

GLID: Pre-training a Generalist Encoder-Decoder Vision Model

Jihao Liu^{1,2*} Jinliang Zheng^{5*} Yu Liu^{2,3}✉ Hongsheng Li^{1,3,4}✉
¹CUHK MMLab ²SenseTime Research
³Shanghai AI Laboratory ⁴CPII under InnoHK
⁵Institute for AI Industry Research (AIR), Tsinghua University

Abstract

This paper proposes a **GeneraLIst encoder-Decoder (GLID)** pre-training method for better handling various downstream computer vision tasks. While self-supervised pre-training approaches, e.g., Masked Autoencoder, have shown success in transfer learning, task-specific sub-architectures are still required to be appended for different downstream tasks, which cannot enjoy the benefits of large-scale pre-training. GLID overcomes this challenge by allowing the pre-trained generalist encoder-decoder to be fine-tuned on various vision tasks with minimal task-specific architecture modifications. In the GLID training scheme, pre-training pretext task and other downstream tasks are modeled as “query-to-answer” problems, including the pre-training pretext task and other downstream tasks. We pre-train a task-agnostic encoder-decoder with query-mask pairs. During fine-tuning, GLID maintains the pre-trained encoder-decoder and queries, only replacing the topmost linear transformation layer with task-specific linear heads. This minimizes the pretrain-finetune architecture inconsistency and enables the pre-trained model to better adapt to downstream tasks. GLID achieves competitive performance on various vision tasks, including object detection, image segmentation, pose estimation, and depth estimation, outperforming or matching specialist models such as Mask2Former, DETR, ViTPose, and BinsFormer.

1. Introduction

Pre-training with self-supervision on large-scale unlabeled images has achieved great success in visual representation learning. The representative Masked Image Modeling (MIM) [1, 2, 13, 15, 18, 38, 41] methods demonstrate that the pre-trained models have impressive scale capabilities and can greatly boost performance on downstream vision tasks, including image classification, object detection, image segmentation, pose estimation, depth estimation, etc.

*Equal contribution




Figure 1. MAE backbone-only pre-training vs. GLID pre-training. The GLID pre-training allows the pre-trained encoder-decoder to be fine-tuned on various vision tasks without task-specific decoder designs and outperforms MAE backbone-only pre-training.

Despite the significant progress made by MIM pre-training, existing approaches still face architectural gaps between the upstream pre-training and downstream fine-tuning. In particular, existing approaches [2, 18, 38] mainly focus on pre-training the vision backbone. However, to tackle the aforementioned downstream vision tasks, task-specific sub-architectures are still needed, e.g., a new detection decoder for object detection [5] and a segmentation decoder for panoptic segmentation [9]. Those task-specific sub-architectures are complicated and need to be trained from scratch on downstream tasks and therefore cannot enjoy the benefits from large-scale pre-training. Such designs lag behind the progress in natural language processing where the same architecture can handle both the pre-training tasks and downstream tasks with minimal differences. More recently, there have been notable advancements in tackling diverse vision tasks with a generalist architecture [9, 19, 20, 23, 31, 37] and can achieve competitive performance compared to the task-specific specialist models. However, only the vision backbones in their generalist architectures are pre-trained while the heavy decoders still need to be trained from scratch on downstream tasks. Consequently, these generalist approaches typically require a large amount of task-specific data to achieve satisfactory performance.

To tackle the aforementioned problems, this paper introduces a self-supervised pre-training method to pre-train a **GeneraLIst encoder-Decoder (GLID)** vision model, which

can be adapted to various vision tasks with marginal task-specific architecture modifications. The key idea behind GLID is to model different tasks as “query-to-answer” problems with the generalist encoder-decoder model, including the pre-training pretext task and the downstream tasks such as object detection, image segmentation, etc. We pre-train the whole generalist encoder-decoder model by tackling a specially designed masked image modeling pretext task on the large-scale ImageNet dataset. In particular, each query is associated with a mask patch location and is set as its positional embedding. The “answer” or the learning target to each query is the pixel values of the corresponding masked patch. The encoder and the decoder can be jointly pre-trained through the reconstruction pretext task. For fine-tuning on downstream tasks, we keep the pre-trained encoder-decoder and queries but replace only the topmost single linear transformation layer with a task-specific linear layer to output the desired targets of each task. In this way, the proposed GLID maintains pre-trained weights as much as possible and minimizes the pretrain-finetune architecture inconsistency, as shown in Fig. 1. In contrast, the conventional MIM approaches [18] discard the pre-trained decoder during fine-tuning and randomly initialize a new task-specific decoder for each new downstream task, causing significant pretrain-finetune gap.

We demonstrate that we can apply the proposed GLID to 6 computer vision tasks without introducing architecture twists. Through extensive experiments, we show that the GLID method, despite its simplicity, can achieve strong performance on downstream vision tasks, outperforming or matching more specialized models such as Mask2Former [9], DETR [5], ViTPose [39], and BinsFormer [25]. Additionally, we demonstrate that GLID is more data-efficient on downstream tasks through the generalist encoder-decoder pre-training.

In summary, our contributions are as follows:

- We propose GLID, a generalist encoder-decoder pre-training method for various vision tasks.
- In GLID, both the pre-training and fine-tuning tasks are modeled as “query-to-answer” problems, thus minimizing the pretrain-finetune architectural gap.
- Through extensive experiments, GLID is shown to outperform or match specialized models on 6 vision tasks and is more data-efficient through the generalist encoder-decoder pre-training.

2. Related Works

Mask Image Modeling (MIM) is a popular pretext task for visual representation learning that aims to reconstruct the masked tokens from a corrupted input. The MAE approach [18] proposes an asymmetric encoder-decoder architecture and uses the encoder to process partially visible tokens and predicts the corresponding pixel values at the

masked positions with a lightweight decoder. In general, the central objective of previous work on MIM is to devise an improved architectural design [8, 18] or reconstruction target [2, 13, 18, 33, 38], with the aim of facilitating the acquisition of a superior representation by the model. However, a common limitation of these works is that only the pre-trained encoder can be transferred to downstream tasks, while the decoder must be replaced with a more complex, task-specific architecture, leading to a significant pretrain-finetune gap.

Unified architecture has been gaining popularity in recent years for solving various vision tasks with one framework. Mask2Former [9] is an example of such an architecture that can handle three segmentation tasks. OneFormer [20] builds on Mask2Former and proposes a task-conditioned joint training strategy to unify image segmentation. Beyond image segmentation, other unified models, such as Unified-IO [31], UniT [19], Pix2seq v2 [7], X-Decoder [43], and UViM [23], can jointly learn various computer vision tasks over multiple datasets, including pose estimation, object detection, image generation, etc. However, these models typically need more task-specific data to achieve satisfactory performance.

Decoder pre-training has been proposed to tackle the problem that the task-specific decoder needs to be trained from scratch on downstream tasks. DDeP [4] proposed a decoder pre-training approach based on denoising, which can be combined with supervised pre-training of the encoder. In UP-DETR [10], a patch feature reconstruction branch is proposed for the decoder, which is jointly optimized with patch detection. Instead of pre-training the encoder and decoder separately with specially designed pre-training tasks, VADeR [32] and DenseCL [36] propose to pre-train the entire architecture through pixel-level contrastive learning and then transfer to dense prediction tasks. In contrast to those works, GLID pre-trains a generalist encoder-decoder model and can better adapt to downstream tasks with minimal architectural gap.

3. Methods

Although pre-training a visual encoder with various strategies has made significant progress, task-specific decoders are still needed for fine-tuning the pre-trained encoder to tackle different downstream tasks [5, 9, 25, 39]. To tackle this problem, we propose to pre-train a **Generalist Encoder-Decoder (GLID)** vision model for better handling various downstream vision tasks and introduce minimal architecture modifications when shifting from pre-training to fine-tuning. We illustrate the framework of the proposed GLID in Fig. 2.

3.1. Generalist Encoder-Decoder Pre-training

To reduce the pretrain-finetune architectural gap and take full advantage of the pre-training, we pre-train a generalist




Figure 2. Overview of GLID. During pre-training, we pre-train a task-agnostic encoder-decoder transformer architecture through masked image modeling. For fine-tuning on downstream tasks, we replace the pre-training linear head with a task-specific linear head. In this way, the proposed GLID minimizes the pretrain-finetune gap and enables the pre-trained architecture to better adapt to downstream tasks.

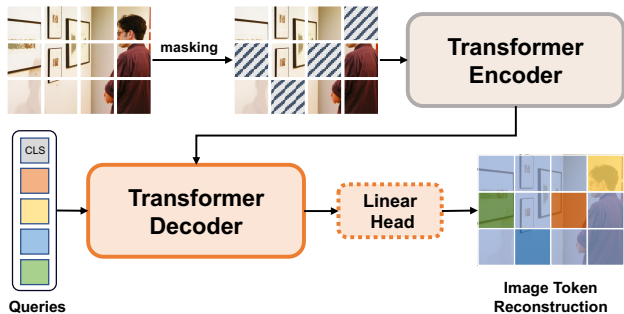


Figure 3. GLID pre-training pipeline.

encoder-decoder architecture that can be directly applied to downstream tasks with a unified formulation for the pre-training and fine-tuning. Specifically, given an input image $x \in \mathbb{R}^{H \times W \times 3}$ where H, W are the image height and width, an image encoder Enc is used to extract visual features, and a query-based transformer decoder Dec to decode out the final features, which are converted to the target predictions P by a linear head H ,

$$P = H(\text{Dec}(Q, \text{Enc}(x))), \quad (1)$$

where Q denotes the queries. All tasks are formulated in this way, including the pre-training and fine-tuning tasks. We use different Q - P pairs to switch between different tasks,

We pre-train the encoder-decoder through a specially designed Masked Image Modeling (MIM) [18, 38] pretext task. Specifically, given the input image, we partition it into image patches and convert them into image tokens, and a random mask M is applied to mask out a proportion of input tokens. Then the visual encoder is utilized to only process the visible tokens following MAE [18]. To align the decoder architecture with the decoders used in downstream tasks, we reformulate the MIM pretext task as a “query-to-answer” [5] problem

and use the query-based transformer decoder [35] to decode out the masked tokens for pre-training. In particular, we introduce a number of mask tokens $[M]$ as the input queries of the decoder, which are added with different positional embedding to indicate a unique masked position. Moreover, we append one additional $[CLS]$ token query [12] at the 0th position, which is not associated with a specific masked location but is used to capture global representations. So during pre-training, we can express Q as $Q = [CLS, M_1, M_2, \dots, M_N]$ where N denotes the number of the masked tokens. After decoding, a linear transformation head is utilized to convert the decoded query representations to target predictions. We use pixel values P^t at the masked locations as the reconstruction targets, so P presents the predicted pixel values in the pre-training stage. We calculate the mean squared error (MSE) loss between P and P^t for pre-training,

$$\mathcal{L}_{rec} = \frac{1}{N} \sum_{i=1}^N \|p_i - p_i^t\|_2^2, \quad (2)$$

where p_i and p_i^t indicate the reconstructed and original pixel values of i th masked tokens. Note that the $[CLS]$ token is not used to calculate the reconstruction loss, but can be optimized during pre-training because the other queries’ predictions are conditioned on the $[CLS]$ token. We illustrate the pre-training pipeline of GLID in Fig. 3.

Our formulation enables the pre-training of a generalist encoder-decoder model, which can be directly fine-tuned on downstream tasks without adding complex task-specific decoders, leading to a minimal pretrain-finetune architectural gap. In contrast, existing pre-training approaches such as MAE-based [18] methods discard their decoders after the pre-training and add a new task-specific decoder for each downstream task’s fine-tuning. Consequently, such methods

require training extra task-specific decoders on downstream tasks from scratch and cannot take full advantage of the pre-training as the decoders are randomly initialized.

3.2. Fine-tuning on Downstream Tasks

Based on the generalist design, we can apply the pre-trained GLID model to various downstream tasks, including object detection, image segmentation, pose estimation, and depth estimation. Following the formulation in Sec. 3.1, all those tasks are also modeled as new “query-to-answer” problems. We can fine-tune GLID to those tasks seamlessly by replacing the topmost linear layer used in pre-training with a task-specific linear layer and assigning different learning targets to the query-answer pairs. We use M to denote the number of queries for downstream tasks, which may be different for different downstream tasks. To better utilize the pre-trained weights, we initialize the M queries with the pre-trained [CLS] token by repeating the [CLS] token for M times and adding M learnable embeddings to the repeated tokens,

$$q_i = [\text{CLS}] + e_i, \quad (3)$$

where q_i and e_i denote the i th query and i th learnable embeddings respectively. Note that the learnable embeddings are zero-initialized to be compatible with the pre-trained weights.

Object detection. For the object detection task, each query is used to represent an object instance following DETR [5]. Two linear layers are utilized to convert the decoded hidden features to a bounding box and class probabilities respectively, $P = \{[bbox_i, cls_i]\}_{i=1}^M$.

Image segmentation. Following Mask2Former [9], each query is utilized to predict a C -dimensional mask embedding and its class, $P = \{[f_i^s, cls_i]\}_{i=1}^M$, where $f_i^s \in \mathbb{R}^C$ and cls_i represent the i th mask embedding and its class. Thus we can obtain the binary mask predictions via a dot product between the mask embeddings and the backbone’s 1/4-scale feature map $f^b \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$ followed by a sigmoid activation,

$$S \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times M} = \text{sigmoid}(\langle f^b, f^s \rangle), \quad (4)$$

where S denotes the binary mask prediction.

Depth estimation. Following AdaBins [3] and BinsFormer [25], the depth regression task is formulated as a “classification-regression” problem where the continuous predictions are obtained via a linear combination of bin centers. Specifically, the queries’ hidden features are converted to bin lengths and bin embeddings by the depth head, $P = \{[l_i, z_i]\}_{i=1}^M$ where $l_i \in \mathbb{R}$, $z_i \in \mathbb{R}^C$ are the length and embedding of the i th bin. We follow the pose process in AdaBins to obtain the depth maps. We first obtain the bins’ centers as:

$$c(l_i) = d_{min} + (d_{max} - d_{min}) \left(\frac{l_i}{2} + \sum_{j=1}^{i-1} l_j \right), \quad (5)$$

where $c(l_i)$, d_{max} , and d_{min} are the center depth of the i th bins, maximum valid depth, and minimum valid depth, respectively. Then we can obtain a probability distribution map P^d via a dot product between the bins embeddings and the backbone’s 1/4-scale feature map followed by a softmax function,

$$P^d \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times M} = \text{softmax}(\langle f^b, z \rangle). \quad (6)$$

where $f^b \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$ denotes the backbone’s feature map. Finally, the depth map can be obtained via a linear combination of bin centers,

$$d = \sum_{i=1}^M c(l_i) p_i^d, \quad (7)$$

where $p_i^d \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4}}$ denotes the i th map of P^d .

Pose estimation. We utilize a heatmap-based approach for pose estimation following ViTPose [39]. In particular, each query is used to output a heatmap of a key point. A pose prediction head is employed to convert the query features to C -dimensional feature vectors for representing the different key points’ heatmaps $P = \{f_i^k\}_{i=1}^M$ with $f_i^k \in \mathbb{R}^C$ denoting the i th heatmap feature vector. The pose heatmaps can be obtained as

$$P^k \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times M} = \langle f^b, f^k \rangle. \quad (8)$$

where $f^b \in \mathbb{R}^{\frac{H}{4} \times \frac{W}{4} \times C}$ denotes the backbone’s feature map.

After obtaining the queries’ predictions, we follow prior arts [5, 9, 25, 39] to use the task loss between the predictions and ground truths for fine-tuning. In particular, we calculate the bipartite matching loss for object detection and image segmentation tasks as for those two tasks, the number of predictions is typically larger than the number of ground truths. For the depth estimation and pose estimation tasks, we use Scale Invariant (SI) regression loss [14] and smooth L_1 loss [16] respectively. When fine-tuning GLID, we inherit majority of the pre-trained architecture weights to take full advantage of the self-supervised pre-training on large-scale dataset.

Architecture. As we aim at applying GLID to various visual tasks, a hierarchical transformer architecture is utilized for better transferability on downstream tasks. In particular, we use Swin Transformer [29] as the visual encoder to produce multi-scale feature maps. A Bi-FPN [27] architecture is utilized for better modeling the interaction across different feature map scales. The multi-scale feature maps are then inputted to the decoder of L cross-attention and self-attention blocks to produce final predictions. Following Mask2Former [9], the decoder transformer processes the multi-scale feature maps with cross-attention layers successively.

Method	Type	Backbone	Data	w/ Labels	COCO		ADE20K			NYUv2		
					AP ^{kp} ↑	AP ^{box} ↑	mIoU↑	mAP↑	PQ↑	REL↓	RMSE↓	
Mask-RCNN [17]	Specialist	ResNeXt-101	ImageNet-1K	✓	63.1	37.1	-	-	-	-	-	
Mask2Former [9]		Swin-B		✓	-	-	52.4	30.7 [†]	44.3 [†]	-	-	
ViTPose [39]		ViT-B		✗	75.8	-	-	-	-	-	-	
ViTPose [39]		ViT-L		✗	78.3	-	-	-	-	-	-	
DETR [5]		ResNet101-DCN		✓	-	44.9	-	-	*	-	-	
Deformable-DETR [42]		ResNet50-DC5		✓	-	46.9	-	-	-	-	-	
BinsFormer [25]		Swin-B		✓	-	-	-	-	-	-	0.104	0.362
BinsFormer [25]		Swin-L		✓	ImageNet-21K	✓	-	-	-	-	0.094	0.330
Pix2seq v2 [6]		Generalist		ViT-B	Objects365	✓	68.0	46.5	-	*	*	-
UniT [19]	ResNet50		COCO, VG, VQAv2	✓	-	40.8	-	-	-	-	-	
UViM [23]	ViT-L		ImageNet-21K	✓	-	-	*	*	*	*	0.467	
Unified-IO [31]	Unified-IO-B		95 datasets	✓	*	*	-	*	-	*	0.469	
Unified-IO [31]	Unified-IO-L		95 datasets	✓	*	*	-	*	-	*	0.402	
Painter [37]	ViT-L		ImageNet-1K	✗	72.1	-	49.9	*	*	0.080	0.288	
GLID	Swin-B		ImageNet-1K	✗	76.7	51.2	52.7	30.9	44.7	0.085	0.293	
GLID	Swin-L	ImageNet-1K	✗	78.5	52.4	53.9	32.5	45.7	0.081	0.287		

Table 1. Comparison with previous specialized and generalized approaches. The ‘‘Data’’ denotes pre-training data. The ‘‘w/ Labels’’ denotes the pre-training uses supervised supervision. ‘‘*’’ denotes the model can perform such task but does not have reported results. ‘‘-’’ denotes the model cannot perform such task. † denotes our implementation with the official code.

Pre-train Method	w/ Labels	COCO		ADE20K			NYUv2	
		AP ^{kp} ↑	AP ^{box} ↑	mIoU↑	mAP↑	PQ↑	REL↓	RMSE↓
Supervised	✓	72.1	44.8	50.0	26.2	40.7	0.182	0.510
EsViT [24]	✗	73.1	45.0	48.7	27.0	41.5	0.160	0.415
MAE [24]	✗	74.2	49.2	51.5	26.5	40.3	0.165	0.340
MixMAE [28]	✗	75.8	47.5	49.7	29.7	43.0	0.106	0.327
SimMIM [24]	✗	75.6	48.9	50.6	25.9	39.2	0.098	0.343
GLID	✗	76.7	51.2	52.7	30.9	44.7	0.085	0.293

Table 2. Comparison with previous pre-training methods on studied tasks. All entries use Swin-B backbone and are pre-trained on ImageNet-1K. The ‘‘w/ Labels’’ denotes pre-training with classification labels of ImageNet-1K.

4. Experiment Setups

We validate the effectiveness of GLID with experiments on various vision tasks. We first pre-train a generalist encoder-decoder on the ImageNet-1K [11] dataset in a self-supervised manner. We then fine-tune the pre-trained encoder-decoder to various downstream tasks, including object detection/pose estimation on COCO [26], semantic/instance/panoptic segmentation on ADE20K [40], and depth estimation on NYU-Depth-v2 [34].

Dataset and evaluation metrics. We use the large-scale ImageNet-1K [11] dataset for the self-supervised pre-training, which contains 1.2M images. After pre-training, we transfer the pre-trained models to various downstream tasks to demonstrate the transferability of our pre-training. In particular, we transfer to the ADE20K [40] dataset to conduct fine-tuning on the semantic, instance, and panoptic segmentation tasks. For those three tasks, we use the mean intersection over union (mIoU), mean average precision (mAP), and panoptic quality (PQ) [22] metrics for performance comparison. We also fine-tune the pre-trained models to the depth estimation task on the NYU-Depth V2 [34] dataset. By default, we report the root mean square error (RMSE) and mean absolute relative error (REL) for performance evaluation. Moreover, we also transfer to the large-scale COCO [26] dataset and conduct fine-tuning on the pose esti-

mation task and object detection task, and report the mean average precision of pose estimation AP^{kp} and object detection AP^{box} respectively for performance evaluation.

Pre-training. By default, we pre-train for 600 epochs with the input size of 192×192 and batch size 2048. The Swin-B and Swin-L [29] are utilized as the backbones. We use 6 transformer blocks in the decoder with a hidden embedding dimension of 512. We randomly mask the input images with a mask ratio of 75%. The AdamW [30] optimizer is utilized with an initial learning rate of 1.2×10^{-3} and weight decay of 0.05. We apply random resizing cropping and horizontal flipping for data augmentation following MAE [18].

Fine-tuning. We conduct fine-tuning experiments on various downstream tasks following previous practice. Note that during fine-tuning, we do not introduce architecture modifications except for task-specific linear head, thus can inherit most of the pre-trained weights. In particular, for the three segmentation tasks, we fine-tune for 160K iterations with batch size of 16 and input size of 640×640 . Following Mask2Former [9], we use 100/200/200 queries for semantic/instance/panoptic segmentation. For depth estimation, we use official data split and fine-tune for 38.4K iterations with batch size of 16. We use 64 queries following the practice in BinsFormer [25]. For object detection, we fine-tune for 50 epochs with batch size of 16. Following H-DETR [21], we use 1,800 queries for training and 300 queries for evaluation.

For pose estimation, we train and evaluate on the cropped regions which contain person instances. We fine-tune for 210 epochs with 17 queries.

Multitask fine-tuning. We also explore fine-tuning with multiple tasks simultaneously. We use the ADE20K [40] dataset and conduct fine-tuning with different segmentation tasks jointly. Except for the task-specific linear head, other parameters are shared. During fine-tuning, the batch samples of different tasks are evenly included. The training hyperparameters are the same as those used for single-task fine-tuning.

5. Main Results

In this section, we compare our GLID to prior arts on various benchmarks. We first conduct comparisons between GLID and previous pre-training approaches in Sec. 5.1. We then compare our results with previous generalized and specialized models in Sec. 5.2. Lastly, we demonstrate the multi-task learning results on segmentation tasks in Sec. 5.3.

5.1. Comparison with pre-training methods

We compare our pre-training method, GLID, with previous pre-training approaches to demonstrate its effectiveness on various visual tasks. Several pre-training approaches are utilized, including the supervised pre-training on ImageNet-1K, the contrastive approach EsViT [24], and the masked-image-modeling approaches [18, 28, 38]. Using the pre-trained backbone of the same architecture, we fine-tune them on the studied downstream tasks and compare their performances in Tab. 2. Our approach outperforms other compared approaches in terms of all reported metrics, demonstrating the effectiveness of our pre-training method.

Particularly, on the three tasks of the ADE20K [40] dataset, we improve the supervised pre-training for an average of 3.8 points. Compared to the recent self-supervised learning methods MixMAE [28] and EsViT, we achieve an average of 1.9 and 3.7 points improvements on these three segmentation tasks. On the NUYv2 [34] depth estimation task, GLID achieves 0.293 RMSE (root mean square error), 0.29 better than the supervised pre-training. In addition, compared to the self-supervised EsViT, our GLID can improve its performance by 0.222 RMSE. On the larger-scale COCO [26] dataset, GLID can still outperform other pre-training approaches. In particular, we improve the supervised pre-training by 4.6 and 6.4 points on AP^{kp} and AP^{box} respectively. Compared to the self-supervised pre-training MixMAE, we surpass its performance by 0.9 AP^{kp} and 3.7 AP^{box} . All the results demonstrate the effectiveness of our generalist encoder-decoder pre-training.

5.2. Comparison with specialists and generalists

Table 1 shows comparisons of our approach with other specialized and generalized approaches. Our generalist ap-

proach can tackle the studied tasks with minimal architectural modifications and achieves strong performance on those tasks. In particular, we only use the ImageNet-1K for self-supervised pre-training, while other approaches use a comparable or stronger setting than ours, such as more pre-training data or supervision.

Specifically, we achieve better performance than Mask2Former [9] on the three segmentation tasks with the same backbone. On the depth estimation task, our large model achieves 0.287 RMSE, demonstrating a significant improvement on the previous state-of-the-art BinsFormer [25]. On the COCO keypoint estimation and object detection task, our approach can surpass the specialized approaches ViT-Pose [39] and DETR [5]. Additionally, compared to the generalized approaches which typically use a larger-scale dataset for pre-training, our approach can still outperform their results by using less pre-training data. In particular, we improve the recent Pix2seq v2 [6] on COCO by 10.5 AP^{kp} and 5.9 AP^{box} . While the generalized approaches typically fail to improve the specialized models, our approach achieves strong performance and can match or outperform the specialized approaches, demonstrating the effectiveness of our method.

In addition, our pre-training method can be scaled up to large models. In particular, our pre-trained Swin-L [29] outperforms Swin-B on all studied tasks. For the three segmentation tasks on ADE20K, the large model has an average of 1.2 mIOU improvements over the base model. On the COCO dataset, the large model improves the base model by 1.8 AP^{kp} and 1.2 AP^{box} .

5.3. Results of multitask fine-tuning

Apart from separately fine-tuning the pre-trained model on various tasks, we also explore the multitask fine-tuning. Specifically, we employ the ADE20K dataset for fine-tuning as its annotations can be used for three types of segmentation tasks. We fine-tune the pre-trained Swin-B on this dataset with different task combinations.

As shown in Tab. 3, we find that the multitask models underperform the single-task models. In particular, under the multitask setting, the best performances for the three tasks are 52.1 mIoU, 29.1 mAP, and 43.2 PQ, which are 0.6, 1.8, and 1.5 worse than the single-task performances. The results are may be caused by task interference. Specifically, compared to training semantic and panoptic tasks, training all three tasks results in 0.3 mIoU and 1.1 PQ drops. However, the results also demonstrate that adding more tasks can also lead to performance improvements. In particular, compared to training instance and panoptic tasks, adding semantic task improves the mAP and PQ for 0.5 and 0.6 points respectively.

Semantic	Instance	Panoptic	mIoU	mAP	PQ
Single task			52.7	30.9	44.7
✓	✓	✗	51.3	29.1	-
✓	✗	✓	52.1	-	43.2
✗	✓	✓	-	28.0	41.5
✓	✓	✓	51.8	28.5	42.1

Table 3. Multitask vs. single task fine-tuning on ADE20K.

Backbone	BiFPN	Decoder	mIoU↑	RMSE↓
✗	✗	✗	34.9	0.753
✓	✗	✗	50.9	0.363
✓	✓	✗	51.1	0.341
✓	✓	✓	52.7	0.293

Table 4. Ablation of loading partial pre-trained weights. The results in row 1 are from scratched training.

# decoder layers	mIoU↑	RMSE↓
3	50.7	0.302
6	52.7	0.293
9	52.7	0.292

Table 5. Ablation of the number of decoder layers.

# pre-training epochs	mIoU↑	RMSE↓
0	34.9	0.753
300	51.5	0.314
600	52.7	0.293
900	52.8	0.290

Table 6. Ablation of the number of pre-training epochs.

6. Ablation Studies

In this section, we conduct ablation studies to evaluate the impact of different design choices of our proposed **GeneralList** encoder-**Gecoder** pre-training (**GLID**) on the final performance. Unless otherwise specified, we use the Swin-B [29] as the encoder and 6 transformer blocks in the decoder and pre-train the encoder-decoder architecture for 600 epochs on ImageNet-1K [11]. We fine-tune the pre-trained models on the ADE20K [40] semantic segmentation task and NYUv2 [34] depth estimation task. For performance comparison, we report mIoU on ADE20K and RMSE on NYUv2.

Load partial pre-trained weights. To investigate the importance of our generalist encoder-decoder pre-training, we conducted an ablation study by loading only partial pre-trained weights for fine-tuning. For the unloaded architecture, we initialize it with random weights. We then fine-tune the model on two downstream tasks and compared their performance with the full pre-trained model.

As shown in Tab. 4, the performance gradually improves as more pre-trained weights are loaded. In particular, compared to only loading the backbone weights as in previous works, loading all pre-trained weights improves the mIoU by 1.8 points and reduces the RMSE by 0.07. Additionally, loading the BiFPN part can also outperform the choice of only loading the backbone weights. All the results demonstrate the importance of the generalist encoder-decoder pre-training.

Masking strategy	mIoU	Masking ratio	mIoU
grid	50.8	0.5	49.8
block	49.3	0.6	51.0
random	51.5	0.75	51.5

Table 7. Ablation of masking strategies and ratios.

% fine-tuning data	Pre-training	mIoU↑	RMSE↓
10	SL	27.0	0.779
	GLID	31.2	0.317
20	SL	31.3	0.681
	GLID	35.0	0.303
50	SL	43.2	0.620
	GLID	46.3	0.295
100	SL	50.0	0.510
	GLID	52.7	0.293

Table 8. Fine-tuning with limited data. “SL” denotes using the supervised pre-training on ImageNet-1K.

Decoder depth. To investigate the impact of decoder depth on our model, we test varying the number of decoder layers in the model. Specifically, we pre-train models with 3, 6, and 9 decoder layers, and fine-tune them on downstream tasks. Note that we use the same number of decoder layers for both pre-training and fine-tuning. As shown in Tab. 5, we find that using 3 layers results in 2.0 mIoU and 0.009 RMSE performance drops compared to those of using 6 layers. Additionally, using more decoder layers does not bring significant performance gains but introduces extra computation costs.

Pre-training epochs. To investigate the effect of the number of pre-training epochs on the performance of the model, we vary the number of pre-training epochs. Specifically, we pre-train with different numbers of epochs and fine-tune the pre-trained models on downstream tasks. As shown in Tab. 6, we find that we can improve the performance by pre-training more epochs. Specifically, compared to pre-training with 300 epochs, we can improve the mIoU by 1.2 and RMSE by 0.021 with extra 300 epochs of pre-training. However, we find that the performance gradually saturates and does not improve much if we pre-train for 900 epochs.

Mask strategy and mask ratio. We investigate the impact of masking strategy and masking ratio and show the results in Tab. 7. We pre-train for 300 epochs and report the mIoU on ADE20K for comparison. In particular, we find that using random masking performs best among other masking strategies. Additionally, using a large mask ratio, *i.e.*, 0.75, leads to the best result.

Fine-tuning with limited data. We investigate the impact of limited data during fine-tuning. Specifically, given the GLID and ImageNet-1K-supervised pre-training models, we fine-tune them on downstream tasks with different percentages of data. As shown in Tab. 8, we find that our approach has larger performance advantage over the supervised pre-training models when using fewer data for fine-tuning. In particular, when using 10% fine-tuning data, our approach




Figure 4. Query’s cross-attention maps on different tasks. Each column shows three-scale attention maps on a specific task. Task-specific queries are utilized for visualization.




Figure 5. Performance curves of loading different parts of pre-trained weights.

surpasses the supervised pre-training by 4.2 mIoU and 0.09 RMSE respectively. When we use 100% data for fine-tuning, the performance gaps are 2.6 mIoU and 0.047 RMSE respectively. The main reason is that for the supervised pre-training approach, we need to train the task-specific decoder from scratch on downstream tasks, In comparison, our unified pre-training approach can pre-train and keep pre-trained encoder-decoder simultaneously, and achieve strong performance even without much task-specific data.

7. Qualitative Evaluation

Query’s cross-attention maps on different tasks. We show the query’s cross-attention maps on various tasks in Fig. 4. We observe that the query can focus on different semantic parts for different tasks. In particular, for the pre-training pretext task, the query learns to attend to the local regions for reconstruction. For the object detection and pose estimation tasks, the query attends to some extreme points, such as the boundary of the object or a specific key point location. Moreover, for the segmentation or depth estimation tasks, the attention maps have a clear separation for different semantic regions.

Convergence curves. We compare the convergence curves of partially loading the pre-trained weights on the semantic segmentation and depth estimation tasks, with results shown in Fig. 5. We observe that loading all the pre-trained weights leads to a much faster convergence speed and can improve the final performance. In particular, compared to only loading the backbone’s weights, loading all the pre-trained

weights can achieve the same performance by using half of the iterations. Unlike the traditional approaches [18, 28] that only load the pre-trained backbone, our generalist pre-training approach allows the pre-trained encoder-decoder to be applied on downstream tasks to obtain better performance.

8. Conclusion

In this paper, we propose GLID, a generalist encoder-decoder pre-training method for better handling various downstream computer vision tasks. In the GLID pre-training scheme, both pre-training and fine-tuning tasks are modeled as “query-to-answer” problems. GLID introduces minimal task-specific architecture modifications and can inherit most pre-trained weights when fine-tuning on downstream tasks, which takes full advantage of the pre-training. We demonstrate that GLID can outperform or achieve comparable performance to specialized models on different vision tasks through extensive experimentation. In addition to better performance, the experimental results indicate that GLID leads to faster convergence speed and data efficiency.

Acknowledgement This project is funded in part by National Key R&D Program of China Project 2022ZD0161100, by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)’s InnoHK, by General Research Fund of Hong Kong RGC Project 14204021. Hongsheng Li is a PI of CPII under the InnoHK.

References

- [1] Alexei Baevski, Wei-Ning Hsu, Qiantong Xu, Arun Babu, Jiatao Gu, and Michael Auli. Data2vec: A general framework for self-supervised learning in speech, vision and language. *arXiv:2202.03555*, 2022. 1
- [2] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. In *ICLR*, 2021. 1, 2
- [3] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 4
- [4] Emmanuel Asiedu Brempong, Simon Kornblith, Ting Chen, Niki Parmar, Matthias Minderer, and Mohammad Norouzi. Denoising pretraining for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4175–4186, 2022. 2
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, pages 213–229. Springer, 2020. 1, 2, 3, 4, 5, 6
- [6] Ting Chen, Saurabh Saxena, Lala Li, David J Fleet, and Geoffrey Hinton. Pix2seq: A language modeling framework for object detection. *arXiv preprint arXiv:2109.10852*, 2021. 5, 6
- [7] Ting Chen, Saurabh Saxena, Lala Li, Tsung-Yi Lin, David J Fleet, and Geoffrey Hinton. A unified sequence interface for vision tasks. *arXiv preprint arXiv:2206.07669*, 2022. 2
- [8] Xiaokang Chen, Mingyu Ding, Xiaodi Wang, Ying Xin, Shentong Mo, Yunhao Wang, Shumin Han, Ping Luo, Gang Zeng, and Jingdong Wang. Context autoencoder for self-supervised representation learning. *arXiv:2202.03026*, 2022. 2
- [9] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1290–1299, 2022. 1, 2, 4, 5, 6
- [10] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pre-training for object detection with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1601–1610, 2021. 2
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, K. Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5, 7
- [12] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 3
- [13] Xiaoyi Dong, Jianmin Bao, Ting Zhang, Dongdong Chen, Weiming Zhang, Lu Yuan, Dong Chen, Fang Wen, and Nenghai Yu. Peco: Perceptual codebook for bert pre-training of vision transformers. *arXiv:2111.12710*, 2021. 1, 2
- [14] David Eigen, Christian Puhirsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. 4
- [15] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023. 1
- [16] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 4
- [17] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 5
- [18] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2021. 1, 2, 3, 5, 6, 8
- [19] Ronghang Hu and Amanpreet Singh. Unit: Multimodal multitask learning with a unified transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1439–1449, 2021. 1, 2, 5
- [20] Jitesh Jain, Jiachen Li, MangTik Chiu, Ali Hassani, Nikita Orlov, and Humphrey Shi. Oneformer: One transformer to rule universal image segmentation. *arXiv preprint arXiv:2211.06220*, 2022. 1, 2
- [21] Ding Jia, Yuhui Yuan, Haodi He, Xiaopei Wu, Haojun Yu, Weihong Lin, Lei Sun, Chao Zhang, and Han Hu. Detsr with hybrid matching. *arXiv preprint arXiv:2207.13080*, 2022. 5
- [22] Alexander Kirillov, Kaiming He, Ross Girshick, Carsten Rother, and Piotr Dollár. Panoptic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9404–9413, 2019. 5
- [23] Alexander Kolesnikov, André Susano Pinto, Lucas Beyer, Xiaohua Zhai, Jeremiah Harmsen, and Neil Houlsby. Uvim: A unified modeling approach for vision with learned guiding codes. *arXiv preprint arXiv:2205.10337*, 2022. 1, 2, 5
- [24] Chunyuan Li, Jianwei Yang, Pengchuan Zhang, Mei Gao, Bin Xiao, Xiyang Dai, Lu Yuan, and Jianfeng Gao. Efficient self-supervised vision transformers for representation learning. In *ICLR*, 2022. 5, 6
- [25] Zhenyu Li, Xuyang Wang, Xianming Liu, and Junjun Jiang. Binsformer: Revisiting adaptive bins for monocular depth estimation. *arXiv preprint arXiv:2204.00987*, 2022. 2, 4, 5, 6
- [26] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5, 6
- [27] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. 4
- [28] Jihao Liu, Xin Huang, Yu Liu, and Hongsheng Li. Mixmim: Mixed and masked image modeling for efficient visual representation learning. *arXiv preprint arXiv:2205.13137*, 2022. 5, 6, 8
- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 4, 5, 6, 7
- [30] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2017. 5

- [31] Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Mottaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. *arXiv preprint arXiv:2206.08916*, 2022. [1](#), [2](#), [5](#)
- [32] Pedro O O Pinheiro, Amjad Almahairi, Ryan Benmalek, Florian Golemo, and Aaron C Courville. Unsupervised learning of dense visual representations. *Advances in Neural Information Processing Systems*, 33:4489–4500, 2020. [2](#)
- [33] Jason Tyler Rolfe. Discrete variational autoencoders. In *ICLR*, 2016. [2](#)
- [34] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. *ECCV (5)*, 7576:746–760, 2012. [5](#), [6](#), [7](#)
- [35] Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [3](#)
- [36] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2021. [2](#)
- [37] Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen, and Tiejun Huang. Images speak in images: A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6830–6839, 2023. [1](#), [5](#)
- [38] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, 2021. [1](#), [2](#), [3](#), [6](#)
- [39] Yufei Xu, Jing Zhang, Qiming Zhang, and Dacheng Tao. Vitpose: Simple vision transformer baselines for human pose estimation. *arXiv preprint arXiv:2204.12484*, 2022. [2](#), [4](#), [5](#), [6](#)
- [40] Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017. [5](#), [6](#), [7](#)
- [41] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. In *ICLR*, 2021. [1](#)
- [42] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159*, 2020. [5](#)
- [43] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. Generalized decoding for pixel, image, and language. *arXiv preprint arXiv:2212.11270*, 2022. [2](#)