309	5.303
	Pix2gestalt	27.621	0.393	0.613	0.732	12.356
	Ours	28.652	0.269	0.698	0.753	13.720
LAION	PD w/o MC	27.479	0.695	0.295	0.354	6.230
	PD-MC	27.573	0.692	0.299	0.346	6.245
	Pix2gestalt	27.260	0.467	0.527	0.691	11.258
	Ours	28.123	0.319	0.657	0.751	13.602
Overall	PD w/o MC	27.922	0.636	0.356	0.411	6.948
	PD-MC	27.984	0.628	0.364	0.413	7.081
	Pix2gestalt	27.417	0.442	0.548	0.714	11.178
	Ours	28.181	0.320	0.646	0.731	12.880

Table 3. Quantitative comparisons across datasets. Our method consistently outperforms other approaches across all evaluated metrics, indicating superior visual and perceptual fidelity in open-world amodal appearance completion.

jects and challenging occlusions with superior realism.

Quantitative Metrics. We assessed completion quality using quantitative metrics: CLIP score for class relevance, and LPIPS, VGG16 feature similarity, and SSIM for appearance quality. As shown in the Tab. 3, our method achieved the highest CLIP score, aligning well with ground-truth object categories and showcasing strong relevance. In appearance metrics, our method also outperformed others, indicating superior visual, semantic and structural consistency. Failure cases where no amodal completion is generated result in LPIPS of 1 and feature similarity/SSIM of 0. While these appearance quantitative metrics provide insight, they are provided for reference only, as the amodal appearance ground-truth is not available.

4.4. Ablation Studies

Inpainting Prompt Variants. We tested three prompt configurations: (1) Q (text query only), (2) T (auto-generated tags only), and (3) $T \cup Q$ (combined tags and query). As shown in Tab. 4, the improvement in all appearance quantitative metrics for $T \cup Q$ indicates its effectiveness in visual, semantic and structural consistency. In our evaluation dataset, Q prompt achieves the highest CLIP score, because it directly aligns with ground-truth target object class labels. However, in practical scenarios, Q may consist of abstract or context-dependent queries rather than specific class labels, making the combined $T \cup Q$ approach more effective by incorporating both object-specific and contextual details.

Background Segments Effectiveness. Background segments (see Sec. 3.1) contribute greatly to improving amodal completion in open-world contexts by handling occlusions

Ablation		Class Relevance	Visual Consistency	Semantic Consistency	Structural Consistency
Inpainting Prompt	Background Segments	↑ CLIP	↓ LPIPS	↑ Feature Similarity	↑ SSIM
Q	✓	28.563	0.327	0.633	0.724
T	✓	28.043	0.324	0.636	0.725
$T \cup Q$	✗	28.071	0.333	0.620	0.713
$T \cup Q$	✓	28.181	0.320	0.646	0.731

Table 4. Ablation study results on the full evaluation dataset for prompt variants and background segmentation effects across class relevance and appearance quantitative metrics.



Figure 6. Amodal completion results from our method, with and without considering background segments as potential occluders.

and ambiguous backgrounds more effectively. As shown in Fig. 6, including background segments in the occlusion reasoning step makes it possible to capture the occlusion relationships between a query object and ambiguous background elements like foliage and ground, making scene understanding more comprehensive. This is particularly effective in the open-world setting, where occlusions are more diverse and complex.

5. Conclusion

We presented a novel framework for open-world amodal appearance completion that reconstructs occluded objects across diverse categories without additional training. Our approach outperforms existing methods in user preference and objective quality metrics over an evaluation dataset of complex object occlusion cases, collected from four different data sources. This amodal completion framework readily supports applications like image editing, novel view synthesis, and 3D reconstruction [30, 31, 39], making it versatile for real-world use.

Limitations. Our reliance on pre-trained image generation models can occasionally introduce artifacts, such as mismatched poses in animal completions (e.g., generating a standing dog when it should be sitting based on the context; see appendix for more failure examples). However, our modular design allows for easy updates as better pre-trained models become available. Future work could improve inpainting accuracy and develop new evaluation metrics for scenarios where ground-truth amodal data is unavailable.

Acknowledgement. This research is supported by Australian Research Council funding (Grant No: DE250100030) and the University of Melbourne’s Research Computing Services and Petascale Campus Initiative.

References

- [1] Jiayang Ao, QiuHong Ke, and Krista A Ehinger. Image amodal completion: A survey. *Computer Vision and Image Understanding*, 229:103661, 2023. 2
- [2] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kashten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *European Conference on Computer Vision*, pages 707–723. Springer, 2022. 1
- [3] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 1
- [4] Helisa Dhamo, Nassir Navab, and Federico Tombari. Object-driven multi-layer scene decomposition from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5369–5378, 2019. 2
- [5] Kiana Ehsani, Roozbeh Mottaghi, and Ali Farhadi. Segan: Segmenting and generating the invisible. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6144–6153, 2018. 2, 6
- [6] Patrick Follmann, Rebecca König, Philipp Härtinger, Michael Klostermann, and Tobias Böttger. Learning to see the invisible: End-to-end trainable amodal instance segmentation. In *2019 IEEE Winter Conference on Applications of Computer Vision*, pages 1328–1336. IEEE, 2019. 6
- [7] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2414–2423, 2016. 6
- [8] Vasileios Gkitsas, Vladimiro Sterzentsenko, Nikolaos Zioulis, Georgios Albanis, and Dimitrios Zarpalas. Panodr: Spherical panorama diminished reality for indoor scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3716–3726, 2021. 1
- [9] Daniel Honerkamp, Martin Büchner, Fabien Despinoy, Tim Welschhold, and Abhinav Valada. Language-grounded dynamic scene graphs for interactive object search with mobile manipulation. *IEEE Robotics and Automation Letters*, 2024. 1
- [10] Yuan-Ting Hu, Hong-Shuo Chen, Kexin Hui, Jia-Bin Huang, and Alexander G Schwing. Sail-vos: Semantic amodal instance level video object segmentation-a synthetic dataset and baselines. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3105–3115, 2019. 6
- [11] Xinyu Huang, Yi-Jie Huang, Youcai Zhang, Weiwei Tian, Rui Feng, Yuejie Zhang, Yanchun Xie, Yaqian Li, and Lei Zhang. Open-set image tagging with multi-grained text supervision. *arXiv e-prints*, pages arXiv–2310, 2023. 4
- [12] Gaetano Kanizsa, Paolo Legrenzi, and Paolo Bozzi. Organization in vision: Essays on gestalt perception. 1979. 1
- [13] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 4
- [14] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International Journal of Computer Vision*, 123:32–73, 2017. 6
- [15] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. Lisa: Reasoning segmentation via large language model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9579–9589, 2024. 2, 3, 6
- [16] Hyunmin Lee and Jaesik Park. Instance-wise occlusion and depth orders in natural scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21210–21221, 2022. 4
- [17] Zhixuan Li, Weining Ye, Tingting Jiang, and Tiejun Huang. 2d amodal instance segmentation guided by 3d shape prior. In *European Conference on Computer Vision*, pages 165–181. Springer, 2022. 1
- [18] Zhixuan Li, Weining Ye, Juan Terven, Zachary Bennett, Ying Zheng, Tingting Jiang, and Tiejun Huang. Muva: A new large-scale benchmark for multi-view amodal instance segmentation in the shopping scenario. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23504–23513, 2023. 6
- [19] Huan Ling, David Acuna, Karsten Kreis, Seung Wook Kim, and Sanja Fidler. Variational amodal object completion. *Advances in Neural Information Processing Systems*, 33: 16246–16257, 2020. 1, 2
- [20] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*. Springer, 2024. 4
- [21] Zhengzhe Liu, Qing Liu, Chirui Chang, Jianming Zhang, Daniil Pakhomov, Haitian Zheng, Zhe Lin, Daniel Cohen-Or, and Chi-Wing Fu. Object-level scene deocclusion. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024. 2
- [22] Yuchen Mo, Hanbo Zhang, and Tao Kong. Towards open-world interactive disambiguation for robotic grasping. In *2023 IEEE International Conference on Robotics and Automation*, pages 8061–8067. IEEE, 2023. 1
- [23] Ege Ozguroglu, Ruoshi Liu, Dídac Surís, Dian Chen, Achal Dave, Pavel Tokmakov, and Carl Vondrick. pix2gestalt: Amodal segmentation by synthesizing wholes. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3931–3940. IEEE Computer Society, 2024. 1, 2, 7
- [24] Dim P Papadopoulos, Youssef Tamaazousti, Ferda Ofli, Ingmar Weber, and Antonio Torralba. How to make a pizza: Learning a compositional layer-based gan model. In *proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8002–8011, 2019. 2
- [25] Giovanni Pintore, Marco Agus, Eva Almansa, and Enrico Gobbetti. Instant automatic emptying of panoramic indoor scenes. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3629–3639, 2022. 1

- [26] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. [5](#), [6](#)
- [27] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded sam: Assembling open-world models for diverse visual tasks. *arXiv preprint arXiv:2401.14159*, 2024. [3](#)
- [28] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. [5](#), [6](#)
- [29] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. [6](#)
- [30] Ruoxi Shi, Hansheng Chen, Zhuoyang Zhang, Minghua Liu, Chao Xu, Xinyue Wei, Linghao Chen, Chong Zeng, and Hao Su. Zero123++: a single image to consistent multi-view diffusion base model. *arXiv preprint arXiv:2310.15110*, 2023. [8](#)
- [31] Petru-Daniel Tudosiu, Yongxin Yang, Shifeng Zhang, Fei Chen, Steven McDonagh, Gerasimos Lampouras, Ignacio Iacobacci, and Sarah Parisot. Mulan: A multi layer annotated dataset for controllable text-to-image generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22413–22422, 2024. [1](#), [2](#), [6](#), [8](#)
- [32] Wenhai Wang, Zhe Chen, Xiaokang Chen, Jiannan Wu, Xizhou Zhu, Gang Zeng, Ping Luo, Tong Lu, Jie Zhou, Yu Qiao, et al. Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. *Advances in Neural Information Processing Systems*, 36, 2024. [2](#)
- [33] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. Cogvlm: Visual expert for pretrained language models. *Advances in Neural Information Processing Systems*, 36, 2024. [2](#)
- [34] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. [6](#)
- [35] Katherine Xu, Lingzhi Zhang, and Jianbo Shi. Amodal completion via progressive mixed context diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9099–9109, 2024. [1](#), [2](#), [4](#), [5](#), [7](#)
- [36] Xiaosheng Yan, Feigege Wang, Wenxi Liu, Yuanlong Yu, Shengfeng He, and Jia Pan. Visualizing the invisible: Occluded vehicle segmentation and recovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7618–7627, 2019. [2](#)
- [37] Guanqi Zhan, Chuanxia Zheng, Weidi Xie, and Andrew Zisserman. Amodal ground truth and completion in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 28003–28013, 2024. [1](#)
- [38] Xiaohang Zhan, Xingang Pan, Bo Dai, Ziwei Liu, Dahua Lin, and Chen Change Loy. Self-supervised scene de-occlusion. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 3784–3792, 2020. [2](#)
- [39] Lvmin Zhang and Maneesh Agrawala. Transparent image layer diffusion using latent transparency. *ACM Trans. Graph.*, 43(4), 2024. [1](#), [2](#), [8](#)
- [40] Ni Zhang, Nian Liu, Junwei Han, Kaiyuan Wan, and Ling Shao. Face de-occlusion with deep cascade guidance learning. *IEEE Transactions on Multimedia*, 25:3217–3229, 2022. [2](#)
- [41] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018. [6](#)
- [42] Chuanxia Zheng, Duy-Son Dao, Guoxian Song, Tat-Jen Cham, and Jianfei Cai. Visiting the invisible: Layer-by-layer completed scene decomposition. *International Journal of Computer Vision*, 129:3195–3215, 2021. [2](#)
- [43] Qiang Zhou, Shiyin Wang, Yitong Wang, Zilong Huang, and Xingang Wang. Human de-occlusion: Invisible perception and recovery for humans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3691–3701, 2021. [2](#)
- [44] Yan Zhu, Yuandong Tian, Dimitris Metaxas, and Piotr Dollár. Semantic amodal segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1464–1472, 2017. [6](#)