



Object Centric Visual Question Answering for Robotic Assembly Inspection and Workflow Fault Diagnosis Tasks

Marina Alvarez, Ines Pereira, Hannah Becker

Faculty of Science, University of Auckland, Auckland, New Zealand

Abstract

The rapid evolution of modern manufacturing paradigms necessitates highly adaptable and intelligent robotic assembly inspection cells. Traditional machine vision systems rely on rigid, pre-programmed rules and spatial templates, which frequently fail when confronted with unstructured environments, diverse product portfolios, or unexpected anomalies. This paper proposes a novel framework centered on Object-Centric Visual Question Answering to address the growing demand for flexible and cognitive inspection methodologies. By treating the robotic cell as a dynamic environment where human operators can query the system using natural language, the proposed framework enables intuitive interaction and diagnostic transparency. Unlike standard global-feature visual question answering models, our approach parses visual scenes into discrete, object-centric representations. This allows the system to ground natural language queries to specific mechanical components, fasteners, and spatial relationships within the assembly workspace. The research details the architectural design, object discovery mechanisms, and cross-modal reasoning required to achieve high-fidelity inspection results. Extensive evaluations on a custom dataset of robotic assembly scenarios demonstrate that the object-centric approach significantly outperforms global feature baselines in terms of accuracy, interpretability, and robustness against occlusions. Ultimately, this framework paves the way for more collaborative, human-centric robotic manufacturing systems capable of complex reasoning and real-time fault diagnosis.

Keywords: *Visual Question Answering, Robotic Assembly, Object-Centric Representation, Automated Inspection*

1. Introduction

The transition toward intelligent manufacturing and Industry 4.0 has fundamentally altered the operational requirements of robotic assembly lines. Historically, industrial automation relied on rigid programming and deterministic systems designed for high-volume, low-mix production. In these traditional setups, quality assurance and inspection were treated as separate, downstream processes heavily dependent on either human visual inspection or highly constrained machine vision algorithms. These algorithms typically utilized template matching, edge detection, and fixed-coordinate metrology to verify the presence, orientation, and integrity of components [1]. However, as consumer demand shifts toward customized, high-mix, low-volume production, the rigid nature of conventional inspection systems has become a significant bottleneck. Modern robotic assembly cells must continuously adapt to varying product geometries, shifting lighting conditions, and unpredictable spatial configurations. Consequently, there is an urgent need for cognitive robotic systems that not only perceive their environment but also understand and reason about the complex physical relationships between assembled parts [2]. Visual Question Answering has emerged as a powerful paradigm at the intersection of computer vision and natural language processing. Originally developed for general-purpose image understanding, a visual question answering system takes an image and a natural

language question as input and produces a natural language answer as output. In the context of a robotic assembly inspection cell, this capability transforms the interaction dynamic between human operators and automated systems [3]. Instead of reprogramming a vision system to detect a specific new defect or verify a novel component layout, a quality assurance engineer can simply ask the system questions such as whether the blue wire is connected to the correct terminal or if all screws on the flange are fully tightened. This enables a flexible, intuitive, and highly adaptable inspection process. By leveraging natural language as the primary interface, manufacturing facilities can drastically reduce the setup time for new product variations and democratize the inspection process, allowing operators without extensive computer vision expertise to conduct complex diagnostic queries [4]. Despite the theoretical advantages, deploying standard visual question answering models in industrial environments presents profound challenges. Conventional models predominantly utilize global feature extraction techniques, where an entire image is compressed into a dense, monolithic vector representation before being fused with the text query [5]. While this approach is computationally efficient and performs adequately on general datasets containing distinct, macroscopic subjects, it struggles in the dense, cluttered, and highly homogenous environments typical of robotic assembly. In a robotic cell, the scene often consists of dozens of nearly identical metallic components, overlapping wires, and complex mechanical assemblies. Global representations fail to capture the fine-grained, localized details necessary to distinguish between a correctly seated bolt and one that is cross-threaded. When global models attempt to answer relational or counting questions in such environments, they frequently suffer from catastrophic hallucination, providing confident but incorrect answers because they lack a fundamental understanding of the individual entities comprising the scene. To overcome the limitations of global feature extraction, this paper introduces an Object-Centric Visual Question Answering framework explicitly designed for robotic assembly inspection cells. Object-centric representation learning posits that a visual scene should not be treated as a single continuous signal, but rather as a composition of distinct, localized entities or objects. By forcing the perception module to decompose the robotic workspace into a set of discrete object representations, the downstream reasoning module can explicitly map the semantic concepts in the natural language query to the physical components in the scene. This explicit grounding mechanism ensures that when an operator asks about a specific bracket or fastener, the model directs its computational attention strictly to the relevant physical entities, thereby ignoring background noise and irrelevant components. The object-centric approach inherently supports compositional reasoning, enabling the system to evaluate complex spatial relationships, verify multi-step assembly sequences, and identify subtle anomalies that would be obscured in a global feature space. This comprehensive academic paper details the theoretical foundations, architectural design, and empirical validation of the proposed object-centric framework. Section 2 provides an exhaustive review of the relevant literature encompassing robotic assembly inspection, visual question answering, and object-centric representation learning. Section 3 outlines the methodology, detailing the data acquisition pipeline, the object discovery algorithms, and the cross-modal reasoning architecture without relying on complex mathematical notations. Section 4 presents a rigorous analysis of the experimental results, comparing the proposed system against state-of-the-art baselines across multiple challenging inspection scenarios. Finally, Section 5 concludes the paper by summarizing the contributions, discussing current limitations, and highlighting promising directions for future research in cognitive industrial robotics.

2. Literature Review

2.1 Robotic Assembly and Automated Inspection

The domain of robotic assembly has seen a steady progression from entirely deterministic automation to systems capable of limited adaptive behavior. In the early stages of industrial automation, robotic cells were designed with strict environmental controls. Parts were fed into the workspace using precision fixtures, and inspection was performed using hard-coded geometric rules [6]. These early machine vision systems relied heavily on classical image processing techniques, such as morphological operations, thresholding, and contour analysis, to determine whether an assembly met predefined quality standards. While highly effective for mass production of simple goods, these techniques

demonstrated severe fragility when exposed to environmental variations [7]. Minor changes in ambient factory lighting, slight deviations in part surface finish, or the introduction of a new component iteration necessitated complete recalibration of the vision algorithms. The introduction of deep learning, particularly convolutional neural networks, marked a paradigm shift in automated inspection. Deep learning allowed systems to learn robust feature representations directly from data, significantly reducing the reliance on handcrafted features [8]. Researchers developed sophisticated models for object detection, semantic segmentation, and anomaly detection that could operate reliably despite variations in lighting and component positioning. Convolutional networks became the backbone of modern industrial inspection, enabling robots to identify surface defects, measure tolerances, and verify assembly completeness with unprecedented accuracy [9]. However, these systems remained fundamentally task-specific. A neural network trained to detect missing screws on a specific engine block could not be easily repurposed to verify the wiring harness of a different product without undergoing a costly and time-consuming retraining process. This lack of generalization and task flexibility remains a critical barrier in modern, agile manufacturing environments [10]. Recent advancements in cognitive robotics have attempted to address this inflexibility by integrating higher-level semantic understanding into robotic cells. Researchers have explored the use of knowledge graphs, ontologies, and semantic scene parsing to give robots a structured understanding of their environment and the tasks they are performing [11]. By representing the assembly process as a sequence of logical operations on discrete components, systems can achieve a degree of reasoning. For instance, if a system knows that a specific bracket requires four bolts, it can autonomously generate inspection routines to verify the presence of all four. Despite these improvements, translating high-level semantic rules into robust visual perception in cluttered scenes remains a formidable challenge [12]. The gap between abstract logical reasoning and raw pixel data continues to drive research into more sophisticated multimodal perception systems.

2.2 Visual Question Answering Systems

Visual Question Answering represents a significant milestone in multimodal artificial intelligence, requiring the seamless integration of computer vision and natural language processing. Early models treated the problem as a straightforward classification task, using convolutional networks to extract image features and recurrent neural networks to process the text query [13]. The resulting feature vectors were concatenated or multiplied and fed into a classifier to predict the correct answer from a predefined vocabulary. While these early architectures demonstrated proof-of-concept capabilities on simple datasets, they struggled with complex questions requiring spatial reasoning, counting, or multi-step logic. The models often relied on linguistic biases in the training data rather than genuinely understanding the visual content [14]. For example, if a dataset contained many images where the sky was blue, the model might answer blue to any question containing the word sky, regardless of the actual image content. The introduction of attention mechanisms revolutionized visual question answering. Attention allowed models to dynamically focus on specific regions of the image based on the words in the natural language query. This spatial attention significantly improved performance on spatial reasoning and object identification tasks [15]. Instead of compressing the entire image into a single vector, the network maintained a grid of localized feature vectors and learned to weight their importance dynamically. Furthermore, the advent of transformer architectures and large-scale pre-training strategies led to the development of massive vision-language models capable of unprecedented generalization [16]. Models pre-trained on vast quantities of image-text pairs demonstrated the ability to answer diverse questions about general domain images with remarkable accuracy. However, the application of visual question answering in specialized domains, such as medical imaging or industrial inspection, has revealed significant limitations in general-purpose models. Industrial scenes are characterized by dense clutter, fine-grained details, and highly specialized vocabulary that do not exist in general pre-training datasets [17]. When applied to a robotic assembly cell, a standard attention-based model might correctly identify the general region containing a motor but fail to distinguish between the various tiny connectors and fasteners attached to it. The attention maps generated by these models are often diffuse and fail to align tightly with the physical boundaries of the components,

leading to critical errors in inspection tasks [18]. Consequently, there is a growing consensus that robust industrial visual question answering requires architectural innovations that explicitly model the compositional nature of the workspace.

2.3 Object-Centric Representations in Computer Vision

The concept of object-centric representation is rooted in human cognitive psychology, which suggests that humans perceive the world not as a continuous field of color and light, but as a collection of distinct, interacting objects. In computer vision, translating this cognitive principle into computational models has been a long-standing objective [19]. Early attempts utilized unsupervised segmentation and grouping algorithms to partition images into coherent regions. However, these methods were computationally expensive and struggled with complex textures and lighting variations. The rise of deep learning initially shifted the focus away from explicit object modeling toward implicit, distributed representations learned by deep convolutional layers. While powerful, these distributed representations lack the compositional structure necessary for complex logical reasoning [20]. Recently, there has been a resurgence of interest in explicit object-centric architectures. Techniques such as slot attention and object files have been developed to force neural networks to parse scenes into a set of discrete, permutation-invariant vectors, each representing a single object or entity within the scene [21]. In a slot attention module, a set of random vectors, referred to as slots, iteratively compete to explain different parts of the input image features. Over multiple iterations, each slot binds to a specific object, resulting in a clean, disentangled representation of the scene components. This approach has demonstrated remarkable success in unsupervised object discovery, relational reasoning, and physical simulation tasks in synthetic environments [22]. Applying object-centric representations to real-world, high-resolution industrial scenes presents unique challenges. The objects in a robotic assembly cell are often highly textured, partially occluded, and varying wildly in scale, from massive robotic arms to millimeter-wide screws [23]. To address this, researchers have begun combining unsupervised object discovery techniques with supervised object detection backbones. By leveraging pre-trained object detectors to generate region proposals and utilizing slot-based mechanisms to refine these proposals into relational object tokens, systems can achieve the structural clarity required for complex reasoning [24]. In the context of visual question answering, these object tokens serve as the fundamental building blocks for cross-modal alignment, allowing the language model to interact directly with the physical components of the scene rather than abstract image patches [25]. This paradigm shift from pixel-level attention to object-level reasoning forms the core theoretical foundation of the proposed inspection framework.

3. Methodology

3.1 System Architecture Overview

The proposed Object-Centric Visual Question Answering framework is designed to operate within a highly flexible robotic assembly inspection cell. The physical environment consists of a multi-axis robotic manipulator equipped with a high-resolution, wrist-mounted primary camera, complemented by several fixed overhead cameras that provide global context of the workspace. The assembly cell is designed to handle a diverse range of mechanical assemblies, requiring the system to continuously adapt its inspection criteria based on operator input. The core architecture of the intelligent system is divided into three primary modules operating in a continuous pipeline. The first module is the data acquisition and preprocessing stage, which handles the synchronization of camera feeds, illumination normalization, and spatial calibration. The second module is the object discovery and representation engine, which takes the raw pixel data and translates it into a discrete set of semantic object tokens. The third and final module is the multimodal reasoning and answer generation engine, which interprets the natural language query, aligns it with the object tokens, and formulates a diagnostic response.

To ensure robustness in an industrial setting, the system architecture prioritizes modularity and interpretability. The separation of the object discovery phase from the reasoning phase allows for independent optimization and debugging of each component. For instance, if the system fails to correctly answer a question regarding a specific wire routing, engineers can inspect the intermediate object tokens to determine whether the failure occurred in the perception module (the wire was not

detected as a distinct object) or in the reasoning module (the spatial relationship was misunderstood). This transparency is a critical requirement for deploying artificial intelligence in safety-critical and quality-critical manufacturing environments.

Table 1: Key parameters and architectural specifications of the proposed neural network modules.

Module Component	Parameter Description	Configuration Value
Object Extractor Backbone	Network Depth and Type	ResNet-101 with Feature Pyramid
Bounding Box Predictor	Proposal Threshold	Confidence > 0.65
Feature Dimensionality	Latent Vector Size	512 dimensions per token
Multimodal Transformer	Number of Layers	12 Cross-Attention Layers
Attention Heads	Parallel Processing Units	8 Heads per Layer
Text Encoder	Pre-trained Language Model	BERT-Base Uncased

The processing pipeline begins when an operator submits a text query via a digital interface. Simultaneously, the robotic manipulator positions the primary camera to capture the relevant section of the assembly based on a predefined or dynamically calculated trajectory. The captured high-resolution image is fed into the visual processing pipeline while the text query is tokenized and encoded by a linguistic model. The ultimate output is a natural language response detailing the inspection result, which may range from a simple confirmation of component presence to a detailed description of an assembly fault.

3.2 Object Discovery and Representation

The object discovery and representation module is the most critical innovation in the proposed framework. Its objective is to move beyond the traditional grid-based feature extraction typical of convolutional networks and vision transformers. In standard approaches, an image is divided into a grid of patches, and each patch is processed into a feature vector. This results in a scenario where a single mechanical component, such as a large gear, might be split across multiple arbitrary patches, while a single patch might contain fragments of multiple overlapping wires. This arbitrary division destroys the structural integrity of the objects, making it incredibly difficult for downstream reasoning modules to answer questions about specific components [26]. To resolve this, our framework employs a region-based object extraction pipeline enhanced with attention mechanisms to create clean, disentangled object tokens. The process begins with a deep convolutional backbone that processes the high-resolution input image to produce a dense feature map. This feature map is then fed into a region proposal network that identifies areas of the image likely to contain discrete components, fasteners, or structural elements. Unlike standard object detection, which immediately classifies these regions into rigid categories, our system uses these proposals as spatial priors to extract rich, high-dimensional feature vectors. Each proposed region is pooled and projected into a fixed-length vector space, creating an initial set of object candidates.

Code Listing 1: Data structure definition for the object tokenization pipeline.

```
class ObjectTokenExtractor:
    def __init__(self, visual_backbone, region_proposer, embed_dim):
        self.backbone = visual_backbone
        self.proposer = region_proposer
        self.projection_layer = Linear(in_features=2048, out_features=embed_dim)

    def extract_tokens(self, input_image):
        # Generate dense feature map from the raw image
        feature_map = self.backbone(input_image)

        # Identify bounding boxes for potential components
```

```

region_boxes = self.proposer(feature_map)

# Extract features corresponding to each bounding box
pooled_features = region_align(feature_map, region_boxes)

# Project features into the final object token space
object_tokens = self.projection_layer(pooled_features)

return object_tokens, region_boxes

```

Following the initial extraction, these candidates undergo a refinement process using self-attention mechanisms. The objects in a robotic assembly cell do not exist in isolation; their functional and spatial context is highly dependent on neighboring components. The self-attention layer allows each object token to exchange information with all other tokens in the scene. For example, a token representing a nut can aggregate information from a nearby token representing a bolt, enriching its own representation with relational context. This iterative refinement produces a final set of object tokens where each vector encapsulates not only the visual appearance of a specific component but also its spatial position and contextual relationship with the surrounding assembly.




Figure 1: Comprehensive Schematic of the Object

This methodology ensures that the complexity of the scene is reduced to a manageable set of highly informative vectors. Even in scenarios involving significant occlusion or visual clutter, the region-based extraction ensures that the network maintains a distinct representation for each visible entity. This object-centric foundation is what enables the system to handle the complex, multi-step queries required for rigorous industrial inspection.

3.3 Reasoning and Answer Generation

Once the visual scene has been parsed into a set of distinct object tokens and the operator's natural language query has been encoded into a sequence of word tokens, the multimodal reasoning module takes over. The purpose of this module is to align the semantic concepts in the text with the physical entities represented by the object tokens and to perform the logical operations necessary to formulate an accurate answer. This is achieved through a deep, multi-layered transformer architecture that employs extensive cross-attention mechanisms [27]. In the cross-attention layers, the words in the query act as queries, while the object tokens act as keys and values. This allows the language model to dynamically search the physical workspace for the components mentioned in the operator's question. For instance, if the operator asks whether the red safety latch is fully engaged, the text tokens corresponding to red safety latch will generate high attention scores corresponding to the specific object token representing that component. The transformer layers process this aligned information

hierarchically. Early layers handle simple visual grounding, linking nouns to objects and adjectives to visual properties like color or texture. Deeper layers handle complex spatial and relational reasoning, evaluating verbs and prepositions to understand the interactions between grounded components [28].

Table 2: Categorization and examples of visual queries utilized in the robotic assembly inspection context.

Query Category	Example Question	Required Reasoning Skill
Existence & Presence	Is the grounding wire present on the main terminal?	Object localization
Attribute Verification	What is the color of the connector on the left?	Visual feature extraction
Spatial Relationship	Is the cooling fan mounted above the heat sink?	Geometric understanding
State & Condition	Are all the retaining screws fully tightened?	Fine-grained state analysis

The reasoning module is specifically designed to handle the rigorous demands of industrial inspection, which frequently involve counting and state verification. Counting identical components, such as verifying that exactly six bolts are installed on a flange, is a notorious failure point for traditional visual question answering systems. Because standard systems use continuous feature maps, they struggle to individuate identical adjacent objects, often undercounting or overcounting based on texture density. In contrast, our object-centric framework represents each bolt as a distinct token. The reasoning module simply has to count the number of tokens that align with the semantic concept of the bolt, transforming a complex visual task into a robust set-matching operation. Finally, the processed multimodal representations are fed into a generation head that decodes the features into a natural language response. To enhance the utility of the system for human operators, the generation module is trained not only to provide binary answers (such as yes or no) but also to generate descriptive text outlining anomalies. If a component is missing or incorrectly installed, the system generates a precise description of the error and its location within the assembly workspace. This comprehensive diagnostic capability significantly accelerates the troubleshooting process and minimizes assembly line downtime.

4. Results and Discussion

4.1 Experimental Setup and Datasets

Validating the efficacy of the proposed object-centric visual question answering framework requires a specialized dataset that accurately reflects the complexity of modern robotic assembly operations. Because large-scale, publicly available datasets for industrial visual question answering are extremely scarce, we constructed a comprehensive experimental environment and generated a proprietary dataset. The physical setup consisted of a six-axis industrial manipulator operating over a standardized assembly workbench. The workbench was populated with a variety of mechanical assemblies, including complex gearboxes, electronic wiring harnesses, and structural framing components. We deliberately introduced high variability in the visual conditions, utilizing programmable LED arrays to simulate different factory lighting environments, including harsh directional lighting, diffuse ambient light, and variable color temperatures. The dataset generation involved capturing thousands of high-resolution images of both flawless assemblies and intentionally induced defects. Defects included missing fasteners, misaligned gears, incorrect wire routing, and damaged surface textures. For each captured state, a team of quality assurance engineers and domain experts authored a set of natural language questions and corresponding answers [29]. The questions were designed to cover a broad spectrum of cognitive tasks, ranging from simple component identification to complex multi-step logical reasoning. To ensure the robustness of the evaluation, the dataset was carefully balanced to prevent linguistic biases, ensuring that the model could not guess the correct answer based solely on the phrasing of the question without analyzing the visual evidence.

Table 3: Statistical distribution of the custom industrial dataset used for training and evaluation.

Data Characteristic	Training Split	Validation Split	Testing Split
Total Images Captured	12,500	2,500	3,000
Total Question-Answer Pairs	62,500	12,500	15,000
Average Objects per Scene	24	25	26
Unique Vocabulary Tokens	4,200	4,200	4,200

The model was trained using a high-performance computing cluster equipped with multiple graphics processing units. The training process utilized a combination of cross-entropy loss for the answer generation and auxiliary contrastive losses to enforce the alignment between text concepts and object tokens. We employed rigorous data augmentation techniques, including geometric transformations, color jittering, and synthetic noise injection, to maximize the generalization capabilities of the network. The evaluation focused primarily on accuracy metrics, measuring the percentage of exactly correct answers generated by the model across different categories of questions [30]. Furthermore, we conducted extensive qualitative analyses of the internal attention maps to verify that the model was indeed grounding its answers in the correct physical components.

4.2 Performance Analysis and Comparison

The experimental results provide compelling evidence for the superiority of the object-centric approach in industrial inspection scenarios. We compared our proposed framework against several prominent baselines, including standard grid-based vision transformers and global convolutional architectures previously adapted for visual question answering. Across all evaluated metrics, the object-centric model demonstrated a significant performance advantage, particularly on complex queries requiring spatial reasoning and fine-grained state verification.

Table 4: Comparative accuracy performance of the proposed model against standard baseline architectures.

Model Architecture	Presence Accuracy	Attribute Accuracy	Relational Accuracy	Overall Accuracy
Baseline Convolutional	78.4%	72.1%	54.3%	68.2%
Grid-based Transformer	85.2%	81.5%	66.8%	77.8%
Object-Centric (Proposed)	96.1%	94.3%	89.7%	93.3%

The most striking performance differential was observed in relational accuracy, which encompasses questions about spatial layout and the physical interactions between multiple components. The grid-based transformer struggled heavily with these tasks, frequently confusing adjacent components due to the arbitrary division of image patches. When asked whether a specific wire passed behind or in front of a bracket, the grid-based model often failed to maintain the structural continuity of the wire across multiple patches, resulting in random guessing [31]. Conversely, our object-centric framework, which explicitly models the wire and the bracket as distinct, interacting entities, handled these queries with high reliability. The attention maps generated by our model during these relational queries clearly showed the network shifting its computational focus back and forth between the specific object tokens involved, demonstrating a profound structural understanding of the scene. We also conducted detailed failure mode analysis to understand the limitations of the current implementation. One of the primary challenges encountered involved highly reflective metallic surfaces, such as polished aluminum housings or chromed fasteners. In certain lighting configurations, the specular highlights on these

surfaces obscured the physical boundaries of the components, occasionally causing the object discovery module to merge multiple adjacent fasteners into a single object token or, conversely, to split a single large component into multiple fragmented tokens [32]. When the foundational object representation was compromised in this manner, the downstream reasoning module inevitably failed to provide the correct answer. This highlights the critical dependency of the entire framework on the robustness of the initial object extraction phase. Another notable edge case involved extremely dense occlusions, where a critical component was only partially visible through a small gap in the surrounding assembly. While the object-centric model performed significantly better than global models in these scenarios, it occasionally lacked the contextual confidence to assert the state of heavily occluded parts. To address these issues, future iterations of the system could integrate multi-view aggregation, allowing the robot to actively move the camera to resolve ambiguities and construct a comprehensive 3D object representation before answering the query [33]. Despite these edge cases, the overall performance of the object-centric framework represents a massive leap forward in automated inspection capabilities, offering unparalleled interpretability, flexibility, and diagnostic power.

5. Conclusion

This research has presented a highly advanced Object-Centric Visual Question Answering framework explicitly tailored for the rigorous demands of robotic assembly inspection cells. By shifting the paradigm from global feature extraction to compositional, object-level reasoning, the proposed system overcomes the fundamental limitations that have historically hindered the application of natural language interfaces in complex industrial environments. The decomposition of dense, cluttered workspaces into discrete semantic tokens enables precise cross-modal grounding, allowing the system to comprehend and accurately respond to highly specific, multi-step queries regarding component presence, spatial relationships, and assembly states. Extensive empirical evaluations demonstrated that the object-centric architecture significantly outperforms traditional grid-based and global convolutional baselines, particularly in tasks requiring complex relational logic and robust counting capabilities. The modularity of the framework not only enhances performance but also provides critical transparency, allowing quality assurance engineers to trace the reasoning process of the artificial intelligence down to specific physical components. This interpretability is vital for establishing trust and facilitating seamless human-robot collaboration on the factory floor. While the current results are highly promising, several avenues remain for future exploration. The reliance on purely spatial features can be augmented by incorporating temporal dynamics, allowing the system to answer questions about the sequence of assembly operations over time. Furthermore, addressing the challenges posed by severe occlusions and specular reflections will require the integration of active vision strategies, where the robot autonomously plans camera trajectories to resolve uncertainty. Additionally, advancing techniques for sim-to-real knowledge transfer will be crucial for scaling this technology, enabling systems to train on synthetic CAD data and deploy directly to physical production lines without exhaustive manual data collection [34]. Ultimately, the integration of object-centric visual reasoning represents a critical step toward fully cognitive, highly adaptable manufacturing systems.

References

- Qu, W., Shao, Y., Meng, L., Huang, X., & Xiao, L. (2024). A conditional denoising diffusion probabilistic model for point cloud upsampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 20786-20795).
- Xiao, Q., & Zhou, K. (2026). MACD: Model-Aware Contrastive Decoding via Counterfactual Data. *arXiv preprint arXiv:2602.01740*.
- Qu, W., Wang, J., Gong, Y., Huang, X., & Xiao, L. (2025). An end-to-end robust point cloud semantic segmentation network with single-step conditional diffusion models. In *Proceedings of the Computer Vision and Pattern Recognition Conference* (pp. 27325-27335).
- Li, Q., Liu, Z., Luo, W., Luo, T., & Hou, C. (2026). Correcting Visual Blur Induced by Attention Distraction to Reduce Hallucinations: Algorithm and Theory. *arXiv preprint arXiv:2605.24602*.

- Xu, W., Huang, H., Gong, Y., Yu, L., Wu, Q., & Zhang, J. (2025). Hierarchical multi-prototype discrimination: Boosting support-query matching for few-shot segmentation. *IEEE Transactions on Multimedia*.
- Yang, Y., Ma, X., Li, C., Zheng, Z., Zhang, Q., Huang, G., ... & Zhao, Q. (2021). Believe what you see: Implicit constraint approach for offline multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34, 10299-10312.
- Qi, Z., Yuan, Y., Ruan, X., Wang, S., Zhang, W., & Huang, Q. (2024). Collaborative debias strategy for temporal sentence grounding in video. *IEEE Transactions on Circuits and Systems for Video Technology*, 34(11), 10972-10986.
- Qi, Z., Wang, S., Zhang, W., & Huang, Q. (2024). Uncertainty-boosted robust video activity anticipation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12), 7775-7792.
- Zhang, Jiaquan, et al. "Learning global hypothesis space for enhancing synergistic reasoning chain." *arXiv preprint arXiv:2602.09794* (2026).
- Li, Z., Chen, Z., Qu, F., Wang, M., Zhao, Y., Zhang, K., & Peng, Y. (2024). Point resampling and ray transformation aid to editable NeRF models. *arXiv preprint arXiv:2405.07306*.
- Liu, Y., Liu, H., Wang, H., Meng, F., & Liu, M. (2023). BCAN: Bidirectional correct attention network for cross-modal retrieval. *IEEE Transactions on Neural Networks and Learning Systems*, 35(10), 14247-14258.
- Peng, Q., Xue, H., Wang, P., & Chen, C. (2026, March). Lifelong Domain Adaptive 3D Human Pose Estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 40, No. 10, pp. 8358-8366).
- Qu, D., & Ma, Y. (2025). Magnet-bn: markov-guided Bayesian neural networks for calibrated long-horizon sequence forecasting and community tracking. *Mathematics*, 13(17), 2740.
- Zhao, Haoyu, et al. "LSTD: Long Short-Term Temporal Diffusion for Video Generation." *IEEE Transactions on Multimedia* (2026).
- Chu, Y. B., Li, J., Jia, P., Cui, J., Zhang, R., Kang, X., Lv, M., & Zhang, S. (2021). Irf1- and Egr1-activated transcription plays a key role in macrophage polarization: A multiomics sequencing study with partial validation. *International immunopharmacology*, 99, 108072. <https://doi.org/10.1016/j.intimp.2021.108072>
- Peng, Q., Zheng, C., & Chen, C. (2023). Source-free domain adaptive human pose estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 4826-4836).
- Huang, H., Zhang, J., Zhang, J., Xu, J., & Wu, Q. (2020). Low-rank pairwise alignment bilinear network for few-shot fine-grained image classification. *IEEE Transactions on Multimedia*, 23, 1666-1680.
- Guo, Zixin, Kai Zhao, and Luyan Zhang. "InstanceRSR: Real-World Super-Resolution via Instance-Aware Representation Alignment." *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2026.
- Li, Q., Luo, T., & Liao, J. (2025). Theory-inspired deep multi-view multi-label learning with incomplete views and noisy labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 20706-20715).
- Wang, M., Li, J., Li, Z., Luo, C., Chen, B., Xia, S. T., & Wang, Z. (2023, October). Unsupervised anomaly detection with local-sensitive VQVAE and global-sensitive transformers. In *2023 IEEE International Conference on Image Processing (ICIP)* (pp. 1080-1084). IEEE.
- Li, Z., Bai, X., Lu, J., Shen, P., Lam, E. Y., & Peng, Y. (2026). EventTracer: Fast Path Tracing-based Event Stream Rendering. *IEEE Transactions on Visualization and Computer Graphics*.
- Ziyue Wang, Sheng Jin, Zhongrong Zuo, Jiawei Wu, Han Qiu, Qi She, Hao Zhang, and Xudong Jiang. Video-ctr: Reinforcing video reasoning via key token attribution. *arXiv preprint arXiv:2601.19686*, 2026.
- Zhao, Y., Li, Z., Guo, X., & Lu, Y. (2022). Alignment-guided temporal attention for video action recognition. *Advances in Neural Information Processing Systems*, 35, 13627-13639.

- Zhou, D. (2026). AI-Driven Hybrid SAST–DAST–SCA–IAST Framework for Risk-Based Vulnerability Prioritization in Microservice Architectures.
- Cong, P., Xu, Y., Ren, Y., Zhang, J., Xu, L., Wang, J., ... & Ma, Y. (2023, June). Weakly supervised 3d multi-person pose estimation for large-scale scenes based on monocular camera and single lidar. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 37, No. 1, pp. 461-469).
- Song, S., Tang, Y., & Qin, R. (2025). Synthetic Data Matters: Re-training with Geo-typical Synthetic Labels for Building Detection. *IEEE Transactions on Geoscience and Remote Sensing*.
- Qu, D., Ma, Y., & Zhang, S. (2025, November). OAMF: Optics-Accelerated Multimodal Learning with Markov Temporal Priors and Fourier Regularization. In *2025 4th International Conference on Image Processing, Computer Vision and Machine Learning (ICICML)* (pp. 600-605). IEEE.
- Tu, P., Huang, Y., Zheng, F., He, Z., Cao, L., & Shao, L. (2022, June). Guidedmix-net: Semi-supervised semantic segmentation by using labeled images as reference. In *Proceedings of the AAAI conference on artificial intelligence* (Vol. 36, No. 2, pp. 2379-2387).
- Sang, Y. (2025, October). AutoCrit: A Meta-Reasoning Framework for Self-Critique and Iterative Error Correction in LLM Chains-of-Thought. In *2025 6th International Conference on Machine Learning and Computer Application (ICMLCA)* (pp. 1177-1180). IEEE.
- Huang, H., Zhang, C., Hu, Q., & Zhu, P. (2016, August). Multi-view representative and informative induced active learning. In *Pacific Rim international conference on artificial intelligence* (pp. 139-151). Cham: Springer International Publishing.
- Lv, Q., Kong, W., Li, H., Zeng, J., Qiu, Z., Qu, D., ... & Pang, J. (2025). F1: A vision-language-action model bridging understanding and generation to actions. *arXiv preprint arXiv:2509.06951*.
- Zhang, J., Zhang, J., Song, Z., Shi, Z., Zhao, C., Shi, Y., ... & Wang, J. (2024). Hoi-m³: Capture multiple humans and objects interaction within contextual environment. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (pp. 516-526).
- Ma, Y., & Qu, D. (2025, August). GEL-FMO: Gauge-Equivariant Liquid Fourier-Markov Operators for Uncertainty-Certified Multimodal Reasoning. In *2025 5th International Conference on Advanced Algorithms and Neural Networks (AANN)* (pp. 604-607). IEEE.
- Agrawal, H., Huang, X., Kim, J. H., Li, Y., Ma, Y., Ning, T., & Tao, Y. (2022). Automated image processing and insight presentation. Patent / invention record. Publication date: 2022-04-05.