Compressed Sensing under Optimal Quantization

Alon Kipnis (Stanford)
Galen Reeves (Duke)
Yonina Eldar (Technion)
Andrea Goldsmith (Stanford)

ISIT, June 2017

Table of Contents

Remote source coding

[Dobrushin & Tsybakov '62]

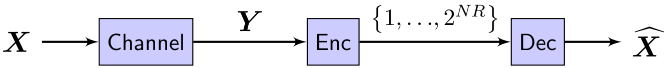

$$D_{\boldsymbol{X}|\boldsymbol{Y}}(R) = \min_{P(\hat{\boldsymbol{x}}|\boldsymbol{y})} \mathbb{E}\left[d(\boldsymbol{X}, \widehat{\boldsymbol{X}})\right]$$

- Estimation under communication constraints
- Learning from noisy data
- Close connection between inference and compression

Two coding schemes:

Estimate-and-compress

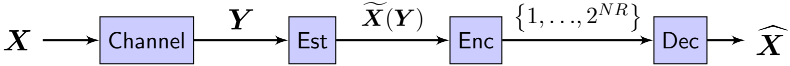

Compress-and-estimate [Kipnis, Rini, Goldsmith '16]

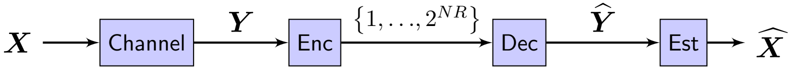

Example: IID source, Gaussian noise, MSE distortion

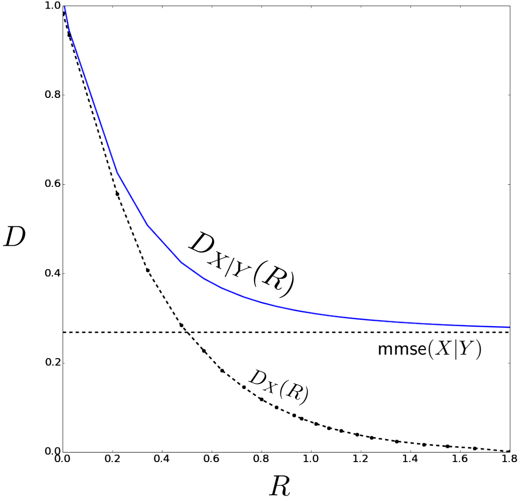

Compressed sensing with quantization

$$\boldsymbol{Y} = \sqrt{\mathsf{snr}}\,\boldsymbol{H}\boldsymbol{X} + \boldsymbol{W}, \qquad \boldsymbol{H} \in \mathbb{R}^{M \times N}$$

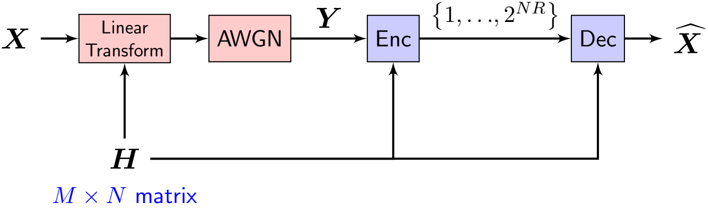

Goal is to understand fundamental tradeoffs between

- Bitrate R
- MSE distortion D
- Sampling rate $\delta = M/N$

Related work on quantization

- Gaussian sources – Kipnis, Goldsmith, Eldar, Weissman '16
- Scaler quantization – Goyal, Fletcher, Rangan '08
- Lasso recovery – Sue, Goyal '09
- Optimal high-bit asymptotic – Wu, Verdu '12, Dai, Milenkovic '11
- 1-bit quantization – Boufounos, Baraniuk '08, Plan, Vershynin '13
- Remote source coding with side information – Guler, MolavianJazi, Yener '15
- Lower bound on optimal quantization – Leinonen, Codreanu, Juntti, Kramer '16
- Sampling rate distortion – Boda, Narayan '17
- Distributed coding of multispectral images – Goukhshtein, Boufounos, Koike-Akino, Draper '17

Fundamental limits of compressed sensing

$$\boldsymbol{Y} = \sqrt{\mathsf{snr}}\, \boldsymbol{H}\boldsymbol{X} + \boldsymbol{W}, \qquad \boldsymbol{H} \in \mathbb{R}^{M \times N}, \qquad M, N \to \infty$$

- Guo & Verdú 2005 analyze large system limit with IID matrices using heuristic replica method from statistical physics.
- Rigorous results for special cases: Verdú & Shamai 1999, Tse & Hanly 1999, Montanari & Tse 2006, Korada & Macris 2010, Bayati & Montanari 20011, R. & Gastpar 2012, Wu & Verdu 2012, Krzakala et al. 2012, Donoho et al. 2013, Huleihel & Merhav 2016
- R. & Pfister 2016 provide rigorous derivation of mutual information and MMSE limits for Gaussian matrices. Proof uses conditional CLT (see Tomorrow's talk)

Characterization of limits via decoupling principle

$$\boldsymbol{Y} = \sqrt{\mathsf{snr}}\, \boldsymbol{H}\boldsymbol{X} + \boldsymbol{W}$$

compressed sensing

$$\widetilde{\boldsymbol{Y}} = \sqrt{s^*}\, \boldsymbol{X} + \widetilde{\boldsymbol{W}}$$

signal plus noise

- Conditional distribution of $\boldsymbol{X}$ given $(\boldsymbol{Y}, \boldsymbol{H})$ is complicated!
- Conditional distribution of small subsets of $\boldsymbol{X}$ given $(\boldsymbol{Y}, \boldsymbol{H})$ characterized by signal plus noise model, i.e. there exists a coupling on $(\boldsymbol{Y}, \boldsymbol{H}, \widetilde{\boldsymbol{Y}})$ such that

$$P_{\boldsymbol{X}_S | \boldsymbol{Y}, \boldsymbol{A}}(\cdot \,|\, \boldsymbol{Y}, \boldsymbol{A}) \approx \prod_{i \in S} P_{X_i | \widetilde{Y}_i}(\cdot \,|\, \widetilde{Y}_i)$$

- Effective SNR given by

$$s^* = \arg\min_s \left\{ I\left(X; \sqrt{s} X + W\right) + \frac{\delta}{2} \left(\log\left(\frac{\delta\, \mathsf{snr}}{s} \right) + \frac{s}{\delta\, \mathsf{snr}} - 1 \right) \right\}$$

Table of Contents

Estimate and compress + decoupling

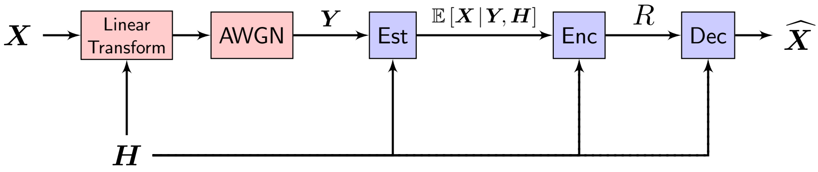

- Idea is to compress conditional expectation using marginal approximation given by signal plus noise model.
- Encoding and decoding do not depend on matrix

Results

Theorem (Achievability via estimate and compress)

For every $\epsilon > 0$, there etxists N large enough and a rate-R quantization scheme such that

$$\mathbb{E}\left[\frac{1}{N}\left\|\boldsymbol{X} - \widehat{\boldsymbol{X}}\right\|^2\right] \le D_{X|\sqrt{s^*}X+W}(R) + \epsilon$$

where s^* is defined by $(P_X, \delta, \mathsf{snr})$.

Theorem (Converse for bounded subsets)

For every $\epsilon > 0$ and fixed subset S, there exists N_0 large enough such that for any $N \ge N_0$ and any quantization scheme using $|S|R$ bits

$$\mathbb{E}\left[\frac{1}{|S|}\left\|\boldsymbol{X}_S - \widehat{\boldsymbol{X}}_S\right\|^2\right] \ge D_{X|\sqrt{s^*}X+W}(R) - \epsilon$$

where s^* is defined by $(P_X, \delta, \mathsf{snr})$.

Bounds described by single-letter DRF

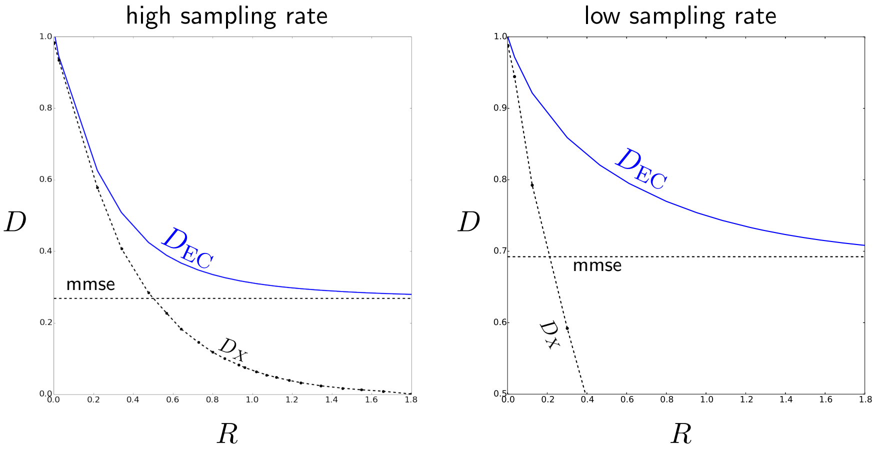

Are we done?

Compress and estimate + decoupling

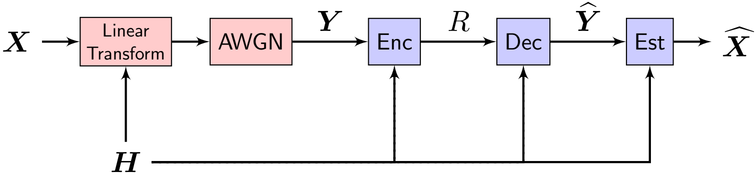

- First compress measurements using Gaussian quantization
- Then estimate signal from reconstructed measurements treating quantization error as additional noise.
- Encoding and decoding do not depend on matrix

Result

Theorem (Achievability via compress-and-estimate)

For every $\epsilon > 0$, there etxists N large enough and a rate-R quantization scheme such that

$$\mathbb{E}\left[\frac{1}{N}\left\|\boldsymbol{X} - \widehat{\boldsymbol{X}}\right\|^2\right] \leq \mathsf{mmse}\left(X|\sqrt{s'}\,X + W\right) + \epsilon$$

where s' is defined by $(P_X, \delta, \mathsf{snr}')$ with

$$\mathsf{snr}' = \mathsf{snr}\frac{1 - 2^{-2R/\delta}}{1 + \mathsf{snr}2^{-2R/\delta}}$$

Comparison of achievability results

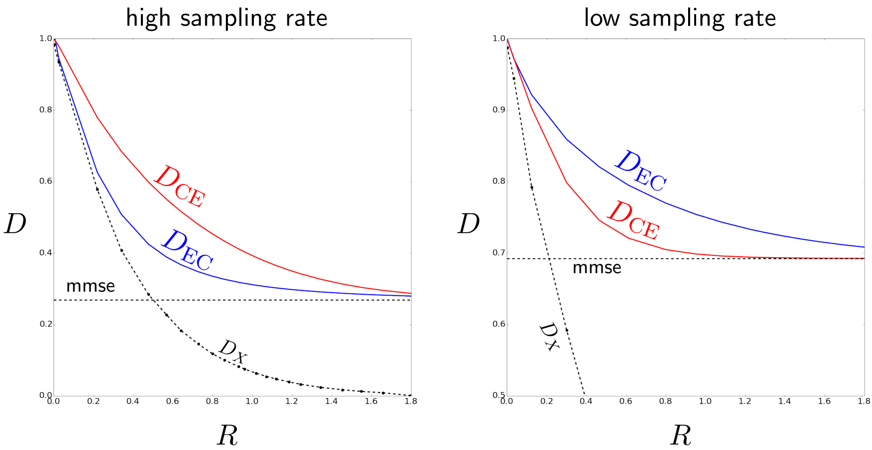

Neither scheme is optimal in general!

Two different quantization schemes

Estimate-and-compress (EC)

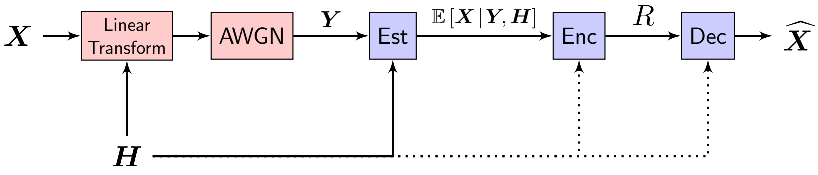

Compress-and-estimate (CE)

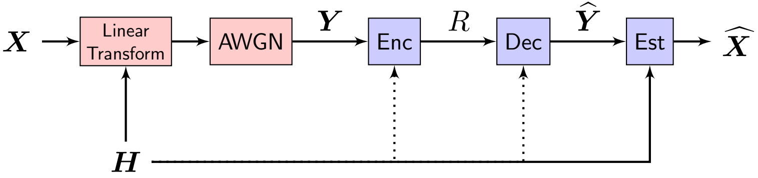

Example: Distortion vs sampling rate

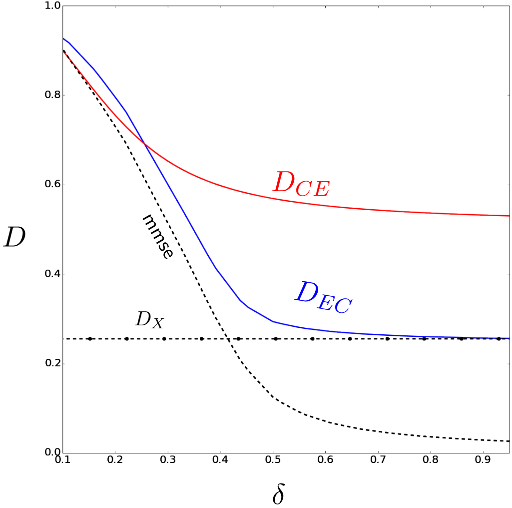

Example: Distortion vs SNR

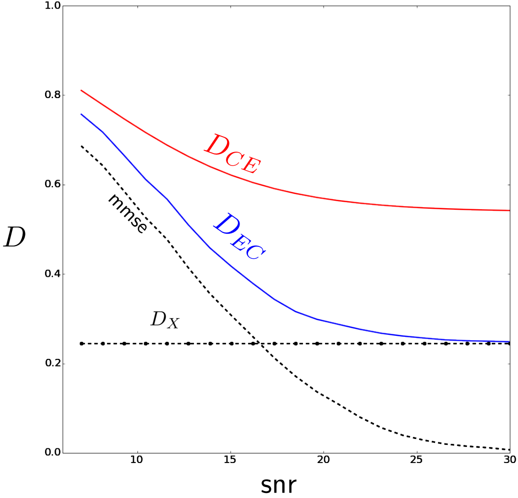

Table of Contents

Summary

- Quantized compressed sensing – tradeoffs between sampling rate, SNR, bitrate, and distortion.
- Conditional distribution of small subsets is described by signal plus noise model. Rigorous results due to R. & Pfister '16.
- Converse for small subsets
- Achievability for Estimate-and-compress
- Achievability for Compress-and-estimate
- Interesting open questions remain...