Project Proposal EE367 Winter 2017

Exploration of Light Field Reconstruction Using Sparse Samples

Philip Lee, Anqi Ji

Motivation:

The goal of our project is to explore sparse reconstruction of the 4D light field reconstruction. Different methods of capturing the 4D light field exist, which balance a tradeoff between spatial and angular resolution, and require additional hardware and capture times. These methods include a microlens array mounted in front of the camera sensor [1], which comes at the cost of spatial resolution, and camera arrays [2] or cameras mounted on a mechanical gantry [3]. Sparse reconstruction of the light field is desirable as it helps to alleviate some of these tradeoffs. Priors that exploit the sparsity of the light field in some dimension allow the entire light field to be reconstructed from a limited number of samples. For example, the light field can be considered to be sparse in the angular domain due to the light field being the same scene acquired from slightly different perspective. The method that we wish to implement exploits sparsity of the light field in the continuous time Fourier domain to reconstruct the entire light field from a smaller set of samples of the light field [4].

Related work:

Multiple methods have been applied to reconstruct the light field from a sparse set of samples. The prior that is applied to the reconstruction is what differentiates these methods.

In [5], a Lambertian scene is assumed and a 3D Gaussian prior is applied to the scene. This prior comes from the observation that the 4D Fourier transform of the 4D ray space contains a 3D subset of entries whose energy is significantly higher than zero. Under this prior, the light field can be recovered using spatial deconvolutions with a depth invariant Point Spread Function. Though this method is easy to compute, it creates artifacts in non-Lambertian scenes.

Compressive sensing is another approach to reconstruction of the light field using sparse samples. By applying a randomly-coded aperture mask to the lens, the authors of [4] were able to reconstruct the light field with a superior angular resolution and SNR compared to sequentially opening and blocking small regions of the aperture. The light field is recovered using a Bayesian reconstruction model with Total Variation and Gaussian priors. A Majorize-Minimization optimization was used to recover the light field.

A third approach applied in [7] recovers the light field using a randomly coded attenuation mask and dictionaries derived from training data acquired from other light field images. The dictionaries are then applied to solve the ill-posed problem of recovering the entire light field from a random linear combination of light field samples provided by the coded mask.

Project overview:

This section describes the reasoning behind exploiting sparsity in the Continuous Time Fourier Domain, and the algorithm that we will apply in our project to reconstruct the light field. This algorithm is identical to the work in [4].

When computing the discrete Fourier transform (DFT) of a signal using a finite window of samples, if the sparsity of the signals is aligned with the sampling grid of DFT, we get a single spike in the discrete Fourier domain. However, in most natural signals, the signal and the sampling grid usually do not align with each other. As a result, we lose the sparsity in the Fourier domain. For example, if we discretize the continuous signal in fourier domain, we get a signal with a sinc tail instead of a single spike. This effect, called the windowing effect, is commonly seen in light field sampling since they light fields are sampled within a limited 2D window. The tails in DFT usually decay very slowly and this makes our reconstruction very difficult. In our project, we plan to apply the combined linear and nonlinear solver proposed in [4] to remove this windowing effect, and reconstruct the light fields from 1D viewpoint trajectories.

The flowchart of the 2D sparse FFT reconstruction algorithm is shown in Figure 1. The initial guess for the non-zero frequency coefficients is calculated using the Fourier slice theorem in a voting scheme similar to the Hough transform. Then we will use gradient descent to recover the position and the coefficients of the missing frequency. Finally, we reconstruct the missing viewpoints using our continuous non-zero frequency coefficients.

In our project, we want to implement this algorithm to recover sparse light fields. We will evaluate the algorithm by 1) comparing our result with reconstructions from the continuous Fourier domain, 2) compare our results in different lighting conditions, and 3) evaluate our algorithm in scenes with motion blur.

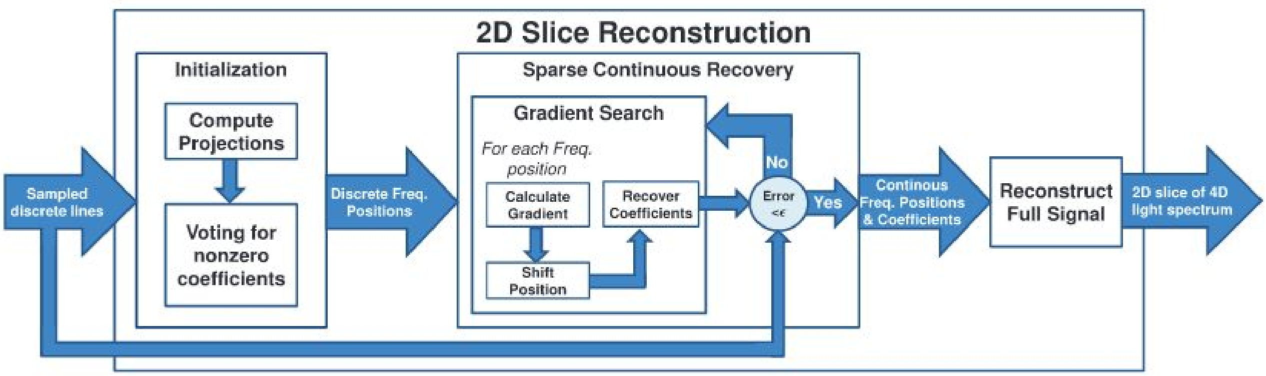

Figure 1. Flowchart of 2D sparse FFT reconstruction algorithm ([4] Fig. 5)

Milestones:

Week 1 (Feb. 12 - 18)	Thoroughly understand [4] and experiment with C++ package provided in [4]. Develop skeleton for MatLab implementation and experiment with batch jobs using FarmShare.
Week 2 (Feb. 19 - 25)	Implement algorithm in [4] and apply to lower resolution single channel images.
Week 3 (Feb. 26 - Mar. 4)	Continue implementing algorithm in [4] and expand to full resolution colour images. Apply algorithm to images in the Stanford light field archive and provided data sets. Compare reconstructed images to ground truth. Experiment with different scenes using different lighting conditions and evaluate performance of reconstruction in the presence of motion.
Week 4 (Mar. 5 - 11)	Complete experiments and comparisons of reconstruction algorithm. If time permits, we will explore additional priors between images in the light field.
Week 5 (Mar. 12 - 17)	Summarize the results and complete poster and report.

References:

[1] R. Ng, “Digital Light Field Photography,” Ph.D. Dissertation, Dept. Elec. Eng., Stanford University, Stanford, CA, 2006.

[2] B. Willburn, “High performance imaging using arrays of inexpensive cameras,” Ph.D. Dissertation, Dept. Elec. Eng., Stanford University, Stanford, CA, 2004.

[3] M. Levoy, P. Hanrahan, “Light field rendering,” in Proc. ACM Conference on Computer Graphics (SIGGRAPH ’96), pages 31–42, New Orleans, USA, August 1996.

[4] L. Shi *et al.,* “Sparsity in Continuous Fourier Domain,” in ACM Transactions on Graphics, Vol. 34, No. 1, Article 12. November, 2014.

[5] A. Levin, F. Durand, “Linear view synthesis using a dimensionality gap light field prior,” in Computer Vision and Pattern Recognition. June, 2010. doi: 10.1109/CVPR.2010.5539854

[6] S. Babacan *et al.,* “Compressive Light Field Sensing,” in IEEE Transactions on Image Processing, Vol. 21, No. 12, pp 4746-4757. December, 2012.

[7] K. Marwah *et al.,* “Compressive light field photography using overcomplete dictionaries and optimized projections,” in ACM Transactions on Graphics - SIGGRAPH 2013, Vol. 32, No. 46. July, 2013.