

Multi-Energy Guided Image Translation with Stochastic Differential Equations for Near-Infrared Facial Expression Recognition

Bingjun Luo¹, Zewen Wang¹, Jinpeng Wang¹, Junjie Zhu¹, Xibin Zhao^{1*}, Yue Gao^{1*}

¹BNRist, KLISS, School of Software, Tsinghua University

luobingjun@gmail.com, {wang-zw19,wjp21,zhujj18}@mails.tsinghua.edu.cn, {zxb, gaoyue}@tsinghua.edu.cn

Abstract

Illumination variation has been a long-term challenge in real-world facial expression recognition (FER). Under uncontrolled or non-visible light conditions, near-infrared (NIR) can provide a simple and alternative solution to obtain high-quality images and supplement the geometric and texture details that are missing in the visible (VIS) domain. Due to the lack of large-scale NIR facial expression datasets, directly extending VIS FER methods to the NIR spectrum may be ineffective. Additionally, previous heterogeneous image synthesis methods are restricted by low controllability without prior task knowledge. To tackle these issues, we present the first approach, called for NIR-FER Stochastic Differential Equations (NFER-SDE), that transforms face expression appearance between heterogeneous modalities to the overfitting problem on small-scale NIR data. NFER-SDE can take the whole VIS source image as input and, together with domain-specific knowledge, guide the preservation of modality-invariant information in the high-frequency content of the image. Extensive experiments and ablation studies show that NFER-SDE significantly improves the performance of NIR FER and achieves state-of-the-art results on the only two available NIR FER datasets, Oulu-CASIA and Large-HFE.

Introduction

Facial expression (FE) is one of the most powerful, natural, and universal means for human beings to convey their emotions. In various human-computer interaction applications, automatic facial expression recognition (FER) can provide more insights into users' emotional states and intentions, which is essential for achieving better behavioral responses in driver fatigue surveillance, sociable robotics, and healthcare. To improve the robustness of FER in these real-world applications, tremendous efforts have been made from different perspectives, such as occluded facial image [34], noisy labels [25, 29], cross-dataset generalization [15], and so on [30, 44]. However, FER under extreme lighting conditions is still an open question.

As illustrated in Fig 1, it's quite difficult for visible (VIS) sensors to capture high-quality images in low lighting or complete darkness conditions. The limitations of these sen-



Figure 1: Comparison of VIS and NIR facial images under extreme lighting. The red frame in the center of the figure is NIR images. Even under low lighting or complete darkness, the geometric and texture details of volunteers in NIR images are still very well preserved.

sors result in the loss of crucial geometric and textural details, especially in key areas such as the cheek or lips. Consequently, even experienced annotators find it arduous to accurately identify the expression type depicted in such images. Hence, it becomes exceedingly challenging to solely rely on algorithmic enhancements to address this issue. Fortunately, the implementation of near-infrared (NIR) imaging techniques offers a low-cost and effective solution to enhance the performance of VIS FER systems in low-light conditions. Furthermore, given their reduced sensitivity to visible (VIS) light illumination variations, NIR face sensors have been widely used in security surveillance within finance, transportation, and other related domains.

However, achieving better performance for NIR FER is still a challenging problem for deep models due to the following two reasons: 1) Insufficient training samples. In the field of VIS, after decades of hard work, more and more high-quality data sets, such as CK+, RAF-DB, KDEF, etc., continue to emerge, and the overall amount of available training data is also growing rapidly. In contrast, there are currently only two NIR facial expression datasets available, Oulu-CASIA (80 subjects) and Large-HFE (360 subjects). 2) The large modal discrepancy. VIS and NIR face images taken of the same individual are obtained through distinct

*Corresponding author.

sensory devices with varying settings, resulting in significant differences in their visual appearances. Thus, directly applying deep models trained on VIS data fails to provide satisfactory results in the NIR spectrum.

To tackle these challenges together, synthesis methods pose opportunities to generate NIR images given input VIS face images so that abundant label information will migrate from VIS to NIR. However, existing methods usually synthesize images without prior task knowledge. Thus, they can not realize customizable domain transfer that allows a controllable appearance for FER. According to (Zhang and Tjondronegoro 2011), the feature of facial expression usually lies in subtle facial movements, which requires detailed facial representations to be preserved in the facial expression translation process.

Inspired by these observations, we present a novel method, called NFER-SDE, to learn and synthesize NIR expression images of good quality from VIS. The synthesized facial images not only preserve the high-frequency content carrying modality-invariant information, but also contain detailed facial representations that are crucial to FER tasks. Specifically, NFER-SDE consists of two novel components, energy functions of task-specific guidance and conditional score matching. First, a novel energy function \mathcal{E}_h is proposed to guide the reverse stochastic differential equations (SDE) process using facial landmark information. Thus, the key facial structure information can be preserved through the reverse SDE process, resulting in lower uncertainty. Encouraged by unique observation in the cross-modality (NIR-VIS) facial expression translation task, the frequency sharing rule is established to preserve more modality-invariant information. Second, conditional score matching is introduced to make sure that the generated NIR image can be more faithful to the original VIS image.

In summary, our main contributions are three-fold:

(1) We develop a novel method, i.e. NFER-SDE, to translate the VIS facial expression sample to NIR modality. To the best of our knowledge, this is the first effort to synthesize NIR facial expression samples for the NIR FER task.

(2) We propose new task-specific energy functions to guide the image generation process with the domain knowledge of FER. We also employ a conditional score-matching network to better focus on facial expression details from the original image.

(3) We conduct extensive experiments on Oulu-CASIA and Large-HFE. The results show that our method has a better performance than the state-of-the-art image synthesis approaches.

Related Work

Facial Expression Recognition The goal of FER is to automatically encode emotional information from subtle facial muscle changes. Inspired by the success of deep neural networks, extensive efforts have been made to address different problems including occluded facial image (Xing et al. 2022), noisy labels (Wang et al. 2022), cross-dataset generalization (Li and Deng 2020), dynamic FER (Zhao and Liu 2021). As discussed in [3], dealing with the challenge of illumination variations is also crucial to FER in real-world

applications. NIR images provide a low-cost and effective solution to enhance the performance of VIS FER systems in uncontrolled or non-visible light conditions. In this paper, we present the first approach that transforms face expression appearance between heterogeneous modalities for NIR FER in extreme light conditions.

Image-to-image Translation Image-to-image translation is an image generation task to transform the images between two different visual domains. In this field, various generative models have made great progress, including Generative adversarial networks (GAN) (Isola et al. 2017; Yang et al. 2020), Variational AutoEncoders (VAE) (Kingma et al. 2014; Sohn, Lee, and Yan 2015; Peng et al. 2021). Recently, diffusion models have attracted much attention due to their remarkable image generation ability in terms of distribution diversity (Croitoru et al. 2023), and thus are utilized in the image-to-image translation task. ILVR (Choi et al. 2021) proposes an iterative latent variable refinement strategy to guide the generative process of diffusion models. EGSDE (Zhao et al. 2022) employs a pretrained energy function to guide the reverse process of score-based diffusion models.

Problem Definition

Let $\mathcal{D}_{\text{HTR}} = \{(\mathbf{x}_V^{(i)}, \mathbf{x}_N^{(i)}, y^{(i)}) | i = 1, \dots, n\}$ denote an NIR-VIS heterogeneous facial expression dataset, where $\mathbf{x}_V^{(i)}$ and $\mathbf{x}_N^{(i)}$ are the VIS and NIR images of the i -th sample respectively, and $y^{(i)}$ is the corresponding expression label. Let $\mathcal{D}_{\text{VIS}} = \{(\mathbf{x}^{(j)}, y^{(j)}) | j = 1, \dots, m\}$ denote a VIS facial expression dataset, where $\mathbf{x}^{(j)}$ is the VIS image of the j -th sample and $y^{(j)}$ is the corresponding label. The goal of NIR facial expression synthesis is to learn a translation model F from the heterogeneous dataset \mathcal{D}_{HTR} , i.e., the model can transform arbitrary facial expression image \mathbf{x}_V of VIS modality to the NIR modality $\mathbf{x}_N = F(\mathbf{x}_V)$. Using this translation model, we can create a synthesized NIR-VIS facial expression dataset $\mathcal{D}_{\text{SYS}} = \{(\mathbf{x}^{(j)}, F(\mathbf{x}^{(j)}), y^{(j)}) | j = 1, \dots, j\}$ from the VIS dataset \mathcal{D}_{VIS} , where $F(\mathbf{x}^{(j)})$ is the synthesized NIR image of the j -th sample. The final goal for NIR facial expression image synthesis is that the FER backbone trained on the augmented dataset $\mathcal{D}_{\text{HTR}} \cup \mathcal{D}_{\text{VIS}}$ can achieve better performance than that only trained on \mathcal{D}_{HTR} .

Preliminaries

Diffusion Models

Diffusion models are a class of deep generative models, which have achieved remarkable performance in image generation over previous state-of-the-art methods like generative adversarial networks (GANs) (Goodfellow et al. 2014). The core idea of diffusion models is to gradually corrupt the training data structure by adding noise to the input, and then learn a reverse network to recover the original data from the corrupted data (Croitoru et al. 2023).

The training process of diffusion models is formulated as two stages: the forward stage and the reverse stage. In the forward stage, random Gaussian noise is iteratively added to the input, with the proportion of noise increasing from 0 to

1. In other words, the input is gradually degraded into pure noise. The reverse stage is a noise-removal process, which is the reversal of the preceding forward noise-adding process. In the reverse stage, the output of the forward stage is recovered iteratively using a denoising neural network that is trained to predict the added noise at each step.

During sampling, random noise is sampled from the Gaussian distribution and then directly fed into the reverse stage to generate the corresponding sample. Therefore, diffusion models can learn the distribution of training data and generate samples from random noise by the reverse network.

Score-based Diffusion Models (SBDMs)

Score-based diffusion models (SBDMs) are a continuous form of the generic diffusion model. By transforming the above discrete diffusion process into a continuous form, SBDMs utilize Stochastic Differential Equation (SDE) to model and solve the diffusion process (Song et al. 2020).

Let $\{\mathbf{x}_t\}_{t=0}^T$ denote the image sequence in the forward stage, where \mathbf{x}_0 is the original image and \mathbf{x}_T is the final degraded image. In SBDMs, the time index t is regarded as a continuous variable $t \in [0, T]$, and the image sequence $\{\mathbf{x}_t\}_{t=0}^T$ is generalized into a continuous function $\mathbf{x}_t = \mathbf{x}(t)$. Thus, the forward stage can be formulated as the corresponding SDE

$$d\mathbf{x}_t = \mathbf{f}(\mathbf{x}_t, t)dt + \sigma(t)d\boldsymbol{\omega} \quad (1)$$

where \mathbf{f} is the drift coefficient of $\mathbf{x}(t)$ that depends on the image $\mathbf{x}(t)$ and time t , σ is the diffusion coefficient of $\mathbf{x}(t)$ that only depends on time t , and $\boldsymbol{\omega}$ is the integral of the white noise Gaussian process, i.e., the standard Wiener process. According to (Song et al. 2020), general DDPM converges to the Variance Preserving SDE as the number of sample steps $T \rightarrow \infty$, where $\mathbf{f}(\mathbf{x}_t, t) = -\frac{1}{2}\beta(t)\mathbf{x}(t)$, $\sigma(t) = \sqrt{\beta(t)}$. Thus, DDPM networks can be utilized to implement SBDMs.

Given the forward SDE in Eq. (1), the corresponding reverse SDE is solved as

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - \sigma(t)^2 \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)]dt + \sigma(t)d\hat{\boldsymbol{\omega}} \quad (2)$$

where $p(\mathbf{x}_t)$ is the marginal data distribution at time t , and $\hat{\boldsymbol{\omega}}$ is the reverse Wiener process from T to 0. To estimate the marginal data distribution, SBDMs propose a score-matching neural network $s(\mathbf{x}_t, t, \Theta_s)$ to parameterize the score function $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t)$, where Θ_s is the parameter of the network. Using the Euler-Maruyama method, the reverse SDE can be further discretized as

$$\Delta \mathbf{x} = [\mathbf{f}(\mathbf{x}_t, t) - \sigma(t)^2 s(\mathbf{x}_t, t, \Theta_s)]\Delta t + \sigma(t)\sqrt{|\Delta t|}z \quad (3)$$

where $\Delta t < 0$ is the reverse step size, and $z \sim \mathcal{N}(0, \mathbf{I})$ is a standard Gaussian noise. Eq. (3) provides the iterative rule for the reverse stage, which enables the computational simulation of SDE for unconditional image generation.

Energy-Guided Conditional SBDMs

To generate images under specific requirements or conditions, the conditional form of SBDMs is proposed to incorporate the conditional information into the diffusion process

(Song et al. 2020). Given a condition \mathbf{c} , the data distribution of \mathbf{x}_t turns into the conditional distribution $p(\mathbf{x}_t|\mathbf{c})$. According to Bayesian rule, the score function of the target data distribution is formulated as $\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t|\mathbf{c}) = \nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t)$. Thus, the reverse SDE in Eq. (2) is modified as

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - \sigma(t)^2 (\nabla_{\mathbf{x}_t} \log p(\mathbf{x}_t) + \nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t))]dt + \sigma(t)d\hat{\boldsymbol{\omega}} \quad (4)$$

where $\nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t)$ is added to the original score function as a modification term. Given the condition representation model $p(\mathbf{c}|\mathbf{x}_t)$, the reverse SDE can be calculated under the guidance of the corresponding condition. The representation model $p(\mathbf{c}|\mathbf{x}_t)$ is exactly matched with some special cases like classifier guidance (Zhao et al. 2022). However, in some more cases like image translation, the conditional distribution $p(\mathbf{c}|\mathbf{x}_t)$ is intractable.

To obtain a more generalized form of guidance, the conditional SBDMs further propose an energy function $\mathcal{E}(\mathbf{x}_t, \mathbf{c}, t)$ to parameterize the modification term $\nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t)$ (Zhao et al. 2022). The energy function $\mathcal{E}(\mathbf{x}_t, \mathbf{c}, t)$ is trained or designed to be minimized when the condition is satisfied. Thus, the modification term is formulated as $\nabla_{\mathbf{x}_t} \log p(\mathbf{c}|\mathbf{x}_t) = -\nabla_{\mathbf{x}_t} \mathcal{E}(\mathbf{x}_t, \mathbf{c}, t)$. Given the score matching network $s(\mathbf{x}_t, t, \Theta_s)$ and the energy function $\mathcal{E}(\mathbf{x}_t, \mathbf{c}, t)$, the reverse SDE is formulated as

$$d\mathbf{x}_t = [\mathbf{f}(\mathbf{x}_t, t) - \sigma(t)^2 (s(\mathbf{x}_t, t, \Theta_s) - \nabla_{\mathbf{x}_t} \mathcal{E}(\mathbf{x}_t, \mathbf{c}, t))]dt + \sigma(t)d\hat{\boldsymbol{\omega}} \quad (5)$$

With the guidance of different energy functions, Eq. (5) can be applied to various conditional image generation tasks.

Method

In this section, we present the proposed method for VIS-NIR facial expression image translation. Firstly, we introduce the framework of the proposed method. Then, we describe the energy function of task-specific guidance and the conditional score matching network in detail.

Model Overview

The framework of the proposed method is illustrated in Fig. 2. Specifically, the proposed NFER-SDE consists of two main components. The first component is the newly designed energy functions of task-specific guidance, which can guide the image-generation process with the domain-specific knowledge of facial expression, including information on the key facial areas and the unique frequency-sharing rule of cross-modality expression translation. The second component is the conditional score-matching network, which can make full use of the subtle facial details in the original facial expression image by explicitly conditioning it as part of the input of the score-matching network.

Energy Functions of Task-Specific Guidance

Conditional diffusion models (Choi et al. 2021; Meng et al. 2021; Zhao et al. 2022) have made great progress and achieved significant results in general image-to-image translation tasks, such as image editing and style transfer. Yet,

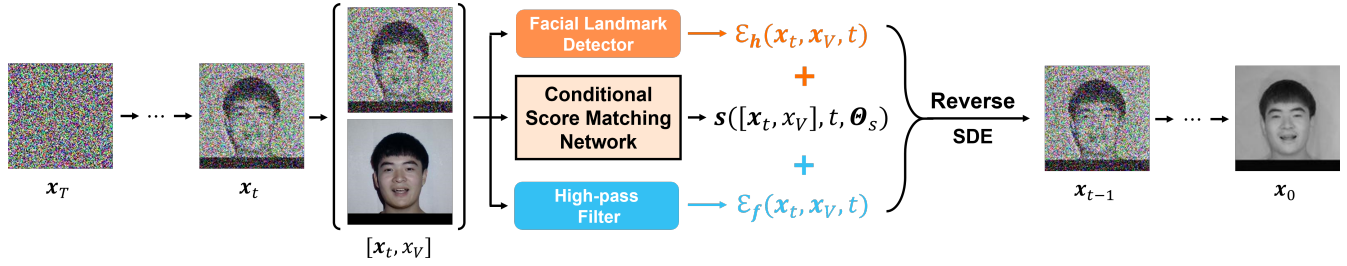


Figure 2: The main framework of the proposed NFER-SDE network for VIS-NIR facial expression translation. The framework mainly consists of two task-specific guidance energy functions and a conditional score-matching network.

when applied to VIS-NIR facial expression translation task, the existing methods do not perform as expected in transferring facial expression attributes. The underperformance is partly due to the feature learning challenges in VIS-NIR facial expression translation task (Wang et al. 2020a), which can be summarized in the following two aspects.

Firstly, compared to general vision targets, facial expressions are usually conveyed through more subtle facial representations in specific facial areas (Zhu et al. 2023). This requires the diffusion models to be well guided by key facial structure information that is related to facial expression. Recently, it has been widely recognized that the facial landmark coordinates can encode the facial structure information, which is crucial for facial expression recognition (Lv et al. 2019; Wang et al. 2020b). Therefore, a novel energy function \mathcal{E}_h is proposed to guide the reverse SDE process using facial landmark information. Specifically, a pretrained time-dependent facial landmark detection network $H(x, t, \Theta_h)$ is proposed to extract the facial landmark heatmap from the input image x and current step t , with its parameter Θ_h fixed during training. Let c_t denote the perturbed source image in the forward SDE and x_t denote the denoised target image in the reverse SDE both at step t . The energy function is formulated as

$$\mathcal{E}_h(x_t, c, t) = \frac{1}{2} \|H(x_t, t, \Theta_h) - H(c_t, t, \Theta_h)\|_2^2 \quad (6)$$

This energy function is designed to be minimized when the facial landmark heatmap of the target image x_t is close to that of the perturbed source image c_t . Thus, the key facial structure information can be preserved through the reverse SDE process.

Secondly, the frequency sharing rule in the cross-modality translation task is unique, as shown in Fig. 3. In the cross-modality facial expression translation task, the high-frequency content of the image carries more modality-invariant information for facial expression translation, as it contains more detailed facial expression representations. This rule is quite different from that in general image translation tasks where the low-frequency content is regarded to carry the domain-invariant information that should be shared between the source and target domains (Choi et al. 2021; Zhao et al. 2022). Thus, the low-pass filters that general image translation tasks adopt to preserve low-frequency content are not suitable for the cross-modality facial expression translation task anymore.

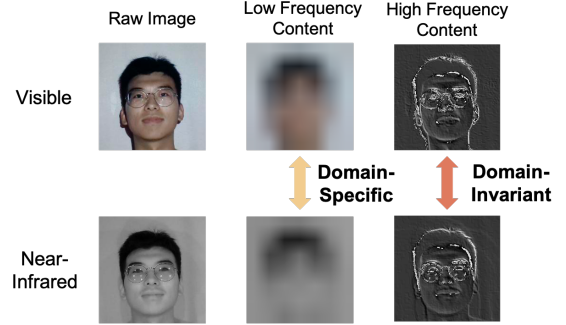


Figure 3: Visualization of frequency sharing rule in cross-modality facial expression translation. It can be observed that the high-frequency content carries more modality-invariant information that should be preserved in cross-modality facial expression translation.

Therefore, we propose a novel energy function based on a high-pass filter to preserve the high-frequency content in the reverse SDE process. Given a high-pass filter $\Phi(\cdot)$ as the domain-invariant feature extractor, the energy function is formulated as

$$\mathcal{E}_f(x_t, c, t) = \|\Phi(x_t) - \Phi(c_t)\|_2^2 \quad (7)$$

With the guidance of \mathcal{E}_f , the high-frequency content can be preserved through the reverse SDE process to preserve the domain-invariant feature in the cross-modality translation task.

To summarize, the above two energy functions are aggregated to form the final energy function

$$\mathcal{E}(x_t, c, t) = \lambda_h \mathcal{E}_h(x_t, c, t) + \lambda_f \mathcal{E}_f(x_t, c, t) \quad (8)$$

where λ_h and λ_f are the weighting hyper-parameters.

Conditional Score Matching

According to (Zhang and Tjondronegoro 2011), the feature of facial expression usually lies in subtle facial movements, which requires detailed facial representations to be preserved in the facial expression translation process. Thus, in the VIS-NIR facial expression translation task, the original VIS image should be fully considered by the denoising network $s(x_t, t, \Theta_s)$ to ensure the faithfulness of the generated NIR image. However, existing SDE-based models

(Meng et al. 2021; Choi et al. 2021; Zhao et al. 2022) do not directly take the original image information into the denoising neural network at each iterative reverse step, which may cause the generated image to lose some key facial details.

To solve the above problem, we propose to condition the denoising neural network with the reference image c in the cross-modality translation process. Inspired by (Saharia et al. 2022), the reference image c is concatenated to the input image x_t in the channel dimension at each step t . The concatenated vector is then fed into the score matching network to generate the output

$$s_t = s([x_t, c], t, \Theta_s) \quad (9)$$

where $s(\cdot, \cdot, \Theta_s)$ is the extension of that in Eq. (5). For the training algorithm of the proposed conditional score matching network, please refer to (Nair and Patel 2023).

Using the additional input c during training, the new conditioned denoising network is supervised to generate better NIR images under the guidance of the detailed information from the VIS source image. Thus, the generated NIR image can be more faithful to the original VIS image.

Experiment Setup

Evaluation Settings

To obtain comprehensive evaluation of NFER-SDE, we follow previous works in heterogeneous translation (Yang et al. 2020; Nair and Patel 2023) and design two settings as follows: improved NIR FER performance and VIS-NIR translation quality. For **improved NIR FER performance**, extra VIS facial expression samples are translated to NIR modality, and then added to the NIR training samples of the original NIR&VIS facial expression dataset. The augmented NIR training samples are utilized to train FER benchmarks. The well-trained FER benchmarks are then evaluated on the original NIR test set and the accuracy and macro F1 score are reported as the improved NIR FER performance. For **VIS-NIR translation quality**, the NIR images generated from the VIS test samples of the original NIR&VIS facial expression dataset are compared with the ground-truth NIR test samples of the original dataset according to the image quality metrics including Learned Perceptual Image Patch Similarity (LPIPS) (Zhang et al. 2018), Peak Signal-to-Noise Ratio (PSNR), Structural Similarity (SSIM).

Datasets

For NIR&VIS facial expression datasets, we select the only two available ones: Oulu-CASIA (Zhao et al. 2011) and Large-HFE. **Oulu-CASIA** is a laboratory-controlled dataset that consists of video sequences captured with both visible and near-infrared lights from 80 subjects. **Large-HFE** is also captured in the laboratory environment, with a larger size of 360 subjects. In both datasets, only the six basic expressions are considered. For each dataset, we employ the subject-independent strategy to split it into a train set of 80% subjects and a test set of 20% subjects.

For extra VIS facial expression dataset, we select CFEE (Du, Tao, and Martinez 2014) dataset. **CFEE** captures facial expression images in the laboratory with visible light

from 230 subjects, in which only the 6 basic expressions are used in our experiments. The full set of CFEE is utilized to be transformed into NIR modality by the proposed model and the transformed samples are added to the training set of Oulu-CASIA or Large-HFE to jointly train the NIR FER benchmark.

Baselines

To comprehensively evaluate the proposed method, We select the following image-to-image translation baselines for comparison, including the GAN-based methods and diffusion-based methods. For GAN-based methods, we select the state-of-the-art methods in image translation, including pix2pix (Isola et al. 2017), CycleGAN (Zhu et al. 2017), HiFaceGAN (Yang et al. 2020). For diffusion-based methods, we select the state-of-the-art methods in image restoration, including two parts: the denoising diffusion probabilistic models (DDPMs) and score-based diffusion models (SBDMs). The DDPM baselines used in this paper include ILVR (Choi et al. 2021), SR3 (Saharia et al. 2022), and T2V-DDPM (Nair and Patel 2023). The SBDM baselines include SDEdit (Meng et al. 2021) and EGSDE (Zhao et al. 2022). Among these baselines, HiFaceGAN and T2V-DDPM are specially designed for heterogeneous face translation, while the others are general image translation methods. For a fair comparison, we train all the baselines on the same training set and evaluate them on the same test set.

Implementation Details

The proposed method is implemented by Pytorch 1.8, and the model training and image sampling procedure is processed on NVIDIA GeForce 4090 card. All images are first cropped to 256×256 using the face detector (Bulat and Tzimiropoulos 2017) and then both NIR&VIS facial expression datasets are divided into the training set and testing set. The training set is used to train the backbone model. The backbone conditional score matching network is implemented by U-Net. During training, the batch size is set to 16 and the learning rate is set to 1×10^{-5} . After the model is ready, we can sample images using our method from CFEE and KDEF for NIR FER performance experiment and from the testing set for VIS-NIR translation quality experiment. The hyperparameter λ_h is set to 100 and λ_f is set to 0.5, which are roughly selected based on the image generation quality. For each experiment setting, the datasets are randomly divided into five parts for five independent experiments, and averaged results are reported in the following part.

Results and Analysis

In this section, we first compare the proposed method with the baselines in terms of NIR FER performance and VIS-NIR translation quality. Then we conduct ablation study to analyze the effectiveness of each component in the proposed method. Finally, we visualize the generated NIR images to have qualitative evaluation of the compared methods.

Method		Performance		Image Quality		
		Acc (%)	F1 (%)	LPIPS (\downarrow)	PSNR (\uparrow)	SSIM(\uparrow)
Without translation		61.11	61.44	0.6369	48.45	0.9069
GAN-based	pix2pix (Isola et al. 2017)	63.06	63.07	0.5392	52.61	0.9372
	CycleGAN (Zhu et al. 2017)	66.57	66.86	0.4582	53.52	0.9507
	HiFaceGAN (Yang et al. 2020) †	66.94	67.00	0.4937	55.29	0.9922
Diffusion-based	ILVR (Choi et al. 2021)	56.94	57.19	0.7163	48.51	0.9072
	SR3 (Saharia et al. 2022)	61.20	60.41	0.7958	51.84	0.9189
	SDEdit (Meng et al. 2021)	64.91	64.72	0.6462	49.16	0.9117
	EGSDE (Zhao et al. 2022)	64.26	64.26	0.6008	48.75	0.9081
	T2V (Nair and Patel 2023) †	66.48	66.46	0.2713	59.70	0.9920
	Ours (proposed)	70.19	70.28	0.1827	65.75	0.9994

Table 1: The performance of the compared methods on Oulu-CASIA. † denotes methods specially designed for heterogeneous face translation. The best results are highlighted in bold.

Method		Performance		Image Quality		
		Acc (%)	F1 (%)	LPIPS (\downarrow)	PSNR (\uparrow)	SSIM(\uparrow)
Without translation		67.07	67.16	0.5364	50.88	0.9417
GAN-based	pix2pix (Isola et al. 2017)	67.17	67.02	0.5657	52.69	0.9589
	CycleGAN (Zhu et al. 2017)	72.49	72.52	0.4801	54.03	0.9677
	HiFaceGAN (Yang et al. 2020) †	69.68	69.75	0.4142	55.26	0.9962
Diffusion-based	ILVR (Choi et al. 2021)	68.67	69.05	0.6053	50.88	0.9420
	SR3 (Saharia et al. 2022)	65.86	65.54	0.5687	54.59	0.9691
	SDEdit (Meng et al. 2021)	69.88	69.76	0.5534	51.11	0.9439
	EGSDE (Zhao et al. 2022)	70.78	70.91	0.5704	50.94	0.9423
	T2V (Nair and Patel 2023) †	72.59	72.87	0.3965	56.12	0.9830
	Ours (proposed)	74.90	74.91	0.3878	58.98	0.9963

Table 2: The performance of the compared methods on Large-HFE. † denotes methods specially designed for heterogeneous face translation. The best results are highlighted in bold.

Comparison with State-of-the-art Methods

To evaluate the proposed method, we compare it with the state-of-the-art methods in the settings of improved NIR FER performance and VIS-NIR translation quality on Oulu-CASIA and Large-HFE. The quantitative results are shown in Table 1 and Table 2 respectively. From the results, we have the observations as follows.

(1) The proposed method shows consistent advantages over the state-of-the-art baselines in terms of improved NIR FER performance and VIS-NIR translation quality on both datasets. Compared to the best baseline, the proposed method achieves the improvement of accuracy by 3.25% and 2.31%, and F1 score by 3.28% and 2.04% on Oulu-CASIA and Large-HFE. This indicates that the proposed method can translate VIS facial expression images to NIR modality with satisfying image quality and joint training with the translated images can effectively enhance the performance of NIR FER benchmark. This is mainly due to introduction of task-specific energy functions and conditional score matching, which guide the model to focus on subtle yet key facial representations that are related to facial expression.

(2) Among the baselines, T2V and HiFaceGAN tend to show consistent effectiveness in image quality on both datasets. This is because T2V and HiFaceGAN are spe-

cially designed for heterogeneous face translation, and thus can generate more realistic NIR images of the facial details than the other baselines. CycleGAN also produces relatively competitive results, since the detailed original VIS image is considered by the GAN-based network. The other diffusion-based models, including ILVR, SDEdit, and EGSDE, perform not as expected, which is mainly due to the ignorance of details in the original image information in translation.

Ablation Study

The proposed method mainly consists of two parts, i.e., the conditional score-matching network and the task-specific energy guidance. To analyze the effectiveness of each component, we conduct ablation study on Oulu-CASIA and Large-HFE under the following settings:

- **Baseline:** Basic SBDM (unconditional score-matching network and no task-specific energy function).
- **Cond:** Conditional score-matching network.
- **Cond+ \mathcal{E}_h :** Conditional score-matching network with energy function \mathcal{E}_h in Eq. (6).
- **Cond+ \mathcal{E}_f :** Conditional score-matching network with energy function \mathcal{E}_f in Eq. (7).
- **Cond+ \mathcal{E}_h + \mathcal{E}_f :** The proposed NFER-SDE method.

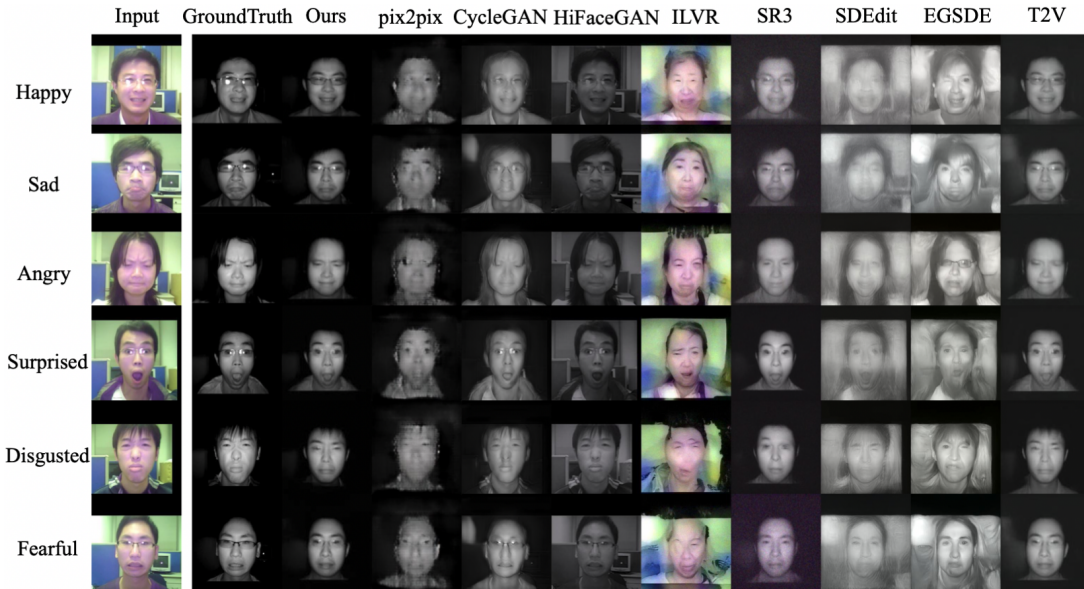


Figure 4: Visualization of generated samples from the compared methods on Oulu-CASIA.

Method	Oulu-CASIA		Large-HFE	
	Acc (%)	F1 (%)	Acc (%)	F1 (%)
Baseline	62.55	62.75	69.12	69.43
Cond	65.76	66.06	71.99	72.17
Cond+ \mathcal{E}_h	67.78	68.98	72.39	72.49
Cond+ \mathcal{E}_f	68.24	68.57	73.59	73.53
Cond+ $\mathcal{E}_h+\mathcal{E}_f$	70.19	70.28	74.90	74.91

Table 3: The ablation study results on Oulu-CASIA and Large-HFE. The best results are highlighted in bold.

The ablation study results are shown in Table 3. Observations can be drawn from the results as follows.

(1) The baseline of unconditional score-matching network and no task-specific energy function performs the worst, which is due to a lack of detailed facial representations and lack of knowledge guidance in the reverse process.

(2) Adding conditional score-matching network significantly improves the performance. This is because the added condition of the source facial expression image provides the score-matching network with key facial details to generate more realistic NIR facial expressions. Adding two task-specific energy functions further improves the performance, which can guide the generation process toward more realistic NIR facial expressions with satisfying quality.

(3) The proposed method that jointly combines the above components achieves the best performance. This indicates that both the task-specific energy guidance and the conditional score-matching network are effective in improving the performance of NIR FER benchmark.

Visualization of Generated Samples

For qualitative evaluation results, we present the translated samples of the proposed method and the baselines on

Oulu-CASIA in Fig. 4. As shown in the figure, the proposed method can generate NIR images with satisfying quality, while the baselines cannot. Methods including ILVR, pix2pix, SDEdit, and EGSDE generate human facial images with varying degrees of distortion on both datasets. Among them, ILVR performs the worst, since even fails to translate VIS images to NIR-like images. All other methods can meet the basic requirements of VIS-NIR translation. However, they still have some disadvantages compared to our proposed method. Images generated from CycleGAN and HiFaceGAN still suffer from black square artifact on Large-HFE, and on Oulu-CASIA, CycleGAN misses expression-related facial details (as in *Angry*), while HiFaceGAN remains the background in VIS images. SR3 and T2V images have a paler color tone compared to real NIR images, which can be seen more clearly on Oulu-CASIA. Furthermore, SR3 fails to generate images of the same person on Large-HFE in some cases (As in *Angry*, *Surprised*, *Disgusted*). In conclusion, the qualitative analysis results are consistent with the quantitative results, which further verifies the effectiveness of the proposed method.

Conclusion

In this paper, we propose a novel framework named NFER-SDE on top of the score-based diffusion models to translate VIS facial expression samples to the NIR domain, so that the performance of NIR FER task can be improved with the translated training images. The framework consists of the energy functions of task-specific guidance and the conditional score-matching network, which guide the score-based diffusion models to capture subtle yet key facial representations for FER. Extensive experiments on two existing NIR FER datasets demonstrate the effectiveness of our method.

Acknowledgments

This research is sponsored in part by the NSFC Program (No. U20A6003, 62076146, 62021002, U19A2062, U1911401, 6212780016) and Industrial Technology Infrastructure Public Service Platform Project "Public Service Platform for Urban Rail Transit Equipment Signal System Testing and Safety Evaluation" (No. 2022-233-225), Ministry of Industry and Information Technology of China.

References

- Bulat, A.; and Tzimiropoulos, G. 2017. How far are we from solving the 2D & 3D Face Alignment problem? (and a dataset of 230,000 3D facial landmarks). In *International Conference on Computer Vision*.
- Choi, J.; Kim, S.; Jeong, Y.; Gwon, Y.; and Yoon, S. 2021. Ilvr: Conditioning method for denoising diffusion probabilistic models. *arXiv preprint arXiv:2108.02938*.
- Croitoru, F.-A.; Hondru, V.; Ionescu, R. T.; and Shah, M. 2023. Diffusion models in vision: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Du, S.; Tao, Y.; and Martinez, A. M. 2014. Compound facial expressions of emotion. *Proceedings of the national academy of sciences*, 111(15): E1454–E1462.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1125–1134.
- Kingma, D. P.; Mohamed, S.; Jimenez Rezende, D.; and Welling, M. 2014. Semi-supervised learning with deep generative models. *Advances in neural information processing systems*, 27.
- Li, S.; and Deng, W. 2020. A deeper look at facial expression dataset bias. *IEEE Transactions on Affective Computing*, 13(2): 881–893.
- Lv, C.; Wu, Z.; Wang, X.; and Zhou, M. 2019. 3D facial expression modeling based on facial landmarks in single image. *Neurocomputing*, 355: 155–167.
- Meng, C.; He, Y.; Song, Y.; Song, J.; Wu, J.; Zhu, J.-Y.; and Ermon, S. 2021. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*.
- Nair, N. G.; and Patel, V. M. 2023. T2V-DDPM: Thermal to Visible Face Translation using Denoising Diffusion Probabilistic Models. In *2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG)*, 1–7. IEEE.
- Peng, J.; Liu, D.; Xu, S.; and Li, H. 2021. Generating diverse structure for image inpainting with hierarchical VQ-VAE. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10775–10784.
- Saharia, C.; Ho, J.; Chan, W.; Salimans, T.; Fleet, D. J.; and Norouzi, M. 2022. Image super-resolution via iterative refinement. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4713–4726.
- Sohn, K.; Lee, H.; and Yan, X. 2015. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2020. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*.
- Wang, H.; Zhang, H.; Yu, L.; Wang, L.; and Yang, X. 2020a. Facial Feature Embedded CycleGAN For Vis-Nir Translation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1903–1907. IEEE.
- Wang, K.; Peng, X.; Yang, J.; Lu, S.; and Qiao, Y. 2020b. Suppressing uncertainties for large-scale facial expression recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 6897–6906.
- Wang, L.; Jia, G.; Jiang, N.; Wu, H.; and Yang, J. 2022. EASE: Robust Facial Expression Recognition via Emotion Ambiguity-SEnsitive Cooperative Networks. In *Proceedings of the 30th ACM International Conference on Multimedia*, 218–227.
- Xing, Z.; Tan, W.; He, R.; Lin, Y.; and Yan, B. 2022. Co-Completion for Occluded Facial Expression Recognition. In *Proceedings of the 30th ACM International Conference on Multimedia*, 130–140.
- Yang, L.; Wang, S.; Ma, S.; Gao, W.; Liu, C.; Wang, P.; and Ren, P. 2020. Hifacegan: Face renovation via collaborative suppression and replenishment. In *Proceedings of the 28th ACM international conference on multimedia*, 1551–1560.
- Zhang, L.; and Tjondronegoro, D. 2011. Facial expression recognition using facial movement features. *IEEE transactions on affective computing*, 2(4): 219–229.
- Zhang, R.; Isola, P.; Efros, A. A.; Shechtman, E.; and Wang, O. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 586–595.
- Zhao, G.; Huang, X.; Taini, M.; Li, S. Z.; and Pietikäinen, M. 2011. Facial expression recognition from near-infrared videos. *Image and vision computing*, 29(9): 607–619.
- Zhao, M.; Bao, F.; Li, C.; and Zhu, J. 2022. Egsde: Unpaired image-to-image translation via energy-guided stochastic differential equations. *Advances in Neural Information Processing Systems*, 35: 3609–3623.
- Zhao, Z.; and Liu, Q. 2021. Former-dfer: Dynamic facial expression recognition transformer. In *Proceedings of the 29th ACM International Conference on Multimedia*, 1553–1561.
- Zhu, J.; Luo, B.; Yang, T.; Wang, Z.; Zhao, X.; and Gao, Y. 2023. Knowledge Conditioned Variational Learning for One-Class Facial Expression Recognition. *IEEE Transactions on Image Processing*.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, 2223–2232.