

Real Deep Research for AI, Robotics and Beyond

Xueyan Zou^{*1}, Jianglong Ye^{*1}, Hao Zhang², Xiaoyu Xiang³, Mingyu Ding⁵, Zhaojing Yang¹,
Yong Jae Lee⁴, Zhuowen Tu¹, Sifei Liu², Xiaolong Wang¹

¹UC San Diego, ²NVIDIA, ³META, ⁴UW-Madison, ⁵UNC,
<https://realdeepresearch.github.io/>

Abstract

With the rapid growth of research in modern AI and robotics—now producing over 10,000 papers annually—it has become increasingly difficult for researchers to stay up to date. Fast-evolving trends, *the rise of interdisciplinary work*, and *the need to explore domains beyond one’s expertise* all contribute to this challenge. To address these issues, we propose a generalizable pipeline capable of systematically analyzing any research area: identifying emerging trends, uncovering cross-domain opportunities, and offering concrete starting points for new inquiry. In this work, we present **Real Deep Research** (RDR)—a comprehensive framework applied to the domains of AI and robotics, with a particular focus on foundation models and robotics advancements. We also briefly extend our analysis to other areas of science. The main paper details the construction of the RDR pipeline, while the appendix provides extensive results across each analyzed topic. We hope this work could shed lights on researchers who works in the filed of AI and beyond.




Figure 1 | Real Deep Research enables: (1) generating surveys for specific research focuses or perspectives; (2) analyzing topic trends over time; (3) mapping interdisciplinary research landscapes; and (4) retrieving high-impact papers relevant to a given topic. (Each dot represents a paper, and each sphere denotes a topic cluster. The cluster keywords and trend information are automatically generated by **RDR**)

* Indicate core contribution.

1. Introduction

The fields of AI and robotics have experienced exponential growth in recent years, while researchers continue to face the constraint of limited time and attention. This work is motivated by the authors’ need to efficiently survey research areas, stay up to date with rapidly evolving trends, identify promising interdisciplinary opportunities, and familiarize themselves with the latest developments on a given topic.

In response to this need, we develop a systematic analysis tool designed to help users quickly navigate and adapt to any research area or topic. We begin by applying our approach to the fields of AI and robotics, conducting an in-depth analysis with a focus on foundation models and robotics research. To broaden our exploration and uncover emerging areas of interest, we also extend our analysis to natural sciences and formal sciences, offering a glimpse into recent developments beyond our core domains.

Although our intentions are well-founded, it is important to acknowledge existing efforts in this space. On the one hand, there are high-quality survey papers written by domain experts [12, 83]; on the other hand, a few recent works have explored automated research pipelines [3, 115]. Expert-written surveys offer depth and accuracy, but require significant manual effort and cannot easily adapt to the fast-paced evolution of research. Meanwhile, current automated approaches often lack domain-specific knowledge and expert insight, limiting their usefulness and relevance to researchers. Our work aims to bridge this gap by combining systematic automation with meaningful, expert-informed analysis.

Therefore, in addition to building an effective pipeline for Real Deep Research, our goal is to make the tool robust and insightful enough to support top-tier researchers in tracking emerging trends and engaging with unfamiliar research areas. A key focus of our work is interdisciplinary exploration—helping researchers identify underexplored intersections between fields that present promising opportunities for cross-domain collaboration.

As shown in Fig. 1, the visualization displays individual papers, clustered research topics, and their corresponding trends. At a glance, it becomes clear that areas such as teleoperation, dexterous manipulation, and open-source robotics are emerging as promising directions, whereas traditional reinforcement learning appears to be declining in momentum. As researchers in the robotics field, we find that these trend insights align well with our domain knowledge and provide valuable guidance for identifying impactful research opportunities. We summarize the key contributions of this paper as follows:

1. We propose the **Real Deep Research (RDR)** pipeline, a systematic framework for exploring and analyzing any research area in depth.
2. Leveraging domain expertise, we deliver high-quality survey outputs in the fields of AI and robotics, providing valuable insights for researchers and practitioners.
3. We quantitatively evaluate the RDR pipeline and demonstrate its advantages over existing commercial large language model tools within the targeted research domains.

2. Related Work

Surveys of Foundation Models. In recent years, a number of survey studies have systematically reviewed foundation models across different domains [12, 55, 83, 141, 234, 298], including natural language processing [20, 288], computer vision [141, 269], graph learning [234], and robotics [153, 242, 246, 267]. However, these surveys require extensive manual effort and become outdated quickly due to the rapid progress of foundation models. Unlike such static surveys, our goal is to design a framework that can automatically analyze thousands of papers and provide an always up to date understanding of different research areas.

LLMs in Scientific Research. Large language models (LLMs) have been applied across various stages of scientific research [145, 161, 190, 217], including idea generation [6, 222], coding [155, 245], paper

reviewing [124, 145], and predicting experimental results [149, 156]. Among these stages, literature analysis plays a central role, involving tasks such as paper retrieval, clustering, and topic trend analysis. However, traditional literature search tools such as Google Scholar rely mainly on lexical matching and struggle with tasks that require deeper semantic reasoning. This has motivated researchers to leverage LLMs for literature analysis [3, 76, 115, 196]. For example, SciLitLLM [115] employs supervised learning to build a specialized LLM for scientific literature understanding; PaSa [76] uses reinforcement learning with synthetic data to train an LLM agent that can answer complex scholarly queries. Unlike prior work that focuses mainly on research question answering, our approach targets a broader and systematic understanding of entire research areas. We highlight not only semantic reasoning over large collections of papers but also automatic analysis of research trends, offering researchers a transparent and evidence-based view of the literature.

Knowledge Organization and Discovery. It has been shown that LLMs are capable of clustering documents [219, 281] and uncovering latent topics [114, 177]. For example, Knowledge Navigator [97] combines LLMs with clustering techniques to organize and structure documents for scientific literature search; SciTopic [114] enhances LLMs in identifying topic structures by refining document embeddings. Beyond knowledge organization, recent research [62, 105, 106] also studies the trend of high-impact research topics. Our work introduces a novel approach by leveraging the reasoning capabilities of LLMs and the embedding representations of foundation models, which leads to more accurate and semantic knowledge organization. Built on this knowledge structure, our framework enables analysis of past and future research trends and supports inspection of connections between topics, providing valuable insights into scientific directions.

3. Method

In the Methods section, we focus specifically on the domains of foundation models and robotics to provide a comprehensive overview of how we conduct Real Deep Research using expert knowledge. As illustrated in Fig. 2, the embedding-based analysis pipeline consists of four main components: (1) Data Preparation (Sec. 3.1), (2) Content Reasoning (Sec. 3.2), (3) Content Projection (also in Sec. 3.2), and (4) Embedding Analysis (Sec. 3.4). This pipeline is powered by a suite of large language and multimodal models (LLMs/LMMs) for content extraction and reasoning, and is designed to be generalizable for the automated analysis of other research domains in the future. The following sections introduce each component in detail.

3.1. Data Preparation.

Selection. To systematically investigate the integration of foundation models and robotics at scale, we focus on emerging trends and research priorities in both academia and industry. To capture the latest developments, we review recent publications from leading conferences in computer vision, robotics, and machine learning. Specifically, we collect papers via web crawling from top conference venues (CVPR, ECCV, ICCV, CoRL, RSS, ICRA, NeurIPS, etc.) as well as from industry research platforms (Nvidia, Meta, and OpenAI, etc.). This curated corpus comprehensively overviews the research contents in foundation models and robotics, highlighting key technical advancements, existing challenges, and future research directions. Specifically, we collect paper titles, authors, abstracts, and PDF links directly from conference and company websites. Then, we apply an area filtering process on paper titles and abstracts using an efficient LLM with a predefined set of criteria to ensure relevance to this study.

Area Filtering. We define the collected paper set as $\mathbf{P}$, while it generally fall under the broad area of vision, language, machine learning, and robotics, it is not guaranteed that each paper directly aligns with the specific focus of our work, such as foundation models ($\mathbf{D}_f$) and robotics ($\mathbf{D}_r$). To address this, we




Figure 2 | Pipeline of the proposed method on filtering and projecting thousands of papers to the embedding space for future analysis.

introduce *Area Filtering*—a step that leverages an efficient LLM with curated prompts—to identify papers relevant to our research scope. To ensure a correct filtering, we first define the scope of foundation models and robotics, clarifying technical boundaries between domains. Below are the prompts that we designed for our research focus:

Foundation Model Definition: "Research involving deep learning models (especially transformer-based) trained on large amounts of data and capable of fitting generalized factual realities. These models typically serve as versatile backbones for a variety of downstream tasks across multiple domains."
Key Indicators:
- Large Multimodal Models (LMM)
- Large Language Models (LLM) ...

Robotics Definition: "Research involving hardware systems equipped with input sensors and mechanical kinematics capable of producing joint movements. These systems are controlled by learning-based algorithms that facilitate automatic or robust mappings from sensory inputs to actuator outputs."
Key Indicators:
- Reinforcement Learning in robotic contexts
- Imitation Learning for physical systems ...

After filtering using an efficient LLM, the resulting set of papers ($\mathbf{P}'$) belongs to either the foundation model domain ($\mathbf{D}_f$), the robotics domain ($\mathbf{D}_r$), or both. Formally we write $\mathbf{P}' = \{p \mid p \in \mathbf{D}_f \cup \mathbf{D}_r\}$.

3.2. Content Reasoning.

Given the filtered papers $\mathbf{P}'$ in the domains of foundation models and robotics, an in-depth analysis is required to narrow the position of each paper. Guided by domain experts in foundation models and robotics, we define perspectives that align with established domain structures, emerging trends, and evolving knowledge. Beyond predefined perspectives, our pipeline supports future user-defined perspectives, allowing adaptation to new research questions. In the following paragraphs, we will depict the general structure, trends, and knowledge of the foundation model and robotics in preparation for analyzing the research works under $\mathbf{P}'$.

Foundation Model. A foundation model's development are systematically analyzed through five fundamental perspectives in this work: Input (**I**), Modeling (**M**), Output (**O**), Objective (**W**), and Learning Recipe (**R**). We have shown some main perspectives examples in Fig. 3. This structured representation facilitates a comprehensive analysis of the foundation model. Below is the formal writing for the procedure:

$$\mathcal{D}_f^{P'} = \bigcup_{p \in \mathbf{P}'} F(p), \quad F(p) = \text{LLM}(p \mid \mathbf{I}, \mathbf{M}, \mathbf{O}, \mathbf{W}, \mathbf{R}),$$

where LLM represents the large multimodal model, and $\mathcal{D}_f^{P'}$ denotes the perspective projection of the given papers in $\mathbf{P}'$, focused on foundation model research. In the following paragraphs, we provide a




Figure 3 | Perspective analysis of foundation model research, which primarily includes (1) Input, (2) Modeling, (3) Output, etc., shown in the figure.

formal definition of each perspective:

Input (I). The input processing for a foundation model generally involves raw data and a tokenization procedure. Standard input sources include images, videos, audio, LiDAR, etc., with tokenization performed through transformations and neural networks.

Modeling (M). With input settled for a foundation model, the modeling part is responsible for extracting critical knowledge from the input, reasoning, and decoding to the output space. It is the critical procedure to transfer input knowledge to output.

Output (O). The task determines the decoding space according to the input and modeling, this is the final step to decode the latent representation to the output used for loss computation or the final interaction.

Objective (W). To fit a foundation model with the corresponding input, and output, the given model architecture is constrained by the learning objective, this fits the model distribution in alignment with the transformation given the task(s).

Recipe (R). The recipe is used as the cookbook on how to tune the model weight with input, output, and objective. It controls the training stage, convergence speed, and updated parameters.

Robotics. For research work in robotics, the core perspective shifts to emphasize hardware and interaction within real-world environments. We define five key perspectives to map each paper within the broader landscape of robotic applications: Input Sensor (**S**), Physical Body (**B**), Joint Output (**J**), Action Space (**A**), and Environment (**E**). An example of core perspectives is illustrated in Fig. 4. These perspectives collectively define how robots perceive, act, and interact within the physical world. The procedure could be formally written as:

$$\mathcal{D}_r^{P'} = \bigcup_{p \in \mathbf{P}'} F(p), \quad F(p) = \text{LMM}(p \mid \mathbf{S}, \mathbf{B}, \mathbf{J}, \mathbf{A}, \mathbf{E}),$$

where $\mathcal{D}_r^{P'}$ represents the perspective projection of the given papers in $\mathbf{P}'$, providing a structured framework for analyzing robotics research. We show the concrete definition in the following:

Input Sensor (S). Input sensors are hardware devices that measure physical quantities or environmental conditions and convert them into digital signals that can be processed by the robot's control system. They serve as the robot's interface with its environment.

Physical Body (P). A physical body in robotics refers to the mechanical structure and architecture that enables physical interaction with the environment. This physical manifestation determines how motor commands translate into real-world forces, movements, and environmental manipulations.

Action Space (A). The action space is the set of all permissible actions a robot can select in a given context, ranging from low-level joint commands to high-level behaviors (e.g., "walk" or "grasp"). Each




Figure 4 | Perspective analysis of robotics research, which primarily includes (1) Input, (2) Modeling, (3) Output, etc., shown in the figure.

chosen action is ultimately executed as a joint output, bridging decision-making to physical movement.

Joint Output (J). Joint output refers to the physical movement or configuration of a robot’s joints resulting from executed motor commands. It translates control signals (e.g., torque or velocity) into mechanical motion, allowing the robot to directly interact with and manipulate its environment.

Environment (E). The environment encompasses the physical space where a robot operates, characterized by its spatial layout, structural features, and contextual elements (e.g., furniture, tools, obstacles) that shape the task-specific challenges and opportunities the robot encounters.

Given the predefined perspective definition, we use the following prompt to extract each perspective from the paper:

```
Can you analyze the paper contents according to the following perspectives: (1)
Definition 1, (2) Definition 2, (3) Definition 3, ...
After analysis, please identify each of the perspectives in the paper, and return
the answer in the following format: {"perspective 1": plain text, "perspective 2":
plain text, "perspective 3": plain text, ...}
```

3.3. Content Projection.

Given the extracted contents from research papers guided by our defined perspective, we aim to project natural language descriptions into an informative latent space. This projection enables large-scale analysis of current research in foundation models and robotics while revealing potential future research directions. Motivated by recent advancements in large language model-based embedding models, we **employ a pre-trained embedding foundation model G** to project $\mathcal{D}_r^{p'}$ (processed robotics papers’ content) and $\mathcal{D}_f^{p'}$ (processed foundation model papers’ content) from natural language space into a more abstract embedding space. The embedding model maps text into a high-dimensional vector space where semantically similar concepts occupy proximate regions.

We formally define this projection procedure as follows: For any text snippet $x \in \mathcal{D}$, its embedding is computed as: $v_x = G(x) \in \mathbb{R}^d$. Our core assumption is that by projecting paper contents through this perspective-aware embedding process and analyzing them in the high-dimensional manifold, we can uncover meaningful patterns, research trends, and potential gaps in the literature through systematic visualization and clustering analysis.

3.4. Embedding Analysis.

The goal of embedding analysis is to structure the understanding of previously extracted embeddings. The pipeline for embedding analysis contains three components: (1) Clustering on the extracted embeddings

and analyzing the main concept from each cluster. (2) Structured the concept for each cluster to formulate an informative table. (3) Given the structured understanding, we trace back to the reference papers.

Clustering for Embeddings. We first embed every paper to obtain its vector representation V and partition the corpus into k clusters. From each cluster, we then draw a random sample of 50 papers and feed their text to a reasoning-based model with the prompt:

Can you summarize the following contents into three distinct keywords: Here is one example output: "keyphrase1, keyphrase2, keyphrase3". The output should be short and precise, with a single output for all papers.

The model returns three compact key phrases that capture the cluster’s core theme, giving every paper both a cluster label and an interpretable set of keywords for subsequent analysis.

Structuring for Thoughts. With clustered embeddings and their associated topic keywords in place, the next step is to generate a structured survey for the given research area. To accomplish this, we leverage the o3 language model, using the clustered keywords as prompts to guide the formulation of the final survey content. Incorporating the clustering results into the prompt ensures that the generated text remains grounded in the actual structure of the research landscape, enhancing both coherence and relevance. We use the following prompt to produce the final output:

Those are summarized keywords for a number of science papers clustered by abstract contents, however they are ambiguous, contents may overlap between clusters, can you summarize the information in a more structured way for audience with the following criteria: ...

4. Analysis

In this section, we conduct a comprehensive qualitative analysis of the conclusions drawn in this work from the following perspectives: (1) Embedding analysis for general research areas. (2) Embedding analysis within specific perspective. (3) Trend analysis of research focus over time. (4) Knowledge graph exploration across different research areas. (5) Retrieval examples based on embeddings. This pipeline will enable a researcher to dive into any research area, identify what to explore, and determine the specific papers to focus on.

Embedding Analysis - General. The output of the embedding analysis is a comprehensive survey tailored to the featured research domain. This survey is organized into major categories and sub-categories, each detailing the specific topics covered. Rather than generating the survey content via LLM, we leverage the clustering results from the embedding analysis to guide its structure and scope. Additionally, for each sub-topic, we include the most relevant citations to provide readers with direct references for further exploration. We have provide the full survey for Foundation Model, Robotics, Computer Vision, Natural language processing, and machine learning in Appendix.


Cat.	Sub-category	What is covered	Typical examples	Cluster
1. Perception & Mapping [71, 168, 185, 277, 294]				
	1.1 Multimodal sensor fusion [14, 123, 224, 256, 303]	Fuse heterogeneous sensors for richer scene understanding	LiDAR–Camera Fusion; Radar–Camera Fusion; V2X Cooperative Perception ...	0,6,7,8,14,16
	1.2 3D reconstruct/occupancy [85, 133, 140, 175, 294]	Build dense or sparse geometric maps for localisation	3-D SLAM & Reconstruction; 3-D Occupancy; Efficient 3-D Representation	0,8,16
	1.3 BEV / top-view mapping [19, 93, 120, 256, 286]	Bird’s-eye or top-down representations for planning	BEV Perception; V2X Collaborative Perception	0,14,16
... ..				

Embedding Analysis - Perspective. After establishing a clear overview of the domain, we analyze

it through a targeted *perspective* to expose structure and problem formulations. As introduced in the Methods section, our perspective analysis uses embedding-based clustering to organize works along a chosen axis. In this study we focus on foundation models and robotics, examining how each community formulates its problems. Below we illustrate the robotics case from the viewpoint of *action space*. This perspective-guided embedding analysis yields a deeper understanding of the domain and a high-level map of how researchers approach and solve its problems. We also provide the full perspective for foundation model and robotics in Appendix.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Continuous Low-Level Actuation				
	1.1 Joint-space commands	Direct numerical inputs to individual joints or actuators, bounded by hardware limits.	joint torques/positions/velocities; high-dimensional joint commands; bounded control inputs; finger-joint configs; parametrised joint trajectories	0, 4, 6, 10, 12, 14, 18
	1.2 Vehicle / body dynamics commands	Low-level controls that change a mobile base, ground-vehicle or aerial body state.	steering angle; throttle / acceleration; braking; linear & angular velocity; body-rate thrust; speed/direction for locomotion; lane-keeping	0, 1, 7, 10, 12, 13, 15
... ..				

Trend Analysis. Once we understand each domain and its key sub-perspectives, the next step is to assess topic momentum. Our trend analysis highlights which areas are accelerating and which have been thoroughly explored in recent years, giving a practical starting point when entering a new field. In robotics (see figure), the trajectories suggest that teleoperation, dexterous manipulation, and low-cost open-source robotics are currently rising, while traditional reinforcement learning and skill-based manipulation appear comparatively mature or show slowing activity. This will guide the researchers to smoothly enter a new field. We provide the full trend analysis in Appendix. A for Computer Vision, NLP, Robotics, and Machine Learning.



Knowledge Graph. Beyond identifying trending topics within individual research areas, an equally important direction for discovery lies in uncovering cross-domain themes—topics that span multiple fields and have the potential to catalyze interdisciplinary breakthroughs. In this work, we analyze the intersections among four major domains: computer vision, natural language processing, machine learning, and robotics. By examining these intersections, we aim to highlight not only where collaboration already exists but also where it could be further cultivated.

As shown in Figure below, the left side of the figure presents a Cross-Domain Topology Graph, where each color corresponds to a specific research domain. Each node (represented as a sphere) signifies a distinct topic cluster derived from our embedding-based analysis, and edges between nodes indicate semantic or topical relationships—especially those that cross domain boundaries. Nodes located at the periphery, with few or no connecting edges, represent domain-specific topics that remain largely isolated

from other fields. In contrast, the densely connected regions at the center of the graph reflect genuinely cross-domain topics, where ideas, methods, or applications from multiple fields converge. This view empowers researchers to discover promising frontiers for collaboration, encourages rethinking isolated problems through new lenses, and supports a more forward-looking approach to scientific inquiry.

Retrieval Examples. Once a target research topic is identified, the next step is to pinpoint concrete entry points. We do this by leveraging the conference-level embeddings inferred earlier to run semantic searches and retrieve the most relevant literature. For example, after surveying robotics, we focus on dexterous manipulation and query the embedding index to surface closely related papers across venues. As shown in the table below, the returned papers align tightly with the query and exhibit meaningful community impact, as reflected by their venues, years, and citation counts.

5. Experiment

Dataset. We curate our dataset from a collection of publicly available, high-impact conference venues, focusing on those central to the fields of artificial intelligence and robotics. As shown in Table 1, the dataset includes papers from venues such as CVPR, ECCV (Computer Vision), ICLR, NeurIPS (Machine Learning), ACL (Natural Language Processing), and CoRL and RSS (Robotics). This selection allows us to capture a broad yet targeted view of the AI and robotics research landscape from 2021 through 2025. To align with our focus on foundation models and

Table 1 | **Paper Distribution Analysis** across different venues, showing the total number of papers.

Model	Rank	General				Foundation Model			Robotics		
		CV	NLP	ML	Robotics	Input	Modeling	Output	Sensor	Body	Action
GPT5	4.80	10.00	17.39	45.45	71.43	44.44	10.00	21.05	22.73	34.78	69.57
GPT5-Thinking	2.75	82.61	59.09	47.83	66.67	55.00	90.91	50.00	88.46	42.86	32.00
GPT5-Research	4.00	42.11	50.00	72.73	63.64	21.05	35.00	50.00	0.00	40.91	52.63
Gemini	4.80	35.00	40.00	15.38	0.00	13.64	54.17	45.83	31.25	45.00	26.32
Gemini-Thinking	3.35	63.64	50.00	56.25	37.50	65.22	45.45	41.67	55.56	56.52	34.78
RDR (Ours)	1.30	58.33	89.47	73.68	77.78	88.46	60.00	94.74	91.30	84.21	89.47

Table 2 | Survey Quality Evaluation among RDR and commercial based methods. We evaluate the pairwise winning rate for each domain and perspective.

robotics, we apply an additional filtering step across all venues. Specifically, we identify and extract 4,424 papers related to foundation models and 1,186 papers focused on robotics, both from the year 2024 onward. This subset enables us to track the most recent developments in these fast-moving areas with higher resolution.

Implementation Details. We do not train any new networks in this work for generating the embedding or survey; instead, we rely on off-the-shelf models. For straightforward tasks—such as classifying research areas—we use the Doubao language model. For reasoning-intensive tasks and complex summarization, we employ the o3 model to achieve stronger performance. To extract text embeddings, we use nvidia/NV-Embed-v2.

Survey Quality. As demonstrated in Sec. 3, our analysis of a research area begins with a broad survey of the existing literature. The analysis pipeline we propose is designed to significantly reduce model hallucination and produce a comprehensive, high-quality survey for a given research direction.

To evaluate the accuracy and quality of the generated surveys, we conducted a user study involving experienced researchers with domain expertise in robotics and foundation models. As a baseline, we prompted a commercial large language model using the following instruction: *“Act as an expert research analyst. Your task is to create a structured map of the research landscape for a given academic or industrial field. The output must be a single, valid JSON object that categorizes the field into its primary research areas and specific sub-topics. For the research area ‘foundation model,’ can you summarize the input perspective with the following definition: The input processing for a foundation model generally involves raw data and a tokenization procedure ...”*

To assess the quality of the generated surveys, we adopted a pairwise comparison methodology rather than asking evaluators to select a single best output. For each comparison, domain experts were presented with two survey outputs and asked to determine which one demonstrated superior quality and accuracy. This evaluation setup helps reduce cognitive load and bias, making the assessment more reliable by avoiding the need for evaluators to recall or rank multiple outputs simultaneously. In total, we collected 8 evaluation entries, each with 80 pairwise comparisons. To quantify performance, we computed the winning rate of each method within its respective domain.

As shown in Tab. 2, our method, Real Deep Research (RDR), achieves the highest overall performance with an average rank of 1.30, outperforming all baselines. RDR leads in key domains such as NLP (89.47), robotics (77.78), and foundation model output (94.74), and also shows strong performance in robotics subfields like sensor (91.30) and action (89.47). While GPT5-Thinking slightly outperforms in CV (82.61) and foundation model modeling (90.91), RDR consistently ranks at or near the top across nearly all categories.

Embedding Quality. Because much of our analysis relies on high-quality embeddings, we evaluate their effectiveness using a simple linear probe trained on top of frozen representations—an approach that best reflects the intrinsic utility of the embeddings themselves. We follow the experimental protocol introduced in SciTopic [114], using the same unsupervised training and evaluation splits to ensure fair comparison. Unlike our method, SciTopic uses pseudo-labels during training, which introduces weak supervision; therefore, we gray out its entry in the results for clarity. As shown in Tab. 2, our method RDR achieves the best performance across both datasets, with an accuracy of 84.86 on AG News and 52.91 on 20 News Groups. RDR also leads in NMI (61.66 and 56.57) and ARI (65.24 and 39.96), outperforming all fully unsupervised baselines, and even surpassing the pseudo-supervised SciTopic model.

Model	AG News			20 News Groups		
	ACC(↑)	NMI(↑)	ARI(↑)	ACC(↑)	NMI(↑)	ARI(↑)
LDA	74.05	47.17	49.01	29.05	31.63	13.34
NMF	34.05	4.59	2.13	12.42	12.86	0.48
ProdLDA	80.93	56.51	60.91	37.42	45.67	23.89
DecTM	55.63	40.04	36.17	36.57	46.18	22.90
ETM	26.14	0.00	0.00	5.35	0.10	0.00
NSTM	26.14	0.01	0.00	16.92	17.02	2.34
TSCTM	79.63	53.91	55.89	40.60	44.06	15.71
ECRTM	78.69	54.05	54.88	25.70	31.00	12.26
Bertopic	35.93	12.88	7.03	29.78	28.57	11.58
FASTopic	83.48	59.10	62.48	51.65	56.32	39.49
SciTopic*	85.29	61.96	65.94	70.88	68.32	55.71
RDR (Ours)	84.86	61.66	65.24	52.91	56.57	39.96

Table 3 | Unsupervised Clustering performance. * indicate using more labels.

A. Appendix

This appendix presents the complete results of our Real Deep Research (RDR) analysis across a wide range of domains. We include detailed domain-level surveys (e.g., AI, robotics, computer vision, natural language processing), perspective-based breakdowns (e.g., input/output modeling in foundation models, sensor/action perspectives in robotics), and trend analyses to track the evolution of research focus over time. These results collectively offer a structured and insightful view of the research landscape, serving as a valuable reference for both new and experienced researchers.

Domain Survey	12
Foundation Model	12
Robotics	13
Computer Vision	14
Natural Language Processing	15
Machine Learning	16
Nature	17
Science	19
Perspective Survey	20
Foundation Model Perspectives	20
Robotic Perspectives	25
Trend Analysis	30
Computer Vision	30
Robotics	31
Natural Language Processing	32
Machine Learning	33

Cat.	Sub-category	What is covered	Typical examples	Cluster
1. Model modalities & representations [32, 112, 195, 214, 230, 238, 241, 255, 272, 305]				
	1.1 Vision-Language [30, 32, 225, 241, 300]	Foundation models that jointly process images/video and natural language.	Vision-Language Models; Vision-Language Robotic FMs	0, 3
	1.2 Multimodal (≥ 3) [29, 204, 247, 295, 296]	Architectures/objectives agnostic to the exact mix of modalities.	Multimodal Foundation Models; Multimodal LLMs; Multimodal Large Language Models	1, 4, 5, 6
	1.3 Open-vocabulary grounding [24, 108, 131, 157, 268]	Linking free-form text to modality-specific regions or anchors.	Open-Vocabulary Grounding	1
	1.4 3D/4D & video reps [25, 68, 94, 252, 299]	Learned neural representations for 3-D/4-D scenes and video.	3D & Multimodal Representation; Diffusion-based 3D/4D Generation; 3D & Video Synthesis	0, 7, 9
	1.5 Neural scene encoding [53, 116, 163, 244, 297]	Representations enabling view-consistent reconstruction.	Multi-view Consistent Reconstruction; Gaussian Splatting; NeRF Representations	9
2. Generative & diffusion techniques [22, 31, 48, 130, 136, 212, 215, 243, 283, 284]				
	2.1 Core diffusion modelling [118, 128, 138, 211, 283]	Diffusion processes used as the primary generative backbone.	Diffusion Generative Modeling	7
	2.2 Control & personalisation [147, 164, 165, 174, 206]	Steering diffusion outputs with prompts, adapters or user profiles.	Controllable Diffusion Personalization; Controllable Efficient Sampling	7, 11
	2.3 Robot policy via diffusion [191, 232, 253, 283, 290]	Using diffusion to learn control policies for robots or manipulators.	Diffusion-Based Policy Learning	3
	2.4 Editing & post-generation [15, 182, 189, 213, 223]	Applying diffusion to edit or refine existing content.	Diffusion-based Generative Editing	4
	2.5 Efficiency & distillation [126, 184, 186, 211, 273]	Speed-ups and compact student models for diffusion.	Diffusion Model Acceleration; Efficient Sampling & Distillation	8
3. Training & adaptation strategies [8, 37, 40, 73, 79, 87, 96, 110, 129, 289]				
	3.1 Self-/pre-training paradigms [1, 17, 41, 154, 218]	Large-scale unsupervised or weakly-supervised pre-training methods.	Elastic Self-Supervised Pre-training	6
	3.2 Prompt/adaptor learning [56, 110, 160, 197, 233]	Lightweight modulation of frozen backbones via prompts or adapters.	Prompt/Adapter Tuning; Prompt/Adapter Learning; Parameter-Efficient Prompt Tuning	0, 1, 5
	3.3 Param-efficient finetune [11, 72, 132, 159, 264]	LoRA/adapters that tune only a small slice of parameters.	Parameter-Efficient Fine-Tuning; Adapter-Efficient Fine-Tuning	10, 11
	3.4 Compression & inference efficiency [23, 50, 102, 169, 231]	Sparsity, low-rank factorisation and runtime acceleration.	Sparse/Low-Rank Model Compression; Efficient Transformer Inference	10
4. Safety, alignment & ethics [18, 65, 111, 142, 148, 151, 166, 240, 292, 304]				
	4.1 Safety alignment [89, 259, 282, 291, 304]	Aligning model behaviour with human or policy constraints.	LLM Safety Alignment; Alignment & Safety	2, 6
	4.2 Bias & harm mitigation [7, 99, 143, 176, 188]	Detecting and reducing social or representational bias.	Safety & Bias Mitigation	11
	4.3 Preference optimisation [18, 166, 236, 240, 279]	Fine-tuning with human preference or RLHF-style signals.	Preference-Optimized Fine-Tuning	2
5. Embodied interaction & robotics [47, 60, 91, 92, 95, 139, 180, 226, 228, 265]				
	5.1 Robotic foundation models [58, 91, 100, 170, 221]	General-purpose models for perception and control on robots.	Vision-Language Robotic Foundation Models	3
	5.2 Embodied agents [33, 80, 90, 209, 262]	Agents acting in simulated or real environments with multimodal inputs.	Embodied Vision-Language Agents	4
	5.3 Intended manipulation [57, 84, 95, 113, 287]	Grounding natural-language instructions into robot actions.	Multimodal Instruction-Guided Manipulation	3
6. Reasoning & agent systems [86, 90, 137, 181, 205, 207, 209, 210, 270, 274]				
	6.1 Multi-agent reasoning [90, 207, 209, 260, 274]	Coordinated planning or dialogue among several learned agents.	Multi-Agent Reasoning	2
7. Generalisation & robustness [5, 21, 38, 45, 63, 66, 101, 172, 192, 285]				
	7.1 Domain robustness [21, 203, 228, 251]	Techniques to maintain performance under domain shift.	Domain-Robust Generalization	5

Table 4 | Domain Survey for Foundation Model.

Cat.	Sub-category	What is covered	Typical examples	Cluster
1. Perception & Mapping [71, 168, 185, 277, 294]				
	1.1 Multimodal sensor fusion [14, 123, 224, 256, 303]	Fuse heterogeneous sensors for richer scene understanding	LiDAR–Camera Fusion; Radar–Camera Fusion; V2X Cooperative Perception ...	0,6,7,8,14,16
	1.2 3D reconstruct/occupancy [85, 133, 140, 175, 294]	Build dense or sparse geometric maps for localisation	3-D SLAM & Reconstruction; 3-D Occupancy; Efficient 3-D Representation	0,8,16
	1.3 BEV / top-view mapping [19, 93, 120, 256, 286]	Bird’s-eye or top-down representations for planning	BEV Perception; V2X Collaborative Perception	0,14,16
2. Manipulation & Grasping [13, 95, 208, 216, 239]				
	2.1 Dexterous grasping [146, 216, 249, 254, 271]	Multi-finger in-hand manipulation	Dexterous Robotic Grasping; Dexterous Grasp & Manipulation	11,12
	2.2 Generalist manipulation [36, 58, 125, 266, 280]	Single policy handles diverse objects/tasks	Generalist Robotic Manipulation; Robotic Manipulation; Humanoid Manipulation	3,4,9
	2.3 Tactile-vision fusion [51, 82, 248, 261]	Combine touch and vision for reactive grasps	Multimodal Tactile-Vision Learning	11
3. Locomotion & Navigation [34, 74, 75, 135, 239]				
	3.1 Legged locomotion control [61, 75, 135, 250, 302]	Whole-body control and adaptation on uneven terrain	Legged Robot Locomotion; Learning-Based Control & Adaptation	17
	3.2 Embodied VL navigation [104, 179, 220, 276]	Language-directed navigation with active mapping	Embodied Vision-Language Navigation; Active 3-D Mapping & Planning	7,19
4. Planning & Control [54, 91, 107, 293, 301]				
	4.1 Language/hierarchical planning [10, 69, 119, 194]	Translate language or high-level goals into executable skills	Language-Guided Planning & Control; Hierarchical Skill Planning & Adaptation	2,18
	4.2 Diffusion/Transformer policies [98, 127, 152, 173, 191]	Trajectory generation with generative sequence models	Diffusion Policies; Diffusion/Transformer Policy Models; Generative Diffusion Models	1,9,14
5. Robot Learning & Adaptation [54, 88, 91, 107, 301]				
	5.1 RL & imitation [52, 78, 109, 258, 275]	Learn skills from rewards, demonstrations or offline data	Robot Reinforcement Learning; Imitation Learning Policies; Sample-Efficient RL ...	3,9,15
	5.2 Sim to real & continual adaptation [39, 59, 103, 150, 257]	Transfer and improve policies across domains over time	Continual Sim-to-Real Adaptation; Sim-to-Real Transfer; Self-Supervised Distillation/Adaptation	0,4,15,16
	5.3 Multitask / generalisable policies [46, 158, 221, 263]	Single policy generalises across many tasks and embodiments	Multitask Generalisable Robotics	1
6. Autonomous Driving [77, 185, 201, 224, 227]				
	6.1 Motion forecasting, perception & simulation [77, 167, 227, 237]	Forecast traffic actors, all-weather perception, long-tail scenario simulation	Motion Forecasting; Trajectory Prediction; Driving Perception; Scenario Simulation	5,6,8,13,14
7. Simulation & World Models [27, 64, 117, 227, 229]				
	7.1 Generative world models [9, 70, 81, 227, 278]	Learn latent physics/world models for planning or RL	Generative World Models	12
	7.2 Self-supervised simulation [4, 16, 49, 121, 178]	Expand synthetic experience using self-supervised signals	Self-Supervised Generative Simulation	5
8. Embodied Language Robotics [10, 179, 183, 235]				
	8.1 LLM-driven robotics [26, 42, 122, 202]	Use large language models for zero-shot policy/reasoning	LLM-Driven Robotics; LLM-Enhanced Driving; LLM-Driven Zero-Shot Planning	2,13,19
	8.2 Vision-language control [10, 44, 179, 265, 287]	Pair vision with text to drive low-level actions	Vision-Language Robotic Control; Hierarchical Skill Planning & Adaptation	18
	8.3 Open-vocabulary mapping [28, 179, 183, 228, 235]	Build scene maps labelled with free-form language	Open-Vocabulary Scene Mapping	19
9. Safety & Robustness [75, 95, 185, 187, 291]				
	9.1 Safety-aware planning [43, 67, 75, 162, 185]	Explicit risk reasoning during motion generation	Safety-Aware Motion Planning	10
	9.2 Runtime monitoring [2, 134, 185, 198, 199]	Detect and mitigate failures on-the-fly	Failure Detection & Runtime Monitoring	18
	9.3 Robust control [144, 171, 193, 200, 306]	Improve stability against disturbances and uncertainties	Safety & Robustness (Locomotion)	17
10. Multi-Robot & Human Collaboration [35, 36, 54, 74, 90]				
	10.1 Multi-agent collaboration [144, 171, 193, 200, 306]	Plan and act with other robots or humans in the loop	Multi-Agent / Human-Robot Collaboration	10

Table 5 | Domain Survey for Robotics.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Robust & Generalizable Learning				
	1.1 Adversarial / OOD Robustness	Defending against or detecting malicious, anomalous or out-of-distribution inputs.	adversarial robustness; deepfake detection; anomaly detection; out-of-distribution detection	0,2,5
	1.2 Domain Adaptation & Generalization	Transferring models across different domains, devices or persons without performance drop.	domain adaptation; domain generalization; test-time adaptation; person re-identification	3,6
	1.3 Low-Data Learning	Learning reliably from scarce, imbalanced or continually arriving data.	few-shot; continual learning; long-tailed recognition; federated learning	0,1,2
2. Representation & Model Efficiency				
	2.1 Representation Learning & Distillation	Un/semisupervised learning and distillation techniques that build informative, explainable features.	self-supervised learning; semi-supervised segmentation; pseudo-label consistency; representation learning; knowledge distillation; explainability	4,5,7,10
	2.2 Efficient Architectures	Designing compact, hardware-friendly or automatically searched neural networks.	efficient ViT; neural architecture compression; sparse/quantized NAS	9
3. Generative Modeling & Editing				
	3.1 2D Generative Imaging	Synthesising or editing images/videos with controllable appearance or compression.	generative adversarial networks; image inpainting; image translation; neural style transfer; diffusion-based image/video generation; controllable generative editing; neural compression	18,21,23
	3.2 3D Neural Rendering & Scene Generation	Generating or reconstructing 3-D scenes via implicit or explicit neural representations.	neural radiance fields; 3D scene generation; 3DGS; dynamic scene reconstruction; neural rendering	25,28,29
4. 2D Perception & Enhancement				
	4.1 Detection & Segmentation	Locating objects or semantic regions in images/videos.	object detection; semantic segmentation; few-shot detection	1,10,16
	4.2 Tracking & Motion Estimation	Following objects or estimating pixel correspondences over time.	object tracking; correspondence; registration; optical flow; UAV	13,20
	4.3 Matting & Transparency	Separating foreground layers or transparency effects in images/videos.	image/video matting; trimap guidance; mask guidance; transformer-based matting models	19
	4.4 Restoration & Enhancement	Improving quality of degraded images/videos or reconstructing HDR.	image/video restoration; diffusion models for restoration; HDR	21,27
5. 3D Perception & Geometry				
	5.1 Depth & Reconstruction	Estimating depth or reconstructing 3-D structure from images.	depth estimation; stereo matching; 3D reconstruction	26
	5.2 LiDAR & 3D Detection	Understanding point clouds for object detection and semantic segmentation.	LiDAR point clouds; 3D object detection; 3D semantic segmentation	16
	5.3 Pose & Localization	Estimating 6-D object poses or localizing cameras in space.	6D pose estimation; visual localization; equivariant features	24
6. Video & Temporal Understanding				
	6.1 Temporal Action & Video Reasoning	Recognising and localising actions or reasoning over temporal video cues.	temporal action localization; video representation; multimodal reasoning	12
7. Multimodal & Vision-Language Systems				
	7.1 Vision-Language Pretraining & Retrieval	Learning cross-modal representations for zero-shot tasks or retrieval.	vision-language pretraining; zero-shot learning; cross-modal retrieval	8
	7.2 Multimodal Large Models & Grounding	Large models that jointly reason over vision and language with grounding.	multimodal large language models; visual grounding; visual reasoning; benchmark datasets; 3D vision-language grounding; scene graph generation	11,14
	7.3 Audio / Sign / Gaze Multimodality	Integrating audio, sign language or gaze with vision tasks.	audio-visual learning; sign language processing; gaze estimation	15
8. Human-Centric Understanding & Animation				
	8.1 Pose & Interaction	Estimating human body pose and modelling human-object interactions.	3D human pose estimation; human-object interaction; transformer-based motion generation	22
	8.2 Avatars & Animation	Building and animating realistic 3-D human avatars.	3D human avatars; pose-driven animation; neural rendering of humans	28
9. Embodied & Autonomous Systems				
	9.1 Embodied Navigation & Mapping	Perception and planning for agents navigating 3-D environments.	embodied navigation; HD-map generation; lane generation	14,17
	9.2 Trajectory Prediction & Traffic Simulation	Forecasting future paths and simulating realistic traffic participants.	trajectory prediction; data-driven traffic simulation	17

Table 6 | Domain Survey for Computer Vision.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Information Extraction				
	1.1 Entity & Relation Extraction	Automatic detection of named entities plus the semantic relations or events connecting them.	Named Entity Recognition, Relation Extraction, Event Extraction	0
2. Text Generation & Summarization				
	2.1 Summarization & Keyphrase Generation	Producing concise summaries or keyphrases from longer documents.	Summarization, Keyphrase, Evaluation	12
	2.2 Controllable & Stylistic Generation	Generating text under user-specified style or attribute constraints.	Style transfer, controllable text generation, representations	27
3. Dialogue & Conversational Systems				
	3.1 Task-oriented Dialogue	Dialogue systems that track state and generate responses to accomplish user goals.	Dialogue systems, Response generation, Dialogue state tracking	14
	3.2 Empathetic & Safe Dialogue	Handling empathy, hate speech and multimodal cues in conversations.	HateSpeechDetection, EmpatheticDialogue, ...	25
4. Multilingual & Cross-lingual NLP				
	4.1 Multilingual Modeling & Transfer	Building models that operate across many (often low-resource) languages and transfer knowledge between them.	low-resource languages, multilingual language models, cross-lingual transfer	16
	4.2 Multilingual Machine Translation	Neural translation among multiple language pairs, often using shared or distilled models.	Neural machine, Knowledge distill, Multilingual modeling	19
	4.3 Multimodal Low-Resource Speech	Speech translation/recognition when data are scarce or involve multiple modalities.	speech translation, multimodal learning, low-resource speech	15
5. Knowledge & Reasoning				
	5.1 Knowledge Graph Reasoning	Embedding and temporal/causal reasoning over structured knowledge graphs.	knowledge graph embedding, event causality reasoning, temporal knowledge reasoning	1
	5.2 Mathematical & Chain-of-Thought Reasoning	Using large language models for step-by-step logical or mathematical reasoning.	Large language models, Mathematical reasoning, Chain-of-thought prompting	10
	5.3 Compositional & Syntactic Generalization	Probing or improving models to generalize compositionally or parse syntax.	syntactic parsing, compositional generalization, lan. model probing	3
6. Retrieval & Question Answering				
	6.1 Dense Retrieval & RAG	Learning dense vector search for open-domain QA and retrieval-augmented generation.	Dense retrieval, open-domain question answering, retrieval-augmented generation	4
	6.2 Table & Structured QA / Generation	Mapping natural language to SQL, answering queries or generating text from structured data.	Text-to-SQL, Table Question Answering, Data-to-Text Generation	2
7. Evaluation, Alignment & Editing				
	7.1 LLM Evaluation & Human Alignment	Designing metrics and feedback loops to align large language models with human intent.	LLM evaluation, alignment methods, human feedback	23
	7.2 Hallucination, Calibration & Knowledge Editing	Detecting/mitigating false outputs and editing or calibrating model knowledge.	hallucination, knowledge editing, calibration	11
	7.3 Evaluation Metrics & Data Augmentation	Developing metrics and synthetic data (incl. figurative language) to assess or improve models.	Evaluation metrics, Data augmentation, Figurative language	13
8. Model Training Paradigms & Efficiency				
	8.1 Continual, In-context & Instruction Tuning	Allowing models to learn new tasks or follow instructions without full retraining.	continual learning, instruction tuning, in-context learning	17
	8.2 Parameter-Efficient & Compressed Models	Reducing training/inference cost via adapters, pruning or lightweight fine-tuning.	parameter-efficient fine-tuning, model compression for LLMs	20
	8.3 Transformer Efficiency & Long-Context Modeling	Architectural or computational methods to scale transformers to longer contexts efficiently.	Transformer efficiency, long-context modeling, adaptive computation	7
	8.4 Sentence & Multilingual Representation Learning	Contrastive or related methods to build versatile sentence embeddings across languages.	multilingual representation learning, sentence embeddings, contrastive learning	6
9. Safety, Bias & Robustness				
	9.1 Social Bias & Fairness	Measuring and mitigating demographic or social biases in NLP systems.	social bias, debiasing, fairness evaluation	26
	9.2 Misinformation & Fact Verification	Detecting false claims, AI-generated text or aligning model values with truthfulness.	fact verification, misinformation detection, evidence retrieval, fake detection, value alignment	18,28
	9.3 Security & Privacy Robustness	Protecting models against adversarial, backdoor or privacy attacks.	adversarial robustness, backdoor attacks, privacy preservation	29
10. Agents & Interactive Reasoning				
	10.1 LLM-based Agents & Planning	Using large language models as autonomous agents capable of interactive planning and theory-of-mind reasoning.	LLM agents, interactive planning, theory of mind	21
11. Code Intelligence				
	11.1 Code Generation & Benchmarks	Generating executable code and evaluating models on coding tasks.	code generation, large language models, benchmark evaluation	24

Table 7 | Domain Survey for Natural Language Processing.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Generative Modelling & Media Synthesis				
	1.1 Image / Video Generation & Editing	Models that create or edit 2-D or temporal visual content via generative techniques.	Generative modeling; Image synthesis/editing; Diffusion-based methods; Optimal Transport; Diffusion Models	1, 3
	1.2 3-D Object & Molecule Generation	Generating 3-D shapes or molecular structures using geometry-aware or equivariant models.	3D shape generation; neural implicit representations; point cloud reconstruction; equivariant GNNs; 3D molecular generation; drug discovery	9, 16
	1.3 Audio & Speech Synthesis	Producing speech or audio from text or multimodal cues, via diffusion models.	text-to-speech; audio-visual; diffusion	2
2. Representation & Transfer Learning				
	2.1 Continual / Few-Shot / Domain Adaptation	Adapting models continually, with few examples, or across shifting domains.	Continual Learning; Few-Shot Learning; Domain Adaptation	0
	2.2 Self-, Contrastive & Retrieval-Augmented Learning	Building rich representations via self/contrastive learning or external retrieval augmentation.	Self-supervised Learning; contrastive learning; disentangled representations; clustering; language-modeling; retrieval-augmentation; representation-learning	5, 6, 15
	2.3 Parameter-Efficient Transfer	Adapting large transformers with minimal new parameters and compute.	Efficient-transformer-architectures; parameter-efficient-fine-tuning; multilingual-adaptation	7
3. Robustness, Security & Privacy				
	3.1 Adversarial & Backdoor Robustness	Defending against adversarial or backdoor manipulations and distribution shifts.	adversarial robustness; backdoor attacks; robustness; knowledge distillation; distribution shift	10, 4
	3.2 Uncertainty & Interpretability	Quantifying model confidence and explaining predictions.	uncertainty estimation; conformal prediction; model interpretability	14
	3.3 Privacy & Machine Unlearning	Ensuring data privacy and enabling deletion or secure distributed learning.	differential privacy; machine unlearning; federated learning; robust optimization	19, 24
4. Model Efficiency & Compression				
	4.1 Pruning, Quantization & Embedding Compression	Compressing networks by pruning, quantizing or embedding reduction for efficient deployment.	Network pruning; Sparse_Network_Pruning; Low-precision quantization; Embedding_Compression; Efficient architecture search; Recommendation_Systems	8, 21
5. Geometric & Graph Learning				
	5.1 Equivariant / Geometric Deep Networks	Networks that respect group symmetries to learn geometric or physical structures.	equivariant neural networks; group symmetry; geometric deep learning	12
	5.2 Graph Neural Network Theory	Theoretical properties, expressivity and robustness of Graph Neural Networks.	Graph Neural Networks; Expressivity; Robustness	20
6. Optimization & Theory				
	6.1 Non-convex & Stochastic Optimization	Algorithms and analysis for nonconvex optimization with stochastic gradients.	Nonconvex optimization; Stochastic gradient methods; Convergence analysis	18
	6.2 Neural Network Theory & Neuroscience Inspiration	Theoretical studies and neuro-inspired modeling of recurrent nets.	recurrent neural networks; neuroscience-inspired modeling; theoretical analysis	17
7. Reinforcement Learning & Embodied Intelligence				
	7.1 Core & Offline Reinforcement Learning	Improving sample efficiency and offline policy learning in RL.	Reinforcement Learning; Offline Learning; Sample Efficiency	29
	7.2 Multi-Agent & Dialogue RL	Learning cooperation, competition or dialogue among multiple agents.	multi-agent reinforcement learning; bandit algorithms; game-theoretic learning; dialogue systems; multi-agent collaboration; reinforcement learning	28, 26
	7.3 Embodied AI & Robotics	Training embodied agents and robots via differentiable simulation and manipulation tasks.	Embodied AI; Robotic manipulation; Differentiable simulation	27
8. Multimodal Perception & Reasoning				
	8.1 Vision-Language & Knowledge Reasoning	Joint reasoning across vision and language plus knowledge graphs.	vision-language reasoning; knowledge graph learning; compositional generalization	11
	8.2 Video Understanding & 3-D Perception	Temporal and 3-D understanding of videos and dynamic scenes.	Video understanding; Temporal modeling; 3D perception	13
9. Scientific & Symbolic Machine Learning				
	9.1 Physics & Differential Equation-guided Learning	Learning operators governed by physical laws and differential equations.	Neural differential equations; Physics-informed operator learning; Spatiotemporal forecasting	23
	9.2 Program Synthesis & Automated Reasoning	Automatically generating code or formal proofs from specifications.	Program synthesis; Code generation; Theorem proving	22
	9.3 Combinatorial, Causal & Bayesian Optimization	Optimization over discrete structures, causal questions or Bayesian objectives.	Combinatorial optimization; Causal inference; Bayesian optimization	25

Table 8 | Domain Survey for Machine Learning.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Life Sciences & Biomedicine				
	1.1 Immuno-oncology & Metabolic Signalling	Immune mechanisms in cancer and metabolic cues that modulate them	T-cell immunity; tumour microenvironment; metabolic signalling	0
	1.2 Cancer Genomics & Epigenetics	Genetic and epigenetic alterations driving oncogenesis and therapy response	Cancer; DNA repair; epigenetics	1
	1.3 Infectious Disease & Microbiome Interactions	Host–pathogen dynamics and microbiome ecology shaping antimicrobial strategies	Host–pathogen interactions; antimicrobial therapeutics; microbiome dynamics	2
	1.4 Neuro-immune Metabolism & Aging	Crosstalk between immune system, metabolism and brain across aging	Neuroimmunology; metabolism; aging	3
	1.5 Genome Editing & Microbial/Plant Immunity	Engineering genomes and decoding microbial	plant defence mechanisms & CRISPR-based genome editing; bacterial anti-phage defence; plant immune signalling	4
	1.6 Protein & RNA Engineering	Designing proteins and regulating chromatin/RNA to control cell function	Protein design; chromatin regulation; RNA processing	5
	1.7 Neural Epigenetics & Disorders	Epigenetic regulation of neural plasticity and neuropsychiatric disease	Neural circuit plasticity; epigenetic regulation; neuropsychiatric disorders	6
	1.8 Population & Single-Cell Genomics	Sequencing-based mapping of genetic variation at population & cellular resolution	Genome sequencing; population genetics; single-cell transcriptomics	7
	1.9 Connectomics & Behaviour	Structural mapping of neural circuits to explain behaviour	Connectomics; neural circuit mapping; behaviour	8
	1.10 Evolutionary Genomics & Paleobiology	Reconstructing evolutionary history using ancient DNA and fossils	Paleogenomics; prehistoric migrations; fossil record	9
2. Chemistry & Materials Science				
	2.1 Catalysis & Green Synthesis	Catalytic and synthetic routes for sustainable chemical production	Advanced catalysis; sustainable chemistry; synthetic methodologies	10
	2.2 Functional Materials for Energy & Electronics	Multifunctional materials for energy storage and flexible devices	Advanced materials; energy storage; flexible electronics	12
	2.3 Perovskite Solar Technologies	High-efficiency perovskite and tandem photovoltaics with interface engineering	Perovskite photovoltaics; tandem solar cells; interface passivation	13
	2.4 Integrated Photonics & Optoelectronic Integration	Integration of perovskites and 2D semiconductors into photonic platforms	Integrated photonics; perovskite optoelectronics; 2D semiconductor integration	14
3. Physics & Quantum Technology				
	3.1 Quantum Materials	Emergent quantum phases in topological and moiré systems, incl. unconventional superconductivity	Topological quantum matter; moiré heterostructures; unconventional superconductivity	16
	3.2 Quantum Computing Hardware & Networks	Scalable, fault-tolerant quantum processors and quantum communication links	Fault-tolerant quantum computing; scalable qubit hardware; quantum networking	17
4. Earth & Environmental Science				
	4.1 Climate Change & Ecosystem Impacts	How climate change alters ecosystems and the global environment	Climate-change; ecosystem-impacts; global-environment	15
5. Astronomy & Astrophysics				
	5.1 Exoplanetary Science with JWST	Characterising exoplanet atmospheres and interiors using JWST observations	Exoplanet atmospheres; planetary interiors; JWST observations	18
	5.2 Early-Universe & Black-Hole Astronomy	Galaxy formation and supermassive black holes in the early Universe probed with JWST	JWST; early-Universe galaxies; supermassive black holes	19
6. Computer Science & Artificial Intelligence				
	6.1 Foundation & Trustworthy AI	Large foundation models, applied AI and methods ensuring fairness & reliability	Foundation models; applied artificial intelligence; fairness and reliability	11

Table 9 | Domain Survey for Natural related Topics.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Earth & Environmental Sciences				
	1.1 Climate & Ecosystem Dynamics	Interactions among climate change, carbon cycling and biodiversity, including conservation responses	Biodiversity loss; Conservation strategies; Climate change impacts; Climate change; Carbon cycle; Environmental impacts	0,1
	1.2 Geophysical Processes	Physical processes shaping Earth's solid and cryospheric systems	Earthquakes; Volcanism; Ice dynamics	2
2. Space Science				
	2.1 Stellar & Space-Plasma Physics	Physics of stars, solar activity and the interstellar medium	Compact objects; Interstellar medium; Solar activity	9
3. Biological Sciences				
	3.1 Evolutionary Genomics	Genetic mechanisms driving adaptation and speciation	Evolutionary genomics; adaptation; speciation	3
	3.2 Molecular & Cellular Regulation	Molecular signaling and structural mechanisms governing development and genome integrity	Hormone signaling; Immune defense; Developmental regulation; Genome stability; Chromosome segregation; Cryo-EM structural biology	4,5
	3.3 Neurobiology & Systems Neuroscience	Gene-to-circuit bases of neural function, plasticity and behaviour	Neuroscience; Gene regulation; Single-cell; neural circuits; synaptic plasticity; behavior	6,8
	3.4 Immunity, Infection & Therapy	Host defence mechanisms and engineered immunotherapies against pathogens and cancer	Antiphage immunity; CRISPR systems; Antibiotic discovery; Immunoregulation; Metabolic signaling; Cancer therapy; Infectious disease; Immunotherapy; Molecular engineering	7,11,13
	3.5 Synthetic & Computational Biology	Design of biomolecules and biological systems using AI and synthetic methods	protein-design; deep-learning; synthetic-biology	12
4. Materials & Chemical Sciences				
	4.1 Catalysis & Chemical Transformations	Selective catalytic methods for constructing organic molecules	Catalytic organic synthesis; Radical-mediated transformations; Selective C–H functionalization	14
	4.2 Advanced Functional Materials	Smart, biointegrated and nanostructured materials with tailored properties	Smart materials; Biointegrated electronics; Soft robotics; Nanostructured materials; Energy storage; Functional properties	15,17
	4.3 Energy Conversion & Separation Materials	Materials enabling electrochemical, thermal and membrane-based energy technologies	Electrocatalysis; Porous framework materials; Membrane separations; Perovskite photovoltaics; Thermoelectric devices; Radiative cooling	16,18
5. Physics & Quantum Technologies				
	5.1 Quantum Materials & Information	Exotic quantum phases and their application to information processing	Topological quantum phases; Quantum information processing; Strongly correlated matter	19
6. Computational & Social Systems Science				
	6.1 Information Dynamics & Society	Computational study of information spread and persuasion in sociotechnical systems	misinformation propagation; democratic polarization; AI-mediated persuasion	10

Table 10 | Structured overview of clustered science research areas.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Pharmacogenomics & Genetics-Guided Therapy				
	1.1 Cytochrome P450 Genotype-Driven Therapy	Links CYP450 genetic variants to drug exposure and response for individualised dosing.	CYP2C19 pharmacogenetics, star-allele variability, precision antithrombotic therapy, antidepressant pharmacogenetics	1,3,4
	1.2 Transporter Pharmacogenetics	Examines genetic variation in drug transporters and its impact on safety and efficacy.	SLCO1B1 variants, statin-associated muscle symptoms	6
	1.3 PGx Implementation & Economic Evaluation	Assesses clinical decision support, workflow integration, and the cost-effectiveness of pharmacogenomics.	pharmacogenomics implementation, clinical decision support, cost-effectiveness	11
	1.4 Oncology / High-Risk Therapy PGx	Uses germline variants to predict toxicity and guide dosing of high-risk or anticancer drugs.	DPYD variants, TPMT variants, NUDT15 variants, chemotherapy toxicity prediction, pharmacogenetic-guided dosing	21
2. Quantitative Pharmacology & Model-Informed Drug Development				
	2.1 Population PK & Dose Optimisation	Applies population PK and exposure–response models to refine dosing in diverse patients.	population pharmacokinetics, precision dosing, anticoagulants, exposure–response modelling, oncology real-world evidence	2,17,25,29
	2.2 Mechanistic PBPK & Special Populations	Uses physiologically based PK models to predict drug disposition in paediatrics, the CNS, pregnancy, and other special populations.	paediatric PBPK modelling, CNS drug delivery, maternal–infant pharmacology, anti-infective therapy	7,19
	2.3 Exposure–Response for Biologic / Cell Therapies	Characterises PK/PD and dose–response of biologics and cell-based therapies.	PK/PD modelling, haematological therapies, biologic PK/PD, immunomodulatory therapies, T-cell engagers	15,26,28
	2.4 Machine-Learning-Assisted Precision Dosing	Integrates machine learning with PK models and real-world factors to individualise therapy.	machine learning, precision pharmacokinetic modelling, transplant immunosuppressant dosing, gut microbiota influence, model-informed drug development	18,20
3. Drug Metabolism, Transport & Interaction Science				
	3.1 Enzyme-Mediated DDIs & Prediction	Investigates cytochrome P450 interactions and uses models to forecast clinical risk.	cytochrome P450, PBPK modelling, QT prolongation	14,5
	3.2 Transporter-Mediated DDIs & Biomarkers	Studies renal and hepatic transporters and endogenous probes to detect interaction liability.	renal transporters, hepatic transporters, endogenous biomarkers of transporter activity	13,24,27
	3.3 Clinical DDIs & Risk Management	Documents real-world interaction scenarios and strategies to mitigate adverse outcomes.	opioid overdose management, nirmatrelvir/ritonavir interactions, EHR-based pharmacovigilance	0,8
4. Regulatory Science & Evidence Generation				
	4.1 Trial Diversity & Health Equity	Promotes representative enrolment and equitable access in clinical research.	clinical trial diversity, health equity, regulatory initiatives	10
	4.2 Real-World Evidence & External Controls	Leverages observational data and synthetic controls to inform regulatory decisions.	real-world evidence, external control trials, regulatory frameworks	9
	4.3 Drug-Lifecycle Oversight & Lag Analysis	Evaluates approval timelines, post-marketing requirements, and regulatory performance.	regulatory science, drug lag, post-marketing studies	16
	4.4 Biomarker / Rare Disease / Biosimilar Qualification	Establishes evidentiary standards for biomarkers, orphan products, and biosimilars.	rare-disease drug development, biomarker qualification, biosimilar development, PK/PD biomarkers, regulatory strategies	22,23
5. Clinical Therapeutics & Outcomes Research				
	5.1 Cardio-Renal & Metabolic Outcomes	Assesses the long-term efficacy and safety of metabolic therapies on cardiovascular and renal endpoints.	SGLT2 inhibitors, cardiovascular–renal outcomes, antidiabetic drug safety	12

Table 11 | Domain Survey for Science related Survey.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Textual inputs				
	1.1 Large-scale tokenized corpora	Massive general-domain text for LM pre-training	Web pages; Wikipedia; books; C4; Pile; WikiText; OpenWebText; SlimPajama	11
	1.2 Prompt & interaction data	User/system prompts and model replies gathered for alignment, RLHF or robustness	Prompts/questions; model responses; preference/reward labels; adversarial triggers; long-context demonstrations	0, 2
	1.3 Problem statements with context	Natural-language tasks paired with explicit structured knowledge or code/data schemas	NL problem + knowledge graph/database schema/code stub; reasoning traces or step-by-step solutions	14
2. Visual inputs (images)				
	2.1 Raw images	Canonical labelled/unlabelled images after basic augmentation	ImageNet, CIFAR, COCO photos; medical scans	19
	2.2 Cued images	Images supplied with auxiliary spatial/sensor cues	Low-light or blurry photos + masks; camera poses; depth/event data; points/boxes	17
	2.3 Patch or region tokens	Visual patches embedded as tokens for transformer processing	ViT/MAE patches from images or single video frames	3
3. Video & motion inputs				
	3.1 Video streams with motion cues	Time-ordered frames plus motion/semantic signals	Video frames; optical flow; 3-D pose; segmentation masks; aligned audio track	13
4. 3-D & spatial inputs				
	4.1 Geometry & depth representations	Explicit 3-D or depth data for spatial reasoning	Point clouds; RGB-D images; TSDF/voxel grids; meshes; camera extrinsics; semantic labels	1
5. Multimodal token sequences				
	5.1 Cross-modal token bags	Tokens from diverse modalities embedded with positional info	Text, audio, vision, graphs, biology tokens with position vectors	12, 18
	5.2 Encoder-fused tokens	Tokens from separate encoders concatenated into one sequence	CLIP/ViT image tokens + BERT/LLaMA text tokens	15
	5.3 Normalized latent embeddings	Modality-specific encoders map data into a shared latent space (may include placeholders)	Text, images, video, audio all $\rightarrow$ joint embeddings (missing modalities allowed)	4
6. Generative-model conditioning				
	6.1 Diffusion noise schedule	Noisy latent sample x_t , timestep token t , optional class/text/geometry conditioning	$x_t + z$; timestep t ; class label; pose map; depth; edges	16
	6.2 Auxiliary generation cues	User-supplied hints steering image generation or editing	Reference image; mask; depth; pose; layout; bounding boxes	10
7. Task-oriented multimodal inputs				
	7.1 Visual observations + NL prompts	Perception frames paired with a natural-language task or edit instruction	Screenshot + “click the red button”; video frame + “highlight the pedestrian”	6
	7.2 Image-text pairs with cues	Captioned/questioned images often carrying region annotations	Image + caption; VQA triplets; bounding-box / mask annotations	7
	7.3 Embodied-agent context	Agent perception, proprioception	history combined with a goal description & RGB-D stream; past actions; goal text (“navigate to the chair”)	8
8. Sequential & trajectory inputs				
	8.1 Offline state-action trajectories	Logged sequences for offline RL or behaviour cloning	Time-series control signals; graphs; 3-D skeleton poses; human preference labels	9
9. Inverse-problem observations				
	9.1 Corrupted measurements with ground truth	Raw measurements transformed by known operators, paired with target outputs	MRI k -space + mask; blurred $\rightarrow$ sharp image pairs; noisy sensor data	5

Table 12 | Structured summary of input types used in foundation-model papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Representation & Architecture				
	1.1 Token & latent representation	Mapping raw data to discrete/continuous tokens or latents	Latent representation learning; token/patch embedding; visual-token projection ...	0, 14, 17, 18, 19
	1.2 Attention & Transformer variants	Architectural changes that make attention cheaper or deeper	Sparse/low-rank attention; spatiotemporal attention; positional scaling ...	2, 10, 11, 14, 17, 18, 19
	1.3 Mixture-of-Experts & routing	Dynamic selection of expert blocks or routes	Modular MoE; dynamic routing; MoE token routing ...	0, 11, 15, 17, 19
2. Generative Paradigms				
	2.1 Diffusion & score-based generation	Noise-to-data generative flows	UNet diffusion; latent diffusion; guided conditional sampling ...	2, 3, 4, 5, 10, 13, 18
	2.2 Energy-based & control formulations	Sampling by minimising energy or solving control processes	Energy-based score learning; optimal-control SDE; solver-accelerated inversion ...	13
	2.3 Probabilistic & masked inference	Non-diffusion probabilistic decoders	Probabilistic generative inference; masked auto-encoding; next-token prediction ...	0, 17, 18, 19
3. Multimodal Alignment & Fusion				
	3.1 Encoders → shared latent space	Separate encoders project each modality into a common space	Modality-specific encoders; projection layers; frozen CLIP backbone ...	1, 6, 12, 18, 19
	3.2 Cross-attention fusion & conditioning	Mechanisms for interaction between modalities	Cross-attention fusion; prompt cross-attention; multimodal concatenation ...	2, 4, 6, 10, 12, 14, 18, 19
	3.3 Vision/Video-language alignment	Aligning paired modalities in latent space	Video-language alignment; contrastive alignment; generative self-supervision ...	1, 8, 12, 14, 18, 19
4. Adaptation & Efficiency				
	4.1 Parameter-efficient adaptation	Updating only a small subset of weights or added modules	LoRA/adapters; low-rank tuning; modular fusion ...	1, 2, 4, 10, 12, 15, 16, 19
	4.2 Prompting & modular extensions	Steering frozen backbones with prompts or plug-ins	Prompt conditioning; chain-of-thought prompting; tool invocation ...	1, 6, 7, 9, 17, 19
	4.3 Compression & efficient training	Reducing compute, memory or training cost	Quantisation; pruning-distillation; communication-efficient sharding ...	2, 5, 10, 15, 16, 19
5. Reasoning & Interaction				
	5.1 Chain-of-thought & tool reasoning	Explicit reasoning traces or calls to external tools	Chain-of-thought reasoning; retrieval-augmented reasoning; self-refinement ...	6, 7, 8, 9
	5.2 RL & preference modeling	Reinforcement or preference-based optimisation	Preference-conditioned policies; RLHF alignment; optimal-transport RL ...	3, 9
	5.3 Multi-agent / planner loops	Multiple interacting agents or explicit planner loops	Multi-agent collaboration; Planner-Actor-Corrector-Verifier loop ...	7, 8
6. Robustness & Domain Shift				
	6.1 Uncertainty & robust optimisation	Estimating confidence and resisting adversarial inputs	Uncertainty quantification; adaptive memory; adversarial robustness ...	0, 9, 16
	6.2 Domain adaptation & model editing	Adapting or editing knowledge post-training	Targeted model editing; synthetic-data adaptation; knowledge probing ...	4, 9, 16, 19

Table 13 | Structured summary of modeling techniques used in foundation-model papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Language-centric outputs				
	1.1 Token probabilities & sequences	Autoregressive LMs: token logits or generated text	next-token probability distributions; generated token sequences	2
	1.2 Aligned LLM responses	Instruction-tuned completions for reasoning, safety, long-context	helpful-harmless-truthful responses; safe refusals; reasoning-enhanced; hallucination-reduced	3, 15
	1.3 Reasoning traces & answers	Chain-of-thought steps plus final decoded result	chain-of-thought reasoning; intermediate outputs; final answers	0
	1.4 Visually-grounded NL outputs	Language grounded in image/video content	captions; visual-QA; reasoning with bounding boxes, masks	9
2. Generative visual & multimodal outputs				
	2.1 Photorealistic images	High-fidelity images from text or prompts	photorealistic; identity-preserving; context-coherent images; text-conditioned images	1, 18
	2.2 Video / motion generation	Consistent video or 3D motion from text	temporally-consistent video; 3-D motion generation/editing	5
	2.3 3D scenes & assets	Meshes, NeRFs, Gaussian fields for rendering/editing	meshes; point clouds; NeRF; editable scene/asset generation	11
	2.4 Multimodal reconstructions	Images, video, audio decoded from latents	reconstructed/generated multi-modal data	19
	2.5 Diffusion samples & noise	Reverse-diffusion outputs with noise estimates	generated or reconstructed samples... plus noise/score estimates	7
3. Predictive & structured outputs				
	3.1 Classification scores	Class labels, probabilities, or logits	class labels / probabilities / logits	4, 13
	3.2 Localization & segmentation	Masks, boxes, poses, or captions pinpointing content	segmentation masks; bounding boxes; 3-D localization/pose; textual grounding/captions	16, 13
	3.3 Structured artefacts	Graphs, coordinates, flows, molecules, causal terms	graphs; poses; flows; molecular/crystal structures; uncertainties; causal/physical parameters	14
	3.4 Downstream embeddings	Transformed features for later task use	transformed feature embeddings; reconstructed/generative outputs	13
	3.5 Control & planning	Predicted actions, plans, or trajectories	action sequences & control commands; task-grounded plans	6
4. Evaluation, improvement & efficiency outputs				
	4.1 Metrics & benchmarks	Accuracy, bias, uncertainty, safety, etc.	accuracy; F1; ROC-AUC; Elo; bias; calibration	17
	4.2 Enhanced / corrected artefacts	Outputs improving other models (predictions, data, signals)	corrected predictions; synthetic/augmented data; anomaly/OOD scores; attribution indicators	10
	4.3 Compressed models	Quantised/tuned checkpoints with reduced cost	efficient, compressed, fine-tuned foundation models	8
5. Embedding-space outputs				
	5.1 Aligned multimodal embeddings	Joint space for text/image/audio enabling retrieval or classification	aligned multimodal embeddings; zero-shot classification	12

Table 14 | Structured summary of output types used in foundation-model papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Language-modeling objectives				
	Next / Masked-token Prediction	Minimize CE on next/masked token.	Next-/masked-token pred.; LM; CE/NLL min.; aux reg.	4
	General LLM Advancement	Improve reasoning, alignment, efficiency, robustness.	Reasoning; alignment; eval.; efficiency; robustness; multi-domain	12
2. Alignment & safety objectives				
	Human-Preference Alignment	Maximize learned reward; limit divergence.	Pref. align; reward max.; safety-divergence reg.	1
	Hallucination & Bias Mitigation	Cut hallucinations/bias via grounding alignment.	Hallucination det./mit.; x-modal align/ground; bias red.	0
	General Safety & Robustness	Losses for safety, explainability, robust autonomy.	Alignment; safety; efficiency; general.; explain.; autonomy	6
	Security & Privacy Defense	Defend attacks, watermark, erase concepts.	Adv. robustness; watermark; backdoor/membership defense; privacy; concept erase; interp.	7
3. Adaptation & continual-learning				
	Prompt / Self-training Adaptation	Prompt/pseudo-label adapt for zero/few-shot.	FM adapt; prompt/pseudo-label; zero-/few-shot OVR; robustness; domain gen.	10
	Retention-Regularized Fine-tuning	Regularize fine-tuning to retain knowledge.	Task loss + retention reg.; preserve knowledge; generalization	17
4. Multimodal objectives				
	Unified Multimodal Representations	Vision-language align, ground, reason.	Unif. multimodal; V-L align; grounding; x-modal reason.; zero/few-shot; cont. adapt.	3
	Contrastive & Masked Alignment	Contrastive+masked for joint embeddings.	X-modal contrast; masked recon.; joint class.; dist. align	13
	3D / Multi-view Generation	Cross-modal loss for 3D-consistent views.	Hi-fid 3D multi-view gen.; sparse 2D/text	19
5. Generative diffusion objectives				
	Core Enhancement	Faster, higher-quality diffusion via guidance.	Accelerate train/inf.; guide/loss opt.; fidelity; diversity; control	16
	Noise-prediction & Score-matching	Train via noise pred., reconstr., ELBO.	Noise pred denoise; recon. fid. min.; score/ELBO opt.	18
	Video / Motion Diffusion	Conditioned diffusion for coherent video.	Hi-fid coherent video/motion synth.; control; prompt align	2
	Controllable Image Diffusion	Steer image diffusion for fairness etc.	Align; personalise; fairness; diversity; spatial; hi fid.; light train	5
	Latent & Denoising Regularization	Extra denoise/latent loss.	Denoise min.; latent align; cond. reg.; dist. fid. train	8
6. Policy-learning & RL				
	Multi-task Policy RL	One policy via cloning+pref. RL.	Multi-task policy; behavior/diffusion cloning; pref.-aligned RL; reward exp.	14
7. Optimization & efficiency				
	Loss & Representation Matching	Minimize task loss, align distributions.	Task combined losses; dist align; repr match; reg. opt.	11, 15
	Compute / Memory Efficiency	Cut compute/memory, keep accuracy.	Min. compute/memory/param cost; train/ft/inf.	9

Table 15 | Structured summary of learning objectives used in foundation-model papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Pre-Training & Representation Learning				
	1.1 Contrastive/masked vision–language pre-training	Learn aligned image–text embeddings before any task-specific tuning.	ViT/CLIP contrastive/masked pretrain; strong aug.; $T=0.07$; 40–600 ep finetune.	5
	1.2 Adapter-aided diffusion image/video pre-training	Freeze released checkpoints; add lightweight adapters to scale.	Frozen ckpt + LoRA/prompt; AdamW + cos LR; prog-res; CF guidance.	13
2. Fine-Tuning & Adaptation				
	2.1 Vision–language instruction tuning	Turn a frozen VLM into an instruction follower.	Image–text pretrain→inst. tune; PEFT; opt. RLHF.	1
	2.2 Parameter-efficient domain adaptation	Keep backbone frozen; adapt via prompts/adapters only.	Prompt/adaptor/LoRA; distill or contrastive shift.	10
	2.3 Instruction SFT + retrieval alignment	Align an LLM with retrieval and preferences.	Multi-stage SFT; retrieval ctx; DPO/RLHF; rerank→generate.	6
	2.4 3-D coarse-to-fine diffusion adaptation	Make diffusion/LLM backbones 3-D consistent.	Alt. SDS/guidance; synth views; render-denoise distill.	4
	2.5 Video-diffusion adapter tuning	Specialise image diffusion for temporal output.	Temp/spatial adapters; latent denoise; low→high-res.	7
	2.6 Controllable diffusion sampling	Add style/identity knobs without retraining core model.	Var-score-recon losses; dyn. guidance; feature mod.	8
	2.7 Layout / prompt-conditioned diffusion	Condition generation on structured layouts or text.	LLM layout cond.; masked-attn sampling; coarse→fine.	11
	2.8 Composite-loss self-supervised fine-tuning	Improve a backbone with multiple unsupervised signals.	Mask/noise; contrast+recon+distill; EMA teacher.	15
	2.9 Pseudo-label self-training	Self-train using synthetic multimodal labels.	Synth labels (Diff/LLM/SAM); filter; adapter FT; contrast/distill.	16
3. Reinforcement Learning & Control				
	3.1 Diffusion-backed policy optimisation	Blend BC and RL signals for policy training.	Traj samp; BC+PPO; Q-guided denoise; self-play.	0
	3.2 Hierarchical planning & embodied control	Combine VLM/LLM skills with robotic policies.	Skill seg; hier plan; RH control; real-time accel.	19
4. Efficiency & Compression				
	4.1 Model compression & quantisation	Shrink models with minimal retraining.	Low-rank+sparse; mixed-prec.; prune+search.	2
	4.2 Transformer training / inference acceleration	Architectural and parallel tricks to cut runtime.	Multi-dev partition; sparse/flash attn; KV prune; stride denoise.	9
	4.3 Hyper-parameter & infrastructure optimisation	Well-tuned schedules and distributed stacks.	AdamW warm-cos LR; FP16/BF16; DeepSpeed; 100k–500k steps.	17
5. Safety & Adversarial Robustness				
	7.1 Jailbreak & adversarial prompt synthesis	Craft inputs that bypass safety guards.	Harmful data; shadow model; grad token opt; synth prompt.	12

Table 16 | Structured summary of training recipes used in foundation-model papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Vision & Imaging Sensors				
	1.1 RGB cameras	Monocular, stereo, multi-view, surround-view or panoramic cameras producing color frames; used for appearance-based perception.	front/side/rear vehicle cameras, egocentric/wrist/-head cameras, aerial/on-board cameras	0, 1, 2, 3, 4, 6, 7, 9, 10, 11, 13, 14, 15, 16, 17, 18, 19
	1.2 RGB-D cameras	Active or structured-light cameras that output synchronized color + depth images.	Intel RealSense, Azure Kinect, panoramic RGB-D rigs	3, 7, 10, 11, 13, 14, 17, 18, 19
	1.3 Event (neuromorphic) cameras	Asynchronous sensors emitting per-pixel brightness changes with micro-second latency.	DVS, DAVIS	9
	1.4 Thermal / LWIR cameras	Passive long-wave IR imagers for temperature or night-vision cues.	Thermal cameras, LWIR DoFP polarization cameras	3, 14
2. Depth & Range Sensors				
	2.1 LiDAR	Spinning or solid-state laser scanners returning 3-D point-clouds.	Multi-beam/spinning LiDAR, PolLidar wavefront lidar	3, 4, 5, 11, 13, 14, 16, 17
	2.2 Time-of-Flight cameras	Pulsed or continuous-wave light cameras computing per-pixel range.	Indirect/monocular ToF depth cameras, AMCW-ToF	9, 14
	2.3 Radar	mmWave / FMCW / 4-D imaging radars measuring range-Doppler or heat-maps.	Automotive mmWave/FMCW, MIMO imaging radar	3, 4, 14
3. Proprioceptive Sensors				
	3.1 Joint & wheel encoders	Optical or magnetic sensors giving joint angle / wheel ticks.	joint encoders, wheel encoders	3, 7, 8, 13, 16
	3.2 IMUs	3-axis accelerometers & gyros providing orientation/velocity.	IMU, pose modules	3, 4, 7, 13, 16, 17
	3.3 Force / torque sensors	Strain-gauge or multi-axis transducers measuring interaction forces.	force-torque sensors, motor-current feedback	7, 13, 16, 19
	3.4 Motor-current sensors	Drive-current read-back for inferred load.	motor-current feedback	19
4. Tactile & Contact Sensors				
	4.1 Vision-based tactile	Camera-in-gel sensors capturing high-resolution surface contact.	GelSight, Soft-Bubble	13
	4.2 Pressure / tactile arrays	Capacitive or resistive skins giving per-taxel pressure maps.	force-torque/pressure arrays, contact sensors	7, 13
5. External Tracking & Global Localization				
	5.1 Optical motion-capture systems	Infra-red camera networks tracking reflective markers.	VICON, optical marker rigs	3, 13, 14, 19
	5.2 Wearable mocap devices	Marker gloves or body suits for fine human-pose capture.	motion-capture gloves, skeletal/hand markers	13, 19
	5.3 Radio-based positioning	Satellite or UWB transceivers returning global coordinates.	GPS, UWB beacons	3, 4, 16
6. Audio Sensors				
	6.1 Microphones / audio arrays	Mono or array microphones for speech / environmental sound.	microphone audio inputs	3, 13, 19

Table 17 | Structured summary of input sensors used in robotic papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Ground-based mobile robots				
	1.1 Small RC / off-road vehicles	1/10-scale cars, ATVs, skid-steer rovers for field tests	RC cars/ATVs; small off-road vehicles	19
	1.2 Kinematic vehicle models	Bicycle/unicycle point-mass models (simulation-only)	Simulated vehicle agents (kinematic/dynamic)	0
2. Aerial robots				
	2.1 Quadrotors / drones	Four-rotor UAVs with cameras, LiDAR, IMU	Quadrotor UAVs; drones	11, 19
3. Legged & humanoid robots				
	3.1 Simulated legged agents	Classic MuJoCo bodies for RL locomotion	Hopper; HalfCheetah; Walker2d; Ant; Quadruped	11
	3.2 Real quadrupeds & hybrids	Torque-controlled ~12-DoF quadrupeds; wheel-leg hybrids	ANYmal; Unitree A1/Go1; MIT Mini-Cheetah; wheel-leg hybrids	12
	3.3 Humanoids	High-DoF bipeds/humanoids, often with articulated hands	Humanoids/bipeds; SMPL-X mesh; simulated avatars	16
4. Manipulators & end-effectors				
	4.1 Standard 6–7 DoF arms	Fixed-base arms with two-finger or suction grippers	UR5e; Sawyer; other 6–7 DoF arms	2
	4.2 Franka-class agile arms	7-DoF Panda-style arms popular in RL/IL	Franka Emika Panda; Robotiq; suction cups	3
	4.3 Mobile / dual-arm manipulators	One or two arms on a wheeled base (bimanual possible)	Mobile bases with dual arms; mobile manipulators	7, 11
	4.4 Arm + dexterous hand	Arms distinguished by multi-finger hands	Shadow; Allegro; Adroit; LEAP; DeltaHand	16, 18
5. Soft & continuum robots				
	5.1 Continuum / soft manipulators	Deformable backbones, pneumatic/tendon actuation, soft skins & grippers; tensegrity frames	Soft continuum arm; soft gripper; compliant tensegrity structures	1

Table 18 | Structured summary of physical bodies used in robotic papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Direct joint-level outputs				
	1.1 Joint state read-outs	Instantaneous articulated joint positions, orientations, angles, velocities	“joint positions & orientations”; “joint angles”; “joint velocities”; “mesh deformations”	0, 1, 9
	1.2 Joint command signals	Low-level motor targets (torque / position / velocity) that drive joint motion	“joint torque/position commands”; “continuous motor control signals”; “PD control torques”	7, 11, 12, 16, 17
	1.3 Joint motion trajectories	Time-indexed sequences of joint states the robot follows	“motion sequences over time”; “planned joint trajectories”; “optimised 6-DoF trajectories”	0, 1, 17
2. Rigid-body / end-effector pose outputs				
	2.1 6-DoF body poses	Position + orientation of whole robots, cameras or objects	“6-DoF poses”; “rigid transformations”; “UAV 3-D position & orientation”	6, 10, 16
	2.2 End-effector pose + gripper	Cartesian pose of manipulator tip plus gripper open/close state	“6-DoF end-effector pose (x, y, z, r, p, y)”; “gripper_state (open/close)”	4
3. Ground-vehicle / mobile-robot control outputs				
	3.1 Steering & pedal commands	Low-level automotive controls for heading and speed	“steering_angle”; “acceleration/throttle”; “brake”	3
	3.2 Wheel / differential-drive velocities	Body-frame linear & angular velocity commands for wheels/actuators	“linear & angular velocity motor commands”; “wheel/actuator motions”	14
	3.3 Motion trajectories	Pre-planned paths or waypoints for vehicle motion	“robot/vehicle motion trajectories”; “position/orientation updates”	19, 16
4. Aerial-rotorcraft control outputs				
	4.1 Rotor thrust & body-rate commands	Per-rotor thrust/speed or body-rate inputs that place a UAV in 3-D space	“rotor thrust/speed commands”; “collective thrust & body-rate inputs”	6

Table 19 | Structured summary of joint outputs used in robotic papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Continuous Low-Level Actuation				
	1.1 Joint-space commands	Direct numerical inputs to individual joints or actuators, bounded by hardware limits.	joint torques/positions/velocities; high-dimensional joint commands; bounded control inputs; finger-joint configs; parametrised joint trajectories	0, 4, 6, 10, 12, 14, 18
	1.2 Vehicle / body dynamics commands	Low-level controls that change a mobile base, ground-vehicle or aerial body state.	steering angle; throttle / acceleration; braking; linear & angular velocity; body-rate thrust; speed/direction for locomotion; lane-keeping	0, 1, 7, 10, 12, 13, 15
2. Mid-Level Pose & Trajectory Control				
	2.1 End-effector & gripper pose	6-DoF goals and time-parameterised trajectories for arms, grippers or aerial manipulators.	continuous 6-DoF poses; pose deltas (Δx , Δy , Δz , $\Delta roll$, $\Delta pitch$, Δyaw); gripper open/close; gripper width/force; grasp trajectories	2, 6, 9, 10, 12, 14, 18
	2.2 Base / waypoint trajectories	Desired paths, way-points or velocity profiles for the robot body or ego vehicle.	waypoint/path-goal selection; future trajectory sequences; base linear & angular velocity commands; lane-change/merge trajectories	0, 1, 7, 10, 15, 19
3. High-Level Discrete Skills & Behaviour Primitives				
	3.1 Manipulation skills	Object-centred primitives that parameterise targets, forces or object states.	grasp/pick; place/drop; push/pull; rotate/open/close; part deformation	0, 10, 18, 19
	3.2 Locomotion & navigation skills	Discrete moves or gait switches for repositioning the whole robot.	move_forward/stop; turn_left/turn_right; gait switch; lane keeping/change; overtaking/merging; "go to X"	0, 1, 10, 15, 19
	3.3 Interaction & instruction skills	Multimodal actions expressed through gesture, speech or scene edits.	gesture actions; speech actions; instructional guidance; scene editing commands	0

Table 20 | Structured summary of action space used in robotic papers.

Category	Sub-category	What is covered	Typical examples	Cluster
1. Autonomous-driving & Mobile-vehicle scenes				
	1.1 On-road urban / suburban / rural driving	Real or simulated road networks with traffic, road rules, and weather variation.	urban roads; highways; intersections; traffic lights/signs; ...	1, 2, 6, 9, 12, 13, 19
	1.2 Off-road, cross-country & planetary terrain	Structured or unstructured natural terrains requiring ground-robot locomotion.	uneven ground; sand; gravel; snow; ...	11
2. Manipulation workspaces				
	2.1 Basic household tabletop	Small cluttered indoor bench for reach-scale manipulation.	cluttered tabletop; household objects; articulated fixtures; ...	0
	2.2 Kitchen & household benchmark suites	Standardised kitchen/tabletop scenes from RLBench, Meta-World, FrankaKitchen, Habitat, Ravens, etc.	kitchen counters; RLBench station; FrankaKitchen; ...	14, 17
	2.3 Assembly & insertion tables	Contact-rich assembly surfaces with precisely shaped parts.	assembly workspace; peg-hole joints; plug-socket joints; ...	18
	2.4 Shared lab / industrial workcells	Planar or 3-D manipulation bays in labs or factories, often human-robot shared.	lab work surfaces; human-robot zones; static & dynamic obstacles; ...	10
3. Embodied navigation & Scene-understanding worlds				
	3.1 Multi-room home / office interiors	Photorealistic or simulated domestic & office floorplans for navigation and light manipulation.	apartments; offices; corridors; dynamic changes; ...	7
	3.2 Large-scale mixed indoor-outdoor simulators	Dynamic 3-D worlds with physics for point-goal, exploration, or social-navigation tasks.	rooms; mazes; multiple agents; partial observability; ...	15
	3.3 Object-rich mixed-reality scene sets	Real + synthetic household, lab, or industrial spaces emphasising clutter & diversity.	household rooms; industrial floors; cluttered indoor scenes; ...	4, 5, 8
4. Physics-centric control benchmarks				
	4.1 Classic locomotion & manipulation suites	Widely-used control benchmarks with domain-randomised dynamics.	MuJoCo tasks; IsaacGym walkers; Robotarium arena; ...	3
	4.2 High-fidelity multi-physics platforms	Environments that model contact, fluids, deformables & human interaction in indoor/outdoor scenes.	rigid-body scenes; deformable objects; fluid interaction; humans; ...	16

Table 21 | Structured summary of environment used in robotic papers.




Figure 5 | Trend Visualization of Computer Vision Research




Figure 6 | Trend Visualization of Robotics Research




Figure 7 | Trend Visualization of NLP Research




Figure 8 | Trend Visualization of Machine Learning Research

References

- [1] S. Addepalli, K. Bhogale, P. Dey, and R. V. Babu. Towards efficient and effective self-supervised learning of visual representations. In *European Conference on Computer Vision*, 2024.
- [2] C. Agia, R. Sinha, J. Yang, Z. Cao, R. Antonova, M. Pavone, and J. Bohg. Unpacking failure modes of generative policies: Runtime monitoring of consistency and progress. In *Conference on Robot Learning*, 2024.
- [3] A. Ajith, M. Xia, A. Chevalier, T. Goyal, D. Chen, and T. Gao. Litsearch: A retrieval benchmark for scientific literature search. *arXiv preprint arXiv:2407.18940*, 2024.
- [4] A. Aubret, C. Teulière, and J. Triesch. Self-supervised visual learning from interactions with objects. In *European Conference on Computer Vision*, 2024.
- [5] M. Augustin, A. Meinke, and M. Hein. Adversarial robustness on in- and out-distribution improves explainability. In *European Conference on Computer Vision*, 2024.
- [6] J. Baek, S. K. Jauhar, S. Cucerzan, and S. J. Hwang. Researchagent: Iterative research idea generation over scientific literature with large language models. *arXiv preprint arXiv:2404.07738*, 2024.
- [7] G. Balakrishnan, Y. Xiong, W. Xia, and P. Perona. Towards causal benchmarking of bias in face analysis algorithms. In *European Conference on Computer Vision*, 2024.
- [8] H. Bansal, A. Hosseini, R. Agarwal, V. Q. Tran, and M. Kazemi. Smaller, weaker, yet better: Training llm reasoners via compute-optimal sampling. In *International Conference on Learning Representations*, 2025.
- [9] L. Barcellona, A. Zadaianchuk, D. Allegro, S. Papa, S. Ghidoni, and E. Gavves. Dream to manipulate: Compositional world models empowering robot imitation learning with imagination. In *International Conference on Learning Representations*, 2025.
- [10] S. Belkhale, T. Ding, T. Xiao, P. Sermanet, Q. Vuong, J. Tompson, Y. Chebotar, D. Dwibedi, and D. Sadigh. Rt-h: Action hierarchies using language. In *Robotics: Science and Systems*, 2024.
- [11] K. Bhardwaj, N. P. Pandey, S. Priyadarshi, V. Ganapathy, S. Kadambi, R. Esteves, S. Borse, P. Whatmough, R. Garrepalli, M. V. Baalen, H. Teague, and M. Nagel. Sparse high rank adapters. In *Neural Information Processing Systems*, 2024.
- [12] R. Bommasani. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.
- [13] S. Brahmbhatt, C. Tang, C. D. Twigg, C. C. Kemp, and J. Hays. Contactpose: A dataset of grasps with object contact and hand pose. In *European Conference on Computer Vision*, 2024.
- [14] T. Broedermann, D. Brüggemann, C. Sakaridis, K. Ta, O. Liagouris, J. Corkill, and L. V. Gool. Muses: The multi-sensor semantic perception dataset for driving under uncertainty. In *European Conference on Computer Vision*, 2024.
- [15] A. Brown, C.-Y. Fu, O. Parkhi, T. L. Berg, and A. Vedaldi. End-to-end visual editing with a generatively pre-trained artist. In *European Conference on Computer Vision*, 2024.
- [16] U. Buchler, B. Brattoli, and B. Ommer. Improving spatiotemporal self-supervision by deep reinforcement learning. In *European Conference on Computer Vision*, 2024.

- [17] J. Cao, Z. Gan, Y. Cheng, L. Yu, Y.-C. Chen, and J. Liu. Behind the scene: Revealing the secrets of pre-trained vision-and-language models. In *European Conference on Computer Vision*, 2024.
- [18] S. Cen, J. Mei, K. Goshvadi, H. Dai, T. Yang, S. Yang, D. Schuurmans, Y. Chi, and B. Dai. Value-incentivized preference optimization: A unified approach to online and offline rlhf. In *International Conference on Learning Representations*, 2025.
- [19] L. Chambon, E. Zablocki, M. Chen, F. Bartoccioni, P. Pérez, and M. Cord. Pointbev: A sparse approach for bev predictions. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [20] Y. Chang, X. Wang, J. Wang, Y. Wu, L. Yang, K. Zhu, H. Chen, X. Yi, C. Wang, Y. Wang, et al. A survey on evaluation of large language models. *ACM transactions on intelligent systems and technology*, 15(3):1–45, 2024.
- [21] P. Chattopadhyay, Y. Balaji, and J. Hoffman. Learning to balance specificity and invariance for in and out of domain generalization. In *European Conference on Computer Vision*, 2024.
- [22] B. Chen, D. M. Monsó, Y. Du, M. Simchowitz, R. Tedrake, and V. Sitzmann. Diffusion forcing: Next-token prediction meets full-sequence diffusion. In *Neural Information Processing Systems*, 2024.
- [23] C. Chen, F. Tung, N. Vedula, and G. Mori. Constraint-aware deep neural network compression. In *European Conference on Computer Vision*, 2024.
- [24] D. Z. Chen, A. X. Chang, and M. Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *European Conference on Computer Vision*, 2024.
- [25] H. Chen*, S. Xie, S.-N. Lim, and A. Shrivastava. Fast encoding and decoding for implicit video representation. In *European Conference on Computer Vision*, 2024.
- [26] M. Chen, Y. Li, Y. Yang, S. Yu, B. Lin, and X. He. Automanual: Constructing instruction manuals by llm agents via interactive environmental learning. In *Neural Information Processing Systems*, 2024.
- [27] Q. Chen, A. Walsman, M. Memmel, K. Mo, A. Fang, D. Fox, and A. Gupta. Urdformer: A pipeline for constructing articulated simulation environments from real-world images. In *Robotics: Science and Systems*, 2024.
- [28] S. Chen, P.-L. Guhur, M. Tapaswi, C. Schmid, and I. Laptev. Learning from unlabeled 3d environments for vision-and-language navigation. In *European Conference on Computer Vision*, 2024.
- [29] X. Chen and J. Yang. X-fi: A modality-invariant foundation model for multimodal human sensing. In *International Conference on Learning Representations*, 2025.
- [30] X. Chen, J. Djolonga, P. Padlewski, B. Mustafa, S. Changpinyo, J. Wu, C. R. Ruiz, S. Goodman, X. Wang, Y. Tay, S. Shakeri, M. Dehghani, D. Salz, M. Lucic, M. Tschannen, A. Nagrani, H. Hu, M. Joshi, B. Pang, C. Montgomery, P. Pietrzyk, M. Ritter, A. Piergiovanni, M. Minderer, F. Pavetic, A. Waters, G. Li, I. Alabdulmohsin, L. Beyer, J. Amelot, K. Lee, A. P. Steiner, Y. Li, D. Keysers, A. Arnab, Y. Xu, K. Rong, A. Kolesnikov, M. Seyedhosseini, A. Angelova, X. Zhai, N. Houlsby, and R. Soricut. On scaling up a multilingual vision and language model. In *Conference on Computer Vision and Pattern Recognition*, 2024.

- [31] Y. Chen, T. Wang, T. Wu, X. Pan, K. Jia*, and Z. Liu. Comboverse: Compositional 3d assets creation using spatially-aware diffusion guidance. In *European Conference on Computer Vision*, 2024.
- [32] Z. Chen, J. Wu, W. Wang, W. Su, G. Chen, S. Xing, M. Zhong, Q. Zhang, X. Zhu, L. Lu, B. Li, P. Luo, T. Lu, Y. Qiao, and J. Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [33] Z. Chen, D. Chen, R. Sun, W. Liu, and C. Gan. Scaling autonomous agents via automatic reward modeling and planning. In *International Conference on Learning Representations*, 2025.
- [34] X. Cheng, Y. Ji, J. Chen, R. Yang, G. Yang, and X. Wang. Expressive whole-body control for humanoid robots. In *Robotics: Science and Systems*, 2024.
- [35] X. Cheng, J. Li, S. Yang, G. Yang, and X. Wang. Open-television: Teleoperation with immersive active visual feedback. In *Conference on Robot Learning*, 2024.
- [36] C. Chi, Z. Xu, C. Pan, E. Cousineau, B. Burchfiel, S. Feng, R. Tedrake, and S. Song. Universal manipulation interface: In-the-wild robot teaching without in-the-wild robots. In *Robotics: Science and Systems*, 2024.
- [37] H. Cho, M. Kato, Y. Sakai, and N. Inoue. Revisiting in-context learning inference circuit in large language models. In *International Conference on Learning Representations*, 2025.
- [38] S. Choi, S. Yang, S. Choi, and S. Yun. Improving test-time adaptation via shift-agnostic weight regularization and nearest source prototypes. In *European Conference on Computer Vision*, 2024.
- [39] W. Choi, J. Park, S. Ahn, D. Lee, and H. Woo. Nesyc: A neuro-symbolic continual learner for complex embodied tasks in open domains. In *International Conference on Learning Representations*, 2025.
- [40] Y. Chow, G. Tennenholtz, I. Gur, V. Zhuang, B. Dai, A. Kumar, R. Agarwal, S. Thiagarajan, C. Boutilier, and A. Faust. Inference-aware fine-tuning for best-of-n sampling in large language models. In *International Conference on Learning Representations*, 2025.
- [41] X. Chu, J. Su, B. Zhang, and C. Shen. Visionllama: A unified llama backbone for vision tasks. In *European Conference on Computer Vision*, 2024.
- [42] A. Curtis, N. Kumar, J. Cao, T. Lozano-Pérez, and L. P. Kaelbling. Trust the proc3s: Solving long-horizon robotics problems with llms and constraint satisfaction. In *Conference on Robot Learning*, 2024.
- [43] A. Curtis, G. Matheos, N. Gothoskar, V. Mansinghka, J. B. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling. Partially observable task and motion planning with uncertainty and risk awareness. In *Robotics: Science and Systems*, 2024.
- [44] P. Ding, H. Zhao, W. Zhang, W. Song, M. Zhang, S. Huang, N. Yang, and D. Wang. Quar-vla: Vision-language-action model for quadruped robots. In *European Conference on Computer Vision*, 2024.
- [45] L. Doorenbos, R. Sznitman, and P. Márquez-Neila. Data invariants to understand unsupervised out-of-distribution detection. In *European Conference on Computer Vision*, 2024.
- [46] R. Doshi, H. R. Walke, O. Mees, S. Dasari, and S. Levine. Scaling cross-embodied learning: One policy for manipulation, navigation, locomotion and aviation. In *Conference on Robot Learning*, 2024.

- [47] J. Duan, W. Yuan, W. Pumacay, Y. R. Wang, K. Ehsani, D. Fox, and R. Krishna. Manipulate-anything: Automating real-world robots using vision-language models. In *Conference on Robot Learning*, 2024.
- [48] A. Eldesokey and P. Wonka. Build-a-scene: Interactive 3d layout control for diffusion-based image generation. In *International Conference on Learning Representations*, 2025.
- [49] E. Elhamifar and D. Huynh. Self-supervised multi-task procedure learning from instructional videos. In *European Conference on Computer Vision*, 2024.
- [50] J. Fang, A. Shafiee, H. Abdel-Aziz, D. Thorsley, G. Georgiadis, and J. H. Hassoun. Post-training piecewise linear quantization for deep neural networks. In *European Conference on Computer Vision*, 2024.
- [51] R. Feng, J. Hu, W. Xia, T. Gao, A. Shen, Y. Sun, B. Fang, and D. Hu. Anytouch: Learning unified static-dynamic representation across multiple visuo-tactile sensors. In *International Conference on Learning Representations*, 2025.
- [52] D. J. Foster, A. Block, and D. Misra. Is behavior cloning all you need? understanding horizon in imitation learning. In *Neural Information Processing Systems*, 2024.
- [53] Y. Fu, S. Liu, A. Kulkarni, J. Kautz, A. A. Efros, and X. Wang. Colmap-free 3d gaussian splatting. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [54] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn. Humanplus: Humanoid shadowing and imitation from humans. In *Conference on Robot Learning*, 2024.
- [55] I. O. Gallegos, R. A. Rossi, J. Barrow, M. M. Tanjim, S. Kim, F. Dornoncourt, T. Yu, R. Zhang, and N. Ahmed. Bias and fairness in large language models: A survey. *Computational Linguistics*, 50: 1097–1179, 2023.
- [56] X. Gao, S. Dong, Y. He, Q. Wang, and Y. Gong. Beyond prompt learning: Continual adapter for efficient rehearsal-free continual learning. In *European Conference on Computer Vision*, 2024.
- [57] H. Geng, S. Wei, C. Deng, B. Shen, H. Wang, and L. Guibas. Sage: Bridging semantic and actionable parts for generalizable articulated-object manipulation under language instructions. In *Robotics: Science and Systems*, 2024.
- [58] D. Ghosh, H. R. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu, J. Luo, Y. L. Tan, L. Y. Chen, Q. Vuong, T. Xiao, P. R. Sanketi, D. Sadigh, C. Finn, and S. Levine. Octo: An open-source generalist robot policy. In *Robotics: Science and Systems*, 2024.
- [59] A. Gomez-Villa, D. Goswami, K. Wang, A. Bagdanov, B. Twardowski, and J. van de Weijer. Exemplar-free continual representation learning via learnable drift compensation. In *European Conference on Computer Vision*, 2024.
- [60] A. Goyal, V. Blukis, J. Xu, Y. Guo, Y.-W. Chao, and D. Fox. Rvt-2: Learning precise manipulation from few demonstrations. In *Robotics: Science and Systems*, 2024.
- [61] R. Grandia, E. Knoop, M. A. Hopkins, G. Wiedebach, J. Bishop, S. Pickles, D. Müller, and M. Bächer. Design and control of a bipedal robotic character. In *Robotics: Science and Systems*, 2024.
- [62] X. Gu and M. Krenn. Forecasting high-impact research topics via machine learning on evolving knowledge graphs. *Machine Learning: Science and Technology*, 6(2):025041, 2025.

- [63] X. Gu, Y. Guo, Z. Li, J. Qiu, Q. Dou, Y. Liu, B. Lo, and G.-Z. Yang. Tackling long-tailed category distribution under domain shifts. In *European Conference on Computer Vision*, 2024.
- [64] X. Gu, Y.-J. Wang, X. Zhu, C. Shi, Y. Guo, Y. Liu, and J. Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. In *Robotics: Science and Systems*, 2024.
- [65] Y. Gui, Y. Jin, and Z. Ren. Conformal alignment: Knowing when to trust foundation models with guarantees. In *Neural Information Processing Systems*, 2024.
- [66] Z. Guo and T. Jin. Smoothing the shift: Towards stable test-time adaptation under complex multimodal noises. In *International Conference on Learning Representations*, 2025.
- [67] J. Günster, P. Liu, J. Peters, and D. Tateo. Handling long-term safety and uncertainty in safe reinforcement learning. In *Conference on Robot Learning*, 2024.
- [68] G. Han*, J. Hur, J. Choi, and J. Kim*. Learning neural deformation representation for 4d dynamic shape generation. In *European Conference on Computer Vision*, 2024.
- [69] M. Han, Y. Zhu, S.-C. Zhu, Y. N. Wu, and Y. Zhu. Interpret: Interactive predicate learning from language feedback for generalizable task planning. In *Robotics: Science and Systems*, 2024.
- [70] N. Hansen, J. S. V, V. Sobal, Y. LeCun, X. Wang, and H. Su. Hierarchical world models as visual whole-body humanoid controllers. In *International Conference on Learning Representations*, 2025.
- [71] D. Harwath, A. Recasens, D. Suris, G. Chuang, A. Torralba, and J. Glass. Jointly discovering visual objects and spoken words from raw sensory input. In *European Conference on Computer Vision*, 2024.
- [72] S. Hayou, N. Ghosh, and B. Yu. The impact of initialization on lora finetuning dynamics. In *Neural Information Processing Systems*, 2024.
- [73] H. He, J. B. Li, X. Jiang, and H. Miller. Smt: Fine-tuning large language models with sparse matrices. In *International Conference on Learning Representations*, 2025.
- [74] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. M. Kitani, C. Liu, and G. Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *Conference on Robot Learning*, 2024.
- [75] T. He, C. Zhang, W. Xiao, G. He, C. Liu, and G. Shi. Agile but safe: Learning collision-free high-speed legged locomotion. In *Robotics: Science and Systems*, 2024.
- [76] Y. He, G. Huang, P. Feng, Y. Lin, Y. Zhang, H. Li, et al. Pasa: An llm agent for comprehensive academic paper search. *arXiv preprint arXiv:2501.10120*, 2025.
- [77] S. Hecker, D. Dai, and L. V. Gool. End-to-end learning of driving models with surround-view cameras and route planners. In *European Conference on Computer Vision*, 2024.
- [78] H. Hu, S. Mirchandani, and D. Sadigh. Imitation bootstrapped reinforcement learning. In *Robotics: Science and Systems*, 2024.
- [79] J. Y.-C. Hu, W.-P. Wang, A. Gilani, C. Li, Z. Song, and H. Liu. Fundamental limits of prompt tuning transformers: Universality, capacity and efficiency. In *International Conference on Learning Representations*, 2025.

- [80] S. Hu, C. Lu, and J. Clune. Automated design of agentic systems. In *International Conference on Learning Representations*, 2025.
- [81] Y. Hu, S. Chai, Z. Yang, J. Qian, K. Li, W. Shao, H. Zhang, W. Xu, and Q. Liu. Solving motion planning tasks with a scalable generative model. In *European Conference on Computer Vision*, 2024.
- [82] B. Huang, Y. Wang, X. Yang, Y. Luo, and Y. Li. 3d-vitac: Learning fine-grained manipulation with visuo-tactile sensing. In *Conference on Robot Learning*, 2024.
- [83] J. Huang and K. C.-C. Chang. Towards reasoning in large language models: A survey. *ArXiv*, abs/2212.10403, 2022.
- [84] W. Huang, C. Wang, Y. Li, R. Zhang, and L. Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. In *Conference on Robot Learning*, 2024.
- [85] Y. Huang, W. Zheng, B. Zhang, J. Zhou, and J. Lu. Selfocc: Self-supervised vision-based 3d occupancy prediction. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [86] Z. Huang, T. Tang, S. Chen, S. Lin, Z. Jie, L. Ma, G. Wang, and X. Liang*. Making large language models better planners with reasoning-decision alignment. In *European Conference on Computer Vision*, 2024.
- [87] J. Hübötter, S. Bongni, I. Hakimi, and A. Krause. Efficiently learning at test-time: Active fine-tuning of llms. In *International Conference on Learning Representations*, 2025.
- [88] T. Ingebrand, A. Thorpe, and U. Topcu. Zero-shot transfer of neural odes. In *Neural Information Processing Systems*, 2024.
- [89] S. Jain, E. S. Lubana, K. Oksuz, T. Joy, P. Torr, A. Sanyal, and P. K. Dokania. What makes and breaks safety fine-tuning? a mechanistic study. In *Neural Information Processing Systems*, 2024.
- [90] U. Jain, L. Weihs, E. Kolve, A. Farhadi, S. Lazebnik, A. Kembhavi, and A. Schwing. A cordial sync: Going beyond marginal policies for multi-agent embodied tasks. In *European Conference on Computer Vision*, 2024.
- [91] V. Jain, M. Attarian, N. J. Joshi, A. Wahid, D. Driess, Q. Vuong, P. R. Sanketi, P. Sermanet, S. Welker, C. Chan, I. Gilitschenski, Y. Bisk, and D. Dwibedi. Vid2robot: End-to-end video-conditioned policy learning with cross-attention transformers. In *Robotics: Science and Systems*, 2024.
- [92] B. Jiang, X. Chen, C. Zhang, F. Yin, Z. Li, G. Yu, and J. Fan*. Motionchain: Conversational motion controllers via multimodal prompts. In *European Conference on Computer Vision*, 2024.
- [93] K. Jiang, J. Huang, W. Xie, J. Lei, Y. Li, L. Shao, and S. Lu. Da-bev: Unsupervised domain adaptation for bird’s eye view perception. In *European Conference on Computer Vision*, 2024.
- [94] H. Jin, H. Jiang, H. Tan, K. Zhang, S. Bi, T. Zhang, F. Luan, N. Snavely, and Z. Xu. Lvsm: A large view synthesis model with minimal 3d inductive bias. In *International Conference on Learning Representations*, 2025.
- [95] Y. Ju, K. Hu, G. Zhang, G. Zhang, M. Jiang, and H. Xu*. Robo-abc: Affordance generalization beyond categories via semantic correspondence for robot manipulation. In *European Conference on Computer Vision*, 2024.

- [96] J. Kang, L. Karlinsky, H. Luo, Z. Wang, J. A. Hansen, J. R. Glass, D. D. Cox, R. Panda, R. Feris, and A. Ritter. Self-moe: Towards compositional large language models with self-specialized experts. In *International Conference on Learning Representations*, 2025.
- [97] U. Katz, M. Levy, and Y. Goldberg. Knowledge navigator: Llm-guided browsing framework for exploratory search in scientific literature. *arXiv preprint arXiv:2408.15836*, 2024.
- [98] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki. 3d diffuser actor: Policy diffusion with 3d scene representations. In *Conference on Robot Learning*, 2024.
- [99] T. Kehrenberg, M. Bartlett, O. Thomas, and N. Quadrianto. Null-sampling for interpretable and fair representations. In *European Conference on Computer Vision*, 2024.
- [100] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. P. Foster, P. R. Sanketi, Q. Vuong, T. Kollar, B. Burchfiel, R. Tedrake, D. Sadigh, S. Levine, P. Liang, and C. Finn. Openvla: An open-source vision-language-action model. In *Conference on Robot Learning*, 2024.
- [101] S. Kim*, B. Jeong, D. Kim, and S. Kwak*. Efficient and versatile robust fine-tuning of zero-shot models. In *European Conference on Computer Vision*, 2024.
- [102] T. Kim, Y. Kwon, J. Lee, T. Kim, and S. Ha. Cprune: Compiler-informed model pruning for efficient target-aware dnn execution. In *European Conference on Computer Vision*, 2024.
- [103] J. Krantz and S. Lee. Sim-2-sim transfer for vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, 2024.
- [104] J. Krantz, E. Wijmans, A. Majumdar, D. Batra, and S. Lee. Beyond the nav-graph: Vision-and-language navigation in continuous environments. In *European Conference on Computer Vision*, 2024.
- [105] M. Krenn and A. Zeilinger. Predicting research trends with semantic and neural networks with an application in quantum physics. *Proceedings of the National Academy of Sciences*, 117(4): 1910–1916, 2020.
- [106] M. Krenn, L. Buffoni, B. Coutinho, S. Eppel, J. G. Foster, A. Gritsevskiy, H. Lee, Y. Lu, J. P. Moutinho, N. Sanjabi, et al. Forecasting the future of artificial intelligence with machine learning-based link prediction in an exponentially growing knowledge network. *Nature Machine Intelligence*, 5(11):1326–1335, 2023.
- [107] N. Kumar, T. Silver, W. McClinton, L. Zhao, S. Proulx, T. Lozano-Pérez, L. P. Kaelbling, and J. L. Barry. Practice makes perfect: Planning to learning skill parameter policies. In *Robotics: Science and Systems*, 2024.
- [108] W. Kuo, F. Bertsch, W. Li, A. Piergiovanni, M. Saffar, and A. Angelova. Findit: Generalized localization with natural language queries. In *European Conference on Computer Vision*, 2024.
- [109] C.-M. Lai, H.-C. Wang, P.-C. Hsieh, Y.-C. F. Wang, M.-H. Chen, and S.-H. Sun. Diffusion-reward adversarial imitation learning. In *Neural Information Processing Systems*, 2024.
- [110] M. Le, C. Nguyen, H. Nguyen, Q. Tran, T. Le, and N. Ho. Revisiting prefix-tuning: Statistical benefits of reparameterization among prompts. In *International Conference on Learning Representations*, 2025.

- [111] S. Lee, S. H. Park, S. Kim, and M. Seo. Aligning to thousands of preferences via system message generalization. In *Neural Information Processing Systems*, 2024.
- [112] F. Li, R. Zhang, H. Zhang, Y. Zhang, B. Li, W. Li, Z. MA, and C. Li. Llava-interleave: Tackling multi-image, video, and 3d in large multimodal models. In *International Conference on Learning Representations*, 2025.
- [113] K. Li, J. Wang, L. Yang, C. Lu, and B. Dai. Semgrasp: Semantic grasp generation via language aligned discretization. In *European Conference on Computer Vision*, 2024.
- [114] P. Li, Z. Wang, X. Zhang, R. Zhang, L. Jiang, P. Wang, and Y. Zhou. Scitopic: Enhancing topic discovery in scientific literature through advanced llm. *arXiv preprint arXiv:2508.20514*, 2025.
- [115] S. Li, J. Huang, J. Zhuang, Y. Shi, X. Cai, M. Xu, X. Wang, L. Zhang, G. Ke, and H. Cai. Scilitllm: How to adapt llms for scientific literature understanding. *arXiv preprint arXiv:2408.15545*, 2024.
- [116] W. Li, P. Wan, P. Wang, J. Li, Y. Zhou, and P. Liu*. Benerf:neural radiance fields from a single blurry image and event stream. In *European Conference on Computer Vision*, 2024.
- [117] X. Li, K. Hsu, J. Gu, O. Mees, K. Pertsch, H. R. Walke, C. Fu, I. Lunawat, I. Sieh, S. Kirmani, S. Levine, J. Wu, C. Finn, H. Su, Q. Vuong, and T. Xiao. Evaluating real-world robot manipulation policies in simulation. In *Conference on Robot Learning*, 2024.
- [118] X. Li, C. Herrmann, K. C. Chan, Y. Li, D. Sun, C. Ma, and M.-H. Yang. A simple approach to unifying diffusion-based conditional generation. In *International Conference on Learning Representations*, 2025.
- [119] Y. Li, Y. Deng, J. Zhang, J. Jang, M. Memmel, C. R. Garrett, F. Ramos, D. Fox, A. Li, A. Gupta, and A. Goyal. Hamster: Hierarchical action models for open-world robot manipulation. In *International Conference on Learning Representations*, 2025.
- [120] Z. Li, W. Wang, H. Li, E. Xie, C. Sima, T. Lu, Y. Qiao, and J. Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European Conference on Computer Vision*, 2024.
- [121] J. Liang, L. Jiang, and A. Hauptmann. Simaug: Learning robust representations from simulation for trajectory prediction. In *European Conference on Computer Vision*, 2024.
- [122] J. Liang, F. Xia, W. Yu, A. Zeng, M. Attarian, M. B. Villalonga, M. Bennice, A. Bewley, A. Dost-mohamed, C. Fu, N. Gileadi, M. Giustina, K. Gopalakrishnan, L. Hasenclever, J. Humplik, J. Hsu, N. J. Joshi, B. Jyenis, J. C. Kew, S. Kirmani, T.-W. E. Lee, K.-H. Lee, A. H. Michaely, J. Moore, K. Oslund, D. Rao, A. Z. Ren, B. Tabanpour, Q. Vuong, A. Wahid, T. Xiao, Y. Xu, V. Zhuang, P. Xu, E. Frey, K. Caluwaerts, T. Zhang, B. Ichter, J. Tompson, L. Takayama, V. Vanhoucke, I. Shafran, M. Mataric, D. Sadigh, N. Heess, K. Rao, N. Stewart, J. Tan, and C. Parada. Learning to learn faster from human feedback with language model predictive control. In *Robotics: Science and Systems*, 2024.
- [123] M. Liang, B. Yang, S. Wang, and R. Urtasun. Deep continuous fusion for multi-sensor 3d object detection. In *European Conference on Computer Vision*, 2024.
- [124] W. Liang, Y. Zhang, H. Cao, B. Wang, D. Y. Ding, X. Yang, K. Vodrahalli, S. He, D. S. Smith, Y. Yin, et al. Can large language models provide useful feedback on research papers? a large-scale empirical analysis. *NEJM AI*, 1(8):AIoa2400196, 2024.

- [125] Y. Liang, K. Ellis, and J. Henriques. Rapid motor adaptation for robotic manipulator arms. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [126] Y. Liang, X. Fang, H. Chen, and Y. Wang. Linear multistep solver distillation for fast sampling of diffusion models. In *International Conference on Learning Representations*, 2025.
- [127] Z. Liang, Y. Mu, H. Ma, M. Tomizuka, M. Ding, and P. Luo. Skilldiffuser: Interpretable hierarchical planning via skill abstractions in diffusion-based task execution. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [128] B. LiChen, S. Shao, zikai zhou, Z. Qi, zhiqiang xu, H. Xiong, and Z. Xie. Zigzag diffusion sampling: Diffusion models can self-improve via self-reflection. In *International Conference on Learning Representations*, 2025.
- [129] C.-H. Lin, S. Gao, J. S. Smith, A. Patel, S. Tuli, Y. Shen, H. Jin, and Y.-C. Hsu. Modegpt: Modular decomposition for large language model compression. In *International Conference on Learning Representations*, 2025.
- [130] H. Lin, J. Cho, A. Zala, and M. Bansal. Ctrl-adapter: An efficient and versatile framework for adapting diverse controls to any diffusion model. In *International Conference on Learning Representations*, 2025.
- [131] Z. Lin, X. Peng, P. Cong, G. Zheng, Y. Sun, Y. HOU, X. Zhu, S. Yang, and Y. Ma*. Wildrefer: 3d object localization in large-scale dynamic scenes with multi-modal visual data and natural language. In *European Conference on Computer Vision*, 2024.
- [132] V. Lingam, A. T. Neerkaje, A. Vavre, A. Shetty, G. K. Gudur, J. Ghosh, E. Choi, A. Dimakis, A. Bojchevski, and S. Sanghavi. Svft: Parameter-efficient fine-tuning with singular vectors. In *Neural Information Processing Systems*, 2024.
- [133] H. Liu, Y. Chen, H. Wang, Z. Yang, T. Li, J. Zeng, L. Chen, H. Li, and L. Wang. Fully sparse 3d occupancy prediction. In *European Conference on Computer Vision*, 2024.
- [134] H. Liu, Y. Zhang, V. Betala, E. Zhang, J. Liu, C. Ding, and Y. Zhu. Multi-task interactive robot fleet learning with visual world models. In *Conference on Robot Learning*, 2024.
- [135] M. Liu, Z. Chen, X. Cheng, Y. Ji, R.-Z. Qiu, R. Yang, and X. Wang. Visual whole-body control for legged loco-manipulation. In *Conference on Robot Learning*, 2024.
- [136] N. Liu, S. Li, Y. Du, A. Torralba, and J. B. Tenenbaum. Compositional visual generation with composable diffusion models. In *European Conference on Computer Vision*, 2024.
- [137] S. Liu*, H. Cheng, H. Liu, H. Zhang, F. Li, T. Ren, X. Zou, J. Yang, H. Su, J. Zhu, L. Zhang, J. Gao, and C. Li*. Llava-plus: Learning to use tools for creating multimodal agents. In *European Conference on Computer Vision*, 2024.
- [138] S. Liu, J. Nam, A. Campbell, H. Stark, Y. Xu, T. Jaakkola, and R. Gomez-Bombarelli. Think while you generate: Discrete diffusion with planned denoising. In *International Conference on Learning Representations*, 2025.
- [139] S. Liu, L. Wu, B. Li, H. Tan, H. Chen, Z. Wang, K. Xu, H. Su, and J. Zhu. Rdt-1b: a diffusion foundation model for bimanual manipulation. In *International Conference on Learning Representations*, 2025.

- [140] W. Liu, Z. Liu, L. Paull, A. Weller, and B. Schölkopf. Structural causal 3d reconstruction. In *European Conference on Computer Vision*, 2024.
- [141] Y. Liu, Y. Zhang, Y. Wang, F. Hou, J. Yuan, J. Tian, Y. Zhang, Z. Shi, J. Fan, and Z. He. A survey of visual transformers. *IEEE transactions on neural networks and learning systems*, 2023.
- [142] Z. Liu, M. Lu, S. Zhang, B. Liu, H. Guo, Y. Yang, J. Blanchet, and Z. Wang. Provably mitigating overoptimization in rlhf: Your sft loss is implicitly an adversarial regularizer. In *Neural Information Processing Systems*, 2024.
- [143] V. S. Lokhande, A. K. Akash, S. N. Ravi, and V. Singh. Fairalm: Augmented lagrangian method for training fair models with little regret. In *European Conference on Computer Vision*, 2024.
- [144] J. Long, W. Yu, Q. Li, Z. Wang, D. Lin, and J. Pang. Learning h-infinity locomotion control. In *Conference on Robot Learning*, 2024.
- [145] C. Lu, C. Lu, R. T. Lange, J. Foerster, J. Clune, and D. Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *arXiv preprint arXiv:2408.06292*, 2024.
- [146] T. G. W. Lum, M. Matak, V. Makoviychuk, A. Handa, A. Allshire, T. Hermans, N. D. Ratliff, and K. V. Wyk. Dextrah-g: Pixels-to-action dexterous arm-hand grasping with geometric fabrics. In *Conference on Robot Learning*, 2024.
- [147] G. Luo, T. Darrell, O. Wang, D. B. Goldman, and A. Holynski. Readout guidance: Learning control from diffusion features. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [148] J. Luo, T. Ding, K. H. R. Chan, D. Thaker, A. Chattopadhyay, C. Callison-Burch, and R. Vidal. Pace: Parsimonious concept engineering for large language models. In *Neural Information Processing Systems*, 2024.
- [149] X. Luo, A. Recharadt, G. Sun, K. K. Nejad, F. Yáñez, B. Yilmaz, K. Lee, A. O. Cohen, V. Borghesani, A. Pashkov, et al. Large language models surpass human experts in predicting neuroscience results. *Nature human behaviour*, 9(2):305–315, 2025.
- [150] J. Lyu, M. Yan, Z. Qiao, R. Liu, X. Ma, D. Ye, J.-W. Yang, Z. Lu, and X. Li. Cross-domain offline policy adaptation with optimal transport and dataset constraint. In *International Conference on Learning Representations*, 2025.
- [151] K. Lyu, H. Zhao, X. Gu, D. Yu, A. Goyal, and S. Arora. Keeping llms aligned after fine-tuning: The crucial role of prompt templates. In *Neural Information Processing Systems*, 2024.
- [152] X. Ma, S. Patidar, I. Haughton, and S. James. Hierarchical diffusion policy for kinematics-aware multi-task robotic manipulation. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [153] Y. Ma, Z. Song, Y. Zhuang, J. Hao, and I. King. A survey on vision-language-action models for embodied ai. *arXiv preprint arXiv:2405.14093*, 2024.
- [154] D. Mahajan, R. Girshick, V. Ramanathan, K. He, M. Paluri, Y. Li, A. Bharambe, and L. van der Maaten. Exploring the limits of weakly supervised pretraining. In *European Conference on Computer Vision*, 2024.
- [155] B. P. Majumder, H. Surana, D. Agarwal, B. D. Mishra, A. Meena, A. Prakhar, T. Vora, T. Khot, A. Sabharwal, and P. Clark. Discoverybench: Towards data-driven discovery with large language models. *arXiv preprint arXiv:2407.01725*, 2024.

- [156] B. S. Manning, K. Zhu, and J. J. Horton. Automated social science: Language models as scientist and subjects. Technical report, National Bureau of Economic Research, 2024.
- [157] E. Margffoy-Tuay, J. C. Perez, E. Botero, and P. Arbelaez. Dynamic multimodal instance segmentation guided by natural language queries. In *European Conference on Computer Vision*, 2024.
- [158] P. Mazzaglia, T. Verbelen, B. Dhoedt, A. Courville, and S. Rajeswar. Genrl: Multimodal-foundation world models for generalization in embodied agents. In *Neural Information Processing Systems*, 2024.
- [159] F. Meng, Z. Wang, and M. Zhang. Pissa: Principal singular values and singular vectors adaptation of large language models. In *Neural Information Processing Systems*, 2024.
- [160] O.-B. Mercea, A. Gritsenko, C. Schmid, and A. Arnab. Time- memory- and parameter-efficient visual adaptation. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [161] L. Messeri and M. J. Crockett. Artificial intelligence and illusions of understanding in scientific research. *Nature*, 627(8002):49–58, 2024.
- [162] J. Michaux, A. Li, Q. Chen, C. Chen, and R. Vasudevan. Safe planning for articulated robots using reachability-based obstacle avoidance with spheres. In *Robotics: Science and Systems*, 2024.
- [163] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, 2024.
- [164] S. Mo, F. Mu, K. H. Lin, Y. Liu, B. Guan, Y. Li, and B. Zhou. Freecontrol: Training-free spatial control of any text-to-image diffusion model with any condition. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [165] W. Mo, T. Zhang, Y. Bai, B. Su, J.-R. Wen, and Q. Yang. Dynamic prompt optimizing for text-to-image generation. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [166] T. Mu, A. Helyar, J. Heidecke, J. Achiam, A. Vallone, I. D. Kivlichan, M. Lin, A. Beutel, J. Schulman, and L. Weng. Rule based rewards for language model safety. In *Neural Information Processing Systems*, 2024.
- [167] M. Najibi, J. Ji, Y. Zhou, C. R. Qi, X. Yan, S. Ettinger, and D. Anguelov. Motion inspired unsupervised perception and prediction in autonomous driving. In *European Conference on Computer Vision*, 2024.
- [168] M. Narasimhan, E. Wijmans, X. Chen, T. Darrell, D. Batra, D. Parikh, and A. Singh. Seeing the un-scene: Learning amodal semantic maps for room navigation. In *European Conference on Computer Vision*, 2024.
- [169] X. Ning, T. Zhao, W. Li, P. Lei, Y. Wang, and H. Yang. Dsa: More efficient budgeted pruning via differentiable sparsity allocation. In *European Conference on Computer Vision*, 2024.
- [170] D. Niu, Y. Sharma, G. Biamby, J. Quenum, Y. Bai, B. Shi, T. Darrell, and R. Herzig. Llarva: Vision-action instruction tuning enhances robot learning. In *Conference on Robot Learning*, 2024.
- [171] A. Oshin, H. Almubarak, and E. Theodorou. Differentiable robust model predictive control. In *Robotics: Science and Systems*, 2024.

- [172] P. Oza, H. V. Nguyen, and V. M. Patel. Multiple class novelty detection under data distribution shift. In *European Conference on Computer Vision*, 2024.
- [173] C. Pan, Z. Yi, G. Shi, and G. Qu. Model-based diffusion for trajectory optimization. In *Neural Information Processing Systems*, 2024.
- [174] R. Parihar, S. VS, S. Mani, T. Karmali, and V. B. Radhakrishnan. Precisecontrol: Enhancing text-to-image diffusion models with fine-grained attribute control. In *European Conference on Computer Vision*, 2024.
- [175] L. Peng, J. Xu, H. Cheng, Z. Yang, X. Wu, W. Qian, W. Wang, B. Wu, and D. Cai. Learning occupancy for monocular 3d object detection. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [176] M. Peychev, A. Ruoss, M. Balunović, M. Baader, and M. Vechev. Latent space smoothing for individually fair representations. In *European Conference on Computer Vision*, 2024.
- [177] C. M. Pham, A. Hoyle, S. Sun, P. Resnik, and M. Iyyer. Topicgpt: A prompt-based topic modeling framework. *arXiv preprint arXiv:2311.01449*, 2023.
- [178] S. Purushwalkam, P. Morgado, and A. Gupta. The challenges of continuous self-supervised learning. In *European Conference on Computer Vision*, 2024.
- [179] Y. Qi, Z. Pan, S. Zhang, A. van den Hengel, and Q. Wu. Object-and-action aware model for visual language navigation. In *European Conference on Computer Vision*, 2024.
- [180] Z. Qi, R. Dong, S. Zhang, H. Geng, C. Han, Z. Ge, L. Yi*, and K. Ma*. Shapellm: Universal 3d object understanding for embodied interaction. In *European Conference on Computer Vision*, 2024.
- [181] C. Qian, Z. Xie, Y. Wang, W. Liu, K. Zhu, H. Xia, Y. Dang, Z. Du, W. Chen, C. Yang, Z. Liu, and M. Sun. Scaling large language model-based multi-agent collaboration. In *International Conference on Learning Representations*, 2025.
- [182] L. Rout, Y. Chen, N. Ruiz, C. Caramanis, S. Shakkottai, and W.-S. Chu. Semantic image inversion and editing using rectified stochastic differential equations. In *International Conference on Learning Representations*, 2025.
- [183] D. Rozenberszki, O. Litany, and A. Dai. Language-grounded indoor 3d semantic segmentation in the wild. In *European Conference on Computer Vision*, 2024.
- [184] H. Ryu, S. Lim, and H. Shim. Memory-efficient fine-tuning for quantized diffusion model. In *European Conference on Computer Vision*, 2024.
- [185] A. Sadat, S. Casas, M. Ren, X. Wu, P. Dhawan, and R. Urtasun. Perceive, predict, and plan: Safe motion planning through interpretable semantic representations. In *European Conference on Computer Vision*, 2024.
- [186] T. Salimans, T. Mensink, J. Heek, and E. Hooeboom. Multistep distillation of diffusion models via moment matching. In *Neural Information Processing Systems*, 2024.
- [187] A. Sampieri, G. M. D. di Melendugno, A. Avogaro, F. Cunico, F. Setti, G. Skenderi, M. Cristani, and F. Galasso. Pose forecasting in industrial human-robot collaboration. In *European Conference on Computer Vision*, 2024.
- [188] M. H. Sarhan, N. Navab, A. Eslami, and S. Albarqouni. Fairness by learning orthogonal disentangled representations. In *European Conference on Computer Vision*, 2024.

- [189] Y. Savani, M. A. Finzi, and J. Z. Kolter. Diffusing differentiable representations. In *Neural Information Processing Systems*, 2024.
- [190] S. Schmidgall, Y. Su, Z. Wang, X. Sun, J. Wu, X. Yu, J. Liu, M. Moor, Z. Liu, and E. Barsoum. Agent laboratory: Using llm agents as research assistants. *arXiv preprint arXiv:2501.04227*, 2025.
- [191] Y. Shaoul, I. Mishani, S. Vats, J. Li, and M. Likhachev. Multi-robot motion planning with diffusion models. In *International Conference on Learning Representations*, 2025.
- [192] R. Shetty, M. Fritz, and B. Schiele. Towards automated testing and robustification by semantic adversarial data generation. In *European Conference on Computer Vision*, 2024.
- [193] F. Shi, C. Zhang, T. Miki, J. Lee, M. Hutter, and S. Coros. Rethinking robustness assessment: Adversarial attacks on learning-based quadrupedal locomotion controllers. In *Robotics: Science and Systems*, 2024.
- [194] L. X. Shi, Z. Hu, T. Z. Zhao, A. Sharma, K. Pertsch, J. Luo, S. Levine, and C. Finn. Yell at your robot: Improving on-the-fly from language corrections. In *Robotics: Science and Systems*, 2024.
- [195] M. Shi, F. Liu, S. Wang, S. Liao, S. Radhakrishnan, Y. Zhao, D.-A. Huang, H. Yin, K. Sapra, Y. Yacoob, H. Shi, B. Catanzaro, A. Tao, J. Kautz, Z. Yu, and G. Liu. Eagle: Exploring the design space for multimodal llms with mixture of encoders. In *International Conference on Learning Representations*, 2025.
- [196] X. Shi, Y. Li, Q. Kou, L. Yu, J. Xie, and H. Zhou. Spar: Scholar paper retrieval with llm-based agents for enhanced academic search. *arXiv preprint arXiv:2507.15245*, 2025.
- [197] C. Si, X. Wang, X. Yang, Z. Xu, Q. Li, J. Dai, Y. Qiao, X. Yang, and W. Shen. Maintaining structural integrity in parameter spaces for parameter efficient fine-tuning. In *International Conference on Learning Representations*, 2025.
- [198] R. Sinha, A. Elhafsi, C. Agia, M. Foutter, E. Schmerling, and M. Pavone. Real-time anomaly detection and reactive planning with large language models. In *Robotics: Science and Systems*, 2024.
- [199] S. Skand, B. Pandit, C. Kim, L. Fuxin, and S. Lee. Simple masked training strategies yield control policies that are robust to sensor failure. In *Conference on Robot Learning*, 2024.
- [200] J. P. Sleiman, M. Mittal, and M. Hutter. Guided reinforcement learning for robust multi-contact loco-manipulation. In *Conference on Robot Learning*, 2024.
- [201] H. Song, W. Ding, Y. Chen, S. Shen, M. Y. Wang, and Q. Chen. Pip: Planning-informed trajectory prediction for autonomous driving. In *European Conference on Computer Vision*, 2024.
- [202] J. Song, Y. Yang, H. Xiao, W. Peng, W. Yao, and F. Wang. Laser: Towards diversified and generalizable robot design with large language models. In *International Conference on Learning Representations*, 2025.
- [203] V. Sreeramdas, R. R. Paleja, L. Chen, S. van Waveren, and M. Gombolay. Generalized behavior learning from diverse demonstrations. In *International Conference on Learning Representations*, 2025.
- [204] S. Srivastava and G. Sharma. Omnivec2 - a novel transformer based network for large scale multimodal and multitask learning. In *Conference on Computer Vision and Pattern Recognition*, 2024.

- [205] K. Stechly, K. Valmeekam, and S. Kambhampati. On the self-verification limitations of large language models on reasoning and planning tasks. In *International Conference on Learning Representations*, 2025.
- [206] N. Stracke, S. A. Baumann, J. Susskind, M. A. Bautista, and B. Ommer. Ctrlralter: Conditional loradapter for efficient 0-shot control & altering of t2i models. In *European Conference on Computer Vision*, 2024.
- [207] V. Subramaniam, Y. Du, J. B. Tenenbaum, A. Torralba, S. Li, and I. Mordatch. Multiagent finetuning: Self improvement with diverse reasoning chains. In *International Conference on Learning Representations*, 2025.
- [208] O. Taheri, N. Ghorbani, M. J. Black, and D. Tzionas. Grab: A dataset of whole-body human grasping of objects. In *European Conference on Computer Vision*, 2024.
- [209] S. Tan, W. Xiang, H. Liu, D. Guo, and F. Sun. Multi-agent embodied question answering in interactive environments. In *European Conference on Computer Vision*, 2024.
- [210] H. Tang, D. Y. Key, and K. Ellis. Worldcoder, a model-based llm agent: Building world models by writing code and interacting with the environment. In *Neural Information Processing Systems*, 2024.
- [211] S. Tang*, Y. Wang, C. Ding, Y. Liang, Y. Li, and D. Xu. Adadiff: Accelerating diffusion models through step-wise adaptive computation. In *European Conference on Computer Vision*, 2024.
- [212] G. Tevet, B. Gordon, A. Hertz, A. H. Bermano, and D. Cohen-Or. Motionclip: Exposing human motion generation to clip space. In *European Conference on Computer Vision*, 2024.
- [213] F. Tian, Y. Li, Y. Yan, S. Guan, Y. Ge, and X. Yang. Postedit: Posterior sampling for efficient zero-shot image editing. In *International Conference on Learning Representations*, 2025.
- [214] S. Tong, E. L. B. II, P. Wu, S. Woo, A. J. IYER, S. C. Akula, S. Yang, J. Yang, M. Middepogu, Z. Wang, X. Pan, R. Fergus, Y. LeCun, and S. Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. In *Neural Information Processing Systems*, 2024.
- [215] H.-Y. Tseng, M. Fisher, J. Lu, Y. Li, V. Kim, and M.-H. Yang. Modeling artistic workflows for image generation and editing. In *European Conference on Computer Vision*, 2024.
- [216] D. Turpin, L. Wang, E. Heiden, Y.-C. Chen, M. Macklin, S. Tsogkas, S. Dickinson, and A. Garg. Grasp’d: Differentiable contact-rich grasp synthesis for multi-fingered hands. In *European Conference on Computer Vision*, 2024.
- [217] R. Van Noorden and J. M. Perkel. Ai and science: what 1,600 researchers think. *Nature*, 621(7980): 672–675, 2023.
- [218] A. Vettoruzzo, L. Braccaioli, J. Vanschoren, and M. Nowaczyk. Unsupervised meta-learning via in-context learning. In *International Conference on Learning Representations*, 2025.
- [219] V. Viswanathan, K. Gashteovski, K. Gashteovski, C. Lawrence, T. Wu, and G. Neubig. Large language models enable few-shot clustering. *Transactions of the Association for Computational Linguistics*, 12:321–333, 2024.
- [220] H. Wang, W. Wang, T. Shu, W. Liang, and J. Shen. Active visual information gathering for vision-language navigation. In *European Conference on Computer Vision*, 2024.

- [221] L. Wang, X. Chen, J. Zhao, and K. He. Scaling proprioceptive-visual learning with heterogeneous pre-trained transformers. In *Neural Information Processing Systems*, 2024.
- [222] Q. Wang, D. Downey, H. Ji, and T. Hope. Scimon: Scientific inspiration machines optimized for novelty. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 279–299, 2024.
- [223] R. Wang, J. Xiang, J. Yang, and X. Tong. Diffusion models are geometry critics: Single image 3d editing using pre-trained diffusion priors. In *European Conference on Computer Vision*, 2024.
- [224] T.-H. Wang, S. Manivasagam, M. Liang, B. Yang, W. Zeng, and R. Urtasun. V2vnet: Vehicle-to-vehicle communication for joint perception and prediction. In *European Conference on Computer Vision*, 2024.
- [225] W. Wang, Q. Lv, W. Yu, W. Hong, J. Qi, Y. Wang, J. Ji, Z. Yang, L. Zhao, S. XiXuan, J. Xu, K. Chen, B. Xu, J. Li, Y. Dong, M. Ding, and J. Tang. Cogvlm: Visual expert for pretrained language models. In *Neural Information Processing Systems*, 2024.
- [226] X. Wang*, Y. Zhang, O. Zohar, and S. Yeung-Levy. Videoagent: Long-form video understanding with large language model as agent. In *European Conference on Computer Vision*, 2024.
- [227] X. Wang, Z. Zhu, G. Huang, C. Xinze, J. Zhu, and J. Lu. Drivedreamer: Towards real-world-driven world models for autonomous driving. In *European Conference on Computer Vision*, 2024.
- [228] X. E. Wang, V. Jain, E. Ie, W. Y. Wang, Z. Kozareva, and S. Ravi. Environment-agnostic multitask learning for natural language grounded navigation. In *European Conference on Computer Vision*, 2024.
- [229] Y. Wang, J. He, L. Fan, H. Li, Y. Chen, and Z. Zhang. Driving into the future: Multiview visual forecasting and planning with world model for autonomous driving. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [230] Y. Wang*, K. Li, X. Li, J. Yu, Y. He, G. Chen, B. Pei, R. Zheng, J. Xu, Z. Wang, Y. Shi, T. Jiang, S. Li, hongjie Zhang, Y. Huang, Y. Qiao*, Y. Wang*, and L. Wang*. Internvideo2: Scaling foundation models for multimodal video understanding. In *European Conference on Computer Vision*, 2024.
- [231] Y. Wang, Y. Lu, and T. Blankevoort. Differentiable joint pruning and quantization for hardware efficiency. In *European Conference on Computer Vision*, 2024.
- [232] Y. Wang, C. Tang, L. Sun, S. Rossi, Y. Xie, C. Peng, T. Hannagan, S. Sabatini, N. Poerio, M. Tomizuka, and W. Zhan. Optimizing diffusion models for joint trajectory prediction and controllable generation. In *European Conference on Computer Vision*, 2024.
- [233] Z. Wang, Z. Zhang, S. Ebrahimi, R. Sun, H. Zhang, C.-Y. Lee, X. Ren, G. Su, V. Perot, J. Dy, and T. Pfister. Dualprompt: Complementary prompting for rehearsal-free continual learning. In *European Conference on Computer Vision*, 2024.
- [234] Z. Wang, Z. Liu, T. Ma, J. Li, Z. Zhang, X. Fu, Y. Li, Z. Yuan, W. Song, Y. Ma, et al. Graph foundation models: A comprehensive survey. *arXiv preprint arXiv:2505.15116*, 2025.
- [235] A. Werby, C. Huang, M. Büchner, A. Valada, and W. Burgard. Hierarchical open-vocabulary 3d scene graphs for language-grounded robot navigation. In *Robotics: Science and Systems*, 2024.

- [236] M. Williams, M. Carroll, A. Narang, C. Weisser, B. Murphy, and A. Dragan. On targeted manipulation and deception when optimizing llms for user feedback. In *International Conference on Learning Representations*, 2025.
- [237] K. Wong, Q. Zhang, M. Liang, B. Yang, R. Liao, A. Sadat, and R. Urtasun. Testing the safety of self-driving vehicles by simulating perception and prediction. In *European Conference on Computer Vision*, 2024.
- [238] S. Wu, H. Fei, X. Li, J. Ji, H. Zhang, T.-S. Chua, and S. YAN. Towards semantic equivalence of tokenization in multimodal llm. In *International Conference on Learning Representations*, 2025.
- [239] Y. Wu, J. Wang, Y. Zhang, S. Zhang, O. Hilliges, F. Yu, and S. Tang. Saga: Stochastic whole-body grasping with contact. In *European Conference on Computer Vision*, 2024.
- [240] Y. Wu, Z. Sun, H. Yuan, K. Ji, Y. Yang, and Q. Gu. Self-play preference optimization for language model alignment. In *International Conference on Learning Representations*, 2025.
- [241] Y. Wu, Z. Zhang, J. Chen, H. Tang, D. Li, Y. Fang, L. Zhu, E. Xie, H. Yin, L. Yi, S. Han, and Y. Lu. Vila-u: a unified foundation model integrating visual understanding and generation. In *International Conference on Learning Representations*, 2025.
- [242] X. Xiao, J. Liu, Z. Wang, Y. Zhou, Y. Qi, S. Jiang, B. He, and Q. Cheng. Robot learning in the era of foundation models: A survey. *Neurocomputing*, page 129963, 2025.
- [243] J. Xing*, M. Xia, Y. Zhang, H. Chen, W. Yu, H. Liu, G. Liu, X. Wang, Y. Shan, and T.-T. Wong. Dynamicrafter: Animating open-domain images with video diffusion priors. In *European Conference on Computer Vision*, 2024.
- [244] D. Xu, Y. Jiang, P. Wang, Z. Fan, H. Shi, and Z. Wang. Sinnerf: Training neural radiance fields on complex scenes from a single image. In *European Conference on Computer Vision*, 2024.
- [245] F. F. Xu, U. Alon, G. Neubig, and V. J. Hellendoorn. A systematic evaluation of large language models of code. In *Proceedings of the 6th ACM SIGPLAN international symposium on machine programming*, pages 1–10, 2022.
- [246] Z. Xu, K. Wu, J. Wen, J. Li, N. Liu, Z. Che, and J. Tang. A survey on robotics with foundation models: toward embodied ai. *arXiv preprint arXiv:2402.02385*, 2024.
- [247] Z. Xu, M. Liu, Y. Shen, J. Rimchala, J. Zhang, Q. Wang, Y. Cheng, and L. Huang. Modality-specialized synergizers for interleaved vision-language generalists. In *International Conference on Learning Representations*, 2025.
- [248] F. Yang, C. Feng, Z. Chen, H. Park, D. Wang, Y. Dou, Z. Zeng, X. Chen, R. Gangopadhyay, A. Owens, and A. Wong. Binding touch to everything: Learning unified multimodal tactile representations. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [249] M. Yang, C. Lu, A. Church, Y. Lin, C. J. Ford, H. Li, E. Psomopoulou, D. A. Barton, and N. F. Lepora. Anyrotate: Gravity-invariant in-hand object rotation with sim-to-real touch. In *Conference on Robot Learning*, 2024.
- [250] R. Yang, Z. Chen, J. Ma, C. Zheng, Y. Chen, Q. Nguyen, and X. Wang. Generalized animal imitator: Agile locomotion with versatile motion prior. In *Conference on Robot Learning*, 2024.
- [251] Y. Yang, B. Huang, F. Feng, X. Wang, S. Tu, and L. Xu. Towards generalizable reinforcement learning via causality-guided self-adaptive representations. In *International Conference on Learning Representations*, 2025.

- [252] S. Yarram* and J. Yuan. Forecasting future videos from novel views via disentangled 3d scene representation. In *European Conference on Computer Vision*, 2024.
- [253] J. Ye, J. Gao, S. Gong, L. Zheng, X. Jiang, Z. Li, and L. Kong. Beyond autoregression: Discrete diffusion for complex reasoning and planning. In *International Conference on Learning Representations*, 2025.
- [254] J. Ye, K. Wang, C. Yuan, R. Yang, Y. Li, J. Zhu, Y. Qin, X. Zou, and X. Wang. Dex1b: Learning with 1b demonstrations for dexterous manipulation. In *Robotics: Science and Systems (RSS)*, 2025.
- [255] J. Ye, H. Xu, H. Liu, A. Hu, M. Yan, Q. Qian, J. Zhang, F. Huang, and J. Zhou. mplug-owl3: Towards long image-sequence understanding in multi-modal large language models. In *International Conference on Learning Representations*, 2025.
- [256] J. Yin, J. Shen, R. Chen, W. Li, R. Yang, P. Frossard, and W. Wang. Is-fusion: Instance-scene collaborative fusion for multimodal 3d object detection. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [257] P. Yin, T. Westenbroek, C.-A. Cheng, A. Kolobov, and A. Gupta. Rapidly adapting policies to the real-world via simulation-guided fine-tuning. In *International Conference on Learning Representations*, 2025.
- [258] Z.-H. Yin and P. Abbeel. Offline imitation learning through graph search and retrieval. In *Robotics: Science and Systems*, 2024.
- [259] J. Yoon, S. Yu, V. Patil, H. Yao, and M. Bansal. Safree: Training-free and adaptive guard for safe text-to-image and video generation. In *International Conference on Learning Representations*, 2025.
- [260] C. Yu, X. Yang, J. Gao, H. Yang, Y. Wang, and Y. Wu. Learning efficient multi-agent cooperative visual exploration. In *European Conference on Computer Vision*, 2024.
- [261] K. Yu, Y. Han, Q. Wang, V. Saxena, D. Xu, and Y. Zhao. Mimictouch: Leveraging multi-modal human tactile demonstrations for contact-rich manipulation. In *Conference on Robot Learning*, 2024.
- [262] S. Yu and C. Lu. Adam: An embodied causal agent in open-world environments. In *International Conference on Learning Representations*, 2025.
- [263] H. Yuan, B. Zhou, Y. Fu, and Z. Lu. Cross-embodiment dexterous grasping with reinforcement learning. In *International Conference on Learning Representations*, 2025.
- [264] S. Yuan, H. Liu, and H. Xu. Bridging the gap between low-rank and orthogonal adaptation via householder reflection adaptation. In *Neural Information Processing Systems*, 2024.
- [265] W. Yuan, J. Duan, V. Blukis, W. Pumacay, R. Krishna, A. Murali, A. Mousavian, and D. Fox. Robopoint: A vision-language model for spatial affordance prediction in robotics. In *Conference on Robot Learning*, 2024.
- [266] Z. Yuan, T. Wei, S. Cheng, G. Zhang, Y. Chen, and H. Xu. Learning to manipulate anywhere: A visual generalizable framework for reinforcement learning. In *Conference on Robot Learning*, 2024.
- [267] F. Zeng, W. Gan, Y. Wang, N. Liu, and P. S. Yu. Large language models for robotics: A survey. *arXiv preprint arXiv:2311.07226*, 2023.

- [268] Y. Zhan, Y. Zhu*, Z. Chen, F. Yang, M. Tang, and J. Wang. Griffon: Spelling out all object locations at any granularity with large language models. In *European Conference on Computer Vision*, 2024.
- [269] C. Zhang, L. Liu, Y. Cui, G. Huang, W. Lin, Y. Yang, and Y. Hu. A comprehensive survey on segment anything model for vision and beyond. *arXiv preprint arXiv:2305.08196*, 2023.
- [270] G. Zhang, Y. Yue, Z. Li, S. Yun, G. Wan, K. Wang, D. Cheng, J. X. Yu, and T. Chen. Cut the crap: An economical communication pipeline for llm-based multi-agent systems. In *International Conference on Learning Representations*, 2025.
- [271] H. Zhang, S. Christen, Z. Fan, O. Hilliges, and J. Song. Graspxl: Generating grasping motions for diverse objects at scale. In *European Conference on Computer Vision*, 2024.
- [272] H. Zhang*, H. Li, F. Li, T. Ren, X. Zou, S. Liu, S. Huang, J. Gao, L. Zhang, C. Li, and J. Yang. Llava-grounding: Grounded visual chat with large multimodal models. In *European Conference on Computer Vision*, 2024.
- [273] H. Zhang, D. Tang, J. Liu, M. Lu, J. Zheng, J. Peng, D. Li, Y. Wang, F. Jiang, L. Tian, S. Tiwari, A. Sirasao, J.-H. Yong, B. Wang, and E. Barsoum. Dip-go: A diffusion pruner via few-step gradient optimization. In *Neural Information Processing Systems*, 2024.
- [274] H. Zhang, Z. Wang, Q. Lyu, Z. Zhang, S. Chen, T. Shu, B. Dariush, K. Lee, Y. Du, and C. Gan. Combo: Compositional world models for embodied multi-agent cooperation. In *International Conference on Learning Representations*, 2025.
- [275] J. Zhang, M. Heo, Z. Liu, E. Biyik, J. J. Lim, Y. Liu, and R. Fakoore. Extract: Efficient policy learning by extracting transferable robot skills from offline data. In *Conference on Robot Learning*, 2024.
- [276] J. Zhang, K. Wang, R. Xu, G. Zhou, Y. Hong, X. Fang, Q. Wu, Z. Zhang, and H. Wang. Navid: Video-based vlm plans the next step for vision-and-language navigation. In *Robotics: Science and Systems*, 2024.
- [277] J. Zhang, G. Zhu, S. Li, X. Liu, H. Song, X. Tang, and C. Feng. Multiview scene graph. In *Neural Information Processing Systems*, 2024.
- [278] L. Zhang, M. Kan, S. Shan, and X. Chen. Prelar: World model pre-training with learnable action representation. In *European Conference on Computer Vision*, 2024.
- [279] W. Zhang, P. Torr, M. Elhoseiny, and A. Bibi. Bi-factorial preference optimization: Balancing safety-helpfulness in language models. In *International Conference on Learning Representations*, 2025.
- [280] X. Zhang and A. Boularias. One-shot imitation learning with invariance matching for robotic manipulation. In *Robotics: Science and Systems*, 2024.
- [281] Y. Zhang, Z. Wang, and J. Shang. Clusterllm: Large language models as a guide for text clustering. *arXiv preprint arXiv:2305.14871*, 2023.
- [282] Y. Zhang, J. Jia, X. Chen, A. Chen, Y. Zhang, J. Liu, K. Ding, and S. Liu. To generate or not? safety-driven unlearned diffusion models are still easy to generate unsafe images ... for now. In *European Conference on Computer Vision*, 2024.
- [283] Y. Zhang*, E. Tzeng, Y. Du, and D. Kislyuk*. Large-scale reinforcement learning for diffusion models. In *European Conference on Computer Vision*, 2024.

- [284] Z. Zhang, H. Liu, J. Chen, and X. Xu. Gooddrag: Towards good practices for drag editing with diffusion models. In *International Conference on Learning Representations*, 2025.
- [285] B. Zhao, S. Yu, W. Ma, M. Yu, S. Mei, A. Wang, J. He, A. Yuille, and A. Kortylewski. Ood-cv: A benchmark for robustness to out-of-distribution shifts of individual nuisances in natural images. In *European Conference on Computer Vision*, 2024.
- [286] T. Zhao, Y. Chen, Y. Wu, T. Liu, B. Du, P. Xiao, S. Qiu, H. Yang, G. Li, Y. Yang, and Y. Lin. Improving bird’s eye view semantic segmentation by task decomposition. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [287] W. Zhao, P. Ding, Z. Min, Z. Gong, S. Bai, H. Zhao, and D. Wang. Vlas: Vision-language-action model with speech instructions for customized robot manipulation. In *International Conference on Learning Representations*, 2025.
- [288] W. X. Zhao, K. Zhou, J. Li, T. Tang, X. Wang, Y. Hou, Y. Min, B. Zhang, J. Zhang, Z. Dong, et al. A survey of large language models. *arXiv preprint arXiv:2303.18223*, 2023.
- [289] Y. Zhao, S. Dang, H. Ye, G. Dai, Y. Qian, and I. Tsang. Second-order fine-tuning without pain for llms: A hessian informed zeroth-order optimizer. In *International Conference on Learning Representations*, 2025.
- [290] Y. Zhao, M. Uehara, G. Scalia, S. Kung, T. Biancalani, S. Levine, and E. Hajiramezanali. Adding conditional control to diffusion models with reinforcement learning. In *International Conference on Learning Representations*, 2025.
- [291] Y. Zhao, W. Zhang, Y. Xie, A. Goyal, K. Kawaguchi, and M. Shieh. Understanding and enhancing safety mechanisms of llms via safety-specific neuron. In *International Conference on Learning Representations*, 2025.
- [292] J. Zheng, H. Wang, A. Zhang, T. D. Nguyen, J. Sun, and T.-S. Chua. Ali-agent: Assessing llms’ alignment with human values via agent-based evaluation. In *Neural Information Processing Systems*, 2024.
- [293] R. Zheng, Y. Liang, S. Huang, J. Gao, H. D. III, A. Kolobov, F. Huang, and J. Yang. Tracevla: Visual trace prompting enhances spatial-temporal awareness for generalist robotic policies. In *International Conference on Learning Representations*, 2025.
- [294] W. Zheng, W. Chen, Y. Huang, B. Zhang, Y. Duan, and J. Lu. Occworld: Learning a 3d occupancy world model for autonomous driving. In *European Conference on Computer Vision*, 2024.
- [295] X. Zheng*, Y. Lyu, jiazhou zhou, and L. Wang*. Centering the value of every modality: Towards efficient and resilient modality-agnostic semantic segmentation. In *European Conference on Computer Vision*, 2024.
- [296] X. Zheng*, Y. Lyu, and L. Wang*. Learning modality-agnostic representation for semantic segmentation from any modalities. In *European Conference on Computer Vision*, 2024.
- [297] Z. Zhong, J. Cao, songen gu, S. Xie, L. Luo, H. Zhao, G. Zhou, H. Li, and Z. Yan*. Structured-nerf: Hierarchical scene graph with neural representation. In *European Conference on Computer Vision*, 2024.
- [298] C. Zhou, Q. Li, C. Li, J. Yu, Y. Liu, G. Wang, K. Zhang, C. Ji, Q. Yan, L. He, et al. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *International Journal of Machine Learning and Cybernetics*, pages 1–65, 2024.

- [299] C. Zhou, X. Liu, F. Luo, and S. Huang. Latent radiance fields with 3d-aware 2d representations. In *International Conference on Learning Representations*, 2025.
- [300] C. Zhou, M. Zhang, P. Chen, C. Fu, Y. Shen, X. Zheng, X. Sun, and R. Ji. Learning interleaved image-text comprehension in vision-language large models. In *International Conference on Learning Representations*, 2025.
- [301] Z. Zhou, P. Atreya, A. Lee, H. R. Walke, O. Mees, and S. Levine. Autonomous improvement of instruction following skills via foundation models. In *Conference on Robot Learning*, 2024.
- [302] Z. Zhuang, S. Yao, and H. Zhao. Humanoid parkour learning. In *Conference on Robot Learning*, 2024.
- [303] W. Zimmer, G. A. Wardana, S. Sritharan, X. Zhou, R. Song, and A. C. Knoll. Tumtraf v2x cooperative perception dataset. In *Conference on Computer Vision and Pattern Recognition*, 2024.
- [304] A. Zou, L. Phan, J. Wang, D. Duenas, M. Lin, M. Andriushchenko, J. Z. Kolter, M. Fredrikson, and D. Hendrycks. Improving alignment and robustness with circuit breakers. In *Neural Information Processing Systems*, 2024.
- [305] X. Zou, L. Li, J. Wang, J. Yang, M. Ding, J. Wei, Z. Yang, F. Li, H. Zhang, S. Liu, A. Aravinthan, Y. J. Lee, and L. Wang. Interfacing foundation models’ embeddings. In *Neural Information Processing Systems*, 2024.
- [306] A. Zouitine, D. Bertoin, P. Clavier, M. Geist, and E. Rachelson. Time-constrained robust mdps. In *Neural Information Processing Systems*, 2024.