ConCL: Concept Contrastive Learning for Dense Prediction Pre-training in Pathology Images

Jiawei Yang^{1,2,\star}, Hanbo Chen^{1,\star\star}, Yuan Liang², Junzhou Huang³, Lei He², and Jianhua Yao^{1,\star\star}

 ¹ Tencent AI Lab
² University of California, Los Angeles
³ University of Texas at Arlington jiawei118@ucla.edu

Abstract Detecting and segmenting objects within whole slide images is essential in computational pathology workflow. Self-supervised learning (SSL) is appealing to such annotation-heavy tasks. Despite the extensive benchmarks in natural images for *dense* tasks, such studies are, unfortunately, absent in current works for pathology. Our paper intends to narrow this gap. We first benchmark representative SSL methods for dense prediction tasks in pathology images. Then, we propose concept contrastive learning (ConCL), an SSL framework for dense pre-training. We explore how ConCL performs with concepts provided by different sources and end up with proposing a simple dependencyfree concept generating method that does not rely on external segmentation algorithms or saliency detection models. Extensive experiments demonstrate the superiority of ConCL over previous state-of-the-art SSL methods across different settings. Along our exploration, we distill several important and intriguing components contributing to the success of dense pre-training for pathology images. We hope this work could provide useful data points and encourage the community to conduct ConCL pre-training for problems of interest. Code is available at https: //github.com/TencentAILabHealthcare/ConCL.

Keywords: Pathology image analysis \cdot Whole slide image \cdot Self-supervised learning \cdot Object detection \cdot Instance segmentation \cdot Pre-training

1 Introduction

Computational pathology is an emerging area in modern healthcare. More whole slide images (WSIs) are now analyzed by deep learning (DL) models [29]. To alleviate the heavy annotation burden required by DL models, reusing weights from pre-trained models has become a common practice. Besides transferring from fully-supervised models, recent attention has been attracted to self-supervised

^{*} Work done during an internship at Tencent AI Lab.

^{**} Corresponding authors



Figure 1: (a) Example pathology images with tissue class names. (b,c) Comparisons of pre-trained models by fine-tuning on the GlaS [28] and CRAG [9] datasets. The detector is Mask-RCNN [14] with a ResNet-18 backbone [15], and an FPN head[20]. Results are averaged over 5 independent runs.

learning (SSL) methods [12,3,10]. They are annotation-free but can achieve comparable or even better performance when transferring.

The breakthrough of SSL methods starts with contrastive learning [11.34.3, 12.4]. where the most popular task is instance discrimination [34]. It requires a model to discriminate among individual instances. To achieve that, it first defines some positive pairs and negative pairs. It then optimizes a model to maximize the representation similarity between positive pairs and minimize it between negative pairs. Later, more SSL methods based on cross-view prediction are proposed, e.q., [2,10,3,38]. However, these methods are optimized for image-level representations and might be sub-optimal for dense prediction tasks such as object detection and instance segmentation. This motivates works for detection-friendly pre-training methods, e.g., DenseCL [33], InsLoc [37], Self-EMD [22], SCRL [27], and more [16,30,35,36]. Despite many interests raised in the natural image domain for dense prediction problems, such studies, which are of important clinical and practical values, are *absent* in the pathology image domain. Our research is intended to bridge the gap between SSL in natural images and pathology images for dense prediction tasks, as well as to distill the key components to the success of dense pre-training in the pathology data.

To that end, we start by presenting a general **Con**cept **C**ontrastive Learning (ConCL) framework. Rather than contrasting image-level representations [34,3,12], it contrasts "concepts" that mark different local (semantic) regions. ConCL is an abstraction of dense contrasting frameworks that can resemble most concurrent related works. We first benchmark current leading image-level SSL methods and a grid-level dense SSL method (*i.e.*, DenseCL [33]) in two public datasets. We observe a considerable performance gap between DenseCL [33] and the others. These gaps indicate the importance of contrasting densely (grid-level) than roughly (image-level). Then, directed by the performance differences and the characteristics of pathology images, we gradually develop and improve ConCL via a series of explorations. Specifically, we explore: 1) what makes the success

of dense prediction pre-training? 2) what kind of concepts are good for pathology images? The nature of having rich low-level patterns in pathology images (see Fig. 1-(a)) gives some surprising and intriguing results, *e.g.*, a randomly initialized model can group meaningful concepts and help dense pre-training. Along the exploration, we distill several key components contributing to the transferring performance for dense tasks. At the end of exploration, the presented ConCL can surpass various state-of-the-art SSL methods by solid and consistent margins across different downstream datasets, detector architectures, fine-tuning schedules, and pre-training epochs. For example, as shown in Figure 1-(b), the 200-epoch pre-trained ConCL wins all the other methods but with $4 \times$ to $8 \times$ fewer epochs. To summarize, this paper makes the following contributions:

- It makes one of the earliest attempts to systematically study and benchmark self-supervised learning methods for dense prediction problems in pathology images, which are of high practical and clinical interest but, unfortunately, *absent* in existing works. We hope this work could narrow the gap between studies in natural images and pathology images.
- It presents ConCL, an SSL framework for dense pre-training. We show how ConCL performs with concepts provided by different sources and find that a randomly initialized model could learn semantic concepts and improve itself without expert-annotation or external algorithms while achieving competitive, if not the best, results.
- It shows how important the *dense* pre-training is in pathology images for dense tasks and provides some intriguing observations that could contribute to other applications such as few-shot and semi-supervised segmentation and detection, or more, in pathology image analysis or beyond.

We hope this work could provide useful data points and encourage the community to conduct ConCL pre-training for problems of interest.

2 Related work

Contrastive learning. The success of deep learning is mainly attributed to mining a large amount of data. When limited data is provided for specific tasks, an alternative is to transfer knowledge by re-using pre-trained models [8,13]. SSL methods learn good pre-trained models from label-free pretext tasks, *e.g.*, colorization [39,40], denoising [31], and thus attract much attention. Recently, contrastive learning [12,4,3,24,34,2], a typical branch of SSL, has made significant progress in many fields, where instance discrimination [11,34,12,4,3] serves as a pretext task. It requires a model to discriminate among individual instances, *i.e.*, image-level representations [34]. MoCo [12,4] and SimCLR[3] are two representatives. Specifically, they generate two views of the same image via random data augmentations (*e.g.*, color jittering, random cropping) and mark them as a positive pair. Then, views from other different images are marked as negative instances or pairs. After that, they learn embeddings by maximizing the similarity between the representations of positive pairs while minimizing it between

the representations of negative pairs. Later methods combine contrasting with clustering, e.g., SwAV[2] proposes to contrast views' cluster assignments, and PCL [19] contrasts instances with cluster prototypes.

Dense prediction pre-training. Despite their success in transferring to classification tasks, good image-level representations do not necessarily result in better performance in dense prediction tasks. Therefore, recent efforts have been made for dense prediction pre-training. Related works are mostly concurrent [33,37,27,30,35,36,22,16]. Among them, DenseCL [33] learns the correspondence among pixels of a positive pair and optimizes a pairwise contrastive loss at a pixel level, yielding a dense contrasting behavior. Self-EMD [22] does dense predicting in a non-contrastive manner as in BYOL [10], *i.e.*, predicting a grid-level feature vector from one view when given its counterpart from another (positive) view. SCRL [27] argues the importance of spatially consistent representations, so it maximizes the similarity of box region features in the intersected area. The most relevant works concurrent to ours are [16,30]. They also optimize contrastive loss over mask-averaged representations. Those masks are generated by external algorithms that are successful for natural images, e.g., Felzenszwalb-Huttenlocher algorithm [7], or models, e.g., MCG[1], BASNet[25], and DeepUSPS[23]. However, the success of such mask generators is unfortunately unverified in pathology images. In this paper, we provide some of their empirical results. Their different performances yield the disparity between natural and pathology images, from which we are motivated to propose a dependency-free concept mask generator. It directly bootstraps the structural concepts inherent in pathology images, learns from *scratch*, and has better potential.

SSL in pathology images. Studying SSL methods in pathology images is still at an early stage. In addition to studies on natural images, SimCLR [3] is also studied and benchmarked for classification, regression, and segmentation tasks in pathology images [6]. Some domain-specific self-supervised pretext tasks, *e.g.*, magnification prediction, JigMag prediction, and hematoxylin channel prediction, are proposed and studied [18]. However, despite interest raised in natural images for dense problems, existing works have not studied, to our knowledge, detection/segmentation-friendly SSL methods in pathology images. Our work aims to bridge this gap and provide our exploration roadmap toward better dense prediction performance for pathology images.

3 Method

3.1 Preliminary: Instance Contrastive Learning

MoCo[12] abstracts the instance discrimination task as a dictionary look-up problem. Specifically, for each encoded query q, there is a set of encoded keys $\{k_0, k_1, k_2, ...\}$ in a dictionary. The instance discrimination task is to pull closer q and its matched positive key k_+ in the dictionary while spreading q away from all



Figure 2: ConCL overview. ConCL has three steps: (1) Given a query view x_q and a key view x_k , their union region is cropped as a reference view x_r . ConCL obtains concept proposals by processing x_r with a "concept generator." (2) For the shared concepts, ConCL computes their representations via masked average pooling (MAP). (3) ConCL optimizes concept contrastive loss (Eq. (2)), and enqueues the concept prototypes from the key encoder to the concept queue.

other negative keys k_{-} . When using the dot-product as similarity measurement, a form of contrastive loss function based on InfoNCE[24] becomes:

$$L_q = -\log \frac{\exp(q \cdot k_+/\tau)}{\exp(q \cdot k_+/\tau) + \sum_{k_-} \exp(q \cdot k_-/\tau)}$$
(1)

where τ is a temperature hyper-parameter [34]. Queries q and keys k are computed by a query encoder and a key encoder, respectively [12,4]. Formally, $q = h(\text{GAP}(f_5(x_q)))$, where h is a MLP projection head as per [3]; GAP(·) denotes global-average-pooling, and $f_5(x)$ represents the outputs from the stage-5 of a ResNet [15]. Keys k are computed similarly using the key encoder. In MoCo [12], the negative keys are stored in a queue to avoid using large batches [3].

3.2 Concept Contrastive Learning

Instance contrastive methods [3,12,34] do well in discriminating among imagelevel instances, but dense prediction tasks usually require discriminating among local details, *e.g.*, object instances or object parts. We abstract such local details, or say, fine-grained semantics as "concepts." A concept does not necessarily represent an object. Instead, any sub-region in an image could be a concept since it contains certain different semantics. From the perspective of dense prediction, it is desirable to build concept-sensitive representations. For example, one WSI patch usually contains multiple small objects, *e.g.*, nucleus, glands, and multiple texture-like tissues, *e.g.*, mucus [29,17]. To successfully detect and segment objects in such images, models need to learn more information from local details. To

this end, we propose a simple but effective framework — Concept Contrastive Learning (ConCL). Figure 2 shows its overview, which we elaborate on below.

Concept discrimination. We first define a pretext task named concept discrimination. Similar to instance discrimination [34,11], concept discrimination requires a model to discriminate among the representations of the same but augmented concepts and the representations of different concepts. We formulate concept discrimination by extending the instance-level queries and keys to concept-level. Specifically, given an encoded query concept q^c and a set of encoded key concepts $\{k_0^c, k_1^c, k_2^c, ...\}$, we derive concept contrastive loss as:

$$L_{c} = -\log \frac{\exp(q^{c} \cdot k_{+}^{c}/\tau)}{\exp(q^{c} \cdot k_{+}^{c}/\tau) + \sum_{k^{c}} \exp(q^{c} \cdot k_{-}^{c}/\tau)}$$
(2)

where τ is the same temperature parameter and k_{-}^{c} are keys in the concept queue — the queue to store concept representations. This objective brings representations of different views of the same concept closer and spreads representations of views from different concepts apart.

Concept mask proposal. We use masks to annotate fine-grained concepts explicitly. Assume a mask generator is given, as diagramed at the bottom of Figure 2; we first pass a reference view x_r , defined as the circumscribed rectangle crop of the union of two views, into the mask generator to obtain a set of concept masks — $\mathcal{M}_r = \{m_i\}_{i=1}^K$, where K is the number of concepts. Since the reference view contains both the query view and the key view, their concept masks \mathcal{M}_q and \mathcal{M}_k are immediately obtained if we restore them in the reference view. Then, we derive concept representations in both views by masked average pooling (MAP) with resized concept masks. Specifically, we compute $q^c = h$ (MAP ($f_5(x_q), m_c$)) and k^c similarly, where MAP $(z, m) = \sum_{ij} m_{ij} \cdot z_{ij} / \sum_{ij} m_{ij}$, and $z \in \mathbb{R}^{CHW}$ denotes feature maps, $m \in \{0, 1\}^{HW}$ is a binary indicator for each concept. Here, only the shared concepts in both views are considered, *i.e.*, $m_c \in \mathcal{M}_q \cap \mathcal{M}_k$.

Our analysis hereafter focuses on 1) What makes the success of dense prediction pre-training? 2) What kind of concepts are good *for pathology images*? Different answers to these two questions reveal the characteristics of pathology images and the disparity between natural and pathology images, as we explore in Section 4. Below, we first introduce the benchmark pipeline and setups.

3.3 Benchmark Pipeline

Despite the extensive benchmarks in natural images for dense tasks, to our knowledge, such studies are unfortunately *absent* in current works for pathology. Note that studying SSL methods in pathology images is still at an early stage. Most current works focus on employing image-level SSL methods for classification tasks. Orthogonal to theirs, we investigate a wider range of SSL methods for object detection and instance segmentation tasks, which are of high clinical value. We hope our work could provide useful data points for future work.

Implementations. For implementation details, please refer to Appendix B. We briefly introduce the datasets here and elaborate on them in Appendix C:

- Pre-training dataset. We use NCT-CRC-HE-100K[17] dataset, referred to as NCT, for pre-training. It contains 100,000 non-overlapping patches extracted from hematoxylin and eosin (H&E) stained colorectal cancer and normal tissues. All images are of size 224×224 at 0.5 MPP ($20 \times$ magnification). We randomly choose 80% of NCT to be the pre-training dataset.
- Transferring dataset. We use two public datasets, the gland segmentation in pathology images challenge (GlaS) dataset [28] and the colorectal adenocarcinoma gland (CRAG) dataset [9], and follow their official train/test splits for evaluation. GlaS [28] collects images of 775×522 from H&E stained slides with object-instance-level annotation; the images include both malignant and benign glands. CRAG [9] collects 213 H&E stained images taken from 38 WSIs with a pixel resolution of 0.55µm/pixel at 20× magnification. Images are mostly of size 1512×1516 with object-instance-level annotation. We study the performance of object detection and instance segmentation.

Experimental setup. We pre-train all the methods on the NCT training set for 200 epochs. For ConCL pre-training, we warm up the model by optimizing instance contrastive loss (Eq. (1)) for the first 20 epochs and switch to concept contrastive loss (Eq. (2)). Then, we use the pre-trained backbones to initialize the detectors, fine-tune them on the training sets of transferring datasets, and test them in the corresponding test sets. Unless otherwise specified, we run all the transferring experiments 5 times and report the averaged performance.

4 Towards Better Concepts: a Roadmap

In this section, we first benchmark some popular state-of-the-art SSL methods for dense pathology tasks. Then, we start with DenseCL [33] and derive better concepts along the way, directed by the questions raised in the previous section.

4.1 Benchmarking SSL methods for Dense Pathology Tasks

Benchmark results. Table 1 (baselines and prior SSL arts) shows the transferring performance for GlaS dataset (left columns) and CRAG dataset (right columns), respectively. We report results using 200-epoch pre-trained models and a $1 \times$ fine-tuning schedule. On the GlaS dataset [28], we observe that the gap between training from randomly initialized models and training from supervised pre-trained models is relatively smaller compared to those in the natural image domain [5,4,10,3]. Nonetheless, state-of-the-art SSL methods all exceed supervised pre-training, meeting the same expectation as in natural images. Yet, on the CRAG dataset [9], most of the pre-trained models, including both the selfsupervised ones and the supervised one, fail to achieve competitive performance

		GlaS			CRAG				
	N. (1 1	Detect		Segment		Detect		Segment	
Category	Methods	AP^{bb}	AP^{bb}_{75}	AP^{mk}	AP^{mk}_{75}	AP^{bb}	AP^{bb}_{75}	$AP^{m\bar{k}}$	AP^{mk}_{75}
Decelines	Rand. Init.	49.8	57.3	52.1	60.7	51.1	57.0	50.6	57.3
Dasennes	Supervised	50.2	56.9	53.2	62.1	49.2	55.2	49.4	55.0
	SimCLR[3]	50.7	56.9	53.6	62.7	49.2	54.8	49.1	54.7
	BYOL[10]	50.9	57.7	53.9	62.6	49.9	55.8	49.3	55.3
Sec. 4.1	$PCL-v2^{\dagger}$ [19]	49.4	55.9	51.9	61.0	51.0	56.6	50.5	56.7
Prior SSL arts	MoCo-v1[12]	50.0	56.2	52.1	59.9	47.2	51.1	47.5	52.0
	MoCo-v2[4]	52.3	60.0	55.3	65.0	50.0	55.7	50.3	56.8
	DenseCL[33]	53.9	62.0	56.5	66.2	52.3	58.2	52.2	59.8
Our differently	instantiated ConCLs:								
Sec. 4.2	(1) g -ConCL(s=3)	54.9	64.1	57.1	66.3	55.4	62.3	54.4	62.0
	(2) g-ConCL($s=5$)	55.4	65.2	57.4	67.2	55.5	62.7	54.6	62.2
Gild concepts	(3) g-ConCL(s= 7)	54.9	63.8	57.0	66.5	55.3	62.5	54.7	62.6
Sec. 4.3	(4) fh-ConCL($s=50$)	55.8	65.6	58.3	68.8	54.8	60.7	54.1	60.7
Natural-image	(5) fh-ConCL($s=500$)	56.2	65.9	57.7	67.9	54.7	61.9	53.8	60.5
priors concepts	(6) bas-ConCL	56.1	66.1	58.1	68.1	54.2	61.1	53.4	60.8
Sec. 4.4									
Bootstrapped	(7) b-ConCL (f_4)	56.8	66.2	58.7	68.9	55.1	62.2	54.1	61.4
concepts	(8) b-ConCL (f_5)	56.1	65.6	57.8	67.7	56.5	63.3	55.3	62.9
Table 1: Main results of object detection and instance segmentation.									

 AP^{bb} : bounding box mAP, AP^{mk} : mask mAP.

compared to training from randomly initialized weights. The only exception is DenseCL [33], a dense contrasting method.

Among the image-level SSL methods, MoCo-v2 [4] performs the best in GlaS and the second-best in CRAG. Enhanced by dense contrasting, DenseCL [33] achieves the best results in both datasets. It should be emphasized that DenseCL [33] gets + 1.6 AP^{bb} for GlaS by using grid-level contrasting. This demonstrates the importance of designing dense pre-training frameworks when transferring to dense tasks since all the stragglers are only optimized for image-level representations. Thus, we here conclude *dense contrasting matters*.

4.2 Correspondence matters

From the previous section, we find dense contrasting is favored in both natural and pathology images, where DenseCL [33] all achieves top performance. The next question is: can we improve the dense contrasting framework? To answer it, we first summarize the overall pipeline of DenseCL [33]. DenseCL computes the dense representations of two views without global average pooling, *i.e.*, $f_5(x_q)$, $f_5(x_k)$, and passes them to a dense projection head to obtain final grid features of size $\mathbb{R}^{128\times7\times7}$. Then it sets the most similar (measured by cosine similarity) grids in two views as positive pairs. As such, the correspondence of positive pairs is learned. However, the reliability of learned correspondence remains questionable and would affect the quality of learned representations.



Figure 3: **Concept descriptors.** (a) Tissue concept illustration. (b) Grid concepts (s: grid number). (c-d) FH concepts (s: scale). (e) Binary saliency concepts, obtained from BASNet [25]. (f-h) Clustering concepts (f_i : ResNet output stage). The image is resized to 448×448 for better visualization.

To address that, we instantiate DenseCL [33] in ConCL by regarding the grid-prior as a form of concept, as shown in Figure 3-(b). We denote this ConCL instance as g-ConCL. Compared with DenseCL [33] (learned matching), ConCL naturally restores the positive correspondence from a reference view (precise matching Fig. 2- x_r). Table 1-(1-3) compares the original DenseCL [33] and ConCL-instantiated g-ConCL. The results indicate that g-ConCL with precise correspondence can boost DenseCL [33] by a large margin. Even with the simplest form of concepts, g-ConCL already has topped entries above it in Table 1. We believe other dense pre-training methods that learn the matching between grids, *e.g.*, *Self*-EMD [22], should perform similarly to DenseCL [33], and g-ConCL could outperform them. Thus, we conclude that *correspondence matters*.

4.3 Natural Image Priors in Pathology Images

ConCL is a general framework for using masks as supervision to discriminate concepts. Some previous works in natural image [42,16,41,30,32] also combines masks with contrastive learning, where the masks are provided by ground truth annotation [42,32,16], or supervised/unsupervised pseudo-mask generation [16,41,30]. The mask generators can be graph-based (*e.g.*, Felzenszwalb-Huttenlocher algorithm [7]), MCG [1], or other saliency detection models [25,23] trained on designated natural image datasets. However, those methods were originally proposed for nature images, and their success for pathology images remains unknown.

Here we instantiate ConCL by using Felzenszwalb-Huttenlocher (FH) algorithm [7] and BASNet [25] as concept generators, dubbed as fh-ConCL and bas-ConCL, respectively. FH [7] is a conventional graph-based segmentation algorithm that relies on local neighborhoods, while BASNet [25] is a deep neural network pre-trained on a curated saliency detection dataset, which only contains daily natural objects. We use these two as representatives to study if these natural image priors win twice in both natural and pathology images.

Specifically, we use the FH algorithm in the scikit-image package and set both "scale" and "size" hyper-parameters to s. We use the pre-trained BASNet provided by [25]. Figure 3-(c-e) shows some examples. Table 1 reports the results.

It is not surprising that the BASNet [25] cannot generate decent concept masks (Fig. 3-(e)) for pathology images since it is pre-trained on curated saliency detection datasets. What is surprising is that bas-ConCL does yield satisfactory

results (Tab. 1-(6)). Similar observations are also found in fh-ConCLs (Tab. 1-(4,5)) that though the generated concept masks are coarse-grained, the resulted transferring performances are unexpectedly good. After inspecting more examples, we find that the generated masks maintain high coherence and integrity despite their coarse-grained nature. That said, each concept contains semantic-consistent objects or textures. For example, Figure 3-(d,e) can be seen as special cases of Figure 3-(a) that merge fine-grained semantics with coarse-grained ones. This property makes the major difference between fh-/bas-ConCLs and g-ConCLs, where the grid-concepts are less likely to have coherent semantics.

Thus, we here conclude that *coherence matters* and natural image priors also work in pathology images, though they mostly provide coarse-grained concepts.

4.4 Pathology Image Priors in Pathology Images

Can we obtain concept masks away with natural image priors? External dependency is not always wanted and sometimes may fail to provide the desired masks (e.g., Fig. 3-(e)). We thus task ourselves to find a dependency-free concept proposal method. One of the key characteristics in pathology images is that they have rich low-level patterns and tissue structures. Can we use that prior instead?

Figure 3-(f-h) shows the clustering visualization from intermediate feature maps generated by a 10-epoch warmed-up MoCo-v2 [4]. Thanks to the rich structural patterns in pathology images, we find that simply clustering over the feature maps provided by a barely trained model can already generate meaningful structural concept proposals. We thus build upon this "free lunch" and use a "bootstrap your own *perception*" mechanism that is similar to the "bootstrap your own latent" mechanism in BYOL [10]. ConCL generates concept proposals from the momentum key encoder's perception while simultaneously improving and refining it via the online query encoder, leading to a "bootstrapping" behavior. Thus, we denote such ConCL as bootstrapped-ConCL (b-ConCL). We provide an additional introduction to BYOL and "bootstrapping" in Appendix A.

b-ConCL. The concept generator is now instantiated as a KMeans grouper. We first pass the reference view x_r to the key encoder to obtain a reference feature map from ResNet stage-*i*: $f_i(x_r) \in \mathbb{R}^{CHW}$. Then, we apply K-Means to group K underlying concepts. b-ConCL relies on neither external segmentation algorithms nor designated saliency detection models for natural images.

Our default setting is K = 8, and clustering from f_4 or f_5 . We postpone the study of hyper-parameters, *i.e.*, the number of clusters in KMeans, and the clustering stage f_i to Section 5.2 and report the main results in Table 1-(7,8). We find b-ConCL tops other entries. Compared to MoCo-v2 [4], our direct baseline, b-ConCL outperforms it by +4.5 AP^{bb} and +3.1 AP^{mk}. Moreover, b-ConCL obtains more gains in terms of AP₇₅ (+6.2 AP^{bb}₇₅, +3.7 AP^{mk}₇₅) compared to MoCo-v2 [4], which means it improves MoCo-v2 [4] by more accurate bounding box regression and instance mask prediction. This aligns with our motivation for ConCL since discriminating local concepts helps shape object borders.

Detector	Drotroin	GlaS D	etection	CRAG Detection		
Detector	rietiam	AP^{bb}	AP_{75}^{bb}	AP^{bb}	AP_{75}^{bb}	
	Rand. Init.	52.9	59.9	49.4	54.2	
MaskRCNN-C4	Supervised	49.1 (-3.8)	55.1(-4.8)	46.1 (-3.3)	50.6 (-2.3)	
	MoCo-v2 [4]	53.6(+0.7)	61.8(+1.9)	48.3 (-1.1)	52.6(-1.6)	
	b-ConCL	55.8(+2.9)	63.6(+3.7)	49.8(+0.4)	54.3(+0.1)	
	Rand. Init.	49.8	57.3	51.1	57.0	
	Supervised	50.2(+0.4)	56.9(-0.4)	49.2 (-1.9)	55.2(-1.8)	
Masknum-FFI	MoCo-v2 [4]	52.3(+2.5)	60.0(+2.7)	50.0 (-1.1)	55.7(-1.3)	
	b-ConCL	56.8(+7.0)	66.2(+8.9)	55.1(+4.0)	62.2(+5.2)	
	Rand. Init.	46.4	51.0	45.2	47.6	
RetinaNet	Supervised	44.7 (-1.7)	48.4(-2.6)	43.1 (-2.1)	44.8(-2.8)	
	MoCo-v2 [4]	47.2(+0.8)	50.9 (-0.1)	43.1 (-2.1)	43.8(-3.8)	
	b-ConCL	52.6(+6.2)	58.6(+7.6)	48.4(+3.2)	51.9(+4.3)	

ConCL: Concept Contrastive Learning for Dense Pre-training in Pathology

Table 2: **Detection performance using different detectors.** Results are averaged over 5 trials.

Closing remarks. So far, we have included: i) dense contrasting matters; ii) correspondence matters; iii) coherence matters; iv) natural image priors, though they might only provide coarse-grained concepts, work in pathology images as well; and find v) a randomly initialized or barely trained convolutional neural network, thanks to the rich low-level patterns in pathology images and good network initialization, can generate good proposals that are *dense*, *fine-grained* and *coherent*, as shown in Figure 3. Though the coarse-grained concepts generated from natural image priors could also help tasks in our studied benchmarks, they might underperform when a fine-grained dense prediction task is given. We hope our closing remarks could be intriguing and guide future works in designing dense pre-training methods for pathology images and beyond.

5 More Experiments

In the previous section, we have explored how we can obtain concepts, what concepts are good, and find b-ConCL to be the best. We here conduct more experiments to study b-ConCL. Some visual comparisons are in Appendix D.

5.1 Robustness to Transferring Settings

Transferring with different detectors. Here we investigate the transferring performance with other detectors, *i.e.*, Mask-RCNN-C4 (C4) [26] and RetinaNet [21]. RetinaNet is a single-stage detector. It uses ResNet-FPN backbone features as Mask-RCNN-FPN but directly generates predictions without region proposal [26]. C4 detector adopts a similar two-stage fashion as Mask-RCNN but uses the outputs of the 4-th residual block as backbone features and re-targets the 5-th block to be the detection head instead of building a new one. These three representative detectors evaluate pre-trained models under different detector architectures. Results together with Mask-RCNN-FPN's are shown in Table 2. b-ConCL performs the best with all three detectors in both datasets. Notably,

training from scratch (Rand. Init.) is one of the top competitors when the C4 detector is used. We conjecture that the pre-trained models are possibly overfitted to their pretext tasks in their 5-th blocks and thus are harder to be tuned than a randomly initialized 5-th block. In CRAG detection, only b-ConCL pretrained models consistently outperform randomly initialized models. In addition, the most significant gap between MoCo-v2[4] and b-ConCL is found in the RetinaNet detector [21]. As also noted by [22], RetinaNet [21] is a single-stage detector, where the local representations from the backbone become more important than other two-stage detectors since results are directly predicted from them. b-ConCL is tasked to discriminate local concepts, and subsequently, the learned representations could be better than other pre-training methods here.

Transferring with different schedules. To investigate if b-ConCL's lead could persist with longer fine-tuning, we fine-tune Mask-RCNN-FPN with $0.5 \times$, $1 \times$, $2 \times$, $3 \times$, and $5 \times$ schedules. Table 3 shows the results. b-ConCL maintains its noticeable gains in longer schedules in both datasets, *e.g.*, b-ConCL achieves 56.2 mAP with a $0.5 \times$ schedule, which is better than MoCo-v2 [4] with a $5 \times$ schedule but costs $10 \times$ less fine-tuning time. Similar observations are also found in CRAG, where the gap between b-ConCL and MoCo-v2 [4] becomes larger (see Δ row). Together, these results confirm b-ConCL's superiority across different fine-tuning schedules.

5.2 Ablation Study

In this section, we ablate the key factors in b-ConCL. Our default setting clusters K = 8 concepts from ResNet stage-4 ($f_4(\cdot)$). Since b-ConCL is built on MoCo-v2 [4], we use it as our direct baseline for comparisons.

Concept loss weight λ . We here study the generalized concept contrastive loss: $L = (1 - \lambda)L_q + \lambda L_c$, where $\lambda \in [0, 1]$ is a concept loss weight parameter. It shows a natural way to combine concept contrastive loss with instance contrastive loss. We start by asking whether instance contrastive loss is indispensable during the training process of b-ConCL. We alter the concept loss

	GlaS dataset				CRAG dataset					
Method	Fine-tuning schedule			Fine-tuning schedule						
	$0.5 \times$	$1 \times$	$2 \times$	$3 \times$	$5 \times$	$0.5 \times$	$1 \times$	$2 \times$	$3 \times$	$5 \times$
Rand. Init.	49.1	49.8	51.4	51.8	52.7	50.2	51.1	51.9	52.4	52.8
Supervised	48.6	50.2	51.4	52.7	54.0	50.0	49.2	50.5	50.1	50.3
MoCo-v2[4]	51.4	52.3	53.7	54.2	55.7	50.2	50.0	50.2	50.8	51.8
b-ConCL	56.2	56.8	57.7	58.3	59.0	54.8	55.1	55.4	55.6	56.0
Δ	+4.8	+4.5	+4.0	+4.1	+3.3	+4.6	+5.1	+5.2	+4.8	+4.2

Table 3: Detection performance under different fine-tuning schedules. Results other than $1 \times$ schedule are averaged over 3 runs. Δ row shows b-ConCL's improvement over MoCo-v2. We report AP^{bb} here.

``	GlaS		CRAG			V	GlaS		CRAG	
λ	AP^{bb}	AP_{75}^{bb}	AP^{bb}	AP_{75}^{bb}	_	K	AP^{bb}	AP_{75}^{bb}	AP^{bb}	AP_{75}^{bb}
0.0	52.3	60.0	50.0	55.7		1	52.3	60.0	50.0	55.7
0.1	53.6	61.1	50.5	55.9		2	54.5	64.1	52.9	60.1
0.3	53.6	61.8	51.7	57.1		4	55.6	64.7	53.4	59.7
0.5	53.6	61.8	51.3	57.0		6	56.3	65.1	53.7	60.2
0.7	55.2	64.1	53.1	59.9		8	56.8	66.2	55.1	62.2
0.9	56.0	65.1	53.6	59.6		10	57.0	66.0	55.1	61.0
1.0	56.8	66.2	55.1	62.2		12	57.4	66.2	54.2	60.1
1.0\w	7. 56.1	66.2	54.0	60.6	-	16	55.7	65.3	54.5	61.3

Pretrain

MoCo.v2

b-ConCL

Rand.

ConCL: Concept Contrastive Learning for Dense Pre-training in Pathology

(a) Concept loss weight.

ν	Gl	aS	CRAG			
n	AP^{bb}	AP_{75}^{bb}	AP^{bb}	AP_{75}^{bb}		
None	52.3	60.0	50.0	55.7		
$f_1(\cdot)$	55.0	65.1	53.3	60.0		
$f_2(\cdot)$	55.0	64.7	53.7	60.4		
$f_3(\cdot)$	56.2	66.4	53.0	59.6		
$f_4(\cdot)$	56.8	66.2	<u>55.1</u>	62.2		
$f_5(\cdot)$	56.1	65.6	56.5	63.3		

57.3Sup. 50.256.947.954.260.0

 AP^{bb}

49.8

52.3

56.8

(b) Number of concepts.

GlaS Detection ResNet-18

 AP_{75}^{bb}

66.2

ResNet-50

 AP_{75}^{bb}

56.1

60.5

65.9

 AP^{bb}

49.9

53.1

57.0

13

(c) Clustering stage.

(d) Backbone capacities.

Table 4: Ablation Study. We study the effect of different hyper-parameters to b-ConCL. Default settings are marked in gray and MoCo-v2 baselines are marked by gray. In (a), "\w." means no warm-up.

weight λ , and Table 4a reports the results. We see a monotonically increasing performance as λ increases in both datasets, which emphasizes the importance of concept loss. When no warm-up is used (last row in Tab. 4a), only a slight performance drop is observed, meaning that warm-up is not the key component of b-ConCL. Warming-up with instance loss (Eq. (2)) is a special case of b-ConCL, where at the early training stage, each instance is regarded as a concept, and we then gradually increase the number of concepts as training goes on. Thus, the overall findings in this ablation support b-ConCL's advance over MoCo-v2 [4].

Number of concepts K. Here, we investigate how the number of concepts clustered during pre-training affects performance in downstream tasks. We report the results of different K in Table 4b. b-ConCL performs reasonably well when $K \ge 4$, with most of performance peaking at K = 8. This demonstrates the robustness of b-ConCL to the choice of K. Note that the best performance for the GlaS dataset is higher than our default setting and outperforms all entries in Table 1, showing the potential room for b-ConCL.

Where to group $f_i(\cdot)$. b-ConCL groups concepts from a model's intermediate feature maps. Our default setting uses feature maps from stage-4 of a ResNet [15], denoted as $f_4(\cdot)$. We now ablate this choice in Table 4c. Clustering concepts from $f_4(\cdot)$ and $f_5(\cdot)$ works similarly well across two datasets. We choose $f_4(\cdot)$ as the default since it achieves top two performance in both datasets under both metrics. Besides, b-ConCL exceeds MoCo-v2 [4], whichever stage it groups concepts from. This again confirms the effectiveness and robustness of b-ConCL.

Longer pre-training. We compare the pre-training efficiency of different SSL methods w.r.t. training epochs in Figure 1-(b,c) with the numerical results in Appendix D. Interestingly, we find SimCLR [3] and BYOL [10] fail to benefit from longer pre-training. This shows the disparity between pathology image data and natural image data. In the latter field, a monotonically increasing performance in downstream tasks is usually observed [16,12,10,2,3]. For MoCo-v1/v2 [12,4], DenseCL [33] and our b-ConCL, we observe the performance consistently improves as the pre-training epoch increases in GlaS dataset [28]. Note that the 200-epoch b-ConCL surpasses the 800-epoch MoCo-v2 [4] and DenseCL[33] by a large margin (Fig. 1-(b)). In the CRAG dataset, we observe all pre-training methods saturate and achieve the best performance in around 200-epoch and 400-epoch. Among them, b-ConCL is still the best (Fig. 1-(c)).

Larger model capacity. Table 4d shows the results of using a larger backbone, ResNet-50. b-ConCL maintains its leading position. For consistency to the previous ablation, a $1 \times$ schedule is also used here, which could put ResNet-50 at a disadvantage since it has more parameters to tune in a relatively short schedule.

6 Conclusion and Broader Impact

In this work, we have benchmarked some of the current SSL methods for dense tasks in pathology images and presented the ConCL framework. Along our exploration, we have distilled several key components to the success of transferring to dense tasks: i) dense contrasting matters, ii) correspondence matters, iii) coherence matters, and more. Finally, we ended up with a dependency-free concept generator that directly bootstraps the underlying concepts inherent in the data and learns from *scratch*. It was shown to be robust and competitive.

While our initial results are presented only for pre-training and fine-tuning, many applications could embrace ConCL. One example is to combine it with few-shot detection or segmentation, where clustering from feature arrays can be an approach for mining latent objects. Another example can be semi-supervised learning, where ConCL can be used as an additional branch for unlabeled data. *Beyond* pathology image analysis, we also hope ConCL would help in speech or tabular data, where little priors can be used. Unsupervised clustering in representation space is likely to be modality-agnostic. Learned by using contrastive learning and clustering, fine-grained "concepts" could also be mined from those data modalities.

References

- Arbeláez, P., Pont-Tuset, J., Barron, J.T., Marques, F., Malik, J.: Multiscale combinatorial grouping. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 328–335 (2014)
- Caron, M., Misra, I., Mairal, J., Goyal, P., Bojanowski, P., Joulin, A.: Unsupervised learning of visual features by contrasting cluster assignments. arXiv preprint arXiv:2006.09882 (2020)
- Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
- Chen, X., Fan, H., Girshick, R., He, K.: Improved baselines with momentum contrastive learning. arXiv preprint arXiv:2003.04297 (2020)
- Chen, X., He, K.: Exploring simple siamese representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15750–15758 (2021)
- Ciga, O., Xu, T., Martel, A.L.: Self supervised contrastive learning for digital histopathology. arXiv preprint arXiv:2011.13971 (2020)
- Felzenszwalb, P.F., Huttenlocher, D.P.: Efficient graph-based image segmentation. International journal of computer vision 59(2), 167–181 (2004)
- Girshick, R., Donahue, J., Darrell, T., Malik, J.: Rich feature hierarchies for accurate object detection and semantic segmentation. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 580–587 (2014)
- Graham, S., Chen, H., Gamper, J., Dou, Q., Heng, P.A., Snead, D., Tsang, Y.W., Rajpoot, N.: Mild-net: Minimal information loss dilated network for gland instance segmentation in colon histology images. Medical image analysis 52, 199–211 (2019)
- Grill, J.B., Strub, F., Altché, F., Tallec, C., Richemond, P.H., Buchatskaya, E., Doersch, C., Pires, B.A., Guo, Z.D., Azar, M.G., et al.: Bootstrap your own latent: A new approach to self-supervised learning. arXiv preprint arXiv:2006.07733 (2020)
- Hadsell, R., Chopra, S., LeCun, Y.: Dimensionality reduction by learning an invariant mapping. In: 2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06). vol. 2, pp. 1735–1742. IEEE (2006)
- He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9729–9738 (2020)
- He, K., Girshick, R., Dollár, P.: Rethinking imagenet pre-training. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4918–4927 (2019)
- He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of the IEEE international conference on computer vision. pp. 2961–2969 (2017)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Hénaff, O.J., Koppula, S., Alayrac, J.B., Oord, A.v.d., Vinyals, O., Carreira, J.: Efficient visual pretraining with contrastive detection. arXiv preprint arXiv:2103.10957 (2021)
- Kather, J.N., Halama, N., Marx, A.: 100,000 histological images of human colorectal cancer and healthy tissue (Apr 2018). https://doi.org/10.5281/zenodo.1214456, https://doi.org/10.5281/zenodo.1214456

- 16 J. Yang et al.
- Koohbanani, N.A., Unnikrishnan, B., Khurram, S.A., Krishnaswamy, P., Rajpoot, N.: Self-path: Self-supervision for classification of pathology images with limited annotations. IEEE Transactions on Medical Imaging (2021)
- Li, J., Zhou, P., Xiong, C., Hoi, S.C.: Prototypical contrastive learning of unsupervised representations. arXiv preprint arXiv:2005.04966 (2020)
- Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2117–2125 (2017)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: Proceedings of the IEEE international conference on computer vision. pp. 2980–2988 (2017)
- 22. Liu, S., Li, Z., Sun, J.: Self-emd: Self-supervised object detection without imagenet. arXiv preprint arXiv:2011.13677 (2020)
- Nguyen, D.T., Dax, M., Mummadi, C.K., Ngo, T.P.N., Nguyen, T.H.P., Lou, Z., Brox, T.: Deepusps: Deep robust unsupervised saliency prediction with selfsupervision. arXiv preprint arXiv:1909.13055 (2019)
- Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
- Qin, X., Zhang, Z., Huang, C., Gao, C., Dehghan, M., Jagersand, M.: Basnet: Boundary-aware salient object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7479–7489 (2019)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. Advances in neural information processing systems 28, 91–99 (2015)
- Roh, B., Shin, W., Kim, I., Kim, S.: Spatially consistent representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1144–1153 (2021)
- Sirinukunwattana, K., Pluim, J.P., Chen, H., Qi, X., Heng, P.A., Guo, Y.B., Wang, L.Y., Matuszewski, B.J., Bruni, E., Sanchez, U., et al.: Gland segmentation in colon histology images: The glas challenge contest. Medical image analysis 35, 489–502 (2017)
- Srinidhi, C.L., Ciga, O., Martel, A.L.: Deep neural network models for computational histopathology: A survey. Medical Image Analysis 67, 101813 (2021)
- Van Gansbeke, W., Vandenhende, S., Georgoulis, S., Van Gool, L.: Unsupervised semantic segmentation by contrasting object mask proposals. arXiv preprint arXiv:2102.06191 (2021)
- Vincent, P., Larochelle, H., Bengio, Y., Manzagol, P.A.: Extracting and composing robust features with denoising autoencoders. In: Proceedings of the 25th international conference on Machine learning. pp. 1096–1103 (2008)
- Wang, W., Zhou, T., Yu, F., Dai, J., Konukoglu, E., Van Gool, L.: Exploring crossimage pixel contrast for semantic segmentation. arXiv preprint arXiv:2101.11939 (2021)
- 33. Wang, X., Zhang, R., Shen, C., Kong, T., Li, L.: Dense contrastive learning for self-supervised visual pre-training. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3024–3033 (2021)
- 34. Wu, Z., Xiong, Y., Yu, S.X., Lin, D.: Unsupervised feature learning via nonparametric instance discrimination. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3733–3742 (2018)
- Xie, E., Ding, J., Wang, W., Zhan, X., Xu, H., Sun, P., Li, Z., Luo, P.: Detco: Unsupervised contrastive learning for object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8392–8401 (2021)

ConCL: Concept Contrastive Learning for Dense Pre-training in Pathology

- Xie, Z., Lin, Y., Zhang, Z., Cao, Y., Lin, S., Hu, H.: Propagate yourself: Exploring pixel-level consistency for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16684–16693 (2021)
- Yang, C., Wu, Z., Zhou, B., Lin, S.: Instance localization for self-supervised detection pretraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3987–3996 (2021)
- Zbontar, J., Jing, L., Misra, I., LeCun, Y., Deny, S.: Barlow twins: Self-supervised learning via redundancy reduction. arXiv preprint arXiv:2103.03230 (2021)
- Zhang, R., Isola, P., Efros, A.A.: Colorful image colorization. In: European conference on computer vision. pp. 649–666. Springer (2016)
- Zhang, R., Isola, P., Efros, A.A.: Split-brain autoencoders: Unsupervised learning by cross-channel prediction. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1058–1067 (2017)
- 41. Zhao, N., Wu, Z., Lau, R.W., Lin, S.: Distilling localization for self-supervised representation learning. arXiv preprint arXiv:2004.06638 (2020)
- Zhao, X., Vemulapalli, R., Mansfield, P.A., Gong, B., Green, B., Shapira, L., Wu, Y.: Contrastive learning for label efficient semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10623–10633 (2021)