

Unsupervised Domain Adaptation for Semantic Segmentation of NIR Images through Generative Latent Search

Prashant Pandey^{*[0000-0002-6594-9685]}, Aayush Kumar Tyagi^{*[0000-0002-3615-7283]}, Sameer Ambekar^[0000-0002-8650-3180], and Prathosh AP^[0000-0002-8699-5760]

Indian Institute of Technology Delhi
getprashant57@gmail.com, aayush16081@iiitd.ac.in,
ambekarsameer@gmail.com, prathoshap@iiitd.ac.in

Abstract. Segmentation of the pixels corresponding to human skin is an essential first step in multiple applications ranging from surveillance to heart-rate estimation from remote-photoplethysmography. However, the existing literature considers the problem only in the visible-range of the EM-spectrum which limits their utility in low or no light settings where the criticality of the application is higher. To alleviate this problem, we consider the problem of skin segmentation from the Near-infrared images. However, Deep learning based state-of-the-art segmentation techniques demands large amounts of labelled data that is unavailable for the current problem. Therefore we cast the skin segmentation problem as that of target-independent Unsupervised Domain Adaptation (UDA) where we use the data from the Red-channel of the visible-range to develop skin segmentation algorithm on NIR images. We propose a method for target-independent segmentation where the ‘nearest-clone’ of a target image in the source domain is searched and used as a proxy in the segmentation network trained only on the source domain. We prove the existence of ‘nearest-clone’ and propose a method to find it through an optimization algorithm over the latent space of a Deep generative model based on variational inference. We demonstrate the efficacy of the proposed method for NIR skin segmentation over the state-of-the-art UDA segmentation methods on the two newly created skin segmentation datasets in NIR domain despite not having access to the target NIR data. Additionally, we report state-of-the-art results for adaption from Synthia to Cityscapes which is a popular setting in Unsupervised Domain Adaptation for semantic segmentation. The code and datasets are available at <https://github.com/ambekarsameer96/GLSS>.

Keywords: Unsupervised Domain Adaptation, Semantic segmentation, Near IR Dataset, VAE

* equal contribution

1 Introduction

1.1 Background

Human skin segmentation is the task of finding pixels corresponding to skin from images or videos. It serves as a necessary pre-processing step for multiple applications like video surveillance, people tracking, human computer interaction, face detection and recognition, facial gesture detection and monitoring heart rate and respiratory rate [7, 32, 33, 38] using remote photoplethysmography. Most of the research efforts on skin detection have focused on visible spectrum images because of the challenges that it poses including, illumination change, ethnicity change and presence of background/clothes similar to skin colour. These factors adversely affect the applications where skin is used as conjugate information. Further, the algorithms that rely on visible spectrum images cannot be employed in the low/no light conditions especially during night times where the criticality of the application like human detection is higher. These problems which are encountered in visible spectrum domain can be overcome by considering the images taken in the Near-infrared (NIR) domain [24] or hyper spectral imaging [35]. The information about the skin pixels is invariant of factors such as illumination conditions, ethnicity etc., in these domains. Moreover, most of the surveillance cameras that are used world-wide are NIR imaging devices. Thus, it is meaningful to pursue the endeavour of detecting the skin pixels from the NIR images.

1.2 Problem setting and contributions

The task of detection of skin pixels from an image is typically cast as a segmentation problem. Most of the classical approaches relied on the fact that the skin-pixels have a distinctive color pattern [13, 19] compared to other objects. In recent years, harnessing the power of Deep learning, skin segmentation problem has been dealt with using deep neural networks that show significant performance enhancement over the traditional methods [21, 30, 40], albeit generalization across different illuminations still remains a challenge. While there exists sufficient literature on skin segmentation in the visible-spectrum, there is very little work done on segmenting the skin pixels in the NIR domain. Further, all the state-of-the-art Deep learning based segmentation algorithms demand large-scale annotated datasets to achieve good performance which is available in the case of visible-spectrum images but not the NIR images. Thus, building a fully-supervised skin segmentation network from scratch is not feasible for the NIR images because of the unavailability of the large-scale annotated data. However, the underlying concept of ‘skin-pixels’ is the same across the images irrespective of the band in which they were captured. Additionally, the NIR and the Red-channel of the visible-spectrum are close in terms of their wavelengths. Owing to these observations, we pose the following question in this paper - Can the labelled data (source) in the visible-spectrum (Red-channel) be used to perform skin segmentation in the NIR domain (target) [37]?

We cast the problem of skin segmentation from NIR images as a target-independent Unsupervised Domain Adaptation (UDA) task [36] where we consider the Red-channel of the visible-spectrum images as the source domain and NIR images as the target domain. The state-of-the-art UDA techniques demand access to the target data, albeit unlabelled, to adapt the source domain features to the target domain. In the present case, we do not assume existence of any data from the target domain, even unlabelled. This is an important desired attribute which ensures that a model trained on the Red-channel does not need any retraining with the data from NIR domain. The core idea is to sample the ‘nearest-clone’ in the source domain to a given test image from the target domain. This is accomplished through a simultaneous sampling-cum-optimization procedure using a latent-variable deep neural generative network learned on the source distribution. Thus, given a target sample, its ‘nearest-clone’ from the source domain is sampled and used as a proxy in the segmentation network trained only on the samples of the source domain. Since the segmentation network performs well on the source domain, it is expected to give the correct segmentation mask on the ‘nearest-clone’ which is then assigned to the target image. Specifically, the core contributions of this work are listed as follows:

1. We cast the problem of skin segmentation from NIR images as a UDA segmentation task where we use the data from the Red-channel of the visible-range of the EM-spectrum to develop skin segmentation algorithm on NIR images.
2. We propose a method for target-independent segmentation where the ‘nearest-clone’ of a target image in the source domain is searched and used as a proxy in the segmentation network trained only on the source domain.
3. We theoretically prove the existence of the ‘nearest-clone’ given that it can be sampled from the source domain with infinite data points.
4. We develop a joint-sampling and optimization algorithm using variational inference based generative model to search for the ‘nearest-clone’ through implicit sampling in the source domain.
5. We demonstrate the efficacy of the proposed method for NIR skin segmentation over the state-of-the-art UDA segmentation methods on the two newly created skin segmentation datasets in NIR domain. The proposed method is also shown to reach SOTA performance on standard segmentation datasets like Synthia [41] and Cityscapes [10].

2 Related Work

In this section, we first review the existing methods for skin segmentation in the visible-range followed by a review of UDA methods for segmentation.

2.1 Skin Segmentation in Visible-range

Methods for skin segmentation can be grouped into three categories, i.e. (i) Thresholding based methods [13, 25, 39], (ii) Traditional machine learning techniques to learn a skin color model [29, 51], (iii) Deep learning based methods

to learn an end-to-end model for skin segmentation [2, 7, 15, 42, 50]. The thresholding based methods focus on defining a specified range in different color representation spaces like (HSV)[34] and orthogonal color space (YCbCr)[3, 18] to differentiate skin pixels. Traditional machine learning can be further divided into pixel based and region based methods. In pixel based methods, each pixel is classified as skin or non-skin without considering the neighbours [45] whereas region based approaches use spatial information to identify similar regions [8]. In recent years, Fully convolutional neural networks (FCN) are employed to solve the problem [30]. [40] proposed a UNet architecture, consisting of an encoder-decoder structure with backbones like InceptionNet[43] and ResNet [14]. Holistic skin segmentation [12] combine inductive transfer learning and UDA. They term this technique as cross domain pseudo-labelling and use it in an iterative manner to train and fine tune the model on the target domain. [15] propose mutual guidance to improve skin detection with the usage of body masks as guidance. They use dual task neural network for joint detection with shared encoder and two decoders for detecting skin and body simultaneously. While all these methods offer different advantages, they do not generalize to low-light settings with NIR images, which we aim to solve through UDA.

2.2 Domain Adaptation for semantic segmentation

Unsupervised Domain Adaptation aims to improve the performance of deep neural networks on a target domain, using labels only from a source domain. UDA for segmentation task can be grouped into following categories:

Adversarial training based methods: These methods use the principles of adversarial learning [16], which generally consists of two networks. One predicts the segmentation mask of the input image coming from either source or target distribution while the other network acts as discriminator which tries to predict the domain of the images. AdaptSegNet [46] exploits structural similarity between the source and target domains in a multi-level adversarial network framework. ADVENT [47] introduce entropy-based loss to directly penalize low-confident predictions on target domain. Adversarial training is used for structural adaptation of the target domain to the source domain. CLAN [31] considers category-level joint distribution and aligns each class with an adaptive adversarial loss. They reduce the weight of the adversarial loss for category-level aligned features while increasing the adversarial force for those that are poorly aligned. DADA [48] uses the geometry of the scene by simultaneously aligning the segmentation and depth-based information of source and target domains using adversarial training.

Feature-transformation based methods: These methods are based on the idea of learning image-level or feature-level transformations between the source and the target domains. CyCADA [1] adapts between domains using both generative image space alignment and latent representation space alignment. Image

level adaptation is achieved with cycle loss, semantic consistency loss and pixel-level GAN loss while feature level adaptation employs feature-level GAN loss and task loss between true and predicted labels. DISE [4] aims to discover a domain-invariant structural feature by learning to disentangle domain-invariant structural information of an image from its domain-specific texture information. BDL [26] involves two separated modules a) image-to-image translation model b) segmentation adaptation model, in two directions namely ‘translation-to-segmentation’ and ‘segmentation-to-translation’.

3 Proposed method

Most of the UDA methods assume access to the unlabelled target data which may not be available at all times. In this work, we propose a UDA segmentation technique by learning to find a data point from the source that is arbitrarily close (called the ‘nearest-clone’) to a given target point so that it can be used as a proxy in the segmentation network trained only on the source data. In the subsequent sections, we describe the methodology used to find the ‘nearest-clone’ from the source distribution to a given target point.

3.1 Existence of nearest source point

To start with, we show that for a given target data point, there exists a corresponding source data point, that is arbitrarily close to, provided that infinite data points can be sampled from the source distribution. Mathematically, let $\mathcal{P}_s(\mathbf{x})$ denote the source distribution and $\mathcal{P}_t(\mathbf{x})$ denote any target distribution that is similar but not exactly same as \mathcal{P}_s (Red-channel images are source and NIR images are target). Let the underlying random variable on which \mathcal{P}_s and \mathcal{P}_t are defined form a separable metric space $\{\mathcal{X}, \mathcal{D}\}$ with \mathcal{D} being some distance metric. Let $\mathcal{S}_n = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \dots, \mathbf{x}_n\}$ be i.i.d points drawn from $\mathcal{P}_s(\mathbf{x})$ and $\tilde{\mathbf{x}}$ be a point from $\mathcal{P}_t(\mathbf{x})$. With this, the following lemma shows the existence of the ‘nearest-clone’.

Lemma 1. *If $\tilde{\mathbf{x}}_S \in \mathcal{S}_n$ is the point such that $\mathcal{D}\{\tilde{\mathbf{x}}, \tilde{\mathbf{x}}_S\} < \mathcal{D}\{\tilde{\mathbf{x}}, \mathbf{x}\} \forall \mathbf{x} \in \mathcal{S}_n$, as $n \rightarrow \infty$ (in \mathcal{S}_n), $\tilde{\mathbf{x}}_S$ converges to $\tilde{\mathbf{x}}$ with probability 1.*

Proof. Let $\mathbb{B}_r(\tilde{\mathbf{x}}) = \{\mathbf{x} : \mathcal{D}\{\tilde{\mathbf{x}}, \mathbf{x}\} \leq r\}$ be a closed ball of radius r around $\tilde{\mathbf{x}}$ under the metric \mathcal{D} . Since \mathcal{X} is a separable metric space [11],

$$\mathbf{Prob}(\mathbb{B}_r(\tilde{\mathbf{x}})) \triangleq \int_{\mathbb{B}_r(\tilde{\mathbf{x}})} \mathcal{P}_s(\mathbf{x}) d\mathbf{x} > 0, \forall r > 0, \quad (1)$$

With this, for any $\delta > 0$, the probability that none of the points in \mathcal{S}_n are within the ball $\mathbb{B}_\delta(\tilde{\mathbf{x}})$ of radius δ is:

$$\mathbf{Prob} \left[\min_{i=1,2,\dots,n} \mathcal{D}\{\mathbf{x}_i, \tilde{\mathbf{x}}\} \geq \delta \right] = [1 - \mathbf{Prob}(\mathbb{B}_\delta(\tilde{\mathbf{x}}))]^n \quad (2)$$

Therefore, the probability of $\tilde{\mathbf{x}}_S$ (the closest point to $\tilde{\mathbf{x}}$) lying within $\mathbb{B}_\delta(\tilde{\mathbf{x}})$ is:

$$\mathbf{Prob}\left[\tilde{\mathbf{x}}_S \in \mathbb{B}_\delta(\tilde{\mathbf{x}})\right] = 1 - [1 - \mathbf{Prob}(\mathbb{B}_\delta(\tilde{\mathbf{x}}))]^n \quad (3)$$

$$= 1 \text{ as } n \rightarrow \infty \quad (4)$$

Thus, given any infinitesimal $\delta > 0$, with probability 1, $\exists \tilde{\mathbf{x}}_S \in \mathcal{S}_n$ (‘nearest-clone’) that is within δ distance from $\tilde{\mathbf{x}}$ as $n \rightarrow \infty$ \square

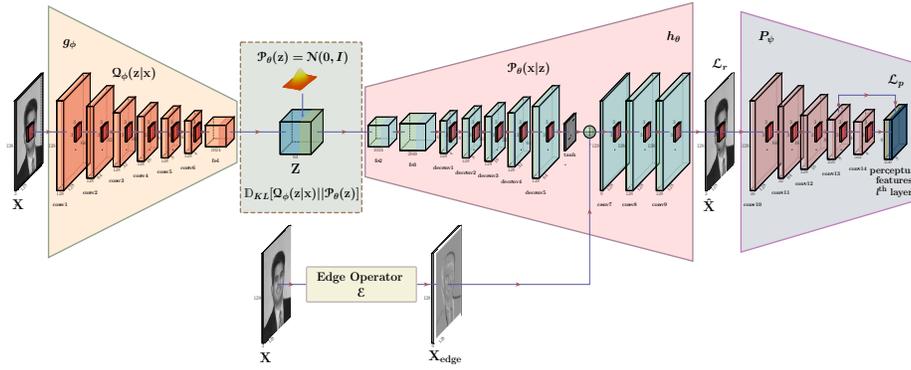


Fig. 1: VAE training. Edges of an input image are concatenated with the features from the decoder h_θ . Encoder and decoder parameters ϕ, θ are optimized with reconstruction loss \mathcal{L}_r , KL-divergence loss \mathcal{D}_{KL} and perceptual loss \mathcal{L}_p . Perceptual model P_ψ is trained on source samples. A zero mean and unit variance isotropic Gaussian prior is imposed on the latent space \mathbf{z} .

While Lemma 1 guarantees the existence of a ‘nearest-clone’, it demands the following two conditions:

- It should be possible to sample infinitely from the source distribution \mathcal{P}_s .
- It should be possible to search for the ‘nearest-clone’ in the \mathcal{P}_s , for a target sample $\tilde{\mathbf{x}}$ under the distance metric \mathcal{D} .

We propose to employ Variational Auto-encoding based sampling models on the source distribution to simultaneously sample and find the ‘nearest-clone’ through an optimization over the latent space.

3.2 Variational Auto-Encoder for source sampling

Variational Auto-Encoders (VAEs) [23] are a class of latent-variable generative models that are based on the principles of variational inference where the variational distribution, $Q_\phi(\mathbf{z}|\mathbf{x})$ is used to approximate the intractable true posterior $\mathcal{P}_\theta(\mathbf{z}|\mathbf{x})$. The log-likelihood of the observed data is decomposed into two terms,

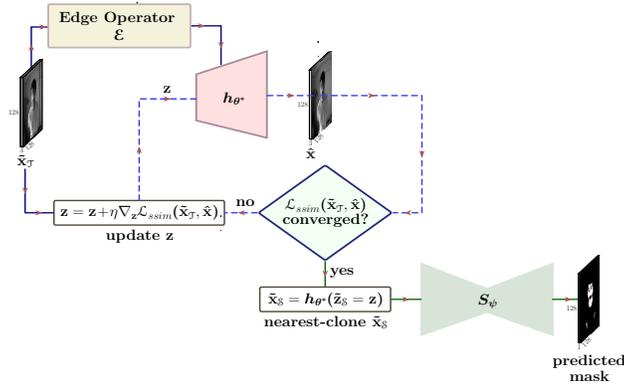


Fig. 2: Latent Search procedure during inference with GLSS. The latent vector \mathbf{z} is initialized with a random sample drawn from $\mathcal{N}(0, 1)$. Iterations over the latent space \mathbf{z} are performed to minimize the \mathcal{L}_{sim} loss between the input target image $\tilde{\mathbf{x}}_T$ and the predicted target image $\hat{\mathbf{x}}$ (blue dotted lines). After convergence of \mathcal{L}_{sim} loss, the optimal latent vector $\tilde{\mathbf{z}}_S$, generates the closest clone $\tilde{\mathbf{x}}_S$ which is used to predict the mask of $\tilde{\mathbf{x}}_T$ using the segmentation network S_ψ .

an irreducible non-negative KL-divergence between $\mathcal{P}_\theta(\mathbf{z}|\mathbf{x})$ and $\mathcal{Q}_\phi(\mathbf{z}|\mathbf{x})$ and the Evidence Lower Bound (ELBO) term which is given by Eq. 5.

$$\ln \mathcal{P}_\theta(\mathbf{x}) = \mathcal{L}(\theta, \phi) + \mathbb{D}_{KL}[\mathcal{Q}_\phi(\mathbf{z}|\mathbf{x})||\mathcal{P}_\theta(\mathbf{z}|\mathbf{x})] \quad (5)$$

where,

$$\mathcal{L}(\theta, \phi) = \mathbb{E}_{\mathcal{Q}_\phi(\mathbf{z}|\mathbf{x})}[\ln(\mathcal{P}_\theta(\mathbf{x}|\mathbf{z}))] - \mathbb{D}_{KL}[\mathcal{Q}_\phi(\mathbf{z}|\mathbf{x})||\mathcal{P}_\theta(\mathbf{z})] \quad (6)$$

The non-negative KL-term in Eq. 5 is irreducible and thus, $\mathcal{L}(\theta, \phi)$ serves as a lower bound on the data log-likelihood which is maximized in a VAE by parameterizing $\mathcal{Q}_\phi(\mathbf{z}|\mathbf{x})$ and $\mathcal{P}_\theta(\mathbf{x}|\mathbf{z})$ using probabilistic encoder g_ϕ (that outputs the parameters $\mu_{\mathbf{z}}$ and $\sigma_{\mathbf{z}}$ of a distribution) and decoder h_θ neural networks. The latent prior $\mathcal{P}_\theta(\mathbf{z})$ is taken to be arbitrary prior on \mathbf{z} which is usually a 0 mean and unit variance Gaussian distribution. After training, the decoder network is used as a sampler for $\mathcal{P}_s(\mathbf{x})$ in a two-step process: (i) Sample $\mathbf{z} \sim \mathcal{N}(0, I)$, (ii) Sample \mathbf{x} from $\mathcal{P}_\theta(\mathbf{x}|\mathbf{z})$. The likelihood term in Eq. 5 is approximated using norm based losses and it is known to result in blurry images. Therefore, we use the perceptual loss [20] along with the standard norm based losses. Further, since the edges in images are generally invariant across the source and target domains, we extract edge of the input image and use it in the decoder of the VAE via a skip connection, as shown in Fig. 1. This is shown to reduce the blur in the generated images. Fig. 1 depicts the entire VAE architecture used for training on the source data.

3.3 VAE Latent Search for finding the ‘nearest-clone’

As described, the objective of the current work is to search for the nearest point in the source distribution, given a sample from target distribution. The decoder h_θ of a VAE trained on the source distribution $\mathcal{P}_s(\mathbf{x})$, outputs a new sample using the Normally distributed latent sample as input. That is,

$$\forall \mathbf{z} \sim \mathcal{N}(0, I), \hat{\mathbf{x}} = h_\theta(\mathbf{z}) \sim \mathcal{P}_s(\hat{\mathbf{x}}) \quad (7)$$

With this, our goal is to find the ‘nearest-clone’ to a given target sample. That is, given a $\tilde{\mathbf{x}} \sim \mathcal{P}_t(\mathbf{x})$, find $\tilde{\mathbf{x}}_S$ as follows:

$$\tilde{\mathbf{x}}_S = h_\theta(\tilde{\mathbf{z}}_S) : \left\{ \mathcal{D}\{\tilde{\mathbf{x}}, \tilde{\mathbf{x}}_S\} < \mathcal{D}\{\mathbf{x}, \tilde{\mathbf{x}}\} \quad \forall \mathbf{x} = h_\theta(\mathbf{z}) \sim \mathcal{P}_s(\mathbf{x}) \right\} \quad (8)$$

Since \mathcal{D} is pre-defined and $h_\theta(\mathbf{z})$ is a deep neural network, finding $\tilde{\mathbf{x}}_S$ can be cast as an optimization problem over \mathbf{z} with minimization of \mathcal{D} as the objective. Mathematically,

$$\tilde{\mathbf{z}}_S = \arg \min_{\mathbf{z}} \mathcal{D}(\tilde{\mathbf{x}}, h_\theta(\mathbf{z})) \quad (9)$$

$$\tilde{\mathbf{x}}_S = h_\theta(\tilde{\mathbf{z}}_S) \quad (10)$$

The optimization problem in Eq. 9 can be solved using gradient-descent based techniques on the decoder network h_{θ^*} (θ^* are the parameters of the decoder network trained only on the source samples \mathcal{S}_n) with respect to \mathbf{z} . This implies that given any input target image, the optimization problem in Eq. 9 will be solved to find its ‘nearest-clone’ in the source distribution which is used as a proxy in the segmentation network trained only on \mathcal{S}_n . We call the iterative procedure of finding $\tilde{\mathbf{x}}_S$ through optimization using h_{θ^*} as the Latent Search (LS). Finally, inspired by the observations made in [17], we propose to use structural similarity index (SSIM) [49] based loss \mathcal{L}_{ssim} for \mathcal{D} to conduct the Latent Search. Unlike norm based losses, SSIM loss helps in preservation of structural information which is needed for segmentation. Fig. 2 depicts the complete inference procedure employed in the proposed method named as the Generative Latent Search for Segmentation (GLSS).

4 Implementation Details

4.1 Training

Architectural details of the VAE used are shown in Fig. 1. Sobel operator is used to extract the edge information of the input image which is concatenated with one of the layers of the Decoder via a *tanh* non linearity as shown in Fig. 1. The VAE is trained using (i) the Mean squared error reconstruction loss \mathcal{L}_r and KL divergence \mathbb{D}_{KL} and (ii) the perceptual loss \mathcal{L}_p for which the features are

extracted from the l^{th} layer (a hyper-parameter) of the DeepLabv3+ [6] (Xception backbone [9]) and the UNet [40] (EfficientNet backbone [44]) segmentation networks. The segmentation network (S_ψ in Fig. 2) is either DeepLabv3+ or UNet and is trained on the source dataset. For training S_ψ , we use combination of binary cross-entropy (\mathcal{L}_{bce}) and dice coefficient loss (\mathcal{L}_{dise}) for UNet with RMSProp (lr = 0.001) as optimizer and binary focal loss (\mathcal{L}_{focal}) [28] with $\gamma = 2.0$, $\alpha = 0.75$ and RMSProp (lr=0.01) as optimizer for DeepLabv3+. For the VAE, the hidden layers of Encoder and Decoder networks use Leaky ReLU and \tanh as activation functions with the dimensionality of the latent space being 64. VAE is trained using standard gradient descent procedure with RMSprop ($\alpha=0.0001$) as optimizer. We train VAE for 100 to 150 epochs with batchsize 64.

4.2 Inference

Once the VAE is trained on the source dataset, given an image $\tilde{\mathbf{x}}_T$ from the target distribution, the Latent Search algorithm searches for an optimal latent vector $\tilde{\mathbf{z}}_S$ that generates its ‘nearest-clone’ $\tilde{\mathbf{x}}_S$ from \mathcal{P}_S . The search is performed by minimizing the SSIM loss \mathcal{L}_{ssim} between the input target image $\tilde{\mathbf{x}}_T$ and the VAE-reconstructed target image, using a gradient-descent based optimization procedure such as ADAM [22] with $\alpha = 0.1$, $\beta_1 = 0.9$ and $\beta_2 = 0.99$. The Latent Search is performed for K (hyper-parameter) iterations over the latent space of the source for a given target image. Finally, the segmentation mask for the input target sample is assigned the same as the one given by the segmentation network S_ψ , which is trained on source data, on the ‘nearest-clone’ $\tilde{\mathbf{x}}_S$. Latent Search for one sample takes roughly 450 ms and 120 ms on SNV and Hand Gesture datasets respectively. Please refer supplementary material for more details.

5 Experiment and Results

5.1 Datasets

We consider the Red-channel of the COMPAQ dataset [21] as our source data. It consists of 4675 RGB images with the corresponding annotations of the skin. Since there is no publicly available dataset with NIR images and corresponding skin segmentation labels, we create and use two NIR datasets (publicly available) as targets. The first one named as the Skin NIR Vision (SNV), consists of 800 images of multiple human subjects taken in different scenes, captured using a WANSVIEW 720P camera in the night-vision mode. The captured images cover wide range of scenarios for skin detection task like presence of multiple humans, backgrounds similar to skin color, different illuminations, saturation levels and different postures of subjects to ensure diversity. Additionally, we made use of the publicly available multi-modal Hand Gesture dataset¹ as another target dataset which we call as Hand Gesture dataset. This dataset covers 16 different hand-poses of multiple subjects. We randomly sampled 500 images in order to cover

¹ https://www.gti.ssr.upm.es/data/MultiModalHandGesture_dataset

illumination changes and diversity in hand poses. Both SNV and Hand Gesture datasets are manually annotated with precision.

5.2 Benchmarking on SNV and Hand Gesture datasets

To begin with, we performed supervised segmentation experiments on both SNV and Hand Gesture datasets with 80-20 train-test split using SOTA segmentation algorithms.

Table 1: Benchmarking Skin NIR Vision (SNV) dataset and Hand Gesture dataset on standard segmentation architectures with 80-20 train-test split.

Method	SNV		Hand Gesture	
	IoU	Dice	IoU	Dice
FPN [27]	0.792	0.895	0.902	0.950
UNet [40]	0.798	0.890	0.903	0.950
DeepLabv3+ [6]	0.750	0.850	0.860	0.924
Linknet [5]	0.768	0.872	0.907	0.952
PSPNet [52]	0.757	0.850	0.905	0.949

Table 1 shows the standard performance metrics such as IoU and Dice-coefficient calculated using FPN [27], UNet [40], LinkNet [5], PSPNet [52], all with EfficientNet [44] as backbone and DeepLabv3+ [6] with Xception network [9] as backbone. It is seen that SNV dataset (IoU \approx 0.79) is slightly complex as compared to Hand Gesture dataset (IoU \approx 0.90).

Table 2: Empirical analysis of GLSS along with standard UDA methods. IoU and Dice-coefficient are computed for both SNV and Hand Gesture datasets using UNet and DeepLabv3+ as segmentation networks.

Models	SNV				Hand Gesture			
	UNet		DeepLabv3+		UNet		DeepLabv3+	
	IoU	Dice	IoU	Dice	IoU	Dice	IoU	Dice
Source Only	0.295	0.426	0.215	0.426	0.601	0.711	0.505	0.680
AdaptSegnet [46]	0.315	0.435	0.230	0.435	0.641	0.716	0.542	0.736
Advent [47]	0.341	0.571	0.332	0.540	0.612	0.729	0.508	0.689
CLAN [31]	0.248	0.442	0.225	0.426	0.625	0.732	0.513	0.692
BDL [26]	0.320	0.518	0.301	0.509	0.647	0.720	0.536	0.750
DISE [4]	0.341	0.557	0.339	0.532	0.672	0.789	0.563	0.769
DADA [48]	0.332	0.534	0.314	0.521	0.643	0.743	0.559	0.761
ours (GLSS)	0.406	0.597	0.385	0.597	0.736	0.844	0.698	0.824

5.3 Baseline UDA Experiments

SNV and Hand Gesture dataset: We have performed the UDA experiments with the SOTA UDA algorithms using Red-channel of the COMPAQ Dataset [21] as the source and SNV and Hand Gesture as the target. Table 2 compares the performance of proposed GLSS algorithm with six SOTA baselines along with the Source Only case (without any UDA). We have used entire target dataset for IoU and Dice-coefficient evaluation. Two architectures, DeepLabv3+ and UNet,

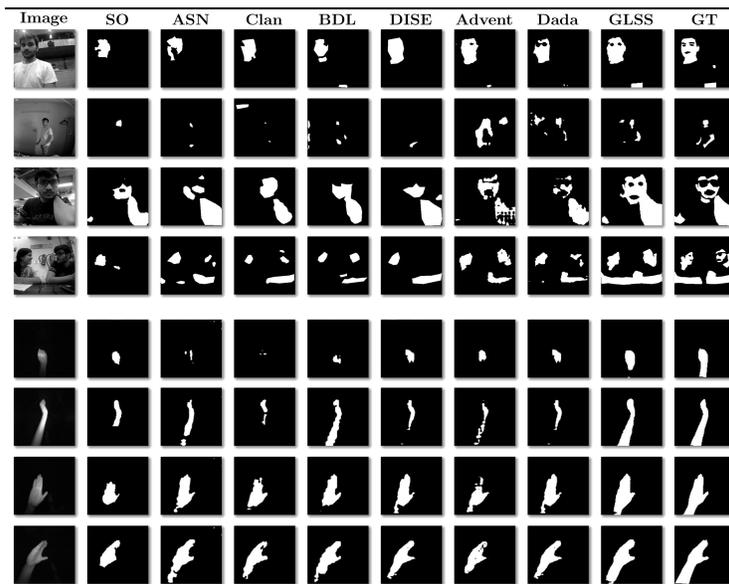


Fig. 3: Qualitative comparison of predicted segmentation skin masks on SNV and Hand Gesture datasets with standard UDA methods. Top four rows shows skin masks for SNV dataset and the last four are the masks for Hand Gesture dataset. It is evident that GLSS predicted masks are very close to the GT masks as compared to other UDA methods. (SO=Source Only, ASN=AdaptSegNet [46], GT=Ground Truth).

are employed for the segmentation network (S_ψ). It can be seen that although all the UDA SOTA methods improve upon the Source Only performance, GLSS offers significantly better performance despite not using any data from the target distribution. Hence, it may be empirically inferred that GLSS is successful in producing the ‘nearest-clone’ through implicit sampling from the source distribution and thereby reducing the domain shift. It is also observed that the performance of the segmentation network S_ψ does not degrade on the source data with GLSS. The predicted masks with DeepLabv3+ are shown in Fig. 3 for

SNV and Hand Gesture datasets, respectively. It can be seen that GLSS is able to capture fine facial details like eyes, lips and body parts like hands, better as compared to SOTA methods. It is also seen that the predicted masks for Hand Gesture dataset are sharper in comparison to other methods. Most of the methods work with the assumption of spatial and structural similarity between the source and target data. Since our source and target datasets do not have similar backgrounds, the methods that make such assumptions perform poorer on our datasets. We observed that for methods like BDL, the image translation between NIR images and Red channel images is not effective for skin segmentation task.

Standard UDA task: We use standard UDA methods along with GLSS on standard domain adaptation datasets such as Synthia [41] and Cityscapes [10]. As observed from Table 3, even with large domain shift, GLSS finds a clone for every target image that is sampled from the source distribution while preserving the structure of the target image.

Table 3: Empirical analysis of GLSS on standard domain adaptation task of adapting Synthia [41] to Cityscapes [10]. We calculate the mean IoU for 13 classes (mIoU) and 16 classes (mIoU*).

Models	mIoU	mIoU*
AdaptsegNet [46]	46.7	-
Advent [47]	48.0	41.2
BDL [26]	51.4	-
CLAN [31]	47.8	-
DISE [4]	48.8	41.5
DADA [48]	49.8	42.6
ours(GLSS)	52.3	44.5

5.4 Ablation Study

We have conducted several ablation experiments on GLSS using both SNV and Hand Gesture datasets using DeepLabv3+ as segmentation networks (S_ψ) to ascertain the utility of different design choices we have made in our method.

Effect of number of iterations on LS: The inference of GLSS involves a gradient-based optimization through the decoder network h_{θ^*} to generate the ‘nearest-clone’ for a given target image. In Fig. 4, we show the skin masks of the transformed target images after every 30 iterations. It is seen that with the increasing number of iterations, the predicted skin masks improves using GLSS as the ‘nearest-clones’ are optimized during the Latent Search procedure. We plot the IoU as a function of the number of iterations during Latent Search as

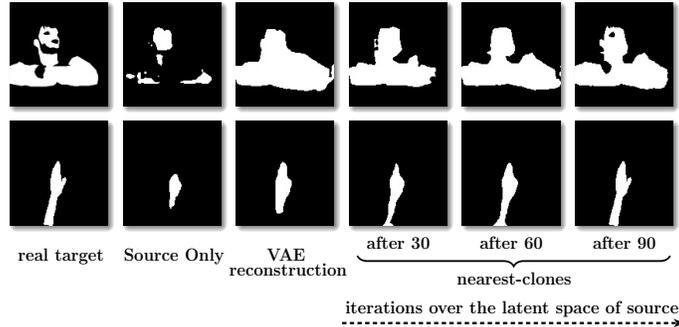


Fig. 4: Illustration of Latent Search in GLSS. Real target is a ground truth mask. Source Only masks are obtained from target samples by training segmentation network S_ψ on source dataset. Prior to the LS, skin masks are obtained from VAE reconstructed target samples. It is evident that predicted skin masks improve as the LS progresses. The predicted masks for the ‘nearest-clones’ are shown after every 30 iterations.

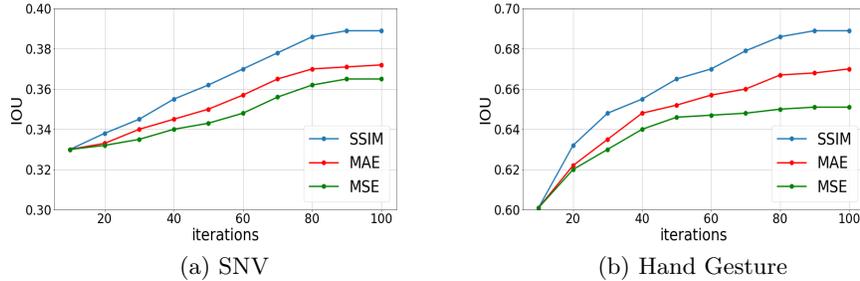


Fig. 5: Performance of gradient-based Latent Search during inference on target SNV and Hand Gesture images using different objective functions; MSE, MAE, SSIM loss. DeepLabv3+ is employed as segmentation network. It is evident that the losses saturate at around 90-100 iterations.

shown in Fig. 5 where it is seen that it saturates around 90-100 iterations that are used for all the UDA experiments described in the previous section.

Effect of Edge concatenation: As discussed earlier, edges extracted using Sobel filter on input images are concatenated with one of the layers of decoder for both training and inference. It is seen from Table 4 that IoU improves for both the target datasets with concatenation of edges. It is observed that without the edge concatenation, the generated images (‘nearest-clones’) are blurry thus the segmentation network fails to predict sharper skin masks.

Table 4: Ablation of different components of GLSS during training and inference; Edge, perceptual loss \mathcal{L}_p and Latent Search (LS).

Edge	\mathcal{L}_p	LS	SNV IoU	Hand Gesture IoU
			0.112	0.227
✓			0.178	0.560
	✓		0.120	0.250
		✓	0.128	0.238
✓	✓		0.330	0.615
	✓	✓	0.182	0.300
✓		✓	0.223	0.58
✓	✓	✓	0.385	0.698

Effect of Perceptual loss \mathcal{L}_p : We have introduced a perceptual model P_ψ , trained on source samples. It ensures that the VAE reconstructed image is semantically similar to the input image unlike the norm based losses. Table 4 clearly demonstrates the improvement offered by the use of perceptual loss while training the VAE.

Effect of SSIM for Latent Search: Finally, to validate the effect of SSIM loss for Latent Search, we plot the IoU metric using two norm based losses MSE (Mean squared error) and MAE (Mean absolute error) for the Latent Search procedure as shown in Fig. 5. On both the datasets, it is seen that SSIM is consistently better than the norm based losses at all iterations affirming the superiority of the SSIM loss in preserving the structures while finding the ‘nearest-clone’.

6 Conclusion

In this paper, we addressed the problem of skin segmentation from NIR images. Owing to the non-existence of large-scale labelled NIR datasets for skin segmentation, the problem is casted as Unsupervised Domain Adaptation where we use the segmentation network trained on the Red-channel images from a large-scale labelled visible-spectrum dataset for UDA on NIR data. We propose a novel method for UDA without the need for the access to the target data (even unlabelled). Given a target image, we sample an image from the source distribution that is ‘closest’ to it under a distance metric. We show that such a ‘closest’ sample exists and describe a procedure using an optimization algorithm over the latent space of a VAE trained on the source data. We demonstrate the utility of the proposed method along with the comparisons with SOTA UDA segmentation methods on the skin segmentation task on two NIR datasets that were created. Also, we reach SOTA performance on Synthia and Cityscapes datasets for semantic segmentation of urban scenes.

References

1. Cycada: Cycle consistent adversarial domain adaptation. In: International Conference on Machine Learning (ICML) (2018)
2. Al-Mohair, H.K., Saleh, J., Saundi, S.: Impact of color space on human skin color detection using an intelligent system. In: 1st WSEAS international conference on image processing and pattern recognition (IPPR'13). vol. 2 (2013)
3. Brancati, N., De Pietro, G., Frucci, M., Gallo, L.: Human skin detection through correlation rules between the ycb and ycr subspaces based on dynamic color clustering. *Computer Vision and Image Understanding* **155**, 33–42 (2017)
4. Chang, W.L., Wang, H.P., Peng, W.H., Chiu, W.C.: All about structure: Adapting structural information across domains for boosting semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1900–1909 (2019)
5. Chaurasia, A., Culurciello, E.: Linknet: Exploiting encoder representations for efficient semantic segmentation. In: 2017 IEEE Visual Communications and Image Processing (VCIP). pp. 1–4. IEEE (2017)
6. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: ECCV (2018)
7. Chen, W., Wang, K., Jiang, H., Li, M.: Skin color modeling for face detection and segmentation: a review and a new approach. *Multimedia Tools and Applications* **75**(2), 839–862 (2016)
8. Chen, W.C., Wang, M.S.: Region-based and content adaptive skin detection in color images. *International journal of pattern recognition and artificial intelligence* **21**(05), 831–853 (2007)
9. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1251–1258 (2017)
10. Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., Schiele, B.: The cityscapes dataset for semantic urban scene understanding. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3213–3223 (2016)
11. Cover, T., Hart, P.: Nearest neighbor pattern classification. *IEEE transactions on information theory* **13**(1), 21–27 (1967)
12. Dourado, A., Guth, F., de Campos, T.E., Weigang, L.: Domain adaptation for holistic skin detection. *arXiv preprint arXiv:1903.06969* (2019)
13. Erdem, C., Ulukaya, S., Karaali, A., Erdem, A.T.: Combining haar feature and skin color based classifiers for face detection. In: 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 1497–1500. IEEE (2011)
14. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
15. He, Y., Shi, J., Wang, C., Huang, H., Liu, J., Li, G., Liu, R., Wang, J.: Semi-supervised skin detection by network with mutual guidance. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 2111–2120 (2019)
16. Hoffman, J., Wang, D., Yu, F., Darrell, T.: Fcns in the wild: Pixel-level adversarial and constraint-based adaptation. *arXiv preprint arXiv:1612.02649* (2016)
17. Hore, A., Ziou, D.: Image quality metrics: Psnr vs. ssim. In: 2010 20th International Conference on Pattern Recognition. pp. 2366–2369. IEEE (2010)

18. Hsu, R.L., Abdel-Mottaleb, M., Jain, A.K.: Face detection in color images. *IEEE transactions on pattern analysis and machine intelligence* **24**(5), 696–706 (2002)
19. Huynh-Thu, Q., Meguro, M., Kaneko, M.: Skin-color-based image segmentation and its application in face detection. In: *MVA*. pp. 48–51 (2002)
20. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: *European conference on computer vision*. pp. 694–711. Springer (2016)
21. Jones, M.J., Rehg, J.M.: Statistical color models with application to skin detection. *International Journal of Computer Vision* **46**(1), 81–96 (2002)
22. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
23. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013)
24. Kong, S.G., Heo, J., Abidi, B.R., Paik, J., Abidi, M.A.: Recent advances in visual and infrared face recognition—a review. *Computer Vision and Image Understanding* **97**(1), 103–135 (2005)
25. Kovac, J., Peer, P., Solina, F.: Human skin color clustering for face detection, vol. 2. *IEEE* (2003)
26. Li, Y., Yuan, L., Vasconcelos, N.: Bidirectional learning for domain adaptation of semantic segmentation. *arXiv preprint arXiv:1904.10620* (2019)
27. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 2117–2125 (2017)
28. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2980–2988 (2017)
29. Liu, Q., Peng, G.z.: A robust skin color based face detection algorithm. In: *2010 2nd International Asia Conference on Informatics in Control, Automation and Robotics (CAR 2010)*. vol. 2, pp. 525–528. *IEEE* (2010)
30. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 3431–3440 (2015)
31. Luo, Y., Zheng, L., Guan, T., Yu, J., Yang, Y.: Taking a closer look at domain shift: Category-level adversaries for semantics consistent domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pp. 2507–2516 (2019)
32. Mahmoodi, M.R.: High performance novel skin segmentation algorithm for images with complex background. *arXiv preprint arXiv:1701.05588* (2017)
33. Mahmoodi, M.R., Sayedi, S.M.: A comprehensive survey on human skin detection. *International Journal of Image, Graphics & Signal Processing* **8**(5) (2016)
34. Moallem, P., Mousavi, B.S., Monadjemi, S.A.: A novel fuzzy rule base system for pose independent faces detection. *Applied Soft Computing* **11**(2), 1801–1810 (2011)
35. Pan, Z., Healey, G., Prasad, M., Tromberg, B.: Face recognition in hyperspectral images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **25**(12), 1552–1560 (2003)
36. Pandey, P., AP, P., Kyatham, V., Mishra, D., Dastidar, T.R.: Target-independent domain adaptation for wbc classification using generative latent search. *arXiv preprint arXiv:2005.05432* (2020)

37. Pandey, P., Prathosh, A., Kohli, M., Pritchard, J.: Guided weak supervision for action recognition with scarce data to assess skills of children with autism. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 34, pp. 463–470 (2020)
38. Prathosh, A., Praveena, P., Mestha, L.K., Bharadwaj, S.: Estimation of respiratory pattern from video using selective ensemble aggregation. *IEEE Transactions on Signal Processing* **65**(11), 2902–2916 (2017)
39. Qiang-rong, J., Hua-lan, L.: Robust human face detection in complicated color images. In: 2010 2nd IEEE International Conference on Information Management and Engineering. pp. 218–221. IEEE (2010)
40. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: International Conference on Medical image computing and computer-assisted intervention. pp. 234–241. Springer (2015)
41. Ros, G., Sellart, L., Materzynska, J., Vazquez, D., Lopez, A.M.: The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3234–3243 (2016)
42. Seow, M.J., Valaparla, D., Asari, V.K.: Neural network based skin color model for face detection. In: 32nd Applied Imagery Pattern Recognition Workshop, 2003. Proceedings. pp. 141–145. IEEE (2003)
43. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
44. Tan, M., Le, Q.V.: Efficientnet: Rethinking model scaling for convolutional neural networks. arXiv preprint arXiv:1905.11946 (2019)
45. Taqa, A.Y., Jalab, H.A.: Increasing the reliability of skin detectors. *Scientific Research and Essays* **5**(17), 2480–2490 (2010)
46. Tsai, Y.H., Hung, W.C., Schuster, S., Sohn, K., Yang, M.H., Chandraker, M.: Learning to adapt structured output space for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7472–7481 (2018)
47. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Advent: Adversarial entropy minimization for domain adaptation in semantic segmentation. In: CVPR (2019)
48. Vu, T.H., Jain, H., Bucher, M., Cord, M., Pérez, P.: Dada: Depth-aware domain adaptation in semantic segmentation. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 7364–7373 (2019)
49. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
50. Wu, Q., Cai, R., Fan, L., Ruan, C., Leng, G.: Skin detection using color processing mechanism inspired by the visual system (2012)
51. Zaidan, A., Ahmad, N.N., Karim, H.A., Larbani, M., Zaidan, B., Sali, A.: On the multi-agent learning neural and bayesian methods in skin detector and pornography classifier: An automated anti-pornography system. *Neurocomputing* **131**, 397–418 (2014)
52. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid scene parsing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2881–2890 (2017)