

## Research

# From patches to objects: exploiting spatial reasoning for better visual representations

Toni Albert<sup>1</sup> · Bjoern Eskofier<sup>1,2</sup> · Dario Zanca<sup>1</sup>

Received: 1 March 2024 / Accepted: 15 April 2024

Published online: 29 April 2024

© The Author(s) 2024 [OPEN](#)

## Abstract

As the field of deep learning steadily transitions from the realm of academic research to practical application, the significance of self-supervised pretraining methods has become increasingly prominent. These methods, particularly in the image domain, offer a compelling strategy to effectively utilize the abundance of unlabeled image data, thereby enhancing downstream tasks' performance. In this paper, we propose Spatial Reasoning, a novel auxiliary pretraining method that takes advantage of a more flexible formulation of contrastive learning by introducing spatial reasoning as an auxiliary task for discriminative self-supervised methods. Spatial Reasoning works by having the network predict the relative distances between sampled non-overlapping patches. We argue that this forces the network to learn more detailed and intricate internal representations of the objects and the relationships between their constituting parts. Our experiments demonstrate substantial improvement in downstream performance in linear evaluation compared to similar work and provide directions for further research into spatial reasoning.

## Article highlights

- We propose a novel auxiliary pretraining method that takes advantage of a more flexible formulation of contrastive learning by introducing spatial reasoning as an auxiliary task for discriminative self-supervised methods.
- Spatial Reasoning enhances Relational Reasoning, improving visual representations with efficient training.
- Spatial reasoning reaches state-of-the-art performance among discriminative self-supervised learning methods.

**Keywords** Self-supervised learning · Contrastive learning · Relational reasoning · Spatial reasoning · Linear evaluation

## 1 Introduction

The rapid growth of deep learning models has led to state-of-the-art architectures containing millions of parameters [1]. Concurrently, the increasing availability of unlabeled data necessitates efficient methods for reducing manual annotation. Self-supervised pretraining enables models to understand relevant domains before finetuning on a smaller labeled dataset [2, 3].

---

✉ Dario Zanca, [dario.zanca@fau.de](mailto:dario.zanca@fau.de); Toni Albert, [toni.albert@fau.de](mailto:toni.albert@fau.de); Bjoern Eskofier, [bjoern.eskofier@fau.de](mailto:bjoern.eskofier@fau.de) | <sup>1</sup>Department Artificial Intelligence in Biomedical Engineering, FAU Erlangen-Nürnberg, Erlangen, Germany. <sup>2</sup>Institute of AI for Health, Helmholtz Zentrum München, Munich, Germany.



Two primary self-supervised learning approaches exist for images [4]: generative and discriminative. Generative methods reconstruct missing image components, yielding superior performance at the cost of larger networks and greater data demands. The complexity of the generative task requires the adoption of large models, such as ViT-B/16, comprising approximately 86 million parameters [5]. In contrast, discriminative self-supervised methods emphasize learning to differentiate between various features or patterns within the data without relying on explicit class labels. Discriminative approaches are successfully employed for pretraining of significantly smaller models like ResNet-32 [6], containing around 0.5 million parameters, thereby promoting suitability for smaller networks in self-supervised scenarios. While large literature exists for self-supervised learning of visual representations, in this work, we concentrate on discriminative methods for convolutional-based models, due to their benefits in network size and data efficiency.

Contrastive learning, a prevalent discriminative approach for self-supervised image learning, has been extensively developed [7–10]. It strives to generate meaningful representations by distinguishing augmented image versions from distinct images. Deep learning models must comprehend and identify image semantics to achieve this, thus creating meaningful representations. Patacchiola and Storkey [11] demonstrate that employing a classification head (relation module) on these representations reformulates the contrastive objective into a classification objective. This eliminates the need for typically large batch sizes and provides a strong supervisory signal, that is not prone to collapsing.

Leveraging the inherent capability of this formulation to flexibly accommodate supplementary objectives, wherein the classification head predicts the relative distance between two randomly selected patches from the same image. We argue that this objective forces the network to recognize the primary object and the spatial relationships among its constituting parts, enabling the creation of meaningful, scaled-domain representations. By concatenating patch-based and full-image representations, we have formulated an expanded representation that outperforms competitors in linear evaluation, even when training images and augmentations are limited. Although additional patch computations during the inference phase lead to increased computational requirements, we demonstrate that the number of patch representations at inference can be flexibly adjusted for more challenging tasks. Furthermore, we present an alternative formulation of spatial reasoning, called *additive-patch-use*, that avoids the cost of additional computing at inference time. Code is available at: <https://github.com/mad-lab-fau/Spatial-Reasoning>.

In summary, our contributions can be summarised as follows:

- We propose Spatial Reasoning and define a straightforward integration into Relational Reasoning architectures.
- We demonstrate better visual representations on several evaluation scenarios, with reduced training phase computation.
- We showcase performance for variable inference computing requirements.
- We present an alternative formulation of Spatial Reasoning, referred to as *additive-patch-use*, circumventing increased inference computation, with minimal loss in performance.
- We provide recommendations on optimal patch sizes and the number of patches to be used during training.

## 2 Related works

There have been significant advances in the creation of meaningful representations through self-supervised pretraining. Self-supervised pretraining methods in the image domain can be divided into discriminative and generative approaches [12]. With modern Vision Transformer (ViT) architectures, significant strides have been made in the generative direction. Atito et al. [13] show state-of-the-art results in linear evaluation among generative methods. This is obtained by training the ViT model to reconstruct transformed versions of the same images. Transformations range from partial grey scale to random replacement of image regions with different images. Modern discriminative methods like MoCo v3 combined with ViT architectures lead to state-of-the-art results as well [8]. Yet these architectures usually require high computational resources, especially during training. The ViT from Atito et al. [13], for example, contains 86 Million parameters in its base configuration, while in Chen et al. [8] variants with 300 million parameters and batch sizes up to 6000 are investigated. Additionally, training of ViT-based models for generative tasks exhibits instabilities [8]. There have been advances in training more data-efficient smaller transformers effectively, for example, reducing the size of ViT to 5 Million parameters in tiny variants, which allows training on 4 GPUs in less than three days on ImageNet [14]. Yet, pretraining transformer architectures seems not feasible for reaching

good downstream performance with smaller datasets and fewer computational resources [15]. In contrast, Convolutional Neural Networks (CNNs), such as the Residual Network (ResNet) architectures, demonstrated competitive performance, even with pretraining on small datasets [6]. Due to their design principles focusing on local receptive fields, shared weights, and spatial hierarchies, CNNs have a more restricted and efficient parameter space, which can be advantageous when training data is scarce [16]. Specifically, ResNet architectures, which introduce “skip connections” or “shortcuts” to allow the gradient to be directly backpropagated to earlier layers, have shown impressive performance in various computer vision tasks [6]. Therefore, convolutional architectures can serve as robust and effective alternatives to transformers in contexts where data or available compute is limited.

Contrastive learning approaches have been proposed for the self-supervised training of visual representations in ResNet architectures. An early approach was proposed by He et al. [7], addressing the problem of contrastive learning as dictionary look-up, and allowing efficient training and competitive results on the linear evaluation protocol on ImageNet. Xie et al. [17] exploited multi-level supervision to intermediate representations and contrastive learning between global image and local patches to improve performance on object detection tasks. Patacchiola and Storkey [11] showed that training a relation head to intra-reasoning and inter-reasoning results in rich visual representations and achieves state-of-the-art performance in pretraining of ResNet architectures via discriminative contrastive learning. The above mentioned approaches rely on a predefined set of image augmentations to define their contrastive learning objective. We adopt the same formulation as in Patacchiola and Storkey [11] and extend these approaches for spatial reasoning on non overlapping image patches.

Prior research has explored the efficacy of incorporating spatial information contained in patches for enhancing representations in self-supervised learning. Specifically, Noroozi and Favaro [18] use jigsaw puzzles and train neural networks to predict the correct combination to solve them. The proposed approach operates on nine patches using a Siamese-like architecture. Several other methods have been proposed for solving different variants of jigsaw puzzles. Kim et al. [19] addressed the problem of solving a damaged jigsaw puzzle by reconstructing a missing piece. Wei et al. [20] generalized the jigsaw problem and proposed an iterative solution. In contrast, our approach utilizes a network that has limited information about the image, with its relation module being fed with only two patches at a time. We argue that this limitation encourages the encoder to reason more deeply about the structure of the partially unseen object. Additionally, the network architecture in our approach is simpler compared to those in previous works and can be readily integrated into existing discriminatory approaches.

While we focus on image classification tasks, other work have been proposed that [21] focuses on image segmentation and introduces a pixel-wise contrastive algorithm, leveraging global context across images to enhance segmentation accuracy for multiple backbones. Similarly, Wang et al. [22] proposes dense contrastive learning at the pixel level, achieving superior performance in dense prediction tasks like object detection and semantic segmentation compared to existing methods. Lastly, Xie et al. [23] introduce pixel-level pretext tasks for unsupervised visual representation learning, surpassing instance-level methods and demonstrating effectiveness in pre-training both backbone and head networks for dense prediction tasks. Recent similar work approaches have been proposed for transformers-based architectures, that relate to the concept of spatial reasoning. Li et al. [24] focus on enhancing feature self-relation in self-supervised training of vision transformers, improving their representation ability for downstream tasks.

A closer approach to our proposal is presented in Doersch et al. [25]. Here the authors randomly sample a patch and then another at 8 adjacent locations. They then use a Siamese-like network to predict the relative positions as a probability over the eight possible patch locations. To avoid trivial solutions, the authors introduce a small location jitter as well as gaps between patches. In contrast, we simplify this approach by sampling at two completely random, non-overlapping positions, and formulate the problem as a contrastive objective. We additionally make use of the standard set of augmentations introduced by Patacchiola and Storkey [11], to combat other trivial solutions.

It is worth mentioning that all the approaches mentioned above generally require larger architectures with shared weights, making them challenging to be used in practice. We solve this problem by leveraging aggregation techniques [11].

### 3 Spatial reasoning

In object recognition tasks, it may, for example, be enough to identify the existence of an elephant skin pattern to discriminate between specific pictures in a dataset. However, this local reasoning might result in poor generalization performance [26]. Traditional contrastive learning approaches use augmentations like cropping to combat this issue.

With Spatial Reasoning, we aim to strengthen the supervisory signal and create more meaningful representations that not only contain the necessary information to identify an object but encode spatial relationships between its constituting parts, or between the main object and the background. As Spatial Reasoning builds on the idea that it is possible to create meaningful representation even from small patches of the image, we set up the objective of predicting the relative distance between  $N$  randomly sampled patches from the same image as an additional auxiliary prediction target. By using a classification head to reformulate typical contrastive losses as a classification problem, the introduction of Spatial Reasoning as auxiliary tasks is straightforward. An overview of the proposed method is given in Fig. 1, and a detailed explanation follows.

### 3.1 Patch and label generation

In standard Relational Reasoning [11], an image is augmented  $K$  times, and each of the augmented versions is then provided to the network, combined with augmented versions of other images from the same batch. The actual batch size thus corresponds to  $K$  times the number of images loaded per batch. We augment this process by creating  $N$  random positions per unaugmented image and extracting square patches at the chosen locations. The first two extracted patches are always guaranteed to be non-overlapping. This reduces the chance of finding a trivial solution. After extracting the patches, they are scaled to the standard input size and saved in the input batch with their corresponding locations as targets. To further reduce the chance of easy solutions, each patch is separately transformed by random colour jittering and grayscaling. We keep the augmentation scheme (colour jitter, grayscale, random resized crop, random horizontal flip) for the full-sized images, as introduced in Relational Reasoning. Saving patches not directly as pairs allow the usage of the algorithm implemented in [11] for live pair aggregation with minimal adjustments.

After generating patches and locations, they are appended to the output of the standard image augmentation step. We additionally test a different version of our approach, in which selected patches are added in their original resolution onto a black image. This can be obtained by padding the patch back to the original size with zeros and is regarded to as *additive-patch-use* in the following discussion. We discuss this approach in more detail in Sect. 4.4 of the Experiment.

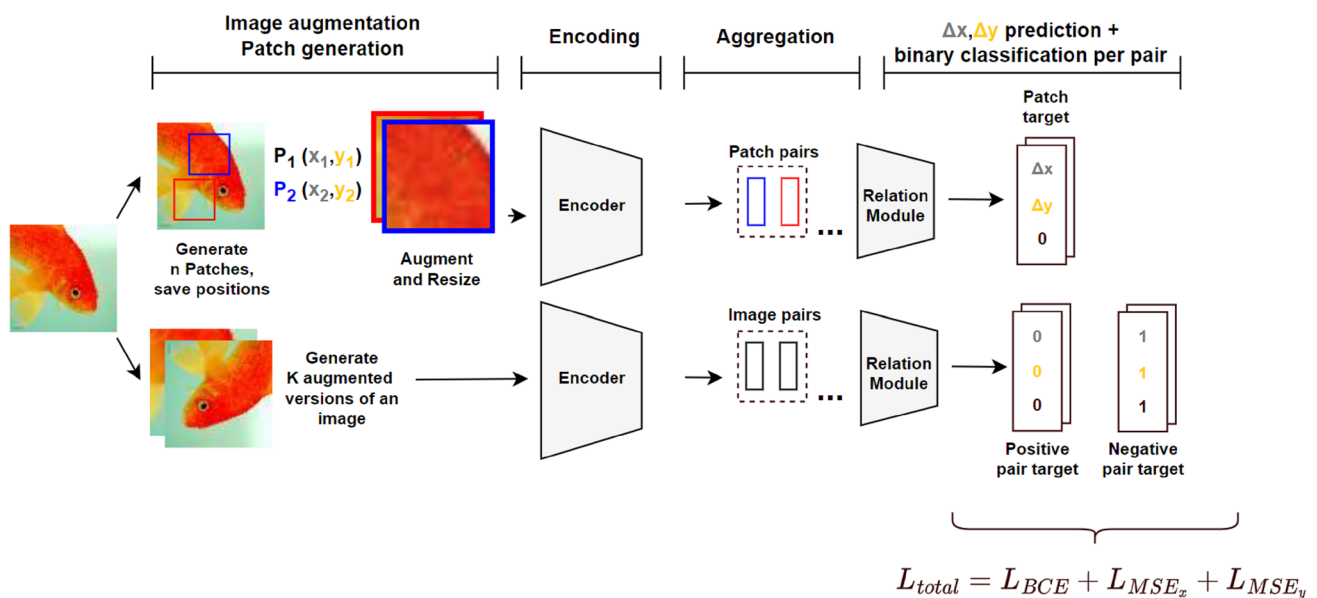


Fig. 1 Spatial reasoning method overview. The image augmentation process is expanded by the selection and rescaling of patches. After the prediction of the representations of both patches and images by the same encoder network, both types of representation are aggregated through specific methods into pairs. The patch aggregation is not using negative samples and defaults to a positive label for the contrastive objective. The image aggregation expands the target by pseudo labels for distance, depending on if the pair is positive or negative

### 3.2 Patch position prediction

Given a mini-batch of size  $M$ , our approach simultaneously processes augmented versions of the original images, and non-overlapping randomly selected patches, each of them randomly augmented once (see Fig. 1). Therefore, the resulting number of representations to be computed is given by

$$P = K \times M + M \times N. \quad (1)$$

where  $K$  is the number of augmentations, and  $N$  is the number of patches generated per image.

Afterward, an aggregation function is used, which concatenates two representations to be fed as input for the classification head. The number of generated pairs depends on the mini-batch size  $M$ , as well as the number of augmentations  $K$  and the number of patches  $N$ . The standard aggregation function creates all possible positive representation pairs by shifting the corresponding vectors. For the case of full images, for each positive pair, we also consider a negative pair generated by shifting to the next image in the mini-batch. The number of total combinations given by just the augmentations is given in Patacchiola and Storkey [11] as  $A = M(K^2 - K)$ . We extend this number to

$$A' = M(K^2 - K) + M \left( \frac{N^2 - N}{2} \right), \quad (2)$$

where  $N$  is the number of patches generated per image, and  $(N^2 - N)/2$  represents the number of unique patch pairs generated by iterating over the total number of patches. Note that, for the case of the patches, we consider positive pairs those constituted by patches belonging to the same image. Negative samples, i.e., pairs of patches belonging to different images, are not used for the patch pairs, therefore the denominator 2 in the second term of the above equation. Our tests show that negative patch pairs do not increase the downstream performance, as a negative patch pair would provide the same signal as a normal negative augmented representation pair from Relational Reasoning, just with less variance in augmentation. Omitting the generation of negative pairs results in a reduced scaling for the total number of aggregated patch representations while not reducing performance.

The relation module is adapted from [11] and expanded to 3 neurons in the final layer. The first neuron is used for the classification of negative and positive samples. The other two neurons predict the  $x$ - and  $y$ -coordinates of the relative distance of the patches or the dummy targets for standard images. In the case the source is an augmented image and not a patch, these units are used as pseudo targets for the classifier. The idea is to set the target to zero if the image pair is positive and to one if the pair contains different images. With this adjustment, the additional two target neurons propagate gradients even for the full-sized augmented image pairs. The patches are aggregated in a similar way, with only minor differences. First, the shifting aggregation is set to the number of generated patches, and the positional labels are subtracted from each other to produce the relative distance. For example, if patch  $p_1$  has a position of (0.2, 0.6) and patch  $p_2$  has a position of (0.4, 0.1), the resulting target would be  $p_1 - p_2 = (-0.2, 0.5)$ .

Let  $L_{BCE}$  denote the binary cross-entropy loss for the first neuron, which classifies negative and positive samples, and let  $L_{MSE_x}$  and  $L_{MSE_y}$  represent the mean-squared-error losses for the second and third neurons, which predict the  $x$ - and  $y$ -coordinates of the relative distance of the patches or the dummy targets. Thus, the total loss  $L_{total}$  can be defined as

$$L_{total} = L_{BCE} + L_{MSE_x} + L_{MSE_y}. \quad (3)$$

### 3.3 Dynamic compute requirements in the evaluation procedure

At inference time, we separate the image into  $n$  patches of the same size used during training and concatenate their representations with the representation from the whole image. It may be noted that most of our experiments use nine patches that cover the whole image in total and lead to a small overlap. The actual impact on downstream performance is investigated in Sect. 4.3. Using  $n$  patches leads to an  $n$  times higher computational cost at inference. On the other hand, our method is mostly trained with two generated patches per image and  $K = 4$  augmentations. While still outperforming the reported results with values for  $K$  ranging from 16 to 32 [11], which leads to a significantly lower computation demand in training.

In essence, for successful spatial reasoning, the neural network is required to generate representations from individual patches that empower the relation module to predict the relative distance between these patches accurately. Given the

fact that our relation module is only comprised of two layers, the encoder network has a crucial role to play: it must effectively discern which part of an object corresponds to the given patch. This integral information, once encoded, is expected to enhance the quality of the final representation when combined with representations from other patches. In other words, by forcing the network to understand the spatial relationships within an image, we encourage the development of more robust and meaningful representations that go beyond simple object identification to include spatial context.

## 4 Experiments

In the process of finding a simple formulation, we tested several different configurations of Spatial Reasoning and the corresponding pretraining architecture. In this section, we elaborate on the results of our experiments regarding the impact of patch size and number of patches on the performance in linear evaluation. Except for the concatenation of representations in the linear evaluation, we generally follow the procedure given in [11]. Specifically, we train the backbone model for 200 epochs using the unlabeled training dataset. Subsequently, a linear classifier is trained for 100 epochs, utilizing the features extracted from the backbone (without performing backpropagation on the backbone weights). The accuracy achieved by this classifier on the test dataset is regarded as the ultimate metric for evaluating the quality of the representations. We also do not change the augmentation strategy, except for leaving out cropping for patches.

In Table 1, we present a comparison of the performance of our method on various benchmarks. The table includes the mean accuracy (in percentage) and standard deviation over three runs using ResNet-32 for all datasets except STL-10, where ResNet-34 is used. The performance of our method is compared across different datasets, such as CIFAR-100, tiny-ImageNet, CIFAR-100-20 (coarse-grained), and STL-10. Additionally, we compare the performance on two cross-domain tasks: 10→100 (training on CIFAR-10 and testing on CIFAR-100) and 100→10 (training on CIFAR-100 and testing on CIFAR-10). The table highlights the effectiveness of our method in direct comparison to previous results in various settings. For tiny-imagenet, we reach a performance of 33.08% with four augmentations compared to the 30.5% reported in [11] with 16 augmentations on a ResNet-32. This is obtained using only two patches per image during training, and the optimal size of 24x24 pixels per patch, see Figs. 2 and 3. In the remaining chapter, we provide insights into how the choice of patch size and number of patches affects the performance in linear evaluation.

### 4.1 Patch size


As evidenced by previous studies [12, 30], the application of cropping and zooming as augmentation step significantly influences the performance of subsequent processes. This observation extends to our own experiments concerning the size of sampled patches, as illustrated in Fig. 2. The best performance for tiny-imagenet is reached for 23 and 24 pixels in width and length. Lower performance with smaller patches is most likely due to the missing amount of relevant information from the main object in the selected patch. This might lead the network to learn more abstract models based on the background and not on the object that is to be identified, which is an instance of the shortcut learning problem

**Table 1** Comparison on various benchmarks


Method	CIFAR-100	tiny-ImgNet	10→100	100→10	CIFAR-100-20	STL-10
Supervised (upper bound)	<b>65.32 ± 0.22</b>	<b>50.09 ± 0.32</b>	<b>33.98 ± 0.71</b>	<b>71.01 ± 0.44</b>	<b>76.35 ± 0.57</b>	<b>69.82 ± 3.36</b>
Random weights (lower bound)	7.65 ± 0.44	3.24 ± 0.43	7.65 ± 0.44	27.47 ± 0.83	16.56 ± 0.48	n/a
DeepCluster [27]	20.44 ± 0.80	11.64 ± 0.21	18.37 ± 0.41	43.39 ± ± 1.84	29.49 ± 1.36	73.37 ± 0.55
RotationNet [28]	29.02 ± 0.18	14.73 ± 0.48	27.02 ± 0.20	52.22 ± 0.70	40.45 ± 0.39	83.29 ± 0.44
Deep InfoMax [29]	24.07 ± 0.05	17.51 ± 0.15	23.73 ± 0.04	45.05 ± 0.24	33.92 ± 0.34	76.03 ± 0.37
SimCLR [12]	42.13 ± 0.35	25.79 ± 0.35	36.20 ± 0.16	65.59 ± 0.76	51.88 ± 0.48	89.31 ± 0.14
Relational reasoning [11]	46.17 ± 0.17	30.54 ± 0.42	41.50 ± 0.35	67.81 ± 0.42	52.44 ± 0.47	89.67 ± 0.33
<b>Spatial reasoning (plain)</b>	<b>50.18 ± 0.49</b>	<b>33.08 ± 0.07</b>	<b>47.93 ± 0.54</b>	<b>75.8 ± 0.76</b>	<b>58.51 ± 0.3</b>	<b>90.25 ± 0.55</b>
<b>Spatial reasoning (additive-patch-use)</b>	<b>48.04 ± 0.17</b>	<b>32.85 ± 0.31</b>	<b>42.29 ± 0.69</b>	<b>68.79 ± 0.49</b>	<b>54.78 ± 0.07</b>	<b>90.08 ± 0.12</b>

Mean accuracy (percentage) and standard deviation over three runs (ResNet-32/ResNet-34 for STL-10). *Spatial Reasoning (plain)* is described in Sect. 3 and requires performing parallel inference on 9 patches. *Spatial Reasoning (additive-patch-use)* is described in Sect. 4.4 and requires a single inference on the full image. Best results in bold, second best underlined

**Fig. 2** Impact of patch size on the linear evaluation performance on Tiny-ImageNet and Cifar10. Standard evaluation procedure with 100 epochs of finetuning on ten combined representations (1 image + 10 patches) and 2 patches extracted per image for each epoch



**Fig. 3** Impact of the number of patches extracted from each image on the linear evaluation performance on tiny-imagenet. Standard evaluation procedure with 100 epochs of finetuning on 10 combined representations (1 image + 9 patches) and 4 augmentations



[31]. The lower performance with larger patches is likely due to lower task difficulty. In fact, if the patches are too large, the number of possible non-overlapping locations decreases substantially. This leads to an easier task that reduces the strength of the supervisory signal produced by Spatial Reasoning.

## 4.2 Number of patches

This section discusses the impact of the number of patches extracted from an image in the training procedure on the linear evaluation performance. The results seen in Fig. 3 must be evaluated in combination with the number  $K$  of augmentations. As we use  $K=4$  in all of our experiments, a patch number of 2 means that every third input image is a patch. This also directly influences the number of possible combinations that are aggregated from the representations. As can be seen from Fig. 3, the optimal number of patches is 3, which means  $3/7$  of the input images are patches, and 192 of 960 total aggregated representation pairs are patch based in an optimal setting.

The lower performance with a rising number of patches can be explained by the higher impact of the *zoomed-in* domain. If too many images differ from the domain of the original dataset, the performance decreases due to the domain generalization problem [32, 33]. As roughly half of all created representation pairs are from patches with a patch count of 6, the performance decays to 33%. The initial improvement by 1% from two to four patches per base image proves the positive impact of spatial reasoning.

## 4.3 Impact of dynamic compute

Prior discussions have underscored the role played by the number of image patches extracted during the training process and its subsequent effect on linear evaluation performance. An integral part of this evaluation procedure involves the concatenation of image representations with their corresponding patch representations. While this process necessitates an increased computational load during the inference phase, it concurrently enables the incorporation of dynamic computing at inference.

Figure 4 demonstrates the impact of varying patch numbers at the inference stage on the accuracy yielded in linear evaluation. This experiment was conducted using two patches and a constant number of augmentations  $K = 4$  during training.

Here, we designate the central patch as the first one, while the patches horizontally adjacent are referred to as patches 2 and 3, respectively. Patches 5 and 6 correspond to the patches situated at the mid-bottom and mid-top positions. When employing five patches, we sample from positions 1 to 5, and in the case of seven patches, two additional random corner patches are sampled.

The initial increase in accuracy to 33.5% until the fifth patch's introduction can be attributed to the incremental availability of more information. However, the subsequent plateauing of performance upon increasing the patch number can be ascribed to the detrimental effects of overfitting. As the standard linear evaluation process excludes the use of any augmentations, the extensive size of the resultant representation over the 100 epochs in the evaluation procedure leads to a training accuracy of 45%, as opposed to a test set accuracy of 33%.

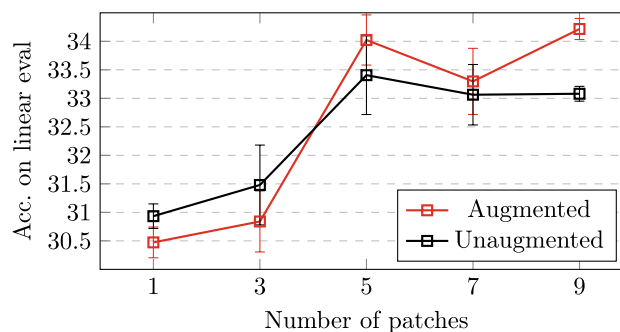
To substantiate this hypothesis, we also present results incorporating additional affine transformations in Fig. 4, configured analogously to those in STL-10, barring flipping and cropping, on the training set, indicated in red. The results show a distinct improvement of up to 1.5% with the use of nine patches. This underscores the potential for mitigating overfitting and fostering enhanced generalization by incorporating additional augmentations alongside larger representations.

#### 4.4 Additive-patch-use

The experiments and evaluations reveal certain limitations associated with spatial reasoning. Specifically, the training process requires a limit on the number of patches to mitigate the impact of domain shift as the number of patches increases. In addition, during evaluation and finetuning, there is a higher computational demand compared to similar methods because the patch representations have to be encoded.

To combat these issues, it is necessary to eliminate the impact of domain shift and make it possible for the network to predict a single representation right away. To reach this goal, we experimented with an adapted version of spatial reasoning. In this approach, instead of scaling the extracted patches, we pad the extracted patch to the size of the base image. We refer to this approach as *additive-patch-use*. The new images can be used for training instead of the resized patch. Additive-patch-use should reduce the domain shift, as the resolution of the patch itself is not changed and the network is able to extract the same features in the same resolution as in the normal base images. This makes the network training more robust to a higher number of patches. After training, we are now able to process an image with a single forward pass. This obviously limits the amount of information that can be extracted from each patch into 64 dimensions but should still take advantage of the deeper understanding of the object.


As shown in Fig. 5, additive-patch-use with 12 generated patches per image results in a 2.3% performance increase in linear evaluation compared to [11]. This is achieved with the same number of images as input in training, as  $K$  still equals 4. The positive impact of the method slightly diminishes with a higher amount of patches. This can be intuitively explained, as more patches lead to more overlap, which could enable trivial solutions. This, in turn, weakens the supervisory signal. In a broader context, the incorporation of additive-patch-use appears to be a straightforward augmentation technique for enhancing representations. Although it exhibits inferior performance compared to the Spatial Reasoning method described earlier, it mitigates the associated increase in inference cost. Additionally, the additive approach demonstrates robustness in relation to hyperparameter selection, particularly the number of patches. In the previously described method of Spatial Reasoning, selecting a significantly larger number of patch pairs compared to the number



**Fig. 4** Comparative influence of dynamic compute with and without augmentations. This figure illustrates the impact of varying patch numbers on the accuracy of linear evaluation, represented in black for unaugmented overfitting and in red for augmented configurations. The error bars represent the range of accuracy values obtained from three independent experiments. It is discernible that the incorporation of additional augmentations in the red graph aids in mitigating overfitting and enhancing accuracy, particularly with larger representations



**Fig. 5** Linear evaluation of additive patch creation. Evaluated with a single representation on tiny-imageNet with  $K = 4$ , 3 runs with different seeds for each



of traditional representation pairs can lead to a decrease in overall performance. In contrast, the additive-patch-use mitigates this issue by being less sensitive to the choice of this hyperparameter, allowing for a more reliable integration without negatively affecting performance.

## 5 Discussion and conclusion

Our work shows that a relation head can be used to design the learning of Spatial Reasoning, formulating it as an auxiliary pretraining objective. Our approach creates better visual representations, even while reducing the necessary compute in training. Based on the results of the linear evaluation, Spatial Reasoning improves the resulting representations significantly in all evaluation scenarios. Further research in this direction could improve the performance of discriminative self-supervised pretraining and reduce the performance gap to generative methods with significantly smaller networks.

Despite reduced computational requirements during training, compared to similar work, our approach requires higher compute cost at inference time, as well as the careful setting of the number of patches as a hyperparameter. We presented an alternative formulation, named additive-patch-use, that demonstrates to be a good solution for these limitations. In fact, it reduces the impact of domain shift as the number of patches increases, and it eliminates the additional computational demand in encoding all patches during inference.

Future work could explore the definition of sampling techniques for the patch size, instead of choosing a constant value. Furthermore, as our approach can easily be adapted to other frameworks, a promising direction is to incorporate Spatial Reasoning into standard contrastive frameworks, like SimCLR, MoCo and BYOL.

## 6 Additional information

In this section, we present further information on our experimental setup, as well as the used datasets.

### 6.1 Experimental setup

All models are trained and evaluated on one of two different nodes. The first node contains a single RTX3080 GPU and a modern 8-core CPU. This node is used for all experiments with a patch count smaller than 3 and every dataset except STL-10. The second node contains an A100 GPU and a similar processor. All models fit on their respective GPUs. The training duration and learning rate as well as the augmentation scheme for all datasets are equivalent to the ones proposed by Patacchiola and Storkey[11]. The only differences in augmentations are the missing random resized crop and random horizontal flip for each patch. The only time the general procedure of linear evaluation is changed by adding affine augmentations refers to the experiments done for Fig. 4 and is denoted in the corresponding chapter. If error bars are present in a figure, each value is evaluated on three different seeds, with the upper limit of the bar denoting the maximum value and the lower limit the minimum value. The marker marks the average of all three runs.

## 6.2 Datasets

To evaluate our method, we used the following datasets: CIFAR-100, CIFAR-100-20, CIFAR-100→10, CIFAR-10→100, tiny-ImageNet, and STL-10. Below, we provide a brief description of each dataset and their characteristics.

- **CIFAR-100** The CIFAR-100 dataset [30] consists of 60,000 32x32 color images, divided into 100 classes with 600 images per class. The dataset is split into a training set of 50,000 images and a test set of 10,000 images. The images in CIFAR-100 are low-resolution, making it a challenging dataset for object recognition tasks.
- **CIFAR-10** The CIFAR-10 dataset [30] has the same structure of 50,000 training images and 10,000 testing images, but only contains 10 classes.
- **CIFAR-100-20** The CIFAR-100-20 is CIFAR 100 with course-grained classes, containing only 20 superclasses from the original 100. This set classes is used to evaluate the course-grained performance of the model on a smaller and less diverse dataset.
- **CIFAR-100→10** In this cross-domain task, the model is trained on the CIFAR-100 dataset and tested on the CIFAR-10 dataset. This task evaluates the model's ability to transfer knowledge from a more complex and diverse dataset to a simpler one.
- **CIFAR-10→100** This cross-domain task is the opposite of the previous one. The model is trained on the CIFAR-10 dataset and tested on the CIFAR-100 dataset. This task assesses the model's capacity to generalize and adapt its learned features from a simpler dataset to a more complex and diverse one.
- **Tiny-ImageNet** The tiny-ImageNet dataset is a downscaled version of the ImageNet dataset [34], consisting of 200 classes, each with 500 training images, 50 validation images, and 50 test images. The images are 64x64 pixels in size. Due to the small size and diversity of the images, tiny-ImageNet presents a challenging benchmark for image recognition algorithms.
- **STL-10** The STL-10 dataset [35] is an image recognition dataset containing 10 classes with 1,300 96x96 color images per class. The dataset is designed for unsupervised and supervised learning, with a training set of 5,000 labeled images, 100,000 unlabeled images, and 8,000 test images. The higher resolution and the presence of both labeled and unlabeled data make STL-10 a suitable dataset for evaluating the performance of self-supervised learning methods.

These datasets provide a diverse set of challenges for our method, allowing us to evaluate its performance and effectiveness across various contexts and tasks.

**Funding** Open Access funding enabled and organized by Projekt DEAL.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

## References

1. Han K, Wang Y, Chen H, Chen X, Guo J, Liu Z, Tang Y, Xiao XuC, Xu Y, Yang Z, Zhang Y, Tao D. A survey on vision transformer. *IEEE Trans Pattern Anal Mach Intell.* 2022. <https://doi.org/10.1109/TPAMI.2022.3152247>.
2. Erhan D, Bengio Y, Courville A, Manzagol P-A, Vincent P, Bengio S. Why does unsupervised pre-training help deep learning? *J Mach Learn Res.* 2010;11(19):625–60.
3. Huh M, Agrawal P, Efros AA. What makes ImageNet good for transfer learning? 2016. arXiv preprint [arXiv:1608.08614](https://arxiv.org/abs/1608.08614)
4. Liu X, Zhang F, Hou Z, Mian L, Wang Z, Zhang J, Tang J. Self-supervised learning: generative or contrastive. *IEEE Trans Knowl Data Eng.* 2021;35(1):857–76.
5. Dosovitskiy A, Beyer L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, et al. An image is worth 16x16 words: transformers for image recognition at scale. 2020. arXiv preprint [arXiv:2010.11929](https://arxiv.org/abs/2010.11929)
6. 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*;2016. p. 770–8.

7. 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; 2020. p. 9729–38.
8. Chen X, Xie S, He K. An empirical study of training self-supervised vision transformers; 2021. <https://arxiv.org/abs/2104.02057>.
9. Chen T, Kornblith S, Swersky K, Norouzi M, Hinton G. Big self-supervised models are strong semi-supervised learners; 2020. <https://arxiv.org/pdf/2006.10029>.
10. Grill J-B, Strub F, Altché F, Tallec C, Richemond P, Buchatskaya E, Doersch C, Avila Pires B, Guo Z, Gheshlaghi Azar M, et al. Bootstrap your own latent—a new approach to self-supervised learning. *Adv Neural Inf Process Syst*. 2020;33:21271–84.
11. Patacchiola M, Storkey AJ. Self-supervised relational reasoning for representation learning. *Adv Neural Inf Process Syst*. 2020;33:4003–14.
12. Chen T, Kornblith S, Norouzi M, Hinton G. A simple framework for contrastive learning of visual representations. In: International conference on machine learning; 2020. p. 1597–607. PMLR.
13. Atito S, Awais M, Kittler J. SiT: self-supervised vision transformer; 2021. <https://arxiv.org/pdf/2104.03602>.
14. Touvron H, Cord M, Douze M, Massa F, Sablayrolles A, Jégou H. Training data-efficient image transformers & distillation through attention; 2021. p. 10347–357. PMLR.
15. Zhai X, Kolesnikov A, Houlsby N, Beyer L. Scaling vision transformers. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR); 2022. p. 12104–13.
16. LeCun Y, Bengio Y, Hinton G. Deep learning. *Nature*. 2015;521(7553):436–44. <https://doi.org/10.1038/nature14539>.
17. 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 (ICCV); 2021. p. 8392–01.
18. Norouzi M, Favaro P. Unsupervised learning of visual representations by solving jigsaw puzzles. In: European conference on computer vision; 2016. p. 69–84. Springer.
19. Kim D, Cho D, Yoo D, Kweon IS. Learning image representations by completing damaged jigsaw puzzles. In: 2018 IEEE winter conference on applications of computer vision (WACV); 2018. p. 793–802. IEEE.
20. Wei C, Xie L, Ren X, Xia Y, Su C, Liu J, Tian Q, Yuille AL. Iterative reorganization with weak spatial constraints: solving arbitrary jigsaw puzzles for unsupervised representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition; 2019. p. 1910–9.
21. Zhou T, Wang W. Cross-image pixel contrasting for semantic segmentation. *IEEE Trans Pattern Anal Mach Intell*. 2024.
22. 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; 2021. p. 3024–33.
23. 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; 2021. p. 16684–93.
24. Li Z-Y, Gao S, Cheng M-M. Sere: exploring feature self-relation for self-supervised transformer. *IEEE Trans Pattern Anal Mach Intell*. 2023.
25. Doersch C, Gupta A, Efros AA. Unsupervised visual representation learning by context prediction. In: Proceedings of the IEEE international conference on computer vision; 2015. p. 1422–30.
26. Geirhos R, Rubisch P, Michaelis C, Bethge M, Wichmann FA, Brendel W. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness; 2018. arXiv preprint [arXiv:1811.12231](https://arxiv.org/abs/1811.12231)
27. Caron M, Bojanowski P, Joulin A, Douze M. Deep clustering for unsupervised learning of visual features. In: Proceedings of the European Conference on Computer Vision (ECCV), pp. 132–149 (2018)
28. Gidaris S, Singh P, Komodakis N. Unsupervised representation learning by predicting image rotations. arXiv preprint [arXiv:1803.07728](https://arxiv.org/abs/1803.07728) (2018)
29. Hjelm R.D, Fedorov A, Lavoie-Marchildon S, Grewal K, Bachman P, Trischler A, Bengio Y. Learning deep representations by mutual information estimation and maximization. arXiv preprint [arXiv:1808.06670](https://arxiv.org/abs/1808.06670) (2018)
30. Krizhevsky A, Hinton G, et al. Learning multiple layers of features from tiny images. University of Toronto (2009)
31. Geirhos R, Jacobsen J.-H, Michaelis C, Zemel R, Brendel W, Bethge M, Wichmann F.A. Shortcut learning in deep neural networks. *Nature Machine Intelligence* 2(11), 665–673 (2020)
32. Torralba A, Efros A.A. Unbiased look at dataset bias. In: CVPR 2011, pp. 1521–1528 (2011). <https://doi.org/10.1109/CVPR.2011.5995347>
33. Ganin Y, Lempitsky V. Unsupervised domain adaptation by backpropagation. In: International Conference on Machine Learning, pp. 1180–1189 (2015). PMLR
34. Le Y, Yang X. Tiny imagenet visual recognition challenge. *CS 231N 7(7)*, 3 (2015)
35. Coates A, Ng A, Lee H. An analysis of single-layer networks in unsupervised feature learning. In: Gordon, G., Dunson, D., Dudík, M. (eds.) Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics. Proceedings of Machine Learning Research, vol. 15, pp. 215–223. PMLR, Fort Lauderdale, FL, USA (2011). <https://proceedings.mlr.press/v15/coates11a.html>

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.