Nonconvex Regularization and Satellite Imagery

Field-based SAR denoising

Rick Chartrand, Descartes Labs

May 25, 2016

A mathematician is a device for turning coffee into theorems.
–Alfréd Rényi

What we really need is a machine to turn some of those theorems back into coffee.
–A. J. Tolland

Descartes Labs

Outline

A bit about Descartes Labs

SAR denoising

Descartes Labs

Who we ~~are~~ were

We're a startup in Los Alamos, NM.

Descartes Labs

Who we ~~are~~ were

We're a startup in Los Alamos, NM.

Descartes Labs

Big data: from space to the cloud

We process petabytes of imagery from many satellites, including:

- Landsats 1–8 (USGS)
- Terra and Aqua (NASA)
- Sentinel 1 and 2 (ESA)
- Dove and RapidEye (Planet Labs)

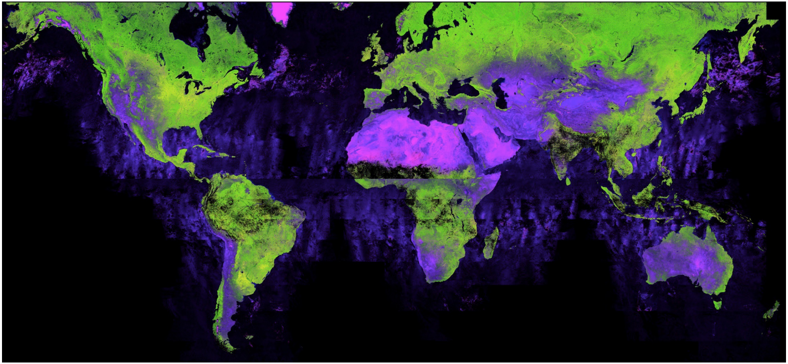

Descartes Labs

Actionable advice for agriculture

Our current focus is on extracting information of value to the agriculture industry.

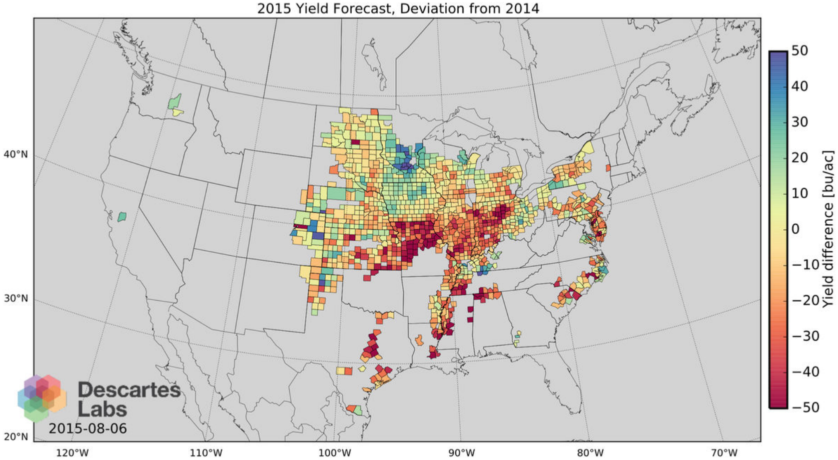

Actionable advice for agriculture

Our current focus is on extracting information of value to the agriculture industry.

Outline

A bit about Descartes Labs

SAR denoising

Sentinel 1

Sentinel 1 uses synthetic-aperture radar, which is unaffected by clouds.

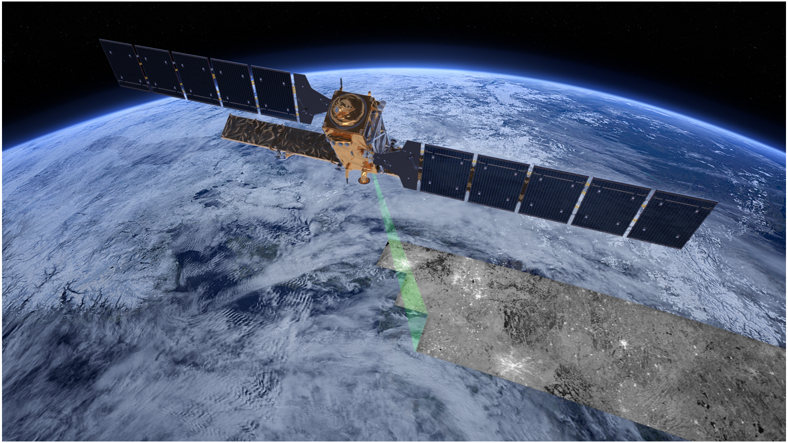

Descartes Labs

SAR image example

A processed Sentinel 1-A image from west-central Iowa: (The colors are VH and VV polarizations, and $\frac{VV-VH}{VV+VH}$ as a third channel.)

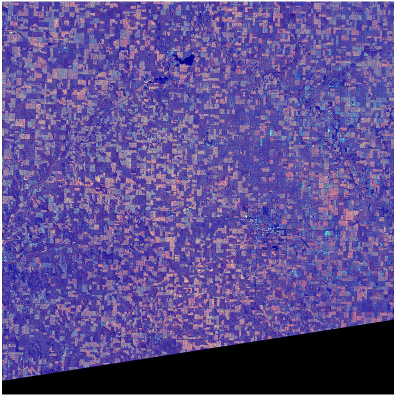

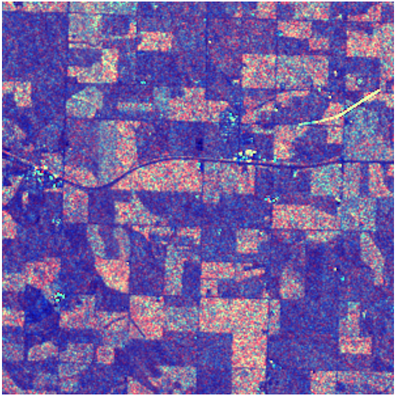

Descartes Labs

Fields

We have lots of other, noiseless imagery over time, allowing us to produce maps of regions of consistent land use (cf. image segmentation).

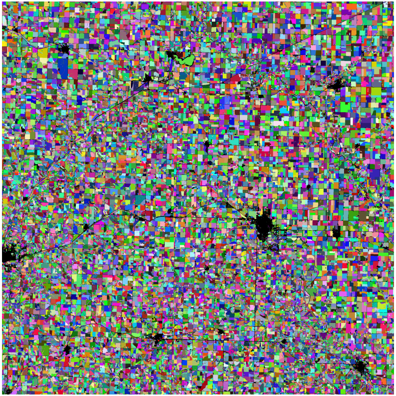

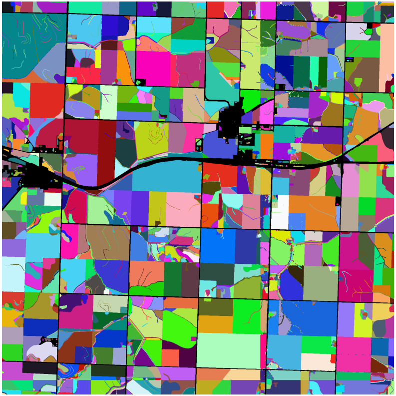

Descartes Labs

A simple approach

Making the image constant on fields is easy and works well. Outliers make the median a good choice. The multiplicative Gaussian noise model makes the mean of the logarithm a good choice.

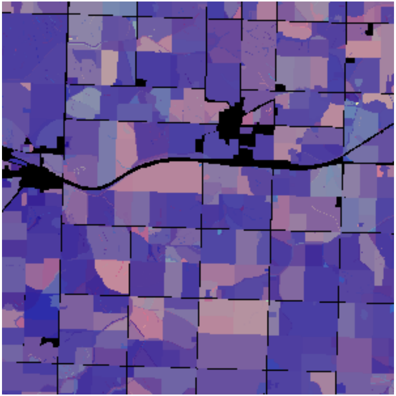

fieldwise median

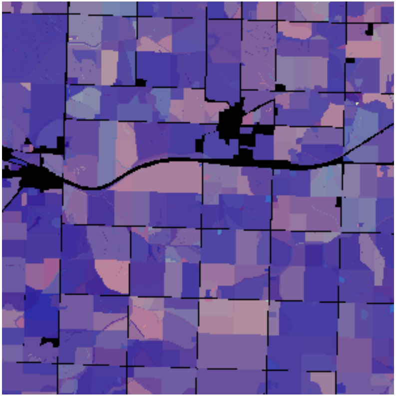

fieldwise exp-mean-log

Descartes Labs

Mumford-Shah functional

If f is a noisy image defined on Ω, we seek an image u and edge-set Γ:

$$\min_{u,\Gamma} \int_{\Omega \setminus \Gamma} |\nabla u|^2 + \frac{\lambda}{2} \int_{\Omega} |u - f|^2 + \mu \mathcal{H}(\Gamma). \qquad (*)$$

The regularization of u is turned off at Γ, which ideally corresponds to the set of edges in the image. This prevents the regularization from blurring the edges.

$(*)$ is difficult to solve.

Field simplification

In our case, we already have a good approximation of the edges. Let Γ be the set of field edges. Now we solve the following:

$$\min_u \int_{\Omega \setminus \Gamma} |\nabla u|^p + \frac{\lambda}{2} \int_{\Omega} |u - f|^2.$$

- Using $p \leq 1$ will help preserve edges that are missing from the fieldmap.
- For $p = 1$, this is total-variation regularization, but with the TV diffusion turned off at edges.
- Using $p < 1$ results in sharper non-field edges and better contrast preservation.

Algorithm and implementation

We use iteratively-reweighted least squares (equivalently, iterative linearization of the gradient, AKA lagged diffusivity). This means iteratively solving:

$$(R\nabla)^T \operatorname{diag}(|\nabla u_{n-1}|^{p-2}) R\nabla u_n + \lambda(u_n - f) = 0,$$

where R is the projection onto the non-edge pixels.

Algorithm and implementation

We use iteratively-reweighted least squares (equivalently, iterative linearization of the gradient, AKA lagged diffusivity). This means iteratively solving:

$$(R\nabla)^T \operatorname{diag}(|\nabla u_{n-1}|^{p-2}) R\nabla u_n + \lambda(u_n - f) = 0,$$

where R is the projection onto the non-edge pixels.

The edge pixels remain noisy, so we repeat the iteration with the edge and non-edge pixels reversed, initializing with the previous result.

Edge maps

This approach requires labeling edge pixels.

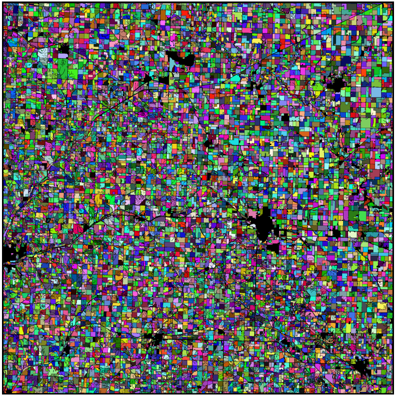

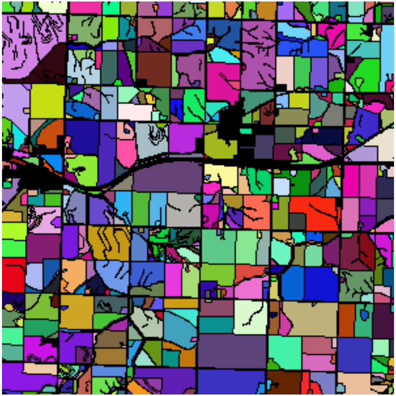

Descartes Labs

Results

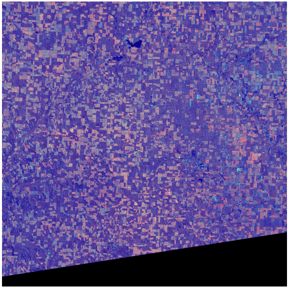

noisy

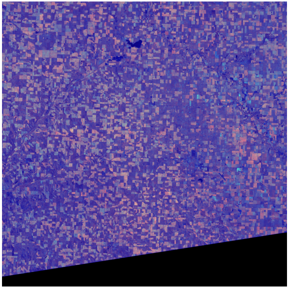

$p = 1, \lambda = 1.0$

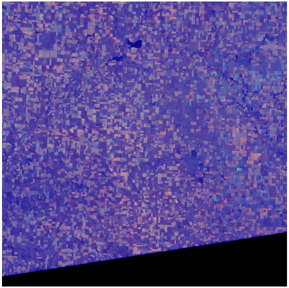

$p = 1, \lambda = 0.1$

Weaker regularization gives a realistic-looking result. Stronger regularization approximates the uniform-field result, while allowing non-field edges.

Results

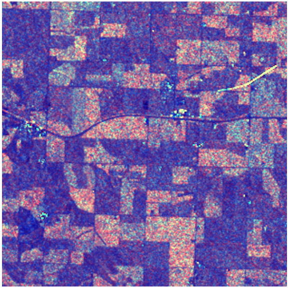

noisy

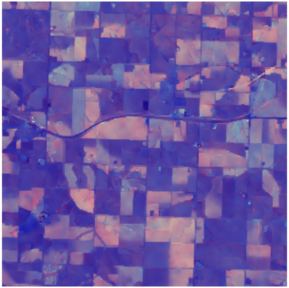

$p = 1, \lambda = 1.0$

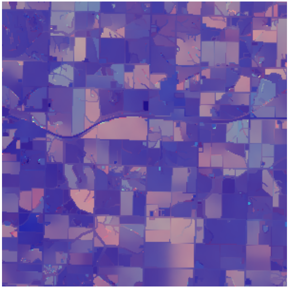

$p = 1, \lambda = 0.1$

Weaker regularization gives a realistic-looking result. Stronger regularization approximates the uniform-field result, while allowing non-field edges.

Descartes Labs

Comparison with edgeless denoising

We get a direct comparison with denoising without edge assistance by simply turning off the edge-set:

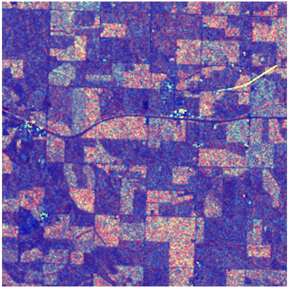

noisy

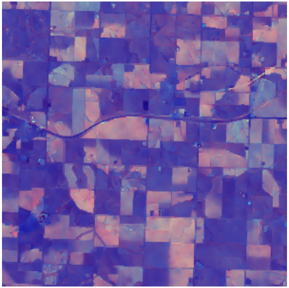

$\lambda = 1.0$, use edges

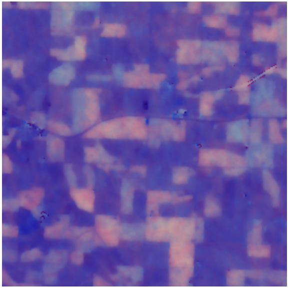

$\lambda = 1.0$, no edges

Without turning the diffusion off at edges, features are blurred.

Comparison with edgeless denoising

We get a direct comparison with denoising without edge assistance by simply turning off the edge-set:

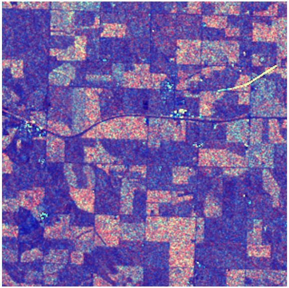

noisy

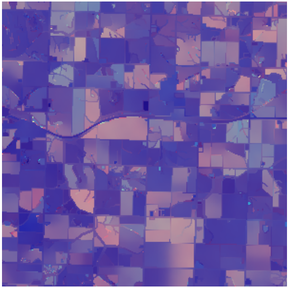

$\boldsymbol{\lambda = 0.1}$, use edges

$\boldsymbol{\lambda = 0.1}$, no edges

Without turning the diffusion off at edges, features are blurred.

Descartes Labs

Comparison with small p

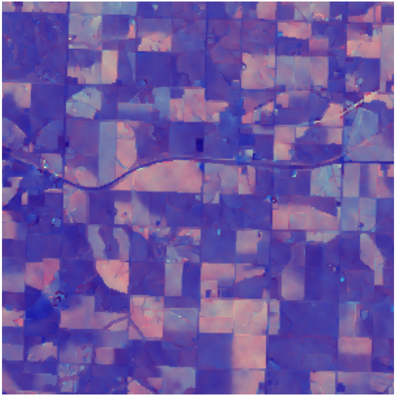

$p = 1$

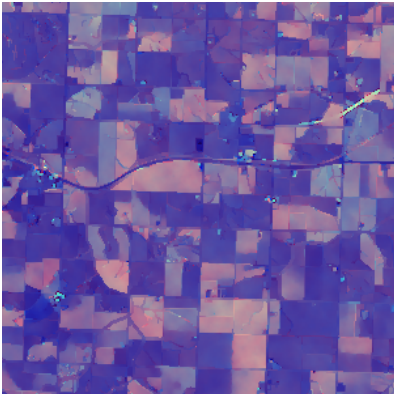

$p = 1/4$

Using $p < 1$ gives sharper non-field edges, and preserves contrast better.

Descartes Labs

A possible storyline

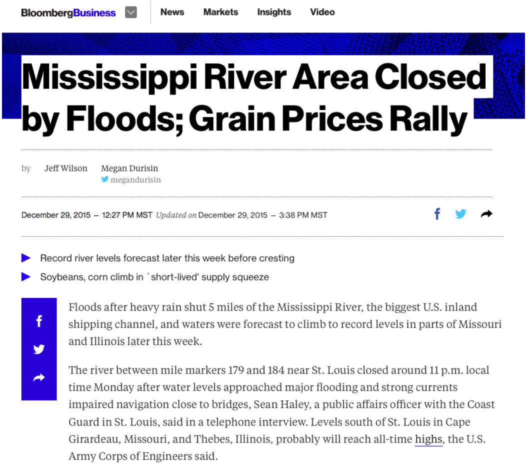

BloombergBusiness | News | Markets | Insights | Video

Mississippi River Area Closed by Floods; Grain Prices Rally

by Jeff Wilson Megan Durisin
megandurisin

December 29, 2015 – 12:27 PM MST *Updated on* December 29, 2015 – 3:38 PM MST

- Record river levels forecast later this week before cresting
- Soybeans, corn climb in `short-lived' supply squeeze

Floods after heavy rain shut 5 miles of the Mississippi River, the biggest U.S. inland shipping channel, and waters were forecast to climb to record levels in parts of Missouri and Illinois later this week.

The river between mile markers 179 and 184 near St. Louis closed around 11 p.m. local time Monday after water levels approached major flooding and strong currents impaired navigation close to bridges, Sean Haley, a public affairs officer with the Coast Guard in St. Louis, said in a telephone interview. Levels south of St. Louis in Cape Girardeau, Missouri, and Thebes, Illinois, probably will reach all-time highs, the U.S. Army Corps of Engineers said.

Maybe we can see barge traffic backed up.

Descartes Labs

The Landsat view

A Landsat 8 image of this area during that time:

The Landsat view

A Landsat 8 image of this area during that time:

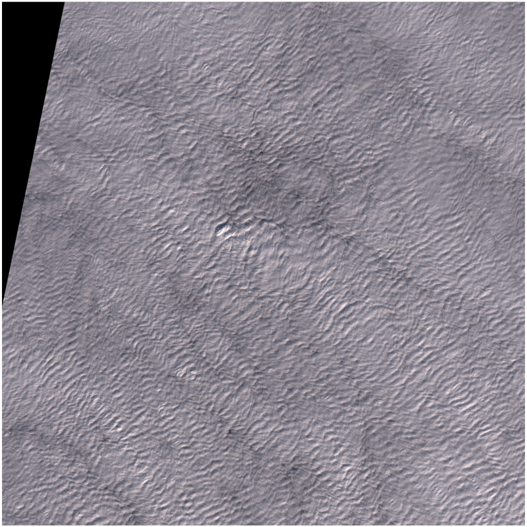

Barges in SAR

Barges pop out in Sentinel-1 images, but pixelwise detection is easier with noise removed. Using $p < 1$ preserves bright pixels better.

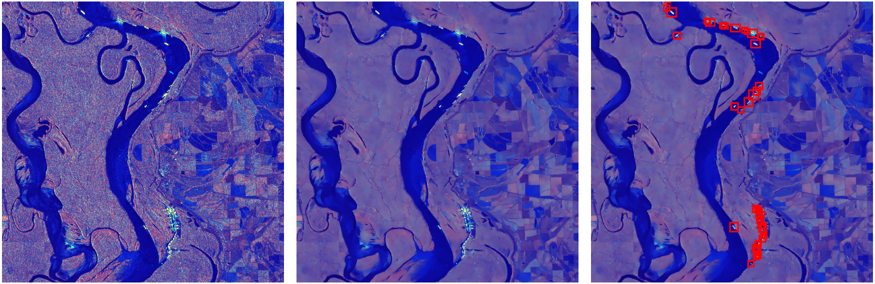

Descartes Labs

Summary

- Descartes Labs: satellite imagery startup. We're hiring!
  `http://www.descarteslabs.com/jobs/`
- SAR imagery from Sentinel 1 lets us see through clouds, but is noisy.
- Using field edges as prior information lets us remove noise with better preservation of features, especially with nonconvex regularization.