UMassAmherst

Autonomous precision error in low-level computer vision tasks

Andrés Corrada-Emmanuel <corrada@cs.umass.edu>

Howard Schultz <hschultz@cs.umass.edu>

University of Massachusetts, Amherst

Aerial Imaging and Remote Sensing Laboratory

Presented by

John Paisley

Duke University

Outline

- *Model-redundant data: An increasingly common scenario*
  - *Lots of data, many ways to process it: each leading to a different model/predictions.*
  - *When predictions go wrong, who is to blame: the data or the algorithm?*
- *Computer-vision task: make maps from aerial images*
  - *Ground truth absent.*
  - *How should multiple terrain models be fused into a single best one? Model weights are a function of the precision error covariance matrix.*
- *Autonomous difference equations and compressed sensing*
  - *Error = accuracy + precision.*
  - *Ground truth is not needed for precision error estimation: a generalization of the bias/variance theorem.*
  - *Given enough models, precision error between models is sparsely correlated.*

The problem: Model-redundant data without ground truth

- *Increasingly common scenario in many scientific fields*
  - *Lots of data.*
  - *Lots of algorithms for processing that data to create different models/predictions of an underlying phenomena.*
  - *Ground truth absent or expensive to obtain.*
- *Questions we would like to answer in these conditions*
  - *What data inputs are good?*
  - *What algorithms are appropriate for processing the data?*
  - *What is the optimal way of fusing all our model predictions?*
  - *What is the minimum uncertainty in a final fused prediction?*

Idea of the talk

- *Precision error of a collection of models/predictions can be reconstructed without any ground truth.*
  - *Total error = accuracy error + precision error.*
  - *No ground truth needed for precision error estimation.*
  - *As the number of models/predictions increases, the number of uncorrelated models increases rapidly – the precision covariance matrix becomes sparse.*
  - *Given enough models/predictions, the covariance matrix of the precision error can be reconstructed using L-1 minimization.*

The Bias-Variance Theorem of Machine Learning

- *Bias/accuracy: how far is the prediction from the ground truth?*
- *Variance/precision: how noisy is your model given the data?*

$$\mathbb{E}[L] = \underbrace{\int (y(x) - \mathbb{E}[t|x])^2 \, p(x) \; dx}_{\text{bias}} + \underbrace{\int (\mathbb{E}[t|x] - t)^2 \, p(x) \; dx}_{\text{variance}}$$

- *Precision independent of ground truth!*
- *Our claim: In many circumstances, precision of a collection of models can be computed without the need for any ground truth.*

Computer-vision task: Digital Elevation Models (DEM)

- *Each pairs of photographs goes through an image matching process to produce a Digital Elevation Map (DEM) represented by*

$$\{Z(x_i, y_i)\}$$

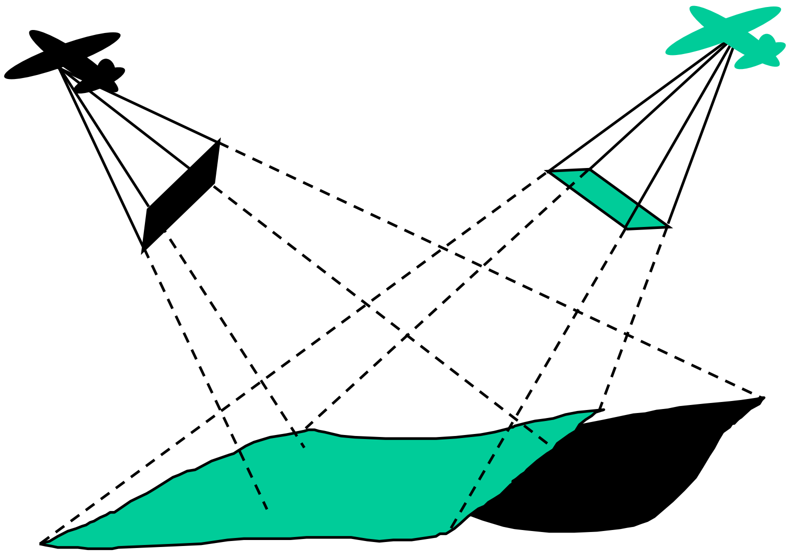

- *Many sources of error:*
  - *Matching on occlusions*
  - *Sensor position and orientation errors*
  - *Etc.*
- *Ground truth is extremely difficult to obtain*

Autonomous precision error estimation

- Given n DEMs we can construct globally invariant quantities that cancel out ground truth.

$$\Delta_{P,Q}(x,y) = \left(\frac{1}{P}\sum_{i}^{P} z_i\right) - \left(\frac{1}{Q}\sum_{j}^{Q} z_j\right)$$

- Any estimate can be written as

$$z_i(x,y) = z_{\text{true}}(x,y) + \delta_i(x,y)$$

- But

$$1/P * (P * z_{\text{true}}) - 1/Q * (Q * z_{\text{true}}) = 0$$

- Ground truth cancels out!

$$\Delta_{P,Q}(x,y) = \left(\frac{1}{P}\sum_{i}^{P} \delta_i\right) - \left(\frac{1}{P}\sum_{j}^{Q} \delta_j\right)$$

A linear system for the average error covariance matrix

- *Square autonomous differences and average over the predictions*

$$\sum_{\text{predictions}} \left(\frac{1}{2}(\delta_1+\delta_2)-\delta_1\right)^2 = \frac{1}{4} < \delta_1^2 > + \frac{-1}{2} < \delta_1\delta_2 > + \frac{1}{4} < \delta_2^2 >$$

- *Under-determined linear system always for covariance matrix entries*

$$\begin{pmatrix} \sum_{(x,y)} \left(\frac{1}{2}(z_1+z_2)-z_1\right)^2 \\ \sum_{(x,y)} \left(\frac{1}{2}(z_1+z_2)-z_2\right)^2 \end{pmatrix} = \begin{pmatrix} \frac{1}{4} & -\frac{1}{2} & \frac{1}{4} \\ \frac{1}{4} & -\frac{1}{2} & \frac{1}{4} \end{pmatrix} \begin{pmatrix} < \delta_1^2 > \\ < \delta_1\delta_2 > \\ < \delta_2^2 > \end{pmatrix}$$

$$S = \Phi\Delta$$

Compressed sensing techniques can recover error signal

- *Sparse signals can be reconstructed with high probability*

$$\min ||\Delta||_1 \text{ subject to } S = \Phi\Delta$$

- *In many situations the signal errors – not the signal – are sparsely correlated*

UMassAmherst

Given enough models, precision error is sparse enough

- *For m models, the number of equations needed to solve for all the covariance matrix entries:*

$$m(m+1)/2$$

- *Number of equations from autonomous difference equations:*

$$m(m+1)/2 - m$$

- *Requires at least m uncorrelated entries.*
- *As m increases, the relative abundance of uncorrelated entries decreases ~ 1/m*

Twenty-Nine Palms Dataset and Maps

- *Four photographs are matched to produced ten DEMs -> 10x10 covariance matrix (blue positive correlation, red negative).*

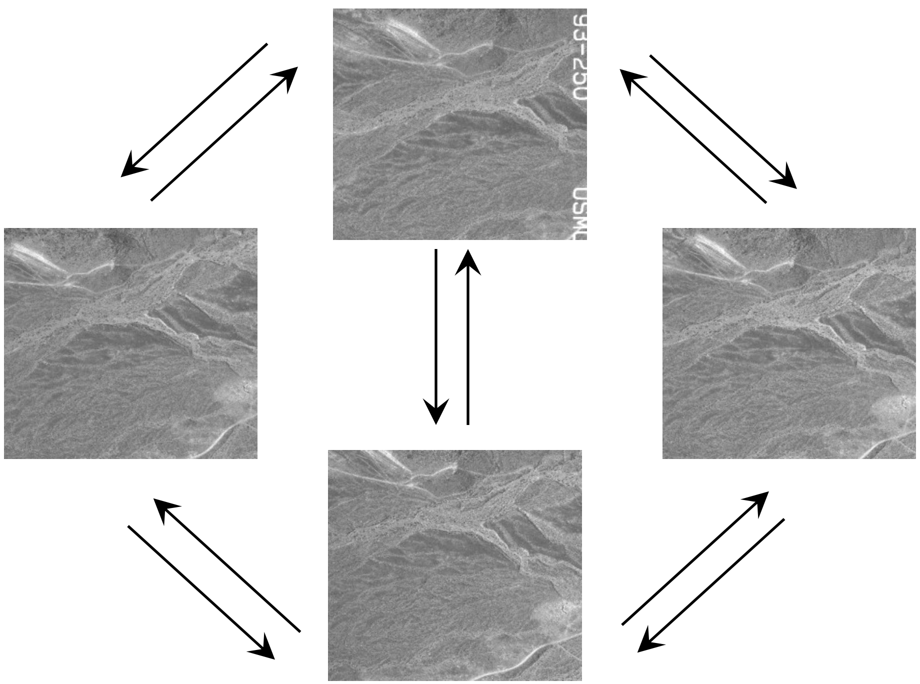

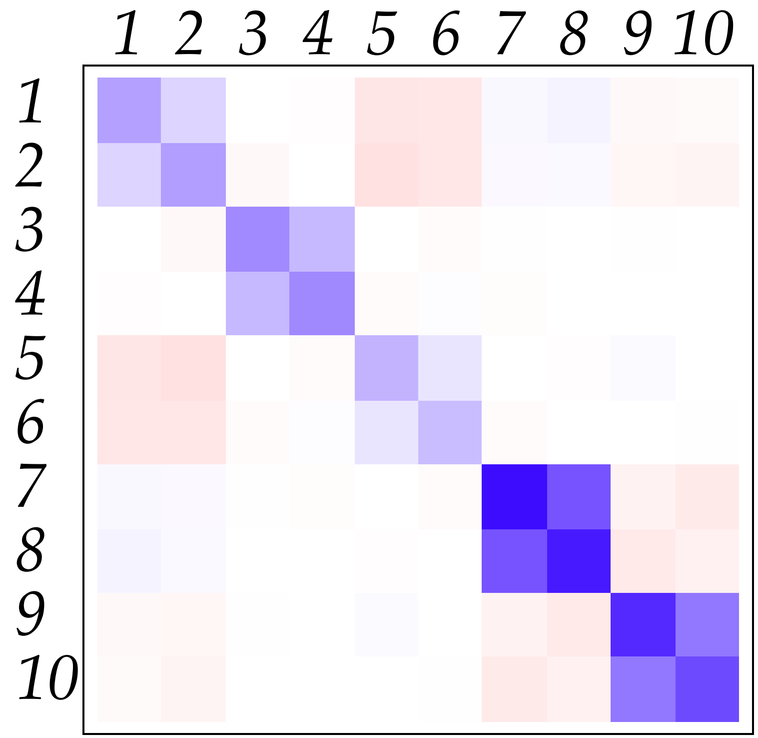

Precision error reconstruction for DEMs

- *Four photographs (A,B,C,D) produced ten DEMs (Z_{AB}, Z_{BA}, Z_{AC},…) and a 10x10 precision covariance matrix.*
- *No assumptions about which DEMs have correlated errors.*
- *Reconstruction discovered an asymmetry. DEM pairs (Z_{AB},Z_{BA}) (Z_{AC},Z_{CA}),… have partially uncorrelated errors. See the 2x2 block structure on the diagonal.*
- *Error covariance matrix has many uncorrelated entries.*
- *Some photographs produced less noisy models – some data inputs not as good as others.*

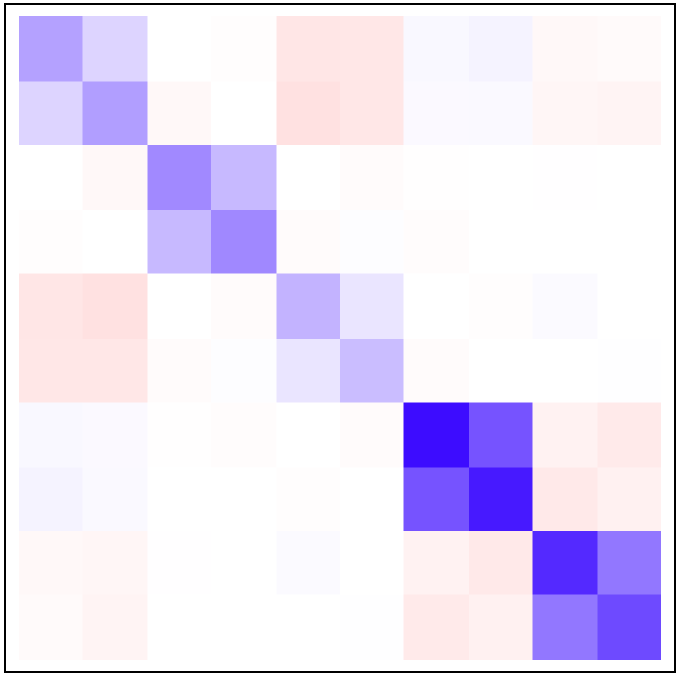

Conclusions

- *Precision error of models can be recovered independent of ground truth when you have "enough" models.*
- *The structure of the covariance matrix contains two fingerprints:*
  - *How the data was processed*
  - *The quality of the data inputs used to create the models*
- *This framework has possible applicability across all of machine learning: precision error of many models applied to data.*

Acknowledgements

This work was supported by grant (IIS-0430742) from the United States National Science Foundation and contract TT0690688 from the Advanced Technology Laboratory, Lockheed-Martin Corporation.