

Looks Too Good To Be True: Analysis of Hallucinations in Generative Restoration Models

NeurIPS, December 2024

Regev Cohen

Idan kligvasser

Ehud Rivlin


Daniel Freedman

regevcohen@google.com



Generative Models for Inverse Problems


Recent generative models have achieved astonishing recovery performance in inverse problems.




* N. Cohen, H. Manor, Y. Bahat, T. Michaeli, *"From Posterior Sampling to Meaningful Diversity in Image Restoration"*.

Generative Models for Inverse Problems

Recent generative models have achieved astonishing recovery performance in inverse problems.




Yet, they are prone to hallucinations that escalates with improved perception, introducing uncertainty.




* N. Cohen, H. Manor, Y. Bahat, T. Michaeli, "From Posterior Sampling to Meaningful Diversity in Image Restoration".

Generative Models for Inverse Problems

Recent generative models have achieved astonishing recovery performance in inverse problems.




Can we design a recovery model of high perception and low uncertainty?



* N. Cohen, H. Manor, Y. Bahat, T. Michaeli, "From Posterior Sampling to Meaningful Diversity in Image Restoration".


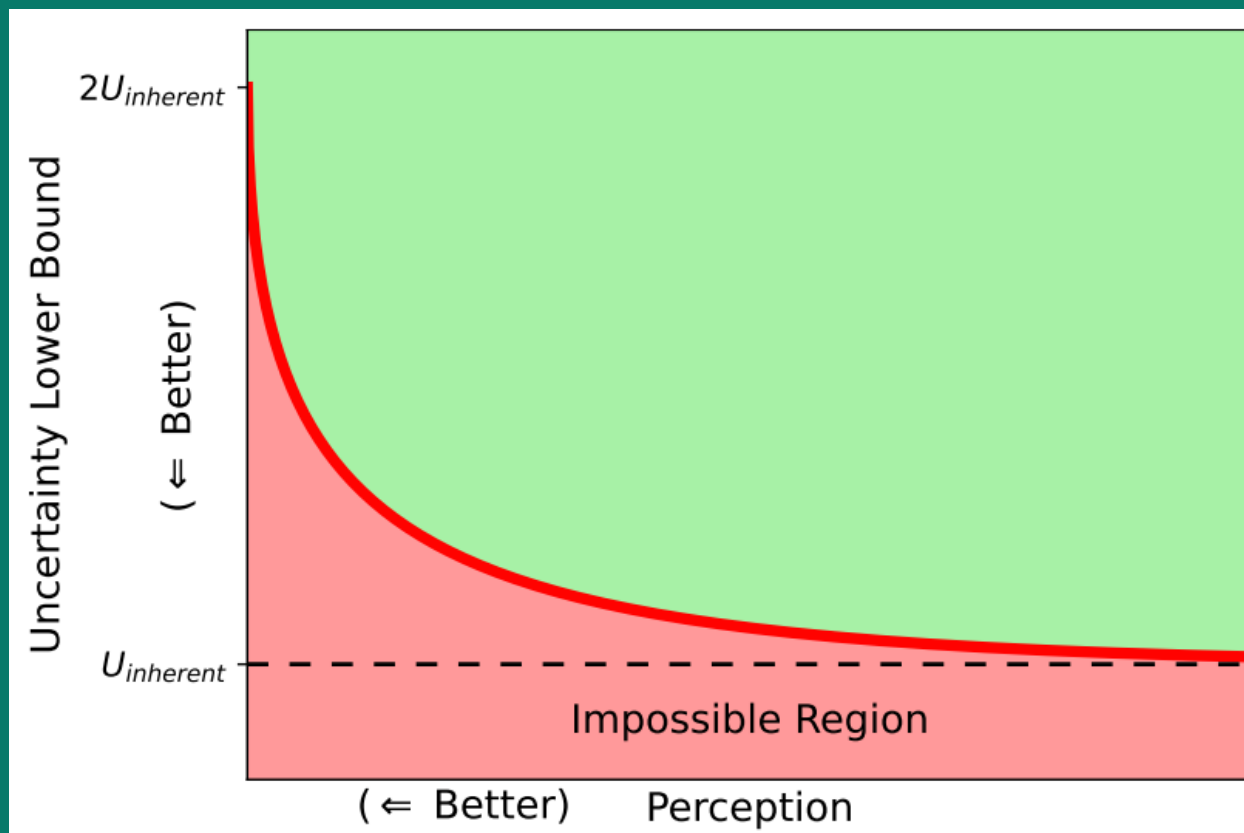
The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?



The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?



The Uncertainty–Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?

The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?

We introduce the **Uncertainty-Perception (UP)** function:

$$U(P) \triangleq \min_{p_{\hat{X}|Y}} \{N(\hat{X} - X|Y) : D_v(X, \hat{X}|Y) \leq P\}$$

The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?

To answer the above, we introduce the **Uncertainty-Perception (UP) function**:

$$U(P) \triangleq \min_{p_{\hat{X}|Y}} \{N(\hat{X} - X|Y) : D_v(X, \hat{X}|Y) \leq P\}$$

In words, the minimal attainable uncertainty for models achieving perceptual quality of at least P .

Perception $D_v(X, \hat{X}|Y)$ - divergence between the true and estimated distributions.

Uncertainty $N(\hat{X} - X|Y)$ - entropy power of the estimation error given the observation information.

The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?

To answer the above, we introduce the **Uncertainty-Perception (UP) function**:

$$U(P) \triangleq \min_{p_{\hat{X}|Y}} \{N(\hat{X} - X|Y) : D_v(X, \hat{X}|Y) \leq P\}$$

The UP function is **monotonically non-increasing** and displays the following important property:

$$\eta(P) \cdot N(X|Y) \leq U(P) \leq \eta(P) \cdot N(X_G|Y)$$

The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?

To answer the above, we introduce the **Uncertainty-Perception (UP) function**:

$$U(P) \triangleq \min_{p_{\hat{X}|Y}} \{N(\hat{X} - X|Y) : D_v(X, \hat{X}|Y) \leq P\}$$

The UP function is **monotonically non-increasing** and displays the following important property:

$$\eta(P) \cdot N(X|Y) \leq U(P) \leq \eta(P) \cdot N(X_G|Y)$$

- $N(X|Y)$ is the **inherent uncertainty** of the inverse problem.

The Uncertainty-Perception Tradeoff

Can we design a restoration model of high perception and low uncertainty?

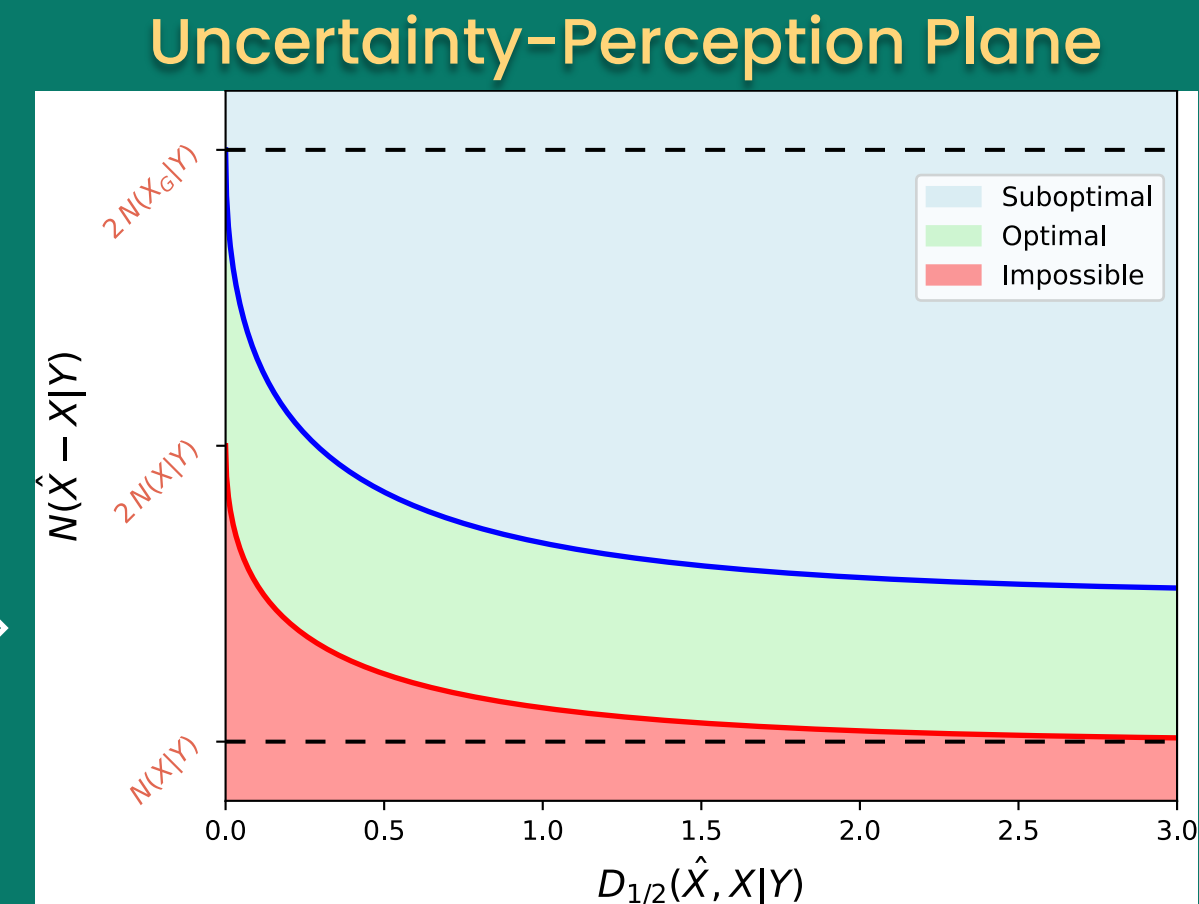
To answer the above, we introduce the **Uncertainty-Perception (UP) function**:

$$U(P) \triangleq \min_{p_{\hat{X}|Y}} \{N(\hat{X} - X|Y) : D_v(X, \hat{X}|Y) \leq P\}$$

The UP function is **monotonically non-increasing** and displays the following important property:

$$\eta(P) \cdot N(X|Y) \leq U(P) \leq \eta(P) \cdot N(X_G|Y)$$


- $N(X|Y)$ is the **inherent uncertainty** of the inverse problem.
- For Rényi divergence, $1 \leq \eta(P) \leq 2$ is a **convex function**,
categorizing algorithms into 3 distinct performance domains \Rightarrow



Impact of Dimensionality

The tradeoff intensifies in higher dimensions!

So minor improvements in perceptual quality leads to dramatic increase in uncertainty



The Distortion-Perception Tradeoff

- * In 2018, Blau & Michaeli proved that inverse methods display a tradeoff between perceptual quality and their reconstruction error for any distortion measure.

* Y.Blau, and T. Michaeli, "*The Perception-Distortion Tradeoff*".

The Distortion-Perception Tradeoff

- * In 2018, Blau & Michaeli proved that inverse methods display a tradeoff between perceptual quality and their reconstruction error for any distortion measure.

Consider the **Distortion-Perception function** for mean squared-error distortion

$$D(P) \triangleq \min_{p_{\hat{X}|Y}} \left\{ \frac{1}{\dim} \mathbb{E} \|\hat{X} - X\|^2 : D_v(X, \hat{X}|Y) \leq P \right\}.$$

* Y.Blau, and T. Michaeli, "The Perception-Distortion Tradeoff".

The Distortion-Perception Tradeoff

- * In 2018, Blau & Michaeli proved that inverse methods display a tradeoff between perceptual quality and their reconstruction error for any distortion measure.

Consider the **Distortion-Perception function** for mean squared-error distortion

$$D(P) \triangleq \min_{p_{\hat{X}|Y}} \left\{ \frac{1}{\dim} \mathbb{E} \|\hat{X} - X\|^2 : D_v(X, \hat{X}|Y) \leq P \right\}.$$

We prove that

$$D(P) \geq U(P).$$

* Y.Blau, and T. Michaeli, "The Perception-Distortion Tradeoff".

The Distortion-Perception Tradeoff

- * In 2018, Blau & Michaeli proved that inverse methods display a tradeoff between perceptual quality and their reconstruction error for any distortion measure.

Consider the **Distortion-Perception function** for mean squared-error distortion

$$D(P) \triangleq \min_{p_{\hat{X}|Y}} \left\{ \frac{1}{\dim} \mathbb{E} \|\hat{X} - X\|^2 : D_v(X, \hat{X}|Y) \leq P \right\}.$$

We prove that

$$D(P) \geq U(P).$$

The uncertainty-perception tradeoff



the distortion-perception tradeoff.

* Y.Blau, and T. Michaeli, "The Perception-Distortion Tradeoff".

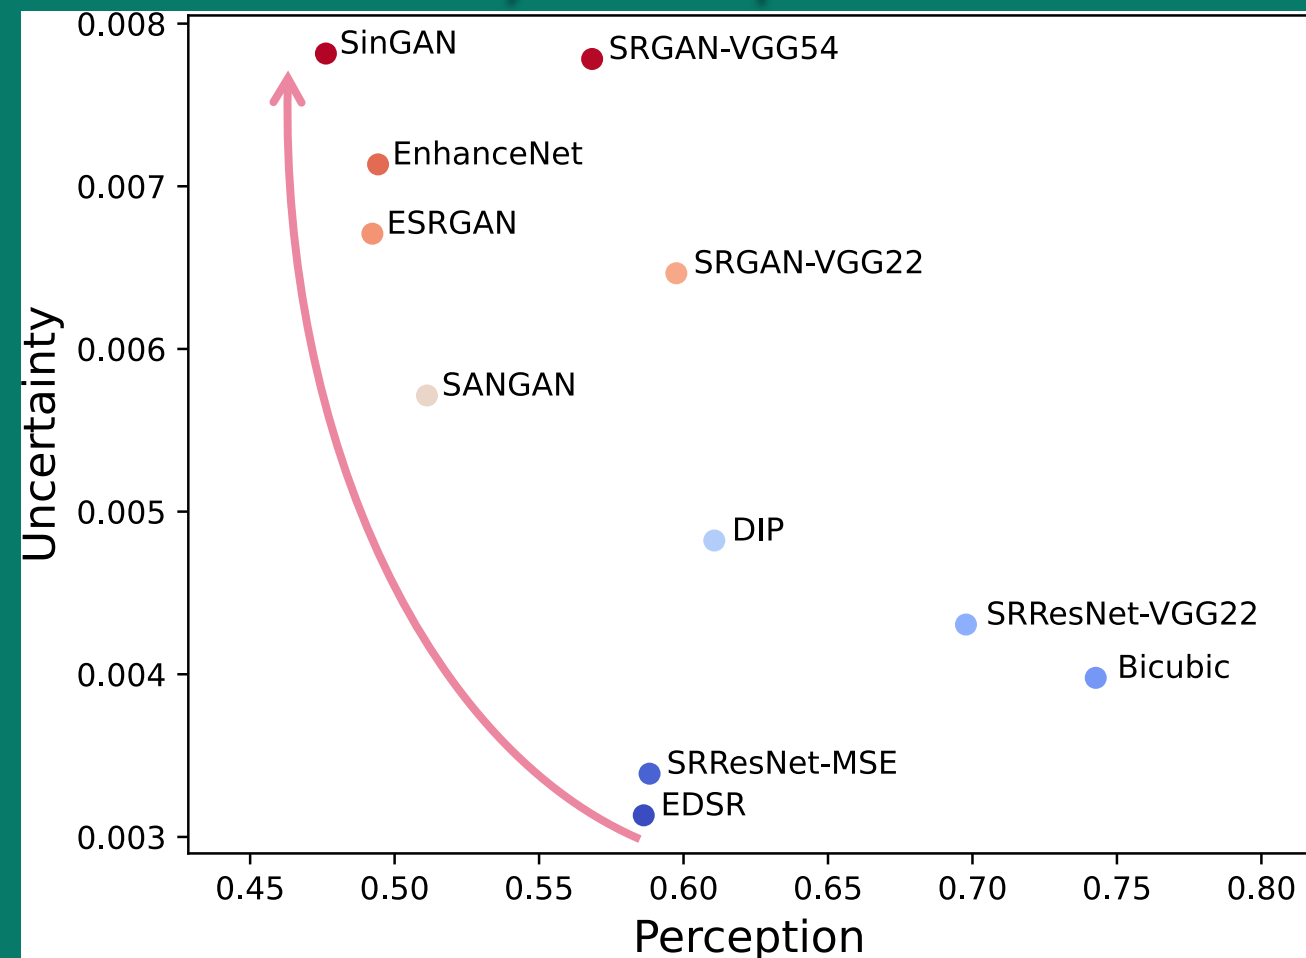
Experiments

To show the UP tradeoff and its connection to distortion in practice, we evaluate well-known restoration algorithms -

Experiments

To show the UP tradeoff and its connection to distortion in practice, we evaluate well-known restoration algorithms -


Uncertainty-Perception Tradeoff




Experiments

To show the UP tradeoff and its connection to distortion in practice, we evaluate well-known restoration algorithms -

Uncertainty-Perception Tradeoff

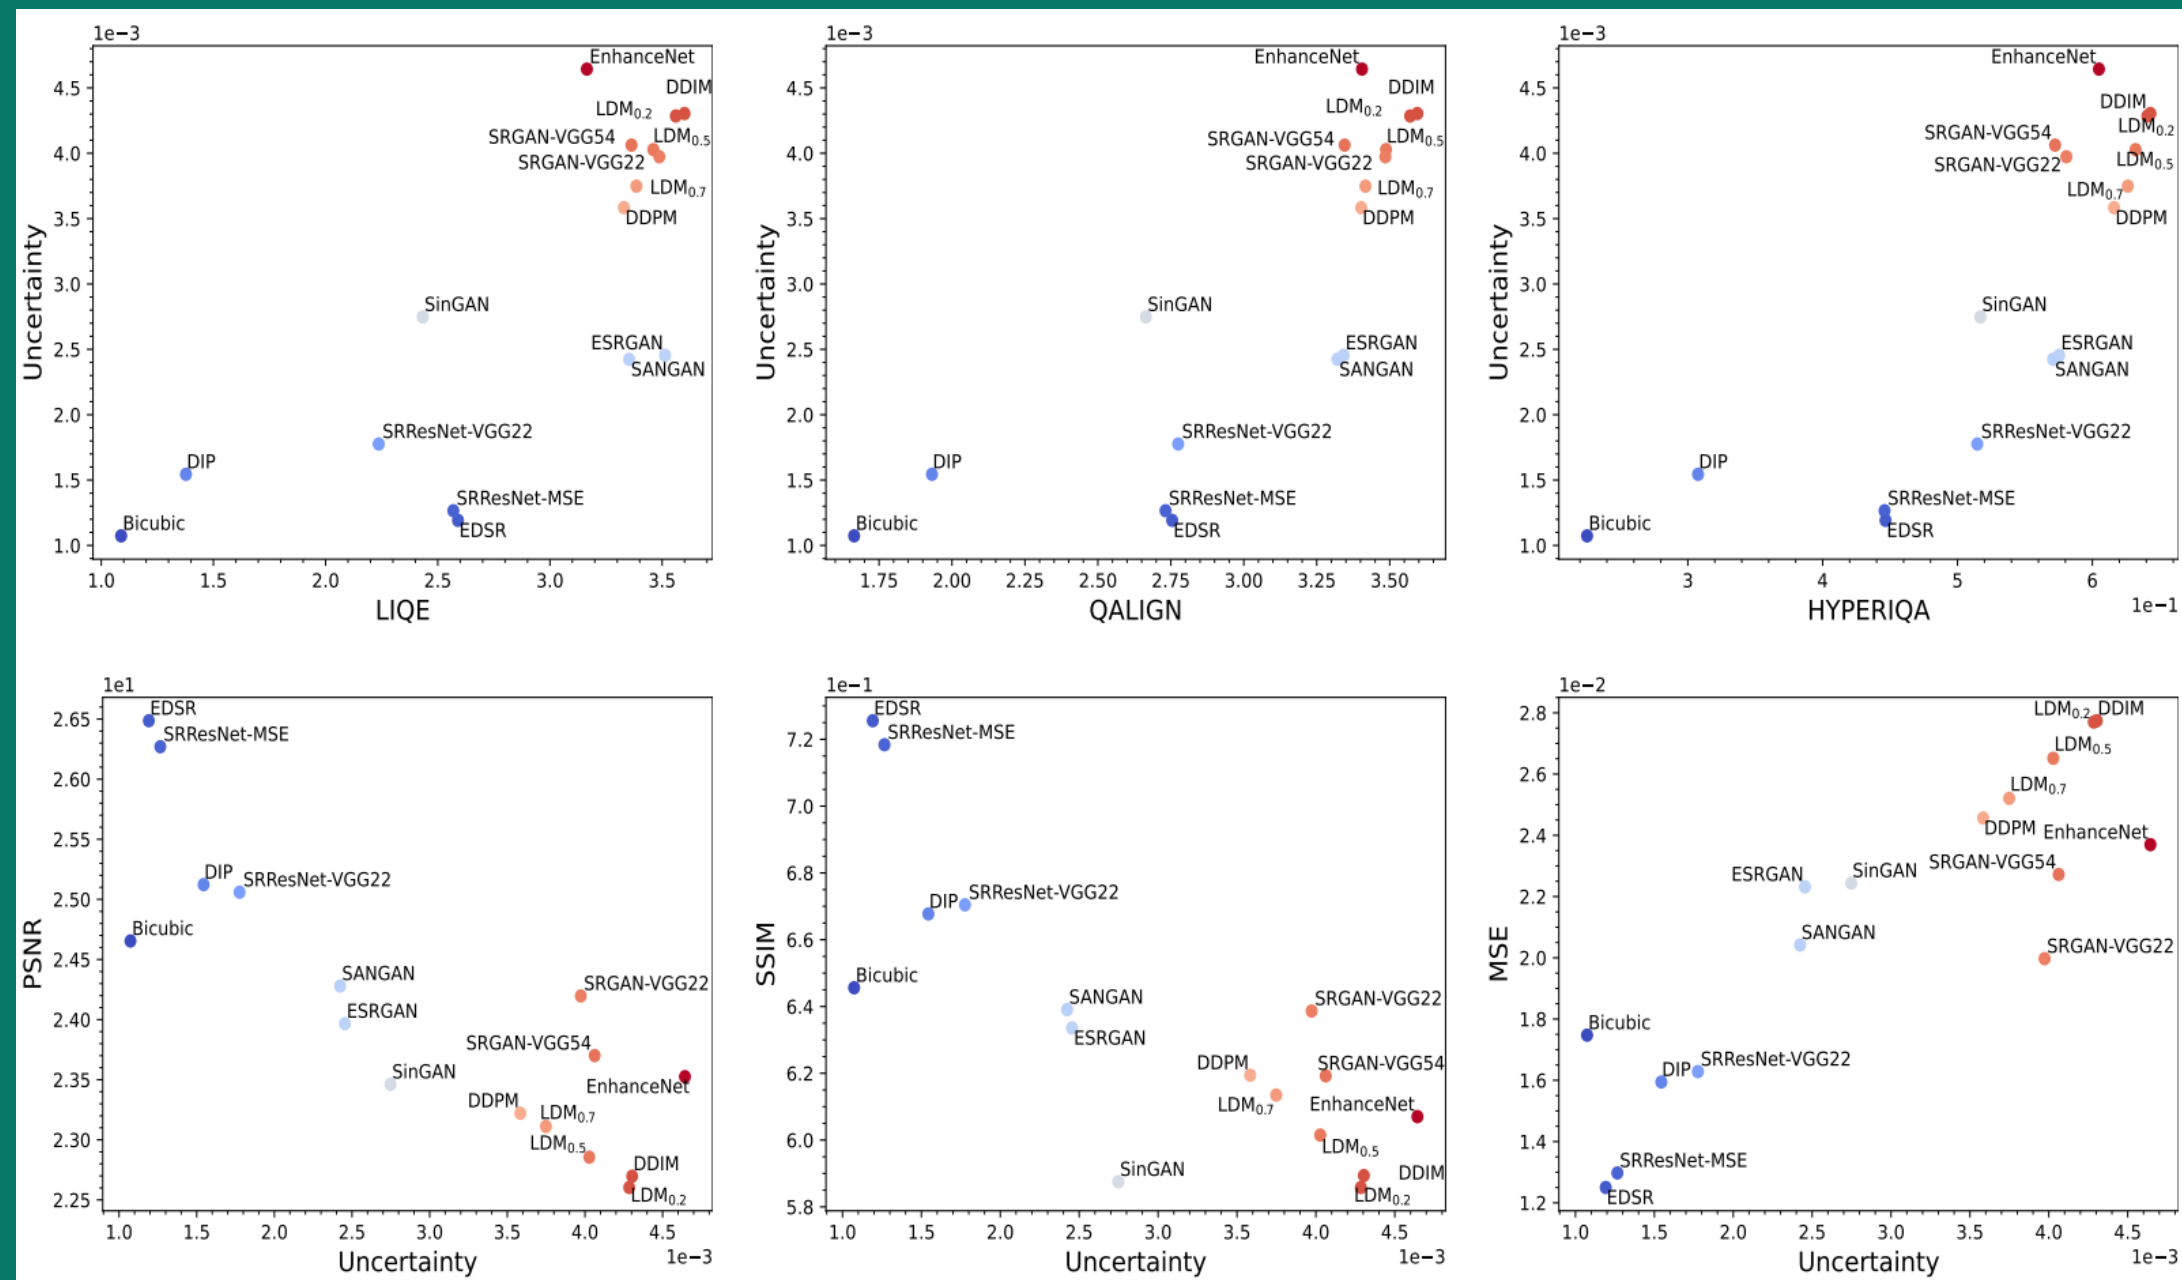


Uncertainty-Distortion Inteplay



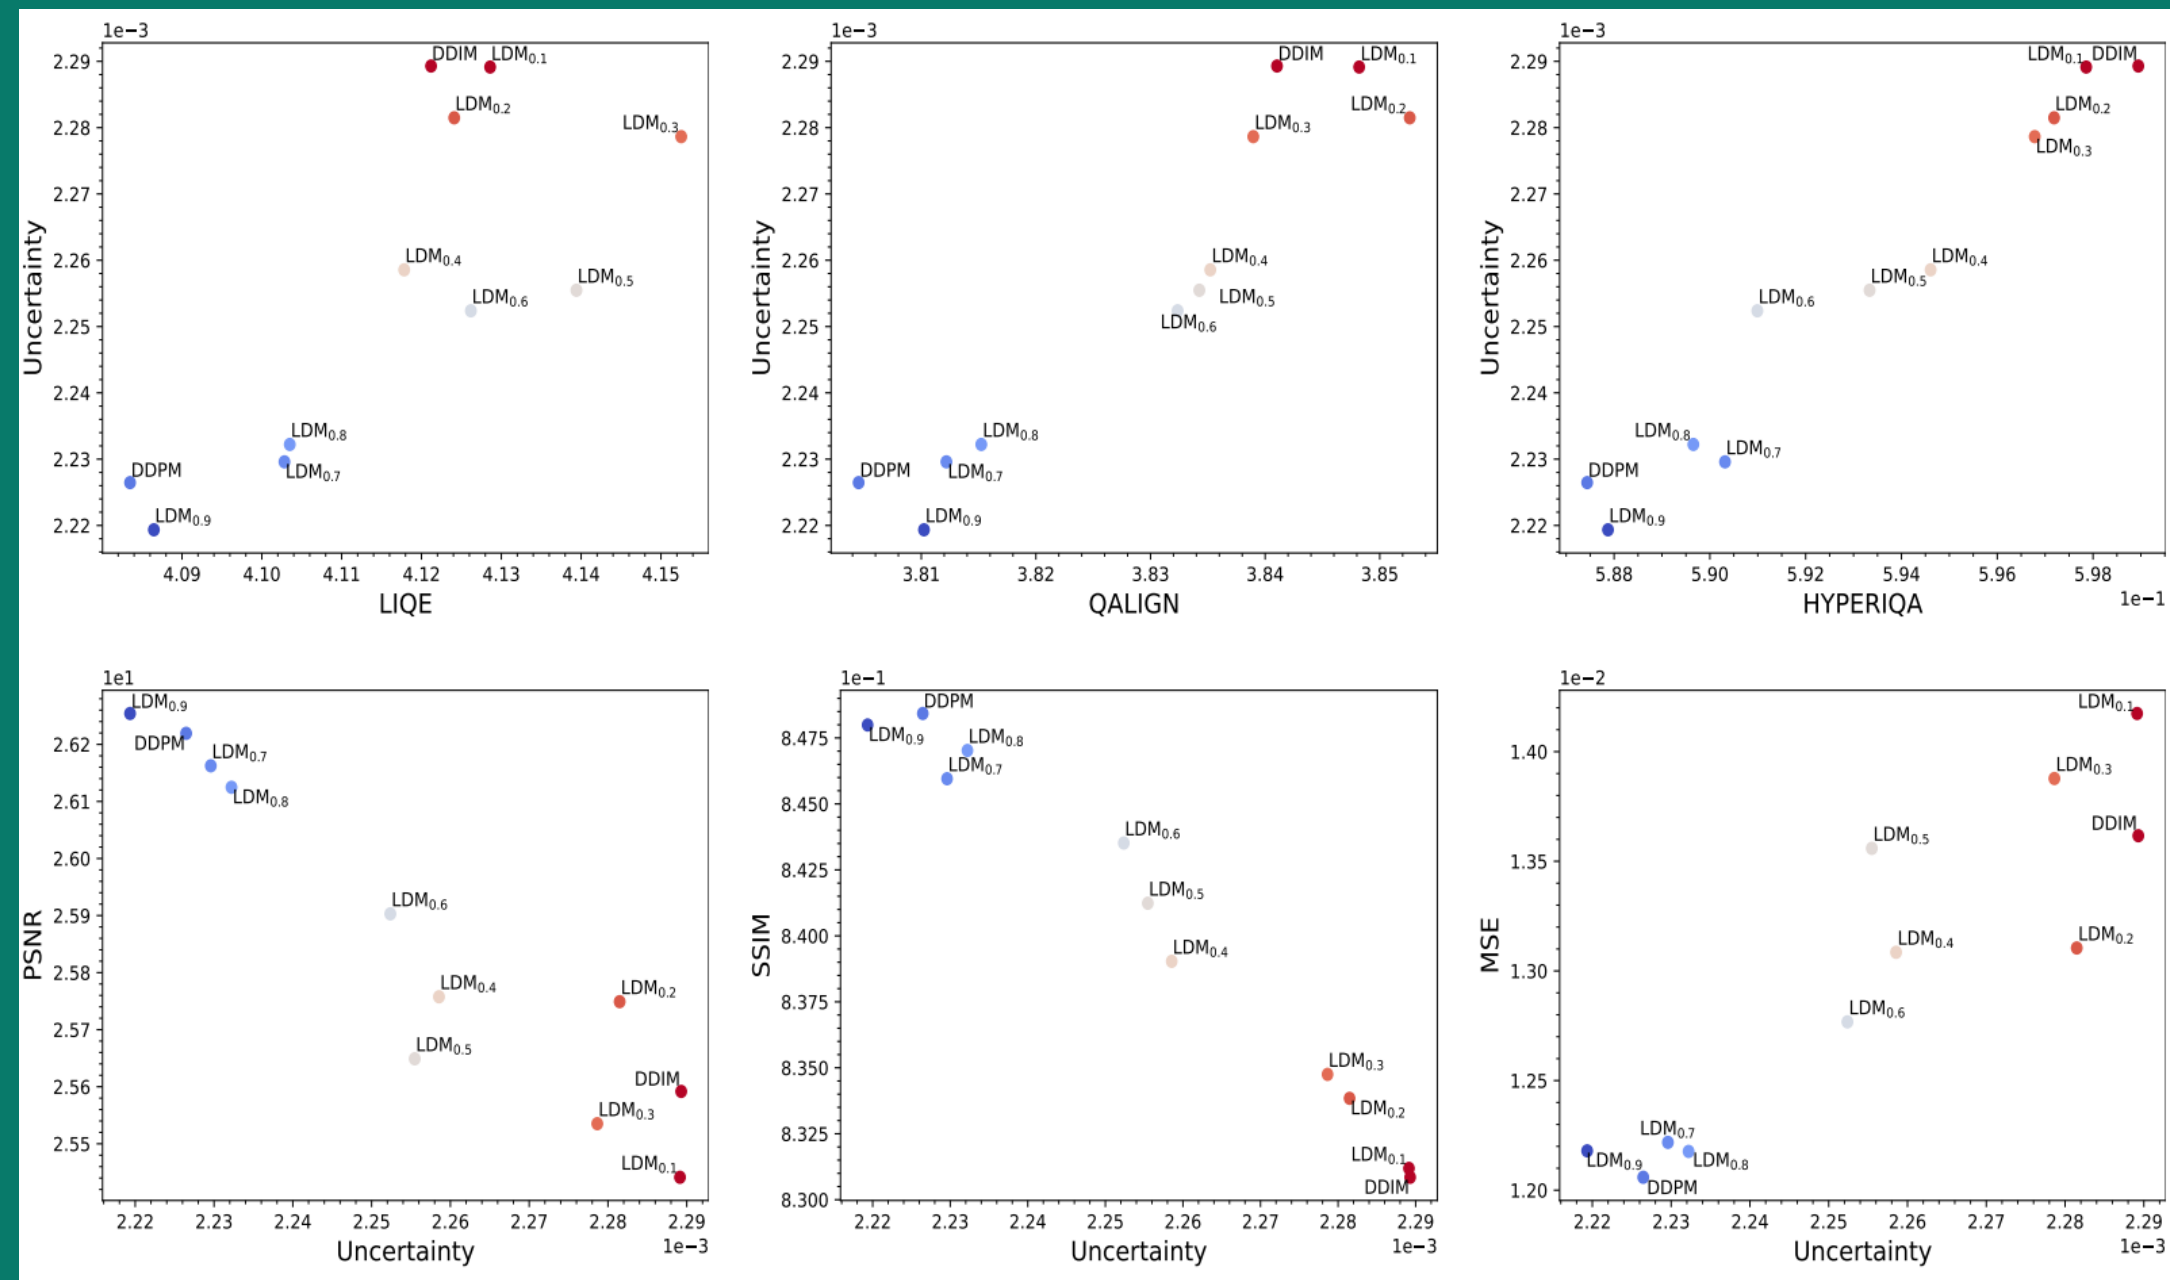
Experiments

To show the UP tradeoff and its connection to distortion in practice, we evaluate well-known restoration algorithms -




Experiments

To show the UP tradeoff and its connection to distortion in practice, we evaluate well-known restoration algorithms -



Experiments

Raising awareness of this tradeoff is crucial for promoting safety and reliability over purely perceptual.



Thank you



regevcohen@google.com