Semixup: In- and Out-of-Manifold Regularization for Deep Semi-Supervised Knee Osteoarthritis Severity Grading From Plain Radiographs

Huy Hoang Nguyen, Simo Saarakkala, Matthew B. Blaschko, and Aleksei Tiulpin

***Abstract*— Knee osteoarthritis (OA) is one of the highest disability factors in the world. This musculoskeletal disorder is assessed from clinical symptoms, and typically confirmed via radiographic assessment. This visual assessment done by a radiologist requires experience, and suffers from moderate to high inter-observer variability. The recent literature has shown that deep learning methods can reliably perform the OA severity assessment according to the gold standard Kellgren-Lawrence (KL) grading system. However, these methods require large amounts of labeled data, which are costly to obtain. In this study, we propose the *Semixup* algorithm, a semi-supervised learning (SSL) approach to leverage unlabeled data. *Semixup* relies on consistency regularization using in- and out-of-manifold samples, together with interpolated consistency. On an independent test set, our method significantly outperformed other state-of-the-art SSL methods in most cases. Finally, when compared to a well-tuned fully supervised baseline that yielded a balanced accuracy (BA) of 70.9 ± 0.8% on the test set, *Semixup* had comparable performance – BA of 71 ± 0.8% ($p = 0.368$) while requiring 6 times less labeled data. These results show that our proposed SSL method allows building fully automatic OA severity assessment tools with datasets that are available outside research settings.**

***Index Terms*— Deep learning, knee, osteoarthritis, semi-supervised learning.**

I. Introduction

OSTEOARTHRITIS (OA) is the most common joint disorder in the world causing enormous burdens at personal and societal levels [1]. OA has an unknown etiology, and its indications at late stages are worn cartilage, bone deformity, and synovitis [2]–[4].

The most common joints affected by OA are knee and hip, and among these, the disease is more prevalent in knee [5]–[8]. At the population level, such factors as sex, body-mass index (BMI), and age are known to be associated with OA [9]–[11]. As such, it was previously shown that people with BMI over 30 have 7-fold higher risk of knee OA than ones with BMI below 25 [12], and a half of elderly people over 65 years of age have OA in at least one joint [13].

From an economic perspective, OA leads to a huge burden in terms of direct costs (e.g, hospitalization, diagnosis, and therapy), and indirect ones (e.g. losses of working days and productivity) [14]. For example, in the United States, OA costs hundreds of billion dollars annually, and is in the top-5 of annual Europe healthcare expenditure [3], [7].

Currently, knee OA diagnosis starts with a clinical examination, and then, a radiographic confirmation takes place when necessary [1], [5]. However, such practice enables knee OA diagnosis only at a late stage when the cartilage is already worn, and the bone deformity is present, which leads to severe pain, and even physical disability [2], [3]. Ultimately, the only remaining option for a patient in that scenario is total knee replacement (TKR) surgery.

The literature shows a large and rapidly growing number of TKR surgeries worldwide [8], [15], [16]. As such, the annual rate of TKR surgeries in the United States has doubled since 2000 for adults of 45-64 years old [6], [17]. Therefore, there is a need for prevention of global disability.

Imaging, in contrast to clinical examination, may enable the detection of early knee OA signs at the stages when behavioral interventions (e.g. exercises and weight loss programs) could

Manuscript received July 16, 2020; accepted August 10, 2020. Date of publication August 17, 2020; date of current version November 30, 2020. This work was supported in part by the strategic funding of the University of Oulu; in part by KAUTE foundation, Finland; in part by Sigrid Juselius Foundation, Finland; in part by the OAI, a public-private partnership, funded by the National Institutes of Health, a branch of the Department of Health and Human Services, and conducted by the OAI Study Investigators, under Grant N01- AR-2-2258, Grant N01-AR-2-2259, Grant N01-AR-2- 2260, Grant N01-AR-2-2261, and Grant N01-AR-2-2262; in part by the MOST, funded by the National Institutes of Health, a branch of the Department of Health and Human Services, and conducted by MOST study investigators, under Grant Felson - AG18820, Grant Torner - AG18832, Grant Lewis - AG18947, and Grant Nevitt - AG19069; in part by the Merck Research Laboratories, in part by Novartis Pharmaceuticals Corporation, in part by GlaxoSmithKline; and in part by Pfizer, Inc. *(Corresponding author: Huy Hoang Nguyen.)*

Huy Hoang Nguyen is with the Research Unit of Medical Imaging, Physics and Technology, University of Oulu, 90570 Oulu, Finland (e-mail: huy.nguyen@oulu.fi).

Simo Saarakkala is with the Research Unit of Medical Imaging, Physics and Technology, University of Oulu, 90570 Oulu, Finland, and also with the Department of Diagnostic Radiology, Oulu University Hospital, 90220 Oulu, Finland (e-mail: simo.saarakkala@oulu.fi).

Matthew B. Blaschko is with the Center for Processing Speech and Images, KU Leuven, 3000 Leuven, Belgium (e-mail: matthew.blaschko@esat.kuleuven.be).

Aleksei Tiulpin is with the Research Unit of Medical Imaging, Physics and Technology, University of Oulu, 90570 Oulu, Finland, also with the Department of Diagnostic Radiology, Oulu University Hospital, 90220 Oulu, Finland, and also with Ailean Technologies Oy, 90100 Oulu, Finland (e-mail: aleksei.tiulpin@oulu.fi).

This article has supplementary downloadable material available at https://ieeexplore.ieee.org, provided by the authors.

Digital Object Identifier 10.1109/TMI.2020.3017007

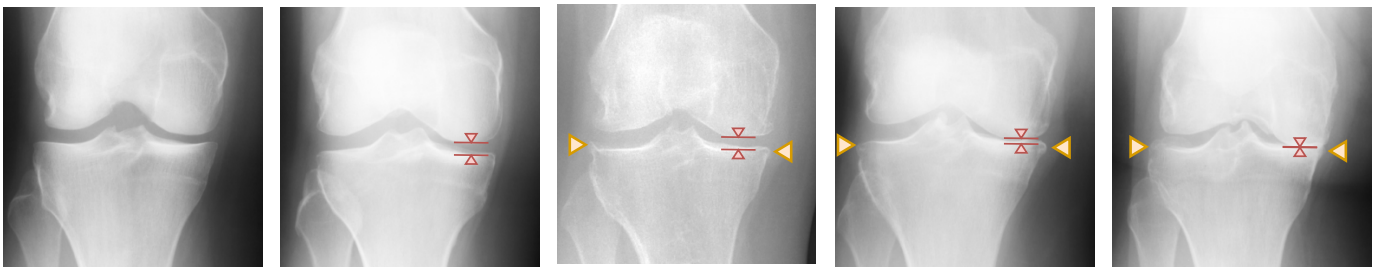

(a) (b) (c) (d) (e)

Fig. 1. Samples of knee radiographs. Joint space narrowing and osteophyte features are indicated by red and yellow marks respectively. (a) KL 0: A healthy knee without OA, (b) KL 1 (Doubtful OA): Potential joint space narrowing, (c) KL 2 (Mild OA): Clear evidences of osteophytes, as well as slight reduction of joint space, (d) KL 3 (Moderate OA): Osteophytes grow and joint space narrowing progresses badly, and (e) KL 4 (Severe OA): Besides osteophytes, joint space is reduced so severely that the tibia and the femur are connected.

slow down the disease progression [18]. Radiographic assessment is the foremost imaging tool for detecting knee OA in primary care, and Kellgren-Lawrence (KL) is one of the most common clinical scales for the assessment of OA severity from plain radiographs (Figure 1). However, visual diagnosis done by a radiologist suffers from low inter-rater agreement [19], [20], thereby introducing large inconsistencies into decision-making. One possible solution to make OA diagnosis more systematic and allow for the detection of knee OA at early stages is to leverage computer-aided methods for image analysis [21]. Deep learning (DL) has become a state-of-the-art approach in this realm, and recent studies [20], [22], [23] have demonstrated that DL-based methods allow fully-automatic KL grading. Furthermore, these studies showed a high level of agreement between the predictions made by DL-based models and the annotations produced by a consensus of radiologists.

Despite good and promising results, all the previously published DL methods in OA domain were based on Supervised-Learning (SL) and required large amounts of labeled data, which are not currently widely available. In practical applications, such datasets as the Osteoarthritis Initiative (OAI, https://nda.nih.gov/oai/) and the Multicenter Osteoarthritis Study (MOST, http://most.ucsf.edu/) are expensive to obtain due to high costs of data acquisition and annotation. As such, the latter needs multiple skilled experts (e.g. radiologists or orthopedists) making the process even more costly. Whereas labeled data are difficult to acquire, unlabeled data are available in large amounts, and can be collected from hospital imaging archives at low cost.

In the natural image recognition domain, it has been shown that leveraging small amounts of labeled and large amounts of unlabeled data in a Semi-Supervised Learning (SSL) setting could potentially resolve the need for large labeled datasets. Recent studies [24]–[26] have developed SSL methods to utilize unlabeled data during the training processes, and have achieved competitive performances in image classification benchmarks using only a small fraction of the labeled data used in fully supervised settings.

Many SSL-based applications in the medical domain have previously been developed for automatic disease diagnosis. However, most of those are related to medical image segmentation [27]–[31], and use generative adversarial networks (GANs) [32] as a core method. To the best of our knowledge, there have been no SSL-based methods developed in the knee OA realm.

In this study, we, for the first time in the OA field, propose to leverage SSL for automatic assessment of knee OA severity from plain radiographs. Inspired by previous research in SSL [24], [25], [33], and the recently developed technique *mixup* [34], we propose a novel SSL method – *Semixup*, providing its systematic empirical comparison with the state-of-the-art approaches. Specifically, our contributions are the following:

1) We enhance the state-of-the-art supervised baseline [20], and propose a novel *Separable Adaptive Max-pooling* (SAM), as a drop-in replacement for the Global Average Pooling (GAP). This allows us to significantly improve over previously reported supervised results.
2) We introduce a novel semi-supervised DL-based method called *Semixup* for automatic KL grading of knee OA from plain radiographs. Our method yields competitive results to a well-tuned SL model trained on over 6 times more labeled data.
3) We systematically compare our method against several state-of-the-art SSL methods, and experimentally show that *Semixup* outperforms them in nearly all data regimes.
4) We follow the guidelines from [35] to conduct a realistic evaluation of our SSL approach, and provide insights into the scalability of our method with respect to the number of labeled examples, together with a tractable amount of unlabeled samples.

II. Related Work

A. Deep Semi-Supervised Learning

There exists a wide variety of DL-based SSL methods in the literature; however, we discuss here only the ones that are close to our method and yield state-of-the-art results on generic image recognition datasets. Such approaches use two main ideas: *consistency regularization* and *pseudo-labeling*.

Consistency regularization is based on the assumption that the predictions of the model $P(y|\mathbf{x})$ and $P(y|T\mathbf{x})$, where $\mathbf{x}$ is a data point and T – class-preserving stochastic data augmentation – should not differ. The methods using the

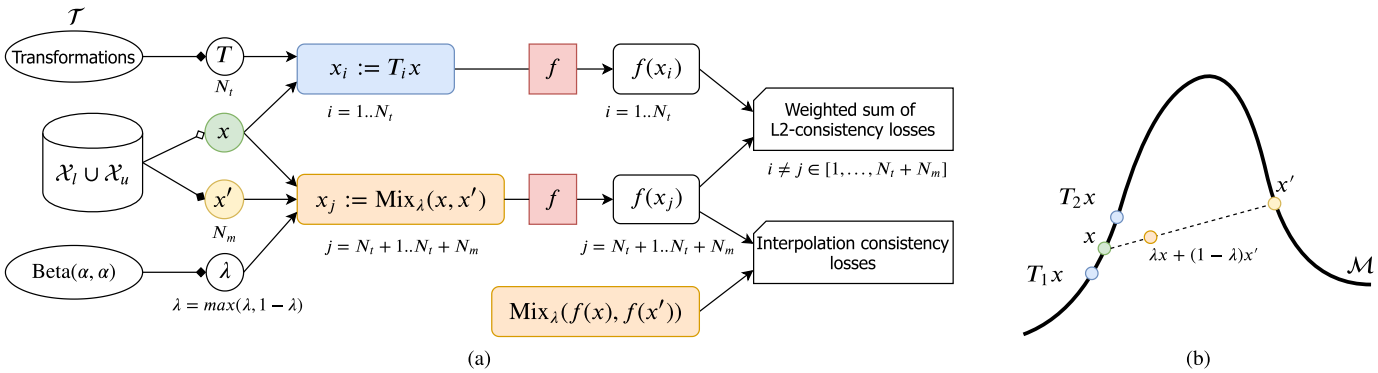

Fig. 2. The main idea of *Semixup*: (a) The workflow of its in- and out-of-manifold consistency regularization branch, and (b) its conceptual illustration. For each image $x \in L \cup U$, we sample N_t transformations, N_m images x', and *mixup* coefficients λ (s.t. $\lambda > 0.5$) to transform and blend with x. Then, we enforce consistency of predictions for every pair x_i, x_j under L2 norm. In addition, we use the interpolation consistency regularizer. Links with ■ denote sampling with replacement, ones with □ indicate sampling without replacement.

technique applied to unlabeled data include the Π-model [24], and Mean Teacher (MT) [36].

Label guessing (or pseudo-labeling) [37] was proposed in a deep learning setting in [38], and uses predicted labels for unlabeled samples with high confidences to update the gradients of the neural network. This technique can also be viewed as entropy regularization which favors low-density separation between classes [39].

The aforementioned SSL techniques have been explored separately; however, Berthelot *et al.* has recently introduced the MixMatch method [26], the idea of which was to combine label guessing and entropy regularization into a holistic framework. The authors of MixMatch made an empirical observation that applying *mixup* [34], which performed convex combinations of 2 arbitrary input data points x and x', and their labels y and y' (i.e. $\lambda x + (1-\lambda)x'$ and $\lambda y + (1-\lambda)y'$, where $\lambda \sim \text{Beta}(\alpha, \alpha)$, for $\alpha \in (0, \infty)$), to labeled data and unlabeled data with guessed labels helps to improve the performance. In our method, we also use *mixup* for both labeled and unlabeled data; however, we avoid using label guessing due to its potential to propagate label errors that can be common in medical domains.

We finally discuss here two main limitations of all the mentioned recent studies on SSL. The first common limitation of those methods is that their evaluations were done on major image recognition benchmarks, namely CIFAR-10 and/or ImageNet without considering real-world problems, such as the ones related to medical image recognition.

The second limitation of the aforementioned studies is that the methods proposed in them were evaluated only with respect to the change of the amount of labeled data. However, none of those studies explored the amount of unlabeled data needed to improve the performance. In our work, we conduct such evaluation.

B. Deep Learning for Knee Osteoarthritis Diagnosis

Fully supervised DL-based methods that use Convolutional Neural Networks (CNN) have recently been used to assess knee OA severity. In particular, in their pioneering work, Antony *et al.* [40] applied transfer learning and showed significant improvements compared to classifiers trained using hand-crafted features. In the follow-up work [22], Antony *et al.* proposed a CNN architecture trained from scratch to classify knee images according to the KL scale. Their new approach outperformed their previous transfer learning results.

The main common limitation of the aforementioned studies was that neither of them utilized an independent test set. In contrast, Tiulpin *et al.* [20] addressed this limitation and also proposed a novel CNN architecture that outperformed the previous methods [22], [40]. Interestingly, that model performed on-par with transfer learning baseline while having significantly less trainable parameters thanks to the inductive bias from using the relative symmetry of visual features in knee images.

More recent studies by Chen *et al.* [41], Norman *et al.* [42], and Górriz *et al.* [43] did not use any independent test set either. The only latest study where an independent test set was used for assessment of the results was by Tiulpin *et al.* [44]. In that study, the authors obtained a KL classification model as a bi-product of their main method; however, they obtained results that are similar to their previous study [20].

Despite having an independent test set and state-of-the-art results, Tiulpin *et al.* [20], [44] used large amounts of labeled data for training. We emphasize here that none of the existing studies in which DL was applied for knee OA diagnosis from radiographs addressed the question of the quality of automatic KL grading as a function of the dataset size. Our work answers this question via a thorough experimental evaluation of both SL and SSL methods.

III. Method

A. Overview

The method proposed in this paper consists of two parts: 1) a novel extension of a previously developed Siamese network developed by Tiulpin *et al.* [20] and 2) a novel deep SSL technique. The utilized Siamese model uses shared branches of the CNN which focus their attention on the medial and the lateral sides of the analyzed knee (see Figure 3). Here, we consider pairs of lateral and medial image patches as

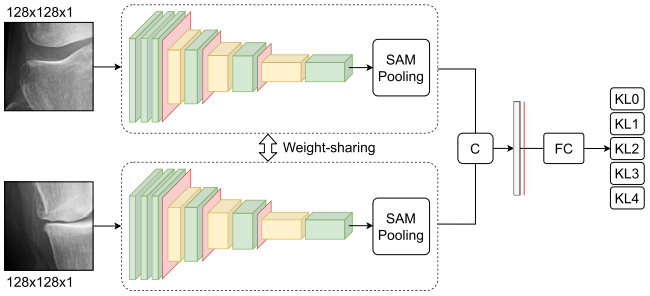

Fig. 3. The Siamese architecture of our model. Green denotes the blocks that use 3×3 convolutions with the stride of 1, and yellow denotes 3×3 convolutions with the stride of 2. Red indicates dropout layers. After pooling the features by Separable Adaptive Max-pooling (SAM) and concatenating (C) them into a single vector, they are passed into a fully-connected (FC) layer that predicts KL grades.

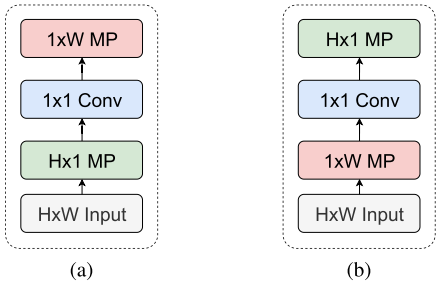

Fig. 4. *Separable Adaptive Max-pooling* configurations. The core idea of this approach is to inject a non-linearity between the pooling steps – vertical-horizontal and horizontal-vertical as displayed in subplots (a) and (b), respectively. Here, H and W indicate the height and the width of the input, respectively.

single data points $x \in \mathbb{R}^{2\times H\times H}$, where H is the size of the image patch. KL grades are the outputs of our model: $y \in \{0, 1, 2, 3, 4\}$. f_θ denotes a Siamese neural network with parameters θ. In our setting, $p(y|\mathbf{x}) = f_\theta(\mathbf{x})$.

Our SSL method aims to perform penalization of local sharpness of the surface loss along the data manifold $\mathcal{M}$ and also within its surroundings. The former is achieved via minimization of $\mathbb{E}_x \|J_\mathcal{M}\|_F^2$, where $J_\mathcal{M}$ denotes the Jacobian along the data manifold $\mathcal{M}$ and $\|\cdot\|_F$ is the Frobenius norm, as well as $\mathbb{E}_x \|J_\theta\|_F^2$, where J_θ denotes the Jacobian in the parameter space, using consistency regularization [24], [45]. To generate out-of-manifold samples, we use *mixup* [34]. Here, we first enforce linear behavior of the model along the *mixup* rays via interpolation consistency training (ICT) [25], and then apply the aforementioned consistency regularization to enforce consistent behavior of the model along the *mixup* rays which are in the close surroundings of $\mathcal{M}$. As our idea for SSL centers around *mixup*, we name our method *Semixup*. We graphically illustrate the process of sample generation for consistency regularization in Figure 2.

B. Network Architecture

We follow the design of the previously developed Siamese model by Tiulpin *et al.* [20], and propose several modifications essential to obtain better KL grading performance in both SL and SSL settings. The schematic illustration of our proposed network is presented in Figure 3.

The basic building block of our model consists of a 3×3 convolution with a zero padding P and a stride S, an instance normalization (IN), and a leaky rectified linear unit (LeakyReLU) activation with a slope of 0.2.

To enrich the representation power of our model, we first start with three consecutive 3×3, $S = 1$, $P = 1$, and one 3×3, $S = 2$, $P = 1$ convolutional blocks (green blocks in Figure 3). Subsequently, we alternate 3×3, $S = 2$, $P = 1$ (yellow blocks in Figure 3), and 3×3, $S = 1$, $P = 1$ blocks until a feature map of size $\frac{1}{8}H \times \frac{1}{8}H$ is obtained.

We use stridden convolutions to perform the downsampling, whereas the original model from [20] used max-pooling layers. The main motivation for our method to use stridden convolutions is that the translation invariance achieved by the use of max-pooling could potentially harm the results by removing the dependencies between the OA-related fine-grained features at higher layers of the network [46].

Similar issues could also arise in the bottleneck of the network, where Tiulpin *et al.* [20] used Global Average Pooling (GAP). We argue that averaging such large feature maps is not the most optimal pooling strategy, and we tackle this problem via our novel SAM pooling (see Section III-C).

All the blocks described above share the weights among the branches of our Siamese CNN, where each branch of the model processes an individual side of the knee image pair (lateral or medial). Similar as GAP, SAM layer in each branch produces 1×1 features, which are concatenated, passed through a dropout, and subsequently fed into a fully-connected layer the OA severity stage. A detailed description of our architecture is provided in Suppl. Table S2.

C. Separable Adaptive Max-Pooling

As mentioned previously, we propose a replacement for GAP to deal with the potential information loss in the bottleneck of the model. The proposed SAM scheme is based on the idea of firstly applying pooling along one direction of the feature map (horizontal or vertical). Secondly, we use a 1×1 convolutional block (with IN and LeakyReLU) to remove unnecessary correlations between the pooled features. Finally, we apply the pooling in the direction orthogonal to the one in the first phase.

As the initial pooling of the features can be done in either horizontal or vertical directions, we suggest two configurations of SAM presented in Figure 4, namely:

1) Max-pooling $\frac{1}{8}H \times 1$ with stride $\frac{1}{8}H \times 1$, 1×1 convolution, IN, LeakyReLU with the slope of 0.2, and $1 \times \frac{1}{8}H$ max-pooling with the stride of 1×1 (SAM-VH).
2) Similar to the above, but the first and the last max-poolings are swapped (SAM-HV).

D. Semi-Supervised Learning

1) Problem Setting: Let $\mathcal{X}_l$ and $\mathcal{X}_u$ be labeled and unlabeled image sets, respectively. Let $\mathcal{Y}$ denote the labels for $\mathcal{X}_l$. In our setting, we optimize the following objective:

$$\min_\theta \mathcal{L}_l(\theta; \mathcal{X}_l, \mathcal{Y}) + \mathcal{L}_u(\theta; \mathbf{w}, \mathcal{X}_u, \mathcal{X}_l), \quad (1)$$

where $\mathcal{L}_l$ is a cross-entropy loss with *mixup*, $\mathcal{L}_u$ is a combination of losses without the involvement of labels, and $\mathbf{w}$ are hyperparameters responsible for weighing unsupervised losses. Here, $\mathcal{L}_u$ acts as a regularizer which leverages data without labels to enhance the robustness of the model f_θ via auxiliary tasks.

2) Supervised Loss: mixup proposed in [34] is a simple and effective technique to improve generalization. It can be viewed as data augmentation, and in a nutshell, it linearly mixes two samples x_i, x_j with a blending coefficient $\lambda \sim \text{Beta}(\alpha, \alpha)$, for $\alpha \in (0, \infty)$:

$$x_{mix} = \text{Mix}_\lambda(x_i, x_j) = \lambda x_i + (1-\lambda)x_j. \tag{2}$$

Having the mixed sample x_{mix}, the following loss is optimized:

$$\mathcal{L}_l(\theta; x_{mix}, y_i, y_j) = \lambda \mathcal{L}_{ce}(x_{mix}, y_i) + (1-\lambda)\mathcal{L}_{ce}(x_{mix}, y_j), \tag{3}$$

where $\mathcal{L}_{ce}$ is a multi-class cross-entropy loss. Here and further, we call the loss in (3) "soft" cross-entropy loss.

After the introduction of *mixup*, it was shown that this technique performs out-of-manifold regularization [47], and in our work, we exploit this property. Specifically, we view *mixup* as a data augmentation technique that generates out-of-manifold samples x_{mix} that belong to the ray between any two points x_i and x_j (Equation (2), and Figure 2b).

3) Consistency Regularization: Penalizing Local Loss Sharpness: The consistency regularization technique used in [24], [36] aims to minimize the following objective w.r.t θ:

$$\mathbb{E}_{x\sim p(x)}\mathbb{E}_{T,T'\sim p(\tau)} \left\| f_\theta(Tx) - f_\theta(T'x) \right\|_2^2, \tag{4}$$

where $p(x)$ is the distribution of training data (labeled and unlabeled, and $p(\tau)$ is the distribution of stochastic transformations (e.g., data augmentations).

According to Athiwaratkun *et al.* [45], regularizing consistency for a model using dropout in convolutional layers implies minimization of two terms: $\mathbb{E}_x \|J_\mathcal{M}\|_F^2$ and $\mathbb{E}_x \|J_\theta\|_F^2$. That idea results in an interesting connection between the consistency-based method and the classic graph-based approaches that use Laplacian regularization for SSL [48].

The explanations provided by Athiwaratkun *et al.* [45] demonstrate that minimization of $\mathbb{E}_x \|J_\theta\|_F^2$ leads to a broader optimum that is presumably helpful for good model generalization. Therefore, the minimization of the regularization term from (4) can lead to better performance which has been supported by experimental evidence in [24], [36], [45].

4) Interpolation Consistency: The *mixup* operator was introduced as an efficient data augmentation for supervised learning regularizing against out-of-manifold samples that lie close to $\mathcal{M}$. Recently, Verma *et al.* [25] utilized it to enforce linear behavior of the model along the *mixup* rays in the ICT method:

$$\mathbb{E}_{\lambda\sim\text{Beta}(\alpha,\alpha)}\mathbb{E}_{x_i,x_j\sim p(x)} \left\| \text{Mix}_\lambda(f_\theta(x_i), f_\theta(x_j)) - f_\theta(x_{mix}) \right\|_2^2. \tag{5}$$

While Verma *et al.* [25] did not consider *mixup* as an out-of-manifold regularizer, we consider this to be an important observation [47].

E. Semixup

1) Motivation: The consistency regularization and ICT aim to maximize the consistency of label assignment within and out of $\mathcal{M}$. However, we also note that while linear behavior of the model is achieved in ICT, it does not aim to minimize the inconsistency of label assignment for e.g. Tx and x_{mix} which can be viewed as in- and out-of-manifold augmented versions of a data point x. In this work, we address this limitation and strengthen the ability of making consistent predictions for out-of-manifold samples that are close to the data manifold $\mathcal{M}$. In summary, besides applying loss (3) over labeled data, *Semixup* optimizes a linear combination of the objectives shown in (4), and (5), as well as the out-of-manifold term described in the following section over both labeled and unlabeled samples. Supplementary Algorithm 1 shows a concrete implementation of our method.

2) Out-of-Manifold Consistency Regularization: We use the aforementioned regularizers from (4) and (5). Subsequently, given the motivation above, we minimize the following additional objective:

$$\mathbb{E}_{\lambda\sim\text{Beta}(\alpha,\alpha)}\mathbb{E}_{x,x'\sim p(x)}\mathbb{E}_{T\sim p(\tau)} \left\| f_\theta(\text{Mix}_\lambda(x, x')) - f_\theta(Tx) \right\|_2^2. \tag{6}$$

The presented objective aims to maximize the consistent label assignment for perturbed data items $Tx \in \mathcal{M}$, and also the ones being out-of-manifold but are close to it. To generate the latter, when sampling λ in *mixup*, we enforce $\lambda = \max(\lambda, 1-\lambda)$.

3) Low-Variance Sampling: Our method uses both labeled and unlabeled data in the unsupervised regularization term of the loss (1). It is motivated by the fact that unlabeled data in medical imaging can come from different device vendors rather than labeled data, thereby the empirical distributions of labeled and unlabeled data might be misaligned.

We note that the objective in (6) uses stochastic augmentations $T \sim p(\tau)$. When using this loss in a combination with consistency regularization and ICT, it can be seen that there exist predictions for two versions of an image x as in (4). Therefore, when optimizing with stochastic gradient descent, we are able to obtain a lower variance stochastic estimate of the objective (6) at low marginal computational cost (see lines 10, 11 of Supplementary Algorithm 1).

IV. Experiments

A. Datasets

We used knee radiographs from two large public cohorts: The OAI and the MOST. The OAI dataset was collected from 4,796 participants whose ages were from 45 to 79 years old. The cohort included a baseline, and follow-up visits after 12, 18, 24, 30, 36, 48, 60, 72, 84, 96, 108, 120, and 132 months. Radiographic imaging for bilateral fixed flexion X-ray images took place at most, but not all follow-ups. We used the data from all the knees without implants that had KL grades available in the metadata except for those imaged during 18, 30, 48, 60, 84, 108, 120, and 132-month follow-ups, where knee radiograph imaging was performed only for small sub-cohorts. We graphically depict the OAI data selection in

TABLE I
DESCRIPTION OF THE OAI AND THE MOST DATASETS

Dataset	Split	# Images	KL grade				
			0	1	2	3	4
OAI	Train/Val	39,902	15,954	7,636	9,617	5,228	1,467
MOST	Test	3,445	1,550	568	520	559	248

Suppl. Figure S4. The MOST dataset had 3,021 participants examined at its baseline, and follow-up visits after 15, 30, 60, 72, and 84 months. Except for the 72-month follow-up, each examination included radiographic imaging.

Radiographs in the OAI dataset were posterior-anterior (PA) bilateral images acquired with the protocol that called for a beam angle of 10 degrees. We utilized OAI data in training and model selection phases.

The MOST dataset included PA, and lateral images. The PA images were acquired with a beam angle of 5, 10, or 15 degrees. In this study, we used the MOST dataset as an independent test set. To ensure the reliability of the labels in MOST, we only kept 10 degree bilateral PA images that were acquired at baseline, had KL and Osteoarthritis Research Society International (OARSI) grades available, and were without implants. Eventually, our training and test data comprised 39,902 and 3,445 knee images from the OAI, and the MOST datasets, respectively. The detailed data distribution by KL grade is presented in Table I.

B. Experimental Setup

1) Data Pre-Processing: We kept our data pre-processing pipeline similar to [20], and consistently applied it to all the methods in our experiments.

The essential step in analyzing knee radiographs is a region of interest (ROI) localization. In this paper, we focused rather on the development of an efficient DL architecture and SSL method than on implementing a full knee X-ray analysis pipeline. Thus, we utilized a Random Forest Regression Voting Constrained Local Model approach implemented in the BoneFinder (BF) tool [49]. For every image, we performed the padding of 300 pixels on each side to account for position precision in the raw data. After our manual inspection of all the cases, we observed that BF localized all the knees in OAI and MOST except two images (image files could not be read by the software). We show data selection in more detail in Suppl. Figure S4.

In each bilateral X-ray, we localized the anatomical landmarks using BF and cropped the ROIs of $140mm \times 140mm$ centered at each knee joint (2 ROIs per image at most). Subsequently, based on the anatomical landmarks, we performed the standardization of each of the ROIs by a horizontal alignment of the tibial plateau.

To standardize the intensity of the histograms, we performed intensity clipping to the 5^{th} and 99^{th} percentiles as well as global contrast normalization prior to converting 16-bit raw DICOM images into an 8-bit intensity range. We then center-cropped the obtained 8-bit images to $110mm \times 110\ mm$ and resized them to 300×300 pixels ($0.37\ mm$ pixel spacing). This second cropping aimed to eliminate more background from input images. We note that it is also feasible to perform the histogram standardization after the second cropping; however, we followed the order proposed in [20] for better reproducibility of their pipeline, which was retrained in the present study.

At the next step of the pre-processing, similar to [20], we flipped the left images to look like the right ones, and cropped the medial and the lateral patches with a size of 128×128 pixels ($H = 128$). These patches were cropped with the common top anchor at one third of the image height. Both patches were cropped at the lateral and medial image sides. After cropping, we flipped the medial patch horizontally (Figure 3). Finally, we normalized the intensities of each item in the obtained pair into the intensity range of $[-1, 1]$. Whereas the previous studies [20], [40] required the statistics of their training sets to normalize input data, we utilized the mean of 0.5 and the standard deviation of 0.5.

2) Training and Evaluation Protocol:

a) Data split: First, we split the OAI dataset into 2 parts so that 25% and 75% were for labeled and unlabeled data, respectively. Here, we stratified the splits by KL grade, and ensured that patient IDs do not belong to both of these parts. We then applied another stratification as above to divide each of the aforementioned splits into 5 folds (illustrated in Suppl. Figure S4). In each fold, we generated 24 data settings, having 4 labeled (50, 100, 500, and 1000 samples per KL grade), and 6 unlabeled data configurations. Here, the amount of unlabeled data was N, $2N$, $3N$, $4N$, $5N$, or $6N$ samples, where N is the corresponding amount of labeled data.

In addition, we also used the whole OAI dataset to assess the performance of our SL baseline. The training and the validation sets of the setting had 31,922 and 7,980 samples, respectively. Therefore, the 4 aforementioned labeled data settings varied from 0.8% to 16% of the whole training set.

b) Architecture selection and training setup: Following the protocol proposed by Oliver *et al.* [35], we firstly tuned the SL setting before comparing it to SSL methods. As such, we considered 6 feasible architectures, each of which was the combination of either the architecture of the SL baseline [20], or ours with each type of pooling layer (GAP, SAM-VH, or SAM-HV). In our experiments, we selected the top-3 architectures for further comparisons based on cross-validation. The best model among those was utilized as the base model of *Semixup* and SSL baselines.

c) SL and SSL comparisons: We investigated effects of SSL methods such as the Π-model [24], MixMatch [26], and *Semixup* without the use of unlabeled data. In such scenario, MixMatch [26] is equivalent to *mixup* [34]. Ultimately, we compared *Semixup* to 3 SSL baselines such as the Π model [24], Interpolation Consistency Training (ICT) [25], and MixMatch [26]. Each method was trained on the 24 data settings of the first fold previously generated from the OAI, and finally evaluated on the independent test set from the MOST.

d) Metrics: In the training phase of all the experiments, the best models were selected based on both balanced multi-class accuracy (BA) [50] and Cohen's quadratic kappa

TABLE II
ABLATION STUDY OF THE SAM POOLING (SL SETTING). THE VALUES INDICATE BALANCED ACCURACIES (%). THE RESULTS WERE COMPUTED OUT-OF-FOLD USING A 5-FOLD CROSS-VALIDATION. THE BOLD AND THE UNDERLINE HIGHLIGHT THE BEST AND THE SECOND BEST RESULTS, RESPECTIVELY.

Base model	Pooling	# data / KL grade				
		50	100	500	1000	Average
Tiulpin *et al.* [20]	GAP	43.7	52.5	62.1	66.8	56.3
	SAM-VH	**54.0**	50.7	62.5	<u>69.1</u>	<u>59.1</u>
	SAM-HV	39.0	51.2	63.7	62.7	54.1
Ours (Sec. III-B)	GAP	46.7	49.6	**67.0**	68.6	58.0
	SAM-VH	47.7	<u>54.6</u>	65.8	65.5	58.4
	SAM-HV	<u>48.3</u>	**56.3**	<u>66.7</u>	**69.7**	**60.3**

TABLE III
SL AND SSL METHODS EVALUATION ON THE TEST SET. WE REPORTED THE TOP-3 SL MODELS WITH GAP*, SAM-VH†, OR SAM-HV‡. OUR MODEL WITH SAM-HV‡ IS THE COMMON ARCHITECTURE FOR ALL SSL METHODS. BOLD HIGHLIGHTS MODELS PERFORMING SUBSTANTIALLY BETTER THAN THE SECOND BEST MODEL IN EACH CATEGORY. $\diamond$ INDICATES NO SUBSTANTIAL DIFFERENCE AMONG THE BEST MODELS

Method	# labels / KL grade			
	50	100	500	1000
Fully SL				
Tiulpin *et al.** [20]	40.5±0.8	49.7±0.9	55.1±0.8	58.5±0.8
Tiulpin *et al.*† [20]	45.2±0.8$^{\diamond}$	48.6±0.9	57.6±0.8	58.5±0.8
Our SL† (Sec. III-B)	45.6±0.8$^{\diamond}$	53.2±0.9$^{\diamond}$	61.5±0.8	63.5±0.8
Our SL‡ (Sec. III-B)	41.5±0.8	52.9±0.9$^{\diamond}$	**64.0±0.8**	**67.5±0.8**
SL and SSL methods - Without unlabeled data				
mixup [34]	39.6±0.8	53.9±0.8	64.5±0.8	67.1±0.8$^{\diamond}$
Π model [24]	**42.3±0.8**	56.3±0.8$^{\diamond}$	64.3±0.8	67.9±0.8$^{\diamond}$
Semixup (Ours)	31.7±0.8	56.1±0.8$^{\diamond}$	**66.6±0.8**	68.0±0.8$^{\diamond}$
SSL methods - Best performing models' comparison				
ICT [25]	47.7±0.9$^{\diamond}$	53.5±0.8	64.1±0.8	67.8±0.8
Π model [24]	45.9±0.8$^{\diamond}$	56.2±0.8	67.1±0.8	69.0±0.8
MixMatch [26]	45.1±0.8$^{\diamond}$	56.0±0.8	67.6±0.8	68.4±0.8
Semixup (Ours)	46.9±0.9$^{\diamond}$	**58.8±0.8**	**69.7±0.8**	**71.0±0.8**
Fully SL - Trained on the full OAI (31,922 samples)				
Our SL‡ (Sec. III-B)	70.9±0.8			

coefficient (KC) [51]. In the final model evaluation and comparison to the baseline methods, we used BA. To be in line with the metrics used in the previous studies, such as [20], [23], [40], we also used confusion matrix, KC, and mean square error (MSE). To assess the performance of detecting radiographic OA (KL $\geq$ 2), we used receiver operating characteristic (ROC) curves, area under the ROC curve (AUC), precision-recall (PR) curves, and average precision (AP).

e) Statistical analyses: In our initial experiments, we noticed that several factors such as data acquisition center, knee side (left or right), and subject ID may affect the validity of the statistical analyses. To verify this, we used a generalized linear mixed effects model [52], and noticed that neither of these factors has an impact on the results. Therefore, for simplicity of evaluation, we used standard error and one-sided Wilcoxon signed-rank test [53]. Here, we split the test into 20 equal-sized chunks (no overlapping patients), and calculated the BA on each of them. Finally, these values were used for the statistical analyses.

C. Supervised Baseline

1) Model Selection: Table II shows that, on average, the models with our base architecture performed 4.1% better compared to the models with [20] in the case of 500 samples per KL grade. With respect to the base architecture of [20], SAM-VH was the most suitable pooling operator, especially in the cases of 50 and 1000 samples per KL grade. On the other hand, our base architecture with SAM-HV yielded the highest average BA across all settings.

Based on the average BAs in Table II, we selected the three best architectures for further comparisons. The best one among those (ours with SAM-HV) was chosen as the base model of all the SSL methods.

2) Performance on the Test Set: In the first group of results, Table III shows that our architecture with SAM outperformed the SL baseline model [20] in all the data settings. Specifically, our architecture with SAM-HV yielded BAs 9% better than the baseline [20] in the data settings of 500 and 1000 images per KL grade. Notably, in the case of 500 samples per KL grade, our model with SAM-HV surpassed by 6% the baseline model that would require a double amount of data. Finally, our model with SAM-HV achieved a BA of 67.5% when trained on only 16% of the full OAI set.

D. Semi-Supervised Learning

1) SSL Methods Without Unlabeled Data: We summarized our results in Table III. The performance of *Semixup* scaled well with the amount of training data, and outperformed its corresponding fully SL model when we had at least 100 samples per KL grade. In the comparison with other SSL baselines, *Semixup* achieved comparable results when having no unlabeled data at all.

2) Semixup Compared to SL and SSL Baselines: After evaluating 24 trained models of each SSL method, we derived maximum BAs grouped by labeled data setting as in the third part of Table III. With 500 labeled samples per KL grade, *Semixup* achieved a BA of 69.7%, which was 5.7% higher than the result of the SL baseline, and surpassed the SL baseline by 2.2% even when trained on 2 times less labeled data than the baseline was. With 1000 labeled samples per KL grade, our SSL method reached the BA of 71% that was comparable to the result of the SL baseline trained on the full OAI setting (over 6 times more data used).

Besides comparing to the SL settings, we also compared *Semixup* to the state-of-the-art SSL baselines. On the test set, our method outperformed the others in the data settings of 100, 500, and 1000 labels per KL grade. In the data setting of 2500 labeled images (7.8% of the full OAI size), *Semixup* yielded a

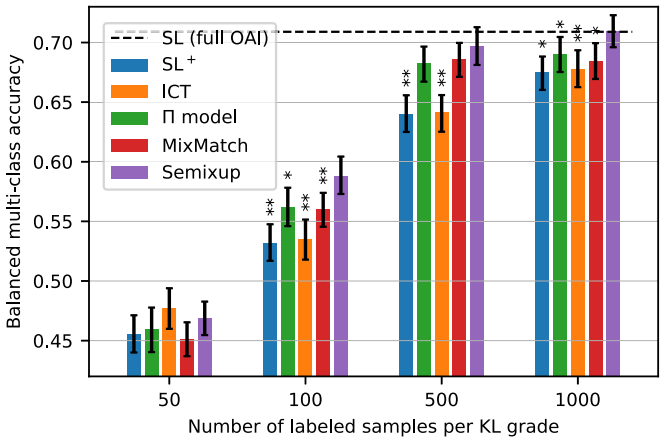

Fig. 5. Graphical comparison of *Semixup* and the baseline methods (MOST dataset). The bars indicate the 95% confidence intervals, SL$^+$ indicates our fully SL models equipped with either SAM-HV or SAM-VH, and * and ** indicate the statistically significant difference ($p < 0.05$ and $p < 0.001$, respectively). The dash line indicates the BA of the fully SL model with SAM-HV trained on the full OAI dataset.

BA 2.1% higher than the best SSL baseline, MixMatch, with the same labeled data amount, and 0.7% higher than the best SSL baseline, the Π model, with twice the amount of labeled data. In addition, *Semixup* achieved the highest average BAs in the data settings having at least 100 labeled images per KL grade (Suppl. Table S3).

The comparisons between *Semixup* and each of the baseline models across different labeled data settings are graphically illustrated in Figure 5. The best models of *Semixup* in the cases of 100 and 1000 labels per KL grade significantly outperformed all the baselines ($p < 0.05$). In the case of 500 labeled samples per KL grade, *Semixup* was significantly better than the SL baseline and ICT with the same amount of labels. Furthermore, our method also surpassed the SL baseline that was trained with twice more labeled data (1000 samples per KL grade). The statistical analyses indicate that SL model trained on the full OAI dataset did not differ significantly from the *Semixup* models that were trained with 500 and 1000 labeled samples per KL grade (p-values were of 0.054 and 0.368, respectively). The details of the statistical testing are presented in Suppl. Table S4.

We finally note here that we also evaluated larger amounts of labeled data. However, we observed a saturation of performance increase with their further addition to the training. Therefore, we omitted these results from the manuscript, but they can be found in Supplement (Suppl. Section C).

3) Detection of Early Radiographic Osteoarthritis: Figure 6 presents the confusion matrices of our 2 best models. Here, we evaluate the accuracy of our method to detect early OA (i.e., KL = 2). With 500 and 1000 labels per KL grade, *Semixup* was able to detect early OA with 58% and 74% accuracies. Notably, without the doubtful OA (KL = 1), we achieved a substantially high BA of 79.25%.

With regard to detection of radiographic OA (KL $\geq$ 2), Figure 7 and Table IV show how our SSL model with 500 labels per KL grade was comparable to the well-tuned SL model trained on the large training set (more than 12 times

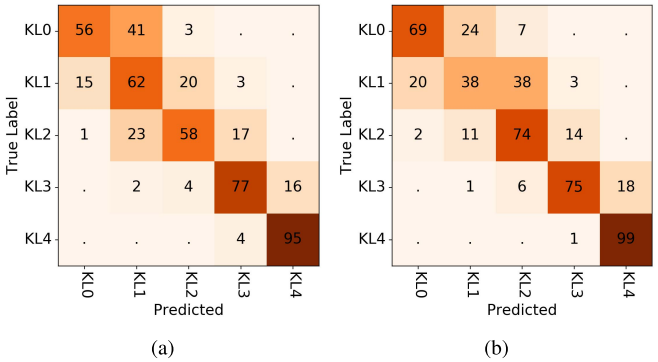

Fig. 6. Confusion matrices showing performance (%) of our models trained with *Semixup* on the test set (MOST dataset) in two labeled data settings (a) and (b).

TABLE IV
COMPARISON OF OUR BEST MODELS TRAINED WITH *Semixup* AGAINST OUR WELL-TUNED SL MODEL WITH SAM-HV

Method	# labels	KC	MSE	AUC (KL $\geq$ 2)	AP (KL $\geq$ 2)
Semixup	250	0.708 [0.692, 0.724]	1.080 [1.000, 1.022]	0.880 [0.869, 0.891]	0.861 [0.849, 0.872]
	500	0.789 [0.776, 0.801]	0.810 [0.764, 0.860]	0.933 [0.926, 0.940]	0.916 [0.906, 0.925]
	2,500	0.877 [0.870, 0.884]	0.442 [0.416, 0.466]	0.967 [0.962, 0.972]	0.956 [0.950, 0.962]
	5,000	0.878 [0.870, 0.885]	0.458 [0.430, 0.487]	0.972 [0.968, 0.976]	0.959 [0.953, 0.965]
SL (full OAI)	31,922	0.881 [0.873, 0.889]	0.440 [0.414, 0.471]	0.974 [0.969, 0.978]	0.963 [0.958, 0.969]

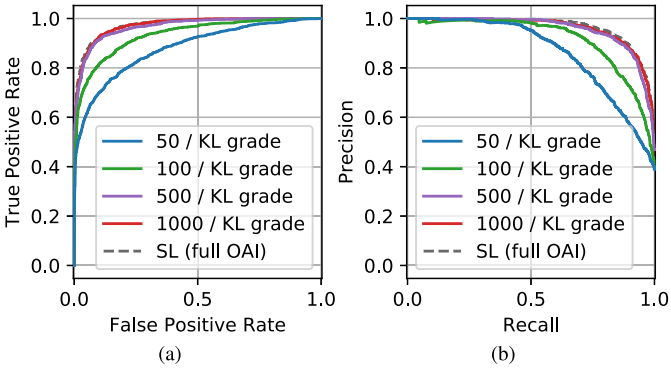

Fig. 7. Comparison of the best models of trained with *Semixup* on radiographic OA detection task (KL $\geq$ 2). The models were trained with 50, 100, 500, or 1000 labeled examples per KL grade.

labeled samples). The detailed comparisons with respect to different amounts of unlabeled data are presented in Suppl. Figure S1 and Suppl. Figure S2.

E. Ablation Study

1) Impact of Consistency Regularization Terms: Because our in- and out-of-manifold regularizer comprises several individual losses, it is essential to understand the impact of each of them onto the method's performance. To assess the contributions of each used regularizer – (4), (5), and (6), we consecutively removed each of them from our loss function. Table V shows this ablation study. Here, we report the best validation performance among 6 different unlabeled data settings. The results show that *Semixup* performs better when all the regularizers are used jointly.

TABLE V
ABLATION STUDY OF REGULARIZATION TERMS USED IN *Semixup*. THE VALUES ARE BAS (%)

Method	# data / KL grade	
	100	500
Semixup w/o in-manifold terms	60.8	73.1
Semixup w/o out-of-manifold term	60.8	72.0
Semixup w/o interpolation consistency	60.8	73.4
Semixup	**65.6**	**74.1**

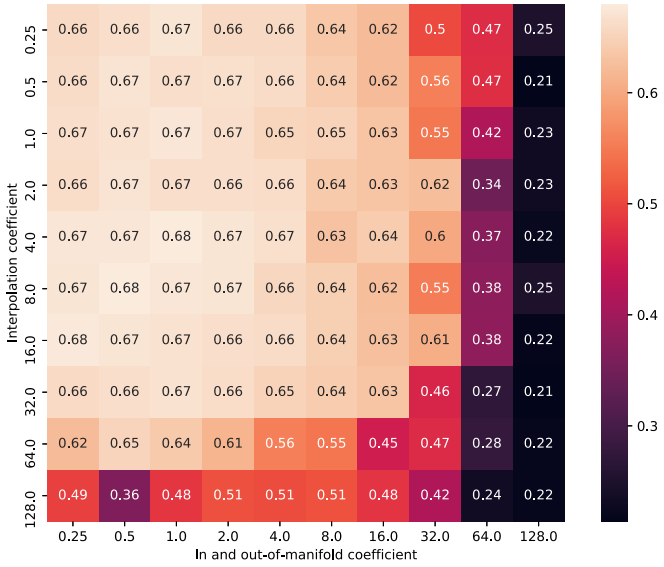

Fig. 8. Sensitivity analysis of balancing interpolation consistency term's coefficient versus the in- and out-of-manifold ones. We report the average balanced accuracy on validation set over 5 random seeds for the case of 2500 labeled, and 2500 unlabeled samples. The weights of in- and out-of-manifold terms were set to the same value.

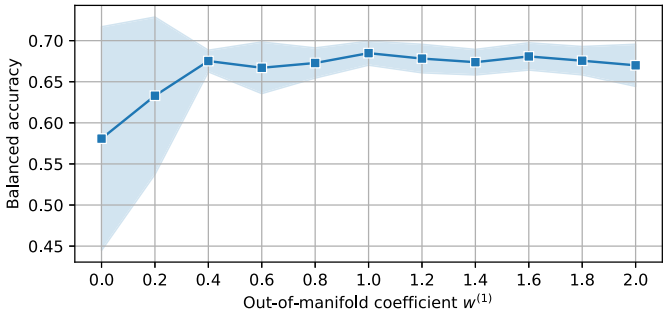

Fig. 9. Relationship between in- and out-of manifold terms. Interpolation consistency term's coefficient was set to 4, and the sum of in- and out-of-manifold consistency terms' weights was set to 2. We evaluated the average balanced accuracy on the validation set over 5 different random seeds for 2500 labeled, and 2500 unlabeled samples. The solid line and the band indicate the mean and standard deviation, respectively.

2) Balancing of Regularization Coefficients: Table V shows that all the consistency regularization terms used in our method in are needed to reach good performance. However, it is unclear what practical effect the balancing of these terms may have on the performance of our method. To gain insights into what value ranges of these coefficients are appropriate, we performed additional sensitivity analyses of these hyperparameters by varying the values used in our experiments.

Firstly, we showed the trade-off between the interpolation consistency term and the rest (Figure 8). Secondly, we showed the effect of balancing in- and out-of-manifold consistency regularization terms (Figure 9). All the experiments in this subsection demonstrate the performance on the validation set in the case of 2500 labeled samples (500 per KL-grade), and the same amount of unlabeled samples.

V. DISCUSSION AND CONCLUSIONS

In this study, we presented a novel SSL method – *Semixup*. This method leverages in- and out-of-manifold consistency regularization, and we demonstrated its application in the task of automatic grading of knee OA severity. Furthermore, we also proposed a novel state-of-the-art architecture for this task.

The core novelty of this work lies in using *mixup* for generating out-of-manifold samples in close proximity of data manifold, and then ensuring consistent predictions made by a neural network on such samples and the ones drawn from the data manifold itself. We experimentally showed that the proposed technique helps to train more robust models in limited data regimes even when unlabeled data is not available.

To be specific about the label efficiency and robustness, we have shown that the *Semixup*-trained model using 5000 labeled samples (1000 per KL-grade) performs comparably to the well-tuned supervised baseline trained on the full OAI dataset (Suppl. Table S4, $p = 0.368$). Furthermore, Table III and Table IV show that our SSL model has overlapping confidence intervals with the SL model according to various metrics, which also supports this statement.

Another novelty of this work is the proposed pooling approach that allows to efficiently process large feature maps, thereby leveraging fine-grained information required for knee OA grading. Our work demonstrated not only the new method for SSL, but also an efficient baseline for supervised setting when the amount of training samples is limited. To our knowledge, this is the first work that systematically studied the problem of learning robust classifiers for knee OA severity assessment in the limited data regime.

We note that while we leveraged the domain knowledge to develop a base CNN architecture (Siamese network), our SSL method is not constrained to any specific domain. Thus, we think that *Semixup* can be used with other data (e.g., in CT or MRI related tasks). However, we also highlight that this requires an experimental evidence, producing which is beyond the scope of our study. To facilitate the applications of our method in other fields, we release the source code at https://github.com/MIPT-Oulu/semixup.

From the methodological point of view, our work has several important insights. Firstly, the results shown in Suppl. Table S3 indicate that none of the evaluated SSL methods consistently scale in performance with scaling the amounts of unlabeled data. Thus, we suggest that it is sufficient

to use as many unlabeled samples. The experiments shown in this work (Figure 9) show that the out-of-manifold consistency regularization is rather more important than the in-manifold one. Here, we point to the recent work by Mao *et al.* [54] in the realm of domain adaptation, which showed that using consistency regularization of mixed-up samples imposes a locally-Lipschitz constraint to the model in the regions of close proximity to the data manifold. To conclude, we believe that further exploiting this type of consistency regularization can help in the other domains.

Besides all the aforementioned benefits, this study has several limitations. Specifically, we did not use EMA in our methods as it is costly to train and the methods using EMA would require significantly more computational power than their competitors. Another limitation of this study is that we did not leverage the power of SAM-VH and SAM-HV in our model. We foresee potential improvements of our results if both of these pooling schemes are used, e.g. by ensembling the results produced by models trained with SAM-VH and SAM-HV, respectively.

A general criticism also applies to our work. While considering SSL to be a powerful tool, we also confirm the concerns previously raised Oliver *et al.* [35]. Specifically, we emphasize that a thorough evaluation of SSL vs SL should always be done, and one might always reach the best possible performance when using the former one. Therefore, we think that a further research on the stability of SSL methods is needed. However, our results show that SSL was still outperforming SL, especially in the low-data regime.

From a clinical point of view, using composite KL grades as a reference can also be considered a limitation, thus we suggest future studies to focus on OARSI grading of the knee joints [44], [55]. Finally, another clinical limitation of this work is that OA is typically graded at a late stage, when it is already present, and we think that the next steps should rather focus on developing models for predicting OA progression [56] using SSL and partially labeled data. These data are available in hospital archives and can be leveraged at a low cost.

To conclude, we would like to highlight the clinical implications of our work. Firstly, we demonstrated that highly accurate KL grading can be done with only small amounts of labeled data, which allows small research teams and medical device vendors to build generalizable models suitable for clinical use. Secondly, the proposed method can significantly reduce routine work done by radiologists, which on the societal level can lead to cost savings while at the same time improving the quality of health care. Thirdly, we expect the proposed method and the network architecture to generalize to other domains, such as hip and hand OA. Finally, we think that this study is an important step towards data efficient medical image recognition, which is currently lacking thoroughly validated methodologies.

ACKNOWLEDGMENT

Private sector funding for the OAI is managed by the Foundation for the National Institutes of Health. This manuscript was prepared using MOST data and does not necessarily reflect the opinions or views of MOST investigators.

Dr. Claudia Lindner is acknowledged for providing BoneFinder. Iaroslav Melekhov is acknowledged for proposing the ablation study of the pooling method. Phuoc Dat Nguyen is acknowledged for discussions about *mixup*.

REFERENCES

[1] D. J. Hunter and S. Bierma-Zeinstra, "Osteoarthritis," *Lancet*, vol. 393, no. 10182, pp. 1745–1759, Apr. 2019.
[2] P. A. Dieppe and L. S. Lohmander, "Pathogenesis and management of pain in osteoarthritis," *Lancet*, vol. 365, no. 9463, pp. 965–973, Mar. 2005.
[3] A. Mobasheri and M. Batt, "An update on the pathophysiology of osteoarthritis," *Ann. Phys. Rehabil. Med.*, vol. 59, nos. 5–6, pp. 333–339, Dec. 2016.
[4] A. Mathiessen and P. G. Conaghan, "Synovitis in osteoarthritis: Current understanding with therapeutic implications," *Arthritis Res. Therapy*, vol. 19, no. 1, p. 18, Dec. 2017.
[5] C. Palazzo, C. Nguyen, M.-M. Lefevre-Colau, F. Rannou, and S. Poiraudeau, "Risk factors and burden of osteoarthritis," *Ann. Phys. Rehabil. Med.*, vol. 59, no. 3, pp. 134–138, Jun. 2016.
[6] B. S. Ferket, Z. Feldman, J. Zhou, E. H. Oei, S. M. A. Bierma-Zeinstra, and M. Mazumdar, "Impact of total knee replacement practice: Cost effectiveness analysis of data from the osteoarthritis initiative," *BMJ*, vol. 356, p. j1131, Mar. 2017.
[7] M. Cross *et al.*, "The global burden of hip and knee osteoarthritis: Estimates from the global burden of disease 2010 study," *Ann. Rheumatic Diseases*, vol. 73, no. 7, pp. 1323–1330, Jul. 2014.
[8] S. Salih and A. Hamer, "Hip and knee replacement," *Surgery*, vol. 31, no. 9, pp. 482–487, 2013.
[9] A. A. Guccione *et al.*, "The effects of specific medical conditions on the functional limitations of elders in the framingham Study.," *Amer. J. Public Health*, vol. 84, no. 3, pp. 351–358, Mar. 1994.
[10] K. N. Murtagh and H. B. Hubert, "Gender differences in physical disability among an elderly cohort," *Amer. J. Public Health*, vol. 94, no. 8, pp. 1406–1411, Aug. 2004.
[11] S. Glyn-Jones *et al.*, "Osteoarthritis," *Lancet*, vol. 386, no. 9991, pp. 376–387, 2015.
[12] A. T. Toivanen *et al.*, "Obesity, physically demanding work and traumatic knee injury are major risk factors for knee osteoarthritis–A population-based study with a follow-up of 22 years," *Rheumatology*, vol. 49, no. 2, pp. 308–314, Feb. 2010.
[13] M. E. Miller, W. J. Rejeski, S. P. Messier, and R. F. Loeser, "Modifiers of change in physical functioning in older adults with knee pain: The observational arthritis study in seniors (OASIS)," *Arthritis Rheumatism*, vol. 45, no. 4, pp. 331–339, 2001.
[14] G. Leardini *et al.*, "Direct and indirect costs of osteoarthritis of the knee," *Clin. Express Rheumatol.*, vol. 22, no. 6, pp. 699–706, 2004.
[15] A. J. Carr *et al.*, "Knee replacement," *Lancet*, vol. 379, no. 9823, pp. 1331–1340, 2012.
[16] A. J. Price *et al.*, "Knee replacement," *The Lancet*, vol. 392, no. 10158, pp. 1672–1682, 2018.
[17] P. Baker, K. Muthumayandi, C. Gerrand, B. Kleim, K. Bettinson, and D. Deehan, "Influence of body mass index (BMI) on functional improvements at 3 years following total knee replacement: A retrospective cohort study," *PLoS ONE*, vol. 8, no. 3, Mar. 2013, Art. no. e59079.
[18] K. Baker and T. McAlindon, "Exercise for knee osteoarthritis," *Current Opinion Rheumatol.*, vol. 12, no. 5, pp. 456–463, 2000.
[19] S. Reichenbach *et al.*, "Prevalence of bone attrition on knee radiographs and MRI in a community-based cohort," *Osteoarthritis Cartilage*, vol. 16, no. 9, pp. 1005–1010, Sep. 2008.
[20] A. Tiulpin, J. Thevenot, E. Rahtu, P. Lehenkari, and S. Saarakkala, "Automatic knee osteoarthritis diagnosis from plain radiographs: A deep learning-based approach," *Sci. Rep.*, vol. 8, no. 1, p. 1727, Dec. 2018.
[21] J. Buckland-Wright, I. Carmichael, and S. Walker, "Quantitative microfocal radiography accurately detects joint changes in rheumatoid arthritis," *Ann. Rheumatic Diseases*, vol. 45, no. 5, pp. 379–383, 1986.
[22] J. Antony, K. McGuinness, K. Moran, and N. E. O'Connor, "Automatic detection of knee joints and quantification of knee osteoarthritis severity using convolutional neural networks," in *Proc. Int. Conf. Mach. Learn. Data Mining Pattern Recognit.* Cham, Switzerland: Springer, 2017, pp. 376–390.

[23] A. J. Antony, "Automatic quantification of radiographic knee osteoarthritis severity and associated diagnostic features using deep convolutional neural networks," Ph.D. dissertation, School Electron. Eng., Dublin City Univ., Dublin, Ireland, 2018.

[24] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," 2016, *arXiv:1610.02242*. [Online]. Available: http://arxiv.org/abs/1610.02242

[25] V. Verma, A. Lamb, J. Kannala, Y. Bengio, and D. Lopez-Paz, "Interpolation consistency training for semi-supervised learning," 2019, *arXiv:1903.03825*. [Online]. Available: http://arxiv.org/abs/1903.03825

[26] D. Berthelot, N. Carlini, I. Goodfellow, N. Papernot, A. Oliver, and C. Raffel, "MixMatch: A holistic approach to semi-supervised learning," 2019, *arXiv:1905.02249*. [Online]. Available: http://arxiv.org/abs/1905.02249

[27] H. Su, Z. Yin, S. Huh, T. Kanade, and J. Zhu, "Interactive cell segmentation based on active and semi-supervised learning," *IEEE Trans. Med. Imag.*, vol. 35, no. 3, pp. 762–777, Mar. 2016.

[28] B. Wang *et al.*, "4D active cut: An interactive tool for pathological anatomy modeling," in *Proc. IEEE 11th Int. Symp. Biomed. Imag. (ISBI)*, Apr. 2014, pp. 529–532.

[29] J. E. Iglesias, C.-Y. Liu, P. Thompson, and Z. Tu, "Agreement-based semi-supervised learning for skull stripping," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, T. Jiang, N. Navab, J. P. W. Pluim, and M. A. Viergever, Eds. Berlin, Germany: Springer, 2010, pp. 147–154.

[30] N. Kumar *et al.*, "Hyperspectral tissue image segmentation using semi-supervised NMF and hierarchical clustering," *IEEE Trans. Med. Imag.*, vol. 38, no. 5, pp. 1304–1313, May 2019.

[31] X. Li, L. Yu, H. Chen, C.-W. Fu, and P.-A. Heng, "Semi-supervised skin lesion segmentation via transformation consistent self-ensembling model," 2018, *arXiv:1808.03887*. [Online]. Available: http://arxiv.org/abs/1808.03887

[32] I. Goodfellow *et al.*, "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst.*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. New York, NY, USA: Curran Associates, 2014, pp. 2672–2680.

[33] M. Belkin, P. Niyogi, and V. Sindhwani, "Manifold regularization: A geometric framework for learning from labeled and unlabeled examples," *J. Mach. Learn. Res.*, vol. 7, pp. 2399–2434, Nov. 2006.

[34] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "Mixup: Beyond empirical risk minimization," 2017, *arXiv:1710.09412*. [Online]. Available: http://arxiv.org/abs/1710.09412

[35] A. Oliver, A. Odena, C. A. Raffel, E. D. Cubuk, and I. Goodfellow, "Realistic evaluation of deep semi-supervised learning algorithms," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 3239–3250.

[36] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 2017, pp. 1195–1204.

[37] X. Zhu and Z. Ghahramani, "Learning from labeled and unlabeled data with label propagation," Carnegie Mellon Univ., Pittsburgh, PA, USA, Tech. Rep., 2002.

[38] D.-H. Lee, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Proc. Workshop Challenges Represent. Learn. (ICML)*, vol. 3, 2013, p. 2.

[39] O. Chapelle and A. Zien, "Semi-supervised classification by low density separation," in *Proc. AISTATS*, 2005, pp. 57–64.

[40] J. Antony, K. McGuinness, N. E. O'Connor, and K. Moran, "Quantifying radiographic knee osteoarthritis severity using deep convolutional neural networks," in *Proc. 23rd Int. Conf. Pattern Recognit. (ICPR)*, Dec. 2016, pp. 1195–1200.

[41] P. Chen, L. Gao, X. Shi, K. Allen, and L. Yang, "Fully automatic knee osteoarthritis severity grading using deep neural networks with a novel ordinal loss," *Comput. Med. Imag. Graph.*, vol. 75, pp. 84–92, Jul. 2019.

[42] B. Norman, V. Pedoia, A. Noworolski, T. M. Link, and S. Majumdar, "Applying densely connected convolutional neural networks for staging osteoarthritis severity from plain radiographs," *J. Digit. Imag.*, vol. 32, no. 3, pp. 471–477, Jun. 2019.

[43] M. Górriz, J. Antony, K. McGuinness, X. Giró-i-Nieto, and N. E. O'Connor, "Assessing knee OA severity with CNN attention-based end-to-end architectures," 2019, *arXiv:1908.08856*. [Online]. Available: http://arxiv.org/abs/1908.08856

[44] A. Tiulpin and S. Saarakkala, "Automatic grading of individual knee osteoarthritis features in plain radiographs using deep convolutional neural networks," 2019, *arXiv:1907.08020*. [Online]. Available: http://arxiv.org/abs/1907.08020

[45] B. Athiwaratkun, M. Finzi, P. Izmailov, and A. Gordon Wilson, "There are many consistent explanations of unlabeled data: Why you should average," 2018, *arXiv:1806.05594*. [Online]. Available: http://arxiv.org/abs/1806.05594

[46] J. Tobias Springenberg, A. Dosovitskiy, T. Brox, and M. Riedmiller, "Striving for simplicity: The all convolutional net," 2014, *arXiv:1412.6806*. [Online]. Available: http://arxiv.org/abs/1412.6806

[47] H. Guo, Y. Mao, and R. Zhang, "Mixup as locally linear out-of-manifold regularization," in *Proc. AAAI Conf. Artif. Intell.*, vol. 33, 2019, pp. 3714–3722.

[48] O. Chapelle, B. Scholkopf, and A. Zien, "Semi-supervised learning (Chapelle, O. et al., Eds.; 2006) [Book reviews]," *IEEE Trans. Neural Netw.*, vol. 20, no. 3, p. 542, Mar. 2009.

[49] C. Lindner, S. Thiagarajah, J. Wilkinson, T. Consortium, G. Wallis, and T. Cootes, "Fully automatic segmentation of the proximal femur using random forest regression voting," *IEEE Trans. Med. Imag.*, vol. 32, no. 8, pp. 1462–1472, Aug. 2013.

[50] R. J. Urbanowicz and J. H. Moore, "ExSTraCS 2.0: Description and evaluation of a scalable learning classifier system," *Evol. Intell.*, vol. 8, nos. 2–3, pp. 89–116, Sep. 2015.

[51] J. Cohen, "A coefficient of agreement for nominal scales," *Educ. Psychol. Meas.*, vol. 20, no. 1, pp. 37–46, Apr. 1960.

[52] N. E. Breslow and D. G. Clayton, "Approximate inference in generalized linear mixed models," *J. Amer. Stat. Assoc.*, vol. 88, no. 421, pp. 9–25, Mar. 1993.

[53] F. Wilcoxon, "Individual comparisons by ranking methods," in *Breakthroughs in Statistics*. New York, NY, USA: Springer, 1992, pp. 196–202.

[54] X. Mao, Y. Ma, Z. Yang, Y. Chen, and Q. Li, "Virtual mixup training for unsupervised domain adaptation," 2019, *arXiv:1905.04215*. [Online]. Available: http://arxiv.org/abs/1905.04215

[55] R. D. Altman and G. E. Gold, "Atlas of individual radiographic features in osteoarthritis, revised," *Osteoarthritis Cartilage*, vol. 15, pp. A1–A56, 2007.

[56] A. Tiulpin *et al.*, "Multimodal machine learning-based knee osteoarthritis progression prediction from plain radiographsand clinical data," *Sci. Rep.*, vol. 9, no. 1, pp. 1–11, 2019.