



HAL
open science

Learning to diagnose cirrhosis from radiological and histological labels with joint self and weakly-supervised pretraining strategies

Emma Sarfati, Alexandre Bone, Marc-Michel Rohe, Pietro Gori, Isabelle Bloch

► To cite this version:

Emma Sarfati, Alexandre Bone, Marc-Michel Rohe, Pietro Gori, Isabelle Bloch. Learning to diagnose cirrhosis from radiological and histological labels with joint self and weakly-supervised pretraining strategies. IEEE ISBI, Apr 2023, Cartagena de Indias, Colombia. hal-03993204

HAL Id: hal-03993204

<https://telecom-paris.hal.science/hal-03993204v1>

Submitted on 16 Feb 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

LEARNING TO DIAGNOSE CIRRHOSIS FROM RADIOLOGICAL AND HISTOLOGICAL LABELS WITH JOINT SELF AND WEAKLY-SUPERVISED PRETRAINING STRATEGIES

Emma Sarfati^{1,2} Alexandre Bône¹ Marc-Michel Rohé¹ Pietro Gori² Isabelle Bloch^{2,3}

¹ Guerbet Research, Villepinte, France

² LTCI, Télécom Paris, Institut Polytechnique de Paris, Saclay, France

³ Sorbonne Université, CNRS, LIP6, Paris, France

ABSTRACT

Identifying cirrhosis is key to correctly assess the health of the liver. However, the gold standard diagnosis of the cirrhosis needs a medical intervention to obtain the histological confirmation, *e.g.* the METAVIR score, as the radiological presentation can be equivocal. In this work, we propose to leverage transfer learning from large datasets annotated by radiologists, which we consider as a weak annotation, to predict the histological score available on a small annex dataset. To this end, we propose to compare different pretraining methods, namely weakly-supervised and self-supervised ones, to improve the prediction of the cirrhosis. Finally, we introduce a loss function combining both supervised and self-supervised frameworks for pretraining. This method outperforms the baseline classification of the METAVIR score, reaching an AUC of 0.84 and a balanced accuracy of 0.75, compared to 0.77 and 0.72 for a baseline classifier.

Index Terms— Deep Learning, Contrastive Learning, Medical Image Classification, Cirrhosis Prediction, Liver.

1. INTRODUCTION

Cirrhosis diagnosis is important for radiologists, as it can support the differential diagnosis of liver masses, such as hepatocellular carcinoma, a liver primary cancer [1]. The presence - or absence - of cirrhosis can more generally be a signal for hepatic pathologies. In clinical routine, the diagnosis of cirrhosis can be performed with three different approaches. First, the gold standard method is the histological analysis obtained by biopsy (or following a resection). However, this method is clinically invasive, hence risky and expensive. Secondly, a clinical examination is done to detect the potential signs of a terminal cirrhosis stage (an interrogation about alcohol consumption, visible signs of jaundice or ascites, related to a swollen belly). Thirdly, CT-scans in portal venous phase can be analyzed by radiologists to find imaging features of the disease, but the task remains difficult and mainly uneven as the diagnosis can change from one radiologist to another, with typically low inter-rate agreement scores [2].

Several methods have been proposed for automatic cirrhosis prediction from medical images. These methods use mainly Deep Convolutional Neural Networks (DCNN) for cirrhosis prediction, considered as a classification or a regression problem [3]. While some methods use large backbones with millions of parameters as encoder such as DenseNet-121 or ResNet-Inception-v2 [4, 5], lighter networks have proved to provide very good results in terms of accuracy [6]. However, these state-of-the-art methods rely on histopathological diagnosis as label features [4], using the METAVIR score classification corresponding to different

stages of fibrosis (F0/F1/F2/F3/F4), or the Inuyama score classification, which is close to the METAVIR one. The majority of these studies use large labeled datasets with hundreds or thousands of histologically-diagnosed patients [7]. Replicating these approaches requires large volumes of images with corresponding biopsy, which are difficult to obtain. By contrast, obtaining large volumes of CT-scans without annotations is much easier, and getting *a posteriori* annotations from radiologists is still possible because no medical intervention is needed. To cope with the limited data availability, deep learning studies have demonstrated the possibility of advantageously pre-training models on large databases to prepare their transfer on a deployment database [8]. Pretraining databases can be labeled (and used in transfer learning [9]), unlabeled (used in self-supervised learning, for instance with SimCLR [10]), or weakly-labeled like in [11], *i.e.* annotated with a label that is close to the reference one, and hence that can be regarded as a proxy for the latter. Pre-training can be beneficial when there are few labeled images and many unlabeled (or weakly/noisily labeled) images.

In this work, we propose to explore several pretraining methods to improve the prediction of a binarized METAVIR score from a small CT-scan dataset, using a large weakly-labeled CT-scan dataset. We compare three different approaches; first, we explore the effect of a standard transfer learning method to improve the prediction of the METAVIR score, *i.e.* we pretrain a supervised model on a weak (or noisy) label and then re-use the weights to predict the strong label. The second part presents the impact of self-supervised pretraining (SimCLR [10]) for the same purpose. Finally, we study the introduction of the radiological label within the self-supervised framework, first using the existing Supervised Contrastive Learning model (SupCon, [9]), then enhancing the latter by proposing a weighted sum of SimCLR and CrossEntropy loss functions. In this article, we denote the radiological labels as “weak” labels, as we suppose that they can be seen as noisy approximations of the reference histological labels.

2. METHOD

Dataset. Two datasets are leveraged in this study. First, \mathcal{D}_{histo} , contains 106 CT-scans from different patients in portal venous phase, with an identified histopathological status obtained by a histological analysis, designated as Y_{histo} . The latter is binarized to indicate the absence or presence of advanced fibrosis [4] obtained by the separation F0/F1/F2 vs. F3/F4. The pathological class contains 78 patients while the healthy one includes 28 patients. The second dataset, \mathcal{D}_{radio} , consists of 2,799 CT-scans of patients in portal venous phase with a radiological annotation, *i.e.* realized by a radiologist, indicating four different stages of cirrhosis: no cirrhosis, mild cirrhosis, moderate cirrhosis and severe. We also bi-

narize this label to obtain no cirrhosis versus mild/moderate/severe (Y_{radio}). This dataset contains 919 pathological (supposedly cirrhotic, i.e. mild/moderate/severe) subjects and 1,880 healthy ones.

All images have a 512x512 size, and we clip the intensity values between -100 and 400. We work with 2D slices rather than 3D volumes. Moreover, we select the slices based on the liver segmentation of the patients. We keep the top 70% most central slices with respect to automatically-computed liver segmentation maps.

Architecture and optimization.

Backbone. Inspired by [6], we propose a baseline backbone based on a simple architecture, illustrated in Figure 1. The 512x512 input images are passed through five convolutional layers, each of them using a 5x5 kernel, followed by a ReLU activation function before a 2x2 MaxPooling operation. The first layer has 32 channels, and the number of channels is doubled at every layer. It then generates a 512x20x20 feature map that is pool-averaged, ending in a 512-dimensional flat vector. Two dense layers end the network, mapping the 512-dimensional vector to a 256-dimensional vector, then to a binary output followed by a softmax. For self-supervised learning, we replace the last linear layer by another one with an output dimension of 128 to be consistent with the original SimCLR method [10]. We denote by $f(\cdot)$ the backbone encoder preceding the dense layers, and $p(\cdot)$ the projector composed of the two dense layers. We obtain either a 128-dimensional or a 2-dimensional output, $z = p(f(x))$. This architecture is used as a basis for all our experiments.




Fig. 1: The DCNN used in our method.

Sampling strategy and loss functions. For the baseline performance, the proposed encoder (Figure 1) is learned on \mathcal{D}_{histo} . In a preliminary experiment, the ratio of leveraged training data is artificially reduced from 100% to 80%, 60% and 40% in order to assess the impact of the number of cases on classification performance. Then, four pretraining experiments are led to improve the baseline performance for cirrhosis classification.

First, we train a supervised model (backbone architecture, Figure 1) to predict the binarized radiological labels present in \mathcal{D}_{radio} , which we consider as a weak label for the gold standard Y_{histo} , and then use a transfer learning strategy to predict the binarized histological labels in \mathcal{D}_{histo} . This first pretraining experiment can be regarded as weakly-supervised for our purpose, hence we denote the binary cross entropy used there as \mathcal{L}_{weak} .

As a second experiment, we leverage a self-supervised pretraining approach, SimCLR, using the original NTXentLoss [10, 12]:

$$\mathcal{L}_{SimCLR} = -\frac{1}{2N} \sum_{i=1}^{2N} \log \frac{\exp(\text{sim}(z_i, z_{j(i)})/\tau)}{\sum_{k=1}^{2N} \mathbb{1}[k \neq i] \exp(\text{sim}(z_i, z_k)/\tau)}$$

with $j(i)$ denoting the positive with respect to i , i.e. the second augmented version of the original image x_i , N denoting the batch size and $z_i = p(f(x_i))$, $z_{j(i)} = p(f(x_{j(i)}))$ denote the output vectors

of the image augmentations x_i and $x_{j(i)}$ passed through two random augmentations modules, \mathcal{T} and \mathcal{T}' . The similarity is defined as $\text{sim}(z_i, z_j) = z_i^T z_j$. We fix the temperature parameter at $\tau = 0.1$.

For the third experiment, we explore SupCon [9] using Y_{radio} to pair samples from the same class together:

$$\mathcal{L}_{SupCon} = \sum_{i=1}^{2N} \frac{-1}{|P(i)|} \sum_{j \in P(i)} \log \frac{\exp(\text{sim}(z_i, z_j)/\tau)}{\sum_{k=1}^{2N} \mathbb{1}[k \neq i] \exp(\text{sim}(z_i, z_k)/\tau)}$$

where $P(i)$ denotes the set of all the indices of samples belonging to the same class as the input image.

Fourth, to maximize the potential information given by Y_{radio} as well as the representation power offered by SimCLR, we propose a new loss function for pre-training, $\mathcal{L}_{weak-SimCLR}$, a simple weighted sum of the binary cross entropy and the NTXentLoss:

$$\mathcal{L}_{weak-SimCLR} = \beta \mathcal{L}_{weak} + (1 - \beta) \mathcal{L}_{SimCLR} \quad (1)$$

where $\beta \in [0, 1]$ is an hyper-parameter. To compute this function, the only change in the training process is that the original image is passed through a third data augmentation module \mathcal{T}'' , before being passed to the backbone and two dense layers, mapping the 512-dimensional representation vector to a 256-dimensional then to a 2-dimensional one. Note that all the weights are shared between the supervised and unsupervised branches, only the last dense layers, due to the difference in output dimensions, differ between both (see Figure 3).

Finally, for sampling, we observe a class imbalance which we try to fix using a weighted sampling during training, and we report the results of the balanced accuracy scores. As we work with 2D slices rather than 3D volumes, we compute the average probability of having the pathology per patient. The evaluation results presented later are based on the patient-level aggregated prediction.

Data augmentation and optimization setting. Unsupervised contrastive learning methods such as SimCLR [10] typically require heavy data augmentations on input images, in order to strengthen the association between positive samples in the representation space [13]. In our work, we leverage three specific types of augmentations: crops, flips and rotations. During our experiments, we also inspected the effect of CutOut [14], which proved not to increase the performances of our models. Data augmentations are computed on the GPU, using the Kornia library [15]. During inference, we remove the augmentation module to only keep the original input images.

We run our experiments on a Tesla V100 with 16GB of RAM and a 6 CPU cores and we used the PyTorch-Lightning library to implement our models. All the models share the same random data augmentation module, with a batch size of $N = 92$ and a fixed number of epochs $n_{epochs} = 200$. For the pretraining experiments, i.e. the models trained on the large dataset \mathcal{D}_{radio} , we fix a learning rate (LR) of $\alpha = 10^{-4}$ and a weight decay of $\lambda = 10^{-4}$. For the classification experiments, i.e. the models trained on the small dataset \mathcal{D}_{histo} , we fix a learning rate (LR) of $\alpha = 10^{-5}$ and a weight decay of $\lambda = 10^{-3}$. For all the experiments, we added a cosine decay learning rate scheduler. Finally, we fix the hyper-parameter β of *weak-SimCLR* at 0.5, unless otherwise specified.

Evaluation protocol. We evaluate our methods using two different procedures. First, we extract the 512-dimensional vectors $f(x)$ from the representation space (see Figure 3) and train a simple logistic regression on the frozen representations with a default regularization parameter of $\lambda = 1$, using scikit-learn. This procedure can be

thought as a linear evaluation in the representation space. Hence, we can denote this first evaluation method a “cross-validated (CV) linear evaluation”, as introduced in Table 1. Secondly, we fine-tune the whole network initializing the latter with the pretrained weights. For both evaluation procedures, we validate the results with a stratified 5-fold cross-validation.

3. RESULTS




Fig. 2: Evolution of the averaged cross-validation AUC with respect to the percentage of the training data.

We first report in Figure 2 the evolution of the averaged cross-validation AUC with respect to the percentage of the training data kept during training, as well as the associated standard deviations. As a reference, [3, 6] reported respectively an AUC of 0.76 for 186 patients and 0.89 for 202 patients in their training sets, while \mathcal{D}_{histo} presents 85 patients for training and 21 for validation.

Pretraining method	AUC		Balanced Accuracy	
	CV linear evaluation	Fine-tuning	CV linear evaluation	Fine-tuning
None	/	0.77 (± 0.07)	/	0.72 (± 0.07)
Supervised	0.64 (± 0.19)	0.65 (± 0.10)	0.56 (± 0.10)	0.64 (± 0.09)
SimCLR	0.78 (± 0.09)	0.78 (± 0.11)	0.68 (± 0.04)	0.69 (± 0.08)
SupCon	0.61 (± 0.05)	0.65 (± 0.13)	0.59 (± 0.05)	0.67 (± 0.14)
Ours	0.84 (± 0.05)	0.81 (± 0.03)	0.75 (± 0.06)	0.73 (± 0.08)

Table 1: Results AUCs and balanced accuracies of our experiments, with a value of $\beta = 0.5$ for *weak-SimCLR*. Best results by column are **underlined**. The standard deviations come from the AUCs within the folds and the AUCs are averaged by fold.

Transfer learning results. Table 1 gathers the AUC and balanced accuracy performance measures for each transfer learning configuration. It first provides results of the supervised classifier directly trained on \mathcal{D}_{histo} , from random initial weights, which is the baseline result. Next, it presents the performances of the four pretraining experiments that were led, ending with the proposed model, *weak-SimCLR*. First, it shows that both SimCLR and *weak-SimCLR* overcome the baseline AUC. In particular, *weak-SimCLR* presents an increase of 7% in cross-validated linear evaluation, and 5% in fine-tuning, which confirms that the representation space built by the proposed model can provide a globally relevant separation. It can be noted that the first evaluation method, training a logistic regression on the frozen representation vectors previously trained on \mathcal{D}_{radio} , is faster and less computationally demanding than the full fine-tuning. For the accuracy, the proposed method slightly outperforms the baseline score, in fine-tuning and transfer learning with respectively 0.75 and 0.73 for balanced accuracy scores, compared to 0.72 for the supervised classifier.

β	AUC		Balanced Accuracy	
	CV linear evaluation	Fine-tuning	CV linear evaluation	Fine-tuning
0 (SimCLR)	/	0.77 (± 0.07)	/	0.72 (± 0.07)
0.2	0.78 (± 0.09)	0.78 (± 0.11)	0.68 (± 0.04)	0.69 (± 0.08)
0.4	0.75 (± 0.08)	0.78 (± 0.10)	0.67 (± 0.05)	0.73 (± 0.13)
0.5	0.82 (± 0.10)	0.81 (± 0.07)	0.71 (± 0.09)	0.75 (± 0.07)
0.8	0.84 (± 0.05)	0.81 (± 0.03)	0.75 (± 0.06)	0.73 (± 0.08)
1 (Supervised)	0.74 (± 0.06)	0.79 (± 0.09)	0.68 (± 0.07)	0.66 (± 0.11)
1 (Supervised)	0.64 (± 0.19)	0.65 (± 0.10)	0.56 (± 0.10)	0.64 (± 0.09)

Table 2: Results AUCs and balanced accuracies of the proposed model, *weak-SimCLR*, with respect to the value of β (supervision) incorporated in the model. Best results by column are **underlined**.

Robustness study. Table 2 gathers the AUC and balanced accuracy performance measures for each value of β of the proposed method, *weak-SimCLR*. For the cross-validated linear evaluation, the value of $\beta = 0.5$ reaches the best AUC and balanced accuracy performances. For the AUC, it can be noted that SimCLR and both $\beta = 0.4$ and $\beta = 0.5$ of *weak-SimCLR* outperform the baseline score of 0.77. Moreover, all the tested values of β , except 0 and 1, provide a better AUC result for fine-tuning. In terms of accuracy, only the value of $\beta = 0.5$ gives a better result than the baseline score with a gain of 3%. For the fine-tuning part, the scores are stable with respect to the CV linear evaluation ones.

4. DISCUSSION AND CONCLUSION

We explored four pretraining methods to improve the cirrhosis classification on a small histologically-labeled dataset, using a large radiologically-labeled dataset. We proposed a method to improve the automatic diagnosis of cirrhosis, using radiological (weak) annotations for pre-training models. To the best of our knowledge, no other work using radiological annotation to improve the automatic prediction of the histological one for cirrhosis prediction with deep learning has been proposed in the literature. The proposed method relies on the combination of weakly-supervised and self-supervised approaches, with the share of each model being a hyper-parameter that can be tuned. It improved the automatic diagnosis of cirrhosis, based on both AUCs and balanced accuracies. Notably, it outperforms two state-of-the-art methods in contrastive learning [10, 9], as well as the baseline classification realized on our small histologically labeled dataset. We obtain a competitive performance with respect to the literature [6], reaching an AUC of 0.84, with only 106 annotated CT-scans.

The first limit of our work is the baseline performance which may be improved by testing different backbones, other than the one we proposed (see Figure 1), such as ResNet-50. The self-supervised block in the proposed model could also be tested with other contrastive or non-contrastive methods (e.g. SimSIAM [16] or BYOL [17]). Second, it could be interesting to evaluate our method on an external public dataset such as the Liver Hepatocellular Carcinoma (LIHC) dataset from the Cancer Genome Atlas [18]. Finally, more evaluation procedures could be added to assess the robustness of the proposed strategy regarding the small number of patients in \mathcal{D}_{histo} . Indeed, the standard deviations in our results are important, due to some disparities between the folds. K-fold cross validation with higher values of K could be tested, or even Leave-One-Out cross-validation.

Compliance with ethical standards. This research study was conducted retrospectively using human data collected from various medical centers, whose Ethics Committees granted their approval.

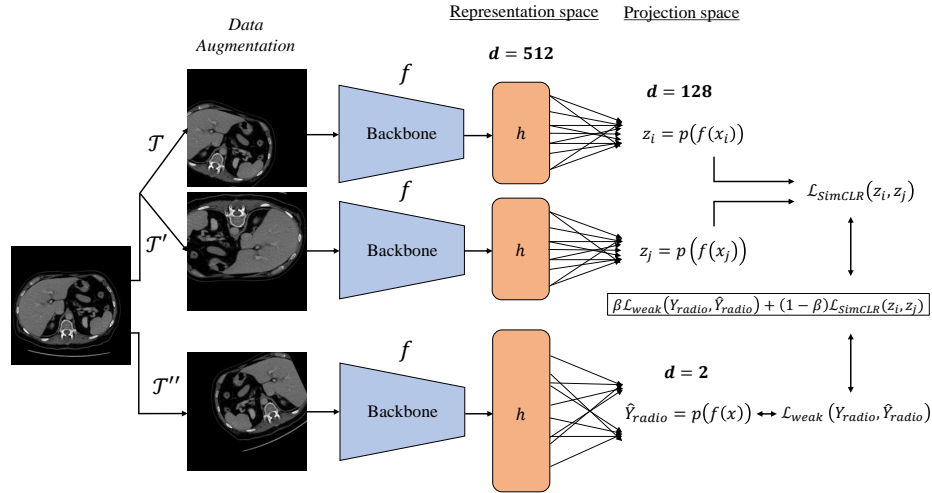


Fig. 3: The proposed pretraining strategy, *weak-SimCLR*, which loss function is a sum of the binary cross-entropy computed with Y_{radio} (\mathcal{L}_{weak}) and of the SimCLR loss function (\mathcal{L}_{SimCLR}).

Data was de-identified and processed according to all applicable privacy laws and the Declaration of Helsinki.

Acknowledgments. This work was supported by Région Île-de-France (ChoTherIA project) and ANRT (CIFRE #2021/1735).

5. REFERENCES

- [1] Kazuo Terao, Akito Nozaki, Takaaki Ikeda, et al., “Real impact of liver cirrhosis on the development of hepatocellular carcinoma in various liver diseases—meta-analytic assessment,” *Cancer Medicine*, vol. 8, 02 2019.
- [2] C. Aubé, P. Bazeries, J. Lebigot, et al., “Liver fibrosis, cirrhosis, and cirrhosis-related nodules: Imaging diagnosis and surveillance,” *Diagnostic and Interventional Imaging*, vol. 98, no. 6, pp. 455–468, 2017.
- [3] Koichiro Yasaka, Hiroyuki Akai, Akira Kunimatsu, et al., “Deep learning for staging liver fibrosis on CT: a pilot study,” *European Radiology*, vol. 28, 05 2018.
- [4] Qiuju Li, Bing Yu, Xi Tian, et al., “Deep residual nets model for staging liver fibrosis on plain CT images,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 15, 06 2020.
- [5] Michał Byra, Grzegorz Styczynski, Cezary Szmigielski, et al., “Transfer learning with deep convolutional neural network for liver steatosis assessment in ultrasound images,” *International Journal of Computer Assisted Radiology and Surgery*, vol. 13, 08 2018.
- [6] Yunchao Yin, Derya Yakar, Rudi Dierckx, et al., “Liver fibrosis staging by deep learning: a visual-based explanation of diagnostic decisions of the model,” *European Radiology*, vol. 31, 05 2021.
- [7] Kyu Choi, Jong Jang, Seung Lee, et al., “Development and validation of a deep learning system for staging liver fibrosis by using contrast agent-enhanced CT images in the liver,” *Radiology*, vol. 289, pp. 180763, 09 2018.
- [8] Zhizhong Li and Derek Hoiem, “Learning without forgetting,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 12, pp. 2935–2947, 2018.
- [9] Prannay Khosla, Piotr Teterwak, Chen Wang, et al., “Supervised contrastive learning,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 18661–18673, 2020.
- [10] Ting Chen, Simon Kornblith, Mohammad Norouzi, et al., “A simple framework for contrastive learning of visual representations,” in *37th International Conference on Machine Learning (ICML)*, 2020.
- [11] Benoit Dufumier, Pietro Gori, Julie Victor, et al., “Contrastive learning with continuous proxy meta-data for 3D MRI classification,” in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2021, pp. 58–68.
- [12] Ting Chen, Simon Kornblith, Kevin Swersky, et al., “Big self-supervised models are strong semi-supervised learners,” in *34th International Conference on Neural Information Processing Systems*, Red Hook, NY, USA, 2020, NIPS’20, Curran Associates Inc.
- [13] Xiao Wang and Guo-Jun Qi, “Contrastive learning with stronger augmentations,” 2021.
- [14] Terrance DeVries and Graham W. Taylor, “Improved regularization of convolutional neural networks with cutout,” 2017.
- [15] Edgar Riba, Dmytro Mishkin, Daniel Ponsa, et al., “Kornia: an open source differentiable computer vision library for pytorch,” 2019.
- [16] Xinlei Chen and Kaiming He, “Exploring simple siamese representation learning,” 2020.
- [17] Jean-Bastien Grill, Florian Strub, Florent Altché, et al., “Bootstrap your own latent - a new approach to self-supervised learning,” in *Advances in Neural Information Processing Systems*, 2020, vol. 33.
- [18] B. J. Erickson, S. Kirk, Lee, et al., “Radiology data from the cancer genome atlas colon adenocarcinoma [tcga-coad] collection,” 2016.