

A Flexible Convolutional Solver with Application to Photorealistic style transfer

Gilles Puy

Technicolor, Cesson-Sévigné, France

Joint work with P. Pérez

Deep Learning: From theory to applications
Cesson-Sévigné, September 06, 2018



Neural style transfer

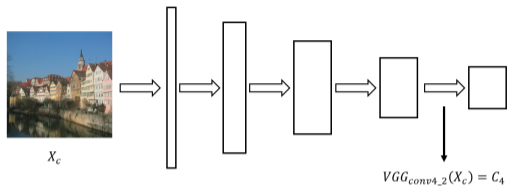
Neural style transfer

Neural style transfer [Gatys *et al.*, CVPR, 2016]

- ▶ Iterative updates to fit statistics of style image and content of target image
- ▶ Statistics: Gram matrices of feature maps at different layers of a pre-trained deep CNN.

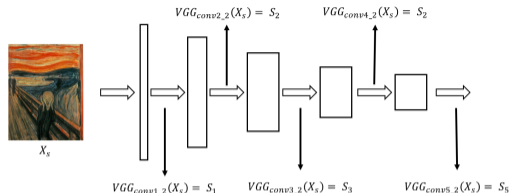
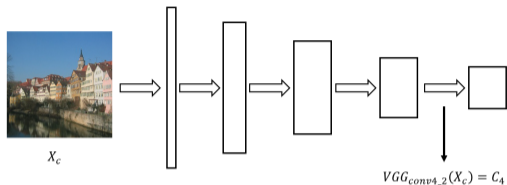


Neural style transfer



$$\text{Content loss: } \mathcal{L}_c(X) = \frac{1}{n_4 c_4} \|F_4 - C_4\|_F^2, \quad \text{where } F_4 = VGG_{conv4_2}(X)$$

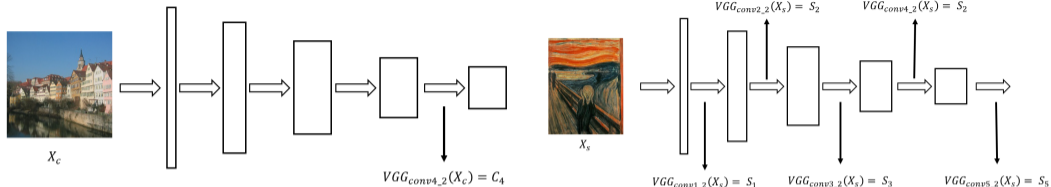
Neural style transfer



$$\text{Content loss: } \mathcal{L}_c(X) = \frac{1}{n_4 c_4} \|F_4 - C_4\|_F^2, \quad \text{where } F_4 = VGG_{conv4_2}(X)$$

$$\text{Style loss: } \mathcal{L}_s(X) = \frac{1}{|\mathcal{I}_s|} \sum_{\ell \in \mathcal{I}_s} \frac{1}{c_\ell^2} \left\| \frac{1}{n_\ell} F_\ell^T F_\ell - \frac{1}{n'_\ell} S_\ell^T S_\ell \right\|_F^2, \quad \text{where } F_\ell = VGG_{conv\ell_2}(X)$$

Neural style transfer



Content loss: $\mathcal{L}_c(X) = \frac{1}{n_4 c_4} \|F_4 - C_4\|_F^2$, where $F_4 = VGG_{conv4_2}(X)$

Style loss: $\mathcal{L}_s(X) = \frac{1}{|\mathcal{I}_s|} \sum_{\ell \in \mathcal{I}_s} \frac{1}{c_\ell^2} \left\| \frac{1}{n_\ell} F_\ell^T F_\ell - \frac{1}{n'_\ell} S_\ell^T S_\ell \right\|_F^2$, where $F_\ell = VGG_{conv\ell_2}(X)$

Complete loss: $\mathcal{L}(X) = \lambda_c \mathcal{L}_c(X) + \lambda_s \mathcal{L}_s(X)$, minimised by gradient descent.

Neural style transfer

The method works well when the style is a painting but not when it is a photograph of a natural scene.

The method works well when the style is a painting but not when it is a photograph of a natural scene.

For photorealistic style transfer, Luan *et al.*, CVPR, 2017 propose to regularize the problem:

$$\mathcal{L}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}),$$

where \mathbf{L} is the matting Laplacian computed on \mathbf{X}_c [Levin *et al.*, IEEE PAMI, 2008].

Neural style transfer



Neural style transfer

Without regularisation

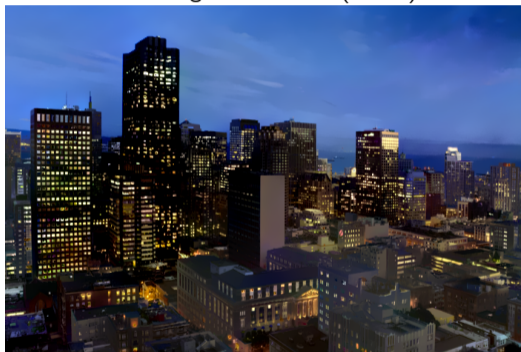


Neural style transfer

Without regularisation



With regularisation $\text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X})$



Neural style transfer

How can we estimate rapidly a solution of the photorealistic style transfer problem ?

$$\mathcal{L}_{\text{ph}}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}),$$

Neural style transfer

How can we estimate rapidly a solution of the photorealistic style transfer problem ?

$$\mathcal{L}_{\text{ph}}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}),$$

Re-use an existing network for fast artistic style transfer and retrain it.

Neural style transfer

How can we estimate rapidly a solution of the photorealistic style transfer problem ?

$$\mathcal{L}_{\text{ph}}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}),$$

Re-use an existing network for fast artistic style transfer and retrain it.

- ▶ Can the network capture the effect of $\text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X})$?
- ▶ Which filters encode the effect of $\text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X})$?
- ▶ If we want to change the regularisation, we need *a priori* to retrain.

Neural style transfer

How can we estimate rapidly a solution of the photorealistic style transfer problem ?

$$\mathcal{L}_{\text{ph}}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}),$$

Re-use an existing network for fast artistic style transfer and retrain it.

- ▶ Can the network capture the effect of $\text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X})$?
- ▶ Which filters encode the effect of $\text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X})$?
- ▶ If we want to change the regularisation, we need *a priori* to retrain.

Following the idea of unrolling optimisation algorithms [Gregor *et al.*, ICML, 2010], we propose a flexible solver for which

- ▶ the network is trained to approximately minimise $\lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X})$, *i.e.*, the artistic style transfer problem;
- ▶ the photorealistic prior can be added at runtime (no retraining).

A Flexible Solver

A Flexible Solver - Gradient descent and Projected gradient descent

A Flexible Solver - Gradient descent and Projected gradient descent

If \mathcal{L} is differentiable, one can solve

$$\min_{\mathbf{X}} \mathcal{L}(\mathbf{X})$$

by gradient descent

$$\mathbf{X}^{(t+1)} = \mathbf{X}^{(t)} - \mu \nabla \mathcal{L}(\mathbf{X}^{(t)}).$$

A Flexible Solver - Gradient descent and Projected gradient descent

If \mathcal{L} is differentiable, one can solve

$$\min_{\mathbf{X}} \mathcal{L}(\mathbf{X})$$

by gradient descent

$$\mathbf{X}^{(t+1)} = \mathbf{X}^{(t)} - \mu \nabla \mathcal{L}(\mathbf{X}^{(t)}).$$

If we add the constraint

$$\min_{\mathbf{X}} \mathcal{L}(\mathbf{X}) \quad \text{subject to} \quad \mathbf{X} \in \mathcal{C},$$

then a solution can be found using the iterations:

$$\mathbf{X}^{(t+1)} = \mathcal{P}_{\mathcal{C}} \left[\mathbf{X}^{(t)} - \mu \nabla \mathcal{L}(\mathbf{X}^{(t)}) \right],$$

where $\mathcal{P}_{\mathcal{C}}$ is the Euclidean projection onto \mathcal{C} .

A Flexible solver - Smooth graph signals

Why is it useful for solving the photorealistic style transfer of Luan *et al.* ?

$$\mathcal{L}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^\top \mathbf{L} \mathbf{X}).$$

A Flexible solver - Smooth graph signals

Why is it useful for solving the photorealistic style transfer of Luan *et al.* ?

$$\mathcal{L}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^\top \mathbf{L} \mathbf{X}).$$

In graph signal processing terms, finding \mathbf{X} for which $\text{Tr}(\mathbf{X}^\top \mathbf{L} \mathbf{X})$ means finding an image that is smooth on the graph with Laplacian \mathbf{L} .

A Flexible solver - Smooth graph signals

Why is it useful for solving the photorealistic style transfer of Luan *et al.* ?

$$\mathcal{L}(X) = \lambda_c \mathcal{L}_c(X) + \lambda_s \mathcal{L}_s(X) + \lambda_L \text{Tr}(X^T L X).$$

In graph signal processing terms, finding X for which $\text{Tr}(X^T L X)$ means finding an image that is smooth on the graph with Laplacian L .

There exists a special class of smooth graph signals: **k -bandlimited graph signals**.

A Flexible solver - Smooth graph signals

The eigendecomposition of L is

$$L = U\Lambda U^T$$

where

- ▶ $U = (\mathbf{u}_1, \dots, \mathbf{u}_n)$ contains the eigenvectors of L (graph Fourier basis).
- ▶ Λ is a diagonal matrix that contains the eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ of L .

A Flexible solver - Smooth graph signals

A k -bandlimited graph signal satisfies

$$\mathbf{x} = \left(\begin{array}{c|c} \mathbf{u}_1 & \dots & \mathbf{u}_k & \mathbf{u}_{k+1} & \dots & \mathbf{u}_n \end{array} \right) \begin{pmatrix} \alpha_1 \\ \vdots \\ \alpha_k \\ 0 \\ \vdots \\ 0 \end{pmatrix} = \mathbf{U}_k \boldsymbol{\alpha}_k,$$

i.e., $\mathbf{x} \in \text{span}(\mathbf{U}_k)$.

A Flexible solver - Photorealistic style transfer

Why is it useful for solving the photorealistic style transfer of Luan *et al.* ?

A Flexible solver - Photorealistic style transfer

Why is it useful for solving the photorealistic style transfer of Luan *et al.* ?

$$\min_{\mathbf{X}} \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) + \lambda_L \text{Tr}(\mathbf{X}^T \mathbf{L} \mathbf{X}).$$

and

$$\min_{\mathbf{X}} \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X}) \quad \text{s.t.} \quad \mathbf{x}_i \in \text{span}(\mathbf{U}_k), \forall i,$$

yields similar solutions (\mathbf{x}_i denotes the i^{th} column of \mathbf{X}).

A Flexible solver

The artistic style transfer problem

$$\min_{\mathbf{X}} \mathcal{L}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X})$$

can be solved by

$$\mathbf{X}^{(t+1)} = \mathbf{X}^{(t)} - \mu \nabla \mathcal{L}(\mathbf{X}^{(t)}).$$

A Flexible solver

The artistic style transfer problem

$$\min_{\mathbf{X}} \mathcal{L}(\mathbf{X}) = \lambda_c \mathcal{L}_c(\mathbf{X}) + \lambda_s \mathcal{L}_s(\mathbf{X})$$

can be solved by

$$\mathbf{X}^{(t+1)} = \mathbf{X}^{(t)} - \mu \nabla \mathcal{L}(\mathbf{X}^{(t)}).$$

The photorealistic style transfer problem

$$\min_{\mathbf{X}} \mathcal{L}(\mathbf{X}) \quad \text{s.t.} \quad \mathbf{x}_i \in \text{span}(\mathbf{U}_k), \forall i,$$

can be solved by

$$\mathbf{X}^{(t+1)} = \mathcal{P}_{\text{span}(\mathbf{U}_k)} \left[\mathbf{X}^{(t)} - \mu \nabla \mathcal{L}(\mathbf{X}^{(t)}) \right].$$

Fast solver for artistic and photorealistic style transfer

Fast solver for artistic and photorealistic style transfer

We propose to learn how to solve the **artistic** style transfer problem by replacing the true gradient with a learned update

$$X^{(t+1)} = X^{(t)} - \mu \nabla \mathcal{L}(X^{(t)}) \quad \longrightarrow \quad X^{(t+1)} = X^{(t)} - g_t(X^{(t)})$$

for $t = 0, \dots, N - 1$ ($N=4$).

Fast solver for artistic and photorealistic style transfer

We propose to learn how to solve the **artistic** style transfer problem by replacing the true gradient with a learned update

$$X^{(t+1)} = X^{(t)} - \mu \nabla \mathcal{L}(X^{(t)}) \quad \longrightarrow \quad X^{(t+1)} = X^{(t)} - g_t(X^{(t)})$$

for $t = 0, \dots, N - 1$ ($N=4$).

And propose to solve the **photorealistic** style transfer problem, **without retraining**, by adding $\mathcal{P}_{\text{span}(U_k)}$ after each update

$$X^{(t+1)} = \mathcal{P}_{\text{span}(U_k)} \left[X^{(t)} - g_t(X^{(t)}) \right].$$

Fast solver for artistic and photorealistic style transfer

We propose to learn how to solve the **artistic** style transfer problem by replacing the true gradient with a learned update

$$X^{(t+1)} = X^{(t)} - \mu \nabla \mathcal{L}(X^{(t)}) \quad \longrightarrow \quad X^{(t+1)} = X^{(t)} - g_t(X^{(t)})$$

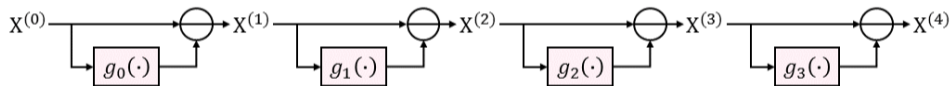
for $t = 0, \dots, N - 1$ ($N=4$).

And propose to solve the **photorealistic** style transfer problem, **without retraining**, by adding $\mathcal{P}_{\text{span}(U_k)}$ after each update

$$X^{(t+1)} = \mathcal{P}_{\text{span}(U_k)} \left[X^{(t)} - g_t(X^{(t)}) \right].$$

There exist fast methods to compute $\mathcal{P}_{\text{span}(U_k)}$.

Architecture of g_t



Architecture of g_t

Continuing with the idea of unrolling gradient descent, we should mimic the architecture of $\nabla \mathcal{L}$ to build g_t :

$$\nabla \mathcal{L} = \lambda_c \nabla \mathcal{L}_c(\mathbf{X}^{(t)}) + \lambda_s \nabla \mathcal{L}_s(\mathbf{X}^{(t)}). \quad (1)$$

Architecture of g_t

Continuing with the idea of unrolling gradient descent, we should mimic the architecture of $\nabla \mathcal{L}$ to build g_t :

$$\nabla \mathcal{L} = \lambda_c \nabla \mathcal{L}_c(\mathbf{X}^{(t)}) + \lambda_s \nabla \mathcal{L}_s(\mathbf{X}^{(t)}). \quad (1)$$

In practice, we do not consider $\nabla \mathcal{L}_c$ and mimic only the architecture of $\nabla \mathcal{L}_s$.

Architecture of g_t

Continuing with the idea of unrolling gradient descent, we should mimic the architecture of $\nabla \mathcal{L}$ to build g_t :

$$\nabla \mathcal{L} = \lambda_c \nabla \mathcal{L}_c(\mathbf{X}^{(t)}) + \lambda_s \nabla \mathcal{L}_s(\mathbf{X}^{(t)}). \quad (1)$$

In practice, we do not consider $\nabla \mathcal{L}_c$ and mimic only the architecture of $\nabla \mathcal{L}_s$.

To compensate, we

1. initialize $\mathbf{X}^{(0)} = \mathbf{X}_c$;
2. use the complete style transfer loss \mathcal{L} for training.

We have

$$\mathcal{L}_s(\mathbf{X}) = \frac{1}{|\mathcal{I}_s|} \sum_{\ell \in \mathcal{I}_s} \frac{1}{c_\ell^2} \left\| \frac{1}{n_\ell} \mathbf{F}_\ell^\top \mathbf{F}_\ell - \frac{1}{n'_\ell} \mathbf{S}_\ell^\top \mathbf{S}_\ell \right\|_{\mathbf{F}}^2,$$

where $\mathbf{F}_\ell = \text{VGG}_\ell(\mathbf{X})$ and $\mathbf{S}_\ell = \text{VGG}_\ell(\mathbf{X}_s)$.

We have

$$\mathcal{L}_s(\mathbf{X}) = \frac{1}{|\mathcal{I}_s|} \sum_{\ell \in \mathcal{I}_s} \frac{1}{c_\ell^2} \left\| \frac{1}{n_\ell} \mathbf{F}_\ell^\top \mathbf{F}_\ell - \frac{1}{n'_\ell} \mathbf{S}_\ell^\top \mathbf{S}_\ell \right\|_{\mathbf{F}}^2,$$

where $\mathbf{F}_\ell = \text{VGG}_\ell(\mathbf{X})$ and $\mathbf{S}_\ell = \text{VGG}_\ell(\mathbf{X}_s)$.

The gradient is obtained in 3 steps

1. Compute each $\mathbf{F}_\ell = \text{VGG}_\ell(\mathbf{X})$;
2. Compute each partial derivative:

$$\frac{\partial \mathcal{L}_s}{\partial \mathbf{F}_\ell} \propto \mathbf{F}_\ell \cdot \left[\frac{1}{n_\ell} \mathbf{F}_\ell^\top \mathbf{F}_\ell - \frac{1}{n'_\ell} \mathbf{S}_\ell^\top \mathbf{S}_\ell \right]; \quad (2)$$

3. Back-propagate these partial derivatives to the input of the VGG-19.

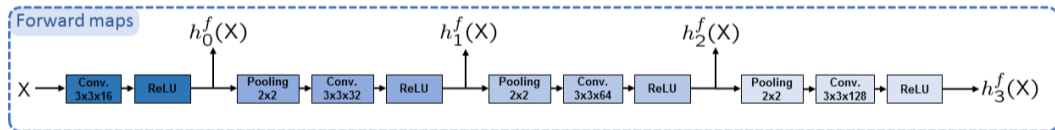
Architecture of g_t

Step 1: Compute each $F_\ell = \text{VGG}_\ell(X)$.

Architecture of g_t

Step 1: Compute each $F_\ell = \text{VGG}_\ell(X)$.

We replace the VGG-19 network with a similar but “simpler” one:



The forward maps $h_\ell^f(X)$ play the role of $\text{VGG}_\ell(X)$. All the filters in $h_\ell^f(X)$ are trained.

Architecture of g_t

Step 2: Compute each partial derivative

$$\mathbf{F}_\ell \cdot \left[\frac{1}{n_\ell} \mathbf{F}_\ell^\top \mathbf{F}_\ell - \frac{1}{n'_\ell} \mathbf{S}_\ell^\top \mathbf{S}_\ell \right] \quad (3)$$

Step 2: Compute each partial derivative

$$\mathbf{F}_\ell \cdot \left[\frac{1}{n_\ell} \mathbf{F}_\ell^\top \mathbf{F}_\ell - \frac{1}{n'_\ell} \mathbf{S}_\ell^\top \mathbf{S}_\ell \right] \quad (3)$$

We directly transform the above expression to

$$h_\ell^f(\mathbf{X}) \cdot \left[\frac{1}{n_\ell} h_\ell^f(\mathbf{X})^\top h_\ell^f(\mathbf{X}) - \mathbf{H}_{\ell,t}^s \right], \quad (4)$$

where, therefore, $\mathbf{H}_{\ell,t}^s$ controls the style.

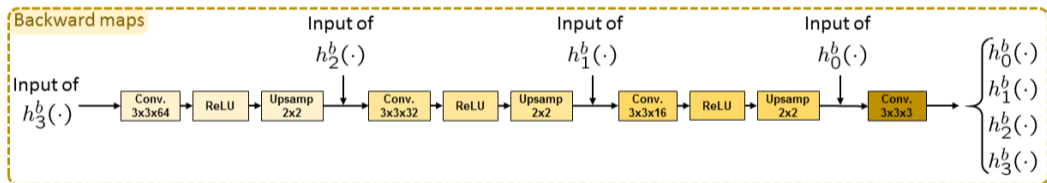
Architecture of g_t

Step 3: Backpropagation

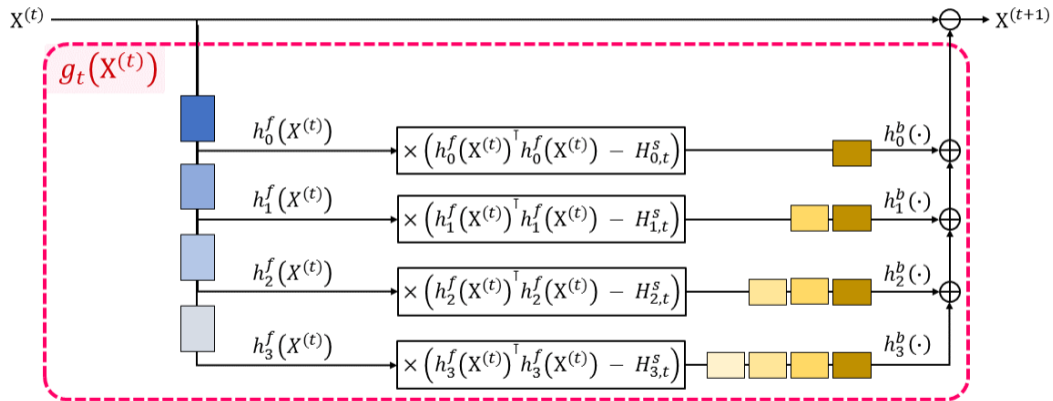
Architecture of g_t

Step 3: Backpropagation

We use backwards maps h_ℓ^b symmetric to h_ℓ^f



Global structure



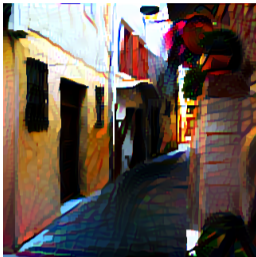
Some results

Artistic style transfer

X_c



[Gatys et al.] - X_c as init.



Ours



Style 1, *Portrait of Dora Maar*,
P. Picasso, 1937.



Style 2, *A Muse*, P. Picasso, 1935.



Artistic style transfer

X_c



[Gatys et al.] - X_c as init.



Ours



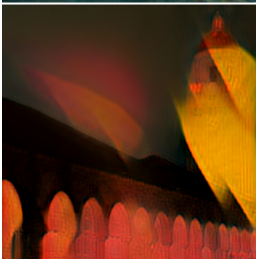
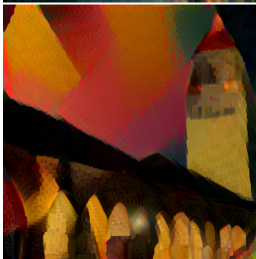
**Style 3, *The scream*,
E. Munch, 1893.**



**Style 4,
Self portrait,
R. Magritte, 1923.**



technicolor

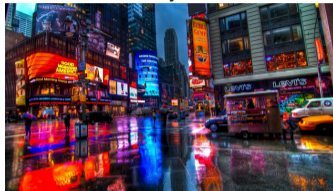



Photorealistic style transfer

X_c



Style



[Luan et al.]



Ours - Graph filt.

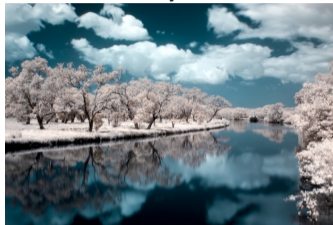


Photorealistic style transfer

X_c



Style



[Luan et al.]



Ours - Graph filt.



Photorealistic style transfer

X_c



[Luan *et al.*]



Style



Ours - Graph filt.



Photorealistic style transfer

X_c



Style



Ours - No Graph filt.



Ours - Graph filt.



Conclusion

Conclusion

- ▶ Deep networks can benefit from the flexibility of existing optimisation algorithms
- ▶ Unrolling algorithms make it easier to interpret the role of the filter in the network
- ▶ References and more details in our technical report [arXiv:1806.05285](https://arxiv.org/abs/1806.05285)
- ▶ Related work: Li *et al.*, “A Closed-form Solution to Photorealistic Image Stylization,” ECCV, 2018.