Cutset Image Sampling for Compression, Reconstruction and Localization

David L. Neuhoff

Univ. of Michigan

CSP Seminar

May 9, 2013

Conventional Image Sampling

sample uniformly at
the points in a lattice

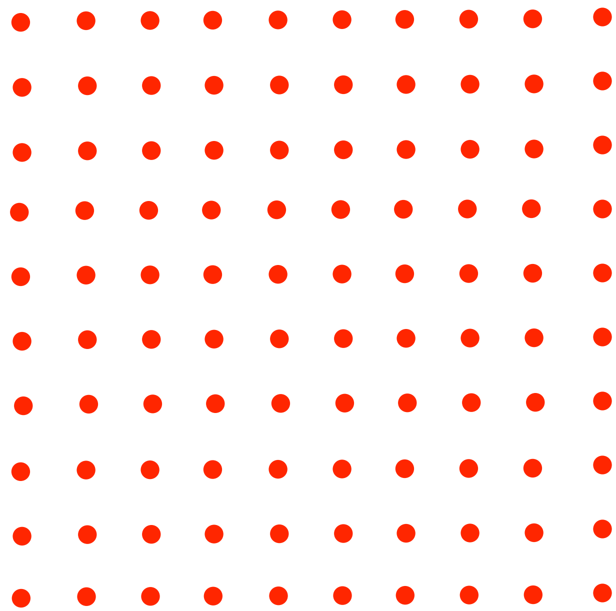

100 samples

Cutset Image Sampling

sample densely
on a grid of lines

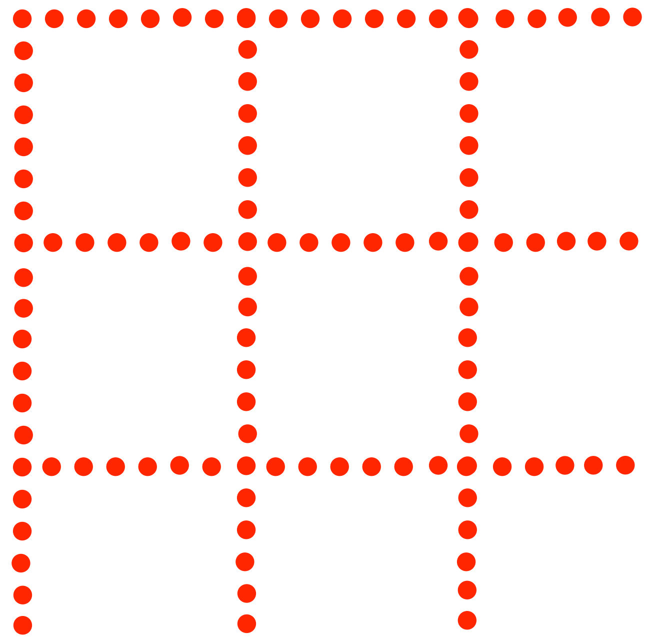

98 samples

Evolution of Cutset Sampling

- Sabbatical at Northwestern
  - My goal: use Markov random field (MRF) models to develop new image compression methods.
  - Result: lossy image compression method for bilevel images based on cutset subsampling and MRF model.
- Later
  - Lossless image encoding for bilevel images based on cutset subsampling, MRF models and arithmetic coding.
  - Cutset sampling as a general approach to image sampling
    - reconstruction algorithms
    - sampling theorems
  - Sensor networks with cutset deployment of sensors
    - Low energy localization algorithms

Collaborators

Thrasos Pappas
Northwestern Univ.
faculty

Matt Reyes
former PhD student

Ashish Farmer
former MS student

Awlok Josan
former PhD student

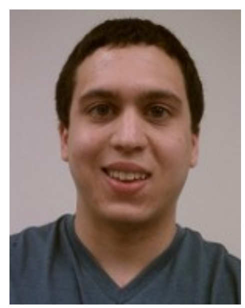

Matt Prelee
current PhD student

Shengxin Zha
current NWU
PhD student

Outline of Talk

- Overview of Markov random fields (MRF)
- Cutsets
- Image Compression based on cutset encoding
  - Overview of image compression
  - Lossy compression of bilevel images based on MRF models
  - Lossless compression of bilevel images based on MRF models an arithmetic coding
- Cutset sampling as a general technique for sampling images
  - Motivation
  - Reconstruction methods
  - Sampling Theorem
- Sensor networks with cutset deployment
  - Low energy localization

Overview of Markov Random Fields

A 1-Dimensional Markov Random Field is a Conventional Markov Chain

- The usual specification:

  X_{n+1} conditionally independent of $X_{n-1}, X_{n-2}, \ldots$ given X_n

- Symmetric specification:

  Given X_n, all random variables before n are conditionally independent of all random variables after n .

  Or: Given $X_n, \ldots, X_{n+m}$, all random variables before n are conditionally independent of all random variables after n+m .

- Graphical representation:

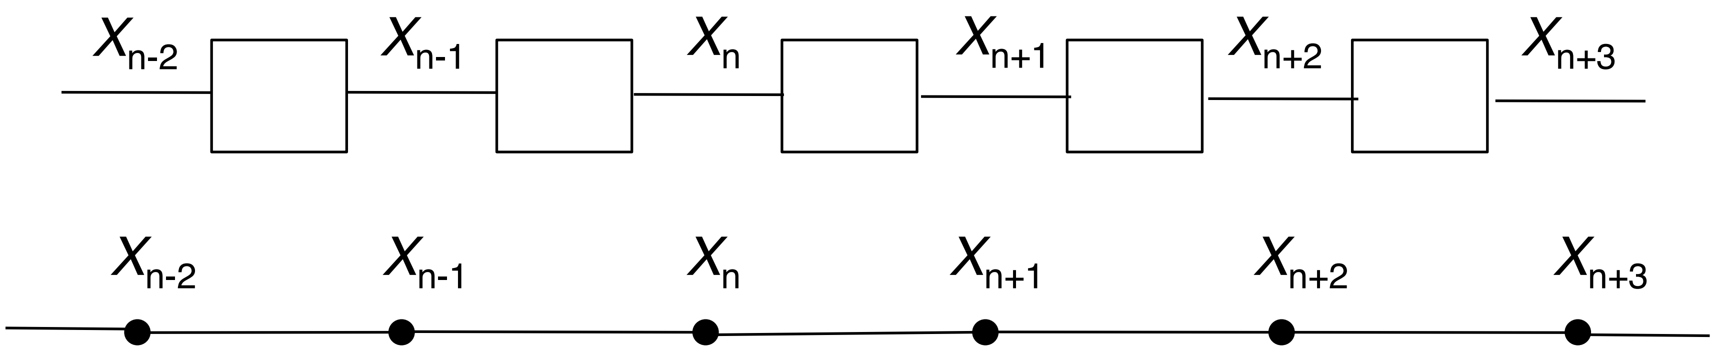

Two-Dimensional Markov Random Field

- Specified in terms of a graph: nodes, undirected edges
- At each node i, there is random variable X_i
- Edge e = (i,j) connecting nodes i and j indicates strong correlation between X_i and X_i
- More precisely, an MRF is defined by any of the following conditional independence properties
  - No edge connects i and j iff X_i and X_j are conditionally independent given all other *X's*
  - If set of nodes **R** surrounds set of nodes **G**, then given X_R, X_G is conditionally independent of X_B.
  - If set of nodes **R** *separates* set of nodes **G** from set of nodes **B**,   then given X_R, X_G is conditionally independent of X_B.

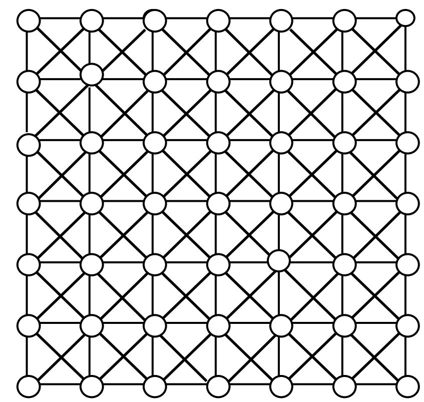

8-way graph

Two-Dimensional Markov Random Field

- Specified in terms of a graph: nodes, undirected edges
- At each node i, there is random variable X_i
- Edge e = (i,j) connecting nodes i and j indicates strong correlation between X_i and X_i
- More precisely, an MRF is defined by any of the following conditional independence properties
  – No edge connects i and j iff X_i and X_j are conditionally independent given all other *X's*
  – If set of nodes **R** surrounds set of nodes **G**, then given X_R, X_G is conditionally independent of X_B.
  – If set of nodes **R** *separates* set of nodes **G** from set of nodes **B**,   then given X_R, X_G is conditionally independent of X_B.

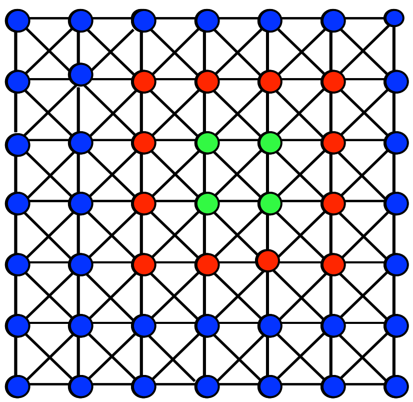

8-way graph

Two-Dimensional Markov Random Field

- Specified in terms of a graph: nodes, undirected edges
- At each node i, there is random variable X_i
- Edge e = (i,j) connecting nodes i and j indicates strong correlation between X_i and X_i
- More precisely, an MRF is defined by any of the following conditional independence properties
  - No edge connects i and j iff X_i and X_j are conditionally independent given all other *X's*
  - If set of nodes **R** surrounds set of nodes **G**, then given X_R, X_G is conditionally independent of X_B.
  - If set of nodes **R** *separates* set of nodes **G** from set of nodes **B**, then given X_R, X_G is conditionally independent of X_B.

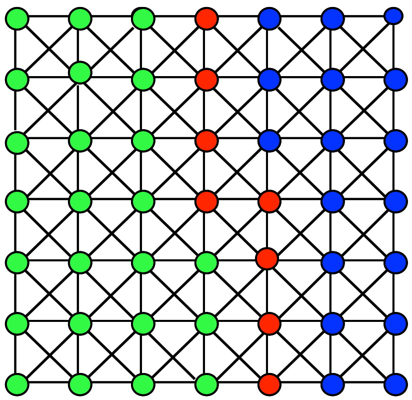

8-way graph

- Hammersley-Clifford Theorem: X is MRF, iff $p(x)$ factors

$$p(x) = \prod_{\text{edges } e=(i,j)} \psi(x_i, x_j) \prod_{\text{nodes } i} \phi(x_i)$$

where $\psi(x,y)$ is an *edge potential function,* $\phi(x)$ is a *node potential function*

Ising MRF Model for Bilevel Image

- $X = [X_1, \ldots, X_N]$, ith pixel X_i = 0 or 1
  wrt some ordering of pixels

$$p(x) = \prod_{\text{edges } e=(i,j)} \psi(x_i, x_j)$$

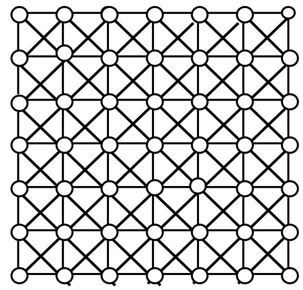

8-way graph

where

$$\psi(x_i, x_j) = \begin{cases} e^{\beta}, & x_i = x_j \\ e^{-\beta}, & x_i \neq x_j \end{cases}$$

- Equivalently,

$$p(x) = \frac{1}{Z} e^{-2\beta\, t(x)}$$

where $t(x)$ = number adjacent pairs of pixels that disagree, i.e., black-white transitions, called *odd bonds*

- x has high prob. iff $t(x)$ small
  i.e., few odd bonds

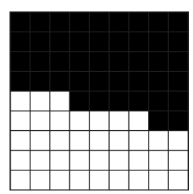

high prob

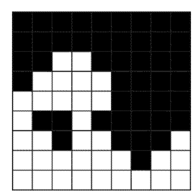

low prob

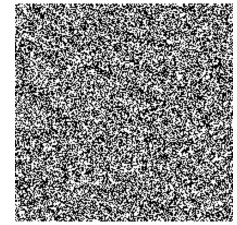

$\beta = 0.0$

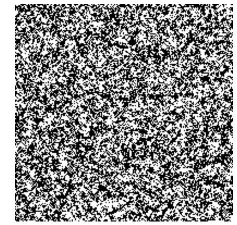

$\beta = 0.1$

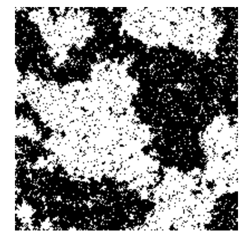

$\beta = 0.2$

$\beta = 0.4$

an actual image

Definition of Cutset

- A set of nodes/pixels that partitions all other pixels into groups such that one cannot go from one group to another without going through cutset pixels.
- Given the cutset pixels, the groups are conditionally independent of one other.
- Prime example: Manhattan grid cutset

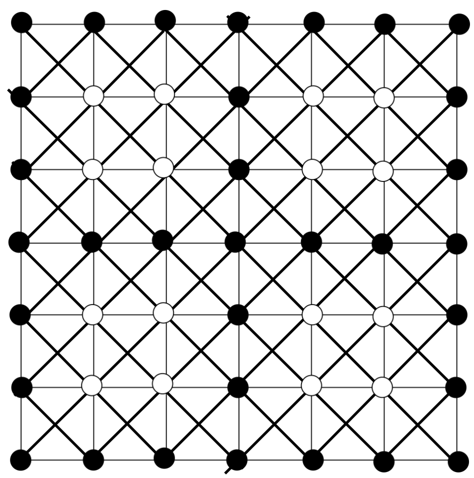

Overview of Image Compression

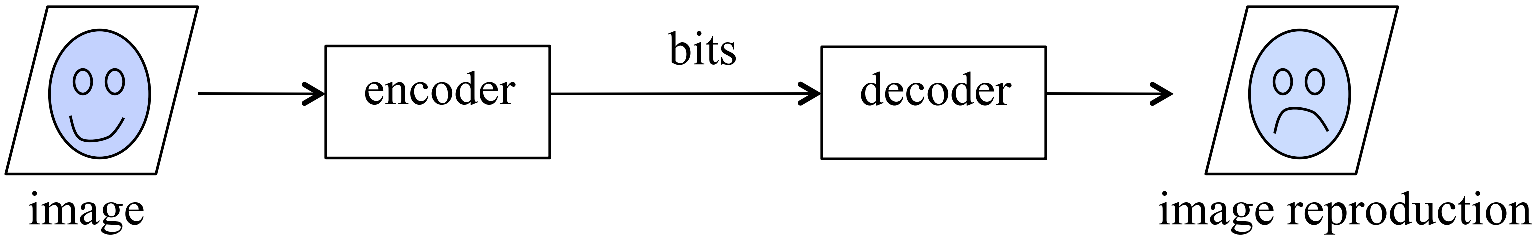

- Image compression system = encoder + decoder
- Performance
  - Coding/compression rate: R = # bits/pixel
  - Reproduction quality: D = distortion in reproduction, e.g. MSE, SNR, PSNR
- Lossless and Lossy Image Coding
- Theory for predicting performance and performance limits
  - Lossless coding -- entropy theory
  - Lossy coding – (a) Shannon rate-distortion theory,
    (b) high-resolution theory

Shannon information theory

Lossy Cutset Coding (LCC) of Bilevel Images

- Encoder: losslessly encodes Manhattan-grid cutset (that's it!)
  - for example, with Arithmetic Coding
  - typically 0.1 to 0.3 bits per cutset which (enabled by closeness of cutset pixels)
- Decoder: reconstructs/ estimates/interpolates block interiors using MRF model and MAP rule

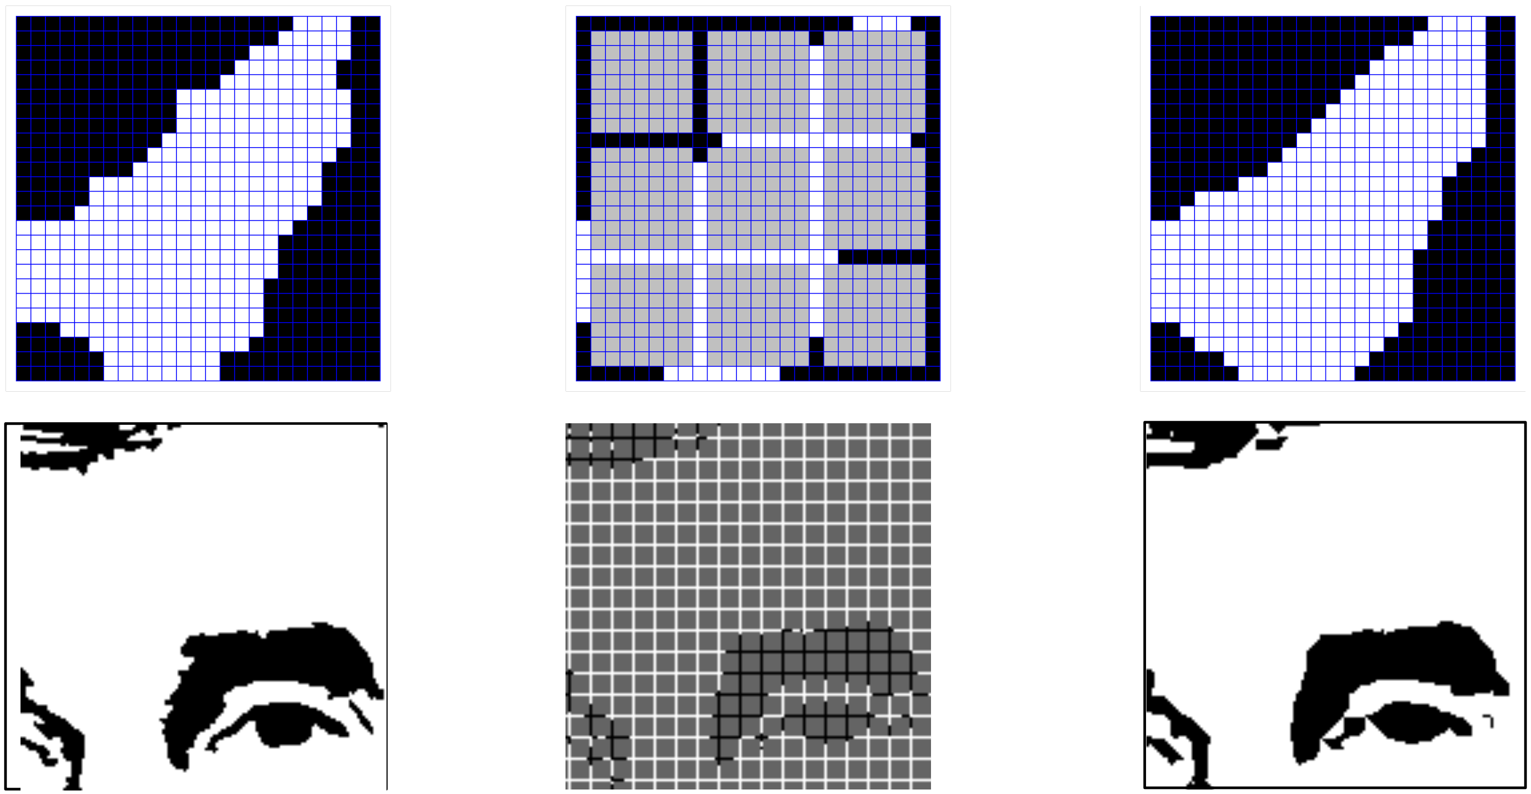

[Reyes et al., ICIP 2007; Reyes, UM PhD Dissertation, 2011]

Decoding: MAP Reconstruction for MRF Model

- Decoding "game": given boundary of a block, find the interior with fewest black-to-white transitions (odd bonds).
- Can be found by iterative algorithm such as loopy belief propagation.
- Can be found analytically for the three most common types of block boundaries.

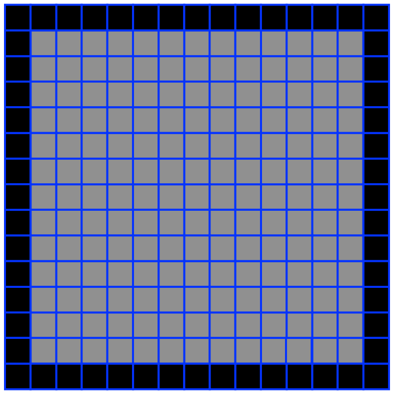

monotone

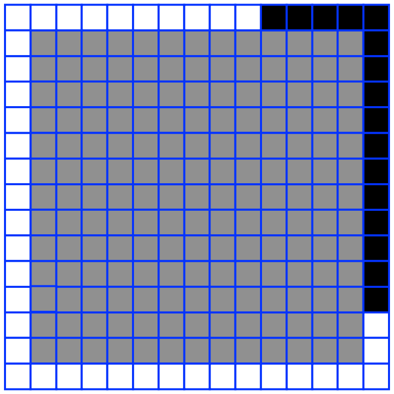

one-run

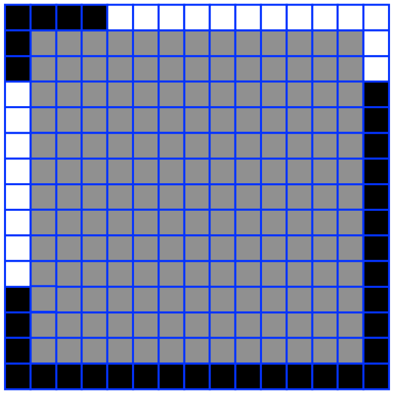

two-run

Key Property of MAP reconstructions

- There is monotone HV path from each interior pixel to a boundary pixel of same color.

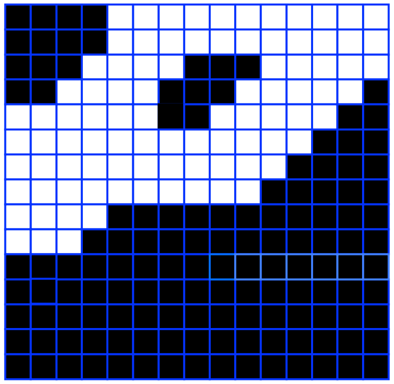

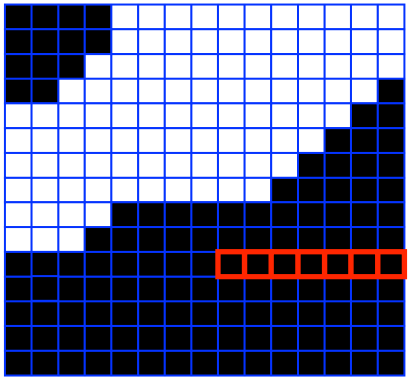

not MAP

could be MAP

- Equivalently, no "islands" in the interior, i.e., every monotone loop must be filled with same color

- Proof:

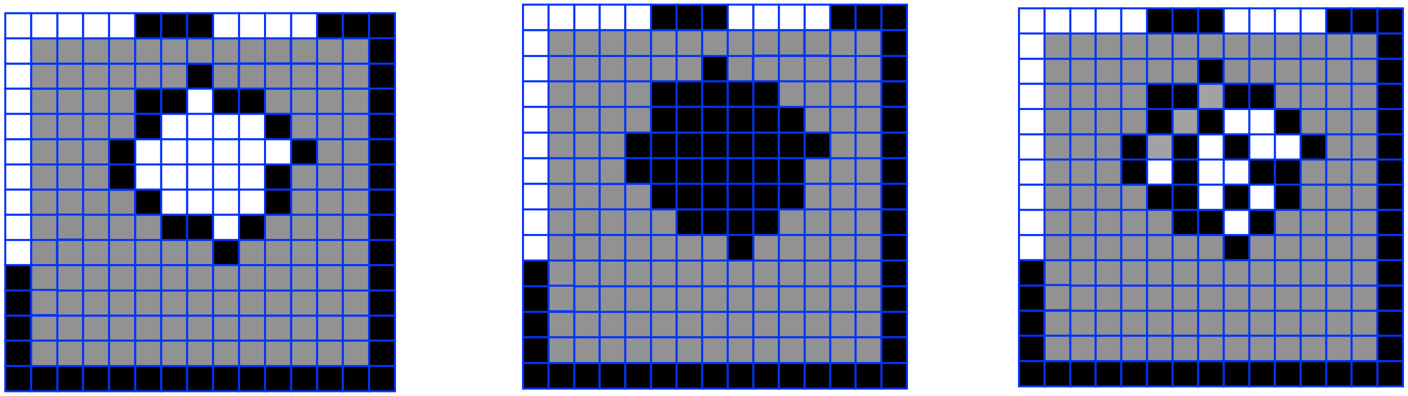

Block MAP Reconstructions for Common Boundaries

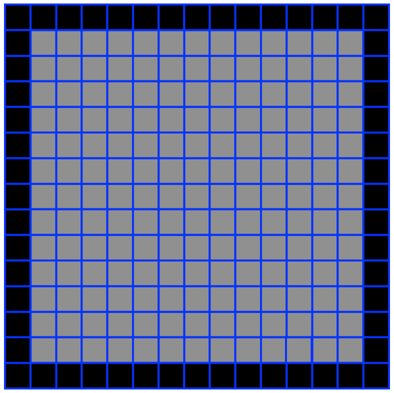

monotone

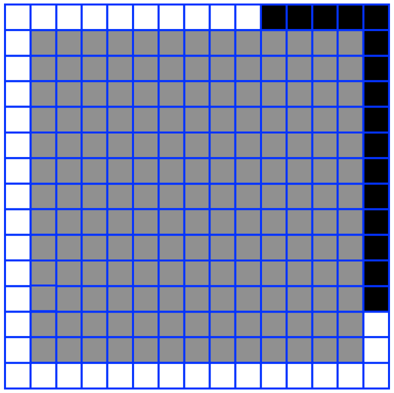

one-run

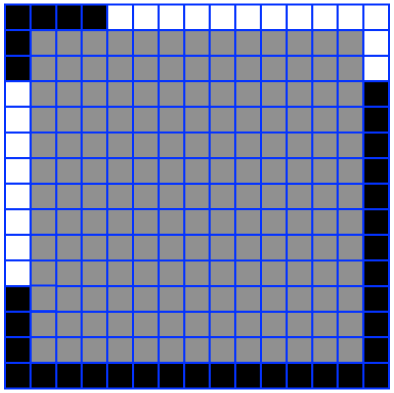

two-run

Monotone Boundary

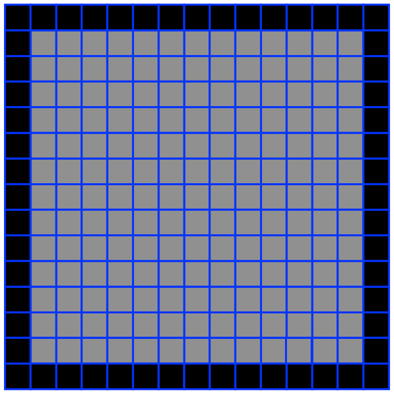

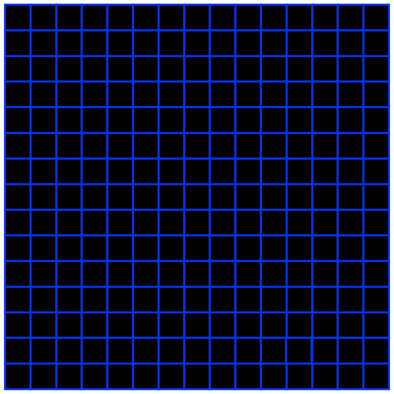

One Run Boundary: Reconstruction Paths

- Since it can have no islands, any MAP reconstruction of one run boundary is determined by a *reconstruction path*, either black or white.

- Therefore, wlog w restrict attention to reconstructions defined by *reconstruction paths,* one for each run.

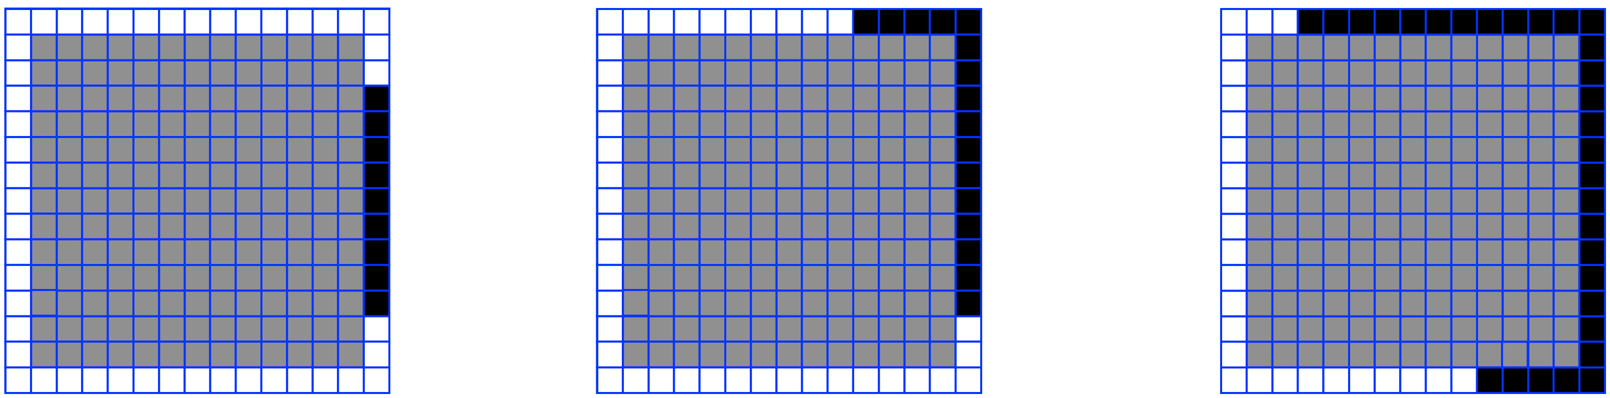

Boundary with One Black Run: MAP paths

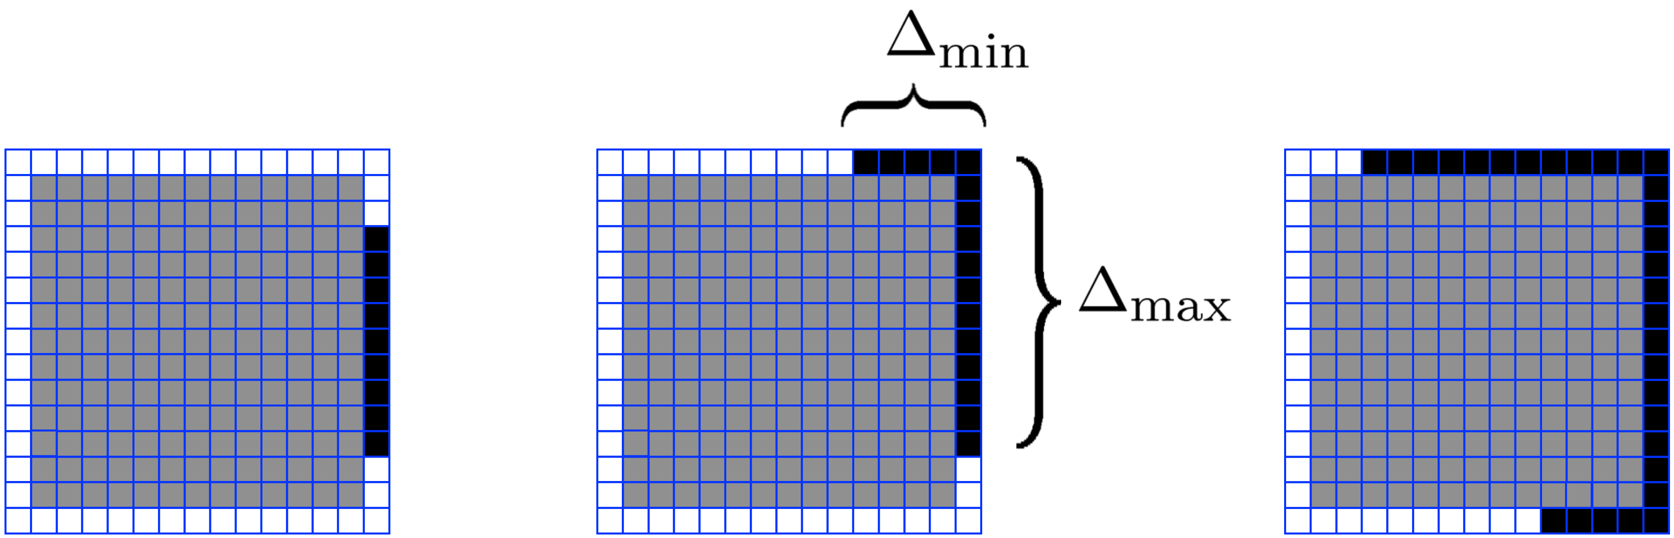

- **Theorem:** For 1-run boundary containing *c* corners and having major and minor differences $\Delta_{\max}$ and $\Delta_{\min}$

(1) all MAP paths have $\Delta_{\max}$ edges

(2) $c \leq 2$: $\binom{\Delta_{\max}}{\Delta_{\min}}$ MAP paths; all *simple*; $3\,\Delta_{\max} + \Delta_{\min} + 5 - 2c$ odd bonds

(3) $c = 3$: $\binom{\Delta_{\max}-1}{\Delta_{\min}-1}$ MAP paths; all *simple*; $3\,\Delta_{\max} + \Delta_{\min} - 1$ odd bonds

(4) $c = 4$: 1 MAP path; *not simple*; $3\,\Delta_{\max} - 1$ odd bonds

Two-Run Boundary

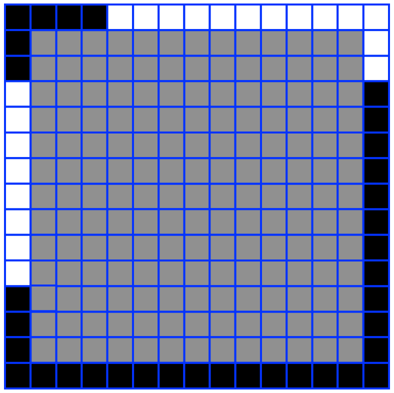

- First thought: merge two 1-run MAP reconstructions
- Not necessarily MAP: deviating from 1-run optimality to make1-run recon's touch, can decrease odd bonds
- Instead find &compare best of 3 types of reconstructions:
  - white HV-connected
  - black HV-connected
  - bi-connected: white and black connected, but neither is HV-connected
- If white- or black HV-connected, merge two 1-run opt. reconstructions
- In bi-connected, "widget" reduces # number of odd bonds

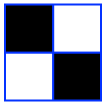

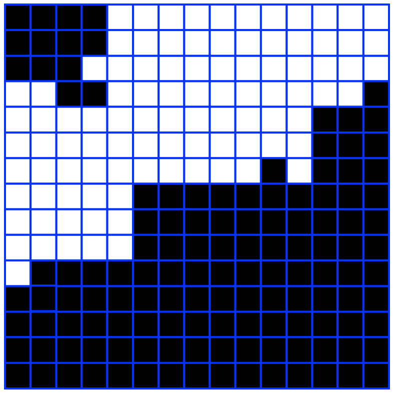

white HV-connected

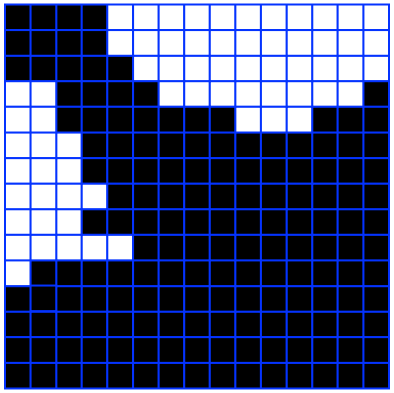

black HV-connected

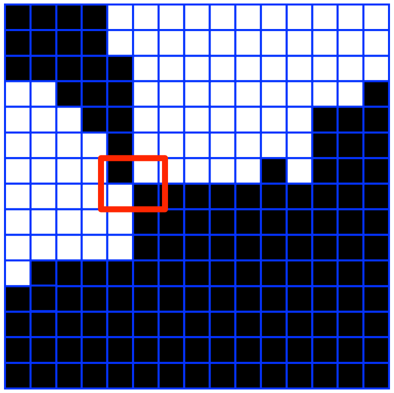

bi-connected

Widget Theorems

- A MAP reconstruction for any boundary can have no widgets in its interior.
- A MAP reconstruction for a two-run boundary can have at most one widget on the boundary

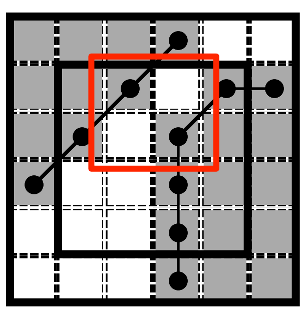

not MAP

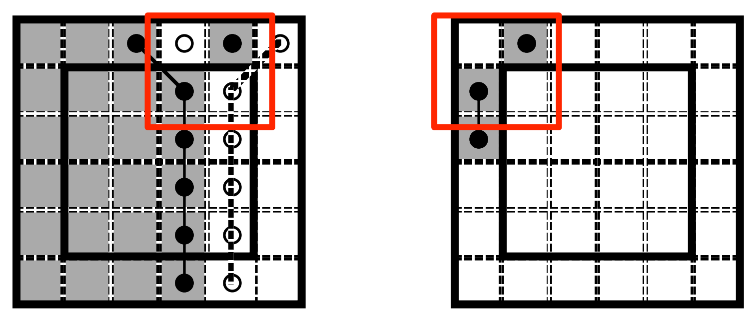

MAP

2-Run MAP Reconstructions

- **Theorem:** Consider block with two boundary runs
  - If one run contains four corners, the only MAP recon is entirely this color.

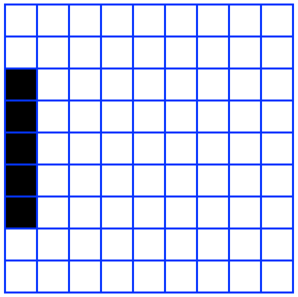

  - If not, and boundary is one of two types shown to right, MAP is bi-connected as shown.

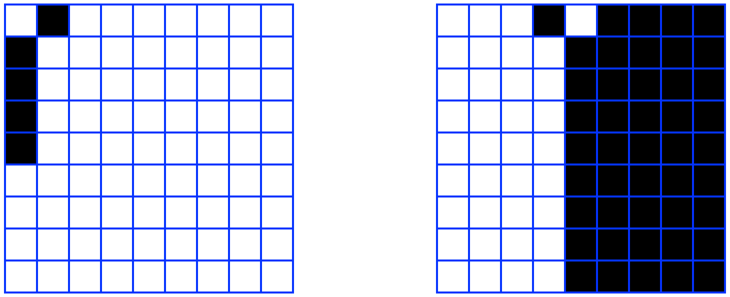

  - If not, MAP is white HV-connected formed by any pair of simple non-touching black reconstruction paths,

    or black HV-connected formed by any pair of simple non-touching white reconstruction paths,

    according to which has fewer odd bonds

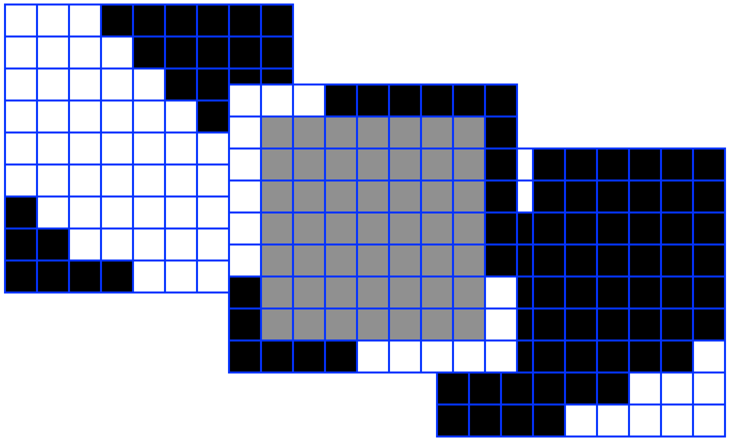

Boundary with More Than Two Runs

- Belief propagation, or
- Simple ad hoc decoding rule
  - Find the color with the two longest runs
  - Change all other pixels to the other color
  - Apply the two-run solution

Sample Decoded Result

original

LCC

blocksize 8

R = 0.05 bpp

D = 3.5%

Decision-bit Coding

- Problem: When thin line passes through block, boundary has 2 runs, and MAP rule sometimes produces white-HV-connected reconstruction, when black-HV-connected is better, and vice versa.

- Fix: Whenever block boundary has two runs, “decision bit coding” tests to see which connection is best, and sends extra bit to tell decoder.

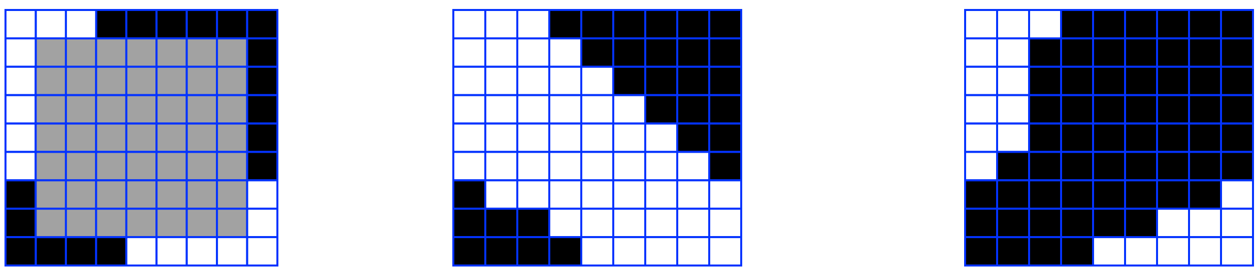

Decision-bit Coding

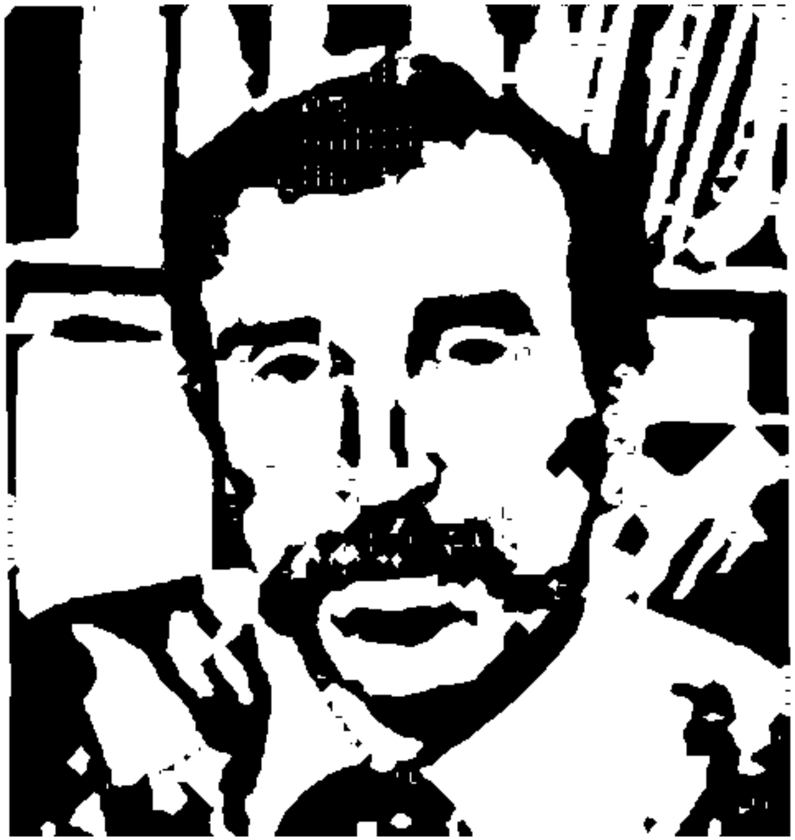

blocksize 8
w/o decision bit

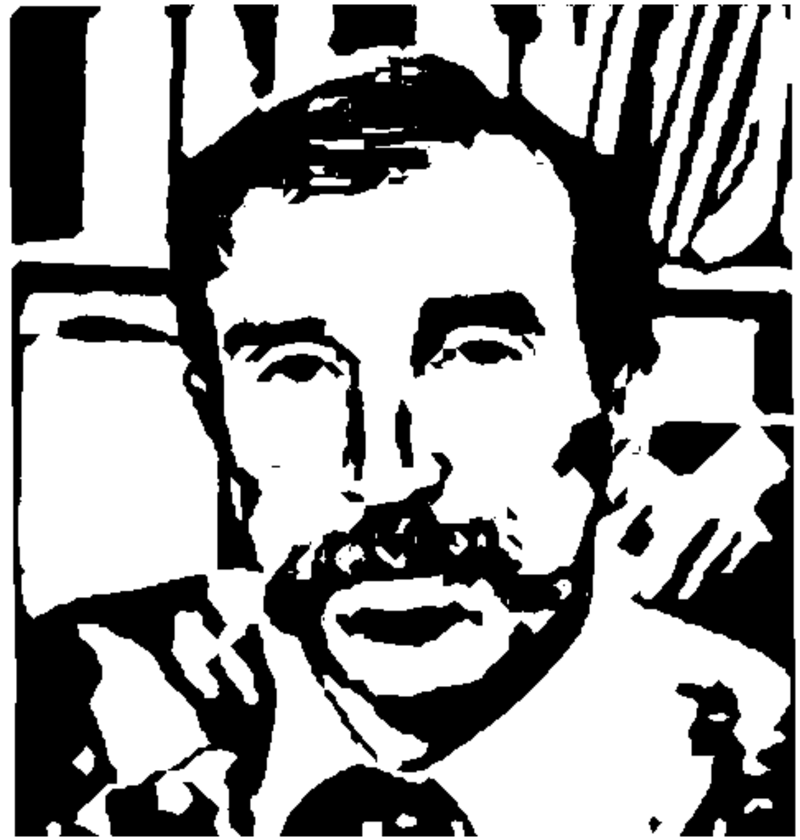

blocksize 8
with decision bit

Rate-Distortion Performance

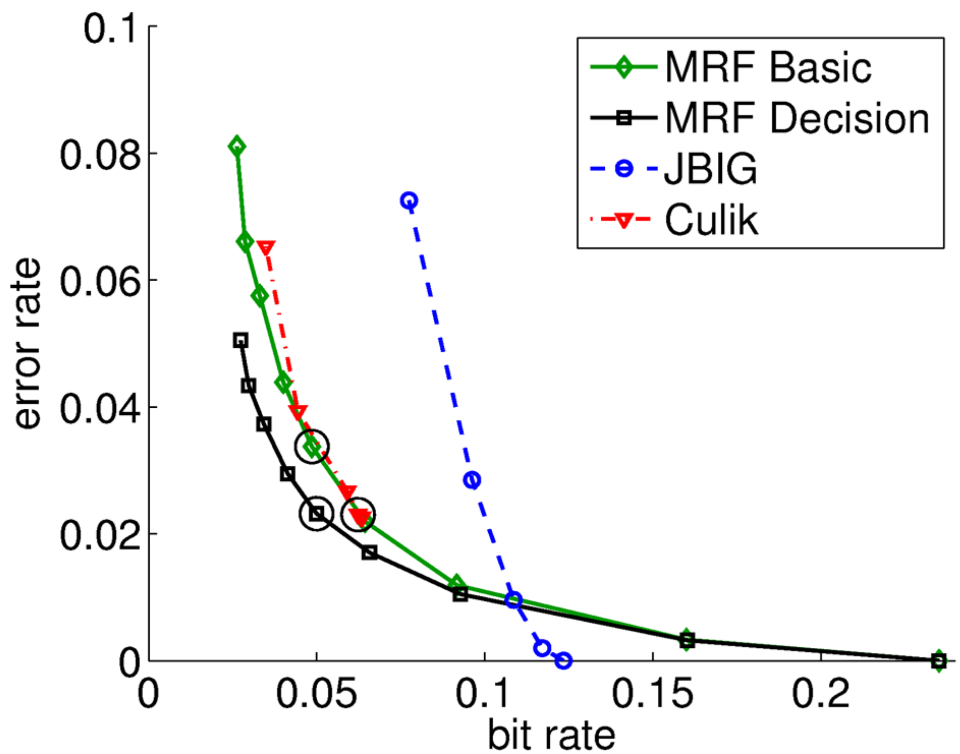

- Compare to Culik and Valenta '97, and nonlinear filter+JBIG
- For *scenic bilevel images* (complex, but not text or halftone) LCC is best method of which we are aware
- Percent error is less than ideal as a distortion measure.
- LCC coded images "look" much better than C&V coded images

Another image

$B = 6$ $B = 10$

basic

decision bit

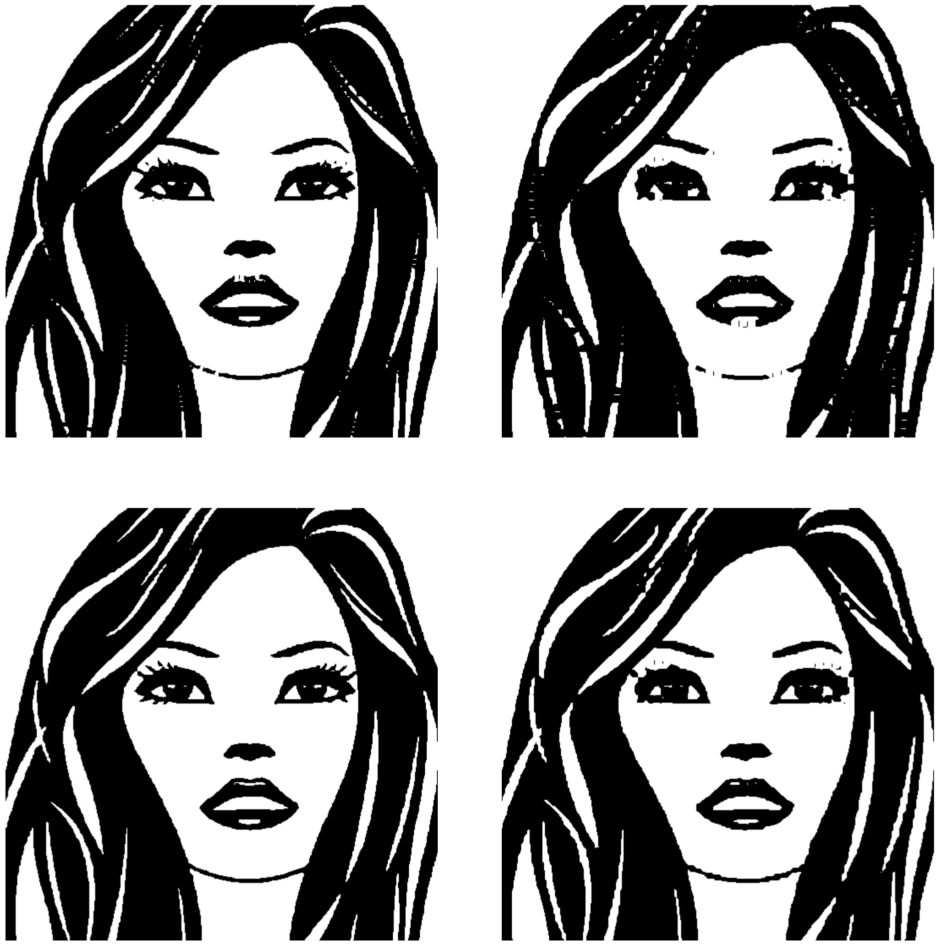

$R = 0.053$ bpp $R = 0.035$ bpp

$D = 1.3$ % $D = 1.8$ %

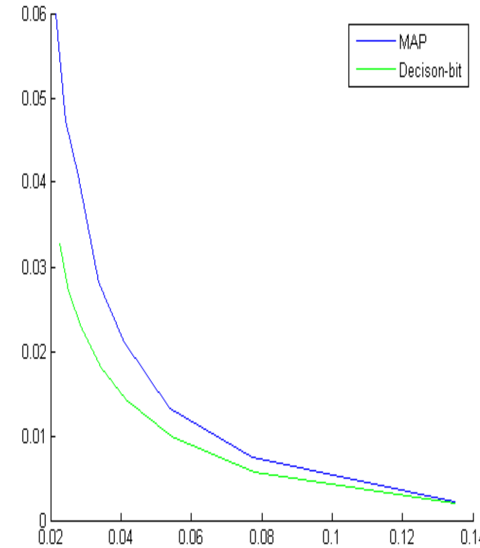

Cutset-First Lossless Image Compression

- Step 1: Encode cutset. Use arithmetic coding (AC) with MRF model guiding coding distributions.
- Step 2: Encode remaining pixels. Use AC to conditionally encode remainder given cutset pixels, again with MRF model guiding coding distributions.
- Choose cutset such that it is feasible to use Belief Propagation (BP) to compute:
  a. approximately optimal "reduced" coding distributions for cutset
  b. opt. conditional coding distributions for remainder

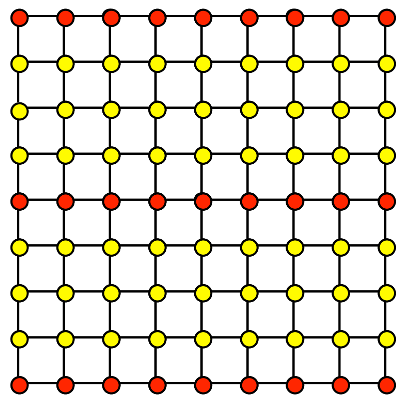

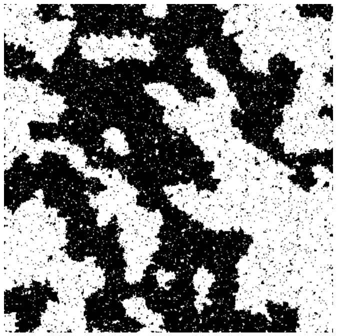

bilevel MRF model, $\beta = 0.5$

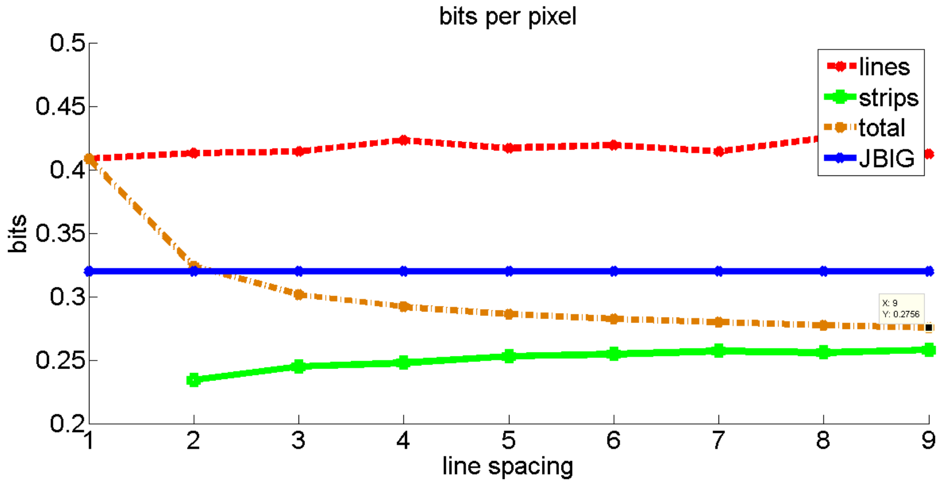

[Reyes, DN, DCC 2010; Reyes, PhD Dissertation, U. Mich., 2011]

Lossless Image Coding with Arithmetic Codes (AC)

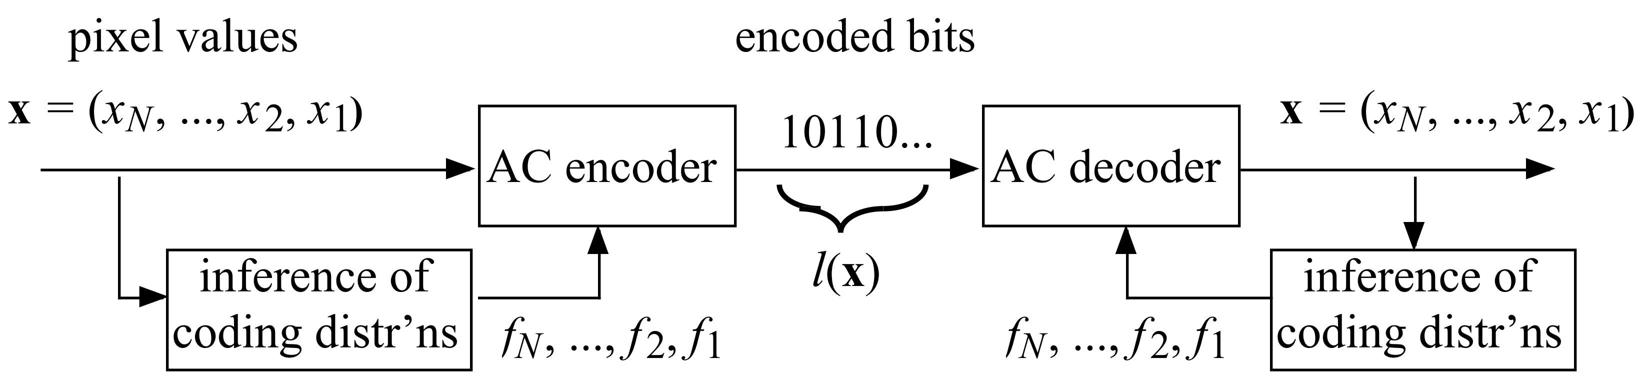

- Encode grid in a 1-D scan order x = $(x_1, \ldots, x_N)$ such that for each pixel except 1st, a horizontally or vertically adjacent pixel, called its *context,* is scanned first.
- Accompany pixel x_i with *coding prob distribution* $\{f_i(0), f_i(1)\}$ (not just $f_i(x_i)$)
- $f_i(0)$ = fraction of previous pixels that are zero and whose context is same as pixel as context of pixel i.
- # bits produced by AC encoder produces:

$$l(\mathrm{x}) \cong \sum_{i=1}^{N} -\log(f_i(x_i)) \cong H(X_2 \mid X_1)$$

Cutset Sampling as General Approach to Sampling Grayscale Images

- Motivation
  - Physical constraints such as sampling from vehicles
  - Potentially better edge preservation
  - Sensor networks – less wire, less energy

- Goals
  - Reconstruct image from cutset samples
  - Identify images that can be perfectly reconstructed, i.e., find a sampling theorem
  - Develop sensor network signal processing algorithms that benefit from cutset deployment
    - source localization, tracking, …

Motivation: Physical Constraint --

Sampling from a Boat, Airplane or Vehicle

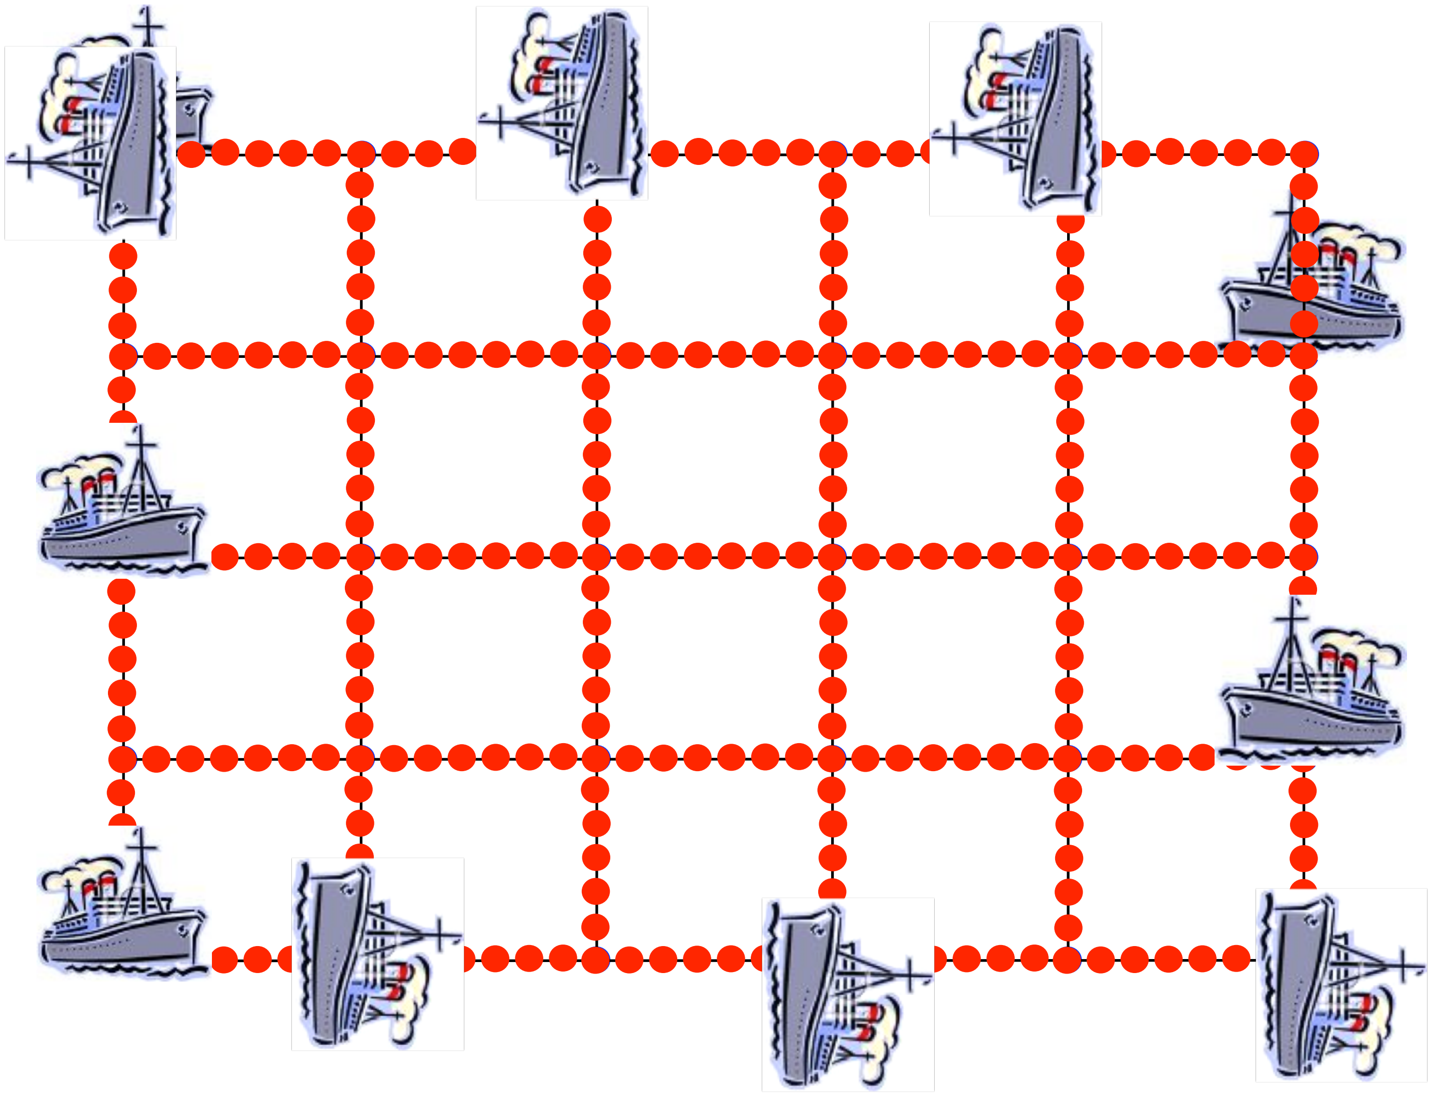

Motivation: Potentially Better Edge Preservation

- Cutset sampling

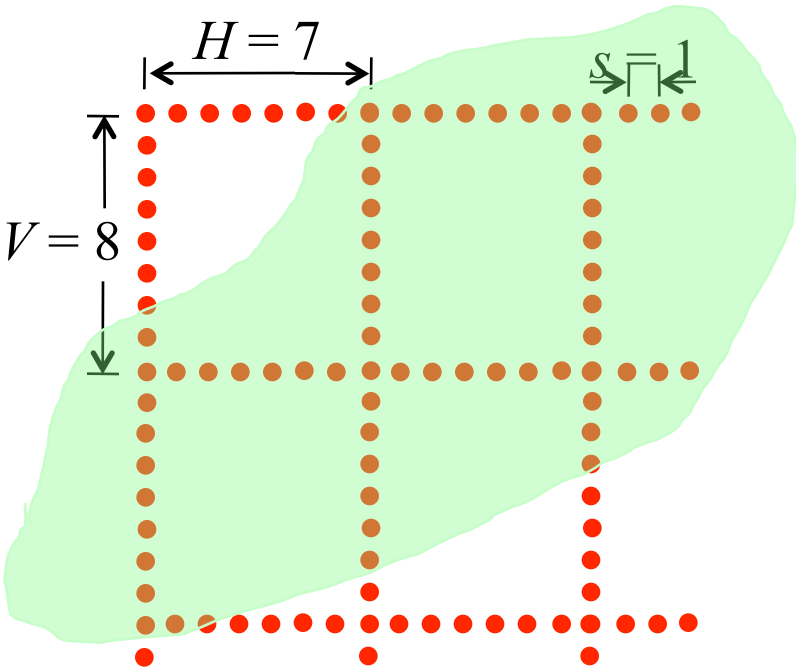

- Conventional sampling

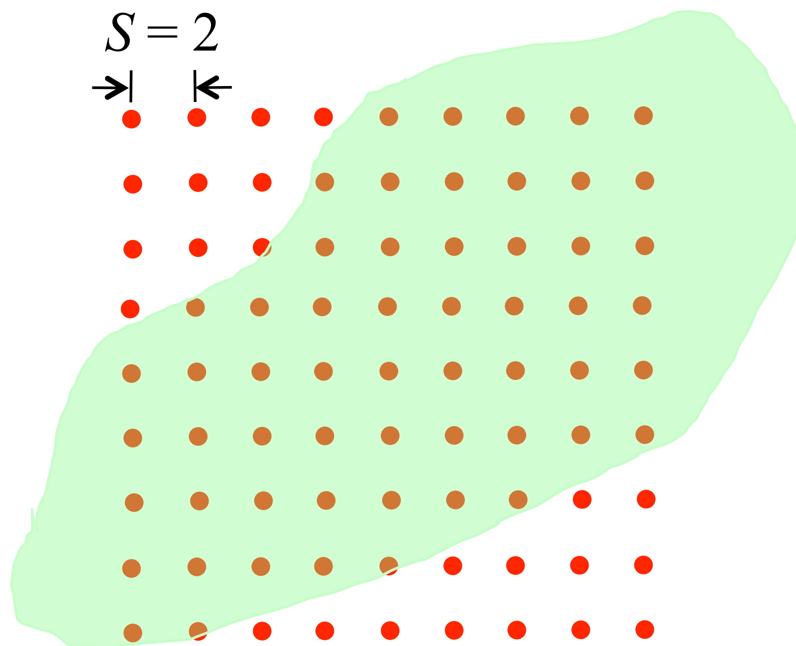

- Samp. density $= \dfrac{(H+V)/s - 1}{HV} = \dfrac{1}{4}$
- Samp. density $= \dfrac{1}{S^2} = \dfrac{1}{4}$
- With same sample density, cutset sensors are closer to each other, with potential to capture edges more accurately.

Motivation: Sensor Networks -- Two Deployments with same density

- Cutset sampling

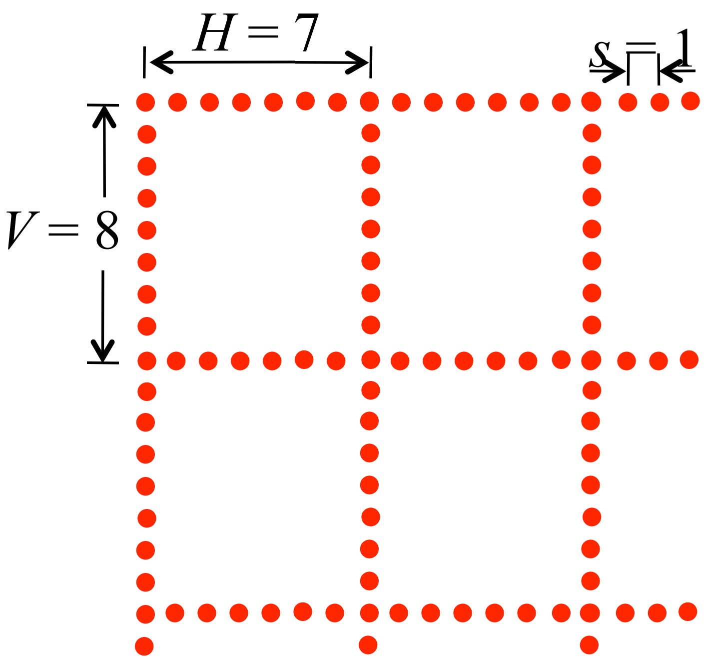

- Conventional sampling

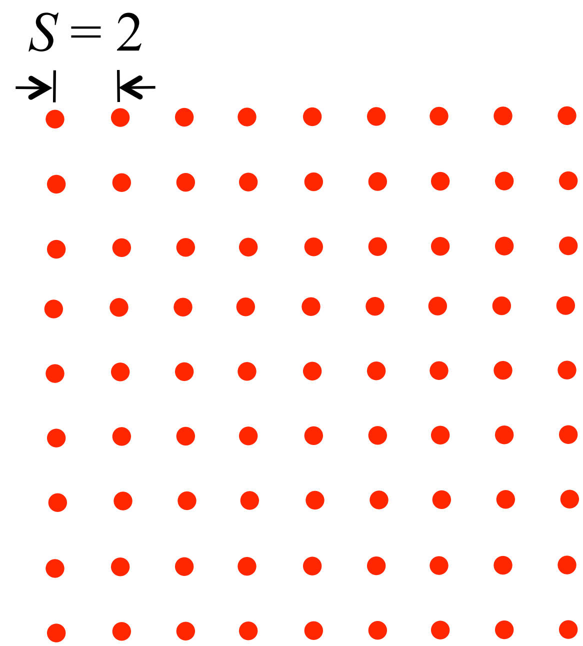

- Samp. density = $\dfrac{(H+V)/s - 1}{HV} = \dfrac{1}{4}$
- Samp. density = $\dfrac{1}{S^2} = \dfrac{1}{4}$
- If $H = V = K\,s$, cutset sampling saves factor of :
  - ○ $\sqrt{2K-1}/2$ in wire for wired network
  - ○ $(K/\sqrt{2K-1})^r$ in energy for wireless network, where $r = 2$ to 4.

Image Reconstruction from Cutset Samples

Discrete Space Image Setting

- Cutset sampling

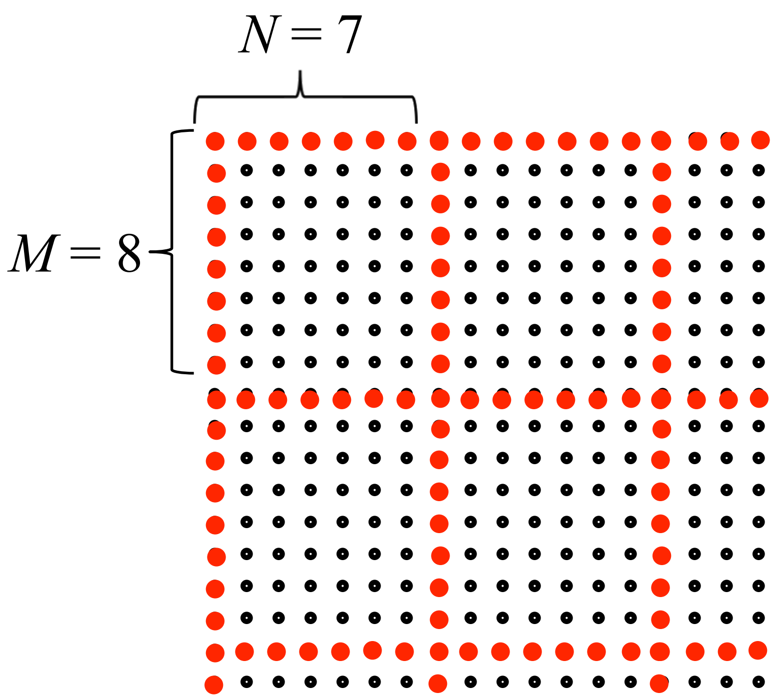

- Conventional sampling

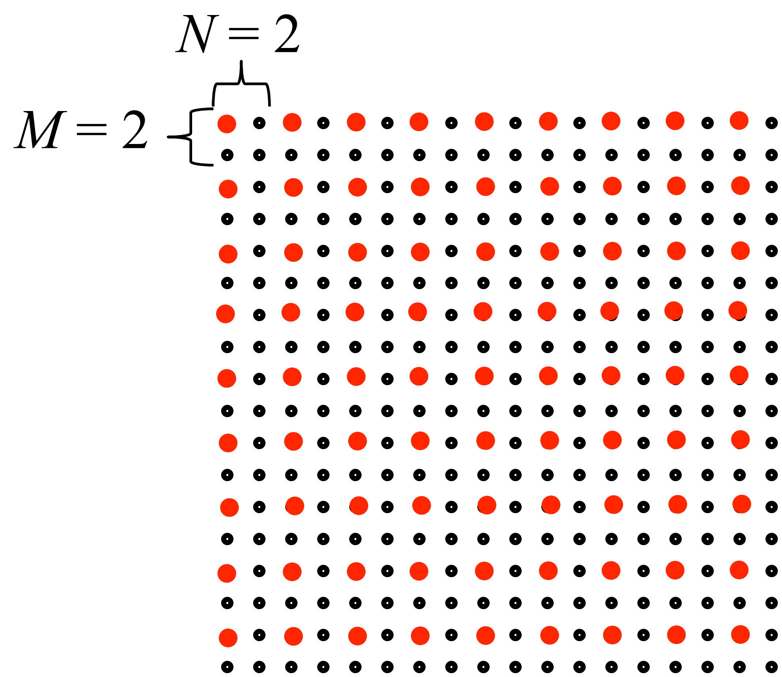

- Density = $\dfrac{M+N-1}{MN}=\dfrac{1}{4}$

- Density = $\dfrac{1}{MN}=\dfrac{1}{4}$

Image Reconstruction

[Farmer et al., ICIP 2011]

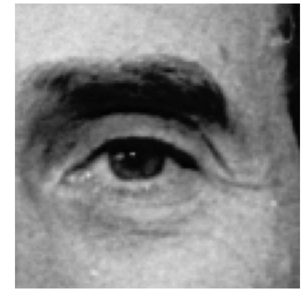

Three-Step Segment-Based Reconstruction

1. Segment cutset
   - Based only on cutset values
   - Criteria: no edges within a segment

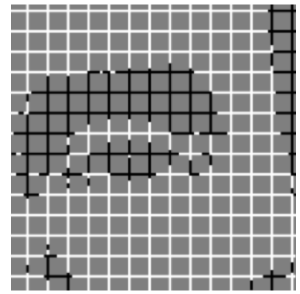

2. Segment block interiors
   - I.e. estimate segmentation of interior
   - Based only on segmented cutset

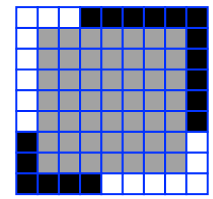

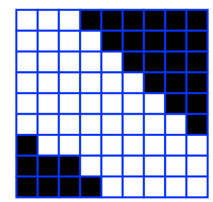

3. Segment-based gray-level interpolation of block interiors
   - Based only on cutset values and segmented blocks and cutset
   - Each pixel is interpolated based only on cutset pixels in same segment

1. Segment Cutset

- We adapt ACA segmentation to cutset segmentation [Pappas `92]

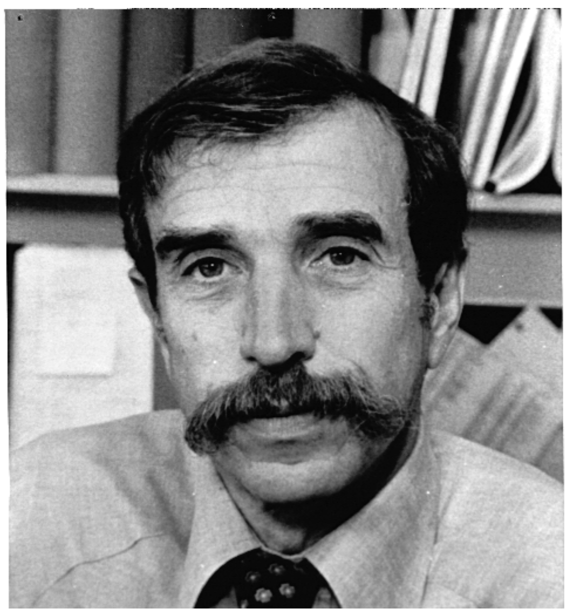

original

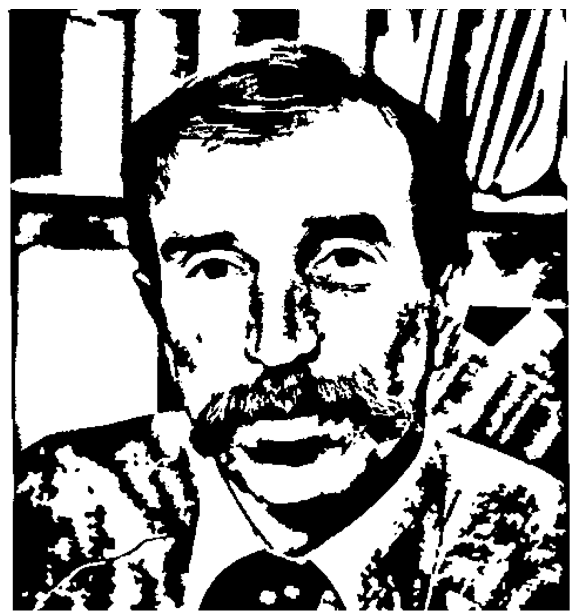

full ACA segmentation

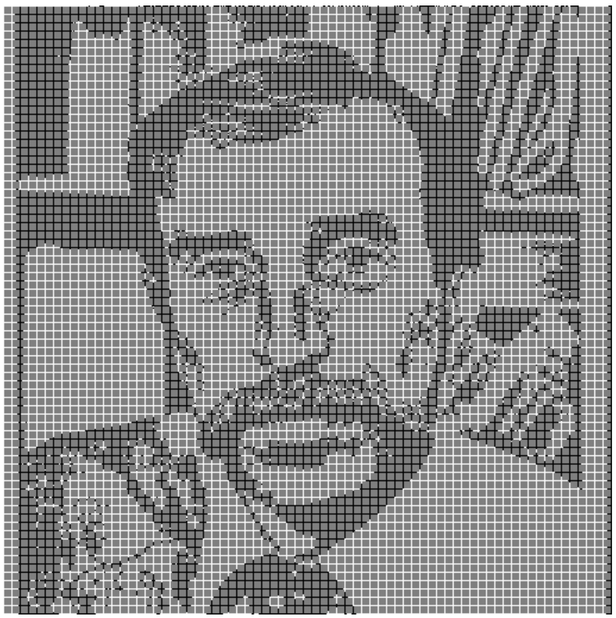

ACA segmentation of 7x7 cutset

2. Segment Block Interiors

- Approach
  - Model ‘ideal segmentation’ of a block as a bilevel MRF.
  - Produce MAP estimate of block interior from segmentation of boundary, using LCC decoder reconstruction algorithm

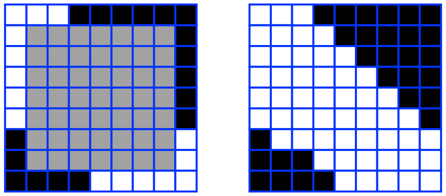

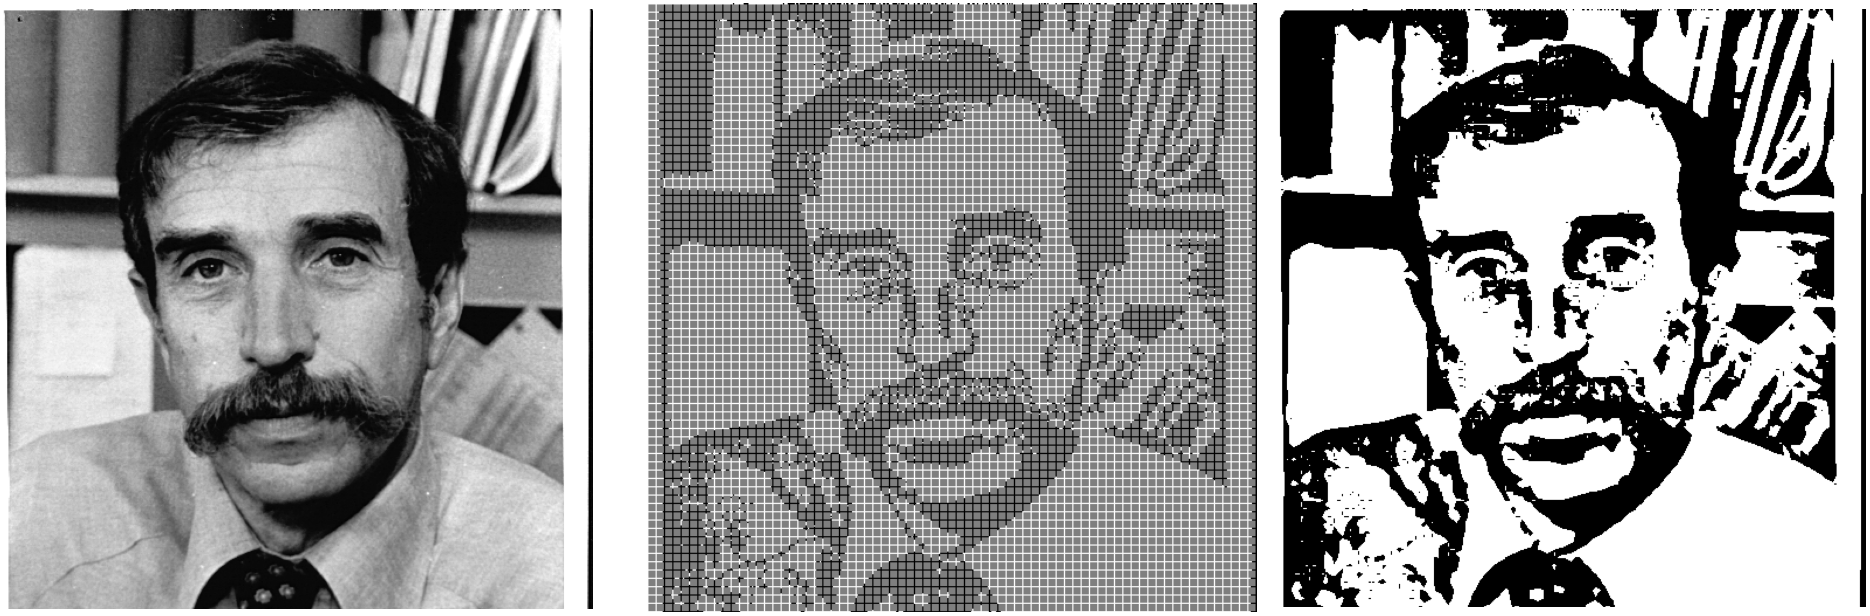

original ACA cutset segmentation cutset + interior segmentation

3. Segment-Based Block Interior Reconstruction

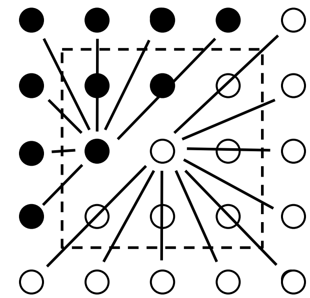

- Key idea: estimate each pixel x_i in block based only on block boundary pixels y_j in the same segment

- MMSE linear estimation:

$$\bar{x}_i = \mu_i + \sum_j a_{i,j} \, (y_j - \mu_i)$$

where

μ_i = mean of segment containing pixel *i*,
estimated as emp. mean of pixels in segment boundary

$\mathbf{A} = [a_{i,j}]$, $\quad \mathbf{A} = K_Y^{-1} K_{YX_i}$

- Assume Gaussian MRF model

Gaussian MRF Model

- $p(x) = \frac{1}{Z} \prod_i \Phi_i(x_i) \prod_{i,j} \Psi_{i,j}(x_i, x_j)$

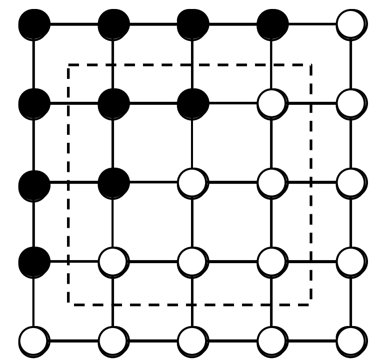

- Node and edge potential functions

$$\Phi_i(x_i) = \exp\left\{ -\tfrac{1}{2} d (x_i - \mu_i)^2 \right\},$$

$$\Psi_{i,j}(x_i, x_j) = \begin{cases} \exp\left\{ -cd(x_i - \mu_i)(x_j - \mu_j) \right\} & \text{if } (i,j) \text{ is edge, \& } i,j \text{ in same seg.} \\ 0, & \text{else} \end{cases}$$

μ_i = seg. mean, est. as emp. mean of pixels in segment $\cap$ boundary

c, d chosen so inverse covariance matrix K^{-1} is positive definite

$$K_{ij}^{-1} = \begin{cases} d, & i = j \\ -cd, & i \neq j,\ (i,j) \text{ an edge} \\ 0, & \text{else} \end{cases}$$

- Either: $\bar{x}_i = \mu_i + \sum_j a_{i,j} (y_j - \mu_i)$ and $\mathbf{A} = K_Y^{-1} K_{YX_i}$
- Or run loopy Belief Propagation on graph.

Example: Sampling Density 1/4

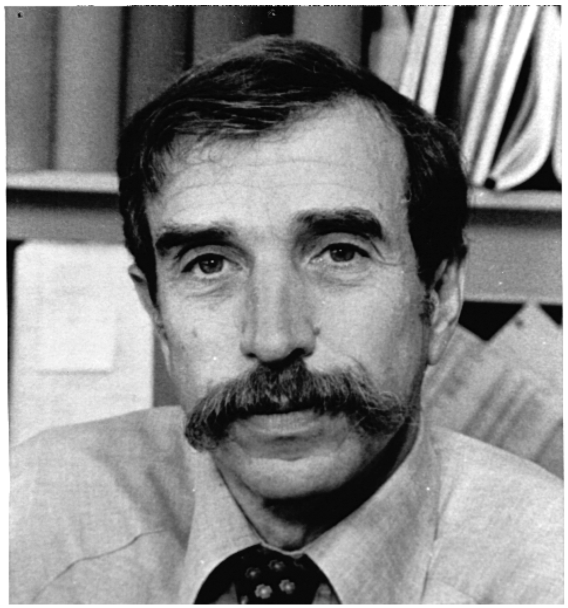

original

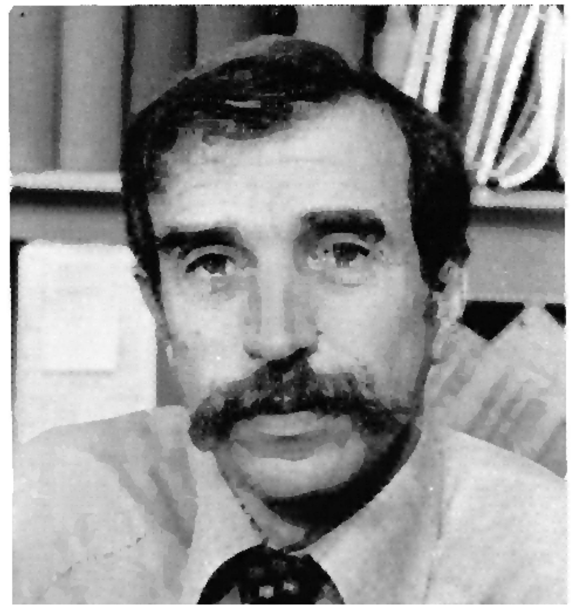

7x7 cutset sampling

full ACA segmentation

PSNR = 28.8 dB

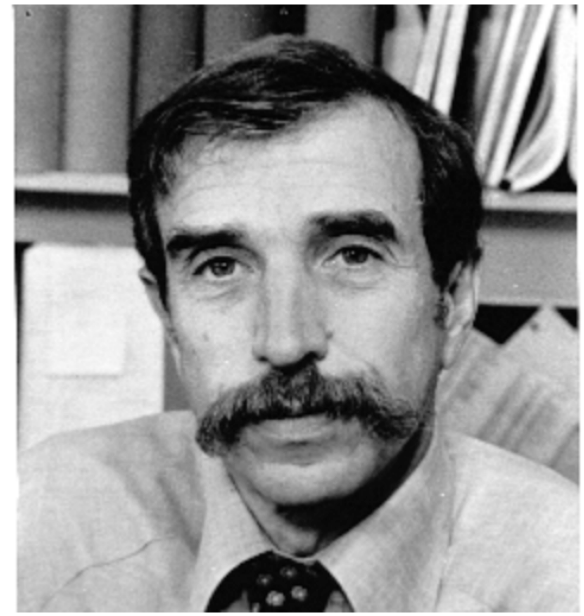

conventional sampling

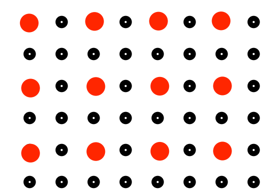

bilinear reconstruction

PSNR = 27.1 dB

Example: Sampling Density 1/2

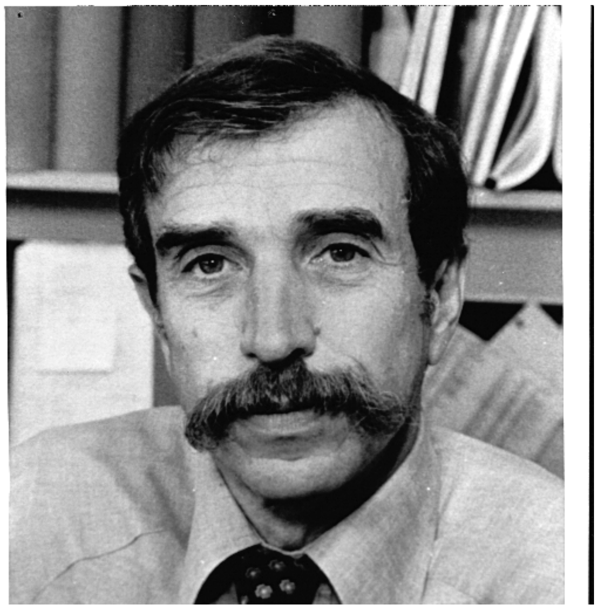

original

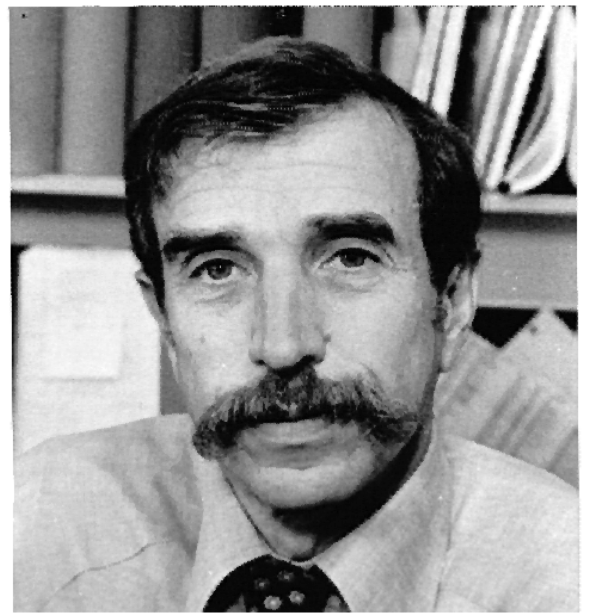

4x3 cutset sampling

ACA cutset segmentation

Gaussian MRF model

MMSE estimation

PSNR = 32.2 dB

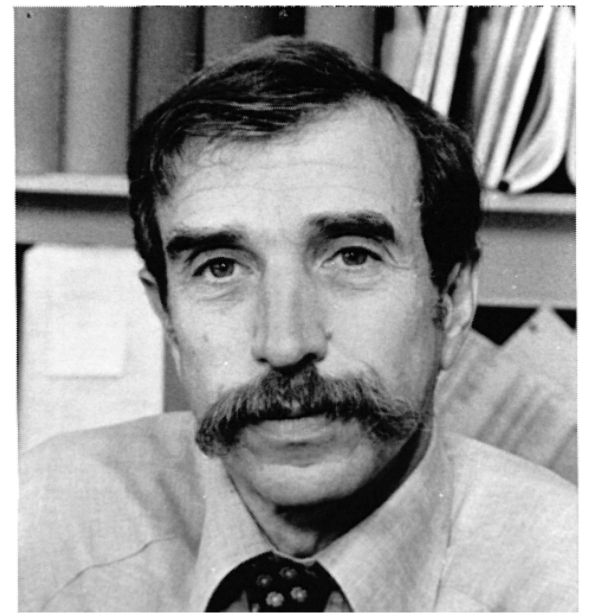

conventional sampling

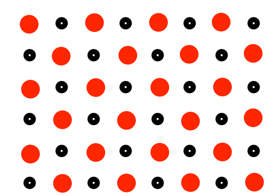

bilinear reconstruction

PSNR = 34.5 dB

PNSR for 'al'

	MRF MMSE			expon. corr. MMSE			distance-based			bilinear
grid	ACA	cut-set	no	ACA	cut-set	no	ACA	cut-set	no	conv'l samplng
7x7	28.8	27.2	27.1	28.2	28.3	28.9	27.3	27.5	26.9	27.1
4x3	33.1	32.2	32.2	33.4	33.8	34.7	32.2	32.2	30.1	34.5

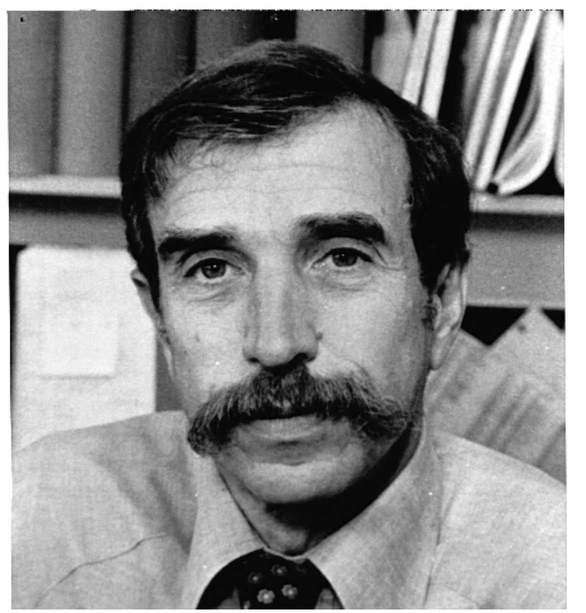

PNSR for ‘tools’

	MRF MMSE			expon. corr. MMSE			distance-based			bilinear
grid	ACA	cut-set	no	ACA	cut-set	no	ACA	cut-set	no	conv’l samplng
7x7	29.1	26.7	27.8	27.9	27.8	29.2	27.3	27.1	26.9	29.9
4x3	33.8	33.2	35.7	33.5	33.7	34.7	33.9	33.8	33.0	38.7

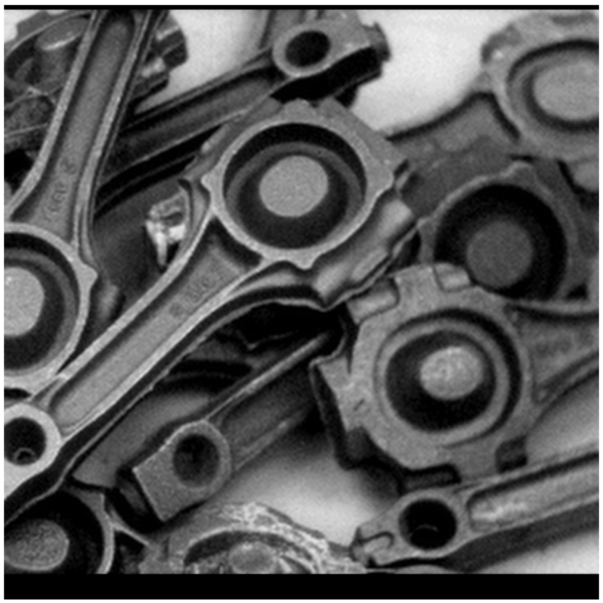

Improved Reconstruction Method

- Matt Prelee et al. ICIP 2012:
  - Image is modeled as piecewise planar plus MRF,
    i.e., as MRF whose mean is piecewise planar
  - For each cutset block, “K-planes algorithm” finds K planes that match image on block boundary, and segments boundary according to the planes associated with it. (typically K=3)
  - Each block interior pixel is associated with one of the K-plans via ad hoc rule, i.e. segmentation extended to the block interior.
  - Each block interior pixel is interpolated as before using MRF model from pixels associated with the same plane (and only these), and with the plane giving the mean of MRF.

Comparison of New and Previous Method

- B=7

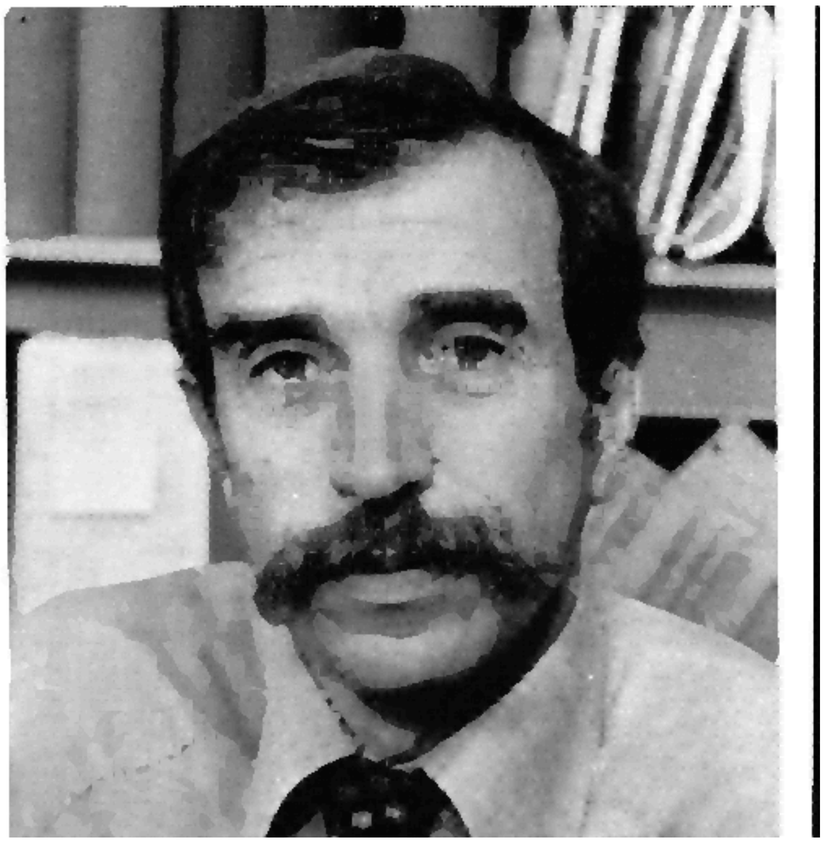

ICIP 2011

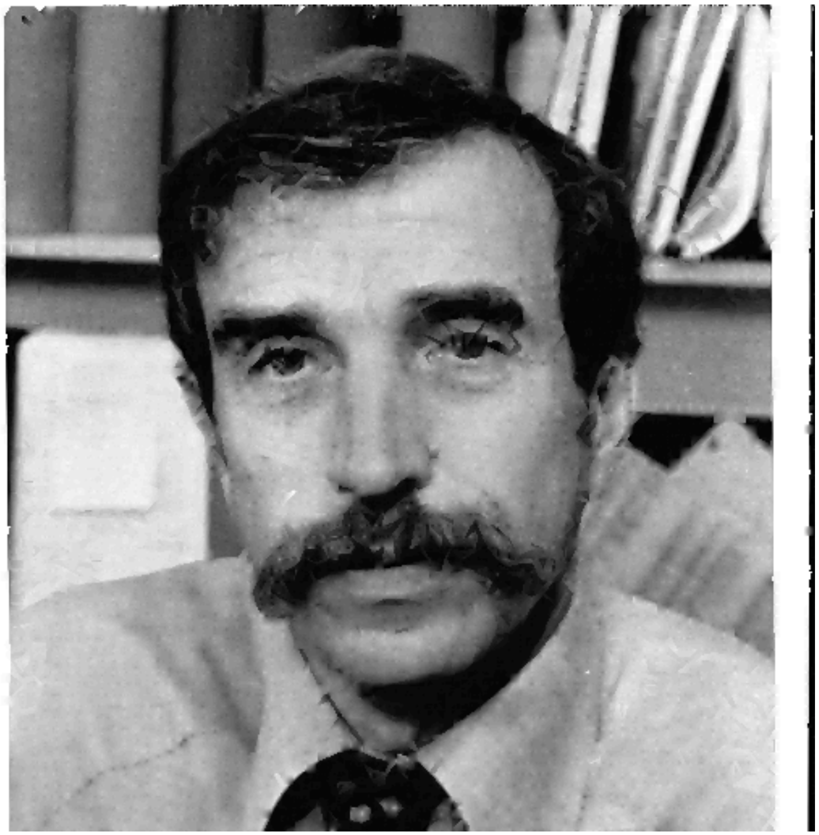

ICIP 2012

Sampling Theorem for Manhattan Cutset Sampling

- If image spectrum is bandlimited to cross-shaped region below,

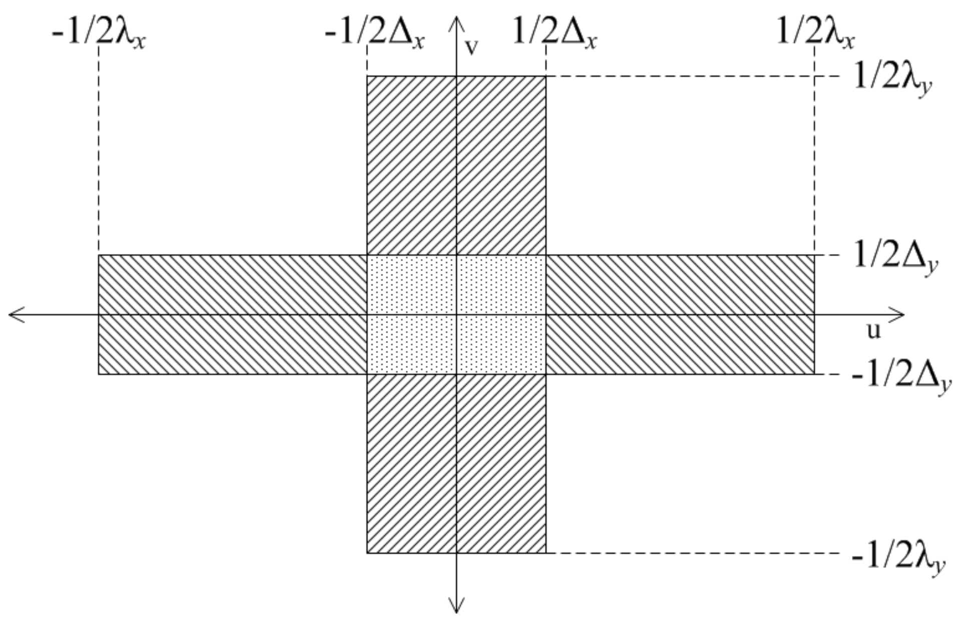

- No larger sampling rate is possible for images bandlimited to this region
- No larger frequency region permits perfect reconstruction

it can be perfectly recovered from Manhattan cutset sampling below

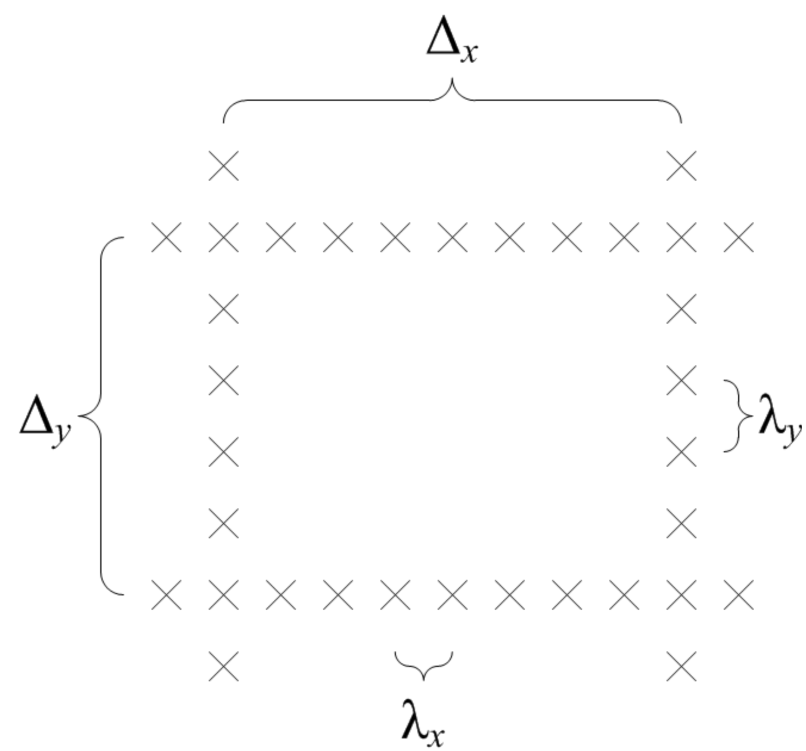

[ICASSP 2012, M. Prelee, DN]

Source Localization

- Goal: Wireless sensor network nodes measures signal strength and collaborate to estimation position.

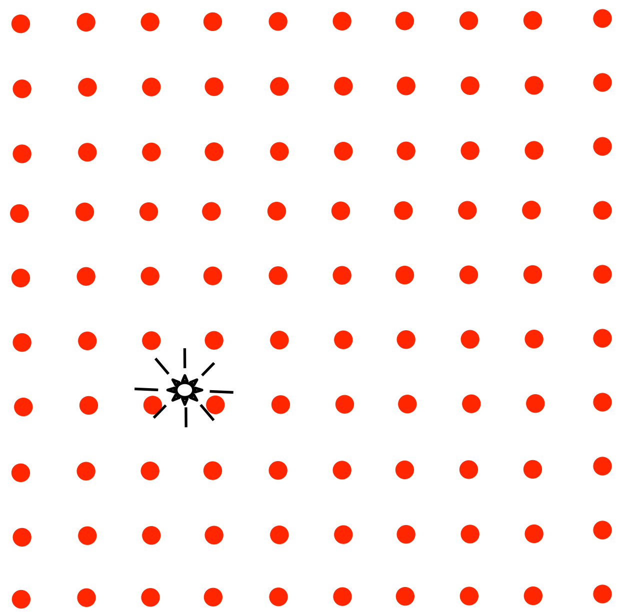

conventional lattice deployment

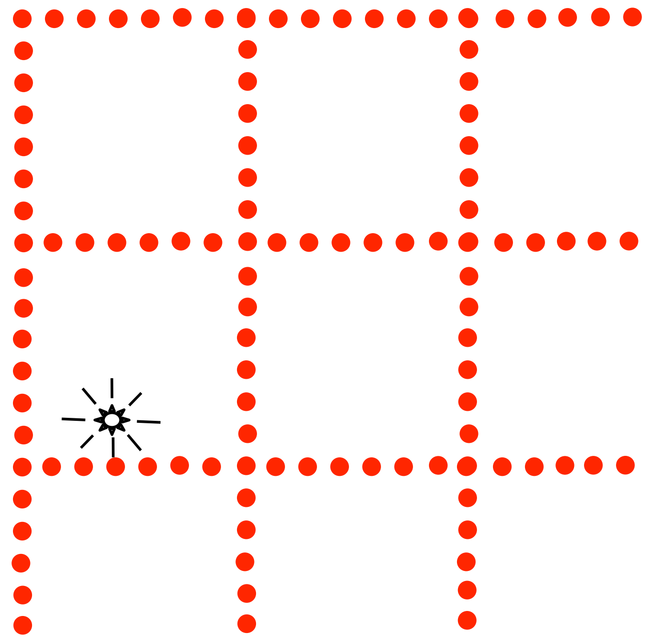

cutset deployment

- Performance: mean-squared position error vs. communication energy

[ICASSP 2012, M. Prelee, DN]

Source Localization Scenario

- N sensors deployed over some geographic region
- Each sensor measures signal strength $y_i = \dfrac{A}{\|x_i - \theta\|^{\beta}} + n_i$
- Sensors whose measurements lie above a threshold communicate and collaborate to make estimate of position.
- Performance measures:
  - Detection rate
  - False alarm rate
  - Mean squared position error
  - Energy required for communication

$$E = \#\text{ bits} \times \#\text{ hops/bit} \times \text{energy per bit per hop}$$

$$\text{energy per bit per hop} = c\,(\text{distance})^{\beta}, \quad \beta \cong 4 \text{ typical}$$

[ICASSP 2013, M. Prelee, DN]

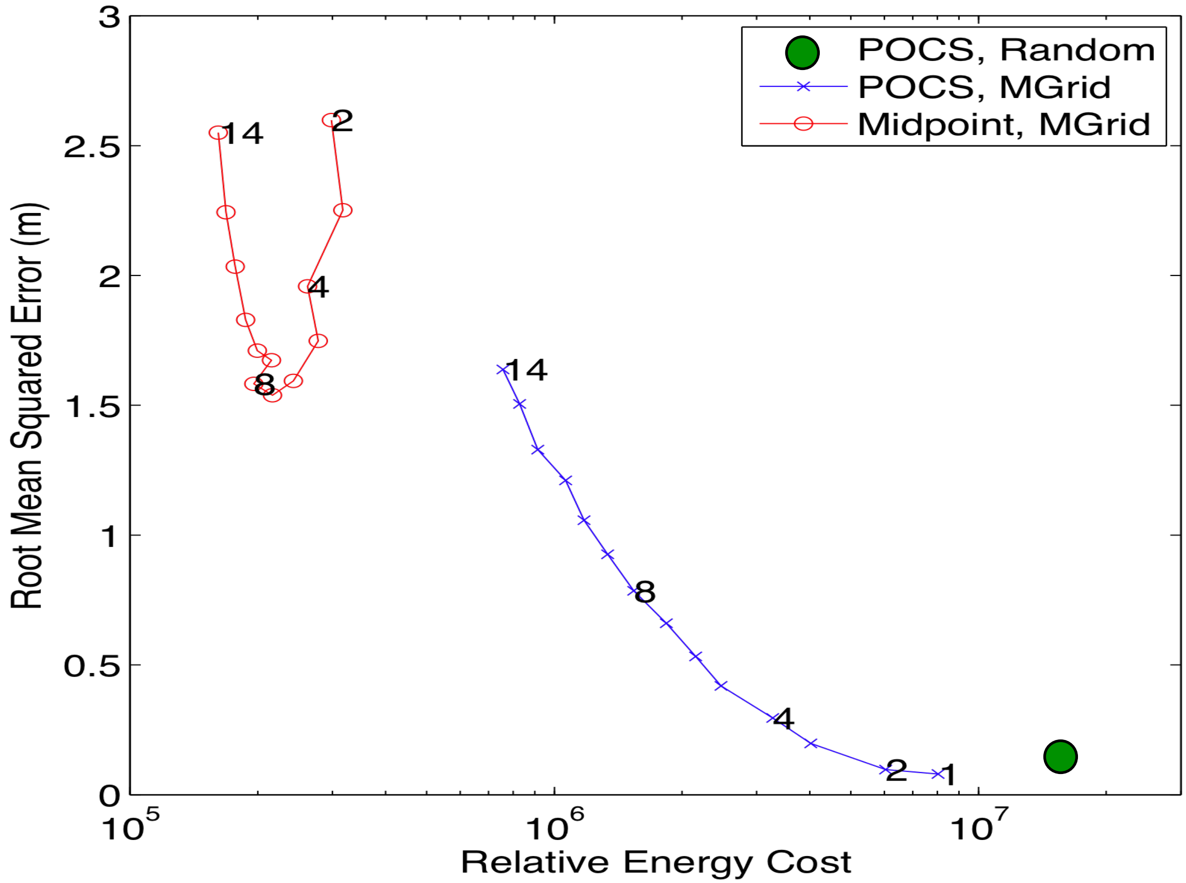

- POCS is Projection onto Convex Sets Method of Blatt, Hero, 2006, for random sensor deployment.
- Midpoint algorithm is very simple, very low energy algorithm that separately estimates horizontal and vertical coordinate of source as midpoint of sensors above threshold.

Ongoing Work and Future Directions

- Hierarchical version of lossy bilevel coding
- Improved reconstruction methods for nonbandlimited images
- Cutset and Manhattan sampling in higher dimensions
  - For video, for example
  - Reconstruction methods
  - Sampling theorems
- Sensor networks with Manhattan grid sensor deployment
  - Localization
  - Communication throughput scaling analysis
  - Other sensor network tasks