

Poster ID: THU-PM-002



# RobustNerf: Ignoring Distractors with Robust Losses

Sara Sabour, Suhani Vora, Daniel Duckworth,  
Ivan Krasin, David Fleet, Andrea Tagliasacchi



# NeRF is amazing at **3D multi view reconstruction from 2D images**




Inputs photos

# NeRF is amazing at **3D rendering from 2D images**



Inputs photos



MipNerf360

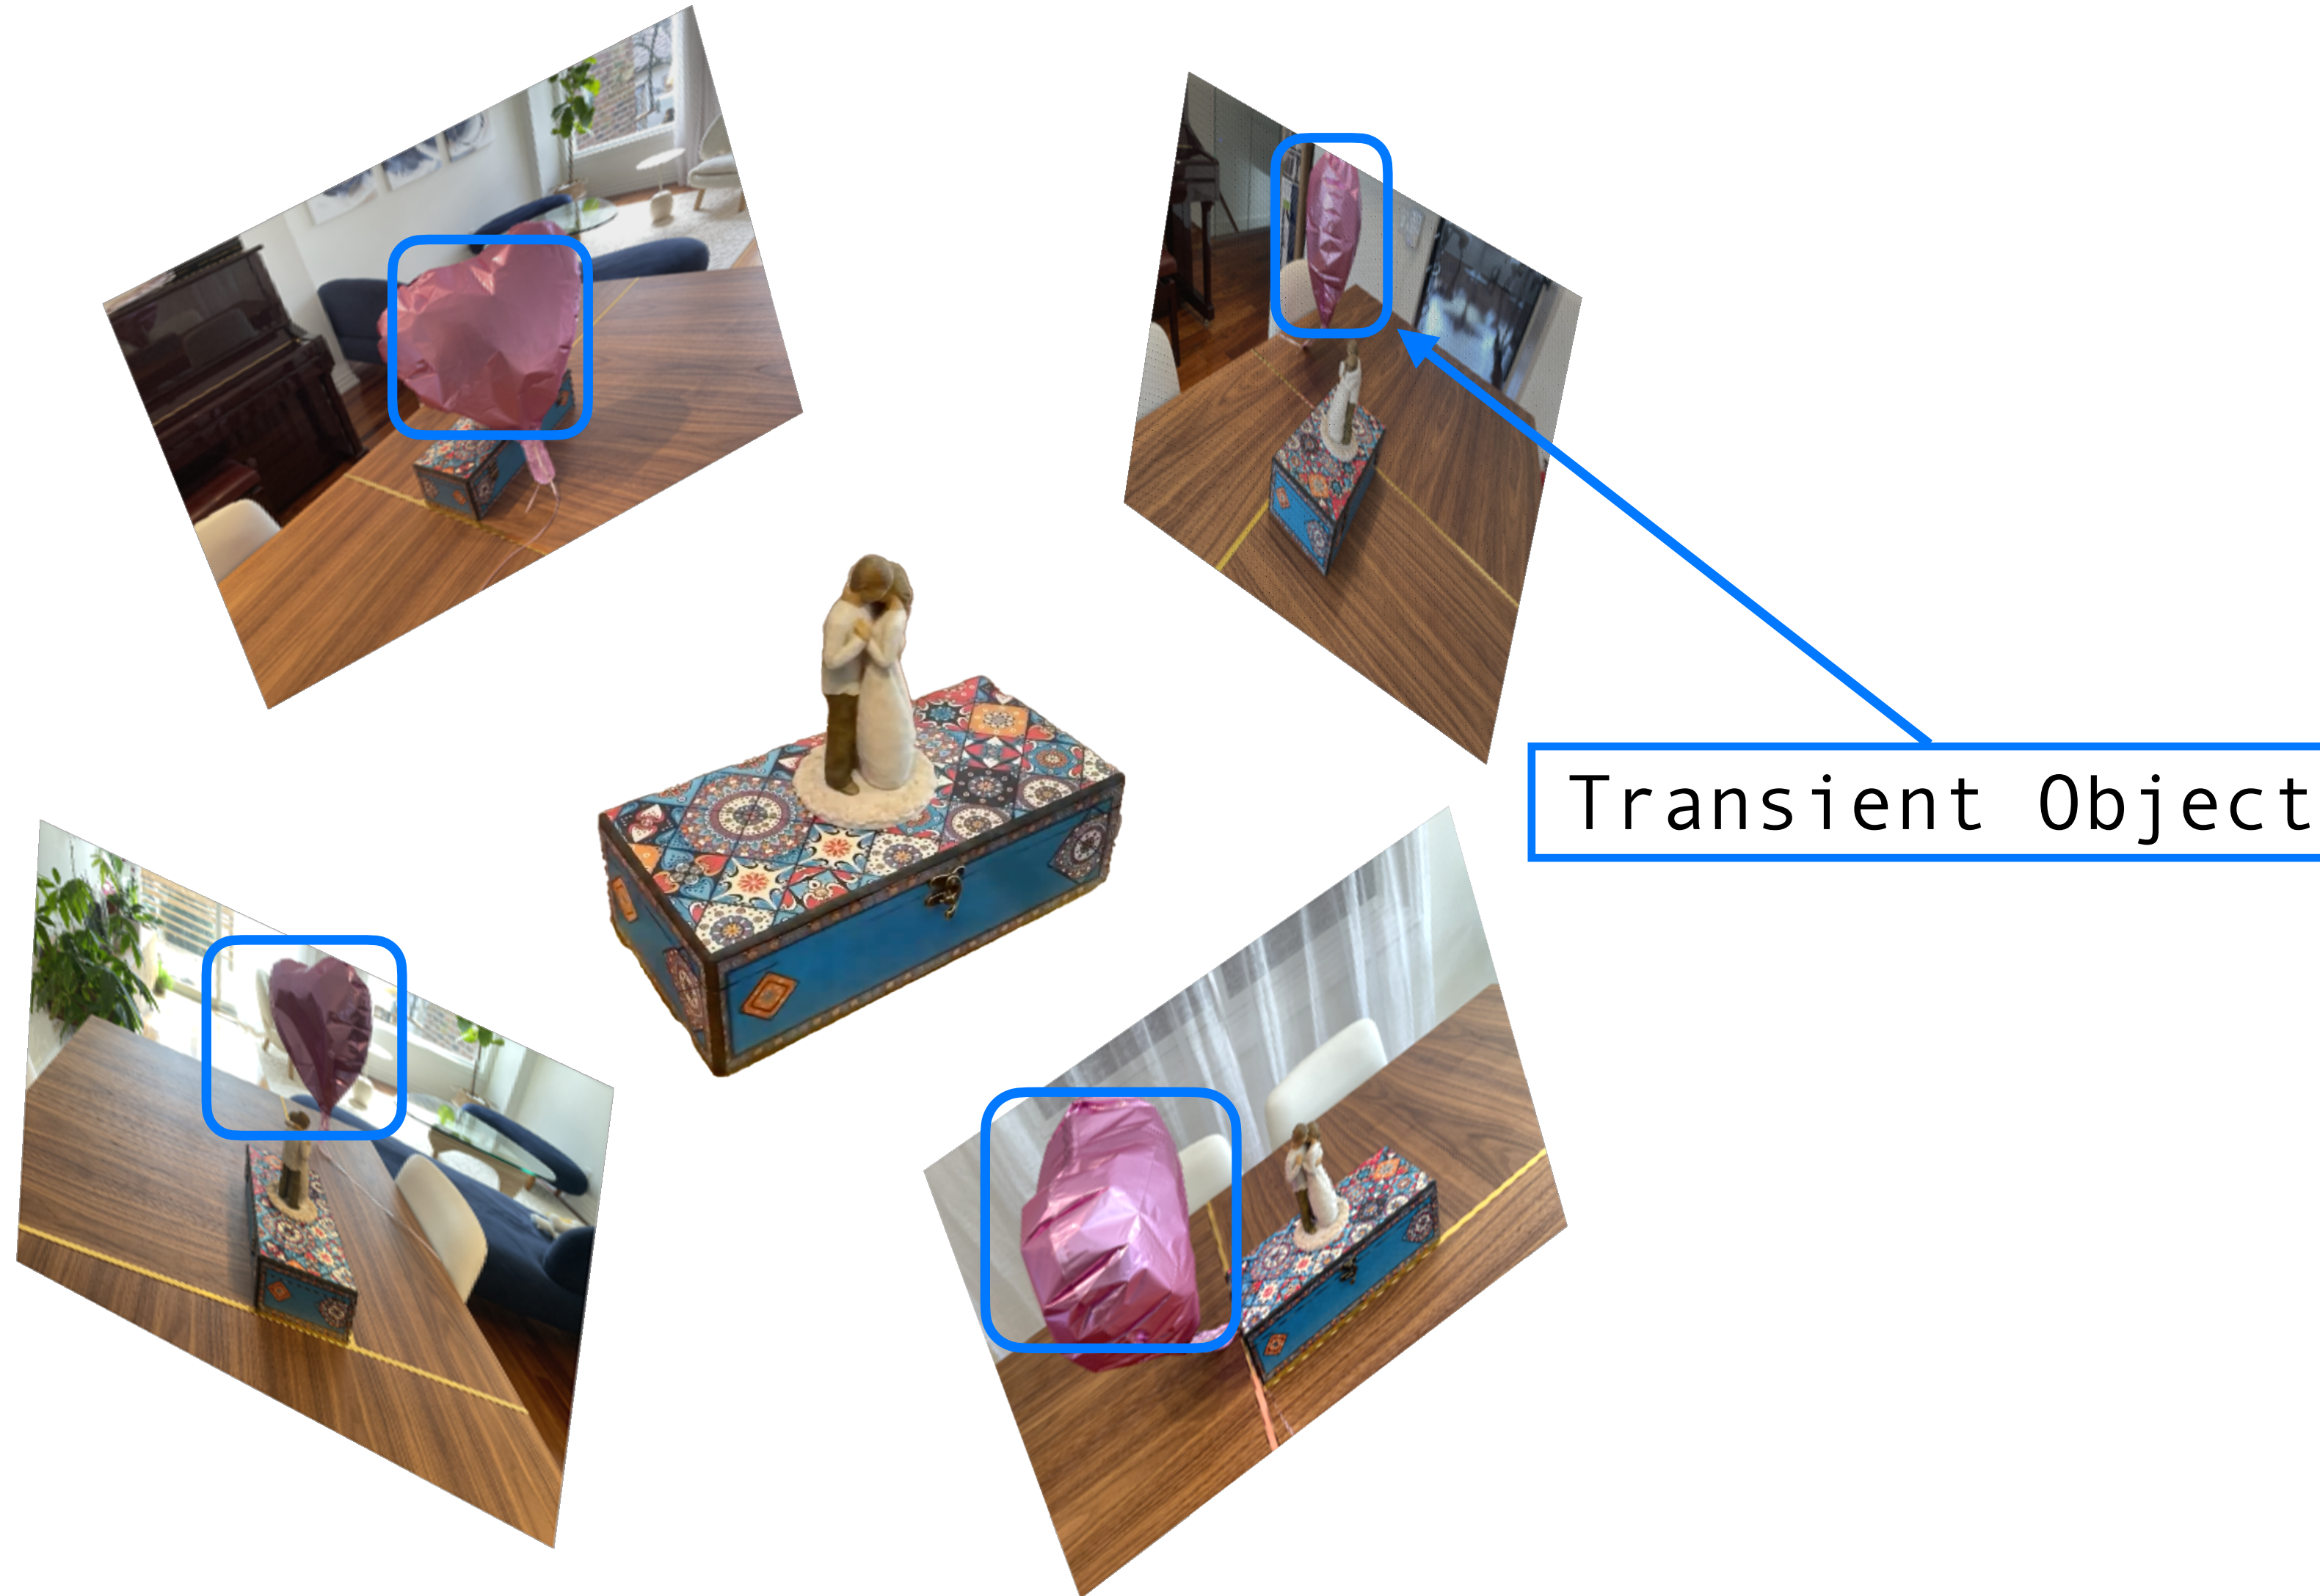


# NeRF struggles with **photometric inconsistencies**



Inputs photos

# NeRF struggles with **photometric inconsistencies**




Inputs photos

# NeRF struggles with **photometric inconsistencies**



Inputs photos




MipNerf360 : Charbonnier Loss



# NeRF struggles with **photometric inconsistencies**



Inputs photos



MipNerf360 : Charbonnier Loss




Cloudy

# NeRF struggles with **photometric inconsistencies**



Inputs photos



MipNerf360 : Charbonnier Loss



Cloudy


Floater



# NeRF struggles with **photometric inconsistencies**



Inputs photos



Clean & Clear

# RobustNerf successfully removes artifacts!

Photometric inconsistencies are outliers in a robust optimization task.



MipNeRF360




RobustNeRF

Poster ID: THU-PM-002 link: <https://robustnerf.github.io>


# Photometric Consistency

- NeRF: optimizes the model parameters  $\theta$  by:
  - Assuming all pictures taken are of **exactly the same scene**
  - It is ideal for **gaussian noise**



# Photometric Consistency

- NeRF: optimizes the model parameters  $\theta$  by:
  - Assuming all pictures taken are of **exactly the same scene**
  - It is ideal for **gaussian noise**



$$\mathcal{L}_{\text{rgb}}^{\mathbf{r},i}(\theta) = \|\mathbf{C}(\mathbf{r}; \theta) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

photometric consistency

# Robust Estimation

- automatically identify non-consistent information
- model distractors as **optimization outliers**

# Robust Estimation


- automatically identify non-consistent information
- model distractors as **optimization outliers**

$$\mathcal{L}_{\text{rgb}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

v.s.

$$\mathcal{L}_{\text{robust}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \kappa(\|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2)$$

Family of robust kernels



# Robust Estimation


- automatically identify non-consistent information
- model distractors as **optimization outliers**

$$\mathcal{L}_{\text{rgb}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

v.s.

$$\mathcal{L}_{\text{robust}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \kappa(\|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2)$$

Family of robust kernels



# Robust Estimation


- automatically identify non-consistent information
- model distractors as **optimization outliers**

$$\mathcal{L}_{\text{rgb}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

v.s.

$$\mathcal{L}_{\text{robust}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \kappa(\|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2)$$

Family of robust kernels





# Robust Estimation


- automatically identify non-consistent information
- model distractors as **optimization outliers**

$$\mathcal{L}_{\text{rgb}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

v.s.

$$\mathcal{L}_{\text{robust}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \kappa(\|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2)$$


Family of robust kernels



MipNerf360: Charbonnier Loss


# Does it work?

Family of robust kernels



# Does it work?


## Family of robust kernels



details preserved, but floaters

# Does it work?

## Family of robust kernels




details preserved, but floaters



floaters gone, but details not preserved

# Are outlier pixels independent in transient objects?

Family of robust kernels



VS



Good at **independent** noise



Bad at **structured** noise

# Robust Optimization w/ IRLS

- IRLS: Iteratively Re-Weighted Least Squares
  - weights are **spatially consistent**
  - weights are in  **$\{0, 1\}$**  (hard weights)



# Robust Optimization w/ IRLS

- IRLS: Iteratively Re-Weighted Least Squares
  - weights are **spatially consistent**
  - weights are in **{0, 1}** (hard weights)



$$\mathcal{L}_{\text{oracle}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \mathbf{S}_i(\mathbf{r}) \cdot \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

v.s.

$$\mathcal{L}_{\text{robust}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \omega(\boldsymbol{\epsilon}^{(t-1)}(\mathbf{r})) \cdot \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

# Robust Optimization w/ IRLS

- IRLS: Iteratively Re-Weighted Least Squares
  - weights are **spatially consistent**
  - weights are in  $\{0, 1\}$  (hard weights)



$$\mathcal{L}_{\text{oracle}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \mathbf{S}_i(\mathbf{r}) \cdot \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

v.s.

$$\mathcal{L}_{\text{robust}}^{\mathbf{r},i}(\boldsymbol{\theta}) = \omega(\epsilon^{(t-1)}(\mathbf{r})) \cdot \|\mathbf{C}(\mathbf{r}; \boldsymbol{\theta}) - \mathbf{C}_i(\mathbf{r})\|_2^2$$

**Residuals at the previous iteration**




# Trimmed / Diffused Least Squares

$$\tilde{\omega}(\mathbf{r}) = \epsilon(\mathbf{r}) \leq \mathcal{T}_\epsilon, \quad \mathcal{T}_\epsilon = \text{Median}_{\mathbf{r}}\{\epsilon(\mathbf{r})\} . \quad \text{threshold the residuals}$$



residuals  $-\epsilon(\mathbf{r})$



inliers  $-\tilde{\omega}(\mathbf{r})$

# Trimmed / Diffused Least Squares

$$\tilde{\omega}(\mathbf{r}) = \epsilon(\mathbf{r}) \leq \mathcal{T}_\epsilon, \quad \mathcal{T}_\epsilon = \text{Median}_{\mathbf{r}}\{\epsilon(\mathbf{r})\} .$$


threshold the residuals

$$\mathcal{W}(\mathbf{r}) = (\tilde{\omega}(\mathbf{r}) \circledast \mathcal{B}_{3 \times 3}) \geq \mathcal{T}_{\circledast}, \quad \mathcal{T}_{\circledast} = 0.5 .$$

diffuse the classification



residuals -  $\epsilon(\mathbf{r})$



inliers -  $\tilde{\omega}(\mathbf{r})$



diffusion -  $\tilde{\omega}(\mathbf{r}) \circledast \mathcal{B}_{3 \times 3}$

# Trimmed / Diffused Least Squares

$\tilde{\omega}(\mathbf{r}) = \epsilon(\mathbf{r}) \leq \mathcal{T}_\epsilon$  ,  $\mathcal{T}_\epsilon = \text{Median}_{\mathbf{r}}\{\epsilon(\mathbf{r})\}$  . **threshold** the residuals

$\mathcal{W}(\mathbf{r}) = (\tilde{\omega}(\mathbf{r}) \circledast \mathcal{B}_{3 \times 3}) \geq \mathcal{T}_{\circledast}$  ,  $\mathcal{T}_{\circledast} = 0.5$  . **diffuse** the classification

$\omega(\mathcal{R}_8(\mathbf{r})) = \mathbb{E}_{\mathbf{s} \sim \mathcal{R}_{16}(\mathbf{r})} [\mathcal{W}(\mathbf{s})] \geq \mathcal{T}_{\mathcal{R}}$  ,  $\mathcal{T}_{\mathcal{R}} = 0.6$  . **threshold** again



residuals –  $\epsilon(\mathbf{r})$



inliers –  $\tilde{\omega}(\mathbf{r})$



diffusion –  $\tilde{\omega}(\mathbf{r}) \circledast \mathcal{B}_{3 \times 3}$



(IRLS) weights –  $\mathcal{W}(\mathbf{r})$

# Trimmed / Diffused Least Squares


$$\tilde{\omega}(\mathbf{r}) = \epsilon(\mathbf{r}) \leq \mathcal{T}_\epsilon, \quad \mathcal{T}_\epsilon = \text{Median}_{\mathbf{r}}\{\epsilon(\mathbf{r})\}. \quad \text{threshold the residuals}$$

$$\mathcal{W}(\mathbf{r}) = (\tilde{\omega}(\mathbf{r}) \circledast \mathcal{B}_{3 \times 3}) \geq \mathcal{T}_{\circledast}, \quad \mathcal{T}_{\circledast} = 0.5. \quad \text{diffuse the classification}$$

$$\omega(\mathcal{R}_8(\mathbf{r})) = \mathbb{E}_{\mathbf{s} \sim \mathcal{R}_{16}(\mathbf{r})} [\mathcal{W}(\mathbf{s})] \geq \mathcal{T}_{\mathcal{R}}, \quad \mathcal{T}_{\mathcal{R}} = 0.6. \quad \text{threshold again}$$



residuals –  $\epsilon(\mathbf{r})$



inliers –  $\tilde{\omega}(\mathbf{r})$



diffusion –  $\tilde{\omega}(\mathbf{r}) \circledast \mathcal{B}_{3 \times 3}$



(IRLS) weights –  $\mathcal{W}(\mathbf{r})$

# Training Progress w/ IRLS



Ground Truth



Train View



Distractors


# Training Progress w/ IRLS



Ground Truth



Train View



Distractors



# Training Progress w/ IRLS



Distractors

.5%

2%

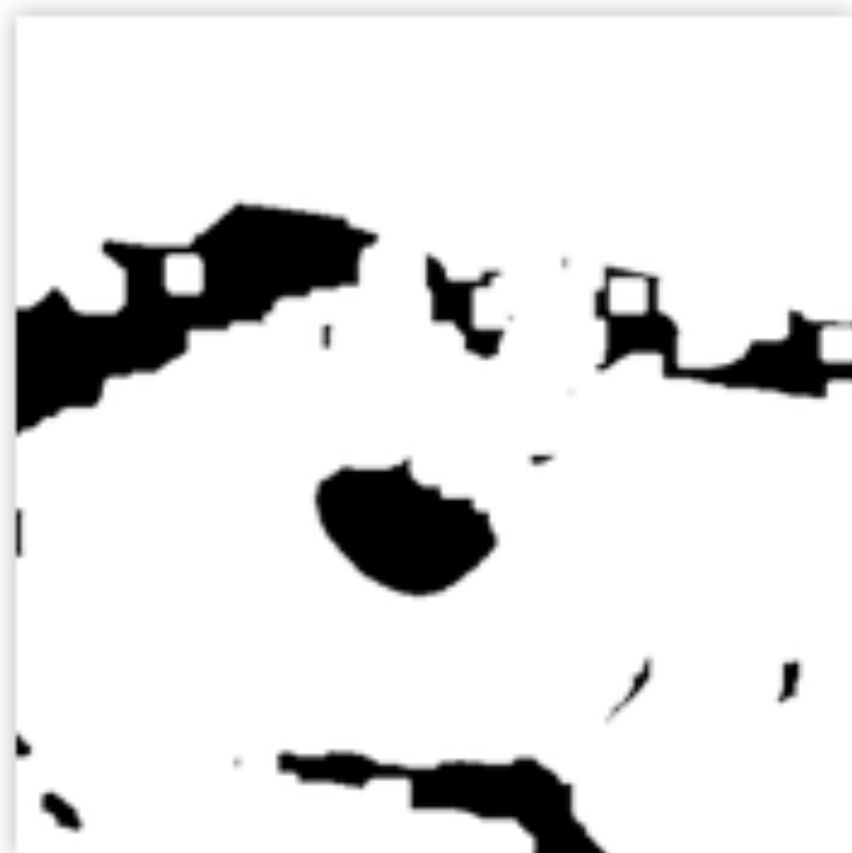
12%

100%

$(t/T)$



residuals  
 $\epsilon(\mathbf{r})$



weights  
 $\mathcal{W}(\mathbf{r})$

# Training Progress w/ IRLS



Distractors


.5%

2%

12%

100%

$(t/T)$




residuals  
 $\epsilon(\mathbf{r})$



weights  
 $\mathcal{W}(\mathbf{r})$



# Training Progress w/ IRLS



Distractors



.5%

2%

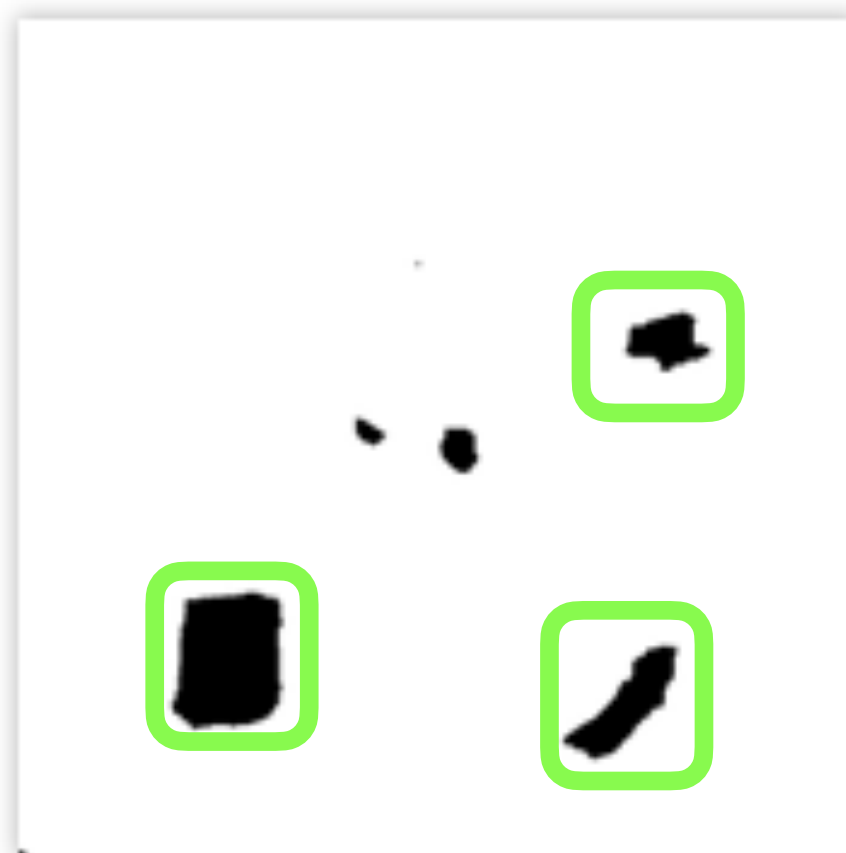
12%

100%

$(t/T)$



residuals  
 $\epsilon(\mathbf{r})$



weights  
 $W(\mathbf{r})$

MipNeRF360 (L2)


MipNeRF360 (L1)

MipNeRF360 (Ch)


D<sup>2</sup>NeRF

RobustNeRF

Single Outlier



Dozens of Outliers



Single Outlier

MipNeRF360 (L2)

MipNeRF360 (L1)



MipNeRF360 (Ch)

D<sup>2</sup>NeRF





RobustNeRF



Dozens of Outliers



# Comparison with MipNeRF360



MipNeRF360

RobustNeRF

# Viewdependent Scene Challenge

Train View



MipNeRF360



RobustNeRF



# Viewdependent Scene Challenge

Train View



MipNeRF360



RobustNeRF



# Viewdependent Scene Challenge

Train View



MipNeRF360



RobustNeRF




# Transient Shadow Challenge

Input

NeRF-W

D<sup>2</sup>NeRF


RobustNeRF






# Side by Side Video on "D2NeRF Pick"


Input



D2Nerf




RobustNerf



1"

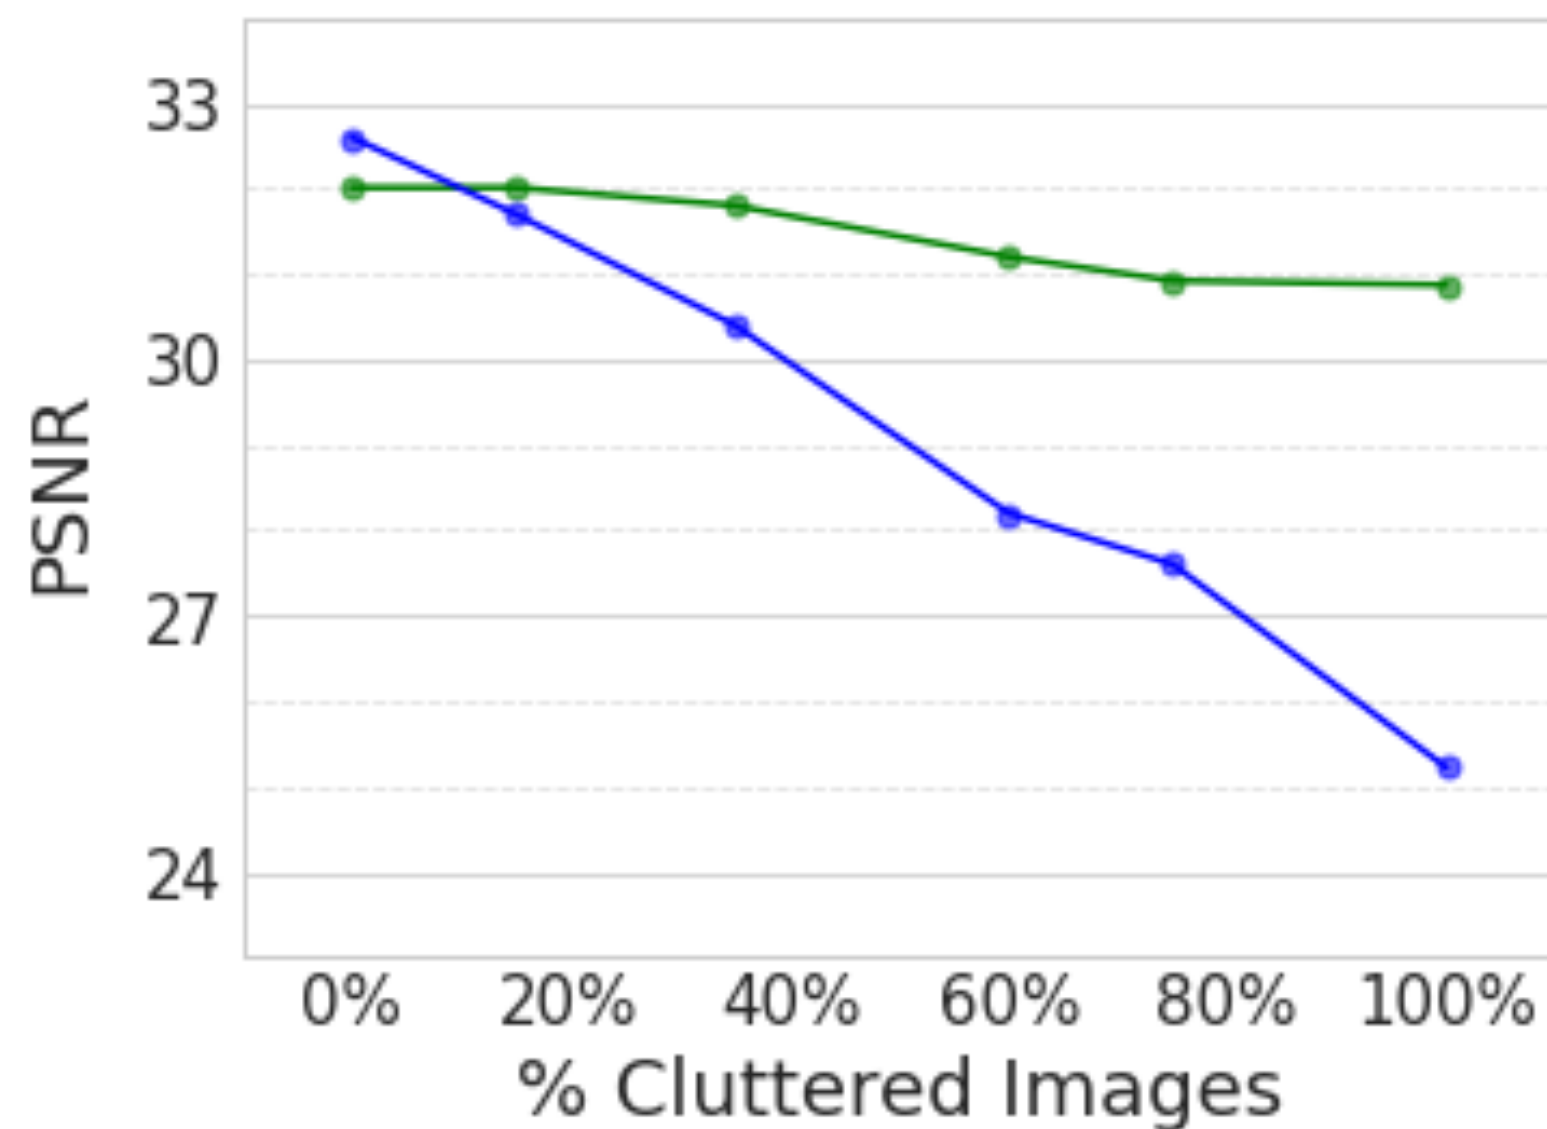

6"

shadow



Floater

# Limitations

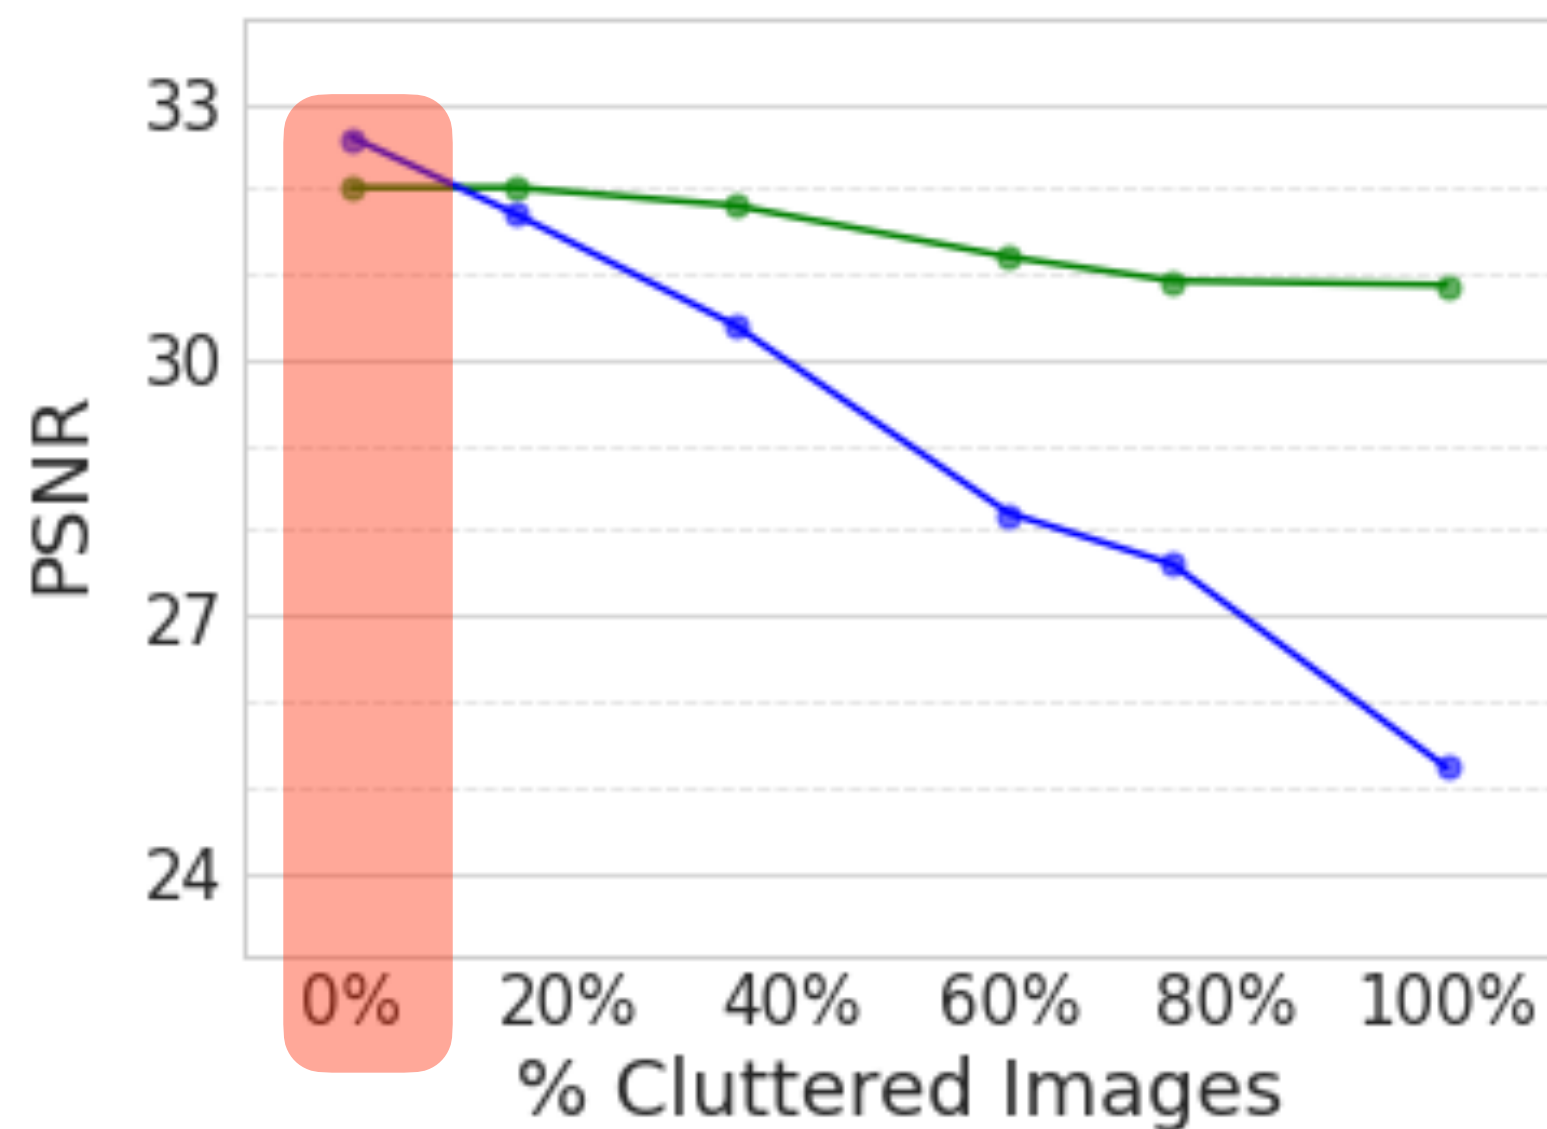



## Statistical inefficiency

Scenes take longer to reach peak accuracy. Clean scenes suffer a bit.

Cause: robust loss always considers a (fixed) portion of data as outliers

# Limitations



## Statistical inefficiency

Scenes take longer to reach peak accuracy. Clean scenes suffer a bit.

Cause: robust loss always considers a (fixed) portion of data as outliers

# Limitations



Lots of transients will confuse Colmap.  
No Colmap No Camera No NeRF.



Transients in the mirror -> The mirrored surface is replaced with blanks.

# RobustNeRF



Poster ID: THU-PM-002

<https://robustnerf.github.io>