New Results in Image Processing based on Sparse and Redundant Representations

Michael Elad
The Computer Science Department
The Technion – Israel Institute of technology
Haifa 32000, Israel

Medical Imaging Workshop
RAMBAM Hospital
25.2.07

Model?

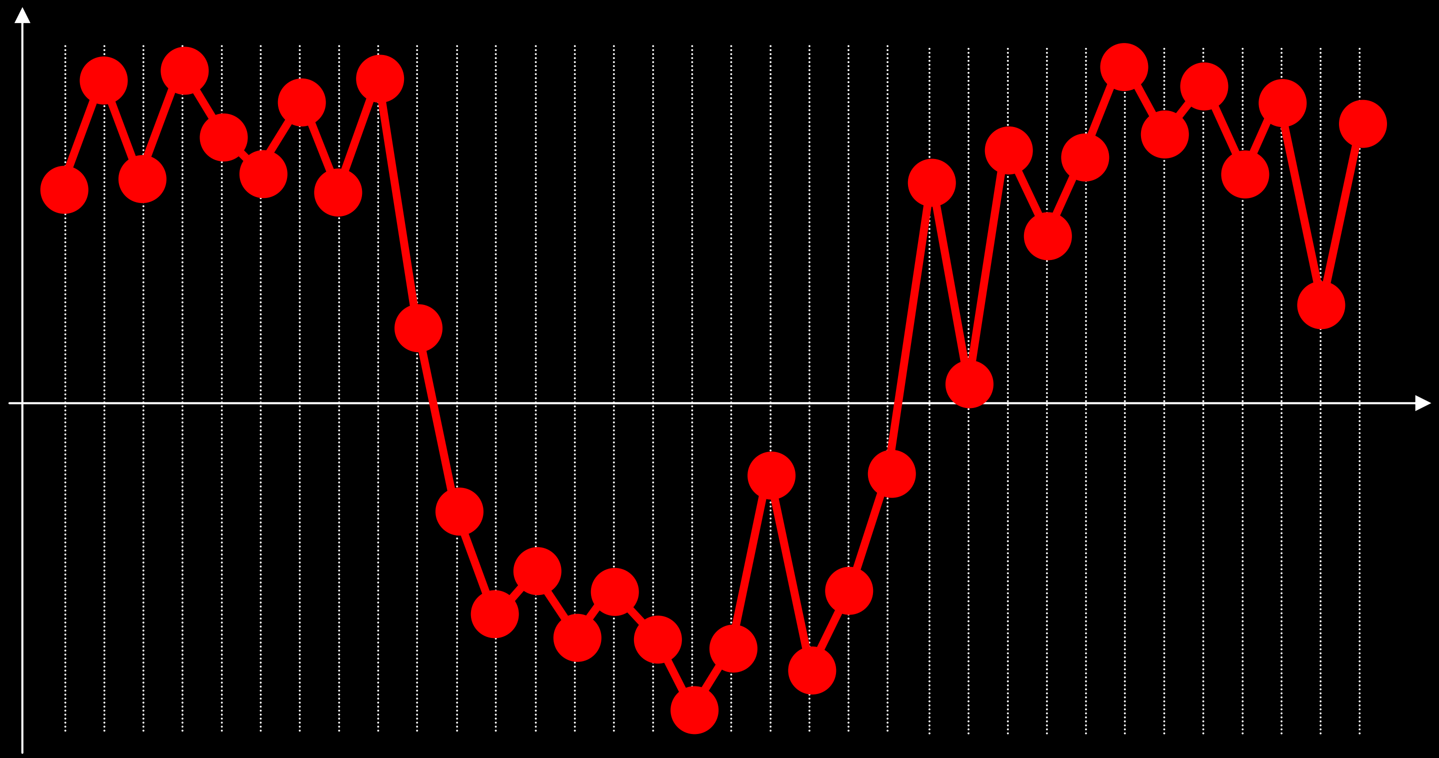

Model?

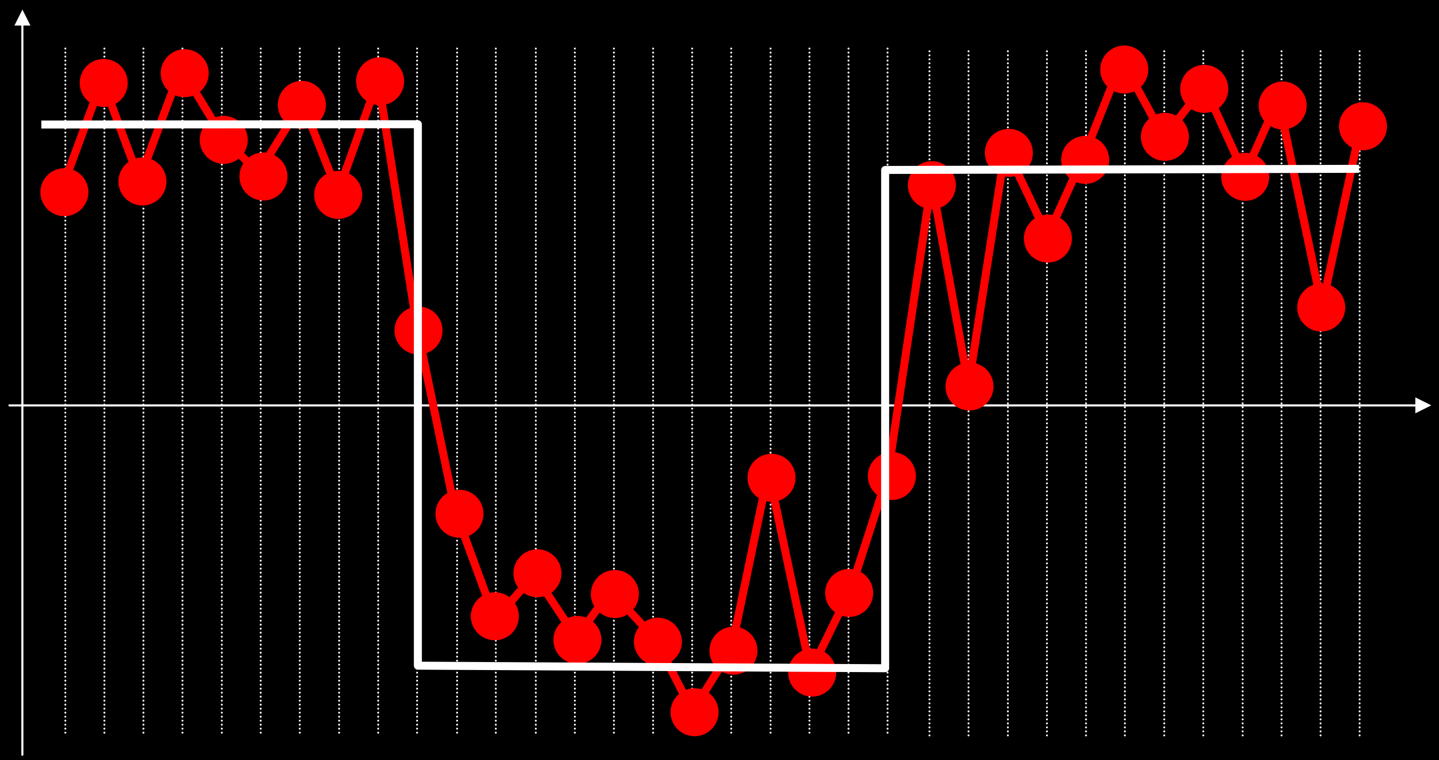

Effective removal of noise (and many other applications) relies on an proper *modeling* of the signal

Michael Elad
The Computer-Science Department
The Technion

Which Model to Choose?

- There are many different ways to model signals and images with varying degrees of success.
- The following is a partial list of such models for images:
- Good models should be simple while matching the signals:

Simplicity ⟷ *Reliability*

Principal-Component-Analysis

Anisotropic diffusion

Markov Random Field

Wienner Filtering

DCT and JPEG

Wavelet & JPEG-2000

Piece-Wise-Smooth

C2-smoothness

Besov-Spaces

Total-Variation

Beltrami-Flow

Sparseland

Multi-Scale Analysis

Wavelet Theory

Approximation Theory

Linear Algebra

Signal Transforms

Optimization Theory

Sparseland

A new Model for Signals / Images

- Blind Source Separation
- Compression
- Inpainting
- Denoising
- Demosaicing
- Super-Resolution
- …

Michael Elad
The Computer-Science Department
The Technion

The *Sparseland* Model for Images

- Task: model image patches of size 10×10 pixels.
- We assume that a *dictionary* of such image patches is given, containing 256 *atom* images.
- The Model: every image patch can be described as a linear combination of *few* atoms.
- This model describes every image patch as a *sparse* combination over a *redundant* dictionary.

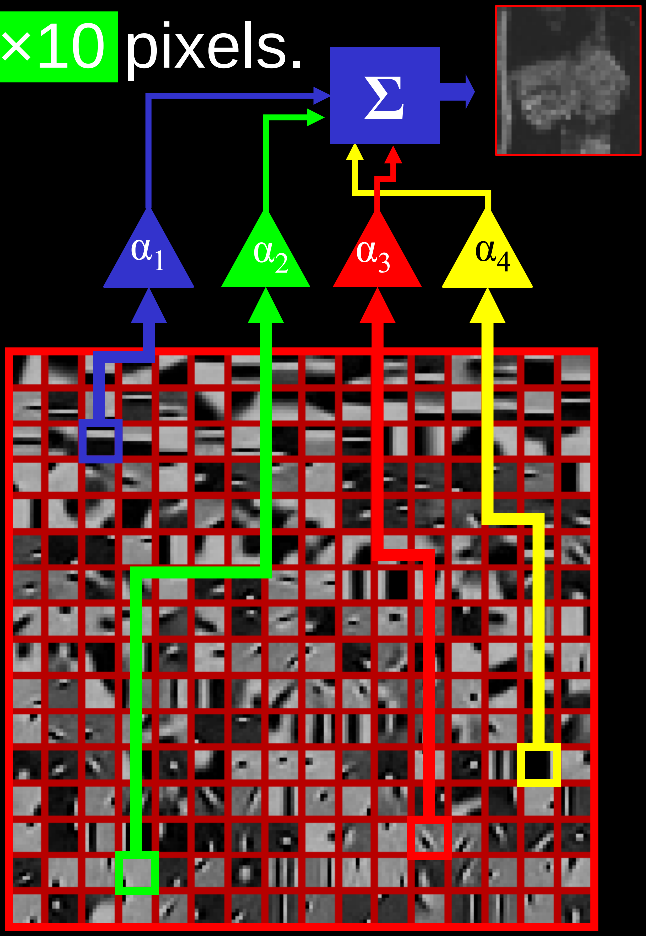

Difficulties With *Sparseland*

- Problem 1: Given an image patch, how can we find its *atom decomposition* ?
- Problem 2: Given a family of signals, how do we find the dictionary to represent it well?
- Problem 3: Is this model flexible enough to describe various sources?

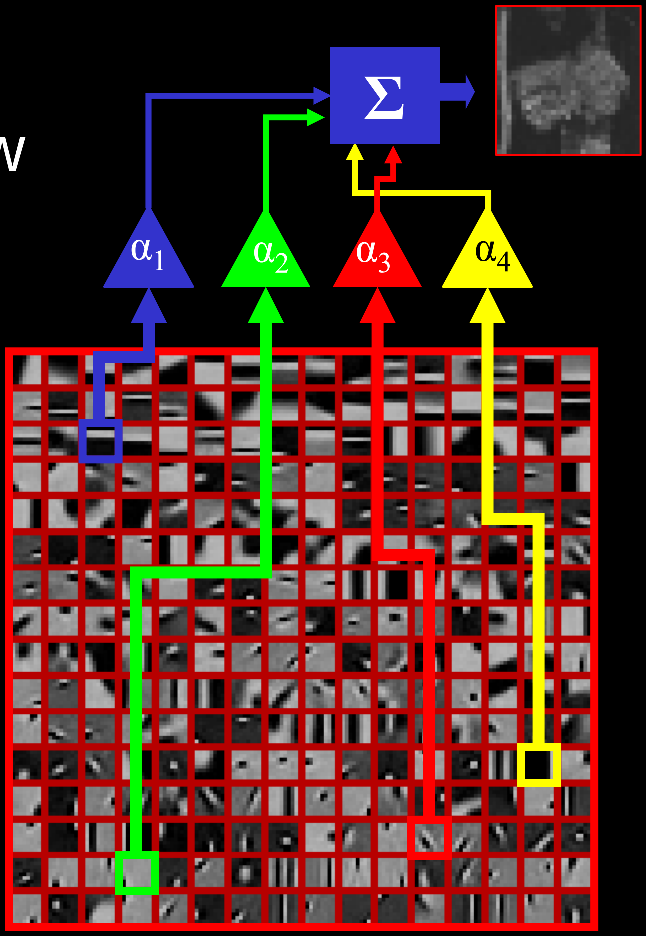

Difficulties With *Sparseland*

- Problem 1: Given an image patch, how can we find its *atom decomposition* ?

- Problem 2: Given a family of signals, how do we find the dictionary to represent it well?

- Problem 3: Is this model flexible enough to describe various sources?

ALL ANSWERED POSITIVELY AND CONSTRUCTIVELY

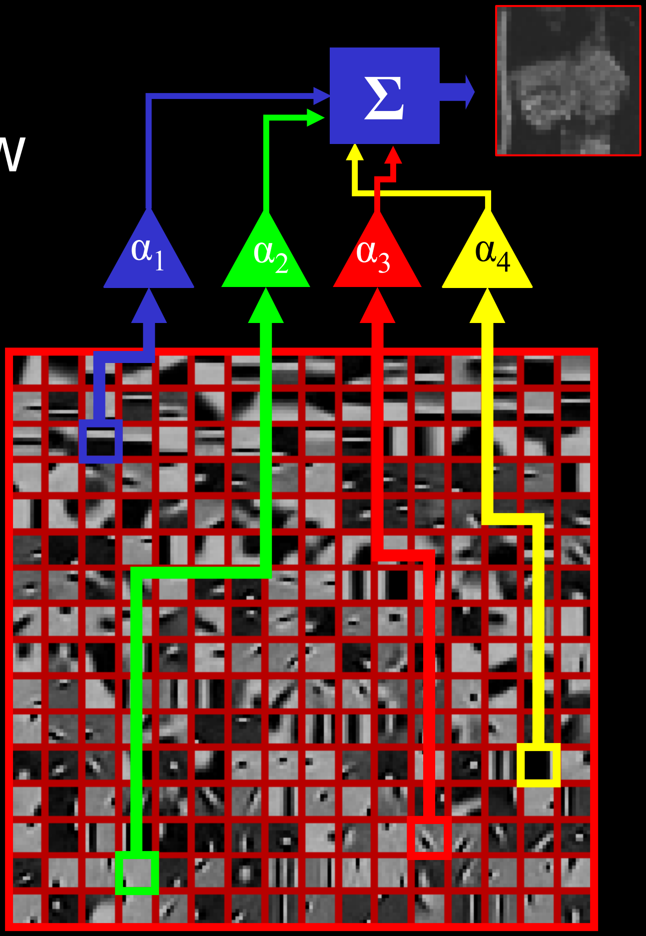

Michael Elad
The Computer-Science Department
The Technion

Image Denoising (Gray) [Elad & Aharon (`06)]

Source

Result 30.829dB

Noisy image

$$\sigma = 20$$

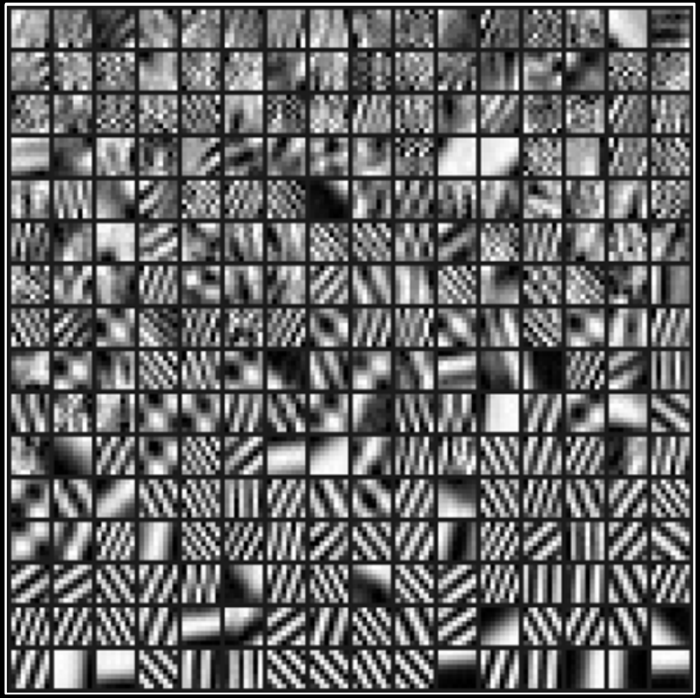

The obtained dictionary after 10 iterations

Image Denoising (Gray) [Elad & Aharon (`06)]

Source

The results of this algorithm compete favorably with the state-of-the-art: E.g.,

- We get ~1dB better results compared to GSM+steerable wavelets [Portilla, Strela, Wainwright, & Simoncelli ('03)].
- Competitive works are [Hel-Or & Shaked (`06)] and [Rusanovskyy, Dabov, & Egiazarian ('07)]. Both also lean on *Sparseland*.

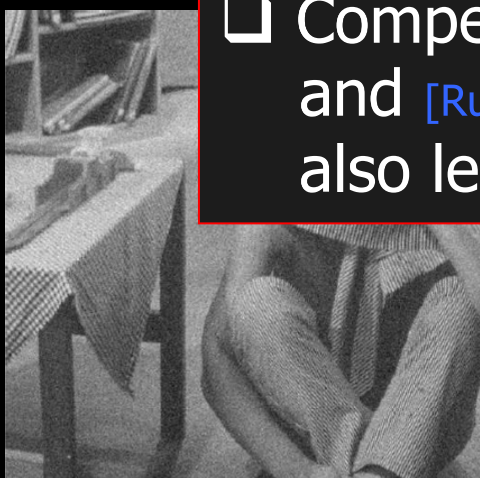

Noisy image

$\sigma = 20$

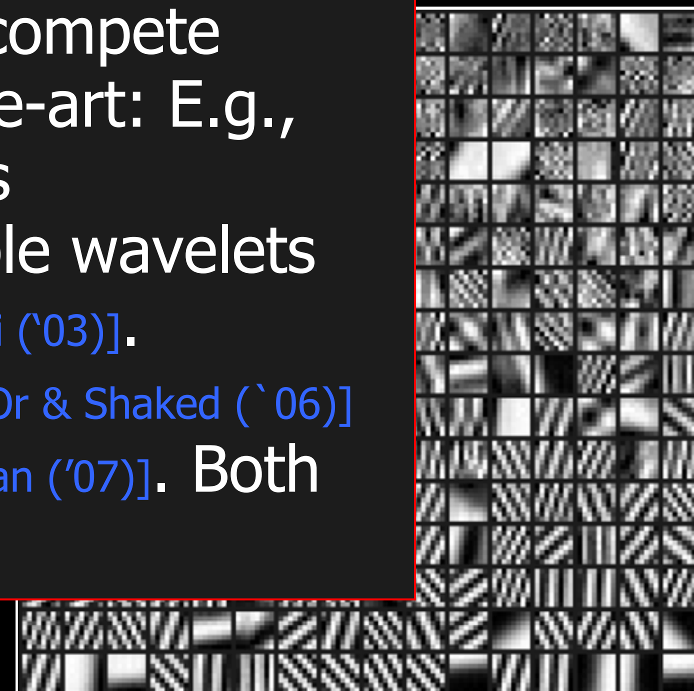

The obtained dictionary after 10 iterations

Denoising (Color) [Mairal, Elad & Sapiro, ('06)]

Our experiments lead to state-of-the-art denoising results, giving ~1dB better results compared to [Mcauley et. al. ('06)] which implements a learned MRF model (Field-of-Experts)

Original

Noisy (12.77dB)

Result (29.87dB)

Inpainting [Mairal, Elad & Sapiro, ('06)]

Our experiments lead to state-of-the-art inpainting results.

Original

80% missing

Result

Inpainting [Mairal, Elad & Sapiro, ('06)]

Our experiments lead to state-of-the-art inpainting results.


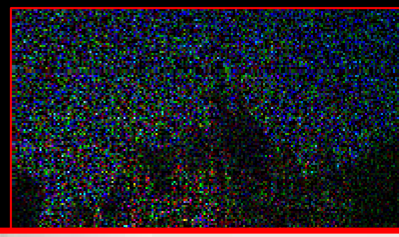
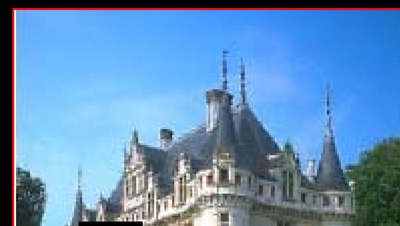

Or

Video Denoising [Protter & Elad ('06)]

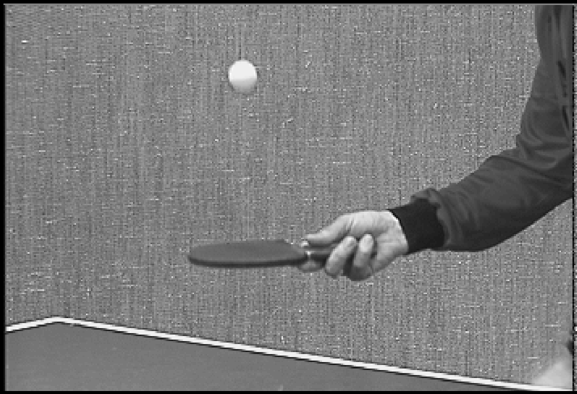

Original

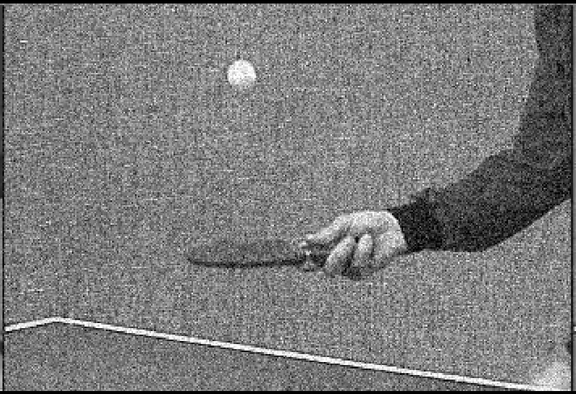

Noisy (σ=25)

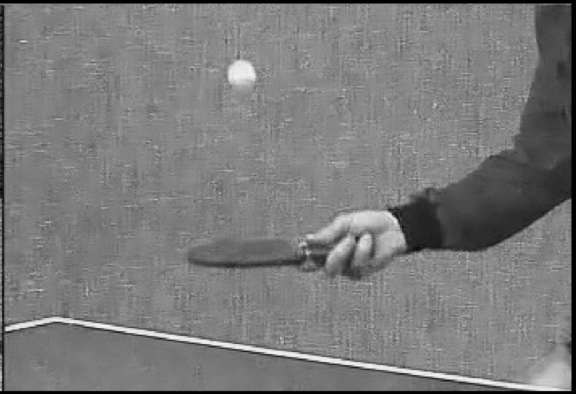

Denoised

Original

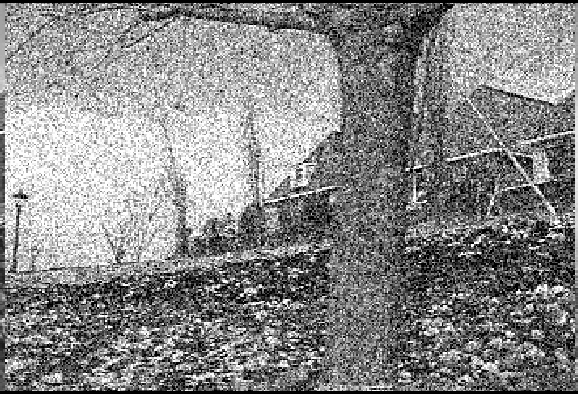

Noisy (σ=50)

Denoised

Video Denoising [Protter & Elad ('06)]

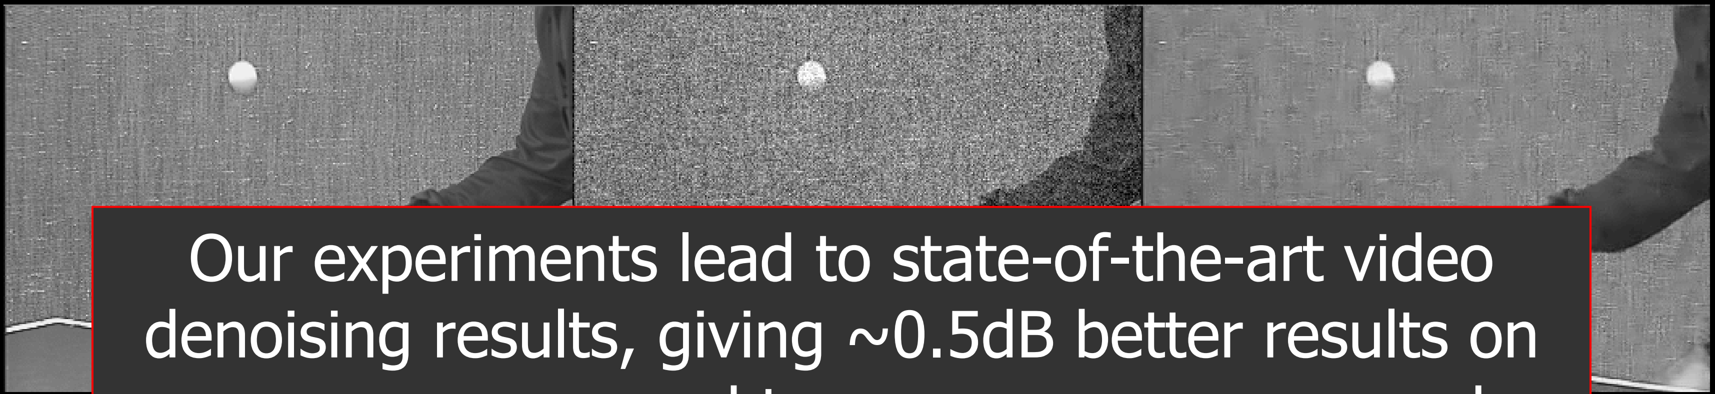

Our experiments lead to state-of-the-art video denoising results, giving ~0.5dB better results on average, compared to [Boades, Coll & Morel ('05)] and [Rusanovskyy, Dabov, & Egiazarian ('06)]

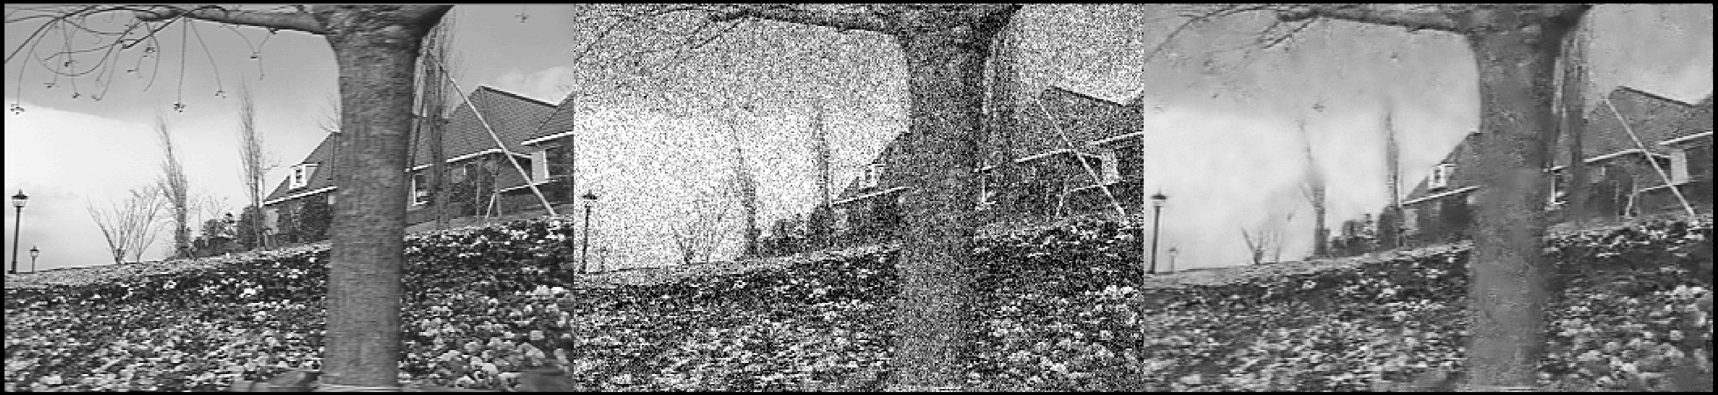

Original Noisy (σ=50) Denoised

Facial Image Compression [Brytt and Elad (`07)]

Results for **550** Bytes per each file

Facial Image Compression [Brytt and Elad (`07)]

Results for **400** Bytes per each file

To Conclude

Effective (yet simple) model for signals/images is key in getting better algorithms for various applications

Which model to choose?

Sparseland is an emerging model with high potential. It is based on sparse and redundant representations of signals, and learned dictionaries

Is it working well?

It has been deployed to many applications, in all leading to state-of-the-art results. More work is required to extend its usability to other domains.