

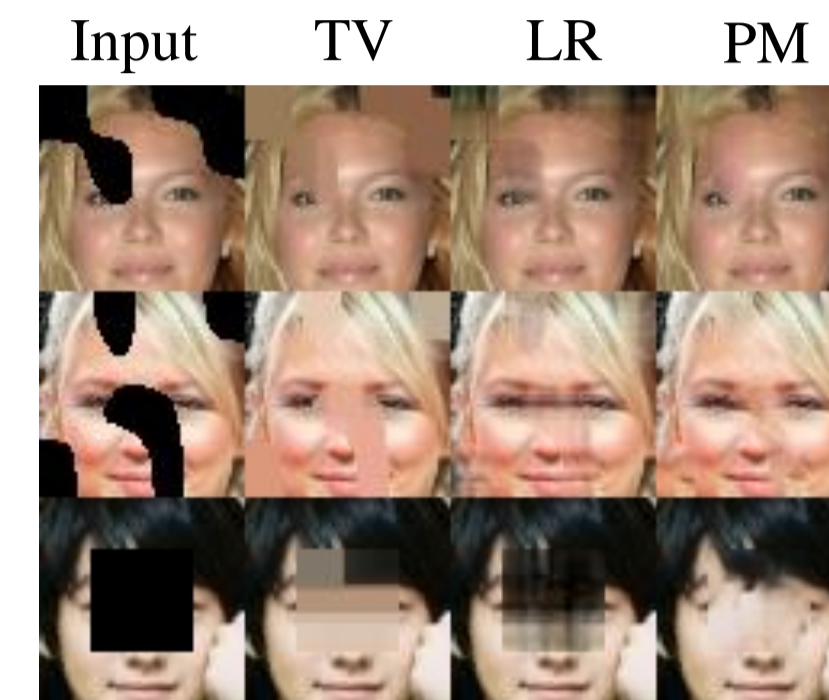
MOTIVATION

Task: Semantic image inpainting (filling large missing regions)

- ill-posed task
- requires strong prior knowledge on the data
- extracting information from only a single image produces unsatisfactory results

Contributions:

- deep generative models produce missing content by conditioning on available data
- inpainting as constrained optimization problem using **context and prior loss**



INTRODUCTION


Problem Formulation:

- Corrupted image: y
- Binary mask: M
- Task: predict uncorrupted version \hat{x}



Baselines:

- Total Variation and Low Rank assume smoothness in the pixel space
- Context Encoder is a deep model which treats inpainting as a regression problem



Instead of explicitly defining the prior, we utilize deep generative models to capture prior information.

Generative Adversarial Networks:

- Generator G : deep net mapping perturbation z to artificial sample
- Discriminator D : deep net discriminating between artificial and real sample, x
- Program:

$$\min_G \max_D V(G, D) = \mathbb{E}_{x \sim p_{data}} [\log(D(x))] + \mathbb{E}_{z \sim p_Z(z)} [\log(1 - D(G(z)))]$$

OUR APPROACH

Intuition of our approach:

- Hypothesis: image that is not from p_{data} (e.g., corrupted data) should not lie on the learned encoding manifold; use manifold can be used as a prior
- Instead of working in the pixel space, we recover the encoding \hat{z} “closest” to the corrupted image while constrained to the manifold

Solving for the “closest” encoding \hat{z} :

$$\hat{z} = \arg \min_z \mathcal{L}_c(z|y, M) + \mathcal{L}_p(z)$$

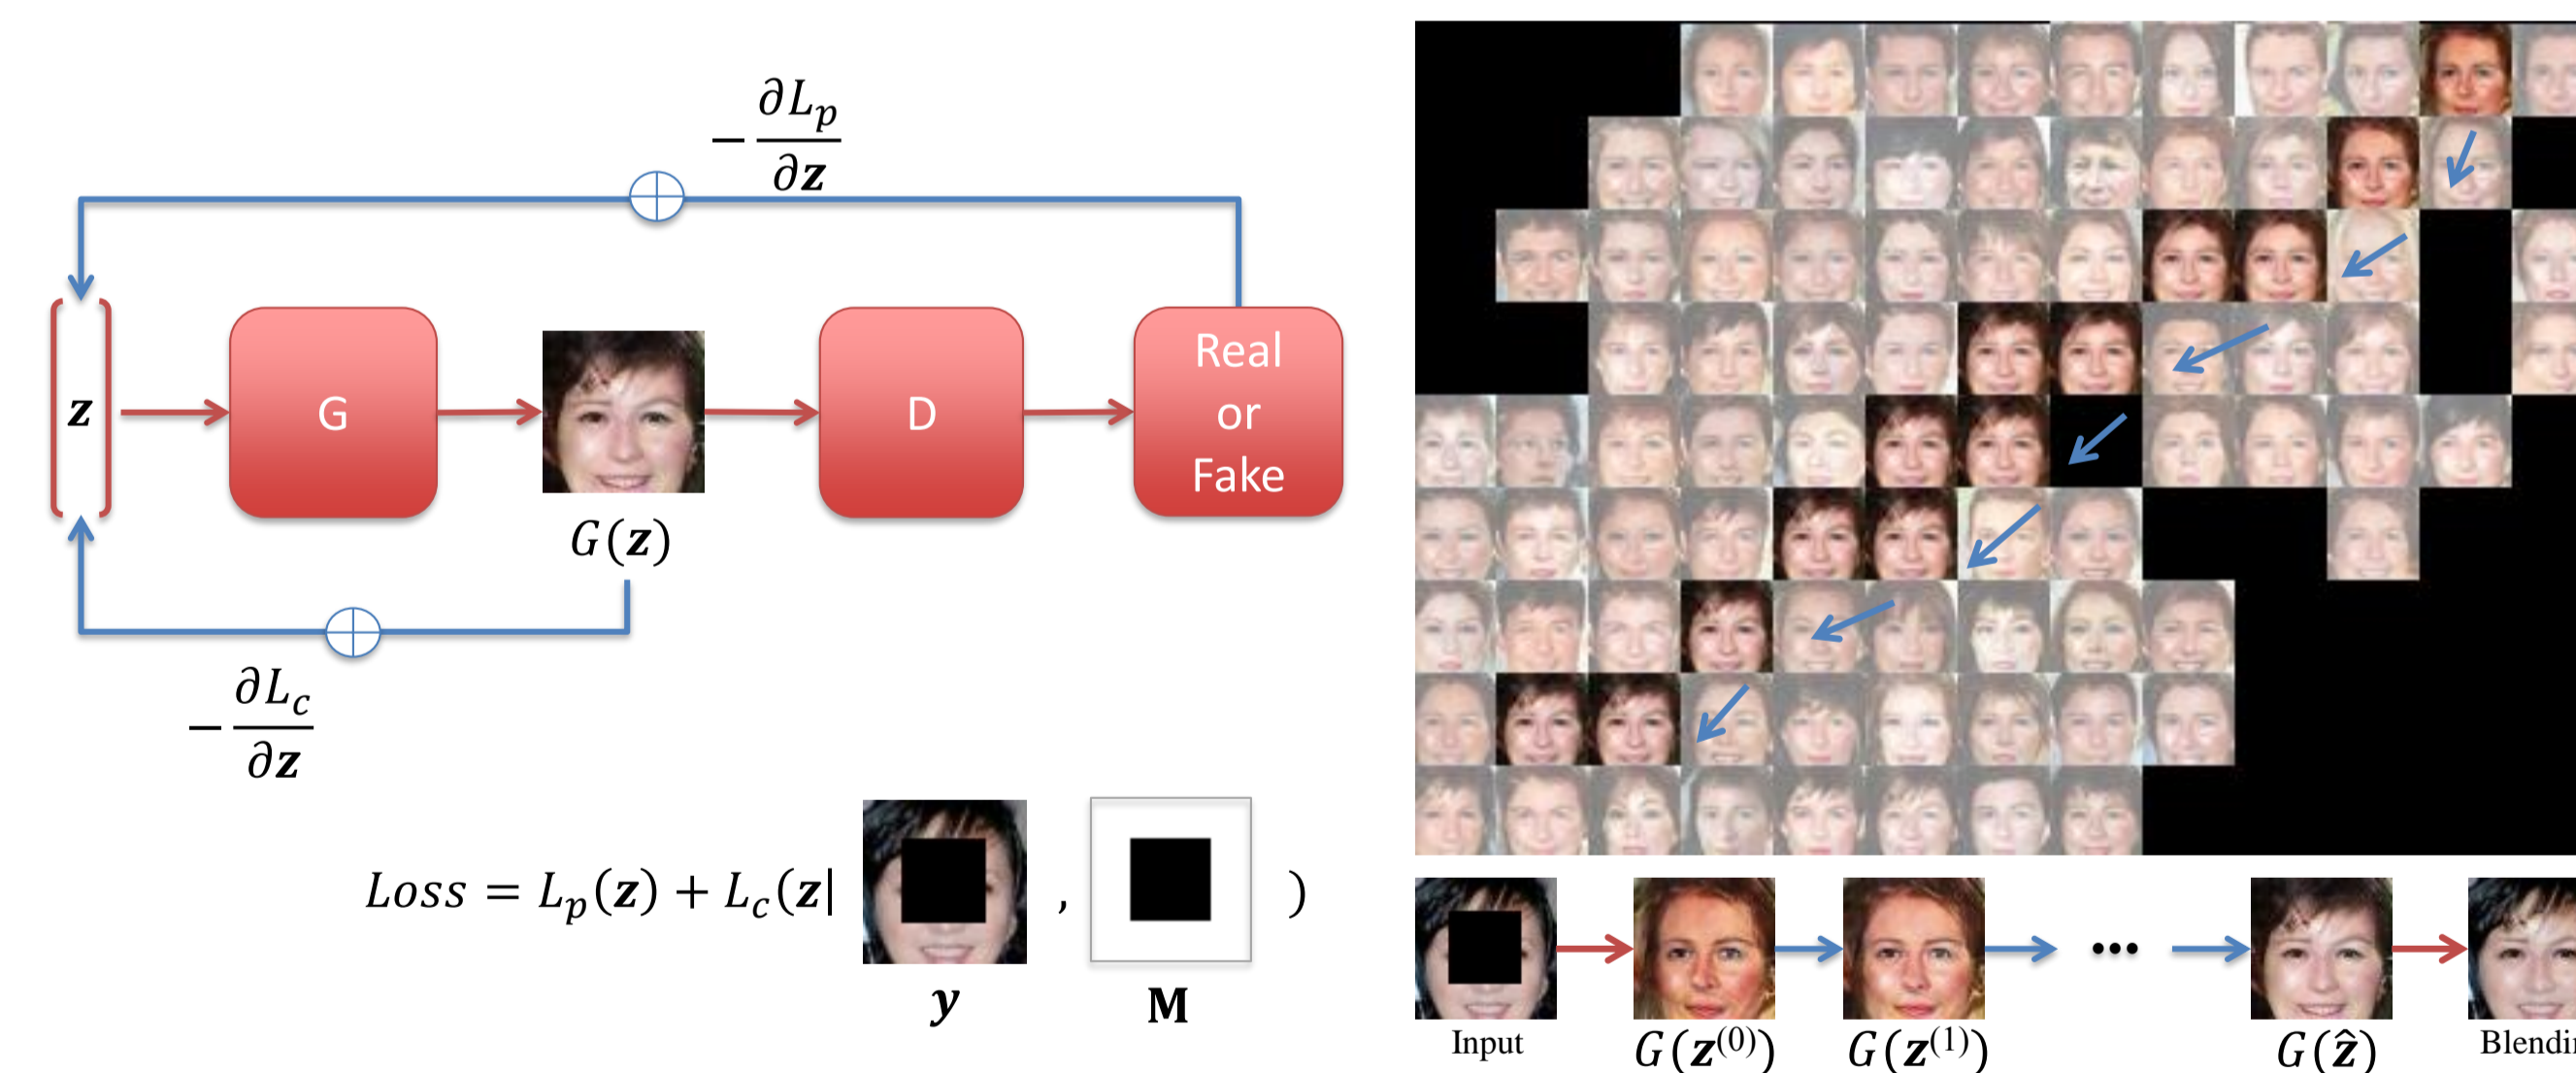
Context Loss: importance weighted metric W to enforce similarity to the uncorrupted regions:

$$\mathcal{L}_c(z|y, M) = \|W \odot (G(z) - y)\|_1$$

Prior Loss: prior penalizing unrealistic images based on the discriminator:

$$\mathcal{L}_p(z) = \lambda \log(1 - D(G(z)))$$

Illustration of the approach:



$$\text{Importance weight metric: } W_i = \begin{cases} \sum_{j \in N(i)} \frac{(1 - M_j)}{|N(i)|} & \text{if } M_i \neq 0 \\ 0 & \text{if } M_i = 0 \end{cases}$$


$N(i)$ defines the neighborhood of i

Recovering prediction \hat{x} via poisson blending rather than simple overlay:

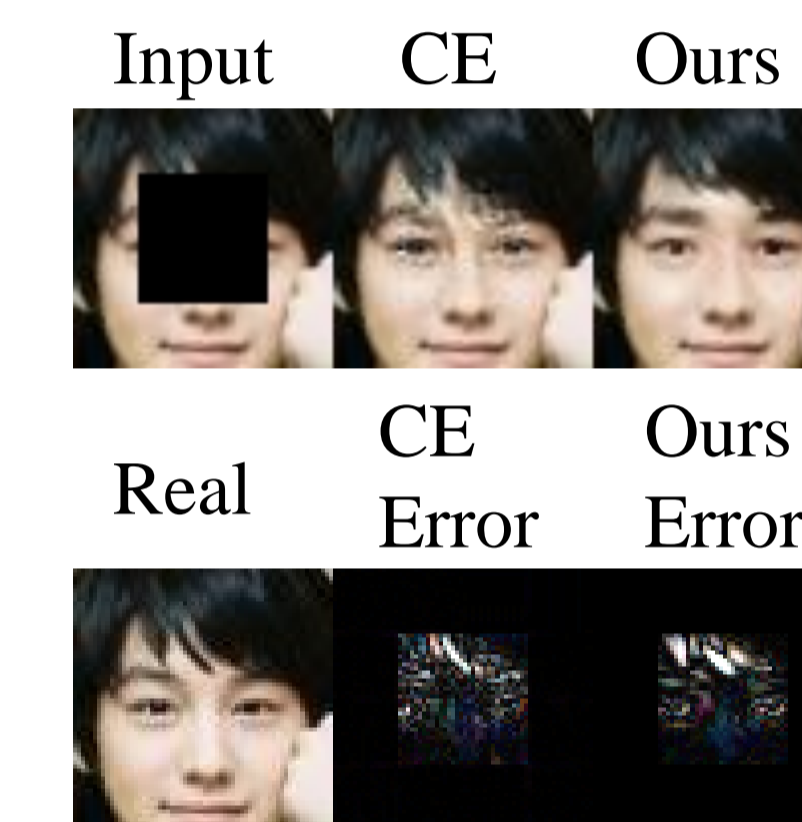
$$\hat{x} = \arg \min_x \|\nabla x - \nabla G(\hat{z})\|_2^2 \quad \text{s.t. } x_i = y_i \text{ for } M_i = 1$$

RESULTS

Comparison: Poisson Blending vs. Overlay:



Quantitative Results:



The PSNR values (dB) on the test sets. Left/right results are by Context Encoder (CE)/ours:

Masks/Dataset	CelebA	SVHN	Cars
Center pattern	21.3/19.4	22.3/19.0	14.1/13.5
random	19.2/17.4	22.3/19.8	14.0/14.1
half	20.6/22.8	24.1/33.0	16.1/18.9

- In the figure above, PSNR for CE is 24.71 dB and ours is 22.98 dB
- Higher PSNR does not mean better visual quality
- The solution is not unique, many hallucinations are reasonable

Qualitative Results:

