CSE 585T/ESE 585A: Sparse Modeling for Imaging and Vision

Ulugbek S. Kamilov
Computational Imaging Group (CIG)
Washington University, St. Louis, USA
cigroup.wustl.edu • @wustlcig • kamilov@wustl.edu

Fall 2018

Three simple stories about sparse models

- **Imaging:** Sparsity revolution in imaging
- **Vision:** Foreground-background separation
- **Learning:** Recommender systems

Three simple stories about sparse models

- **Imaging:** Sparsity revolution in imaging
- Vision: Foreground-background separation
- Learning: Recommender systems

A contemporary paradox:
we acquire more data than we end up using

A contemporary paradox: we acquire more data than we end up using

iPhone 7 with 12 MP camera

Raw: 26 MB

JPEG: 500 KB

- Large amount of measured data
- Most of this data is thrown away afterwards

Washington University in St. Louis

A contemporary paradox:
we acquire more data than we end up using

Energy of natural images is highly concentrated

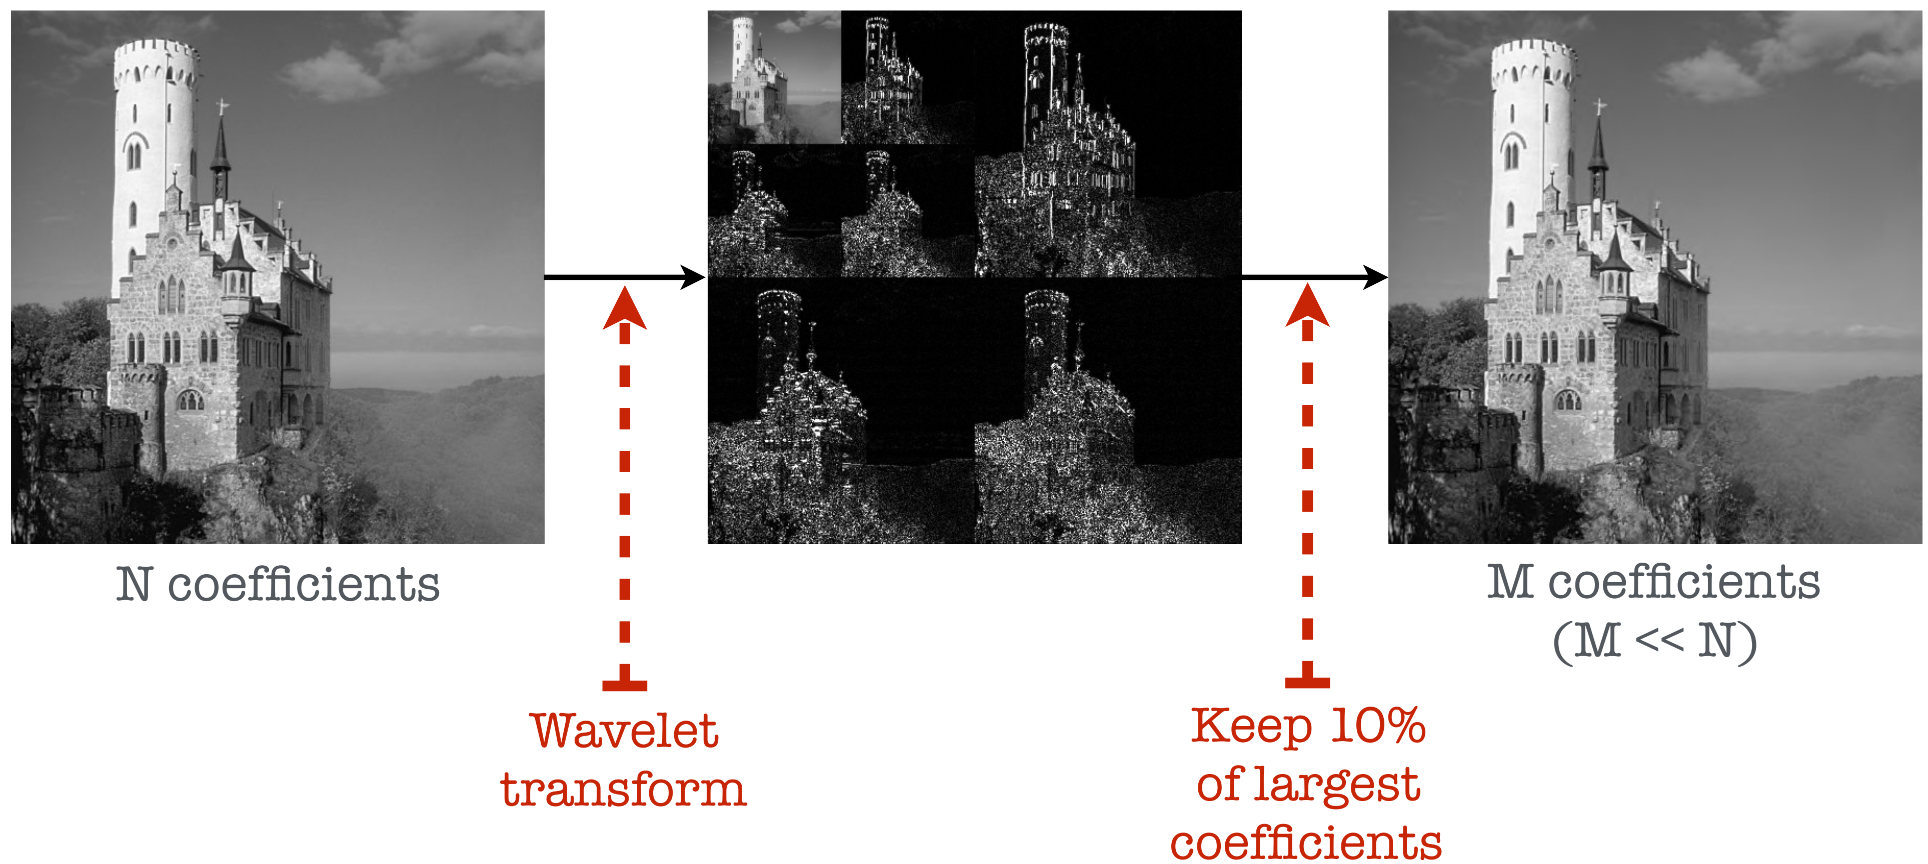

Wavelet transform is at the heart of JPEG-2000 image compression standard

the jpeg 2000 alliance

Washington University in St. Louis

A contemporary paradox: we acquire more data than we end up using

A long established sensing pipeline in imaging

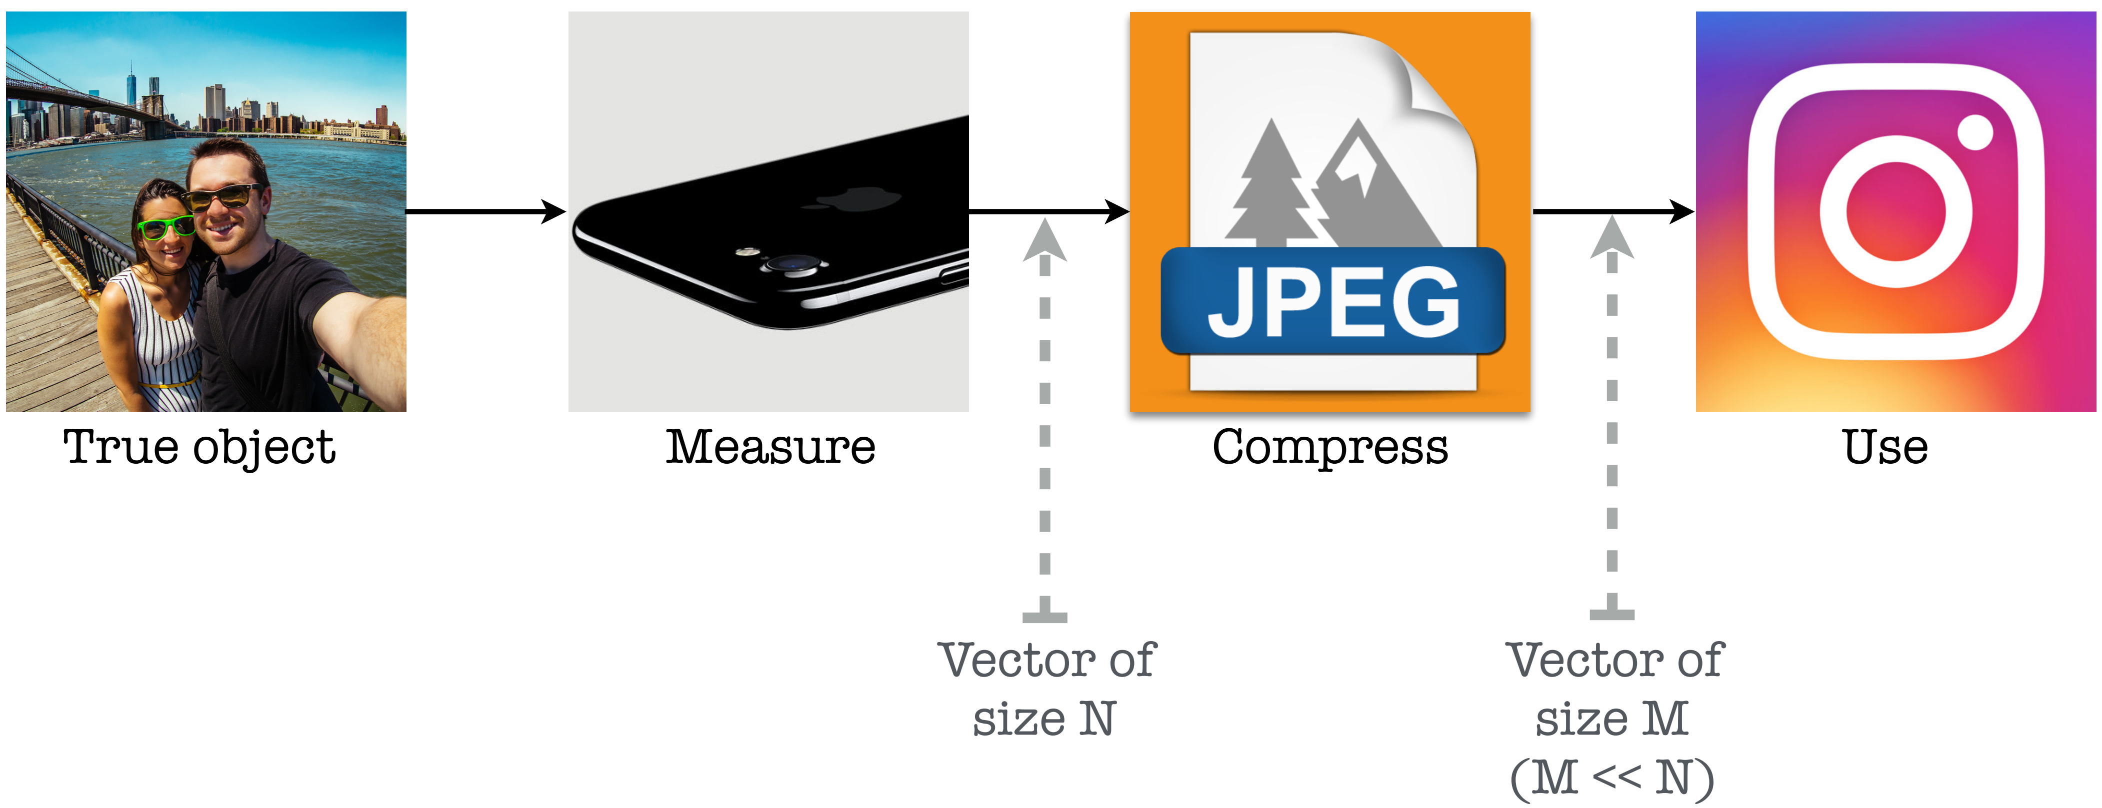

Washington University in St. Louis

A contemporary paradox: we acquire more data than we end up using

A long established sensing pipeline in imaging

True object

Measure

Compress

Use

Vector of size N

Vector of size M (M << N)

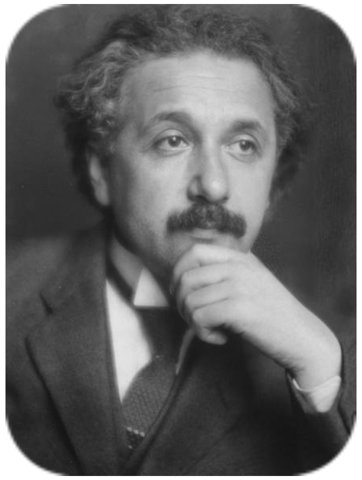

What are possible limitations of this pipeline?

Compressive imaging requires some theory and advanced algorithms

Compressive imaging requires some theory and advanced algorithms

What happens if we simply take less measurements?

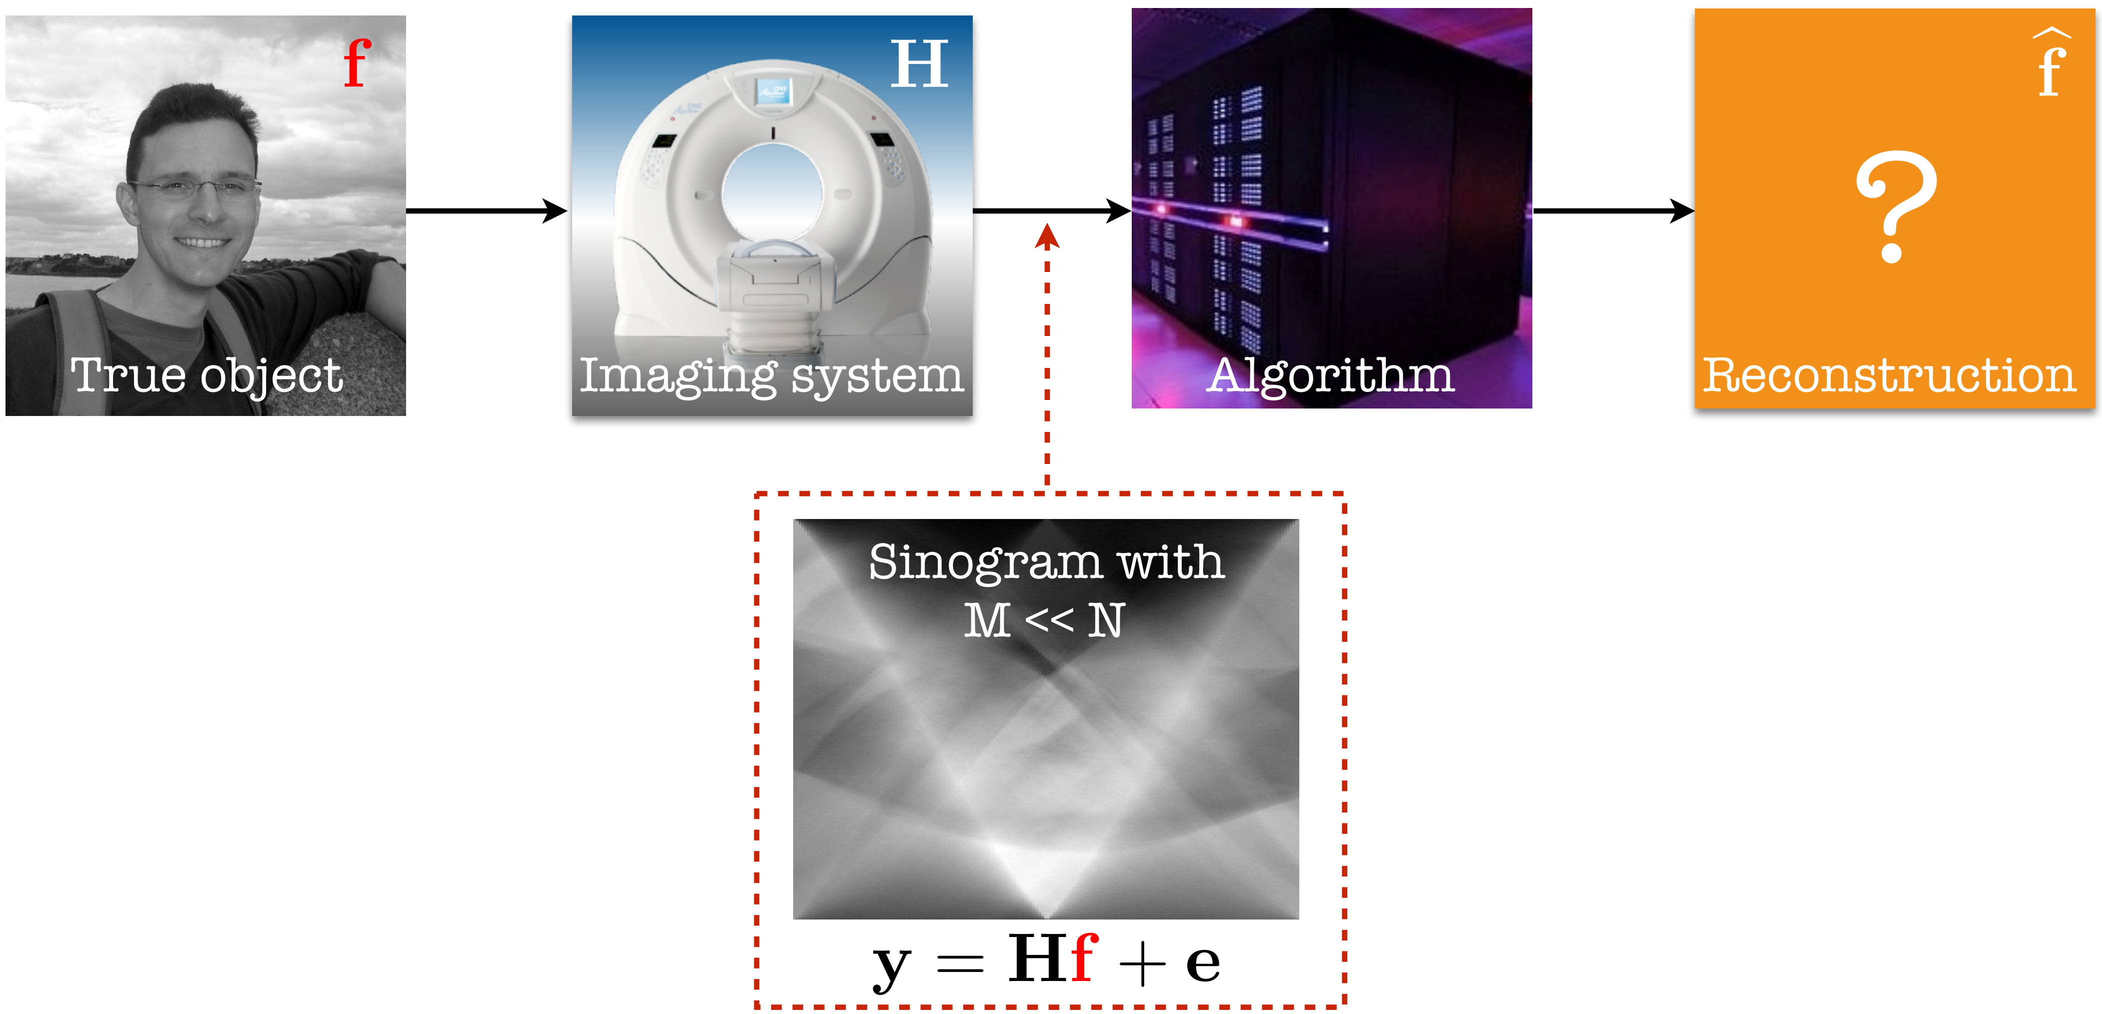

Washington University in St. Louis

Compressive imaging requires some theory and advanced algorithms

What happens if we simply take less measurements?

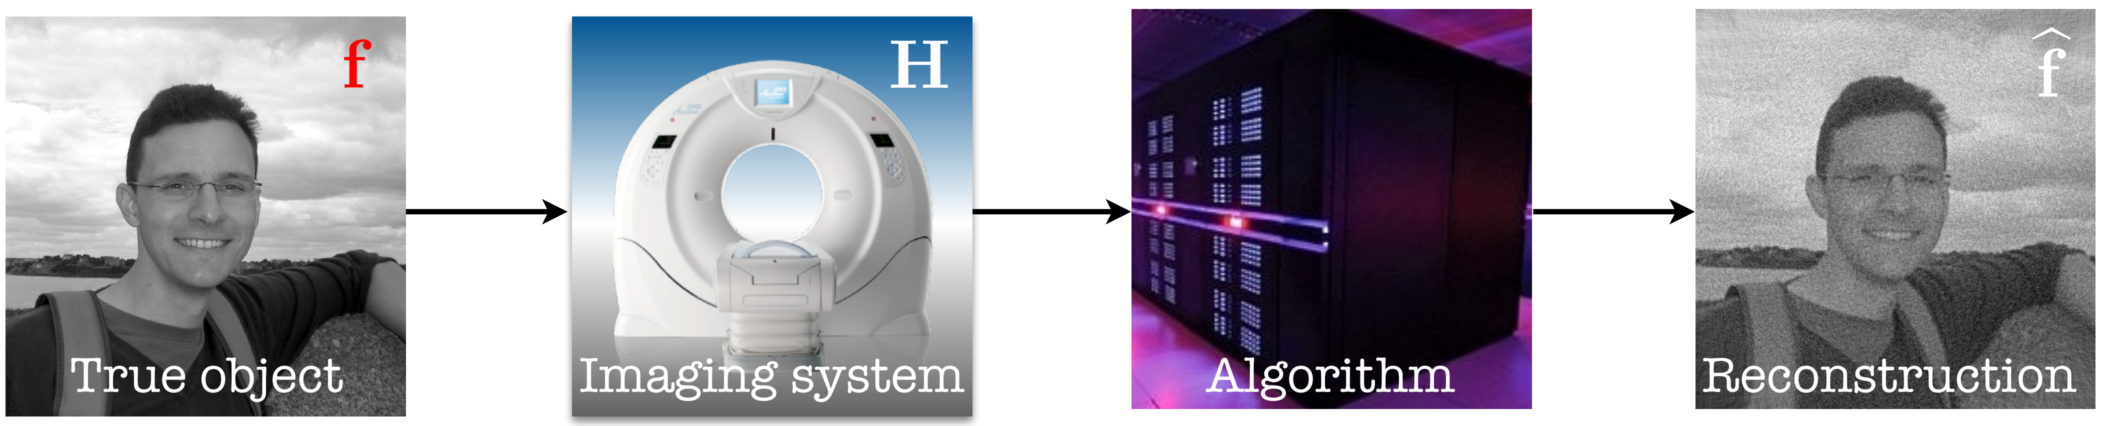

Least-squares solution (*i.e.*, basic reconstruction)

$$\min_{\mathbf{f} \in \mathbb{C}^N} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{f}\|_{\ell_2}^2 \right\}$$

Remarks:

1. Noise amplification
2. Bad for compressive imaging

Compressive imaging requires some theory and advanced algorithms

What happens if we simply take less measurements?

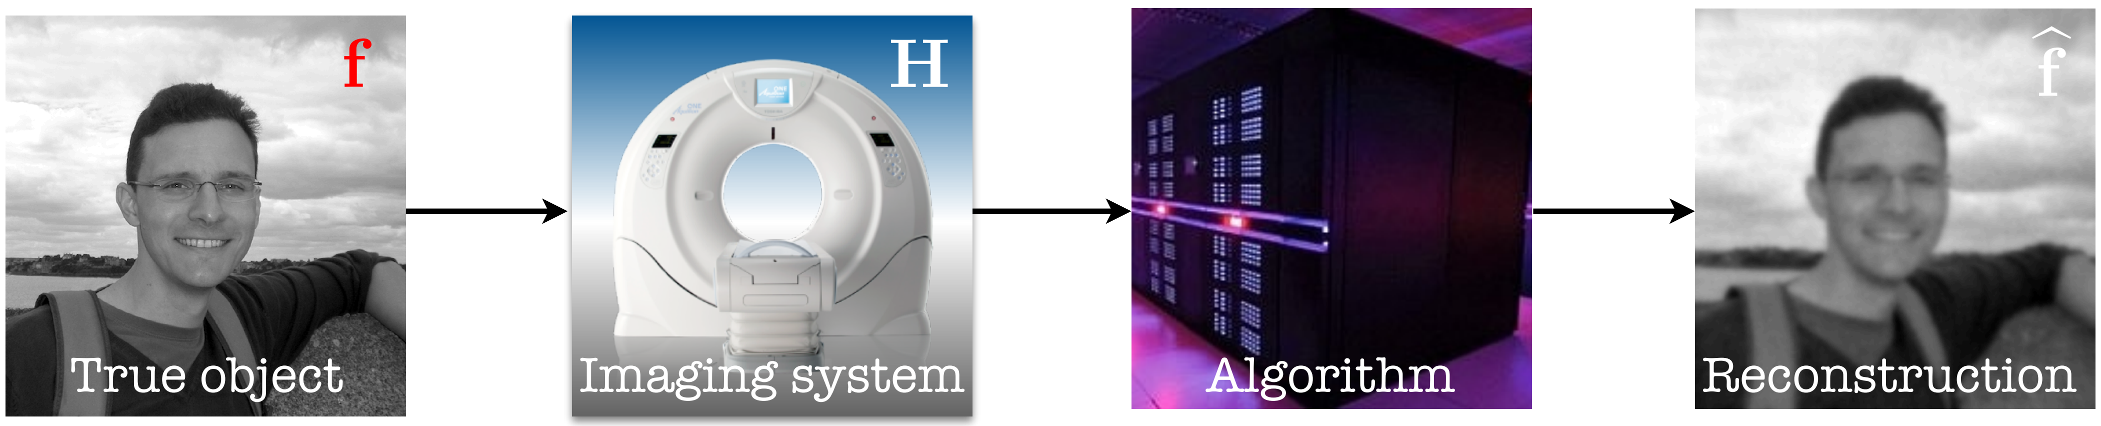

Tikhonov regularization (*i.e.*, 20th century technology)

$$\boxed{\min_{\mathbf{f}\in\mathbb{C}^N}\left\{\frac{1}{2}\|\mathbf{y}-\mathbf{H}\mathbf{f}\|_{\ell_2}^2+\lambda\|\mathbf{D}\mathbf{f}\|_{\ell_2}^2\right\}}$$

Remarks:
1. Blurry images
2. Linear solution

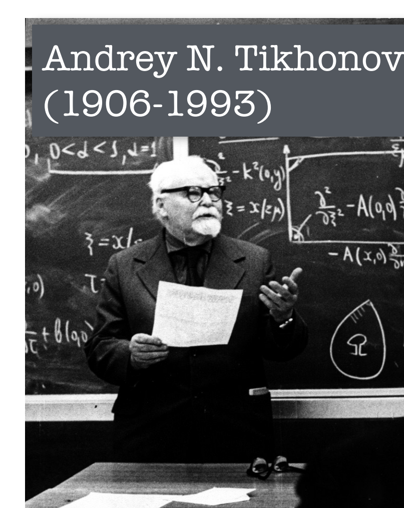
Andrey N. Tikhonov (1906-1993)

Compressive imaging requires some theory and advanced algorithms

What happens if we simply take less measurements?

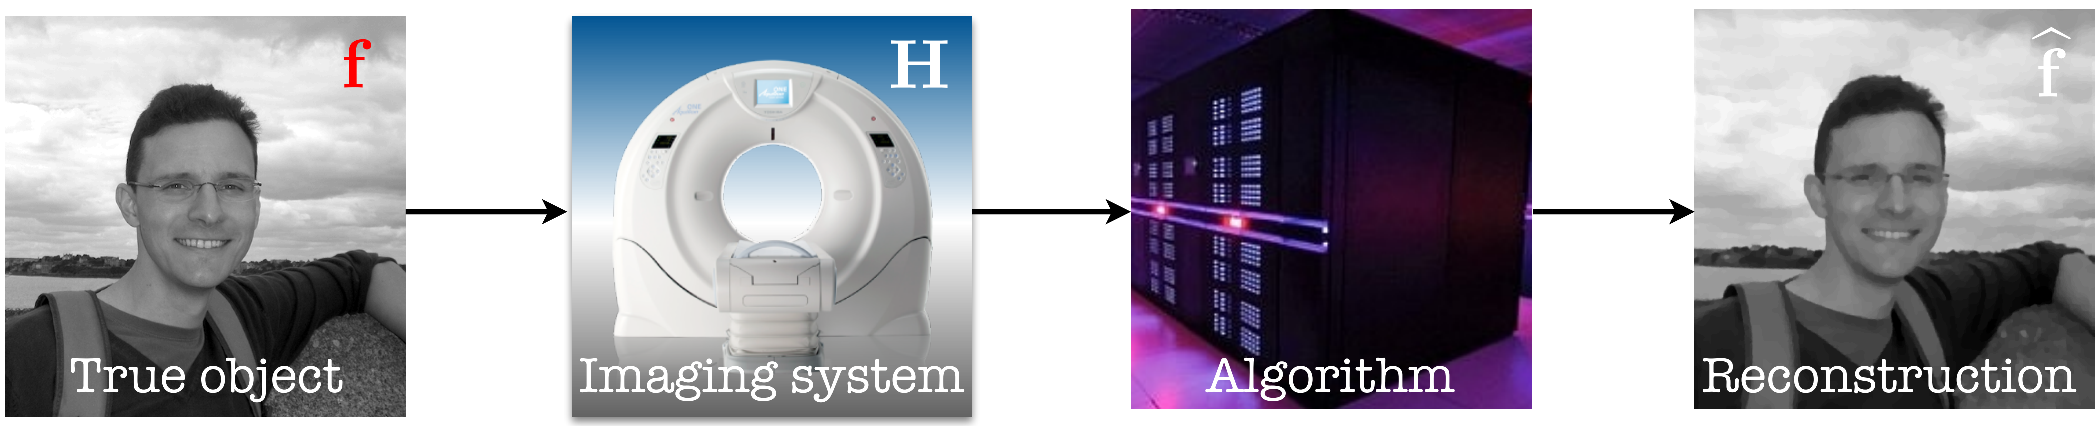

Sparse regularization: (20th century) $\ell_2 \rightarrow \ell_1$ (21st century)

$$\boxed{\min_{\mathbf{f} \in \mathbb{C}^N} \left\{ \frac{1}{2} \|\mathbf{y} - \mathbf{H}\mathbf{f}\|_{\ell_2}^2 + \lambda \|\mathbf{D}\mathbf{f}\|_{\ell_1} \right\}}$$

Remarks:

1. Compressive acquisition
2. Nonlinear algorithms

Compressive imaging requires some theory and advanced algorithms

First part of the course will cover the theory of compressive imaging:

- How to optimally measure?
- How many measurements are needed?
- How to reconstruct?

David Donoho

Emmanuel Candes

Wired, March 2010

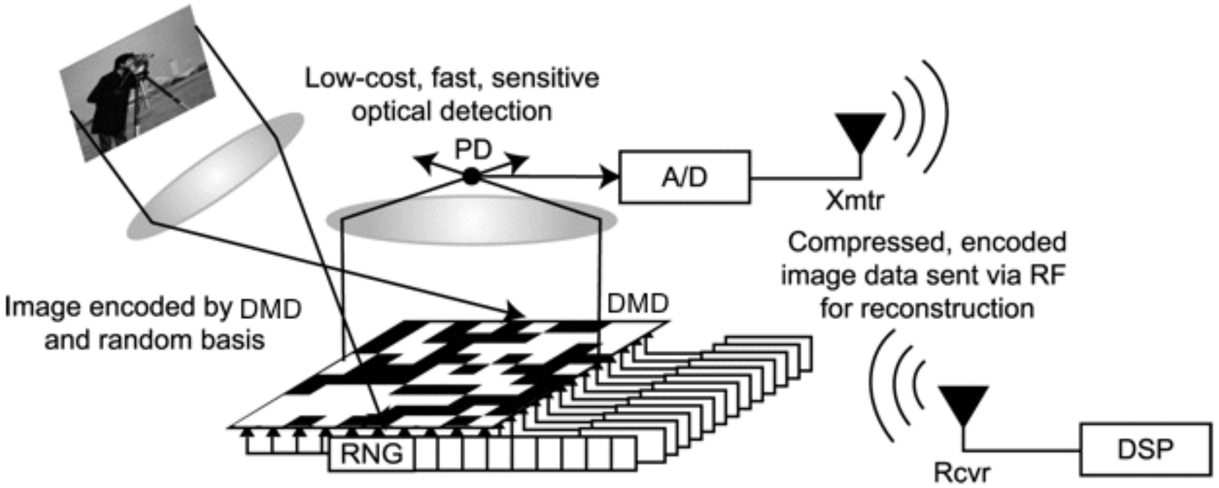

Single Pixel Camera

Candès, Romberg, Tao, *IEEE Trans. Inf. Theory*, 2006
Donoho, *IEEE Trans. Inf. Theory*, 2006

Compressive imaging requires some theory and advanced algorithms

Application: MRI done 10x faster

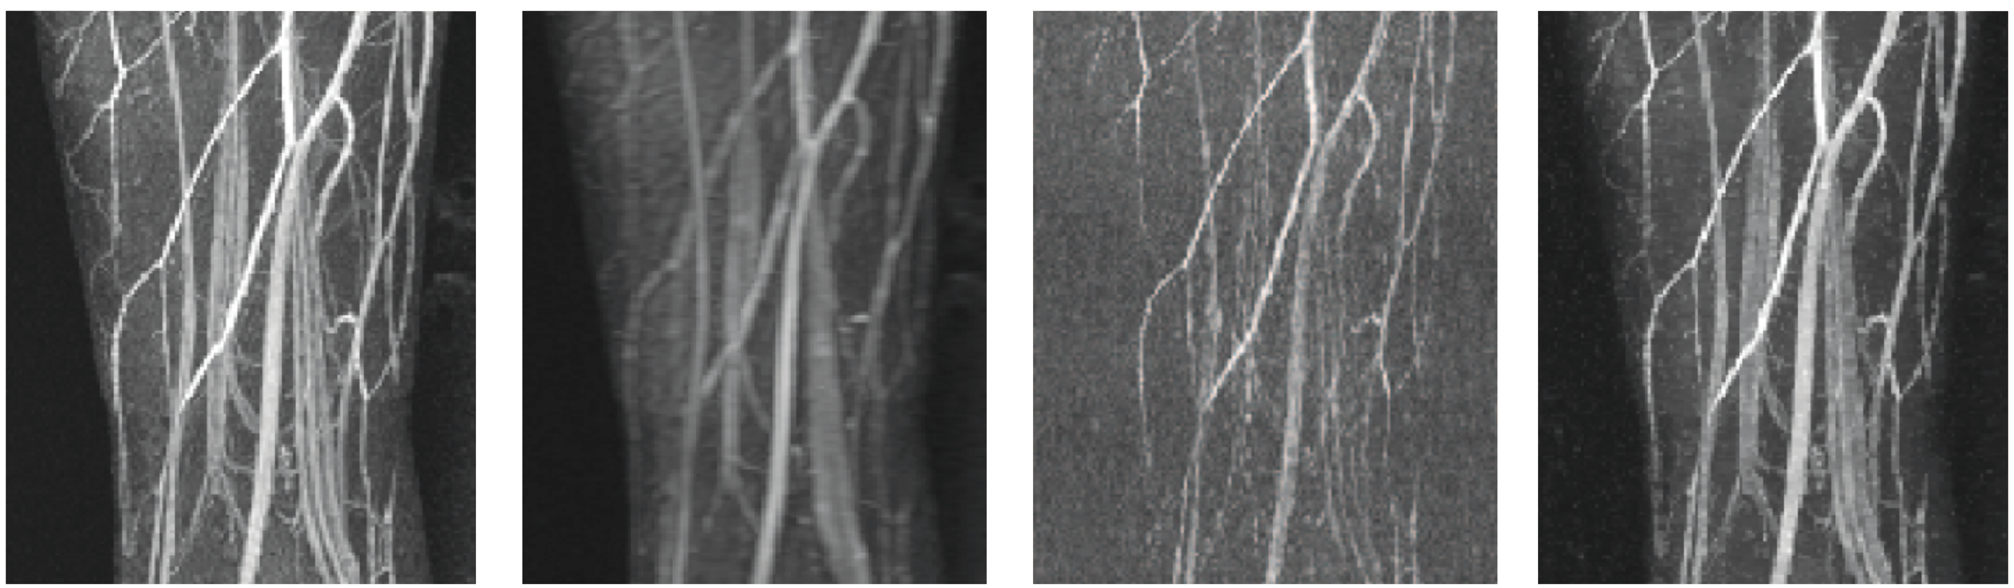

Full sampling (slow) | Low-resolution acquisition | Linear reconstruction | Sparse reconstruction

Big interest from leading companies:
Siemens, GE, Phillips, and *etc.*

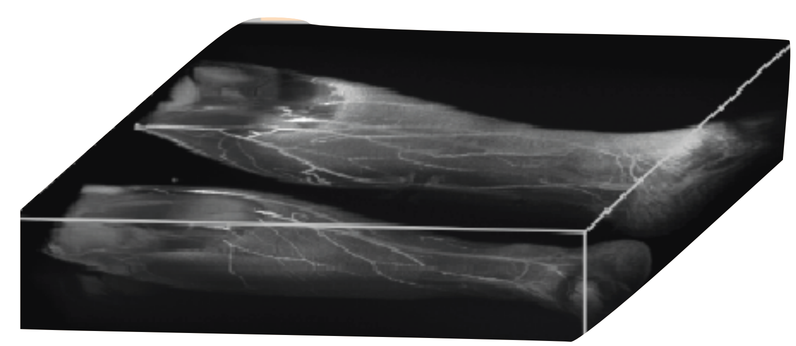

3D MRI image

Lustig *et al.*, *IEEE Signal Process. Mag.*, 2008

Compressive imaging requires some theory and advanced algorithms

Application: Autonomous driving

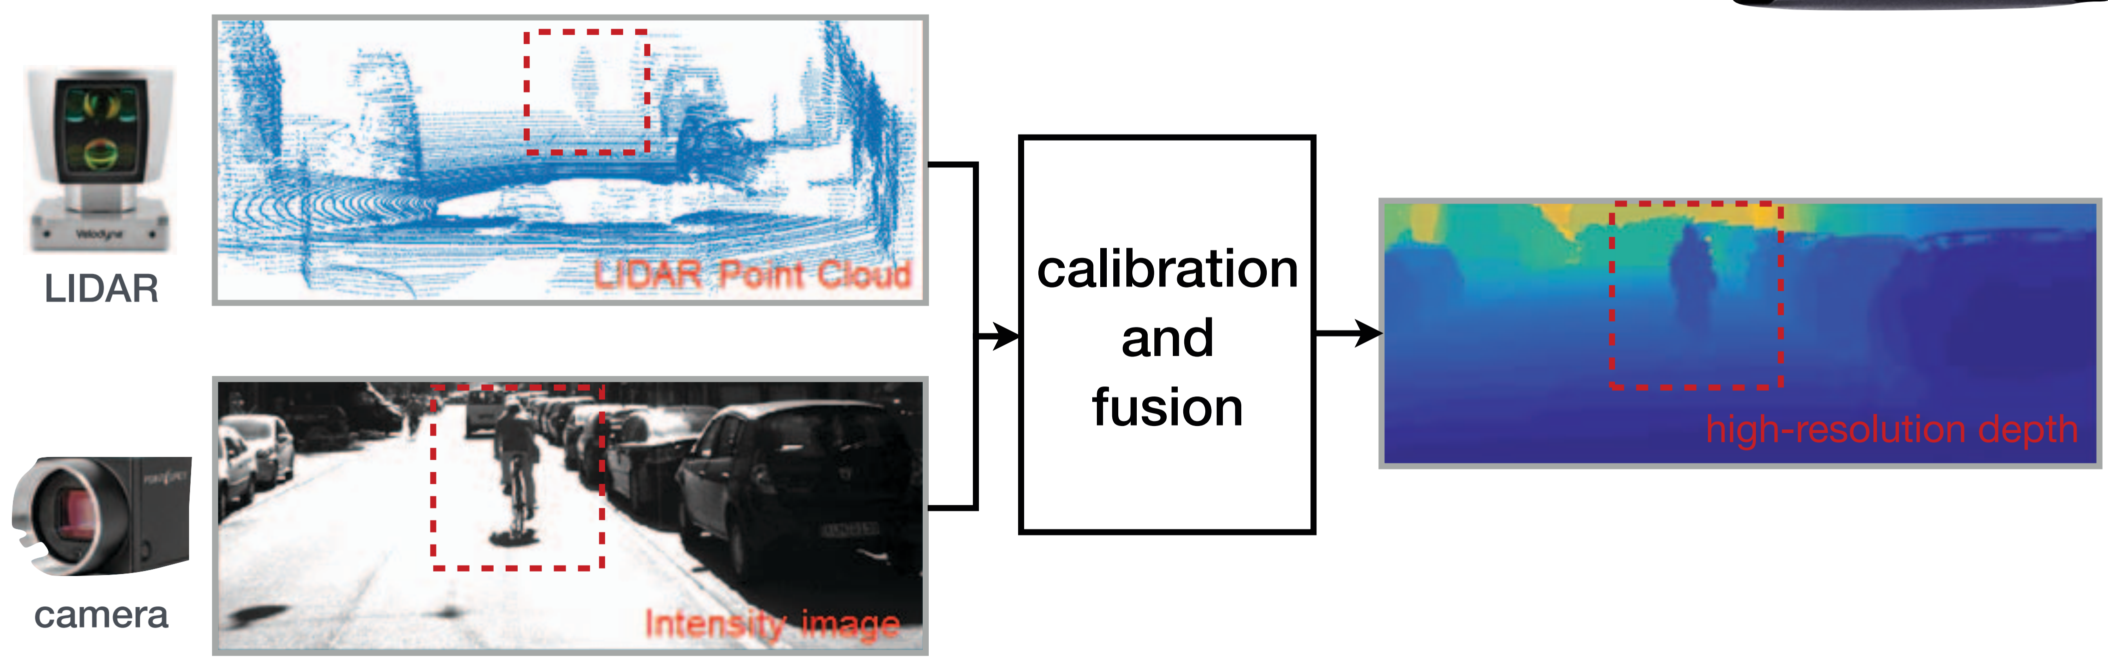

Big interest from leading companies: Google, Apple, Nvidia, Mitsubishi, etc.

Degraux, Kamilov, Boufounos, and Liu., *IEEE ICIP,* 2018

Three simple stories about sparse models

- Imaging: Sparsity revolution in imaging
- **Vision:** Foreground-background separation
- Learning: Recommender systems

Sparse modeling for computer vision

Application: advanced surveillance technology

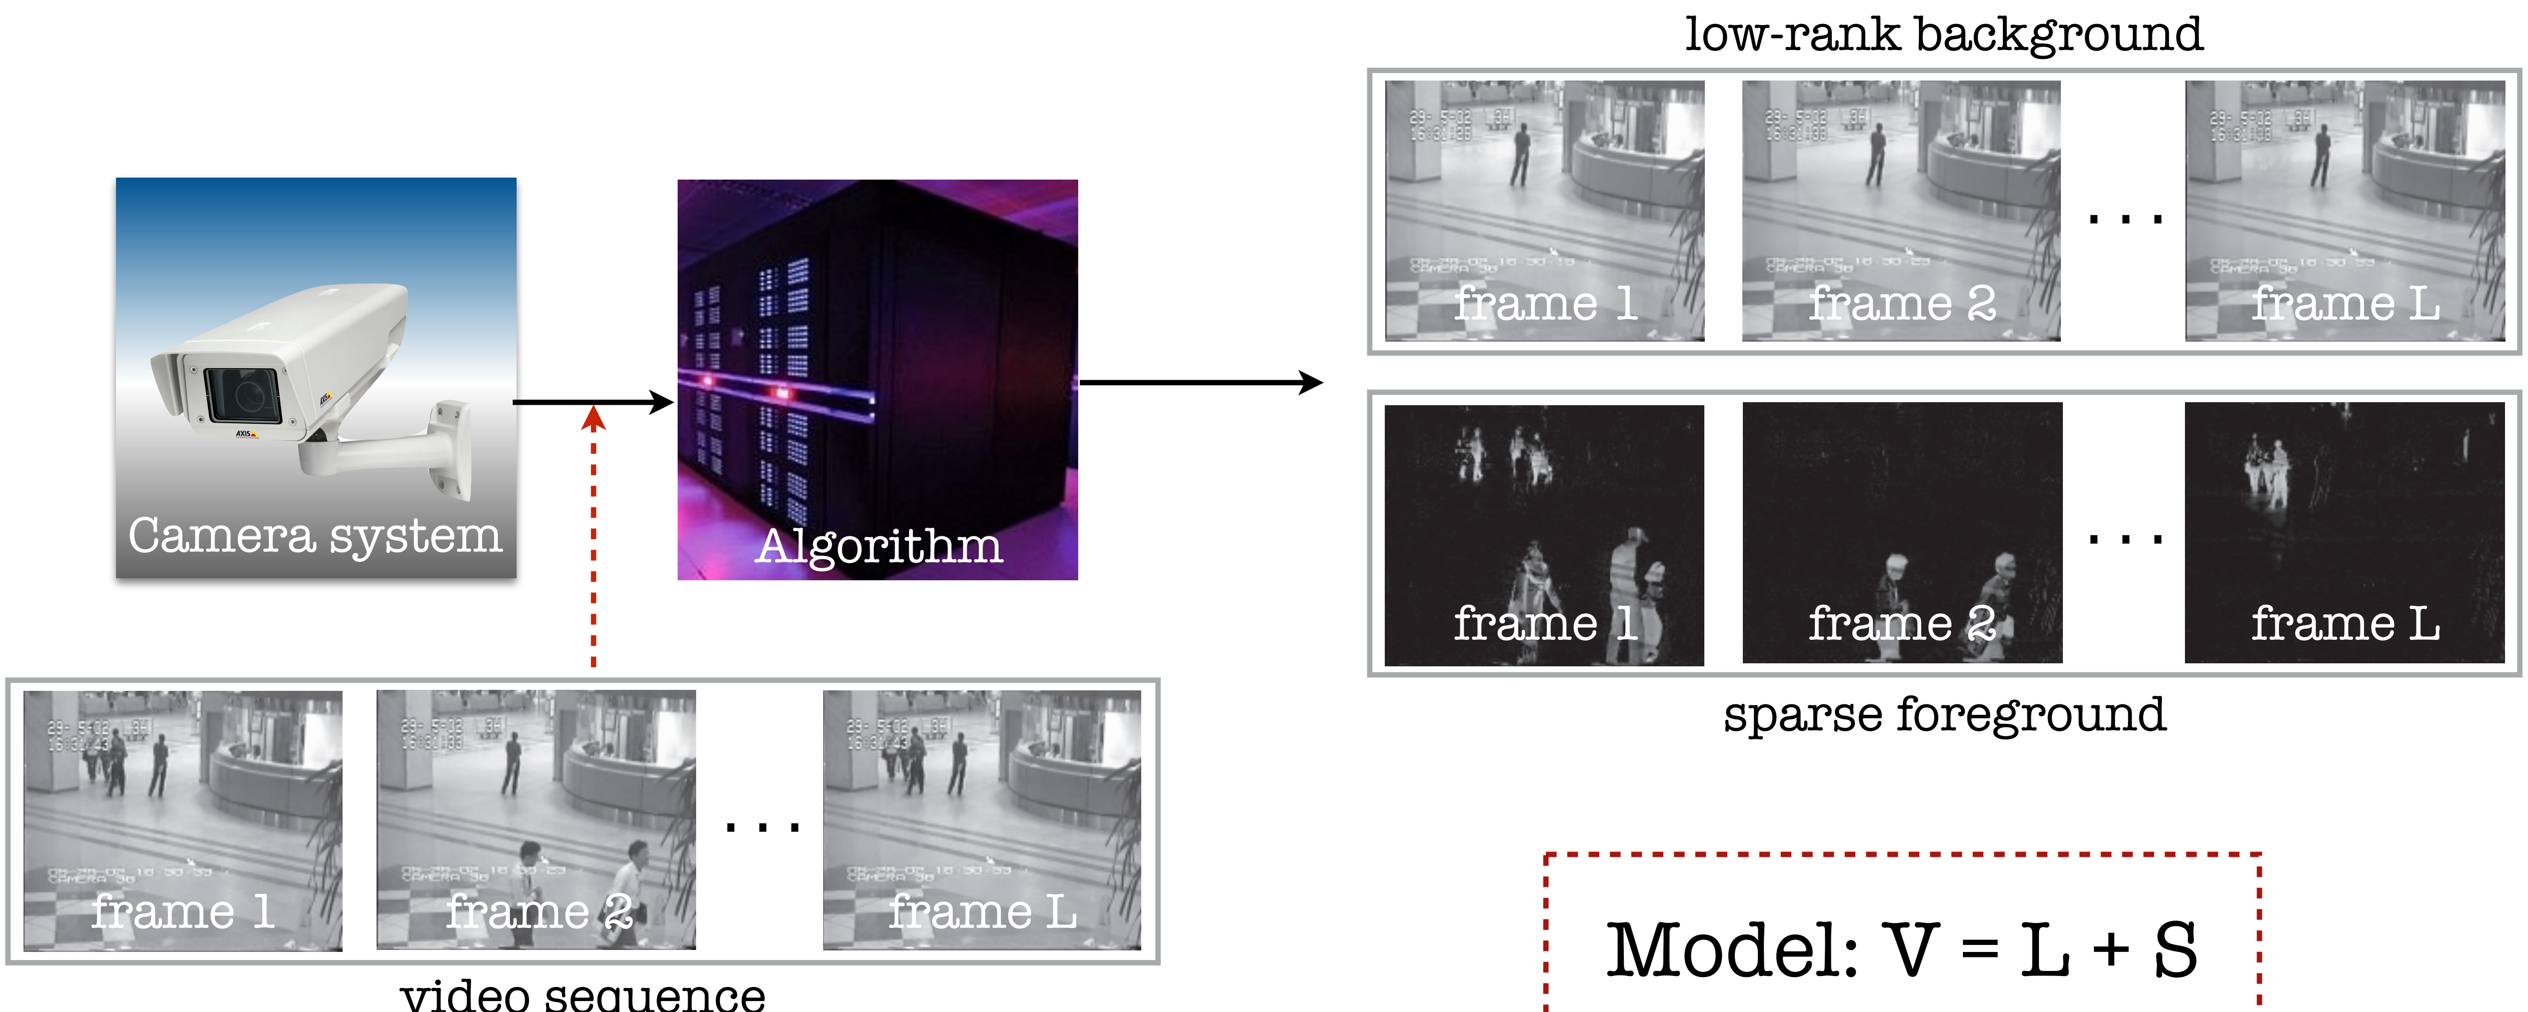

Model: V = L + S

Wright *et al.*, *Proc. NIPS*, 2009

Sparse modeling for computer vision

Application: advanced surveillance technology

Model: V = L + S

Wright *et al.*, *Proc. NIPS*, 2009

Sparse modeling for computer vision

Application: removing shadow and specularities from faces

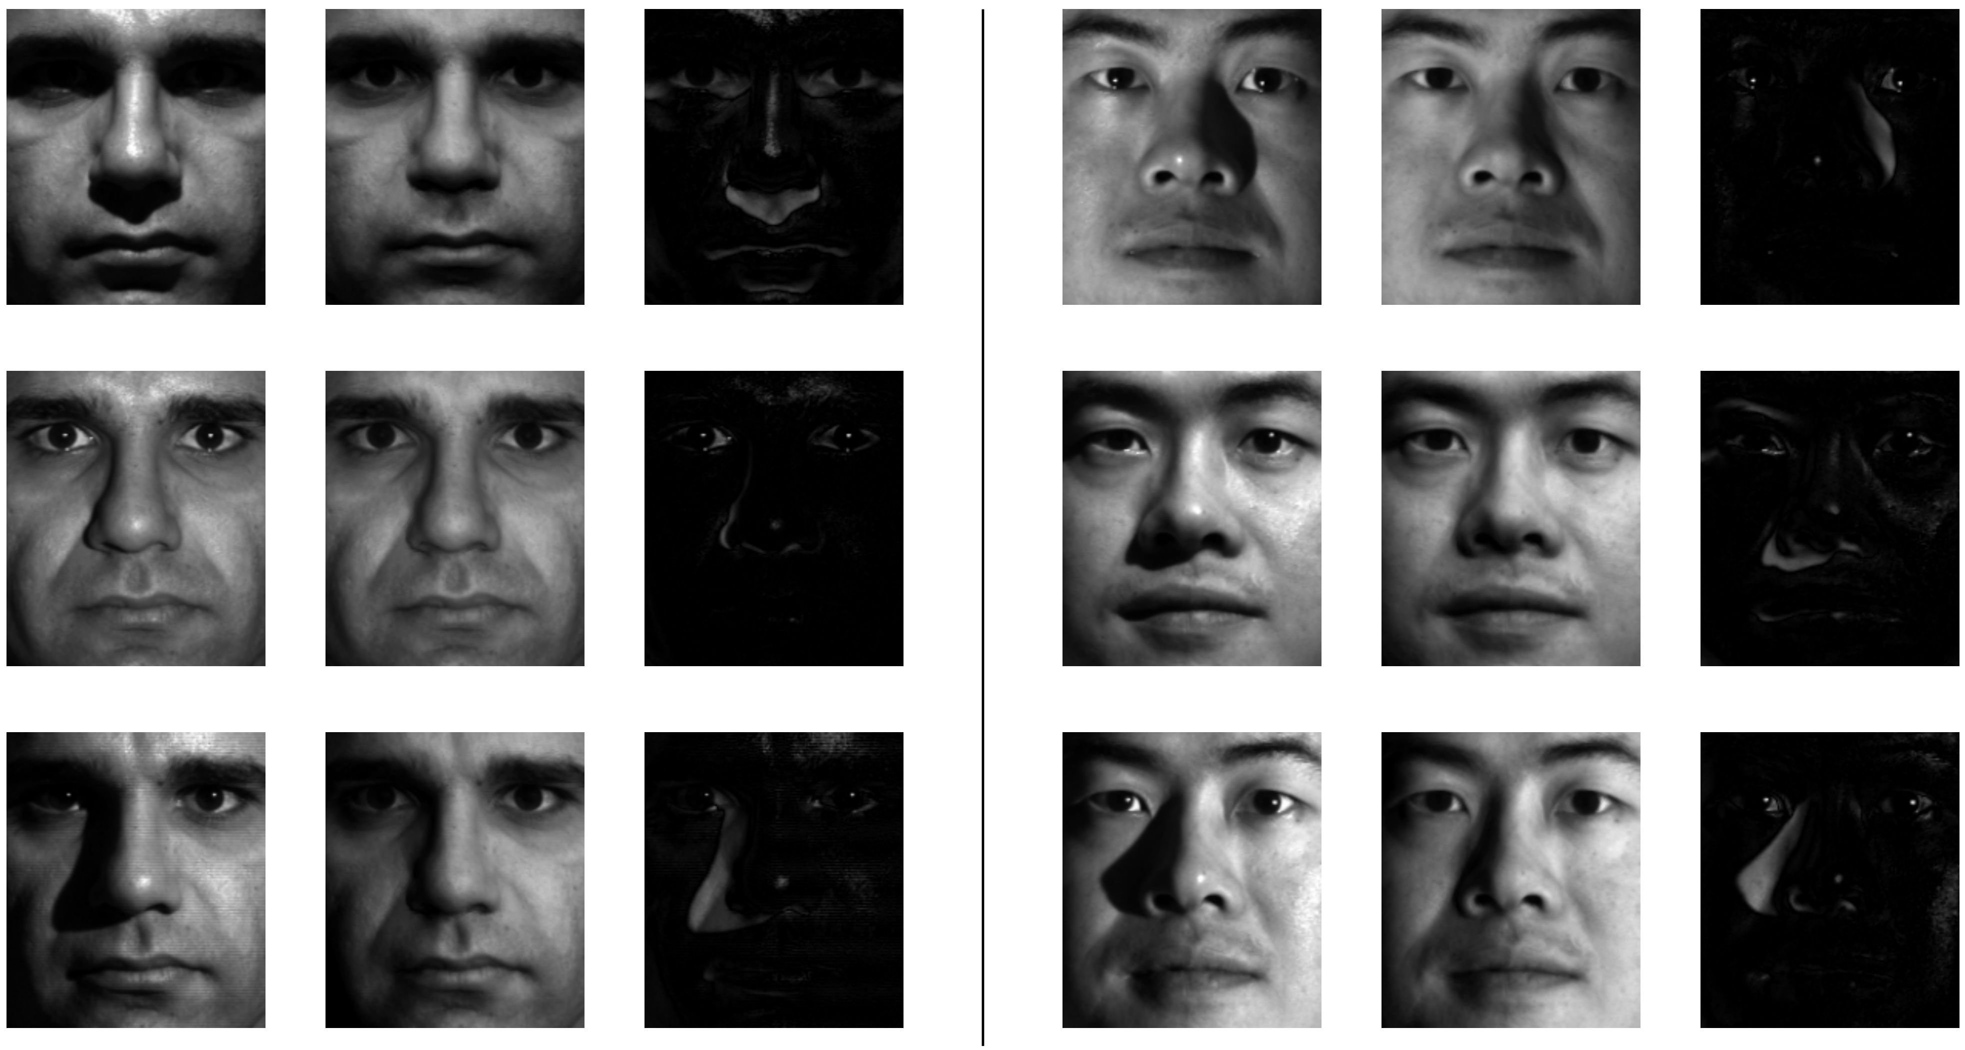

face low-rank sparse | face low-rank sparse

$\boxed{\text{Model: } V = L + S}$

Wright *et al.*, *Proc. NIPS*, 2009

Three simple stories about sparse models

- Imaging: Sparsity revolution in imaging
- Vision: Foreground-background separation
- **Learning:** Recommender systems

Sparse modeling for machine learning

Application: Recommender systems

$$
\begin{array}{cc}
 & \text{Movies: 18K} \\
\text{Users: 0.5M} & \begin{bmatrix}
\times & ? & ? & ? & \times & ? \\
? & ? & \times & \times & ? & ? \\
\times & ? & ? & \times & ? & ? \\
? & ? & \times & ? & ? & \times \\
\times & ? & ? & ? & ? & ? \\
? & ? & \times & \times & ? & ?
\end{bmatrix}
\end{array}
$$

$1 million dollar
"Netflix" problem

Practical information

Practical information

Course website:
https://cigroup.wustl.edu/teaching/cse-585t-2018

Recommended reading:

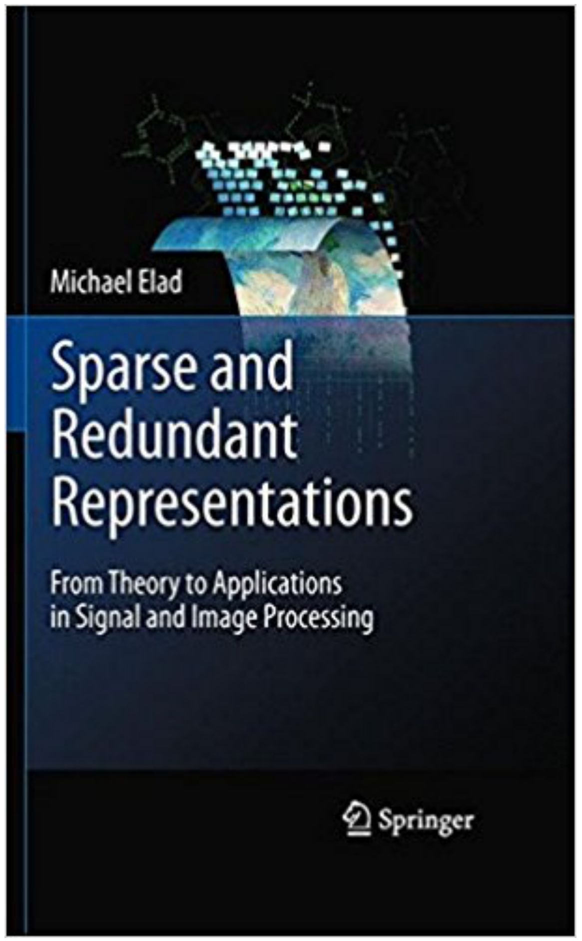

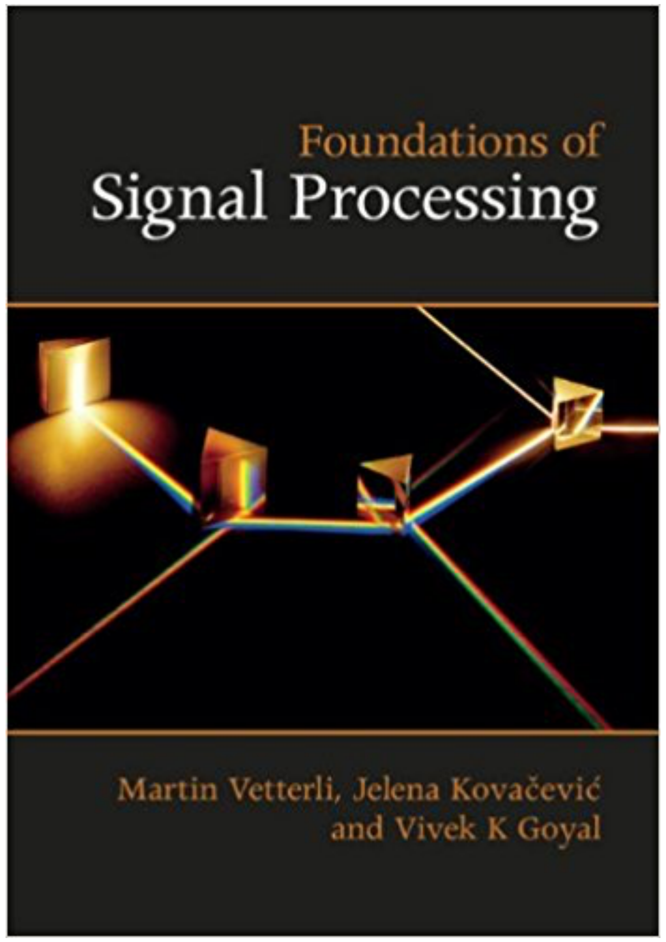

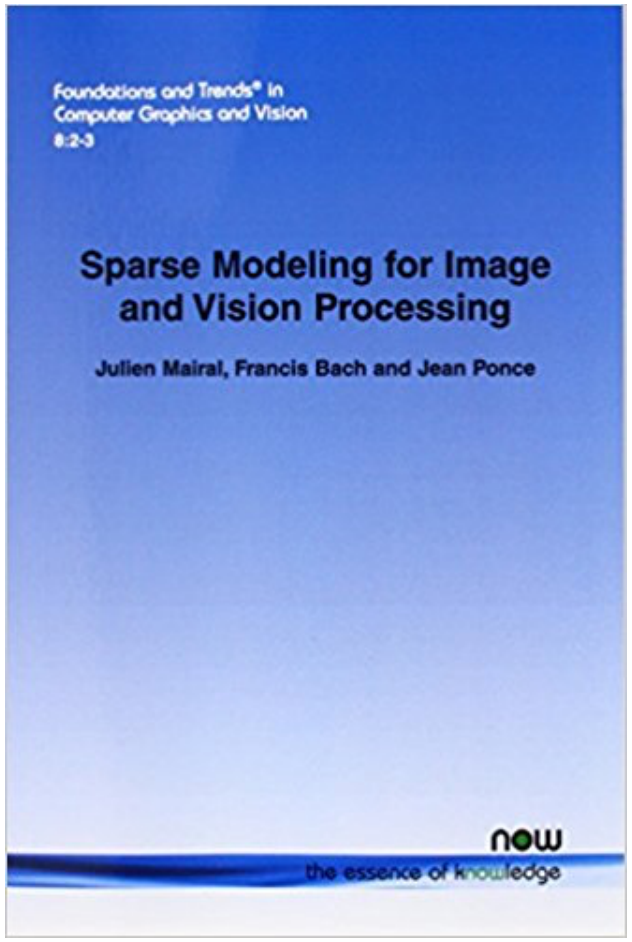

Practical information

High-level learning plan:

- Introduction, motivation, basics of linear algebra
- Theory: recovery guarantees and compressive sensing
- Algorithms: iterative thresholding, splitting methods, and stochastic optimization
- Applications: imaging, learning, and vision

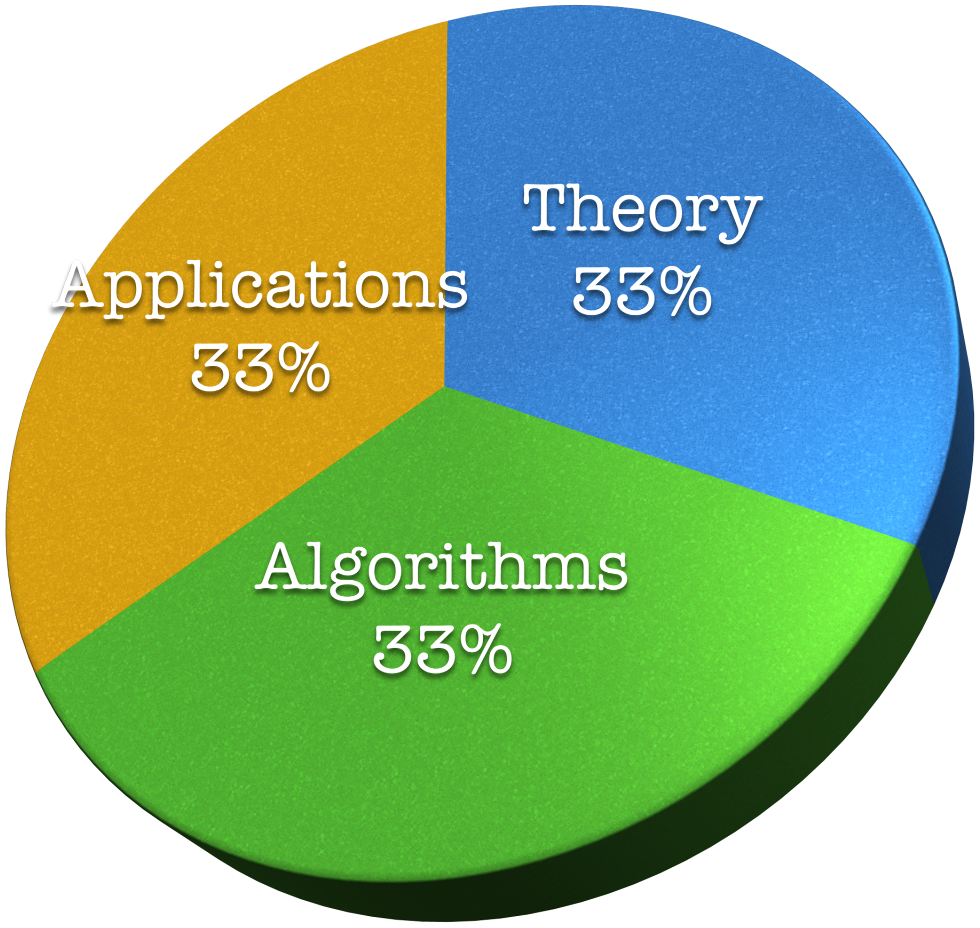

Practical information

By the end of the semester, hopefully you will be able to

- Understand the concepts behind sparse modeling
- Comment on algorithms and do some implementations
- Read most publication in the field
- Apply most relevant techniques in your careers

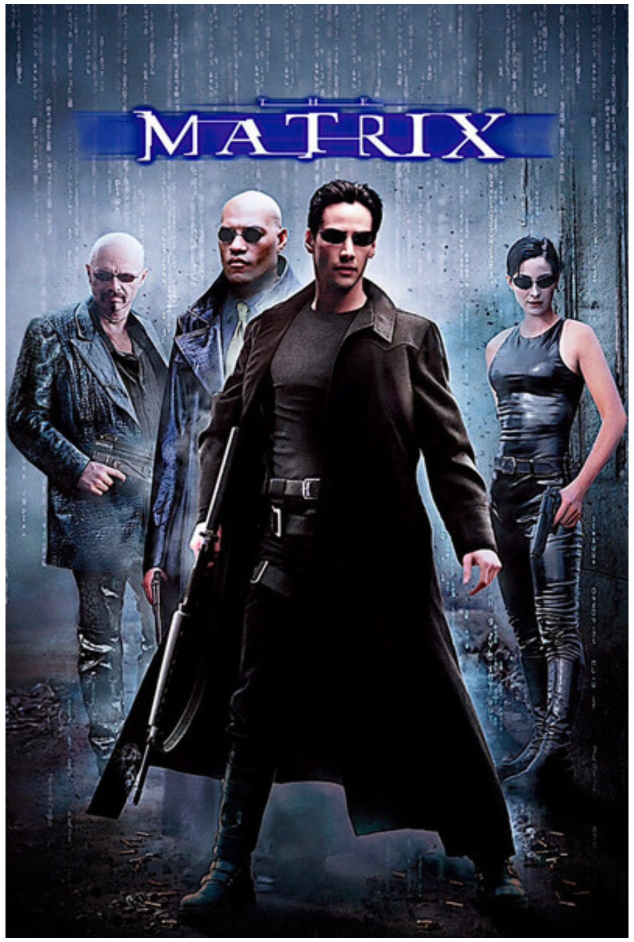

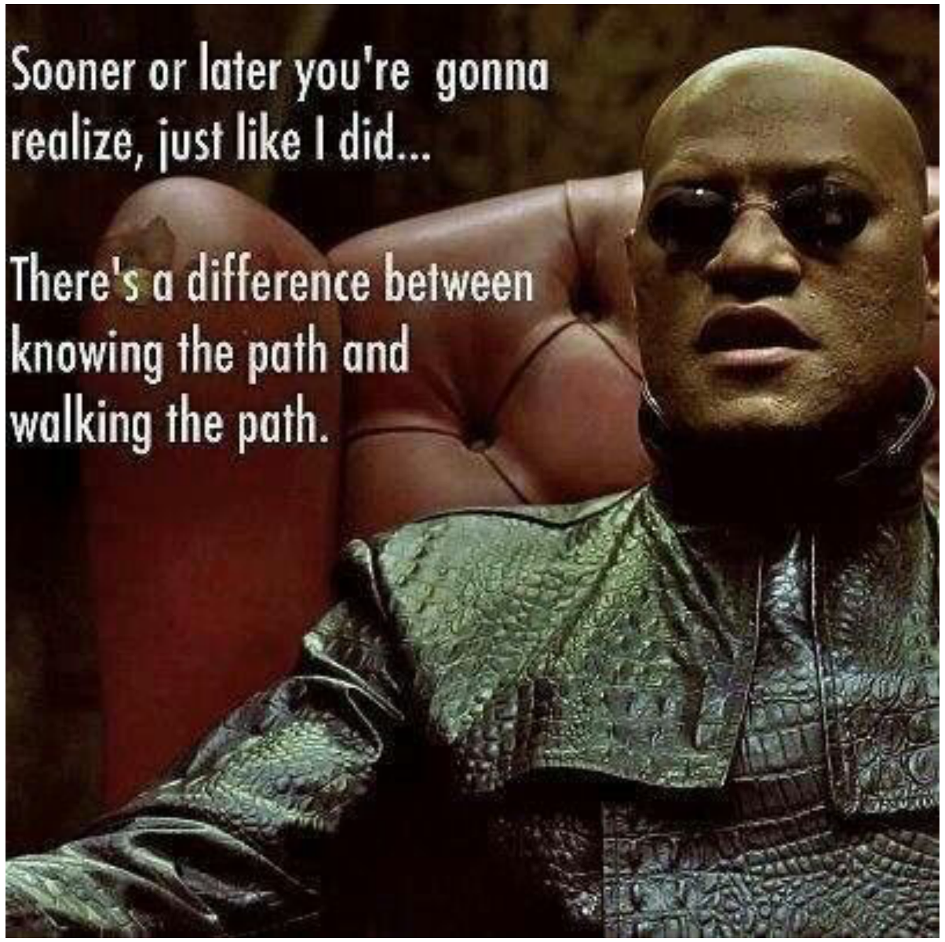

Practical information

Grading: Attendance (10%)

➡ Important since we have no homework and exams

➡ No attendance sheet, but I will know who is here

Practical information

Grading: Class participation (30%)

- ➡ Regular student lectures
- ➡ Four students work together each time
- ➡ Check the list of topics on the course page
- ➡ Topics will include most popular publications in the field
- ➡ Understand the basics and provide your comment
- ➡ Extra appreciation for demo and implementations

Practical information

Grading: Class participation (30%)

- ➡ Topic 1: Compressive sensing
- ➡ Topic 2: Sparse MRI
- ➡ Topic 3: Reweighted l1-minimization
- ➡ Topic 4: Compressive sensing and compression
- ➡ Topic 5: Fast ISTA
- ➡ Topic 6: Variable splitting and ADMM
- ➡ Topic 7: Total variation
- ➡ Topic 8: Total generalized variation
- ➡ Topic 9: Dictionary learning
- ➡ Topic 10: Convolutional dictionary learning

Practical information

Grading: Class participation (30%)

- ➡ Topic 11: Online learning
- ➡ Topic 12: Learned ISTA
- ➡ Topic 13: Super-resolution CNN (SR-CNN)
- ➡ Topic 14: Trainable nonlinear reaction diffusion (TNRD)
- ➡ Topic 15: Plug-and-play priors (PnP)
- ➡ Topic 16: Regularization by denoising (RED)
- ➡ Topic 17: Bayesian compressive sensing
- ➡ Topic 18: Approximate message passing (AMP)
- ➡ Topic 19: Using trained CNNs as priors
- ➡ Topic 20: Deep image prior (DIP)

Practical information

Grading: Project (60%)

- ➡ Group sizes between 1 and 4 students
- ➡ Each student submits individual proposal
- ➡ Each student submits individual report
- ➡ Group presentation
- ➡ Any topic related to the course is acceptable

Conclusion

Sparse modeling is now extensively used in industry and research

The goal of CSE 585T is to help you understand and apply the basics

This is an active research area with many open questions

CONTACT INFO

Ulugbek S. Kamilov
Computational Imaging Group (CIG)
Washington University in St. Louis
kamilov@wustl.edu
http://kamilov.info
@ukmlv