



# Optimizing Transformations for Contrastive Learning in a Differentiable Framework

Camille Ruppli<sup>1,3</sup>(✉), Pietro Gori<sup>1</sup>, Roberto Ardon<sup>3</sup>, and Isabelle Bloch<sup>2,3</sup>

<sup>1</sup> LTCI, Télécom Paris, Institut Polytechnique de Paris, Paris, France

<sup>2</sup> Sorbonne Université, CNRS, LIP6, Paris, France

<sup>3</sup> Incepto Medical, Paris, France

camille.ruppli@incepto-medical.com

**Abstract.** Current contrastive learning methods use random transformations sampled from a large list of transformations, with fixed hyper-parameters, to learn invariance from an unannotated database. Following previous works that introduce a small amount of supervision, we propose a framework to find optimal transformations for contrastive learning using a *differentiable* transformation network. Our method increases performances at low annotated data regime both in supervision accuracy and in convergence speed. In contrast to previous work, no generative model is needed for transformation optimization. Transformed images keep relevant information to solve the supervised task, here classification. Experiments were performed on 34000 2D slices of brain Magnetic Resonance Images and 11200 chest X-ray images. On both datasets, with 10% of labeled data, our model achieves better performances than a fully supervised model with 100% labels.

**Keywords:** Contrastive learning · Semi-supervised learning · Transformations optimization

## 1 Introduction

When working with medical images, data are increasingly available but annotations are fewer and costly to obtain. Self-supervised methods have been developed to take full advantage of the non-annotated data and increase performances in supervised tasks at low annotated data regime. As part of self-supervised methods, contrastive learning methods [1, 2, 11, 12] train an encoder on non-annotated data to learn invariance between transformed versions of images. Contrastive learning methods are also used with medical images. For instance, the authors of [1] learn local and global features invariance while those of [5] introduce a kernel to take metadata into account in contrastive pretraining.

In most works, the transformations used to learn invariance are randomly sampled from a given list. While many works study the impact of removing some transformations on supervised task performance [2, 12], not much investigation has been done on optimizing the transformations and their hyper-parameters. Some authors

[11, 15] focus on the role of transformations but without explicit transformations optimization. The work of [11] proposes a formal analysis of transformations composition to select admissible transformations while [15] explores the latent spaces of specific transformations. The authors of [16] introduce a generative network to learn transformations distribution present in the data to use complementary transformations in self-supervised tasks. Unlike our work (see Sect. 2) they need a pre-training step before the contrastive one to learn transformations distribution.

Within supervised training (not self-supervision), some works have proposed to optimize data augmentation. In [4], a pre-training step using reinforcement learning is required. The work of [17] shows that data augmentation should be applied on both discriminator and generator optimization steps but no optimization is performed on augmentation choice. The authors of [8, 9] learn a vector containing augmentations probability. They also present a transformations optimization strategy. Unlike our approach (see Sect. 2), transformation parameters are discretized. Optimization is performed on the probability of choosing a family of transformations and a set of parameters.

While supervision is also introduced in contrastive learning in [6, 18], few authors used it in order to influence the choice of transformations. Among them, the authors of [14] introduce a transformation generator (a flow-based model based on [7]) to generate transformed images in new color spaces minimizing mutual information while keeping enough information for the supervised task. As transformations only impact color spaces, their application to gray scale images, in particular medical images, is very limited. Furthermore, consistently synthesizing anatomically relevant images with generative models can be difficult [3]. To the best of our knowledge, the approach in [14] is the only existing method optimizing a transformation generator for contrastive learning.

As in [14], the present work uses a small amount of supervision (10%) for transformation optimization. We introduce a differentiable framework on transformations that needs no pre-training, and, unlike [14], is applicable to both color and gray scale images. Our contributions are the following:

- We propose a semi-supervised differentiable framework to optimize the transformations of contrastive learning.
- We demonstrate that our method finds relevant transformations for the downstream task, which are easy to interpret.
- We show that our framework has better performances than fully supervised training at low data regime and contrastive learning [2] without supervision.

## 2 Transformation Network

Contrastive learning methods train an encoder to bring close together latent representations of positive pairs of images while pushing away representations of negative pairs of images. As in simCLR [2], positive pairs are two transformations of the same image while negative pairs are transformed versions of different images.

Transformations used in most methods are chosen at random from a fixed given list. However, as shown in [14], using positive (transformed) images, that are very similar to each other (i.e., high mutual information), might entail a sub-optimal solution since it would not bring additional information to the encoder. By using a small amount of supervision, transformations can be optimized in order to contain relevant information for the targeted supervised task.

In this work, we focus on classification tasks. We introduce a transformation network ( $M$ ) that minimizes the mutual information between images of a positive pair without compromising the supervised task performance. For each image of the training set,  $M$ , implemented as a neural network, outputs a set of parameters ( $\Lambda$ ) defining the transformations to apply ( $T_{\Lambda_M}$ ). As in [2, 15], the latent space of the encoder ( $f$ ) is optimized using a projection head ( $g$ ) into a lower dimension space where a contrastive loss function ( $I_{NCE}$ ) is minimized. Supervision is added on the latent space using a linear classifier ( $p$ ) that minimizes a classification loss function ( $\mathcal{L}$ ). Figure 1 shows a schematic view of the architecture used ( $X$  denotes an image from the training set and  $X_M$  its transformed version).



**Fig. 1.** Proposed architecture (red color indicates a trainable element, blue color indicates a non-trainable element). (Color figure online)

## 2.1 Optimizing Transformations

We consider a finite set of intensity and geometric transformations acting on images. Each transformation is parameterized by a vector of parameters (for example, the parameter vector of a rotation around a fixed point only contains its angle). The transformation function ( $T_{\Lambda_M}$ ) is the composition of transformations applied in a fixed order. The transformation network ( $M$ ) outputs the transformation function parameters. We propose to train  $M$  to find the optimal transformations for the semi-supervised contrastive problem. The network  $M$  maps an image to the space of parameter vectors, normalized to  $[0, 1]$ . The order of the transformations in the composition is not optimized, but the impact of this order has been studied and results are shown in Sect. 3.

Let  $\lambda_k$  be the vector of parameters for a given transformation, then the transformation function, noted as  $T_{\Lambda_M}$ , is parameterized by  $\Lambda = [\lambda_1, \dots, \lambda_K]$  (where  $K$  is the number of transformations considered).

The optimal transformations for the semi-supervised contrastive problem is then obtained via  $M$ , which is thus responsible for finding the optimal  $\Lambda_M^*$ . In contrast with [2], we only transform one version of the image batch. Our experiments show better results in this setting. The optimization goes as follows.

**Transformation Network Optimization Steps:** (i)  $M$  generates a batch of  $\Lambda_M$  vectors defining a transformation  $T_{\Lambda_M}$ . For every image  $X$  in a batch, a transformed version is generated:  $X_M = T_{\Lambda_M}(X)$ . (ii) The transformed and untransformed data batches are passed through the encoder  $f$ , the projection head  $g$  and the linear classifier  $p$ . (iii) The contrastive loss  $-I_{NCE}$  (see below, Eq. 2) gradient is computed to update the weights of the network  $M$  aiming to minimize mutual information and classification loss function.

**Encoder Optimization Steps:** (i) From the previous optimization steps of  $M$ , one transformed version of the data is generated. Latent projections of the transformed and untransformed data are generated using encoder  $f$  and projection head  $g$ . (ii) The contrastive loss gradient is computed and parameters of  $f$ ,  $g$  and  $p$  are updated. This brings closer positive pairs and further away negative ones, and ensures that transformed images are properly classified.

Formally, these steps aim to solve the following coupled optimization problem where contrastive and classification loss functions are taken into account:

$$\begin{cases} \min_M \alpha_0 I_{NCE}(g \circ f(X_M), g \circ f(X)) + \alpha_1 \mathcal{L}(p \circ f(X_M), y) \\ \min_{f,p,g} -\alpha_2 I_{NCE}(g \circ f(X_M), g \circ f(X)) + \alpha_3 \mathcal{L}(p \circ f(X_M), y) \\ \quad + \alpha_4 \mathcal{L}(p \circ f(X), y) \end{cases} \quad (1)$$

where  $\alpha_i$  are weights balancing each loss term and  $y$  are the classification labels when available. The term  $I_{NCE}$  is the contrastive loss function as in [2]:

$$I_{NCE}(X_{M_i}, X_i) = - \sum_i \log \left( \frac{e^{\text{sim}(g(f(X_{M_i})), g(f(X_i)))}}{\sum_{j, j \neq i} e^{\text{sim}(g(f(X_{M_i})), g(f(X_j)))}} \right) \quad (2)$$

where the index  $i$  defines positive pairs,  $j$  negative ones, and  $\text{sim}$  is a similarity measure defined as  $\text{sim}(x, x') = \frac{x^T x'}{\tau}$  where  $\tau$  is a fixed scalar, here equal to 1. Finally,  $\mathcal{L}$  is the binary cross entropy loss function for the supervised constraint.

## 2.2 Differentiable Formulation of the Transformations

A fundamental difference of the proposed transformation optimization, compared to [8, 9, 14], is the use of explicit transformations differentiation. During training, gradient computations of Eq. 1 involve the derivative of  $T_{\Lambda_M}$  with respect to the weights ( $w$ ) of  $M$ :  $d_w(T_{\Lambda_M}) = dT_{\Lambda_M} \circ d_w M$ . This requires the explicit computation of the derivatives of  $T$  with respect to its parameters  $\Lambda$  and the differential calculus for each transformation composing  $T$ . Thus, we introduce specific formulations and normalized parameterization for the transformations used in our experiments.

We use the following transformations: crop ( $Crop$ ), Gaussian blur ( $G$ ), additive Gaussian noise ( $N$ ), rotation ( $R$ ) around the center of the image, horizontal

( $Flip_0$ ) and vertical ( $Flip_1$ ) flips. Table 1 lists the expressions of these transformations. The final transformation function is defined as:

$$T_A = (R \circ Flip_1 \circ Flip_0 \circ Crop \circ N \circ G)(X, A) \quad (3)$$

and  $T_A$  thus depends on 7 parameters (the crop has 2 parameters) which are generated by  $M$ .

**Table 1.** Differentiable expressions of the transformations used, parameterized by  $\lambda \in [0, 1]$ , where  $S$  is the sigmoid function,  $s$  the size of our images,  $\text{erfinv}$  the inverse of the error function  $2\pi^{-\frac{1}{2}} \int_x^\infty e^{-u^2} du$ ,  $\mathcal{U}$  the uniform distribution and  $x$  is a point of the image grid. We fix the maximum Gaussian blur standard deviation to  $\sigma_{max} = 2.0$  and the maximum additive noise standard deviation to  $\bar{\sigma}_{max} = 0.1$ .

Flip around axis $e$	$Flip(X, \lambda, e)(x) = (1 - \lambda)X(x) + \lambda X(x - 2\langle x, e \rangle e)$
Crop centered at $c_\lambda = [\lambda_1 s, \lambda_2 s]$	$Crop(X, \lambda)(x) = X(x) \times S(\frac{s}{8} - \ x - c_\lambda\ _\infty)$
Gaussian blur with kernel $g_{\lambda\sigma_{max}}$	$G(X, \lambda) = g_{\lambda\sigma_{max}} * X$
Rotation	$R(X, \lambda)(x) = X \left( \begin{pmatrix} \cos(\lambda 2\pi) & -\sin(\lambda 2\pi) \\ \sin(\lambda 2\pi) & \cos(\lambda 2\pi) \end{pmatrix} x \right)$
Additive Gaussian noise	$N(X, \lambda) = X + \lambda \bar{\sigma}_{max} \times \sqrt{2} \text{erfinv}(\mathcal{U}[-1, 1])$

## 2.3 Experimental Settings

**Dataset.** Experiments were performed on BraTs MRI [10] and Chest X-ray [13] datasets. The Chest X-ray dataset is composed of 10000 images. BraTs volumes were split along the axial axis to get 2D slices. Only slices with less than 80% of black pixels were kept. This resulted in 34000 slices. For both datasets, we studied the supervised task of pathology presence classification (binary classification, present/not present). In medical imaging problems, it is common to have labels only for a small part of the dataset. We thus choose 10% of supervision in all of our experiments. We randomly selected three hold-out test sets of 1000 slices for BraTs experiments. With the Chest dataset, we used the provided test set of 1300 images, from [13], evenly split in three to evaluate variability.

**Implementation Details.** For every experiment with the BraTs dataset, the encoder  $f$  is a fully convolutional network composed of four convolution blocks with two convolutional layers in each block. Following the architecture proposed in [13], the encoder  $f$  for experiments on the Chest dataset is a Densenet121. The network  $M$  is a fully convolutional network composed of two convolutional blocks with one convolutional layer. The projection head  $g$  is a two-layer perceptron as in [2]. On BraTs dataset (resp. Chest dataset), we train with a batch size of 32 (resp. 16) for 100 epochs. In each experiment, the learning rate of  $f$  is set to  $10^{-4}$ . When optimizing  $M$  with (resp. without) supervision,  $M$  learning rate is set to  $10^{-3}$  (resp.  $10^{-4}$ ). When using 10% of labeled data for the supervision task, on relatively small databases ( $10^5$  images), there is a risk of overfitting on

the classification layer ( $p$  in Eq. 1). Contrastive and supervision loss terms need to be carefully balanced while optimizing both the encoder and the transformation generator. To evaluate the impact of hyper-parameters, we carried out experiments with  $(\alpha_0, \alpha_2) \in \{1, 0.1\}$  and  $(\alpha_1, \alpha_3, \alpha_4) \in \{1, 10\}$ . Linear evaluation results (see Sect. 2.4) on BraTs dataset after convergence are summarized in Table 2. Results in Sect. 3 are shown with the best values found for each method.

**Table 2.** 3-fold cross validation mean linear evaluation AUC after convergence with different  $\alpha_i$  values (standard deviation in parentheses).

	$\alpha_i$ values	AUC
Optimizing M	$\alpha_{0,2} = 1, \alpha_{3,4} = 1, \alpha_1 = 10$	<b>0.884</b> (0.042)
	$\alpha_0 = 0.1, \alpha_{1,3,4} = 10, \alpha_2 = 0.1$	0.868 (0.030)
	$\alpha_0 = 0.1, \alpha_1 = 10, \alpha_2 = 1, \alpha_{3,4} = 1$	<b>0.887</b> (0.013)
Random M	$\alpha_2 = 1, \alpha_{3,4} = 1$	0.874 (0.000)
	$\alpha_2 = 0.1, \alpha_{3,4} = 10$	0.820 (0.037)
	$\alpha_2 = 1, \alpha_{3,4} = 10$	<b>0.883</b> (0.003)
Base simCLR [2]		0.730 (0.020)

The fully supervised experiments described in Sect. 3 are optimized with the same encoder architecture and one dense layer followed by a sigmoid activation function for the classification task. For the fully supervised experiments we used a learning rate of  $10^{-4}$ .

**Computing Infrastructure.** Optimizations were run on Tesla NVIDIA V100 cards.

## 2.4 Linear Evaluation

To evaluate the representation quality learned by the encoder, we follow the linear evaluation protocol used in the literature [2, 12, 14]. The encoder is frozen with the weights learned with our framework. One linear layer is added, after removing the projection head ( $g$ ), and trained with a test set of labeled data, not used in the previous training phase. This means that we first project the test samples in the latent space of the frozen model and then estimate the most discriminative linear model. The rationale here is that a good representation should make the classes of the test data linearly separable.

## 3 Results and Discussion

To assess the impact of each term in Eq. 1 we performed optimization using the following strategies:

**Random** (without  $M$ , without supervision): each image is transformed with parameters generated by a uniform distribution:  $A = \mathcal{U}([0, 1]^7)$ , and  $\alpha_{1,3,4} = 0$ .

**Random with supervision** (without  $M$ , with supervision): we add the supervision constraint to the random strategy. We set  $\alpha_2 = 1$  and  $\alpha_{3,4} = 10$ .

**Self-supervised** (with  $M$ , without supervision): while setting  $\alpha_1, \alpha_3$  and  $\alpha_4$  to 0, we optimize Eq. 1.

**Self-supervised with supervision constraint** (with  $M$  and supervision): setting  $\alpha_1 = 10$  and  $\alpha_{0,2,3,4} = 1$ , we optimize Eq. 1.


We split the data into pre-training and test sets. Data from the pre-training set are further split into training and validation sets for the perturbator/encoder optimization. For optimizations with supervision constraint (self-supervised and random), all pre-training data are used for self-supervision and a small set of labeled data is used for the supervision constraint. For variability analysis, three optimizations were performed by changing the supervision set. With the BraTs dataset, as slices come from 3D volumes, we split the data ensuring that all slices of the same patient were in the same set.

**Linear evaluation** was performed on the four optimization strategies with the hold-out test set. Performances were evaluated with the weights obtained at different epochs. We aim to evaluate if our method outputs better representations during training. In Fig. 2, we show performances (mean and standard deviation) on three different test sets for both datasets. We also trained the encoder on the classification task in a fully supervised setting with 10% and 100% labeled data. For the fully supervised training, we used data augmentation composing the tested transformations randomly. Each transformation had a 0.5 probability of being sampled. We performed linear evaluation on the frozen encoder with the hold-out test set and report the obtained AUC as horizontal lines in Fig. 2. Figure 2 also reports linear evaluation results of the base simCLR optimization as in [2] where only one image is transformed by a random composition of the tested transformations. As with the fully supervised experiments, each transformation had a 0.5 probability of being sampled.

Figure 2 shows that optimizing  $M$  with supervision helps to have better representations for both datasets. It also shows that optimizing with only 10% of labeled data allows us to reach the same quality of representation as the fully supervised training with 100% of labels.


To investigate the impact of the supervised loss function, we launched an experiment with the supervised contrastive loss introduced in [6] using only 10% of labeled data. After convergence, we obtained a mean AUC of  $0.52 \pm 0.12$  compared to  $0.93 \pm 0.01$  with our method.

On the Chest X-ray database, strong results were obtained in [13] using a network pretrained on ImageNet. Optimizing  $M$  with 10% supervision on this ImageNet pretrained network has a smaller impact compared to random transformations ( $0.96 \pm 0.001$  for both approaches). However, ImageNet pretrained networks can only be used with 2D slices whereas our strategy could be easily extended to 3D volumes.



**Fig. 2.** Linear evaluation results comparing with other methods (left BraTs dataset with batch size 32, right Chest dataset with batch size 16).

**Relevance.** When optimizing without supervision, the network  $M$  needs to minimize the mutual information and it can therefore generate transformations that create images that are very different from the untransformed images but that do not contain relevant information for the downstream task, in particular for medical images. Without the supervision constraint, the optimal crop can be found, for instance, in a corner, leading to an image with a majority of zero values (i.e., entirely black), thus useless for the supervised task. The supervision constraint helps  $M$  to generate relevant images that keep pathological pixels (see some examples in Fig. 3).



**Fig. 3.** Two examples (row 1 and 2) of generated transformations in the BraTs dataset with different optimization strategies (red contour highlights the tumor). (Color figure online)

**Runtime.** The addition of the network  $M$  increases the training computational time of around 20–25% which is balanced by a performance gain.



**Transformation Composition Order.** As in [2], the transformation order is fixed. We launched an additional experiment with a different transformation order for both simCLR and our method. Linear evaluation results after convergence are respectively:  $0.730 \pm 0.020$  and  $0.760 \pm 0.027$  for simCLR and  $0.926 \pm 0.020$  and  $0.923 \pm 0.021$  for our method. The transformation order has thus little impact on our results and, above all, our method substantially outperforms simCLR in both experiments.

## 4 Conclusions and Perspectives

We proposed a method to optimize usual transformations employed in contrastive learning with very little supervision. Extensive experiments on two datasets showed that our method finds more relevant transformations and obtains better latent representations, in terms of linear evaluation. Future works will try to optimize the transformations composition order. Furthermore, in a weakly-supervised setting, we could also investigate constraining latent space representations of non labeled data with pseudo-labels and nearest neighbor clustering.

## References

1. Chaitanya, K., Erdil, E., Karani, N., Konukoglu, E.: Contrastive learning of global and local features for medical image segmentation with limited annotations. In: *NeurIPS*, vol. 33, pp. 12546–12558 (2020)
2. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *37th International Conference on Machine Learning. Proceedings of Machine Learning Research*, vol. 119, pp. 1597–1607. PMLR (2020)
3. Cohen, J.P., Luck, M., Honari, S.: Distribution matching losses can hallucinate features in medical image translation. In: *MICCAI* (2018)
4. Cubuk, E.D., Zoph, B., Mane, D., Vasudevan, V., Le, Q.V.: AutoAugment: learning augmentation strategies from data. In: *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 113–123 (2019)
5. Dufumier, B., et al.: Contrastive learning with continuous proxy meta-data for 3D MRI classification. In: *MICCAI* (2021)
6. Khosla, P., et al.: Supervised contrastive learning. In: *Advances in Neural Information Processing Systems*, vol. 33, pp. 18661–18673 (2020)
7. Kingma, D.P., Dhariwal, P.: Glow: generative flow with invertible 1x1 convolutions. In: *NeurIPS*, vol. 31 (2018)
8. Li, Y., Hu, G., Wang, Y., Hospedales, T., Robertson, N.M., Yang, Y.: Differentiable automatic data augmentation. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.-M. (eds.) *ECCV 2020. LNCS*, vol. 12367, pp. 580–595. Springer, Cham (2020). [https://doi.org/10.1007/978-3-030-58542-6\\_35](https://doi.org/10.1007/978-3-030-58542-6_35)
9. Liu, A., Huang, Z., Huang, Z., Wang, N.: Direct differentiable augmentation search. In: *2021 IEEE/CVF International Conference on Computer Vision (ICCV)* (2021)
10. Menze, B.H., et al.: The multimodal brain tumor image segmentation benchmark (BRATS). *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2015)

11. Patrick, M., et al.: On compositions of transformations in contrastive self-supervised learning. In: IEEE/CVF International Conference on Computer Vision (ICCV), pp. 9577–9587 (2021)
12. Perakis, A., Gorji, A., Jain, S., Chaitanya, K., Rizza, S., Konukoglu, E.: Contrastive learning of single-cell phenotypic representations for treatment classification. In: MLMI - MICCAI, pp. 565–575 (2021)
13. Tang, Y., et al.: Automated abnormality classification of chest radiographs using deep convolutional neural networks. *NPJ Digit. Med.* **3**, 1–8 (2020)
14. Tian, Y., Sun, C., Poole, B., Krishnan, D., Schmid, C., Isola, P.: What makes for good views for contrastive learning? In: *NeurIPS* (2020)
15. Xiao, T., Wang, X., Efros, A.A., Darrell, T.: What should not be contrastive in contrastive learning. In: *International Conference on Learning Representations* (2021)
16. Yang, S., Das, D., Chang, S., Yun, S., Porikli, F.M.: Distribution estimation to automate transformation policies for self-supervision. In: *Advances in Neural Information Processing Systems* (2021)
17. Zhao, S., Liu, Z., Lin, J., Zhu, J.Y., Han, S.: Differentiable augmentation for data-efficient GAN training. In: *Advances in Neural Information Processing Systems*, vol. 33, pp. 7559–7570 (2020)
18. Zhao, X., et al.: Contrastive learning for label efficient semantic segmentation. In: IEEE/CVF International Conference on Computer Vision (ICCV), pp. 10623–10633 (2021)