Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Stars do not make a noise

...to be silent is to be beautiful

– James Stephens

15! :)

A ~~25~~ mins journey from the classics to the deep learning era of image denoising

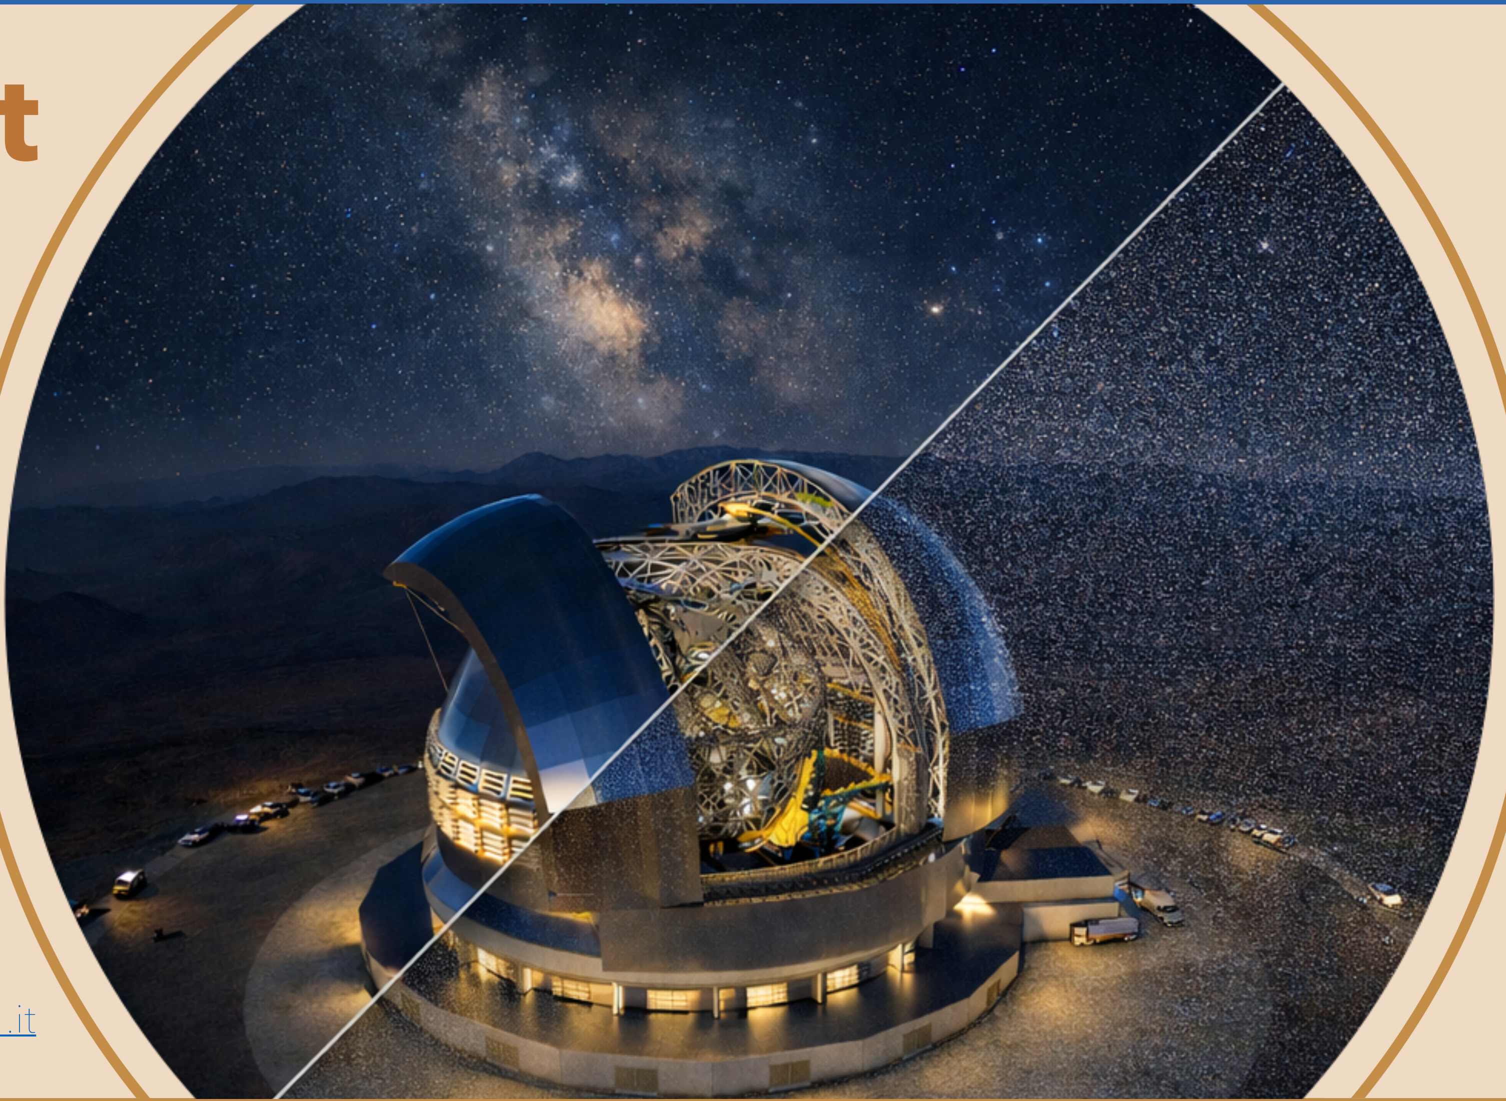

Pierpaolo Brutti | Department of Statistics
pierpaolo.brutti@uniroma1.it

PNRR
Missione 4 • Componente 2
Investimento 3.1

STILES – IR0000034
CUP C33C22000640006

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Pierpaolo Brutti
Department of Statistics Science
Sapienza University of Rome

PNRR
Missione 4 • Componente 2
Investimento 3.1

STILES – IR0000034
CUP C33C22000640006

Spot the intruder

Not exactly a navigated astronomer

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

…it's 2026, so….

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

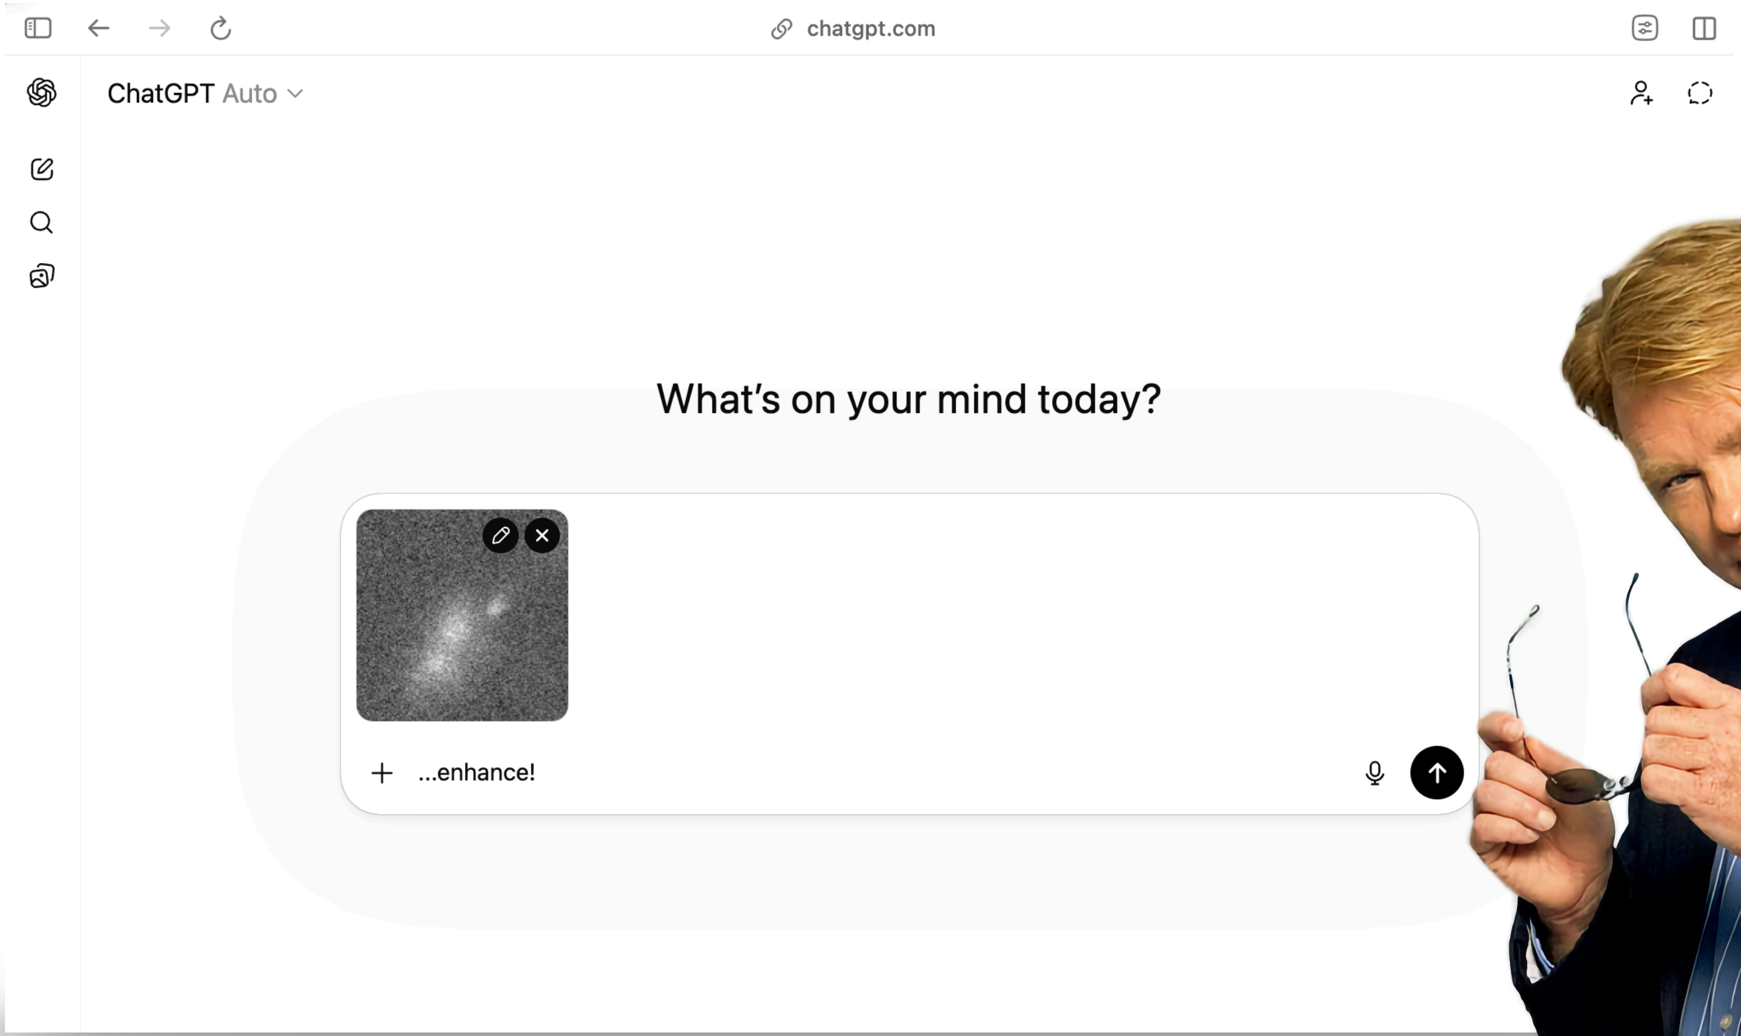

PNRR
Missione 4 • Componente 2
Investimento 3.1

STILES – IR0000034
CUP C33C22000640006

https://tinyurl.com/27e7ffdr

chatgpt.com

ChatGPT Auto

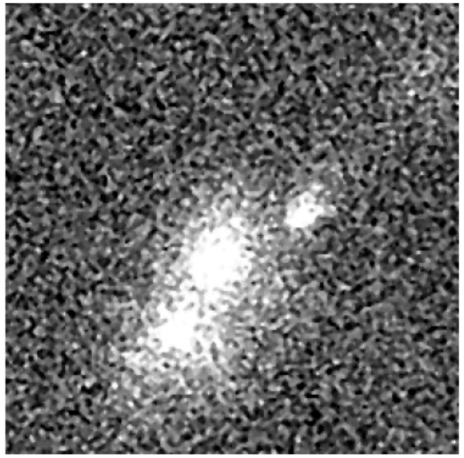

I enhanced the image by reducing noise, increasing contrast, and sharpening the bright structures so the central features stand out more clearly.

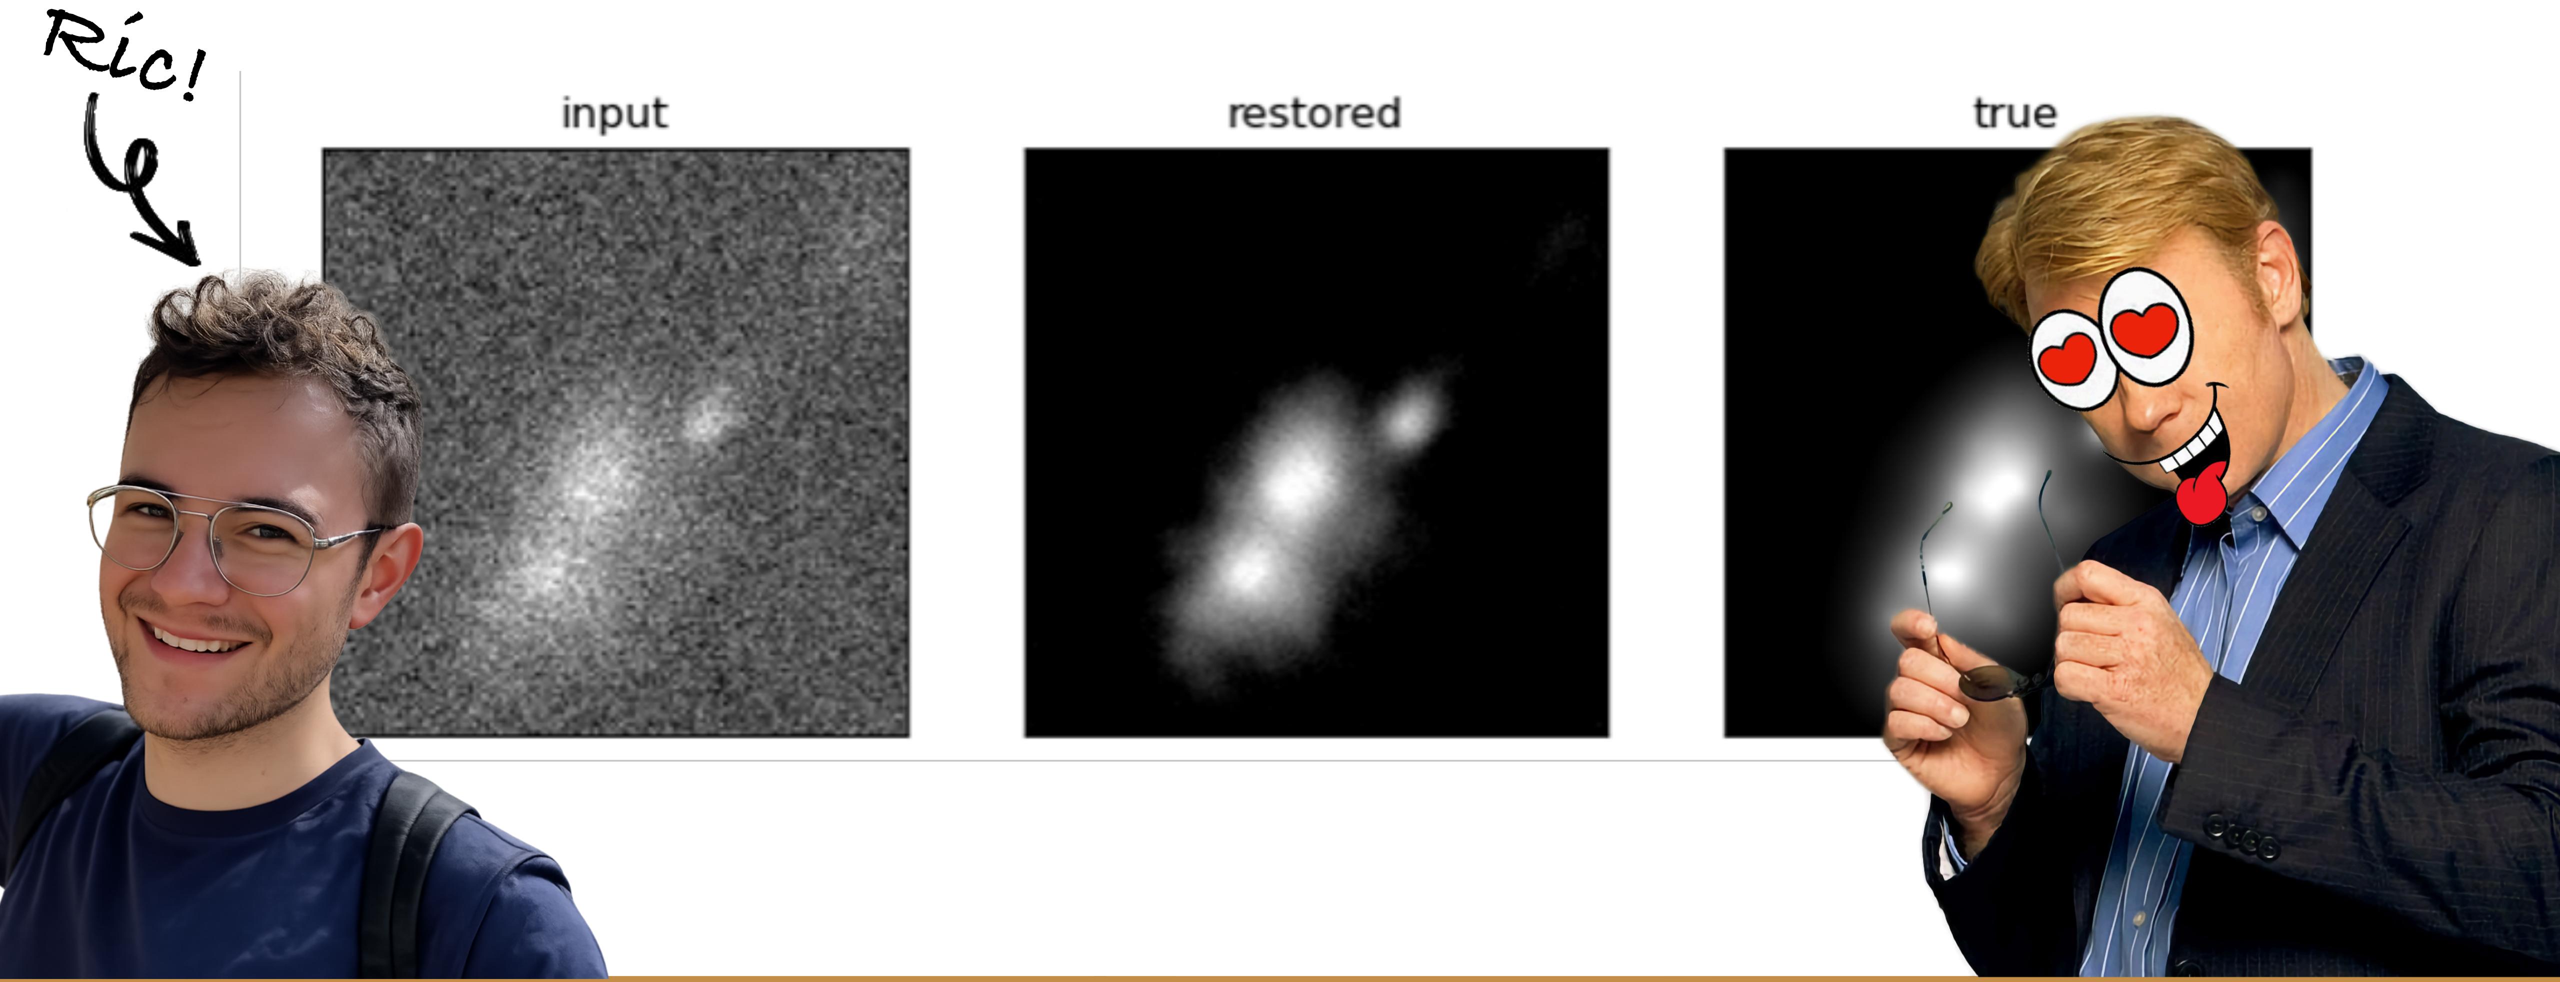
Ric!
input
restored
true

Ric!

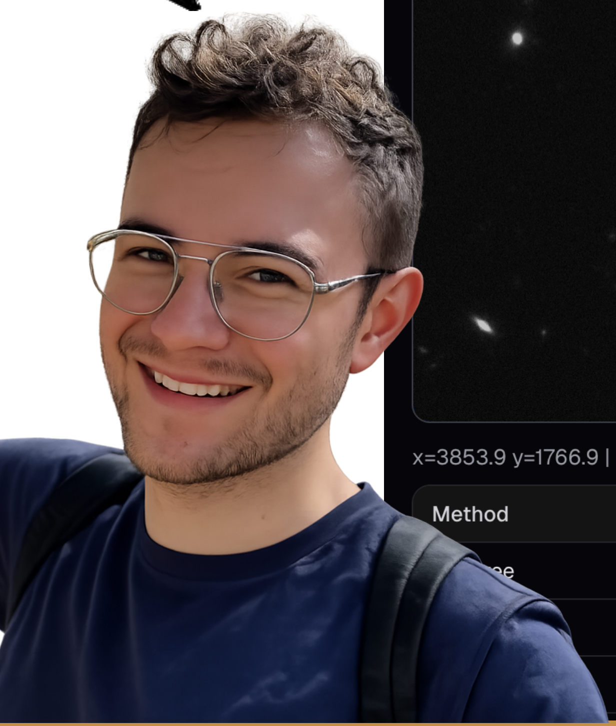

Astro Denoising (FORECAST)

linear zscale

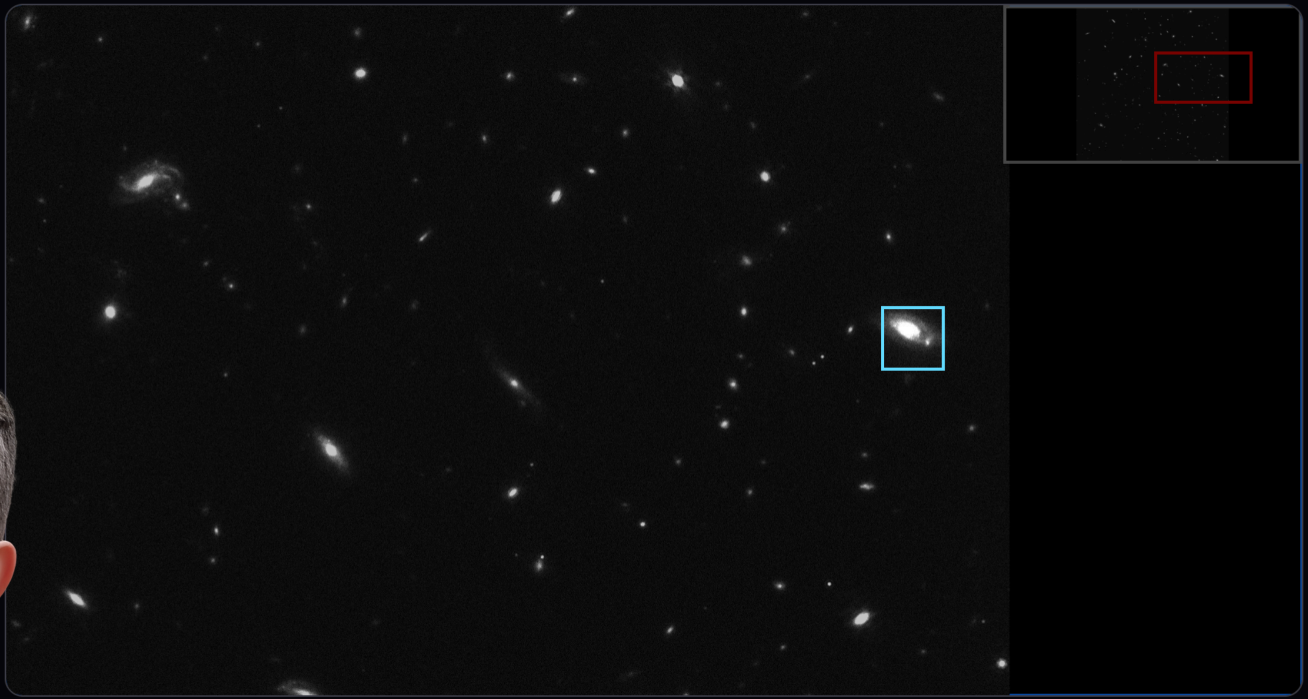

x=3853.9 y=1766.9 | patch=(3840, 1760) | pid=6995

Method	MSE	PSNR	SSIM
ee	1.687e-6	39.44	0.9530
	4.044e-6	35.64	0.9180
	5.928e-6	33.98	0.7347

Patch 128x128

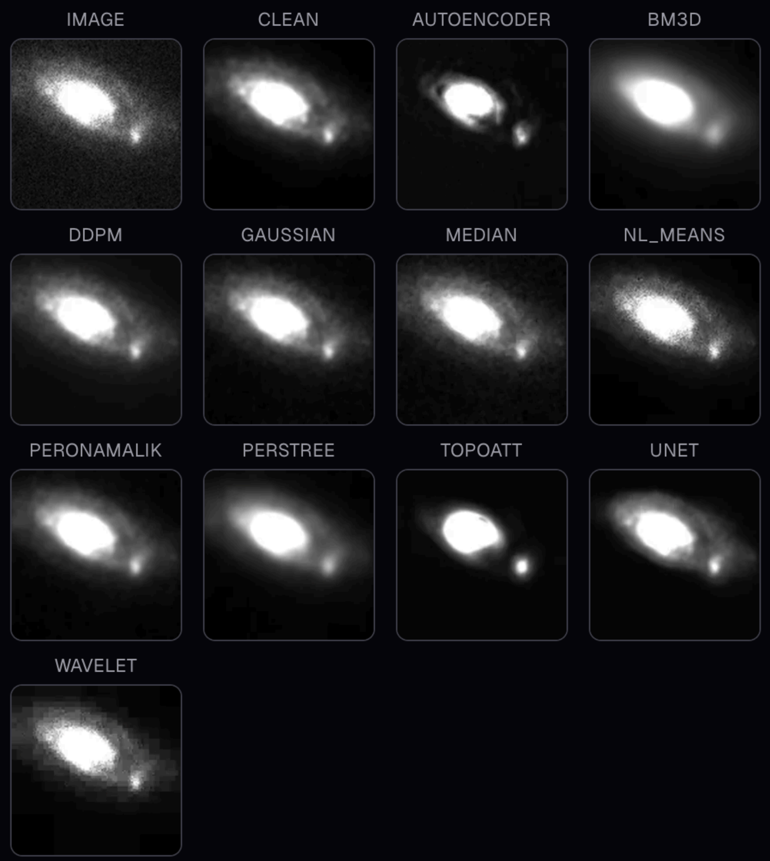

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Ric!

https://tinyurl.com/2puyc73y

…let's cut to the chase…

Main source: Elad et al. (2023). Image Denoising: The Deep Learning Revolution and Beyond – A Survey Paper

Disclaimer:

I'll make tons of approximations, be patient, thx!

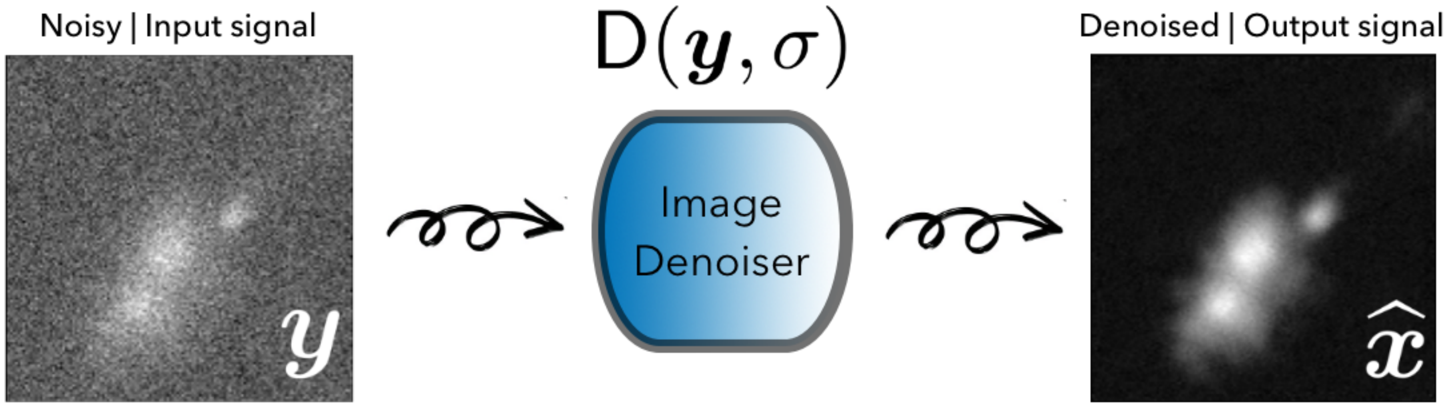

Image Denoising

Finanziato dall'Unione europea
NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA
INAF
ISTITUTO NAZIONALE DI ASTROFISICA

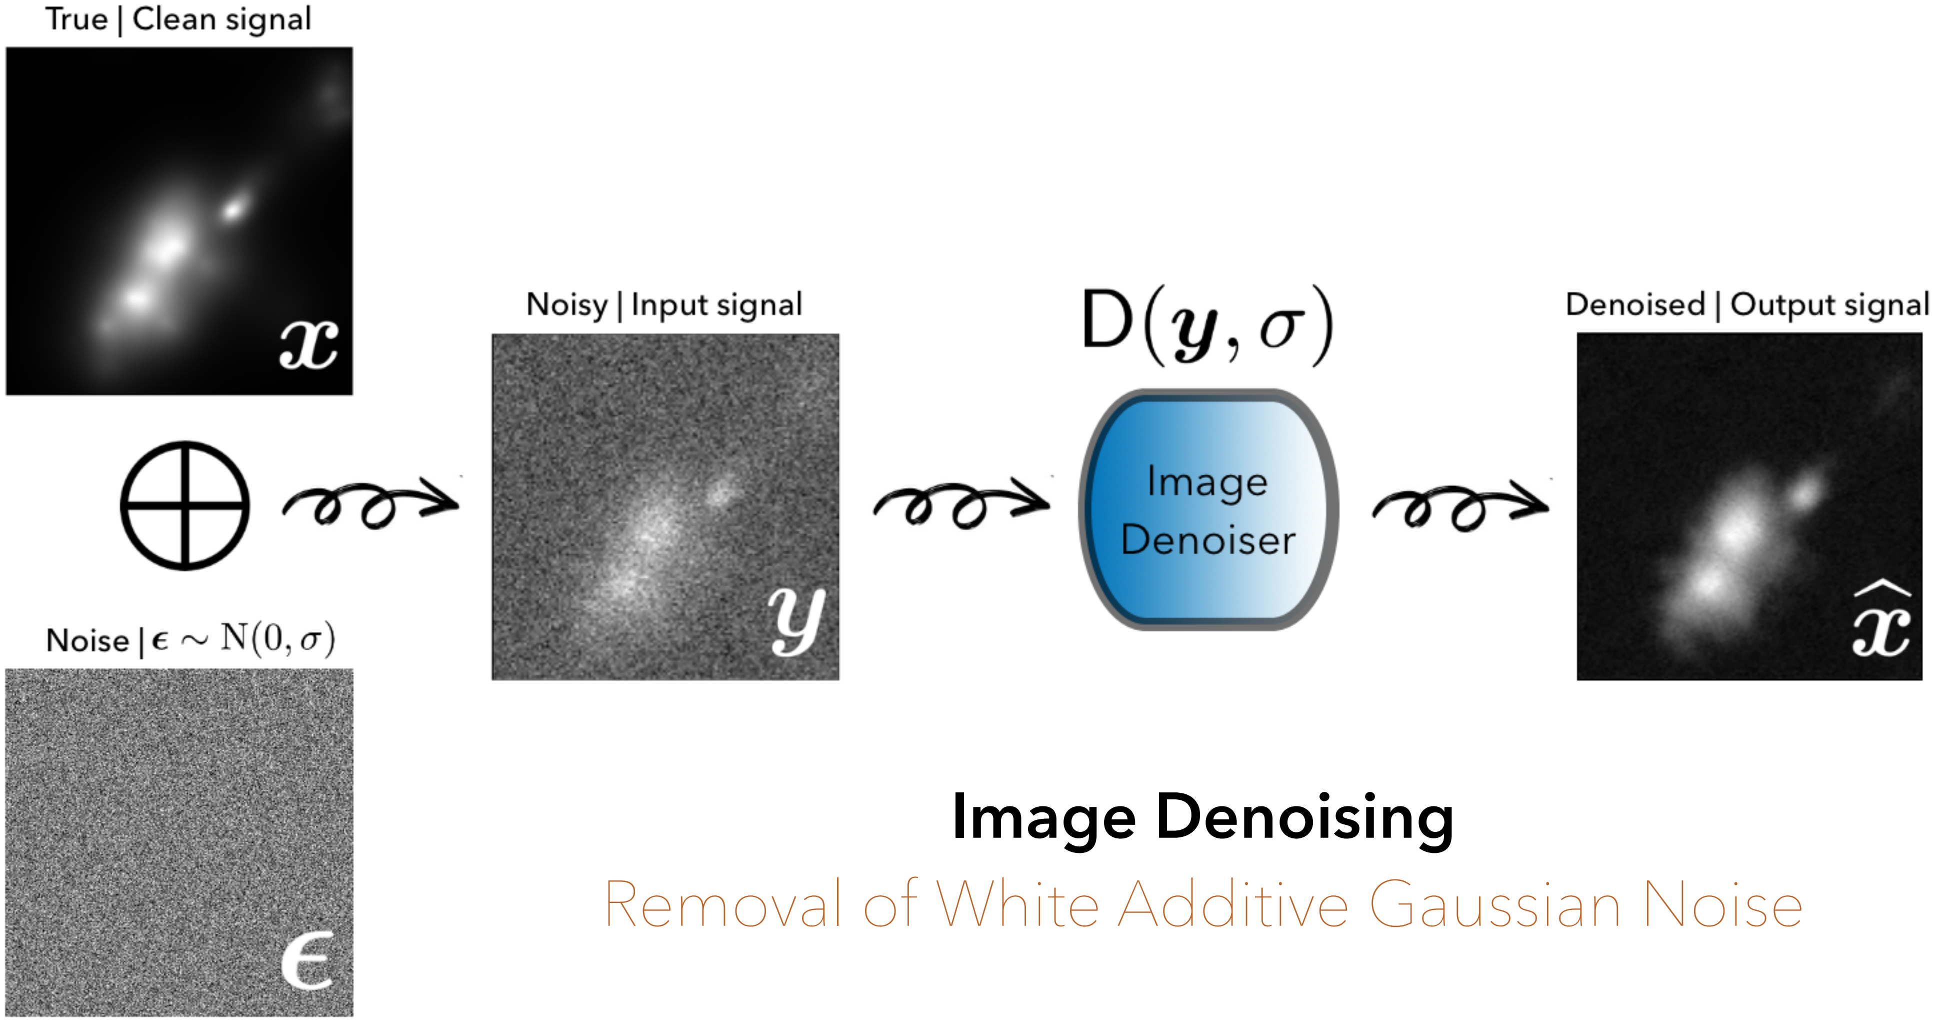

Image Denoising

Removal of White Additive Gaussian Noise

- **Poisson noise** useful but well approximated by Gaussian in high/low photon count scenarios, possibly Anscombe's variance-stabilizing transformation.
- Many developments under gaussianity can be converted to other noise models, also of mixed type (Poisson-Gaussian)
- Optimal denoisers under gaussianity play a key role in modern, **generative models**!

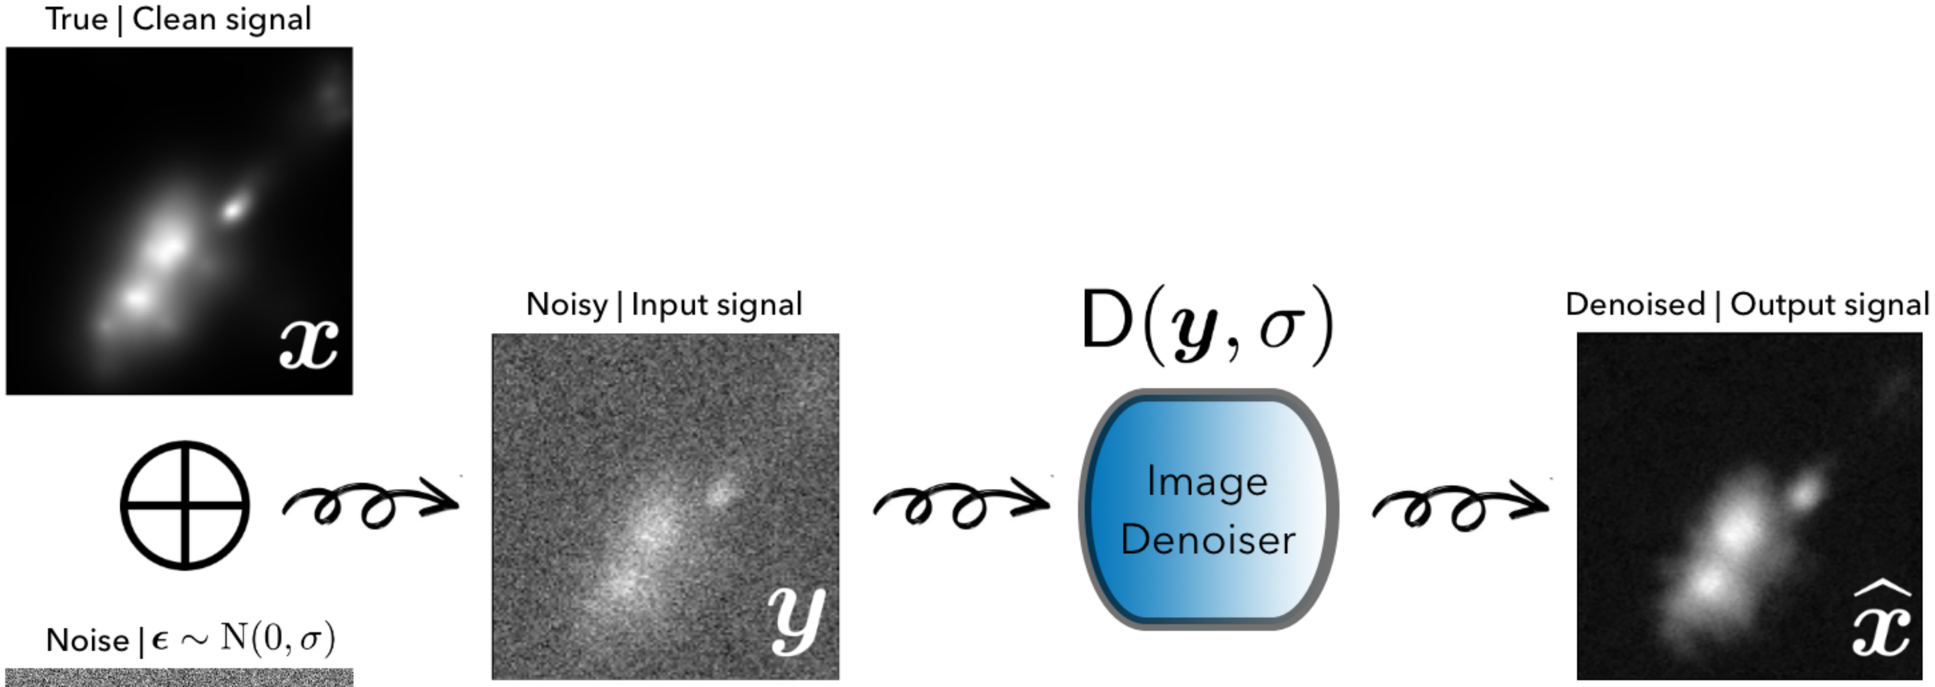

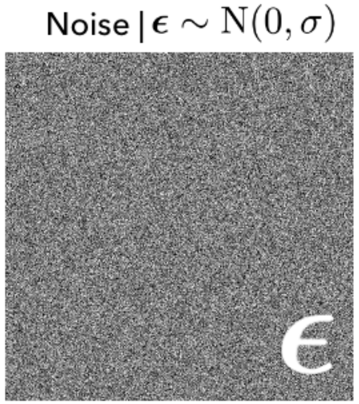

Image Denoising

Removal of White Additive Gaussian Noise

Recap of the basics

Observation Model

$$\boldsymbol{y} = \boldsymbol{x} + \boldsymbol{\epsilon} \quad \boldsymbol{\epsilon} \sim \mathrm{N}(\mathbf{0}, \sigma^2 \mathbb{I})$$

Ideal: MSE-Optimal denoiser (MMSE)

$$\widehat{\boldsymbol{x}}_{\texttt{MMSE}} = \underset{\texttt{All } \widehat{\boldsymbol{x}}}{\operatorname{argmin}} \, \mathbb{E}\left(\|\boldsymbol{x} - \widehat{\boldsymbol{x}}\|^2\right) = \mathbb{E}(\boldsymbol{x}|\boldsymbol{y}) = \int \boldsymbol{x} \underbrace{p(\boldsymbol{x}|\boldsymbol{y})}_{\texttt{unknown}} \mathrm{d}\boldsymbol{x}.$$

Likelihood:

$$p(\boldsymbol{y}|\boldsymbol{x}) \propto \exp\left\{-\frac{\|\boldsymbol{y} - \boldsymbol{x}\|^2}{2\sigma^2}\right\}$$

Posterior:

$$p(\boldsymbol{x}|\boldsymbol{y}) \propto p(\boldsymbol{y}|\boldsymbol{x}) \, p(\boldsymbol{x})$$

where $p(\boldsymbol{x})$ is the **image prior**

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Recap of the basics

Observation Model

$$\boldsymbol{y} = \boldsymbol{x} + \epsilon \qquad \epsilon \sim \mathrm{N}(\mathbf{0}, \sigma^2 \mathbb{I})$$

Ideal: MSE-Optimal denoiser (MMSE)

$$\widehat{\boldsymbol{x}}_{\texttt{MMSE}} = \underset{\texttt{All } \widehat{\boldsymbol{x}}}{\operatorname{argmin}} \, \mathbb{E}\left(\|\boldsymbol{x} - \widehat{\boldsymbol{x}}\|^2 \right) = \mathbb{E}(\boldsymbol{x}|\boldsymbol{y}) = \int \boldsymbol{x} \, \underbrace{p(\boldsymbol{x}|\boldsymbol{y})}_{\texttt{unknown}} \, \mathrm{d}\boldsymbol{x}.$$

Maximum Likelihood:

$$\widehat{\boldsymbol{x}}_{\texttt{MLE}} = \underset{\texttt{All } \boldsymbol{x}}{\operatorname{argmax}} \, p(\boldsymbol{y}|\boldsymbol{x}) = \boldsymbol{y}$$

No denoising $\rightsquigarrow$ we need a prior.

Maximum a Posterior:

$$\widehat{\boldsymbol{x}}_{\texttt{MAP}} = \underset{\texttt{All } \boldsymbol{x}}{\operatorname{argmax}} \, p(\boldsymbol{x}|\boldsymbol{y})$$

$$= \underset{\texttt{All } \boldsymbol{x}}{\operatorname{argmin}} \, \frac{1}{2}\|\boldsymbol{y} - \boldsymbol{x}\|^2 - \log p(\boldsymbol{x})$$

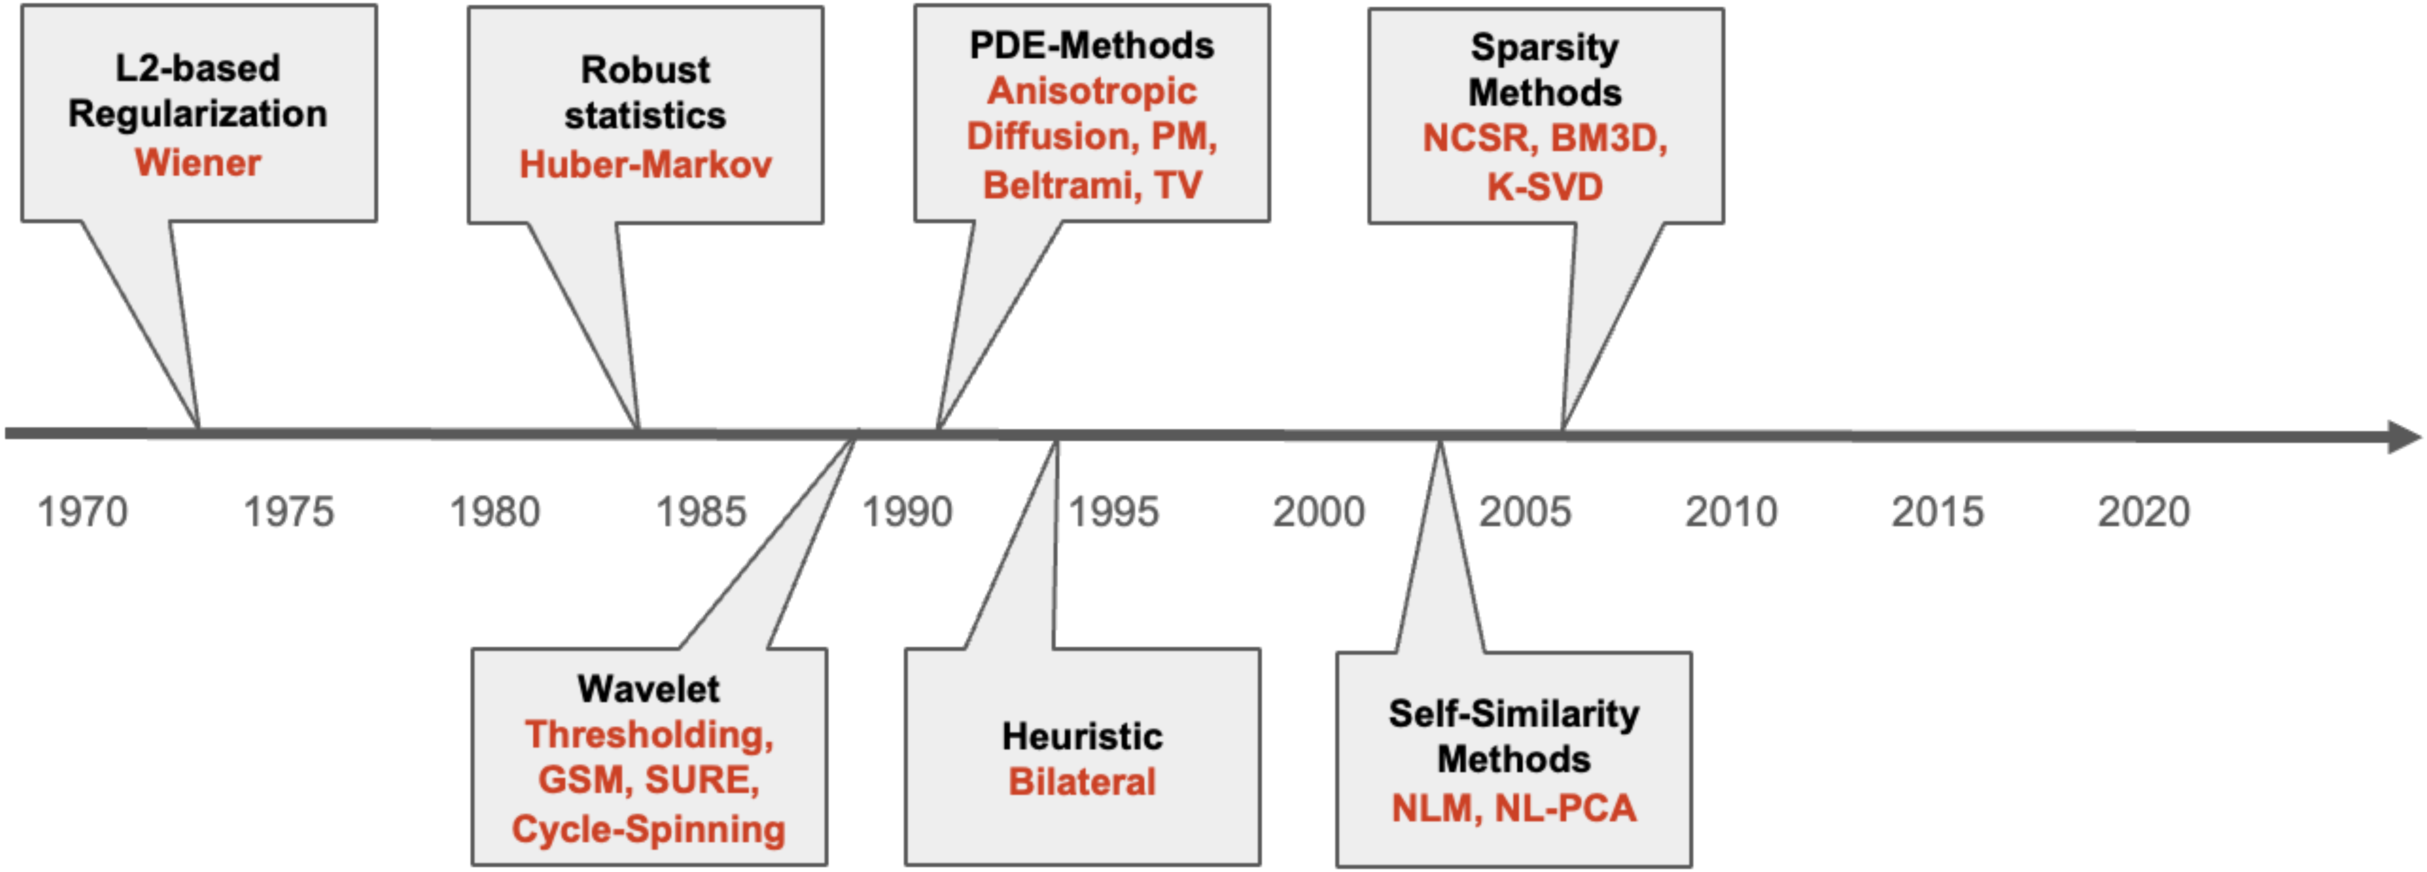
L2-based Regularization
Wiener
Robust statistics
Huber-Markov
PDE-Methods
Anisotropic Diffusion, PM, Beltrami, TV
Sparsity Methods
NCSR, BM3D, K-SVD
1970
1975
1980
1985
1990
1995
2000
2005
2010
2015
2020
Wavelet
Thresholding, GSM, SURE, Cycle-Spinning
Heuristic
Bilateral
Self-Similarity Methods
NLM, NL-PCA

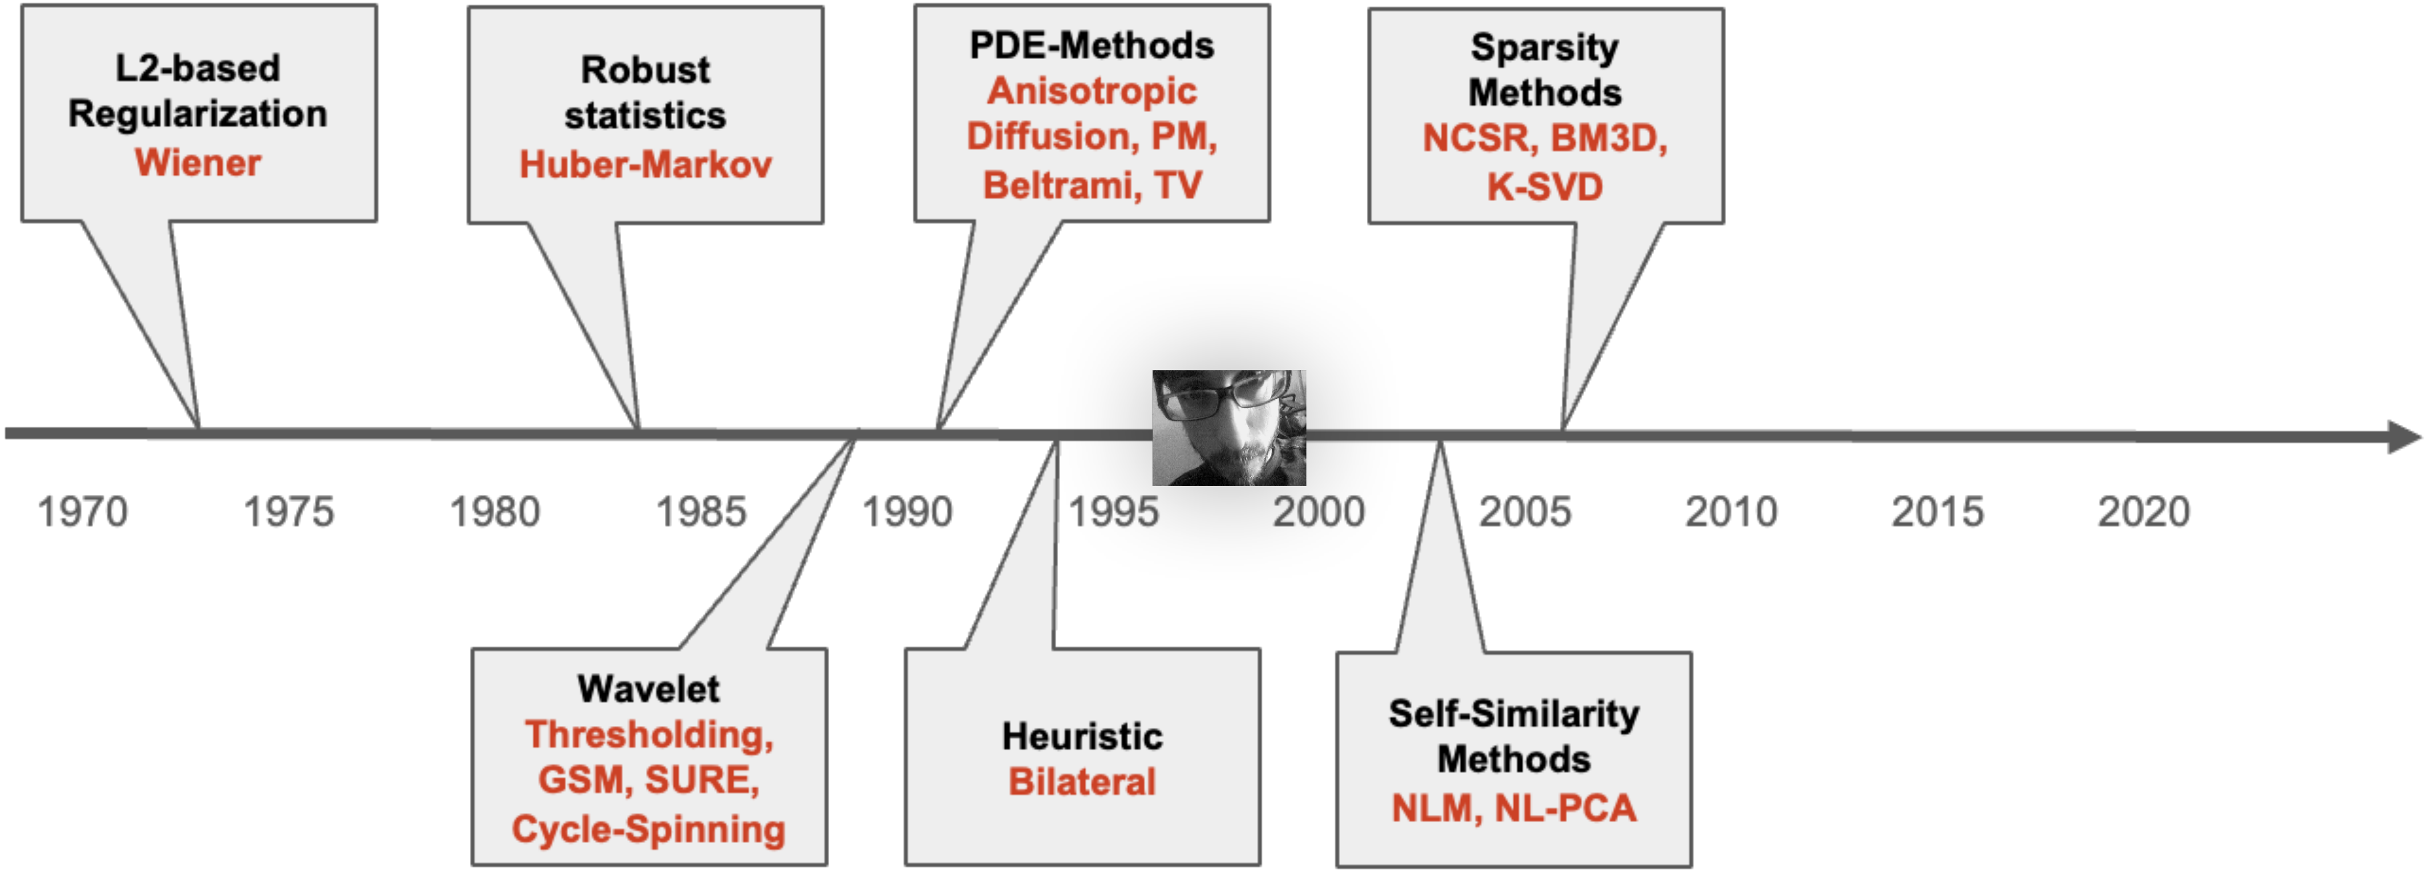
L2-based Regularization
Wiener
Robust statistics
Huber-Markov
PDE-Methods
Anisotropic Diffusion, PM, Beltrami, TV
Sparsity Methods
NCSR, BM3D, K-SVD
1970
1975
1980
1985
1990
1995
2000
2005
2010
2015
2020
Wavelet
Thresholding, GSM, SURE, Cycle-Spinning
Heuristic
Bilateral
Self-Similarity Methods
NLM, NL-PCA

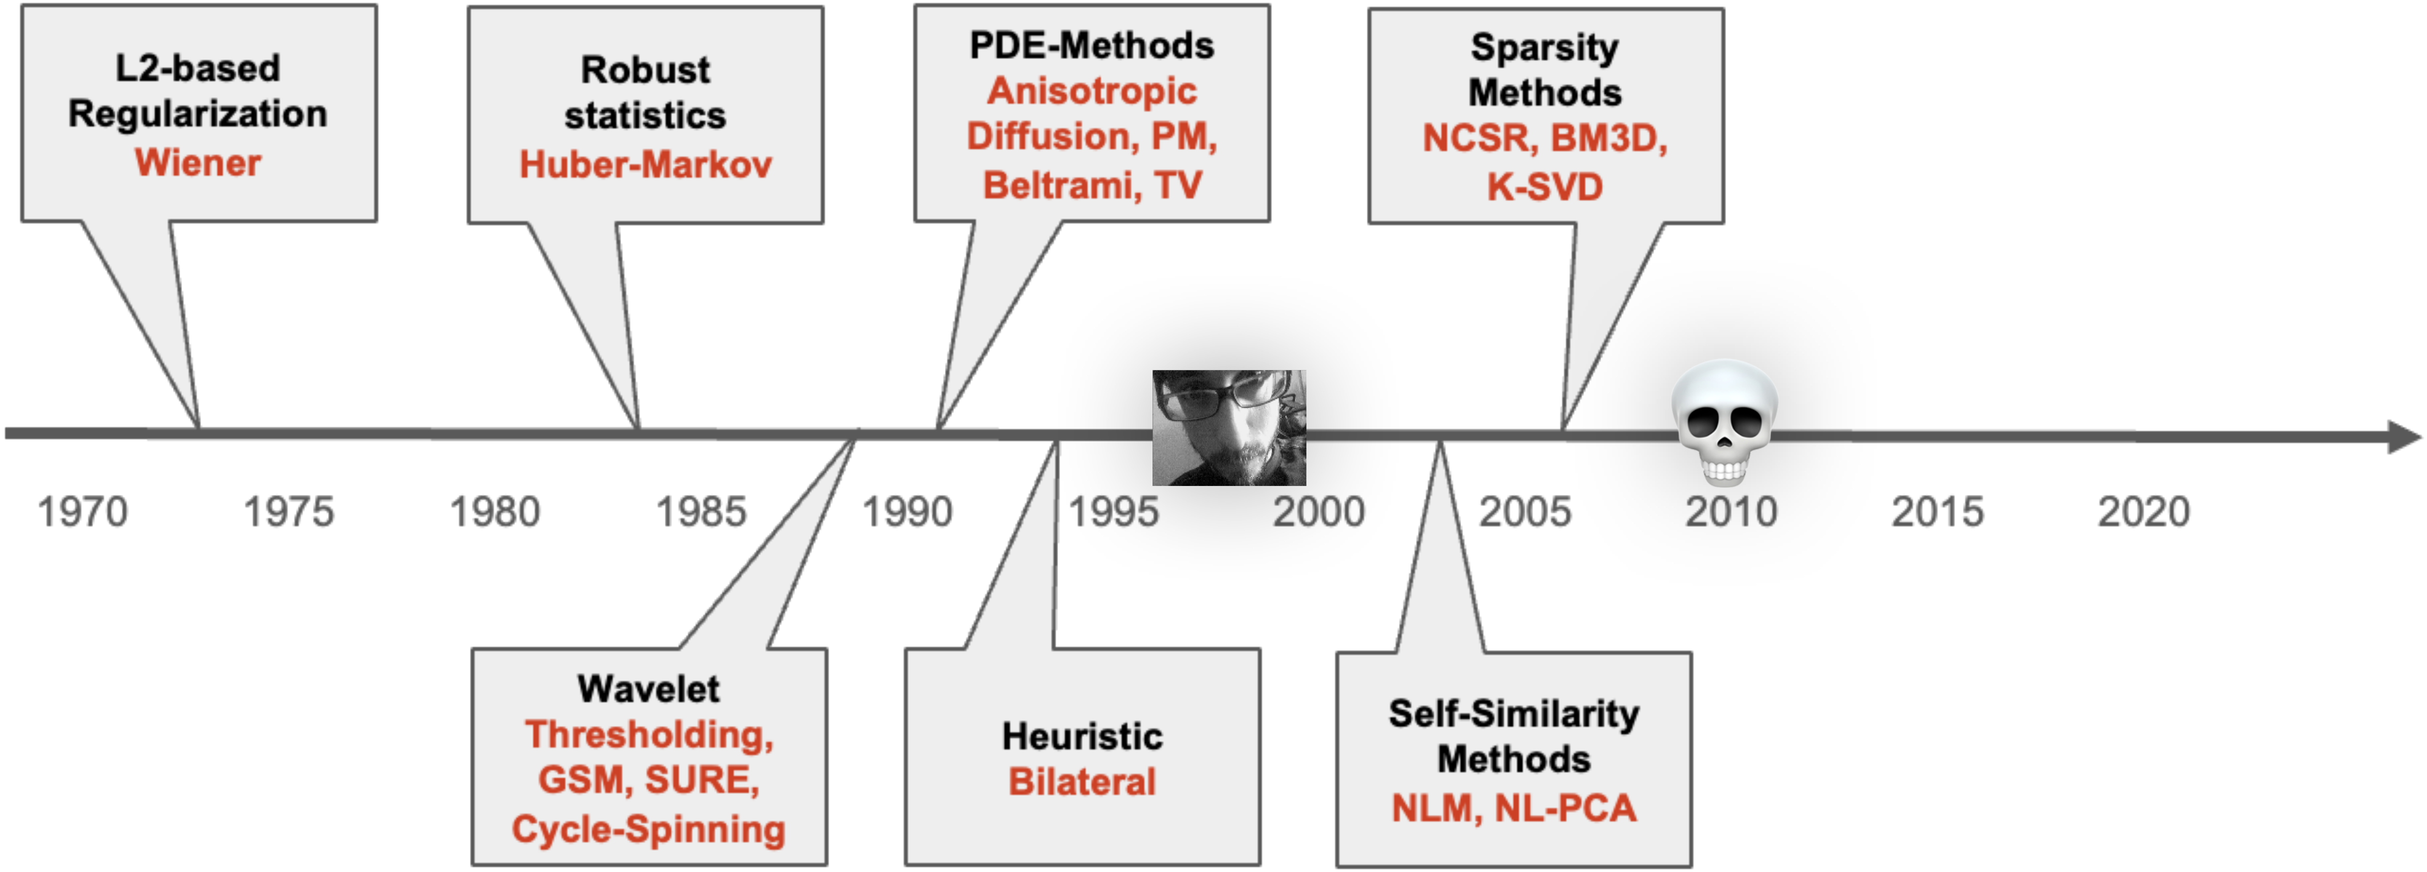
L2-based Regularization
Wiener
Robust statistics
Huber-Markov
PDE-Methods
Anisotropic Diffusion, PM, Beltrami, TV
Sparsity Methods
NCSR, BM3D, K-SVD
1970
1975
1980
1985
1990
1995
2000
2005
2010
2015
2020
Wavelet
Thresholding, GSM, SURE, Cycle-Spinning
Heuristic
Bilateral
Self-Similarity Methods
NLM, NL-PCA

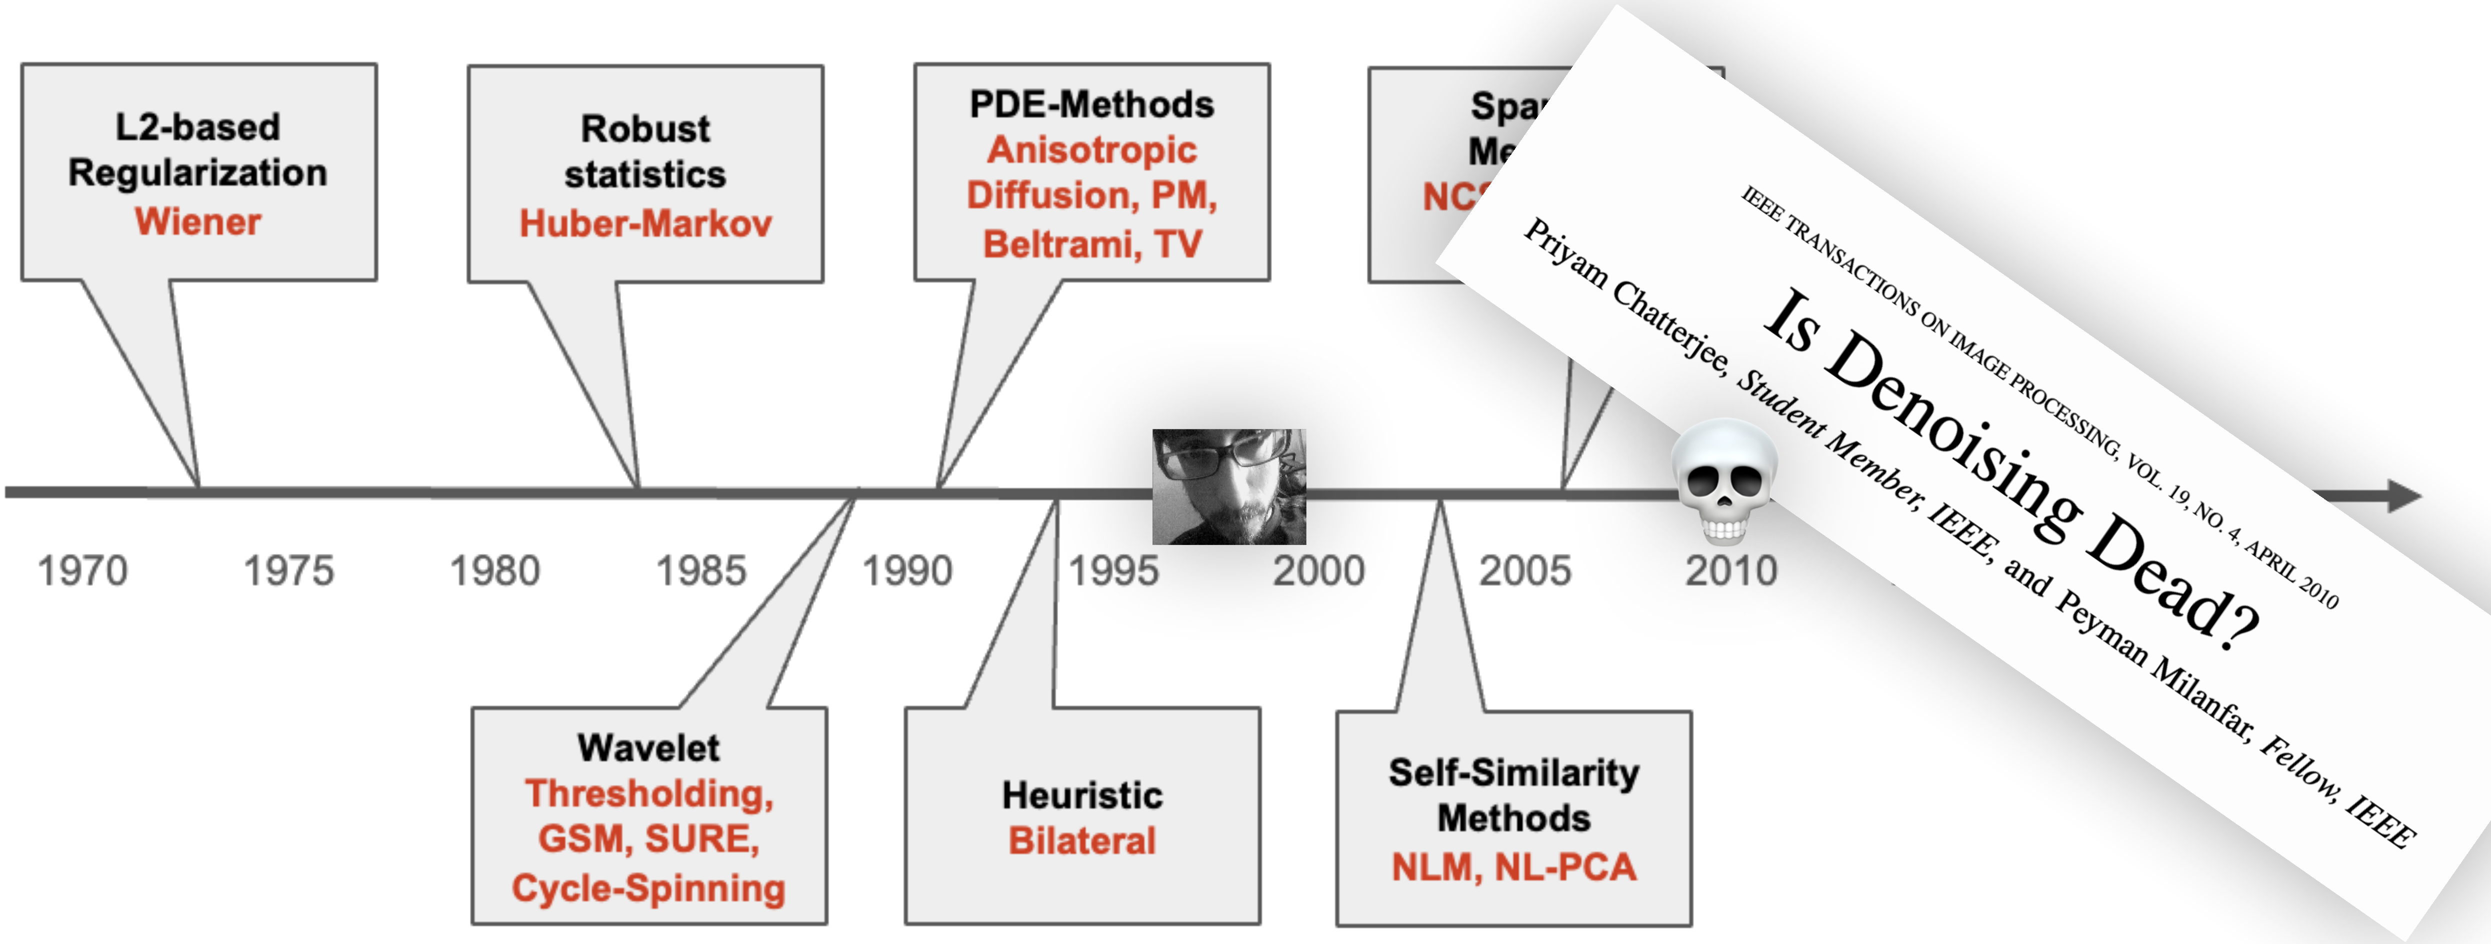
L2-based Regularization
Wiener
Robust statistics
Huber-Markov
PDE-Methods
Anisotropic Diffusion, PM, Beltrami, TV
Spa
Me
NC
Wavelet
Thresholding, GSM, SURE, Cycle-Spinning
Heuristic
Bilateral
Self-Similarity Methods
NLM, NL-PCA
1970
1975
1980
1985
1990
1995
2000
2005
2010
IEEE TRANSACTIONS ON IMAGE PROCESSING, VOL. 19, NO. 4, APRIL 2010
Is Denoising Dead?
Priyam Chatterjee, Student Member, IEEE, and Peyman Milanfar, Fellow, IEEE

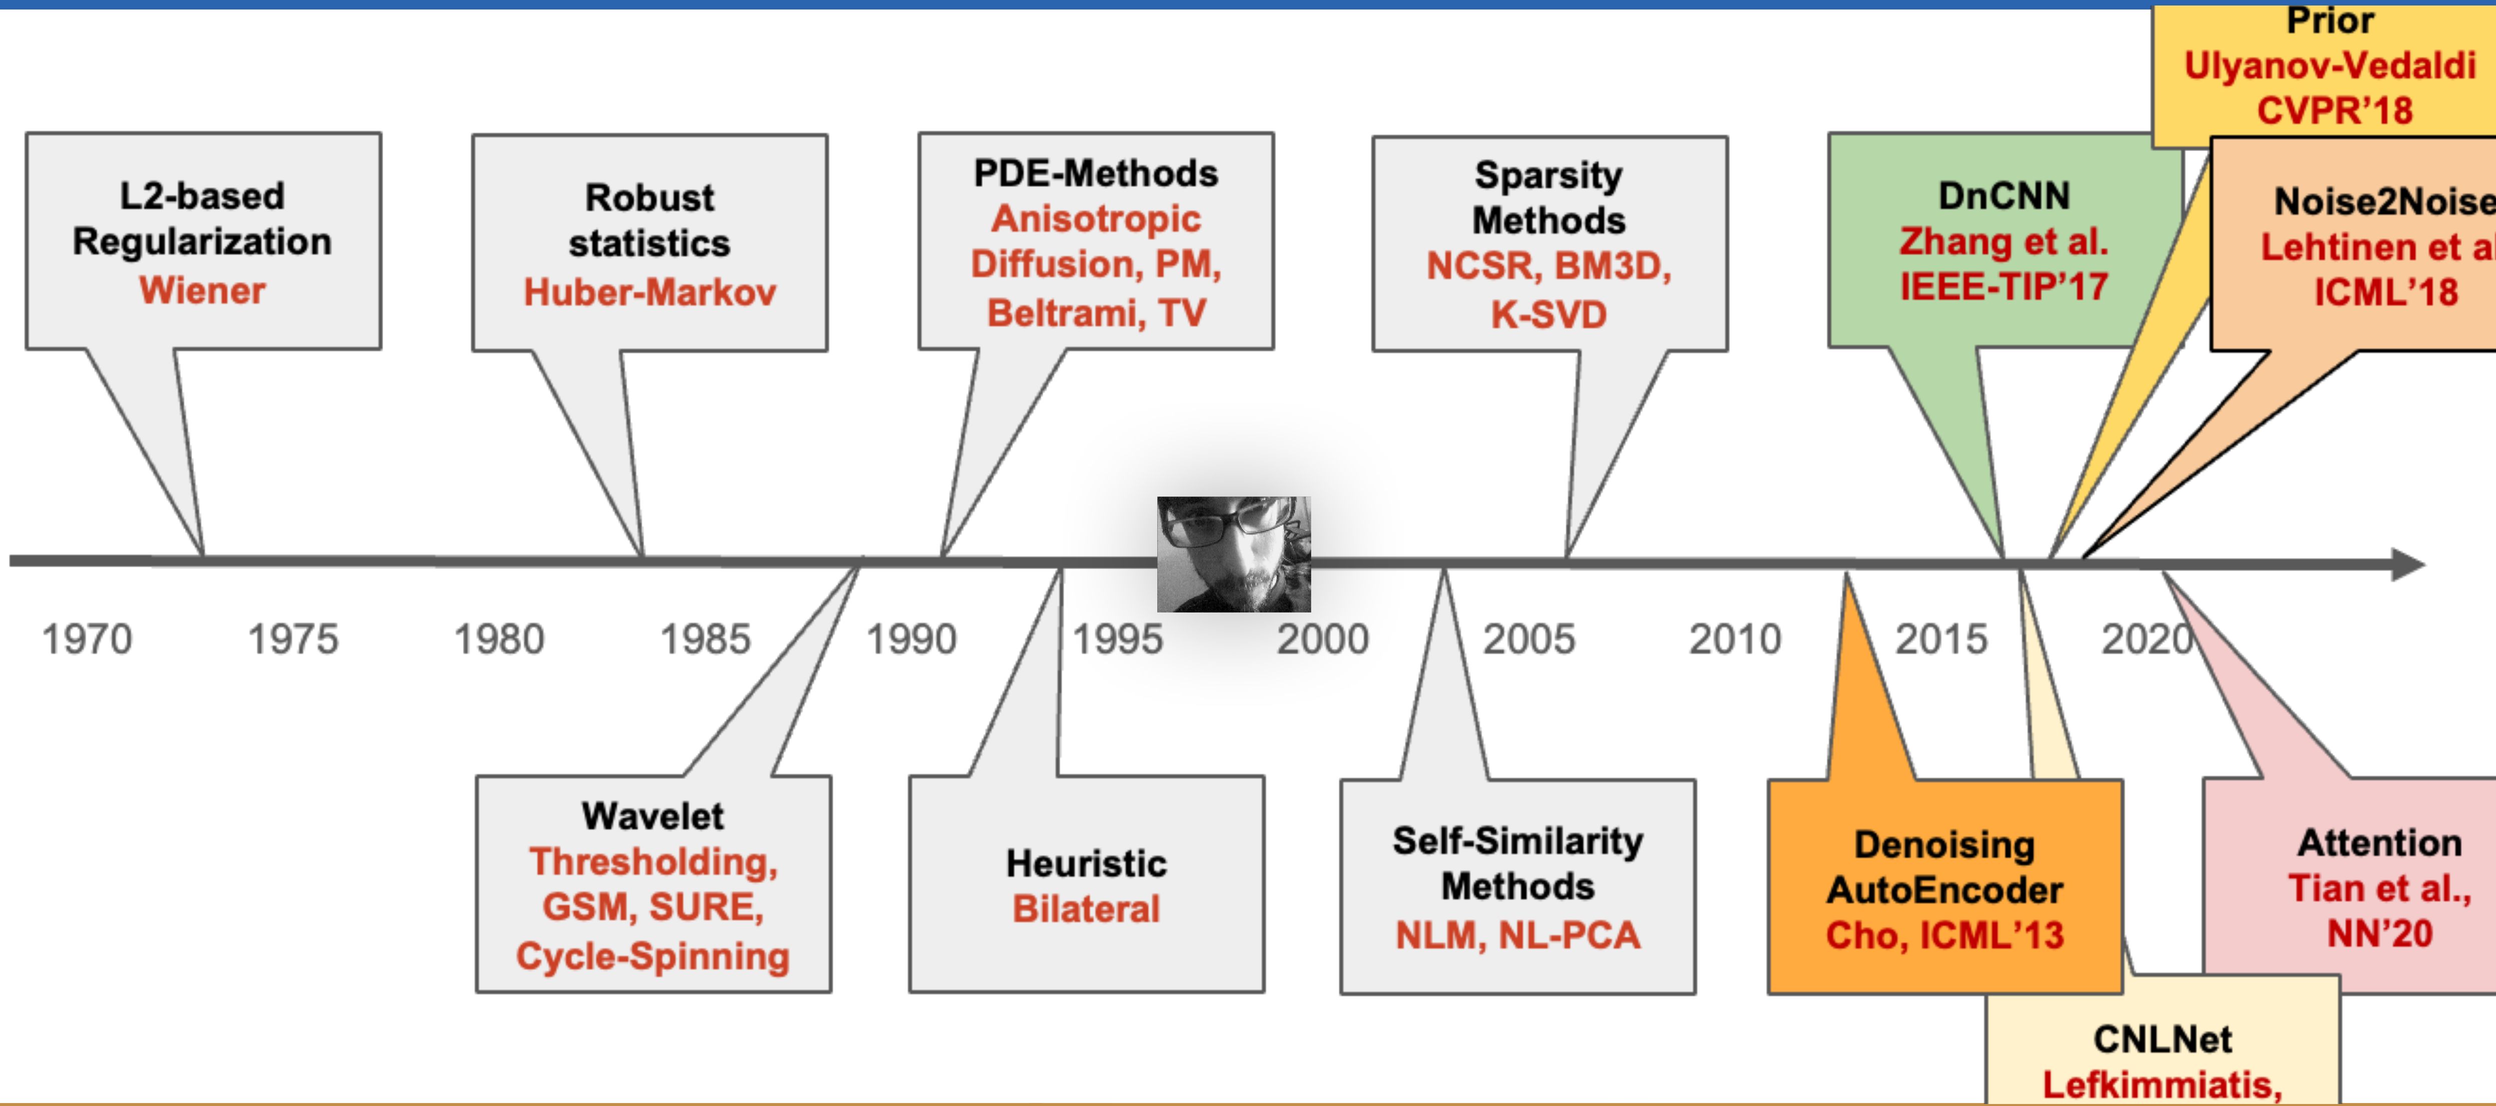
L2-based Regularization
Wiener
Robust statistics
Huber-Markov
PDE-Methods
Anisotropic Diffusion, PM, Beltrami, TV
Sparsity Methods
NCSR, BM3D, K-SVD
DnCNN
Zhang et al. IEEE-TIP'17
Prior
Ulyanov-Vedaldi CVPR'18
Noise2Noise
Lehtinen et al
ICML'18
1970
1975
1980
1985
1990
1995
2000
2005
2010
2015
2020
Wavelet
Thresholding, GSM, SURE, Cycle-Spinning
Heuristic
Bilateral
Self-Similarity Methods
NLM, NL-PCA
Denoising AutoEncoder
Cho, ICML'13
Attention
Tian et al., NN'20
CNLNet
Lefkimmiatis,

Classical Principles

- Low dimensionality
- Sparse Representation
- Local Smoothness
- Non-local self-similarity
- Low Rank

Years	Core concept	Formulae for prior
~1970	Energy regularization	$\\|\boldsymbol{x}\\|_2^2$
1975–1985	Spatial smoothness	$\\|L\boldsymbol{x}\\|_2^2$ or $\\|D_v\boldsymbol{x}\\|_2^2 + \\|D_h\boldsymbol{x}\\|_2^2$
1980–1985	Optimally learned transform	$\\|T\boldsymbol{x}\\|_2^2 = \boldsymbol{x}^\top R^{-1}\boldsymbol{x}$ (via PCA)
1980–1990	Weighted smoothness	$\\|L\boldsymbol{x}\\|_{2,W}^2$
1990–2000	Robust statistics	$\mathbf{1}^\top \mu\{L\boldsymbol{x}\}$ (e.g. Huber–Markov)
1992–2005	Total Variation (TV)	$\int_{v\in\Omega} \lvert\nabla\boldsymbol{x}(v)\rvert\, dv = \mathbf{1}^\top\sqrt{\lvert D_v\boldsymbol{x}\rvert^2 + \lvert D_h\boldsymbol{x}\rvert^2}$
1987–2005	Other PDE-based options	$\int_{v\in\Omega} g\big(\nabla\boldsymbol{x}(v), \nabla^2\boldsymbol{x}(v)\big)\, dv$
2005–2009	Field-of-experts	$\sum_k \lambda_k \mathbf{1}^\top \mu_k\{L_k\boldsymbol{x}\}$
1993–2005	Wavelet sparsity	$\\|W\boldsymbol{x}\\|_1$
2000–2010	Self-similarity	$\sum_k \sum_{j\in\Omega(k)} d\{R_k\boldsymbol{x}, R_j\boldsymbol{x}\}$
2002–2012	Sparsity methods	$\\|\boldsymbol{\alpha}\\|_0$ s.t. $\boldsymbol{x} = D\boldsymbol{\alpha}$
2010–2017	Low-rank assumption	$\sum_k \\|X_{\Omega(k)}\\|_\star$

Classical Principles

ss

Non-lo -similarity

nk

"just"…

Supervised Training

Finanziato dall'Unione europea
NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA
INAF
ISTITUTO NAZIONALE DI ASTROFISICA

...still, plenty to be happy about!

...still, plenty to be happy about!

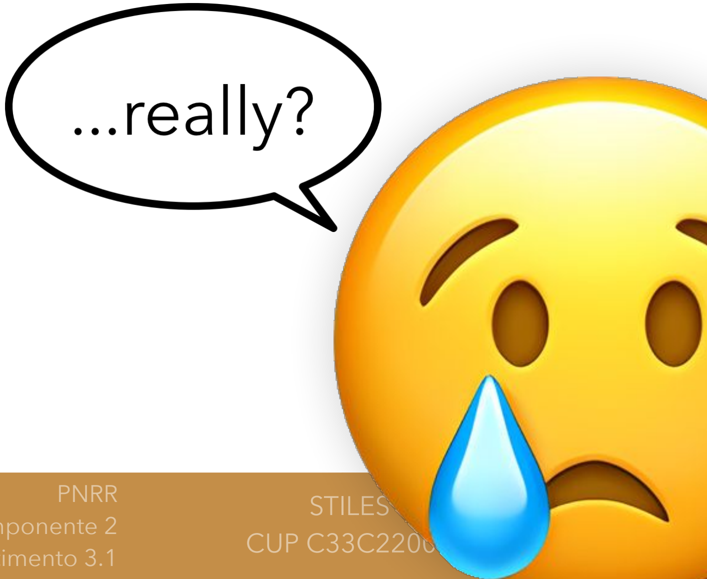

From **denoising** to **sampling** and back...

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) steps:
  1. Start with (or *get to*) a random noise image.
  2. "Climb" to a more probable image by **Langevin Dynamics.**
  3. …and look back at the Bayesian literature of the '60s…surprise!

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) steps:
  1. Start with (or *get to*) a **random noise** image.
  2. "Climb" to a more probable image by **Langevin Dynamics.**
  3. …and look back at the Bayesian literature of the '60s…surprise!

Diffusion Models Beat GANs on Image Synthesis

Prafulla Dhariwal*
OpenAI
prafulla@openai.com

Alex Nichol*
OpenAI
alex@openai.com

Abstract

We show that diffusion models can achieve image sample quality superior to the current state-of-the-art generative models. We achieve this on unconditional image synthesis by finding a better architecture through a series of ablations. For conditional image synthesis, we further improve sample quality with classifier guidance: a simple, compute-efficient method for trading off diversity for fidelity using gradients from a classifier. We achieve an FID of 2.97 on ImageNet 128×128, 4.59 on ImageNet 256×256, and 7.72 on ImageNet 512×512, and we match BigGAN-deep even with as few as 25 forward passes per sample, all while maintaining better coverage of the distribution. Finally, we find that classifier guidance combines well with upsampling diffusion models, further improving FID to 3.94 on ImageNet 256×256 and 3.85 on ImageNet 512×512. We release our code at `https://github.com/openai/guided-diffusion`.

1 Introduction

Midjourney

Stable Diffusion

DALL-E

Adobe Firefly

*Remark: in these examples the generation is conditional to a prompt.

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) steps:
  1. Start with (or *get to*) a **random noise** image.
  2. "Climb" to a more probable image by **Langevin Dynamics.**
  3. …and look back at the Bayesian literature of the '60s…surprise!

- Pretend

  We can access the **true** image distribution $p(\boldsymbol{x})$

- Goal

  Draw a sample from $p(\boldsymbol{x})$

PNRR
Missione 4 • Componente 2
Investimento 3.1
STILES – IR0000034
CUP C33C22000640006

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) steps:
  1. Start with (or *get to*) a **random noise** image.
  2. "Climb" to a more probable image by **Langevin Dynamics.**
  3. …and look back at the Bayesian literature of the '60s…surprise!

- Pretend

  We can access the **true** image distribution $p(\boldsymbol{x})$

- Goal

  Draw a sample from $p(\boldsymbol{x})$

Step0: start with a purely random noise image $\widehat{\boldsymbol{x}}_0$

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) steps:
  1. Start with (or *get to*) a **random noise** image.
  2. "Climb" to a more probable image by **Langevin Dynamics.**
  3. …and look back at the Bayesian literature of the '60s…surprise!

- Pretend

We can access the **true** image distribution $p(\boldsymbol{x})$

- Goal

Draw a sample from $p(\boldsymbol{x})$

Step0: start with a purely random noise image $\widehat{\boldsymbol{x}}_0$

Step1: climb to a more probable image by iterating

$$\widehat{\boldsymbol{x}}_{t+1} = \widehat{\boldsymbol{x}}_t + \alpha \cdot \nabla \log p(\widehat{\boldsymbol{x}}_t) \qquad \text{[stuck in local max]}$$

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
INAF ISTITUTO NAZIONALE DI ASTROFISICA

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) steps:
  1. Start with (or *get to*) a **random noise** image.
  2. "Climb" to a more probable image by **Langevin Dynamics.**
  3. …and look back at the Bayesian literature of the '60s…surprise!

- Pretend

We can access the **true** image distribution $p(\boldsymbol{x})$

- Goal

Draw a sample from $p(\boldsymbol{x})$

Step0: start with a purely random noise image $\widehat{\boldsymbol{x}}_0$

Step1: climb to a more probable image by iterating

$$\widehat{\boldsymbol{x}}_{t+1} = \widehat{\boldsymbol{x}}_t + \alpha \cdot \nabla \log p(\widehat{\boldsymbol{x}}_t) + \underbrace{\beta \cdot \boldsymbol{\eta}_t}_{\text{random perturbation}} \qquad \text{[Langevin]}$$

- At the turn of this decade, the idea of using a good **denoiser** to **synthesize** (*hallucinate*) new images emerged.

- In essence:
  Use the denoiser as a **projector** onto the *image manifold*.

- Main (backward) step
  1. Start with (…
     ran…

  …ack at the
  …sian literature of
  the '60s…surprise!

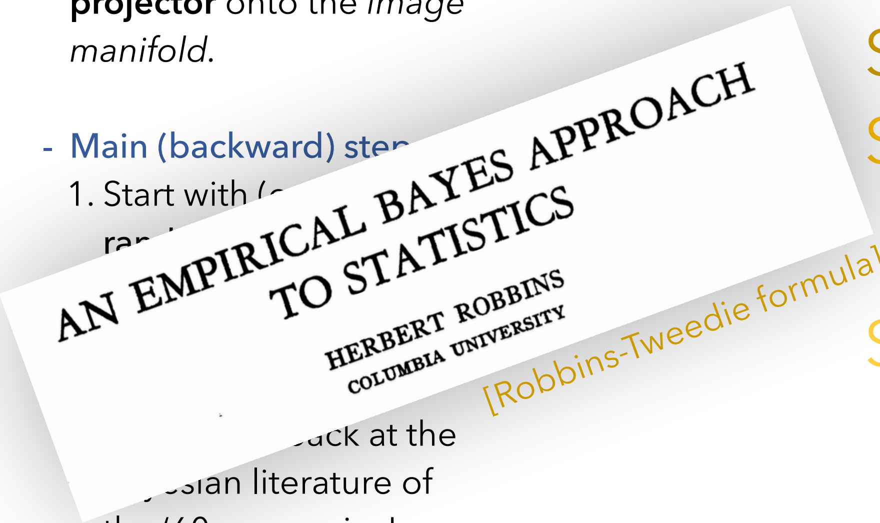
AN EMPIRICAL BAYES APPROACH TO STATISTICS
HERBERT ROBBINS
COLUMBIA UNIVERSITY

[Robbins-Tweedie formula]

- Pretend

  We can access the **true** image distribution $p(\boldsymbol{x})$

- Goal

  Draw a sample from $p(\boldsymbol{x})$

Step0: start with a purely random noise image $\widehat{\boldsymbol{x}}_0$

Step1: climb to a more probable image by iterating

$$\widehat{\boldsymbol{x}}_{t+1} = \widehat{\boldsymbol{x}}_t + \alpha \cdot \underbrace{\nabla \log p(\widehat{\boldsymbol{x}}_t)}_{\text{score function | unknown}} + \beta \cdot \boldsymbol{\eta}_t \qquad \text{[Langevin]}$$

Step2: look at the Empirical Bayes literature to see that

$$\nabla \log p(\widehat{\boldsymbol{x}}_t) \propto \left[\widehat{\boldsymbol{x}}_t - \mathrm{D}(\widehat{\boldsymbol{x}}_t, \sigma)\right] \qquad \text{[small σ]}$$

implicit link between MMSE denoisers and priors

- "Technical" caveat
  In practice, instead of the plain Langevin with a fixed (and small) σ we use the Annealed Langevin that considers a sequence of blurred priors associated with a decreasing sequence of noise levels:

$$\sigma_0 > \sigma_1 > \cdots > \sigma_N$$

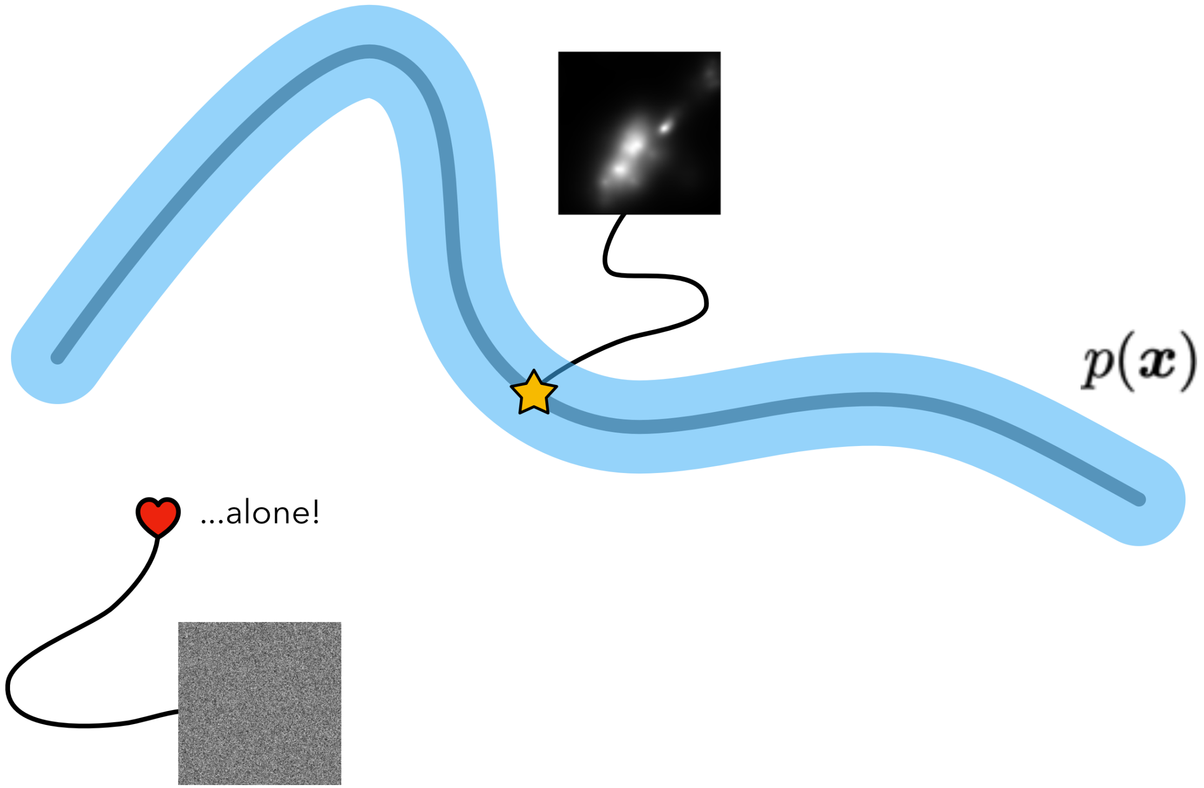

- "Technical" caveat
  In practice, instead of the plain Langevin with a fixed (and small) σ we use the Annealed Langevin that considers a sequence of blurred priors associated with a decreasing sequence of noise levels:

$$\sigma_0 > \sigma_1 > \cdots > \sigma_N$$

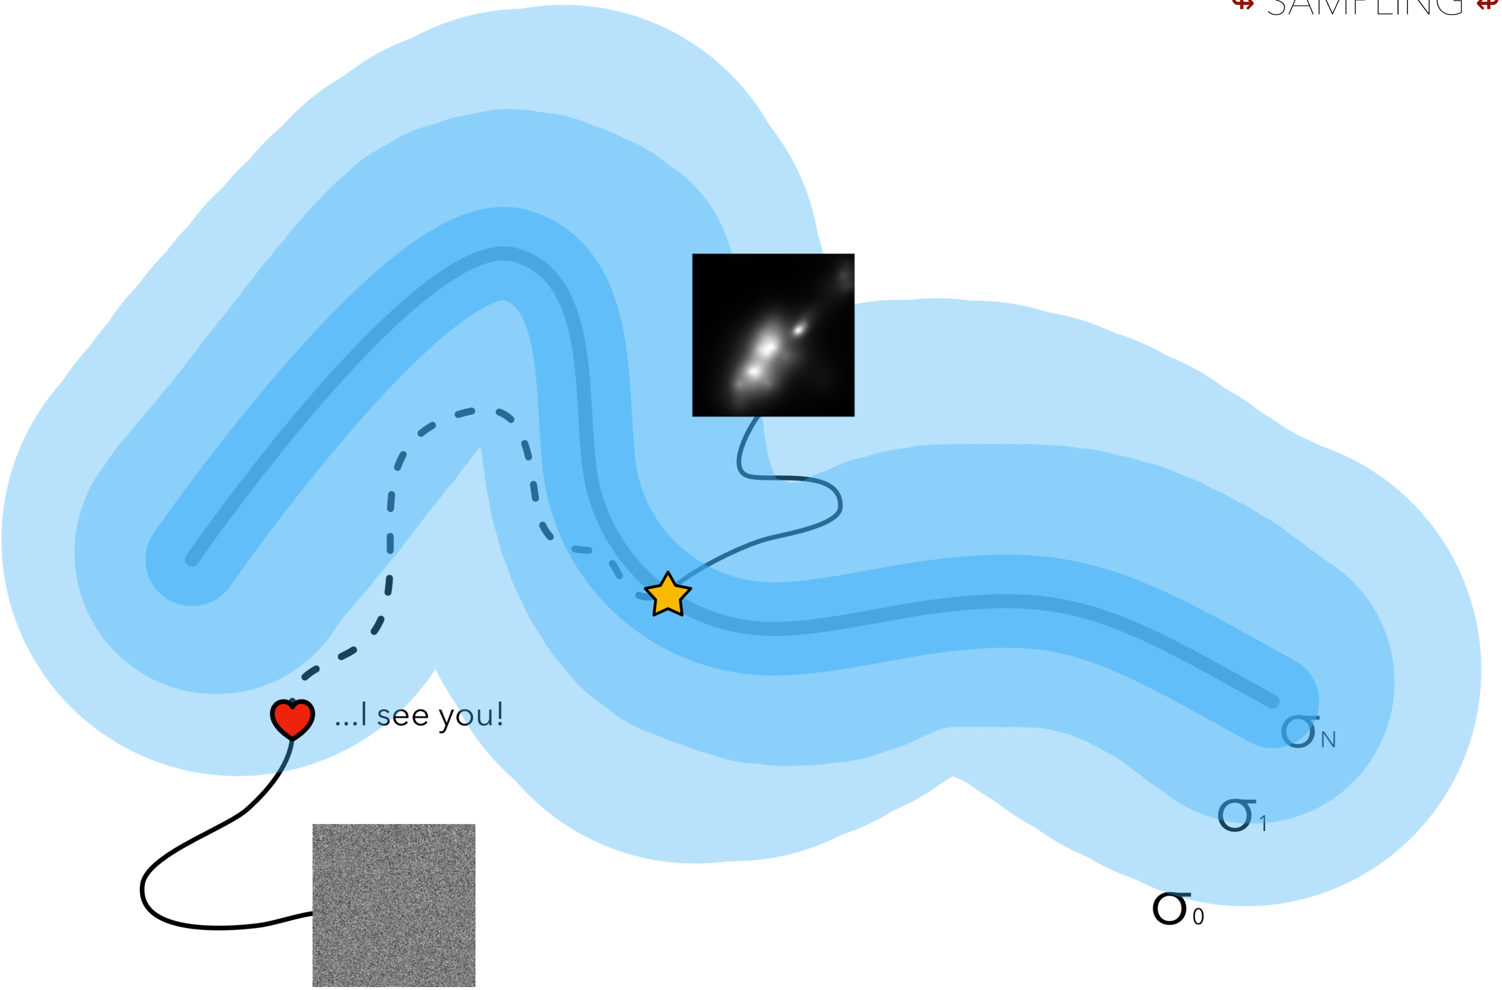

Finanziato dall'Unione europea NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani PIANO NAZIONALE DI RIPRESA E RESILIENZA
INAF ISTITUTO NAZIONALE DI ASTROFISICA

- **Denoising by Sampling:** The Annealed Langevin algorithm can be twisted to design a **Stochastic Image Denoiser** that recover a sequence of gradually less noisy images by targeting the posterior distribution

$$p(\boldsymbol{x} \mid \boldsymbol{y})$$

Solution is **not** unique.

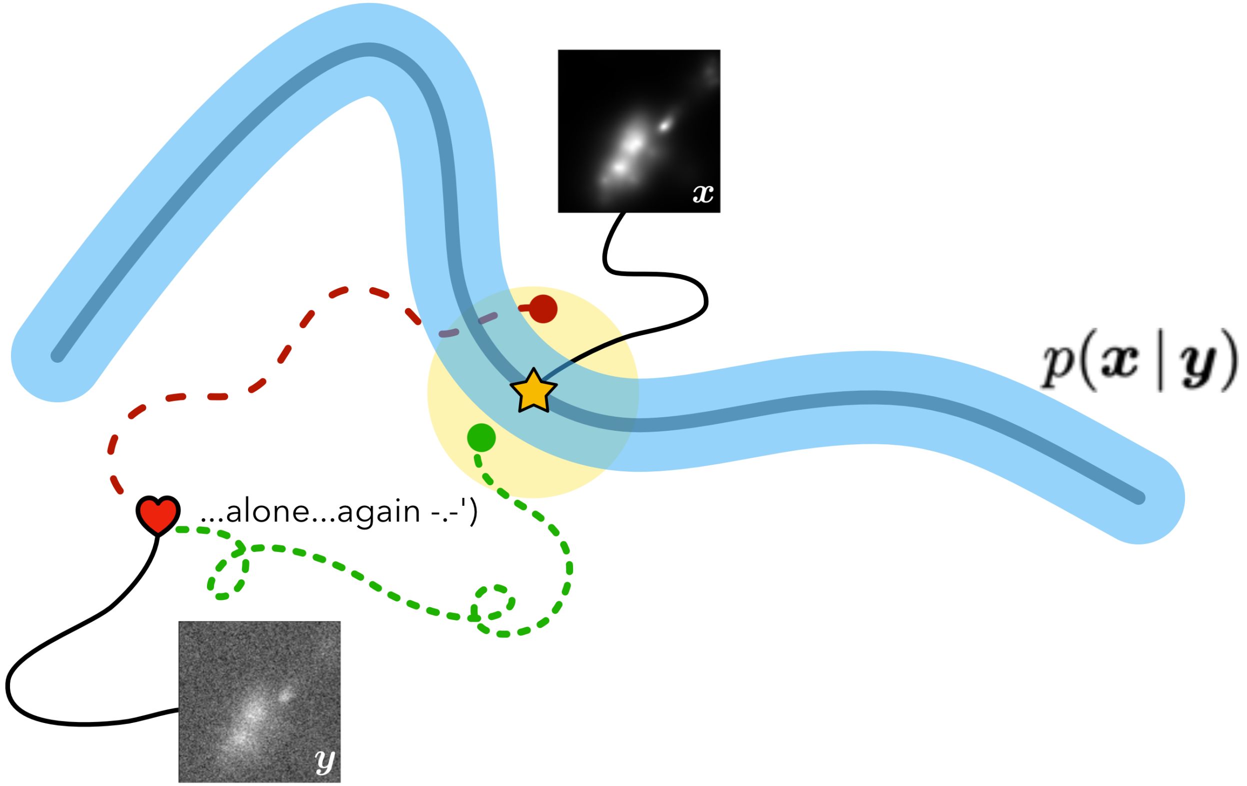

- ...and beyond!

  Pretty much the same Langevin-based paradigm can be adapted to deal with a variety of **inverse problems** in order to recover a signal from **corrupted measurements**:

$$\boldsymbol{y} = \mathrm{H}\boldsymbol{x} + \boldsymbol{\epsilon}$$

  - De-blurring
  - In-Painting
  - De-Mosaicing
  - Super-Resolution

- **...and beyond!**
  Pretty much the same Langevin-based paradigm can be adapted to deal with a variety of **inverse problems** in order to recover a signal from **corrupted measurements**:

$$\boldsymbol{y} = \mathrm{H}\boldsymbol{x} + \boldsymbol{\epsilon}$$

  - De-blurring
  - In-Painting
  - De-Mosaicing
  - Super-Resolution

- **Power is nothing without control**
  - Often we'd to ensure the generated data not only looks realistic but also adheres to desired **connectivity** constraints.
  - With **topological steering** we guide the denoising process to produce outputs with specific structural properties (e.g. connected components, holes, or specific branching structures)
  - Typical tools **persistent homology** or **topological loss** functions, although they are quite **computationally intensive**, **but...**

Finanziato dall'Unione europea
NextGenerationEU
Ministero dell'Università e della Ricerca
Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA
INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Ric Strikes back!

- **Persistent homology** is expensive on large images.
- Complexes grow quickly in **time** and **memory**.
- Standard pipelines are impractical for real surveys.
- PixHomology makes PH feasible on real images.

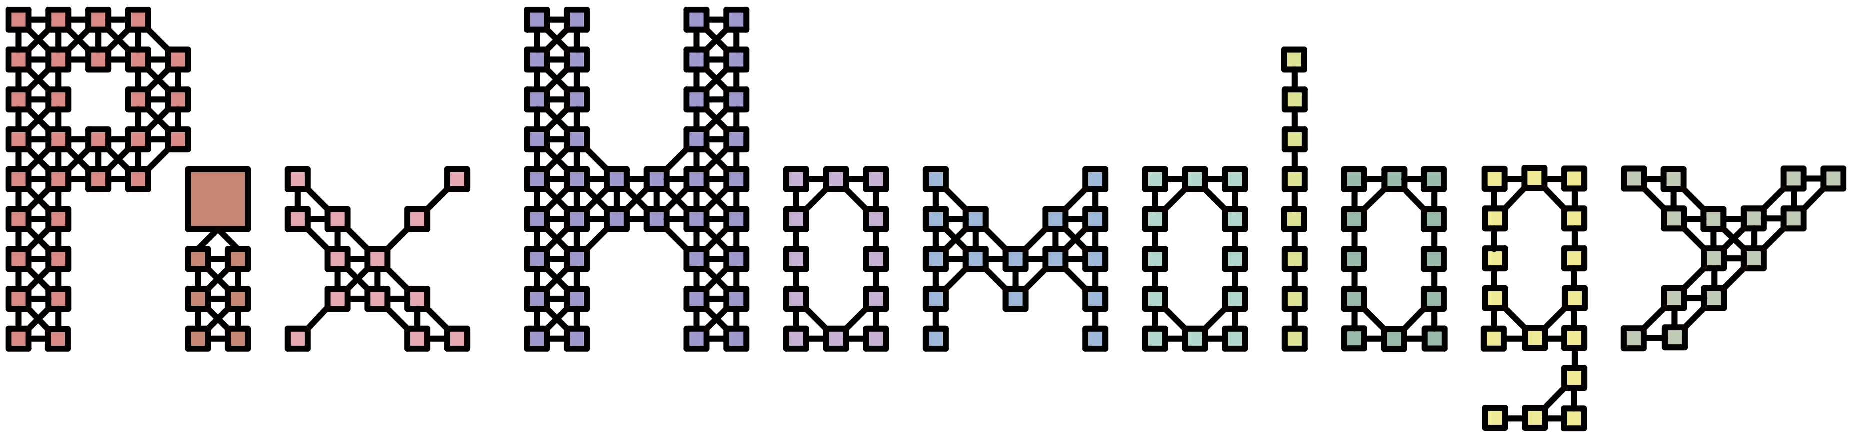

https://github.com/riccardoc95/PixHomology

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

https://teralab.ai/

Welcome to TeraLab

HPC for Statistics and AI

Finanziato dall'Unione europea
NextGenerationEU

Ministero dell'Università e della Ricerca

Italiadomani
PIANO NAZIONALE DI RIPRESA E RESILIENZA

INAF
ISTITUTO NAZIONALE DI ASTROFISICA

Thanks!

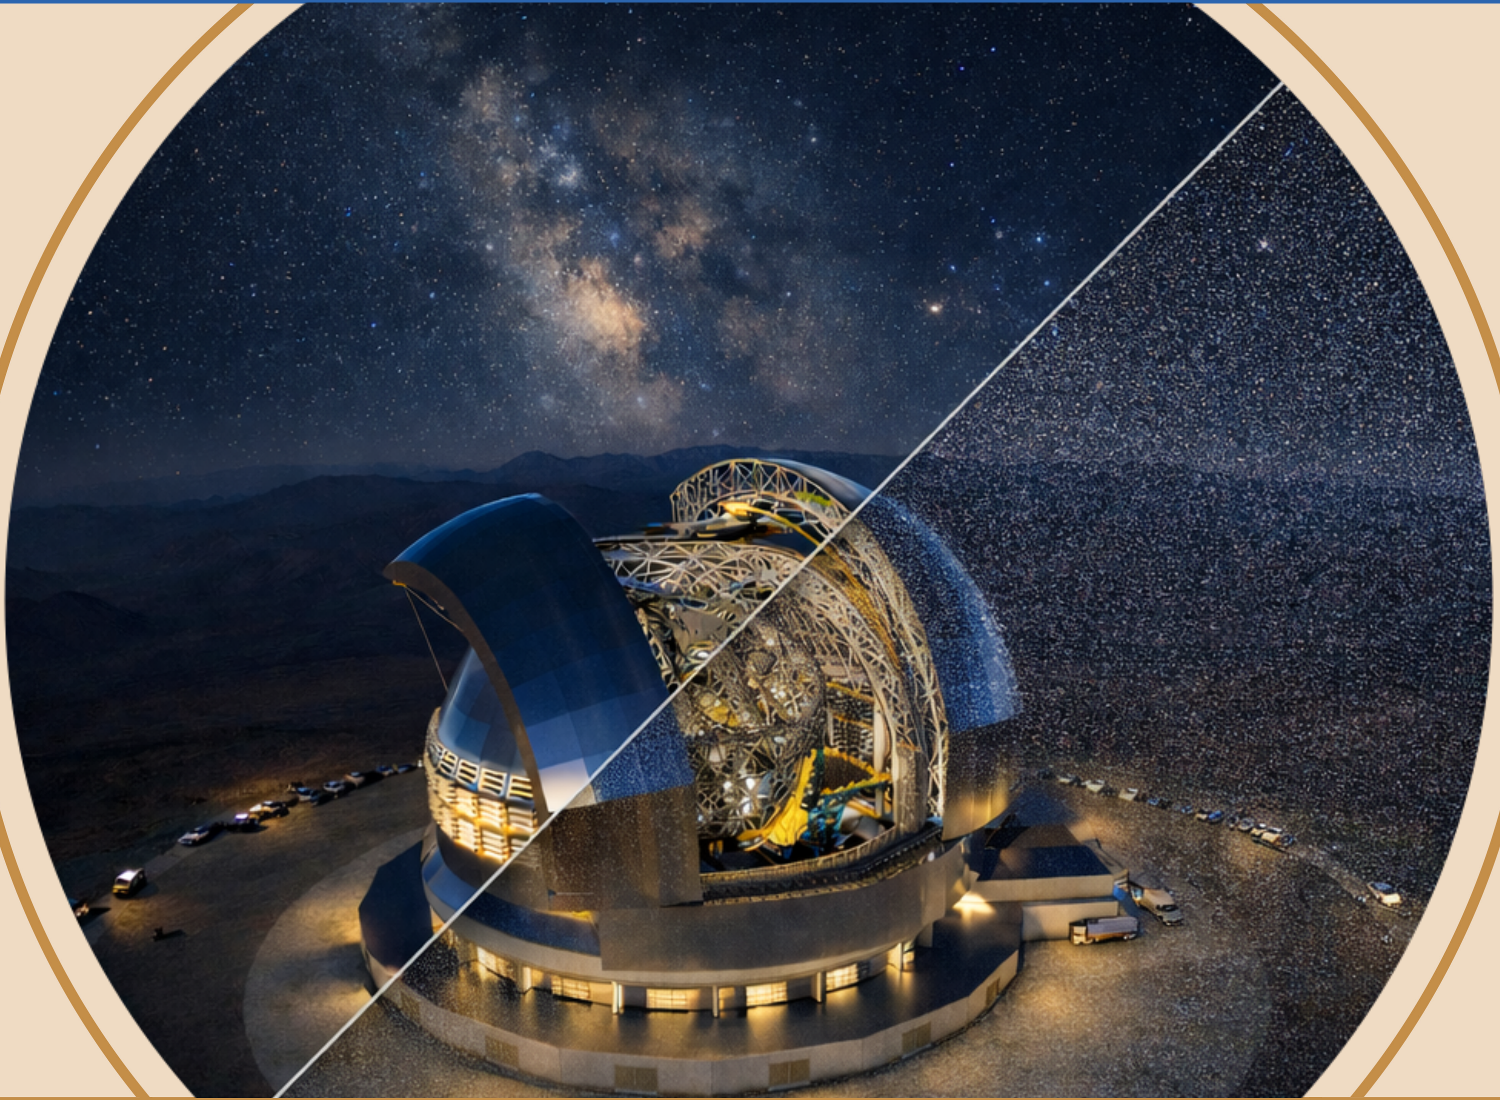

*See (some of) you in a future project…maybe ^.^)