

Rethinking Interactive Image Segmentation with Low Latency, High Quality, and Diverse Prompts

Qin Liu Jaemin Cho Mohit Bansal Marc Niethammer
University of North Carolina at Chapel Hill

<https://github.com/uncbiag/SegNext>

Abstract

The goal of interactive image segmentation is to delineate specific regions within an image via visual or language prompts. Low-latency and high-quality interactive segmentation with diverse prompts remain challenging for existing **specialist** and **generalist** models. Specialist models, with their limited prompts and task-specific designs, experience high latency because the image must be recomputed every time the prompt is updated, due to the joint encoding of image and visual prompts. Generalist models, exemplified by the Segment Anything Model (SAM), have recently excelled in prompt diversity and efficiency, lifting image segmentation to the foundation model era. However, for high-quality segmentations, SAM still lags behind state-of-the-art specialist models despite SAM being trained with $\times 100$ more segmentation masks. In this work, we delve deep into the architectural differences between the two types of models. We observe that dense representation and fusion of visual prompts are the key design choices contributing to the high segmentation quality of specialist models. In light of this, we reintroduce this dense design into the generalist models, to facilitate the development of generalist models with high segmentation quality. To densely represent diverse visual prompts, we propose to use a dense map to capture five types: clicks, boxes, polygons, scribbles, and masks. Thus, we propose **SegNext**, a next-generation interactive segmentation approach offering low latency, high quality, and diverse prompt support. Our method outperforms current state-of-the-art methods on HQSeg-44K and DAVIS, both quantitatively and qualitatively.

1. Introduction

Interactive image segmentation is a long-standing computer vision task that aims to precisely delineate specific image regions, typically using visual or linguistic prompts. Low-latency interactive segmentation is crucial for prompt decision-making in real-time applications [16, 25, 38]. With

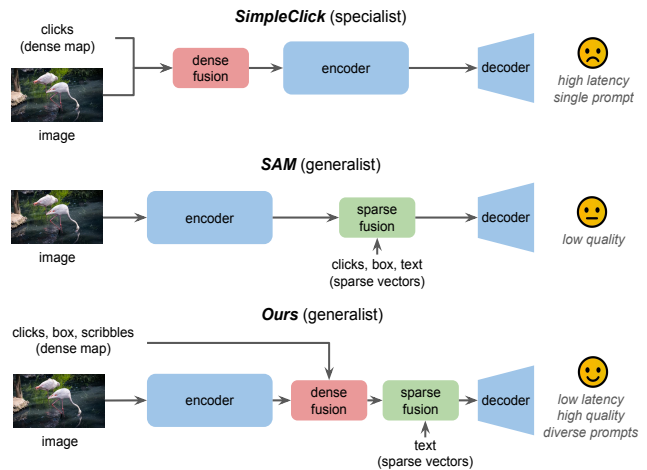


Figure 1. Conceptual comparison between our approach **SegNext** and prior state-of-the-art methods, SimpleClick [27] and SAM [16], for the interactive segmentation task. Our method combines the best of both worlds for interactive segmentation with low latency, high quality, and diverse prompts.

rapid advancements in camera and display technologies, image resolution has been significantly elevated, making 4K and 6K formats increasingly becoming the norm [34]. Therefore, the importance of high-quality interactive segmentation for high-resolution images has become more pronounced. On top of these, an effective and practical interactive segmentation approach that supports diverse prompts is key to improving user experience.

However, these requirements remain challenging for existing interactive segmentation approaches, including specialist [4, 13, 27, 36] and generalist [14, 16, 43, 47] models. Specialist models, represented by FocalClick [4] and SimpleClick [27], suffer from the high latency issue due to the joint encoding of image and visual prompts. Although InterFormer [13] circumvents this issue with a decoupled design, it still lacks diversity in the prompts, as only clicks are explored. Generalist models, represented by SAM [16] and SEEM [47], have recently achieved great success in

addressing the two issues, lifting image segmentation into the era of foundation models. Nonetheless, in the high-quality segmentation regime, SAM remains a step behind the leading specialist models when more clicks are provided, despite being trained with $\times 100$ more segmentation masks. To sum up, designing an optimal architecture for low-latency, high-quality interactive segmentation with diverse prompts remains an unresolved challenge in the field.

To tackle this challenge, we delve deep into the architectural distinctions between specialist and generalist models (Fig. 1 briefly summarizes the major distinctions). We hypothesize that the sparse representation of visual prompts, akin to the treatment of linguistic prompts, could potentially dampen the performance of generalist models for high-quality segmentation. This hypothesis is derived from our answer to a fundamental **question**: *What are the key distinctions between visual and linguistic prompts in the context of interactive image segmentation?* **Our answer**: *Visual prompts provide detailed spatial information of image objects, whereas linguistic prompts offer semantic insights without spatial detail. Therefore, linguistic prompts excel in general object recognition, while visual prompts are essential in refining objects for high-quality segmentation.* However, in generalist models, the sparse representation of these two distinct types of prompts fails to acknowledge their inherent differences; what remains unknown is what specific information is compromised in the representation of visual prompts as sparse vectors. A possible explanation suggests that representing visual prompts as sparse vectors, achieved either by interpolating positional embeddings as in SAM [16] or by sampling from feature space as in SEEM [47], amounts to a basic approximation of spatial information. Conversely, a dense map matching the image’s spatial dimensions better preserves the detailed spatial attributes of visual prompts. Nevertheless, an in-depth analysis of this issue falls outside the scope of this work.

Driven by our previous analysis, we propose incorporating the dense design, common in specialist models, into generalist models. We propose a three-channel dense map to represent five diverse visual prompts: clicks, boxes, polygons, scribbles, and masks (Sec. 3.1). This strategy has not yet been extensively explored. We encode the dense map of visual prompts into the image embedding space, followed by element-wise addition and a lightweight fusion module (Sec. 3.3 and 3.4).

In summary, we propose an approach that can serve as an alternative design of a generalist model for interactive image segmentation with low latency, high quality, and diverse prompts. Extensive evaluations on HQSeg-44K and DAVIS show that our approach outperforms prior state-of-the-art methods both quantitatively and qualitatively.

2. Related Work

Low-latency interactive segmentation. Low-latency segmentation is a field of ongoing research with continuous advancements in model architecture. FCN [28] is a seminal work that introduces deep convolutional networks for efficient dense predictions. Originally designed for medical image segmentation, U-Nets [33, 46] are highly efficient due to their hierarchical feature representation. While primarily a classification backbone, EfficientNet [37] can be adapted for segmentation. Its scalable architecture makes it a versatile choice for various segmentation tasks. Other advancements include DeepLab [2], HRNet [39], and ResNet [11]. Vision Transformers (ViT) have recently gained popularity in image segmentation. Swin Transformer [35] introduces shifted windows, allowing for efficient local and global feature modeling. SegFormer [40] uses a hierarchical transformer to reduce computation costs, making it suitable to be deployed on low-end devices.

Built upon efficient backbones and large-scale datasets, various advanced interactive segmentation approaches have recently been proposed [4, 7, 22–27, 36, 45]. RITM [36] adopts HRNet [39] as the backbone and proposes a combination of COCO and LVIS for training, achieving remarkable balance between low latency and state-of-the-art segmentation quality. SimpleClick [27] first uses a large pre-trained plain ViT as the backbone to outperform previous state-of-the-arts such as RITM [36] and FocalClick [4]. However, SimpleClick suffers from a high-latency issue due to jointly encoding image and visual prompts with the heavy backbone. InterFormer [13] circumvents this issue by decoupling image encoding and prompt fusion. This decoupled design is also adopted by the Segment Anything Model (SAM) [16] for promptable segmentation with low latency. Our method differs significantly from InterFormer and SAM. InterFormer uses multi-scale cross-attention modules to coarsely fuse prompt and image features, while ours uses a much simpler single-scale self-attention module for dense fusion. Besides, InterFormer is a specialist model which does not explore diverse prompts. SAM represents and fuses visual prompts sparsely, while ours uses a dense representation of visual prompts for high-quality segmentation.

High-quality interactive segmentation. High-quality segmentation refers to accurately dividing an image into meaningful and distinct segments or regions. It has been explored in various segmentation tasks, including semantic segmentation [20, 42], instance segmentation [15, 44], panoptic segmentation [6], and entity segmentation [30]. However, this field remains under-explored for interactive segmentation approaches. This is because high quality and low latency are often at odds, and interactive segmentation approaches must balance the trade-off. Ensuring sufficient quality while maintaining a responsive and real-time ex-

perience is crucial for interactive segmentation approaches. Recently, HQ-SAM [14] proposed efficient token learning for accurate mask predictions to improve the segmentation quality of SAM [16]. Instead of directly refining a coarse mask as in previous methods, HQ-SAM predicts a new high-quality mask directly by reusing SAM’s image encoder and mask decoder. Our method aims not to refine any specific approaches but to reflect on and improve the design choices in existing approaches. Our key intuition is that visual prompts in interactive image segmentation should be treated differently than language prompts, as the two types of prompts have different granularities and semantic meanings. Guided by intuition, we adopt dense representations for visual prompts and sparse ones for linguistic cues, marking the key distinction from HQ-SAM.

Segment anything and beyond. The Segment Anything Model (SAM) [16] is a recent breakthrough in interactive segmentation. It proposes the Segment Anything Task (SAT) as the “holy grail” of interactive image segmentation: a model segments the desired region given a single pixel in the foreground [18]. SAM tackles the ambiguity in this ill-posed task by 1) learning more prior knowledge of object appearance from SA-1B, the largest labeled segmentation dataset, and 2) predicting multiple candidate masks simultaneously. MobileSAM [43] and HQ-SAM [14] are two SAM variants to achieve more memory-efficient and higher-quality segmentations. With the great success of SAM and its variants, the attention in this field is shifting from *specialist* models to *generalist* models. More recently, SEEM [47] has taken the first step to unify all segmentation tasks in a single model via a universal interface. The key difference between our method and these generalist models lies in the way of representing visual prompts. We adopt a dense representation for visual prompts to keep the spatial information uncompromised. We observe that this dense design leads to high-quality segmentation.

3. SegNext

We propose an interactive image segmentation approach that supports diverse prompts for low-latency and high-quality segmentation. Our method encodes visual and language prompts separately. We use a three-channel dense map to encode five visual prompts: clicks, boxes, scribbles, polygons, and masks. We encode the language prompt as a sparse vector via the CLIP model [32]. To allow for low-latency segmentation, we encode the image *once* to obtain the image embeddings and fuse the prompt features with the image embeddings via lightweight fusion modules. In the following sections, we describe our dense representation strategy for visual prompts (Sec. 3.1) and how we fuse the visual and language prompts with the image embedding (Sec. 3.2). We conclude this section with implementation details for training and inference (Sec. 3.3 and 3.4).

3.1. Visual Prompts Representation

We encode five visual prompts (clicks, boxes, polygons, scribbles, and masks) via a three-channel dense map with an identical spatial resolution to the input image (*i.e.*, 1024×1024). We encode these five visual prompts as follows: 1) **A click** is encoded as a binary disk with a predefined radius. Positive clicks, which should be put in the foreground, are encoded in the first channel; negative clicks, which should be put in the background, are encoded in the second channel. 2) **A box** is encoded as four negative clicks, representing the box’s four corners. We mandate that the corner points lie outside the foreground area, and we have found that a couple of negative clicks sufficiently convey the object’s location. The same logic applies to polygons as well. 3) **A polygon** is encoded as a set of negative clicks, similar to the box. 4) **A scribble** is represented by a set of clicks. Similarly to the clicks, a positive scribble is encoded in the first channel, while a negative scribble is encoded in the second. 5) **A mask** is encoded in the third channel. Unlike the first two channels, the information in this channel can be ambiguous due to false positives or negatives in the initial segmentation mask. Therefore, we encode the mask prompt in a separate channel. The mask prompt allows users to use previous segmentation as the input for the next-round segmentation for better performance.

3.2. Visual and Language Prompts Fusion

We encode the three-channel dense map via a convolutional layer (*i.e.*, patch embedding layer). The resulting embeddings are element-wisely added to the image embeddings, followed by self-attention blocks for dense fusion. The text prompt can be naturally incorporated after the dense fusion. We use CLIP [32] to encode the text prompt into a vector and query the aforementioned fused embeddings via cross-attention blocks, as introduced in SAM [16]. In this work, we primarily concentrate on visual prompts, exploring text prompts only in the appendix as a proof of concept. We remove the text prompt in our main experiments as no language annotations are provided in the main training sets.

3.3. End-to-end Training

Since clicks can represent all visual prompts, we find that a model trained only with clicks generalizes well to other types of visual prompts. Therefore, we only simulate clicks during training, though it remains unexplored whether training with the other interaction types improves overall performance. Once the model is trained, we can prompt it with “unseen” visual prompts, such as boxes and scribbles.

Clicks simulation. Given the ground truth and current segmentation, we can automatically simulate clicks for training. To boost the performance, we follow RITM [36] to use a combination of random and iterative click simulation strategies. The random click simulation strategy generates

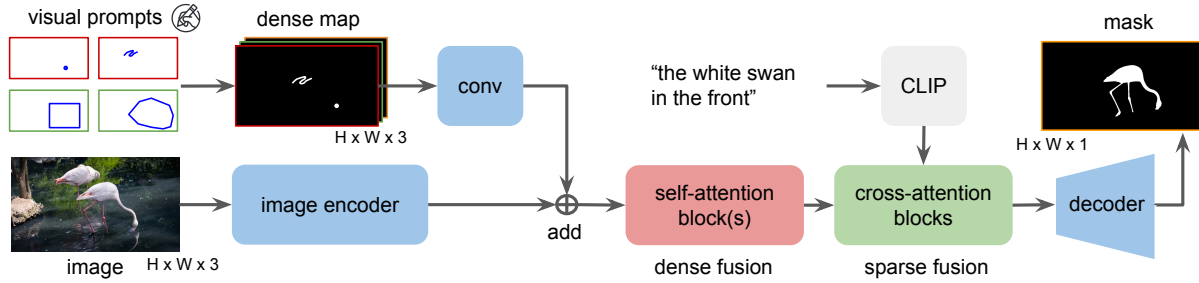


Figure 2. *SegNext* overview. We use a three-channel dense map to represent five diverse visual prompts: clicks, boxes, polygons, scribbles, and masks. The embeddings of image and visual prompts are fused by element-wise addition, followed by an enhanced fusion via one or two self-attention blocks. The language prompt is encoded as a vector by CLIP [32], followed by querying the image embedding via cross-attention blocks for the mask embeddings. A lightweight decoder processes the mask embeddings for segmentation.

a set of positive and negative clicks without considering the order between them. In contrast, the iterative simulation strategy generates clicks sequentially: a new click is generated based on the erroneous region of a prediction produced by a model using the set of previous clicks.

Loss function. We use normalized focal loss [36], allowing faster convergence and better accuracy than binary cross-entropy loss for interactive segmentation tasks as demonstrated in previous works [4, 36]. Similar training pipelines have been proposed by previous works [4, 27, 36].

3.4. Implementation Details

Image encoder. We use ViT-Base (ViT-B/16) [8] as the image encoder that takes as input an image of size 1024×1024 . Following previous works [13, 16, 27], we use readily available pre-trained weights (*i.e.*, MAE [12]) to initialize the model for stable and efficient training. The image encoder only encodes each image once. This process can be done offline before prompting the model.

Prompt encoder. We use two types of prompts: *sparse* (text) and *dense* (clicks, boxes, scribbles, polygons, masks). We represent all dense prompts in a three-channel dense map and free-form text with a CLIP model [32] (ViT-L/14@336px), an off-the-shelf text encoder. Dense prompts (*i.e.*, visual prompts) are encoded via convolutions and added element-wisely with the image embedding.

Fusion modules. After the summation of image embeddings and the embeddings of dense prompts, we use self-attention blocks to enhance the features further. For the sparse prompt (*i.e.*, text), we use the same cross-attention blocks as used in SAM [16].

Decoder. Our decoder can be any segmentation decoder. In our implementation, we use the same lightweight decoder introduced in SimpleClick [27] for its simplicity and efficiency. It has a simple feature pyramid and two MLP layers. The simple feature pyramid consists of four parallel convolutional layers, to convert the single-scale image feature to a multi-scale one. The MLP layers upsample and combine

the multi-scale representation to the output mask, from a multi-scale representation (*i.e.*, $\{\frac{1}{32}, \frac{1}{16}, \frac{1}{8}, \frac{1}{4}\}$ of the image resolution) to a single scale representation (*i.e.*, $\{\frac{1}{4}\}$ of the image resolution). The output is resized to the image’s resolution to obtain the final segmentation mask.

Training and inference settings. During training and inference, the clicks are encoded as a disk map with a fixed radius of 5 if not otherwise specified. We train our models on COCO [21]+LVIS [10] for 90 epochs with batch size 16 and fine-tune the models on HQSeg-44K [14] for an additional epoch. We observe that finetuning longer does not improve the performance. We use 4 A6000 GPUs for training and 1 A6000 for inference. The initial learning rate is set to 5×10^{-5} and decreases to 5×10^{-6} after epoch 50. We use the following data augmentation strategies: 1) uniform random resizing, 2) random horizontal and vertical flips, 3) random 90-degree rotations, 4) shift scale rotate transforms, and 5) random brightness contrasts. During training, we resize the image’s longest side to 1024 and pad the image to 1024×1024 . During inference, we directly resize the test image to 1024×1024 without padding.

4. Experiments

4.1. Benchmarks

We use COCO+LVIS, a combination of COCO [21] and LVIS [10], for training. We mainly evaluate our method on two public benchmarks: HQSeg-44K [14] and DAVIS [29]. HQSeg-44K is designed to evaluate the performance of high-quality segmentation. DAVIS contains high-quality annotations and has been widely adopted for evaluating interactive image segmentation performance. We also conduct out-of-domain evaluation on two medical datasets: ssTEM [9] and BraTS [1].

Datasets. 1) **COCO+LVIS.** COCO contains 118K training images (1.2M instances); LVIS shares the same images with COCO but has a much higher segmentation quality. 2) **HQSeg-44K,** a collection of six existing image




Figure 3. *Click to segmentation evaluation on HQSeg-44K and DAVIS.* With varying numbers of clicks, our method consistently outperforms existing competitive approaches. The metric is mean Intersection over Union (mIoU).

Method	Backbone	Training data	SAT Latency (s)↓	HQSeg-44K				DAVIS			
				5-mIoU ↑	NoC90 ↓	NoC95 ↓	NoF95 ↓	5-mIoU ↑	NoC90 ↓	NoC95 ↓	NoF95 ↓
<i>Specialist models</i>											
RITM [36]	HRNet32 ₄₀₀	COCO+LVIS	22.4	77.72	10.01	14.58	910	89.75	5.34	11.45	139
FocalClick [4]	SegF-B3-S2 ₂₅₆	COCO+LVIS	36.5	84.63	8.12	12.63	835	90.82	5.17	11.42	155
FocalClick [4]	SegF-B3-S2 ₃₈₄	COCO+LVIS	51.0	85.45	7.03	10.74	649	91.22	4.90	10.40	123
SimpleClick [27]	ViT-B ₄₄₈	COCO+LVIS	70.5	85.11	7.47	12.39	797	90.73	5.06	10.37	107
InterFormer [13]	ViT-B ₁₀₂₄	COCO+LVIS	24.3	82.62	7.17	10.77	658	87.79	5.45	11.88	150
<i>Generalist models</i>											
SAM [16]	ViT-B ₁₀₂₄	SA-1B	7.0	86.16	7.46	12.42	811	90.95	5.14	10.74	154
MobileSAM [43]	ViT-T ₁₀₂₄	SA-1B	6.6	81.98	8.70	13.83	951	89.18	5.83	12.74	196
HQ-SAM [14]	ViT-B ₁₀₂₄	SA-1B+HQ	8.3	89.85	6.49	10.79	671	91.77	5.26	10.00	136
Ours (SA×1)	ViT-B ₁₀₂₄	COCO+LVIS	13.3	85.41	7.47	11.94	731	90.13	5.46	13.31	177
Ours (SA×2)	ViT-B ₁₀₂₄	COCO+LVIS	17.6	85.71	7.18	11.52	700	89.85	5.34	12.80	163
Ours (SA×2)	ViT-B ₁₀₂₄	COCO+LVIS+HQ	17.6	91.75	5.32	9.42	583	91.87	4.43	10.73	123

Table 1. *Quantitative comparison with existing methods on HQSeg-44K and DAVIS.* We compare two types of baselines: specialist and generalist models. Our model achieves comparable performance to the specialist baselines but with significantly lower latency; our model achieves comparable performance to the generalist models in terms of latency and segmentation quality despite being trained with much fewer segmentation data. “HQ” denotes the HQSeg-44K dataset; “SA×2” denotes the model has two self-attention blocks for dense fusion.

datasets, including DIS [31] (train set), ThinObject-5K [19] (train set), FSS-1000 [17], ECSSD [35], MSRA10K [5], and DUT-OMRON [41]. Each of them contains extremely fine-grained image masks. The average number of masks for each dataset is 7.4K. The training set contains 44320 images; the validation set contains 1537 images. We report evaluation results on the validation set. 3) **DAVIS**, a high-quality and high-resolution densely annotated video segmentation dataset. Following previous works, we only use 345 frames for evaluation. Our out-of-domain evaluation datasets are: 1) **ssTEM** [9] contains 20 high-resolution medical images, and 2) **BraTS** [1] contains 69 magnetic resonance image (MRI) volumes; we test on the same 369 slices used in SimpleClick [27].

Evaluation metrics. Following previous work, we report both the number of clicks (NoC) and the number of failure cases (NoF) metrics. We also introduce a new metric mea-

suring the latency of Segmentation Anything Task (SAT). Our evaluation metrics include:

- **mIoU** measures the average intersection over union (IoU) given a fixed number of consecutive interactions. We use clicks as the default interaction type for this metric. For example, 5-mIoU measures the average IoU given five consecutive clicks. Each click is automatically simulated based on the error of the previous prediction. For the first click, the previous segmentation is empty.
- **NoC** measures the number of clicks required to achieve a predefined IoU. We set two target IoUs: 90% and 95%. The corresponding metrics are denoted as NoC%90 and NoC%95, respectively. Note that this metric is affected by the maximum number of clicks for evaluation. We set this maximum number to 20. Should a scenario necessitate more clicks than this maximum number, it will be considered a failure case (as further explained below).

- **NoF** measures the number of failure cases. As defined above, a failure case requires more than 20 clicks to achieve a predefined IoU. For example, NoF90 denotes the number of cases requiring more than 20 clicks to achieve 90% IoU.
- **SAT Latency** measures the latency for the Segment Anything Task (SAT). Given a grid of points, we measure how long it takes to prompt the model for SAT. For a fair comparison, we fix the input image size at 1024×1024 and prompt the model with a grid of 16×16 points. This metric favors generalist models where the image only needs to be encoded once, and multiple prompts amortize the encoding cost. Conventional speed metrics such as seconds per click (SPC) are not set to capture the overall latency of SAT, as they only measure one-prompt latency. The capacity of hardware affects this metric dramatically. In this work, we use an A6000 GPU and an Intel Xeon Gold 6226R CPU to obtain SAT latency results in Tab. 1.

4.2. Baselines

We compare our method with various interactive segmentation models, including four *specialist* and three *generalist* models, as follows:

- **RITM** [36] is a light-wight model state-of-the-art model that first proposes using COCO+LVIS for training. The other three specialist models also use this training set. The best-performing RITM model uses HRNet-32 [39] as the backbone. We use this model for comparison.
- **FocalClick** [4] uses a coarse-to-fine mechanism to achieve practical interactive image segmentation. FocalClick is the only multi-stage approach that requires global segmentation, local refinement, and progressive merging compared to other baselines. For comparison, we use FocalClick-B3-S2, its best-performing model trained on COCO+LVIS. We set the input size to 256×256 for the segmenter and refiner. A larger input size (*e.g.* 384×384 as reported in Tab. 1) can achieve better performance with more computation.
- **SimpleClick** [27] uses MAE-pretrained backbones to achieve state-of-the-art performance. We use its ViT-base model trained on COCO+LVIS for comparison.
- **InterFormer**. [13] decouples image encoding and prompt fusion for efficient segmentation. We use its ViT-base model trained on COCO+LVIS for comparison. This model was trained with input size 512×512 , but using 1024×1024 for inference achieves better performance.
- **SAM** [16] is trained with SA-1B, the largest labeled segmentation dataset so far ($\times 100$ more masks than COCO+LVIS); we use its ViT-base model for comparison. **MobileSAM** [14] is variant of SAM for efficient segmentation on mobile devices; we use its ViT-tiny model for comparison. **HQ-SAM** [47] is a variant of SAM for high-quality segmentation; we use its ViT-base model.

Method	Backbone	Zoom-in	ssTEM	BraTS
			10-mIoU \uparrow	10-mIoU \uparrow
CDN [3]	ResNet-34	✓	88.46	80.24
RITM [36]	HRNet32	✓	94.11	88.34
FocalClick [4]	SegF-B0-S2	✓	92.62	86.02
FocalClick [4]	SegF-B3-S2	✓	93.61	88.62
SimpleClick [27]	ViT-B	✓	93.72	86.98
SAM [16]	ViT-B	✗	91.58	87.03
Ours (SA \times 1)	ViT-B	✗	90.86	86.50
Ours (SA \times 2)	ViT-B	✗	92.87	87.29

Table 2. *Out-of-domain evaluation* on two medical image datasets: ssTEM [9] and BraTS [1]. All models are trained on COCO+LVIS. Our models generalize well on the two datasets, without test-time “zoom-in”.

4.3. Comparisons and Analysis

Quantitative comparisons. If not otherwise specified, we only use clicks as the default interaction for quantitative comparison. We compare our method with baselines on the HQSeg-44K, and DAVIS benchmarks in Tab. 1. Our method performs best across all segmentation metrics on HQSeg-44K while maintaining low latency. We observe that training on HQSeg-44K significantly boosts our model’s segmentation quality. Compared with generalist models, our method performs best while being trained with $\times 100$ fewer masks than the mask for training SAM. Compared with specialist models, our method performs competitively while enjoying the lowest latency. Histogram analysis in Fig. 4 shows the distribution of the segmentation quality. **Quantitative analysis** Although our best model achieves superior performance on DAVIS, there are still 123 failure cases among all 345 test cases. We show some failure patterns in Fig. 6.

4.4. Out-of-Domain Evaluation

We evaluate the generalizability of our models on two medical image datasets: ssTEM [9] and BraTS [1]. We compare our method with both specialist and generalist baselines. Our models generalize very well on the two datasets, without the “zoom-in” strategy that has been widely adopted in specialist models [3, 4, 13, 27, 36]. “Zoom-in” is a test-time optimization strategy that focuses on cropping specific local areas, which are determined based on the estimated locations of objects, to achieve finer segmentation. While it is efficient, this approach also incurs extra computational expenses. Tab. 2 reports the evaluation results on the two datasets. Overall, our models generalize well on the out-of-domain datasets, without test-time “zoom-in”.

4.5. Diverse Prompts Evaluation

In the previous experiments, we focused solely on using clicks as the interaction mode. However, our method is compatible with a variety of prompts. This section will

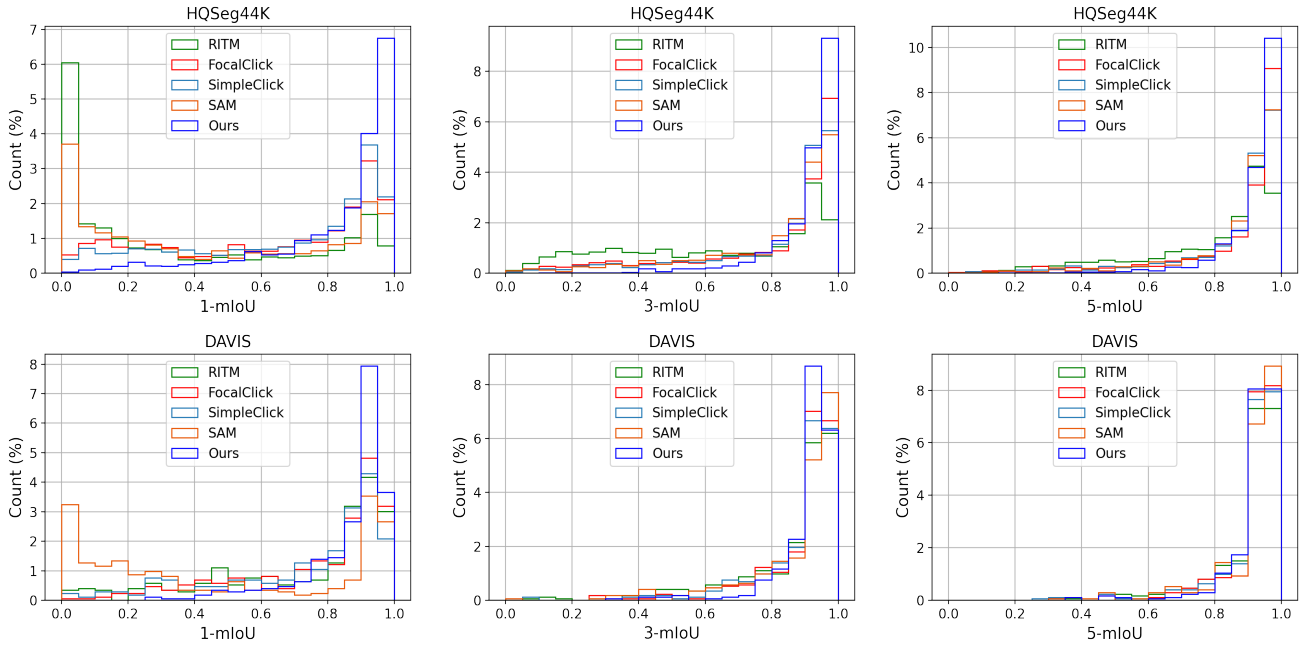


Figure 4. *Histogram analysis* of segmentation IoUs given a predefined k clicks. We report analysis on HQSeg-44K and DAVIS. Compared with the two baselines, our models achieve higher-quality segmentation with fewer failure cases.




Figure 5. *Qualitative results with diverse prompts*. Left: an example from DAVIS. Right: three examples from HQSeg-44K. The results are achieved by a user providing all the prompts using our best-performing model.

assess how our model performs with different types of prompts. In Tab. 3, we compare our method with two strong baselines, SimpleClick and SAM, on DAVIS. We evaluate four common types of prompts: click, box, scribble, and polygon. The metric “1-IoU” indicates the mean IoU of all images with a single click. The outcomes highlight our method’s robust adaptability, establishing it as a solid foundation for future investigations into various prompts. Fig. 5

shows qualitative results on DAVIS and HQSeg-44K with diverse prompts. The supplementary files also include human evaluation videos to give readers an insight into how our method operates in real applications. These videos provide more details on human evaluations involving natural and medical images.




Figure 6. *Failure patterns analysis.* We observed two failure patterns of our method: 1) difficulty with thin structures (as seen on the left) and 2) challenges with cluttered occlusions (as shown on the right). In the case of thin structures, our method tends to overlook fine details of an object, such as missing out on furry textures. In scenarios with cluttered occlusions, our method may struggle to accurately differentiate between the foreground and background.

Method	Click	Box	Scribble	Polygon
	1-mIoU \uparrow	1-mIoU \uparrow	1-mIoU \uparrow	1-mIoU \uparrow
SimpleClick [27]	72.41	69.21	76.63	68.16
SAM [16]	48.66	44.87	30.42	66.49
Ours (SA \times 1)	71.17	70.37	74.70	41.33
Ours (SA \times 2)	73.61	69.08	74.26	64.04

Table 3. *Evaluation of diverse prompts on DAVIS.* Although our method is not trained with diverse prompts, it generalizes reasonably well on different “unseen” visual prompts. We also observe that SAM fails to handle dense visual prompts, such as scribbles. This serves as additional evidence of our hypothesis that dense visual prompts should be represented densely.

4.6. Ablations

We conduct ablation studies to validate the efficacy of our design choices, as detailed in Tab. 4. *No dense fusion* is a model variant that removes the self-attention blocks for dense fusion. The prompt embeddings and the image embeddings are fused only by element-wise addition. The segmentation quality of the model drops significantly without the dense fusion module. *No disk* is a model version in which we represent each click as a point on the image instead of a disk. The model’s performance drops slightly with this modification. *Weak dense fusion* removes one of the two self-attention blocks for dense fusion. The performance also drops slightly. All the model variants are trained on COCO+LVIS for 90 epochs and tested on HQSeg-44K.

5. Limitations

Our dense presentation of visual prompts is more costly than a sparse representation used in specialist models. We believe a downsampled dense representation may alleviate this issue. Our text prompt is tentative and not completely stable, although we are confident that it can be enhanced

Method	5-mIoU \uparrow	NoC90 \downarrow	NoC95 \downarrow	NoF95 \downarrow
No dense fusion	65.34	12.27	15.81	959
No disk	83.72	7.94	12.65	882
Weak dense fusion	85.41	7.47	11.94	731
Full	85.71	7.18	11.52	700

Table 4. *Ablation study on HQSeg-44K [14].* *No dense fusion* is a model variant that removes the self-attention blocks for dense fusion. *No disk* is a model version in which we represent each click as a point on the image instead of a disk. *Weak dense fusion* removes one of the two self-attention blocks for dense fusion.

through further dedication. Our method may fail in some challenging scenarios, as shown in Fig. 6. Finally, our work was confined to employing the ViT-Base as the foundational backbone, notwithstanding the availability of more powerful pretrained alternatives such as ViT-Large and ViT-Huge. However, this limitation does not diminish the significance of our findings. Indeed, our results offer valuable insights into the potential for scaling existing models to more substantial backbones or to accommodate larger datasets.

6. Conclusion

We proposed SegNext for next-generation interactive image segmentation with low latency, high quality, and diverse prompts. We reintroduced the dense representation and fusion of visual prompts, common in specialist models, into the generalist models to facilitate high-quality segmentation. We unified five visual prompts, including clicks, boxes, scribbles, polygons, and masks, on a dense map. We observed that dense representation and fusion of visual prompts are the key design choices contributing to high-quality segmentation. Our method outperformed existing state-of-the-art approaches on challenging benchmarks, both quantitatively and qualitatively.

References

- [1] Ujjwal Baid, Satyam Ghodasara, Suyash Mohan, Michel Bilello, Evan Calabrese, Errol Colak, Keyvan Farahani, Jayashree Kalpathy-Cramer, Felipe C Kitamura, Sarthak Pati, et al. The rsna-asnr-miccai brats 2021 benchmark on brain tumor segmentation and radiogenomic classification. *arXiv preprint arXiv:2107.02314*, 2021. [4](#), [5](#), [6](#)
- [2] Liang-Chieh Chen, George Papandreou, Iasonas Kokkinos, Kevin Murphy, and Alan L Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):834–848, 2017. [2](#)
- [3] Xi Chen, Zhiyan Zhao, Feiwu Yu, Yilei Zhang, and Manni Duan. Conditional diffusion for interactive segmentation. In *ICCV*, pages 7345–7354, 2021. [6](#)
- [4] Xi Chen, Zhiyan Zhao, Yilei Zhang, Manni Duan, Donglian Qi, and Hengshuang Zhao. Focalclick: Towards practical interactive image segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1300–1309, 2022. [1](#), [2](#), [4](#), [5](#), [6](#)
- [5] Ming-Ming Cheng, Niloy J Mitra, Xiaolei Huang, Philip HS Torr, and Shi-Min Hu. Global contrast based salient region detection. *IEEE transactions on pattern analysis and machine intelligence*, 37(3):569–582, 2014. [5](#)
- [6] Daan de Geus and Gijs Dubbelman. Intra-batch supervision for panoptic segmentation on high-resolution images. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3165–3173, 2023. [2](#)
- [7] Henghui Ding, Scott Cohen, Brian Price, and Xudong Jiang. Phraseclick: toward achieving flexible interactive segmentation by phrase and click. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 417–435. Springer, 2020. [2](#)
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [4](#)
- [9] Stephan Gerhard, Jan Funke, Julien Martel, Albert Cardona, and Richard Fetter. Segmented anisotropic sstem dataset of neural tissue. *figshare*, pages 0–0, 2013. [4](#), [5](#), [6](#)
- [10] Agrim Gupta, Piotr Dollár, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, pages 5356–5364, 2019. [4](#)
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [2](#)
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16000–16009, 2022. [4](#)
- [13] You Huang, Hao Yang, Ke Sun, Shengchuan Zhang, Liujuan Cao, Guannan Jiang, and Rongrong Ji. Interformer: Real-time interactive image segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22301–22311, 2023. [1](#), [2](#), [4](#), [5](#), [6](#)
- [14] Lei Ke, Mingqiao Ye, Martin Danelljan, Yifan Liu, Yu-Wing Tai, Chi-Keung Tang, and Fisher Yu. Segment anything in high quality. *arXiv preprint arXiv:2306.01567*, 2023. [1](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [15] Alexander Kirillov, Yuxin Wu, Kaiming He, and Ross Girshick. Pointrend: Image segmentation as rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9799–9808, 2020. [2](#)
- [16] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4026, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [17] Xiang Li, Tianhan Wei, Yau Pun Chen, Yu-Wing Tai, and Chi-Keung Tang. Fss-1000: A 1000-class dataset for few-shot segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2869–2878, 2020. [5](#)
- [18] Zhuwen Li, Qifeng Chen, and Vladlen Koltun. Interactive image segmentation with latent diversity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 577–585, 2018. [3](#)
- [19] Jun Hao Liew, Scott Cohen, Brian Price, Long Mai, and Jiashi Feng. Deep interactive thin object selection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 305–314, 2021. [5](#)
- [20] Guosheng Lin, Anton Milan, Chunhua Shen, and Ian Reid. Refinenet: Multi-path refinement networks for high-resolution semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1925–1934, 2017. [2](#)
- [21] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. [4](#)
- [22] Zheng Lin, Zhao Zhang, Lin-Zhuo Chen, Ming-Ming Cheng, and Shao-Ping Lu. Interactive image segmentation with first click attention. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13339–13348, 2020. [2](#)
- [23] Zheng Lin, Zheng-Peng Duan, Zhao Zhang, Chun-Le Guo, and Ming-Ming Cheng. Focuscut: Diving into a focus view in interactive segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2637–2646, 2022.
- [24] Zheng Lin, Zheng-Peng Duan, Zhao Zhang, Chun-Le Guo, and Ming-Ming Cheng. Knifecut: Refining thin part segmentation with cutting lines. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 809–817, 2022.
- [25] Qin Liu, Zhenlin Xu, Yining Jiao, and Marc Niethammer. isegformer: interactive segmentation via transformers with application to 3d knee mr images. In *International Confer-*

- ence on Medical Image Computing and Computer-Assisted Intervention, pages 464–474. Springer, 2022. 1
- [26] Qin Liu, Meng Zheng, Benjamin Planche, Srikrishna Karanam, Terrence Chen, Marc Niethammer, and Ziyan Wu. Pseudoclick: Interactive image segmentation with click imitation. In *European Conference on Computer Vision*, pages 728–745. Springer, 2022.
- [27] Qin Liu, Zhenlin Xu, Gedas Bertasius, and Marc Niethammer. Simpleclick: Interactive image segmentation with simple vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22290–22300, 2023. 1, 2, 4, 5, 6, 8
- [28] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully convolutional networks for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3431–3440, 2015. 2
- [29] Federico Perazzi, Jordi Pont-Tuset, Brian McWilliams, Luc Van Gool, Markus Gross, and Alexander Sorkine-Hornung. A benchmark dataset and evaluation methodology for video object segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 724–732, 2016. 4
- [30] Lu Qi, Jason Kuen, Tiancheng Shen, Jiuxiang Gu, Wenbo Li, Weidong Guo, Jiaya Jia, Zhe Lin, and Ming-Hsuan Yang. High quality entity segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4047–4056, 2023. 2
- [31] Xuebin Qin, Hang Dai, Xiaobin Hu, Deng-Ping Fan, Ling Shao, and Luc Van Gool. Highly accurate dichotomous image segmentation. In *European Conference on Computer Vision*, pages 38–56. Springer, 2022. 5
- [32] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 3, 4
- [33] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. 2
- [34] Tiancheng Shen, Yuechen Zhang, Lu Qi, Jason Kuen, Xingyu Xie, Jianlong Wu, Zhe Lin, and Jiaya Jia. High quality segmentation for ultra high-resolution images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1310–1319, 2022. 1
- [35] Jianping Shi, Qiong Yan, Li Xu, and Jiaya Jia. Hierarchical image saliency detection on extended cssd. *IEEE transactions on pattern analysis and machine intelligence*, 38(4): 717–729, 2015. 2, 5
- [36] Konstantin Sofiiuk, Ilya A Petrov, and Anton Konushin. Reviving iterative training with mask guidance for interactive segmentation. In *2022 IEEE International Conference on Image Processing (ICIP)*, pages 3141–3145. IEEE, 2022. 1, 2, 3, 4, 5, 6
- [37] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. 2
- [38] Jue Wang, Pravin Bhat, R Alex Colburn, Maneesh Agrawala, and Michael F Cohen. Interactive video cutout. *ACM Transactions on Graphics (ToG)*, 24(3):585–594, 2005. 1
- [39] Jingdong Wang, Ke Sun, Tianheng Cheng, Borui Jiang, Chaorui Deng, Yang Zhao, Dong Liu, Yadong Mu, Mingkui Tan, Xinggang Wang, et al. Deep high-resolution representation learning for visual recognition. *IEEE transactions on pattern analysis and machine intelligence*, 43(10):3349–3364, 2020. 2, 6
- [40] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems*, 34:12077–12090, 2021. 2
- [41] Chuan Yang, Lihe Zhang, Huchuan Lu, Xiang Ruan, and Ming-Hsuan Yang. Saliency detection via graph-based manifold ranking. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3166–3173, 2013. 5
- [42] Yuhui Yuan, Jingyi Xie, Xilin Chen, and Jingdong Wang. Segfix: Model-agnostic boundary refinement for segmentation. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 489–506. Springer, 2020. 2
- [43] Chaoning Zhang, Dongshen Han, Yu Qiao, Jung Uk Kim, Sung-Ho Bae, Seungkyu Lee, and Choong Seon Hong. Faster segment anything: Towards lightweight sam for mobile applications. *arXiv preprint arXiv:2306.14289*, 2023. 1, 3, 5
- [44] Gang Zhang, Xin Lu, Jingru Tan, Jianmin Li, Zhaoxiang Zhang, Quanquan Li, and Xiaolin Hu. Refinemask: Towards high-quality instance segmentation with fine-grained features. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6861–6869, 2021. 2
- [45] Shiyin Zhang, Jun Hao Liew, Yunchao Wei, Shikui Wei, and Yao Zhao. Interactive object segmentation with inside-outside guidance. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12234–12244, 2020. 2
- [46] Zongwei Zhou, Md Mahfuzur Rahman Siddiquee, Nima Tajbakhsh, and Jianming Liang. Unet++: A nested u-net architecture for medical image segmentation. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4*, pages 3–11. Springer, 2018. 2
- [47] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *arXiv preprint arXiv:2304.06718*, 2023. 1, 2, 3, 6