

# In the light of feature distributions: moment matching for Neural Style Transfer

Nikolai Kalischek

Jan D. Wegner

Konrad Schindler

EcoVision Lab, Photogrammetry and Remote Sensing, ETH Zürich

{nikolai.kalischek, jan.wegner}@geod.baug.ethz.ch, schindler@ethz.ch



Figure 1: Style transfer results: artists' portraits rendered in their own painting style by our method.

## Abstract

*Style transfer aims to render the content of a given image in the graphical/artistic style of another image. The fundamental concept underlying Neural Style Transfer (NST) is to interpret style as a distribution in the feature space of a Convolutional Neural Network, such that a desired style can be achieved by matching its feature distribution. We show that most current implementations of that concept have important theoretical and practical limitations, as they only partially align the feature distributions. We propose a novel approach that matches the distributions more precisely, thus reproducing the desired style more faithfully, while still being computationally efficient. Specifically, we adapt the dual form of Central Moment Discrepancy (CMD), as recently proposed for domain adaptation, to minimize the difference between the target style and the feature distribution of the output image. The dual interpretation of this metric explicitly matches all higher-order centralized moments and is therefore a natural extension of existing NST methods that only take into account the first and second moments. Our experiments confirm that the strong theoretical properties also translate to visually better style transfer, and better disentangle style from semantic image content.*

## 1. Introduction

In 2017 *Loving Vincent* was released, the first fully painted feature film with >65,000 frames. Indeed, every

single frame is an oil painting drawn by one of over 100 artists. The creation of the movie was split into two steps. First, the entire movie was produced with real actors in front of a green screen, which was then replaced by Van Gogh paintings. In a second step, each frame was painted over by an artist with the techniques and style of Van Gogh, which took over six years to complete.

Attempts to automate this form of texture synthesis, termed *style transfer*, date back to at least the mid-90s [12]. More recently, Gatys *et al.* [9] pioneered the idea of *Neural Style Transfer (NST)*. It is based on the idea that the deep layers of a pre-trained Convolutional Neural Network (CNN) encode high-level semantic information and are insensitive to the actual appearance, whereas shallow layers learn low-level features such as color, texture and brush patterns. A fundamental question that arises in this context is how to define style. Li *et al.* [25] proved that the loss introduced in [9] can be rewritten as a Maximum Mean Discrepancy (MMD), offering an interpretation of style transfer as aligning feature distributions. In fact, most existing methods can be interpreted in this way. This has led to a series of works all centered around aligning feature distributions of CNNs, linking style transfer to Domain Adaptation (DA). Here we look deeper into that interpretation. By translating NST to distribution matching, it becomes amenable to a suite of tools developed to measure the divergence between probability distributions, such as integral probability metrics,  $f$ -divergences and Optimal Transport (OT).

Divergences  $d(P, Q)$  between two distributions, respectively probability measures, are in general not metrics, but they should fulfil the weaker conditions of (i) non-negativity:  $d(P, Q) \geq 0$ ; and (ii) identity of indiscernibles:  $d(P, Q) = 0$  iff  $P = Q$ . However, *in the light of feature distributions, existing style transfer methods suffer from rather elementary theoretical limitations*. Broadly, there are two schools. Either the distributions are unrestricted, but the discrepancy between them is measured without adhering to the law of indiscernibles [9, 25, 15, 32]; or the distributions are approximated roughly with simple functions, so that they admit closed-form solutions [29, 19, 24, 27].

Here, we show how to overcome these limitations with the help of the recently proposed framework of Central Moment Discrepancy (CMDs) [39]. That (pseudo-)metric is based on the representation of distributions as moment sequences on compact intervals. In the limit, CMD is an integral probability metric on the set of compactly supported distributions, so it complies with the law of indiscernibles (as well as non-negativity) by definition. Importantly, in its dual formulation the CMD is computationally efficient, and approximations can be seamlessly justified with an upper bound on the central moments [38]. In summary, we make the following contributions: (i) We systematically categorize existing NST methods according to their way of aligning distributions; (ii) we make explicit underlying approximations and highlight the corresponding limitations; (iii) We propose a novel NST algorithm based on the Central Moment Discrepancy. To our knowledge, our method is the first one that aligns style distributions in a rigorous and computationally efficient manner, with theoretically grounded approximation bounds. Empirically, the method achieves a more perspicuous separation between artistic style and semantic content, and enables visually more compelling style transfer according to a user study with >50 participants.

## 2. Related work

**Style Transfer** has been an active research topic in computer vision for at least two decades. Until recently it was based on hand-crafted features and styles. This includes stroke-based rendering [20] to repaint an image with a set of brush strokes [13], image quilting [5] where texture is synthesized in small patches according to a segmentation map, or image analogies [14] that learn style filters in a supervised fashion. The shift to CNNs has given rise to Neural Style Transfer. Current NST techniques can be categorized as being based on either image optimization or model optimization [16]. Methods in the first group iteratively transfer style to each new output image, following the seminal paper of [9]. That work first introduced the idea to match feature statistics of intermediate layers in a CNN. Subsequent

works explored different directions to improve the quality of stylization. Risser *et al.* [32] circumvent instabilities of the optimization by incorporating additional histogram and total variation losses. To further enhance the preservation of low-level content such as edges, Li *et al.* [22] add a Laplacian loss. In order to transfer style between semantically matching patches (*e.g.*, from eyes of a dog to eyes of a cat), [28] defines a loss that compares regions with similar semantic meaning. Similarly, [21] use MRFs to find the nearest-neighbor patch in the feature space of the style image. Both require similar shapes and boundaries in the content and style images. Gatys *et al.* [10] also went on to add user-control for perceptual factors such as color or scale, *e.g.*, by transferring style only in the luminance channel to preserve color. Recently, Kolkin *et al.* [19] also incorporate user-defined spatial constraints, via appropriate weights in the cost function.

Iterative optimization per image is comparatively slow. Model optimization methods instead employ feed-forward networks [17, 36] trained offline on large datasets, to achieve real-time style transfer. Initially they were restricted to a fixed set of styles [35, 36, 4, 23]. Later they were extended to handle unseen styles. Huang and Belongie [15] propose an adaptive instance normalization layer that normalizes the content image with affine parameters from the style image, Chen and Schmidt [2] define a swap layer that replaces content feature patches with matching style feature patches. However, there is a price to pay for fast feed-forward inference, as it does not reach the quality of iterative methods. Recently it has been shown that adaptive instance normalization, as well as the whitening color transform [24] are special cases of an OT map between Gaussian distributions, thus providing some theoretical foundation for feed-forward models [27, 29].

**Domain Adaptation** is a particular instance of transfer learning, *i.e.*, distilling and transferring knowledge across different domains. Domain Adaptation (DA) utilizes supervision in a source domain to guide the learning for a target domain where no labeled data is available [3]. The principle is that the shift between the source and target domains can be measured, and therefore also minimized. Several authors have noted the close relation to NST [25, 1]. A common approach is to learn a joint feature space by aligning the distributions in the latent feature space with measures such as Kullback-Leibler divergence [40], Maximum Mean Discrepancy [26] or correlation alignment [34]. Also related to style transfer, another approach to DA is to directly learn the mapping between the source and target domains, *e.g.*, using GANs [1]. For an overview of DA, see [3, 37]. Here we make use of yet another idea originally aimed at DA, emphasizing its close relation to style transfer.

### 3. Method

We first briefly review the core ideas of *Neural Style Transfer*. In that context, we revisit several existing methods and classify them into three categories. By taking the view of distribution alignment to its logical end, we then go on to provide an alternative loss function that has strong theoretical guarantees, is efficient to compute, and delivers visually appealing results (c.f. Fig. 1).

#### 3.1. Neural style transfer

The fundamental idea of NST is to use a pretrained, deep neural network to generate an image  $I_o$  with the content-specific features of a content image  $I_c$  and the style-specific features from a style image  $I_s$ . Typically, one minimizes a convex combination of a content and a style loss:

$$\mathcal{L} = \alpha \mathcal{L}_{\text{content}} + (1 - \alpha) \mathcal{L}_{\text{style}}. \quad (1)$$

We further specify those losses following the notation of [29]. Let  $g$  be a deep encoder, say VGG-19 [33]. For a specific layer  $l$  with corresponding output feature map of spatial dimension  $H_l \cdot W_l = n_l$  and channel depth  $C_l$ , we denote the  $j^{\text{th}}$  component of the feature map as a (reshaped) function  $F_j^l : \mathbb{R}^d \rightarrow \mathbb{R}^{C_l}$ ,  $j \in [n_l]$ . We write  $\mathbf{F}^l = (F_j^l)_{j \in [n_l]} \in \mathbb{R}^{C_l \times n_l}$  and call  $\mathbf{F}^l(I)$  the  $l^{\text{th}}$  (reshaped) feature map of image  $I$ . I.e., the  $L^{\text{th}}$  feature map of image  $I$  is the activation map after applying all layers  $l = 1, \dots, L$  to  $I$ . Then, the content loss is proportional to

$$\mathcal{L}_{\text{content}}(I_o, I_c) \propto \sum_l \|\mathbf{F}^l(I_o) - \mathbf{F}^l(I_c)\|^2, \quad (2)$$

where  $l$  iterates over a set of layers of  $g$ . Commonly, only a single, deep layer is used to compute the content loss; whereas the style loss is an average over multiple layers, shallow and deep, with hyper-parameters  $w_l$ :

$$\mathcal{L}_{\text{style}}(I_o, I_s) = \sum_l w_l \mathcal{L}_{\text{style}}^l(I_o, I_s). \quad (3)$$

#### 3.2. Style as feature distribution

Losses proposed for  $\mathcal{L}_{\text{style}}^l$  can be categorized according to how they align distributions. We first need some additional definitions, again following [29]. To obtain a distribution, we view the feature map  $\mathbf{F}^l(I)$  as a  $C_l$ -dimensional empirical distribution measure over  $n_l = H_l \cdot W_l$  samples. Note, by regarding the  $n_l$  samples as an unordered set we explicitly discard the spatial layout. This corresponds to the intuition that style attributes like color, strokes and texture are independent of the location. More formally, we define

$$\nu^l : \mathbb{R}^d \rightarrow \mathcal{P}(\mathbb{R}^{C_l}) \quad , \quad I \mapsto \frac{1}{n_l} \sum_{i=1}^{n_l} \delta_{F_j^l(I)}, \quad (4)$$

where  $\mathcal{P}(\mathbb{R}^{C_l})$  is the space of empirical measures on  $\mathbb{R}^{C_l}$ . We abbreviate  $\nu_j^l = \nu^l(I)$  and drop the layer index when not needed. With these definitions we now review existing style transfer methods in the light of distribution alignment.

**MMD-based optimization.** Already the first NST paper [9] used statistics of feature maps to extract style-specific attributes of  $I_s$ , via the Gram matrix  $G$ . The Gram matrix contains 2<sup>nd</sup>-order statistics, in our case correlations between corresponding channels in the feature map. The link to aligning distributions may not be obvious, but Li *et al.* [25] show that the style loss in [9] can be rewritten as an unbiased empirical estimate of the MMD [11] with a polynomial kernel  $k(x, y) = (x^T y)^2$ :

$$\mathcal{L}_{\text{style}}^l(I_o, I_s) \propto \text{mmd}^2[\mathbf{F}^l(I_o), \mathbf{F}^l(I_s)]. \quad (5)$$

Under the assumption that the reproducing kernel Hilbert space (RKHS) is *characteristic* [7], the MMD vanishes if and only if the two distributions are the same. By treating the feature maps of  $I_o$  and  $I_s$  as samples, minimizing the objective (5) is the same as minimizing the discrepancy between  $\nu_{I_o}$  and  $\nu_{I_s}$ .

**Moment-based optimization** approaches explicitly minimize the difference between style distributions. Theoretical support for these methods comes from moment generating functions (MGFs). It is known that a distribution is uniquely characterized by its moments if the MGF is finite in an open interval containing zero. Hence, if two distributions with finite MGF have equal moments, they are identical.

Besides relating style transfer to distribution alignment, Li *et al.* [25] also introduced a style loss based on *batch normalization* statistics. That loss is the first to explicitly match moments in feature space, namely the means  $\mu_{\mathbf{F}^l(I)}$  and the standard deviations  $\sigma_{\mathbf{F}^l(I)}$ :

$$\mathcal{L}_{\text{style}}^l(I_o, I_s) \propto \sum_{i=1}^{C_l} [(\mu_{\mathbf{F}^l(I_o)}^i - \mu_{\mathbf{F}^l(I_s)}^i)^2 + (\sigma_{\mathbf{F}^l(I_o)}^i - \sigma_{\mathbf{F}^l(I_s)}^i)^2]. \quad (6)$$

Interestingly, moment alignment can also produce reasonable results when applied in feed-forward mode, without iterative optimization. Based on ideas from [36, 4], Huang and Belongie [15] align the mean and variance with a transformation layer. In summary, matching the mean and variance of the content image's feature space to that of the style image reduces the divergence between  $\nu_{I_o}$  and  $\nu_{I_s}$  – but discrepancies due to higher-order moments remain.

**Optimal Transport-based optimization** provides a principled framework to minimize the discrepancy between distributions, notably taking into account the geometry of the underlying spaces. When working in the space of probability measures  $\mathcal{P}_p(\mathbb{R}^d)$  with bounded  $p^{\text{th}}$  moment, the *Wasser-*




Figure 2: Illustration of distribution matching in 1D. The source  $\sim \text{Beta}(2, 3)$  and target  $\sim \text{Beta}(0.5, 0.45)$  cannot be aligned with MMD, MM or OT (which in 1D is the same as MM). On the contrary, CMD aligns them well already with five moments, and the residual error decreases asymptotically as more moments are added. See text for details.

stein distance for  $P, Q \in \mathcal{P}_p(\mathbb{R}^d)$  is defined as

$$\mathcal{W}_p(P, Q)^p = \inf_{\Gamma(P, Q)} \int \|x - y\|^p d\pi(x, y). \quad (7)$$

We can use the Wasserstein distance for back-propagation to minimize the discrepancy between  $\nu_{I_o}$  and  $\nu_{I_s}$ . In general, computing the OT has complexity  $O(n_l^3 \log n_l)$  and is not suitable for iterative optimization schemes. However, restricting the distributions to Gaussians,  $\tilde{\nu}_{I_o} := \mathcal{N}(\mu_{\nu_{I_o}}, \Sigma_{\nu_{I_o}})$  and  $\tilde{\nu}_{I_s} := \mathcal{N}(\mu_{\nu_{I_s}}, \Sigma_{\nu_{I_s}})$  admits a closed form solution,

$$\begin{aligned} \mathcal{W}_2(\tilde{\nu}_{I_o}, \tilde{\nu}_{I_s})^2 &= \|\mu_{\nu_{I_o}} - \mu_{\nu_{I_s}}\|_2^2 \\ &+ \text{Tr} \left( \Sigma_{\nu_{I_o}} + \Sigma_{\nu_{I_s}} - 2(\Sigma_{\nu_{I_o}}^{\frac{1}{2}} \Sigma_{\nu_{I_s}} \Sigma_{\nu_{I_o}}^{\frac{1}{2}})^{\frac{1}{2}} \right). \end{aligned} \quad (8)$$

This is similar to matching the first and second moments as in moment-based optimization (higher-order moments of Gaussians are constant w.r.t. mean and variance). Conveniently, the OT map can also be directly derived. If one is willing to accept the Gaussian approximation, the style features can be aligned by iteratively minimizing  $\mathcal{W}_2$ , or by integrating the OT map into the encoder-decoder network [29, 19, 27, 24]. It has been shown [29, 27] that adaptive instance normalization can be seen as OT of Gaussians with diagonal covariances.

### 3.3. Motivation

From a statistical perspective all three categories of methods contradict, to some extent, the goal of optimally aligning feature distributions.

Methods based on **MMD** rely on simplistic (typically, linear or quadratic) kernels [9, 25]. Previously, [32] already identified instabilities during training, as different distributions result in the same MMD. They point out that changes in mean and variance can compensate each other, giving rise to the same Gram matrix (and thus the same MMD with quadratic kernel), since the Gram matrix is related to *non-central* second moments. We offer an alternative explanation

why the Gram matrix violates the identity of indiscernibles: the quadratic kernel is non-characteristic, *i.e.*, the map  $p \rightarrow \mathbb{E}_{x \sim p}[k(x, \cdot)]$  is not injective and the distribution  $p$  has no unique embedding in the RKHS. Moreover, the quadratic kernel (resp. Gram matrix) is obviously restricted to 2<sup>nd</sup> moments. It is highly unlikely that those are sufficient statistics for deep feature activations, so  $\text{MMD}(p, q) = 0$  almost certainly does not imply  $p = q$ .

A similar argument can be made about existing methods based directly on **Moment Matching (MM)**, since they match only the means and variances. It is trivial to define two distinct distributions with the same variances – *e.g.*, a Gaussian  $\mathcal{N}(0, \sqrt{2})$  and a Laplace distribution  $\mathcal{L}(0, 1)$ .

While **OT** is a powerful framework at the conceptual level, it is hobbled by high computation cost. The Gaussian approximation makes OT tractable, but at the cost of losing information. There is no evidence that the distributions  $\nu_{I_o}$  and  $\nu_{I_s}$  are (approximately) Gaussian – in fact it is very unlikely, unless one artificially constrains them, thus seriously restraining the deep network’s expressive power. We claim that OT, at least in its prevalent, restricted form, also mostly reduces to matching the first and second moments – the approximations in (8) are completely defined in terms of means and covariances.

Finally, we point out the *mean over-penalization* effect: [39] found instabilities of distribution alignment during DA training under small perturbations, which arise from the use of raw instead of centralized moments (as in MMD with standard polynomial kernel and non-centralized integral probability metrics). For details, please refer to [39].

### 3.4. CMD for neural style transfer

Instead of only matching first- and second-order moments, we propose to make use of a suitable integral probability metric, the Central Moment Discrepancy [38]. At its core, that metric utilizes the dual representation of compactly supported distributions as moment sequences. The translation to central moments leads to natural geometric relations such as variance, skewness and kurtosis. Not that

the idea of matching higher moments has been investigated in early work on texture synthesis [31], but so far has been disregarded in NST.

In Fig. 2, we illustrate the enhanced expressive power of CMD. In our toy example, the source and target are univariate *Beta*-distributions with different parameters, *i.e.*, their third and fourth moments are non-zero. We represent each distribution with 10,000 samples and minimize the respective alignment loss with gradient descent. The example confirms that none of the three approaches based on first and second moments can align the two distributions (note that for the 1D case MM and OT are identical). On the contrary, CMD aligns them nicely.

The CMD between two compactly supported distributions  $P$  and  $Q$  is defined as follows [39]:

$$\text{cmd}_k(P, Q) := \sum_{i=1}^k a_i \|c_i(P) - c_i(Q)\|_2 \quad , \quad \text{where} \quad (9)$$

$$c_i(X) = \begin{cases} \mathbb{E}_X[x] & i = 1 \\ \mathbb{E}_X[\eta^{(i)}(x - \mathbb{E}_X[x])] & i \geq 2 \end{cases}$$

with  $a_i \geq 0$ . The  $\eta^{(i)}(x)$  are monomial vectors of order  $i$  defined as

$$\eta^{(i)} : \mathbb{R}^m \rightarrow \mathbb{R}^{\frac{(i+1)^m - 1}{(m-1)!}} \quad (10)$$

$$x \mapsto (x_1^{r_1} \cdots x_m^{r_m})_{\substack{(r_1, \dots, r_m) \in \mathbb{N}_0^m \\ r_1 + \dots + r_m = k}}$$

By construction the CMD is non-negative, respects the triangle inequality, and if  $P = Q$  then  $\text{cmd}_k(P, Q) = 0$ . Furthermore, [38, Theorem 1] states that  $\text{cmd}_k(P, Q) = 0$  implies  $P = Q$  for  $k \rightarrow \infty$ , so CMD is a metric on compactly supported distributions.

For practical applications computing  $\text{cmd}_\infty$  is obviously not possible, and we have to bound  $k$  to  $K < \infty$  from above. Compared to other approximations used for style transfer [29, 19], the bounded  $\text{cmd}_K$  has a natural theoretical justification. It can be shown [39, Proposition 1] that the  $i^{\text{th}}$  term in the summation of equation 9 is bounded by an upper bound that strictly decreases with the order  $i$ . *I.e.*, the contribution of higher-order moment terms in equation (9) converges monotonically to 0. To keep the implementation efficient we only compute the marginal moments, by restricting the monomial vectors to  $\eta^{(i)}(x) = (x_1^i, \dots, x_m^i)$ .

Adapting CMD to our style feature distributions is straight-forward. To fulfill the requirements, we wrap a sigmoid function  $\sigma(\cdot)$  around each feature output so as to restrict the support of the empirical distribution to  $[0, 1]$ . With a slight abuse of notation we write  $\sigma(\nu^l)$  for the  $\nu^l$  computed from sigmoid-transformed features and define

$$\mathcal{L}_{\text{style}}^l(I_o, I_s) := \text{cmd}_k(\sigma(\nu_{I_o}^l), \sigma(\nu_{I_s}^l)), \quad (11)$$

for layer  $l$ . The moments are simply the moments of the empirical measure, *i.e.* powers of  $\mathbb{E}[\mathbf{F}^l(I) - \mu_{\mathbf{F}^l(I)}] \in \mathbb{R}^{C_l}$ . By adopting CMD we have an integral probability metric for NST at our disposal that not only has favourable theoretical properties, but is also easy to implement, computationally efficient, and able to handle complex feature distributions with significant higher-order moments.

## 4. Results

In this section, we compare our results with existing methods from each of the categories. After summarizing details of the implementation, we qualitatively evaluate the effects of aligning the style features with CMD. Beyond visual comparisons, we report quantitative results from an user-study, which supports our hypothesis that higher-order moments carry important style information and should not be ignored. Lastly, we further investigate the impact of different moments in an ablation study.

### 4.1. Experimental setup

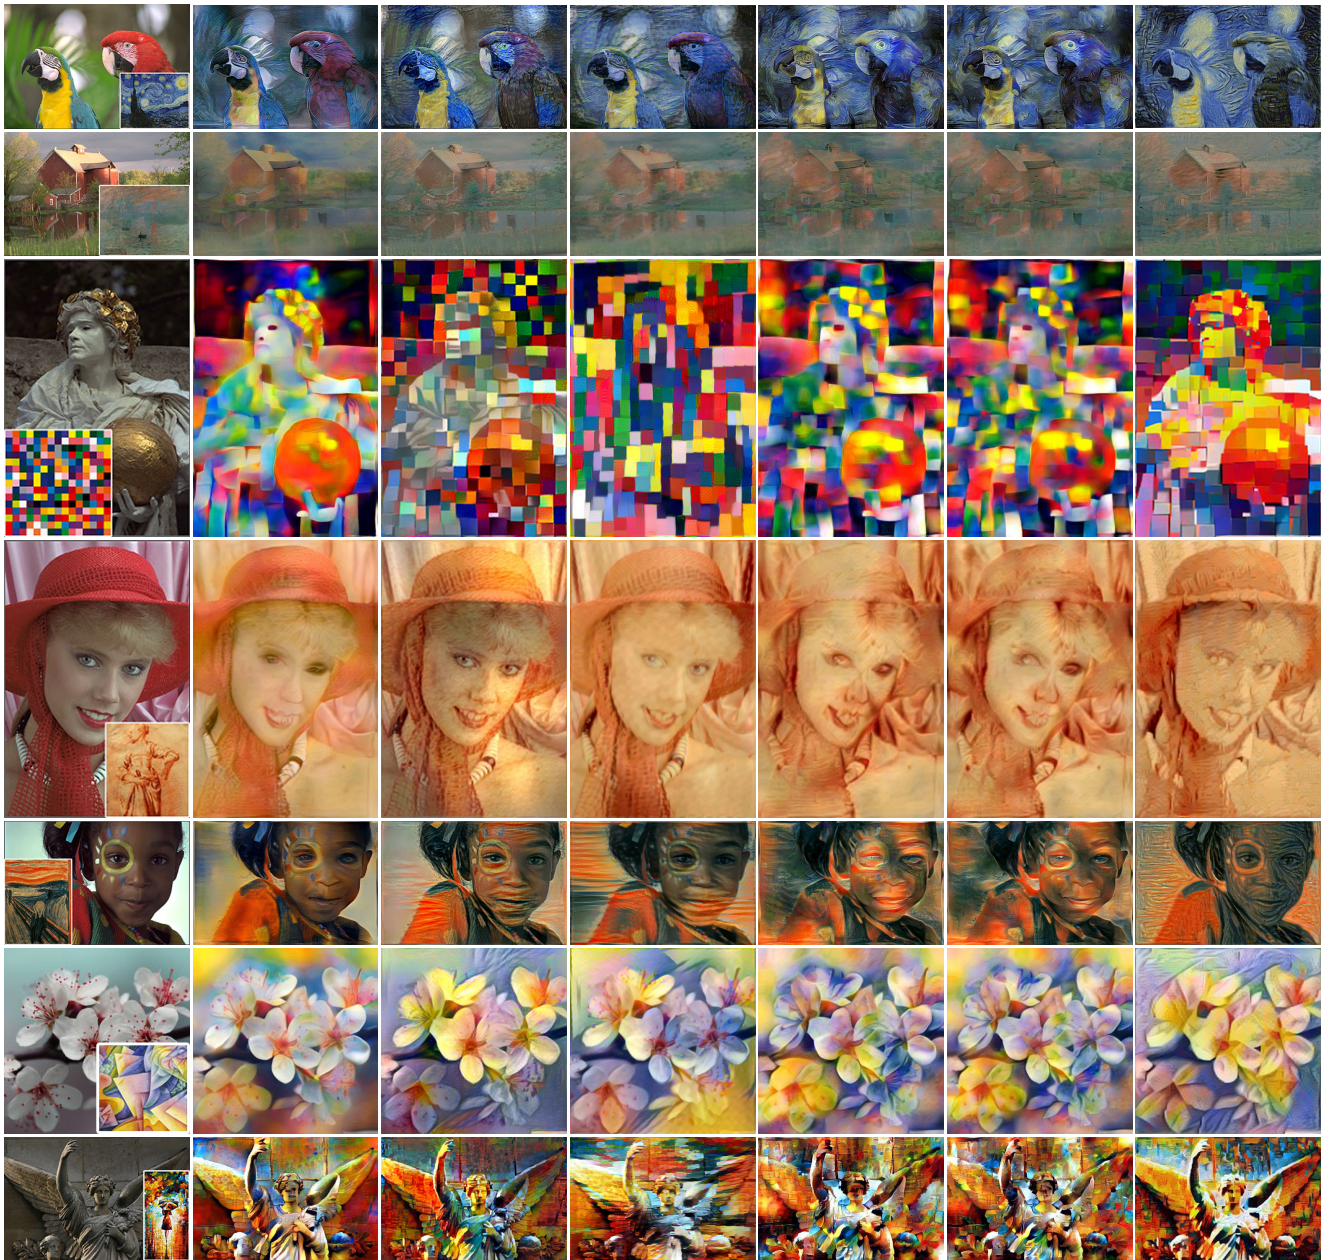
We employ VGG-19 [33] as feature encoder and read out feature maps at layer levels  $l \in \{1\_1, 2\_1, 3\_1, 4\_1, 5\_1\}$ . Deviating slightly from the commonly used NST setting, we work with the raw convolution outputs *conv-l* rather than their rectified versions *relu-l*, since we clamp them to  $[0, 1]$  with sigmoid activations for computing the CMD, see (11). The content loss is computed on *conv4\_1*, for the individual layers in the style loss we use the same weighting scheme as proposed in [9]. Optimization is performed with Adam [18]. Instead of blindly stopping after a fixed number of iterations, we implement a stopping criterion based on the difference of the current style loss and a moving average of the style loss. We compare our algorithm to five baselines: one from the MMD group [9], two based on direct moment differences [25, 15] and two based on OT [24, 27]. We use the existing open-source implementations<sup>1</sup> and keep all hyper-parameters as proposed in the original papers, respectively source codes. Our implementation is based on PyTorch [30] and is also publicly available.<sup>2</sup> For our experiments we bound the order of the moments to  $K=5$ , as higher orders have little influence.

### 4.2. Qualitative results

We have pinpointed theoretical limitations of previous NST methods in Sec. 3.3. To see how these translate to concrete visual differences, we analyze how well the stylized images preserve three different style attributes, *color, texture and stroke, shape*. See Fig. 3, and further results in the supplementary material.

<sup>1</sup>For [9, 25, 27], original implementations by the authors; for [15, 24], implementation provided by the authors of [27].

<sup>2</sup>Code: [https://github.com/D1noFuZi/cmd\\_styletransfer](https://github.com/D1noFuZi/cmd_styletransfer)



(a) Input (b) AdaIN [15] (c) Gatys [9] (d) MM [25] (e) OST [27] (f) WCT [24] (g) Ours

Figure 3: Style transfer results of our algorithm and of previous methods from all three categories. Best viewed on screen. Please zoom in to appreciate style details.

**Color and brightness.** This paper is concerned with fully automatic NST, without additional user control. Hence, the output should have the color palette of the style image. *I.e.*, only the semantic content of the content image should be retained, but colors should be replaced by those representative of the style, and in particular the two color spaces should not be mixed. Looking at the 1<sup>st</sup> row of Fig. 3, the red of the right parrot strongly leaks into the results of AdaIN,

Gatys and MM, and traces are also visible in WCT. Besides our method, those based on OT fare best in terms of color palette, but OT has a tendency towards exaggerated brightness variations not warranted by the content, *e.g.*, the girl’s face in row 5 and the background in row 6. Indeed, it appears that local color and intensity information is to some degree hidden in higher-order moments. That observation is also supported by the ablation study in Sec. 4.4.

**Texture and stroke.** Maintaining strokes and textures is especially important when it comes to artistic style transfer, to preserve the concomitant individual painting techniques. We find that the proposed CMD method is particularly good at replicating granular canvas, oriented brush strokes, *etc.* Clear cases in point are rows 1 and 5 of Fig. 3, as well as the reflections on the lake in row 2. We also point out the particularly challenging example in the 4<sup>th</sup> row. Zooming in on the style image, we can see the rough texture of the paper, as well as a preference for oriented shading strokes. While none of the methods is perfect on this difficult instance, the only ones to even partially pick up those patterns are our method and to some degree Gatys (but with strong color artifacts). In general, we observe that oriented high-frequency patterns appear to benefit from higher (particularly, odd) moments, but further research is needed to explore the relation in depth.

**Shape.** Lastly, we turn our attention to shape. That attribute is somewhat more complex, as ornamental and decorative shape elements such as the square pattern in row 3 of Fig. 3 are part of the style, whereas semantically meaningful elements of similar size are part of the content, like the eyes in row 4 or the make-up in row 5. CMD manages to disentangle these two aspects and preserve important boundaries and details of the content rather well, while still imposing the characteristic shape features of the style. Perhaps the most convincing example is row 3. But also in other cases the delicate balance between imposing the style and preserving salient content features appears to benefit from higher-order moments, *e.g.*, rows 4, 5, 6.

### 4.3. Quantitative results

**User study.** There is no clear consensus how to quantitatively evaluate NST. The question what constitutes a “correct” output is clearly ill-posed, and even the judgment how “good” a given stylization is depends on aesthetic preferences and must remain subjective. In fact one can, with the *same* method, generate very different results only by changing the relative weights of the style and content losses, and it depends on the application and on personal taste which one is preferred.

The current consensus is to perform user studies where participants are shown results without revealing how they were generated, and to collect statistics of user preferences. We note that, while we agree that aesthetic quality is hard to measure, people can usually pick their favorite among a handful of alternative stylizations without much hesitation, which lends some support to these studies: at the very least, they are a guideline which one among the available methods will deliver the result that the relatively largest share of the user group likes best. We conduct a user study with the same methods as above: AdaIN [15], Gatys [9], Moment Matching [25], OST [27], WCT [24] and the pro-

posed CMD method. The study uses parts of the Kodak image dataset [6] and additional content images widely used in NST, showing a variety of scenes, objects and humans. The style dataset is made up by paintings and drawings commonly used for NST, from a range of artists including Picasso, Kandinsky, Van Gogh and others. In total we exhaustively combine 31 content images and 20 style images, resulting in 620 stylized images per algorithm. For the study, the six stylization results were displayed side-by-side in random order, along with the underlying content and style images. Users were asked to pick a single image that would best transfer style aspects such as shape, textures and colors using their own judgement.

Overall, we have collected >2700 votes from 56 different participants. The scores are reported in Tab. 1. The study reveals some interesting insights. Indeed, our proposed CMD method performs favorably, with  $\approx 10\%$  more votes than the closest competitor. The classical NST of [9] attains the second-highest number of votes. This supports our claim that iterative methods still have an edge in terms of quality, as one-shot approaches trade quality for speed.

AdaIN*	Gatys	MM	OST*	WCT*	Ours
155	533	443	523	463	<b>587</b>
5.7%	19.7%	16.3%	19.3%	17.1%	<b>21.7%</b>

Table 1: Number of votes each method received in our user study. \* denotes one-shot feed-forward methods.

### 4.4. Ablation studies

In our method it is possible to individually reweight or turn off moments. We have conducted an ablation study to better understand the effects of different moments, see Fig. 4. Note that this tuning knob is orthogonal to user control in the spirit of [10], where one isolates a specific attribute like color in preprocessing and applies the stylization selectively. Figure 4 shows style transfer results with different combinations of moments. Only a single moment corresponding to the row/column index is used on the diagonal. Then higher-order moments are progressively added along the rows, so for instance position (2, 2) corresponds to only the second moment (weight vector  $a = [0, 1, 0, 0, 0]$ ) and element (2, 4) corresponds to the the 2<sup>nd</sup>, 3<sup>rd</sup> and 4<sup>th</sup> moments (weight vector  $a = [0, 1, 1, 1, 0]$ ). As was to be expected there is no obvious, “pure” correspondence between moments and visual attributes. Still, the study illustrates some interesting relations. First, one can immediately see that even the 5<sup>th</sup> order still contributes significant style elements, for instance on the chin and the cap in the first row. Odd moments appear to primarily modulate overall brightness and contrast, whereas even ones tend to change colors and high-frequency texture.



Figure 4: Ablation study using only selected moments. See text for details.

Our CMD method changes only the loss function for distribution alignment and can be seamlessly combined with other extensions of NST. For instance, the user can still control how strongly the style is imprinted on the image content, by adjusting the relative weight of the style and content losses. To illustrate this, we stylize with our CMD method and linearly interpolate the weight  $\alpha$  in eq. (1). Figure 5 shows an example how putting more weight on the content loss produces increasingly weaker "partial stylizations" that stay closer to the content image.

## 5. Limitations and future work

There are currently two conceptual directions in NST: iterative optimization techniques and one-shot feed-forward approaches. Our algorithm belongs to the former. While iterative methods arguably still produce better results, they are too slow for real-time applications. Our method inherits that shortcoming, *e.g.*, it could not be used for (near) real-time video synthesis.

At the conceptual level, we had to make two simplifying approximations to take the step from the mathematical formalism of CMD to a practical implementation. On the one hand, we limit the order of the central moments to a finite, in practice small  $K$ . At least in principle the impact of




Figure 5: Varying the strength of style transfer by varying the relative influence  $\alpha$  of the content loss (*c.f.* eq. (1)).

that restriction can be kept as small as desired by increasing  $K$ , because the influence of additional central moments provably converges  $\rightarrow 0$  with increasing order.

On the other hand, and perhaps more importantly, we only utilize the *marginal* central moments in our loss. We take this shortcut for computational reasons, but it effectively means that we only achieve exact distribution matching when the marginal distributions are independent. There is currently no evidence that this is the case, and we do not see a simple way to gauge how much information might be lost due to the approximation.

## 6. Conclusion

We have revisited the interpretation of Neural Style Transfer as aligning feature distributions. After categorizing existing methods into three groups based on MMD, moment matching and OT, we show that all of them, in practice, only match first and second moments. We then went on to propose a novel approach based on Central Moment Discrepancy. Our method can be interpreted alternatively as minimizing an integral probability metric, or as matching all central moments up to a desired order. Our method has both theoretical and practical benefits. In terms of theory it comes with strong approximation guarantees. On the practical side it offers a computationally efficient way to account for higher-order moments of complex feature distributions, and achieves visually better transfer of many artistic styles. On a broader scale, even though Portilla and Simoncelli proposed higher order matching to texture synthesis [31], Gatys *et al.* [8, 9] disregarded all but second-order moments when pioneering Neural Style Transfer. In this regard, our method reintroduces higher order matching to NST.



## References

- [1] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *CVPR*, 2017.
- [2] Tian Qi Chen and Mark Schmidt. Fast patch-based style transfer of arbitrary style. *arXiv preprint arXiv:1612.04337*, 2016.
- [3] Gabriela Csurka. Domain adaptation for visual applications: A comprehensive survey. *arXiv preprint arXiv:1702.05374*, 2017.
- [4] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*, 2016.
- [5] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *ACM SIGGRAPH*, 2001.
- [6] Rich Franzen. Kodak lossless true color image suite. <http://r0k.us/graphics/kodak/>.
- [7] Kenji Fukumizu, Arthur Gretton, Xiaohai Sun, and Bernhard Schölkopf. Kernel measures of conditional dependence. In *NeurIPS*, 2008.
- [8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. *arXiv preprint arXiv:1505.07376*, 2015.
- [9] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *CVPR*, 2016.
- [10] Leon A Gatys, Alexander S Ecker, Matthias Bethge, Aaron Hertzmann, and Eli Shechtman. Controlling perceptual factors in neural style transfer. In *CVPR*, 2017.
- [11] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *Journal of Machine Learning Research*, 13(1):723–773, 2012.
- [12] David J Heeger and James R Bergen. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 229–238, 1995.
- [13] Aaron Hertzmann. Painterly rendering with curved brush strokes of multiple sizes. In *ACM SIGGRAPH*, 1998.
- [14] Aaron Hertzmann, Charles E Jacobs, Nuria Oliver, Brian Curless, and David H Salesin. Image analogies. In *ACM SIGGRAPH*, 2001.
- [15] Xun Huang and Serge Belongie. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*, 2017.
- [16] Yongcheng Jing, Yezhou Yang, Zunlei Feng, Jingwen Ye, Yizhou Yu, and Mingli Song. Neural style transfer: A review. *IEEE Transactions on Visualization and Computer Graphics*, 2019.
- [17] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *ECCV*, 2016.
- [18] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [19] Nicholas Kolkin, Jason Salavon, and Gregory Shakhnarovich. Style transfer by relaxed optimal transport and self-similarity. In *CVPR*, 2019.
- [20] Jan Eric Kyprianidis, John Collomosse, Tinghuai Wang, and Tobias Isenber. State of the ”art”: A taxonomy of artistic stylization techniques for images and video. *IEEE Transactions on Visualization and Computer Graphics*, 19(5):866–885, 2012.
- [21] Chuan Li and Michael Wand. Combining Markov random fields and convolutional neural networks for image synthesis. In *CVPR*, 2016.
- [22] Shaohua Li, Xinxing Xu, Liqiang Nie, and Tat-Seng Chua. Laplacian-steered neural style transfer. In *ACM Multimedia*, 2017.
- [23] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Diversified texture synthesis with feed-forward networks. In *CVPR*, 2017.
- [24] Yijun Li, Chen Fang, Jimei Yang, Zhaowen Wang, Xin Lu, and Ming-Hsuan Yang. Universal style transfer via feature transforms. In *NeurIPS*, 2017.
- [25] Yanghao Li, Naiyan Wang, Jiaying Liu, and Xiaodi Hou. Demystifying neural style transfer. *arXiv preprint arXiv:1701.01036*, 2017.
- [26] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *ICML*, 2017.
- [27] Ming Lu, Hao Zhao, Anbang Yao, Yurong Chen, Feng Xu, and Li Zhang. A closed-form solution to universal style transfer. In *ICCV*, 2019.
- [28] Roey Mechrez, Itamar Talmi, and Lihi Zelnik-Manor. The contextual loss for image transformation with non-aligned data. In *ECCV*, 2018.
- [29] Youssef Mroueh. Wasserstein style transfer. In *AISTATS*, 2020.
- [30] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019.
- [31] Javier Portilla and Eero P Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *International journal of computer vision*, 40(1):49–70, 2000.
- [32] Eric Risser, Pierre Wilmot, and Connelly Barnes. Stable and controllable neural texture synthesis and style transfer using histogram losses. *arXiv preprint arXiv:1701.08893*, 2017.
- [33] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [34] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *ECCV*, 2016.
- [35] Dmitry Ulyanov, Vadim Lebedev, Andrea Vedaldi, and Victor S Lempitsky. Texture networks: Feed-forward synthesis of textures and stylized images. In *ICML*, 2016.
- [36] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Instance normalization: The missing ingredient for fast stylization. *arXiv preprint arXiv:1607.08022*, 2016.

- [37] Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- [38] Werner Zellinger, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. Central moment discrepancy (CMD) for domain-invariant representation learning. *arXiv preprint arXiv:1702.08811*, 2017.
- [39] Werner Zellinger, Bernhard A Moser, Thomas Grubinger, Edwin Lughofer, Thomas Natschläger, and Susanne Saminger-Platz. Robust unsupervised domain adaptation for neural networks via moment alignment. *Information Sciences*, 483:174–191, 2019.
- [40] Fuzhen Zhuang, Xiaohu Cheng, Ping Luo, Sinno Jialin Pan, and Qing He. Supervised representation learning: Transfer learning with deep autoencoders. In *IJCAI*, 2015.