Adaptive Plenoptic Sampling

Christopher Gilliam, Pier Luigi Dragotti and Mike Brookes
Communications & Signal Processing Group
Imperial College London

14th September 2011

Overview

- Motivation and Image Based Rendering
- Plenoptic Sampling Theory
- Spectral Analysis of a Slanted Plane
- Adaptive Sampling Algorithm
- Conclusions and Future Work

Motivation

Image Based Rendering (IBR) $\Longrightarrow$ Rendering new viewpoints of a scene from a multi-view image set

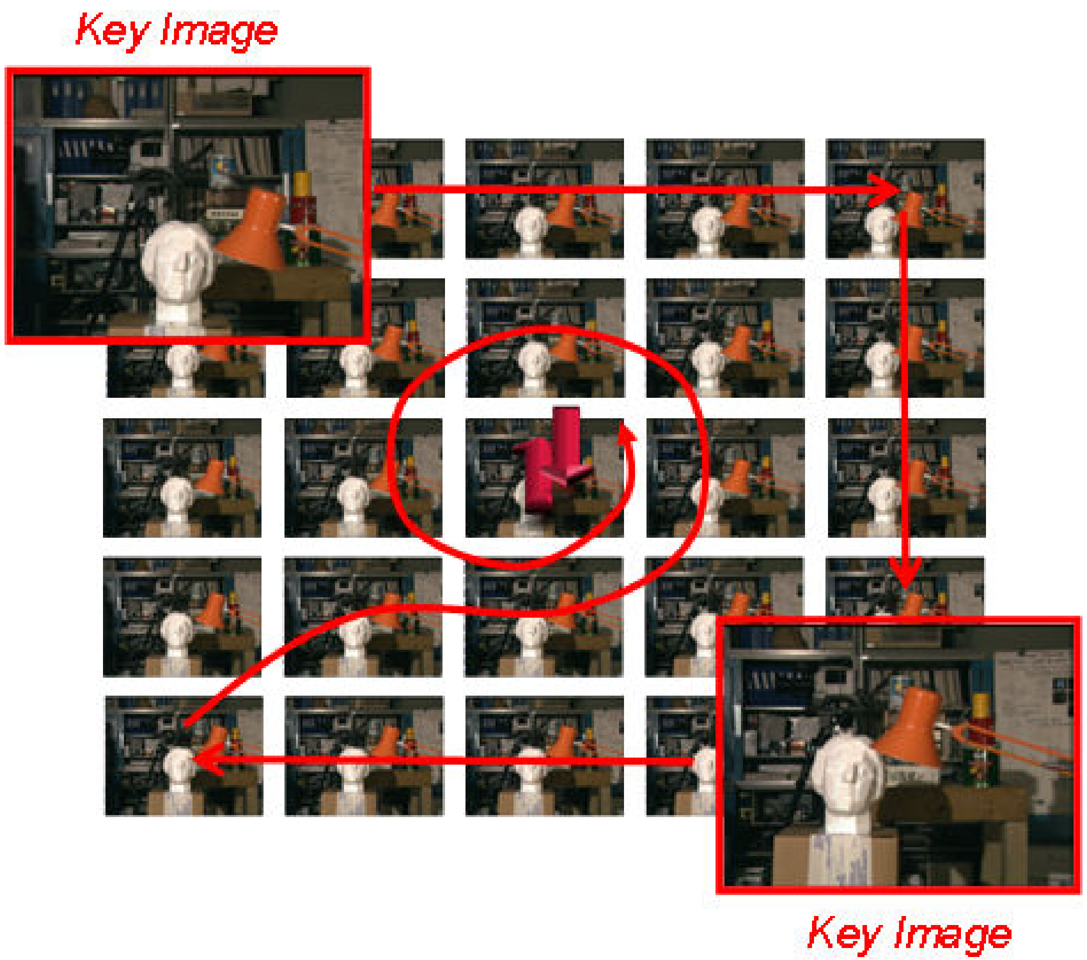

Courtesy of James Pearson [1]

$\looparrowright$ Decide the optimum location of the samples

The Plenoptic Function

IBR in more detail:

- Images sample a set of lights rays from the scene to the camera
- New rendering interpolated from captured light rays
- Lights ray modelled using the 7D *Plenoptic Function* [2]

$\looparrowright$ IBR viewed as the Sampling and Reconstruction of the Plenoptic Function

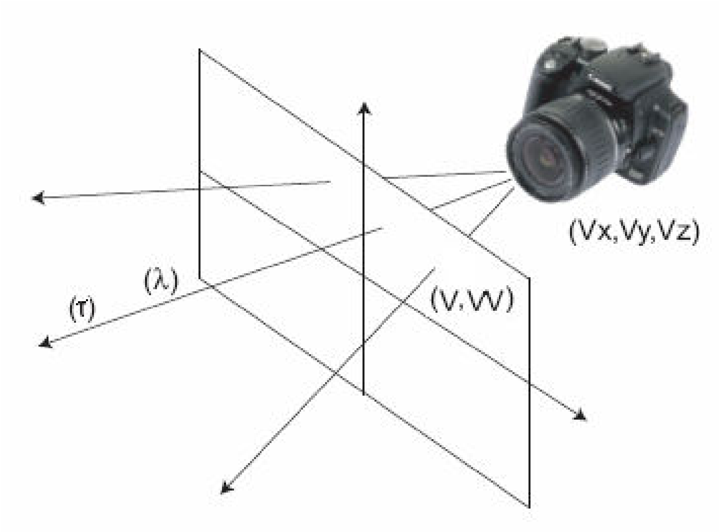

- Camera centre location (v_x, v_y, v_z),
- Viewing direction (v, w),
- Wavelength ν,
- Time τ.

Plenoptic Function and the Epipolar Plane Image

Consider the 2D Plenoptic Function, $p(t, v)$, known as the Epipolar Plane Image (EPI) [3]

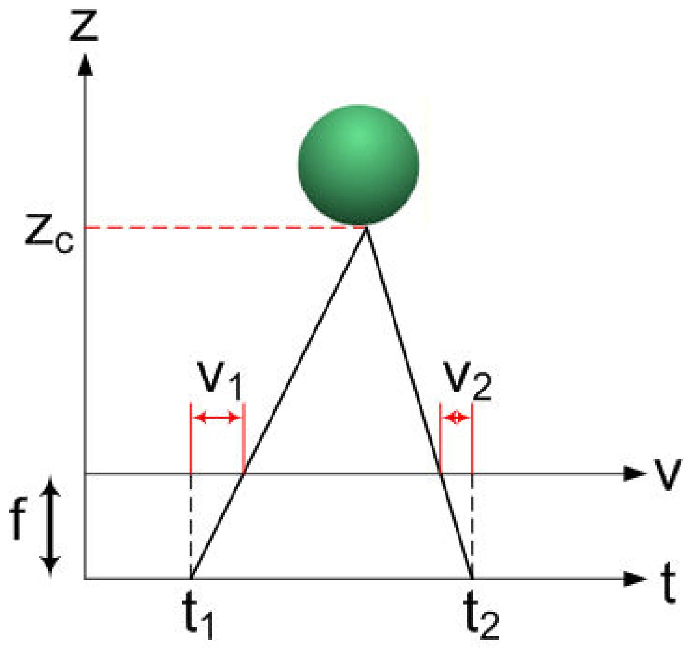

(a) Scene

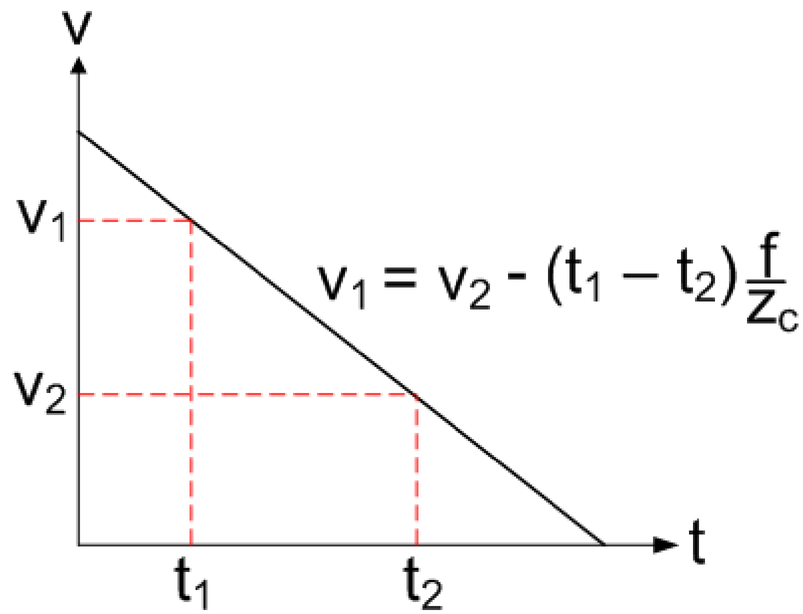

(b) EPI

- Point in the scene $\Longrightarrow$ Line in the EPI plane where the slope depends on the depth
- Fixing a camera position t_1 $\Longrightarrow$ 1D image signal
- Fixing a pixel v_1 $\Longrightarrow$ 1D signal of the pixel captured by all cameras

Plenoptic Spectral Analysis

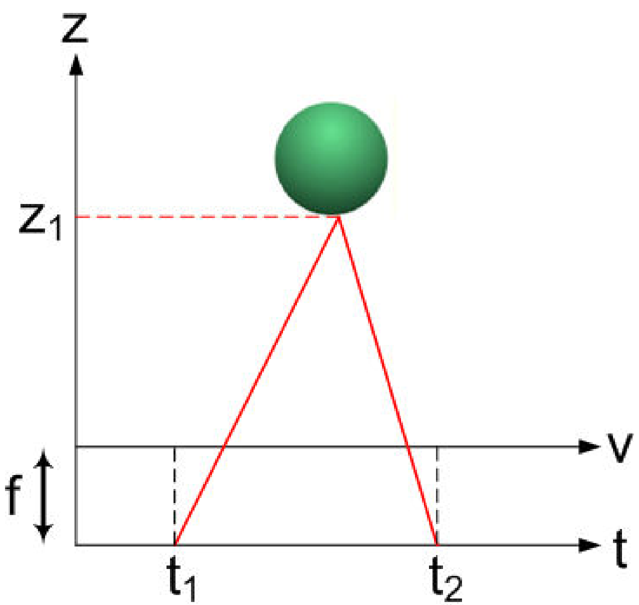

(a) Scene

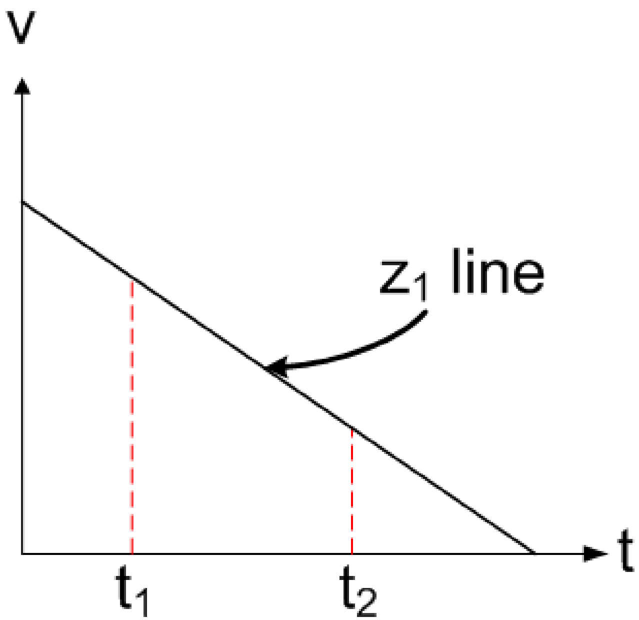

(b) EPI

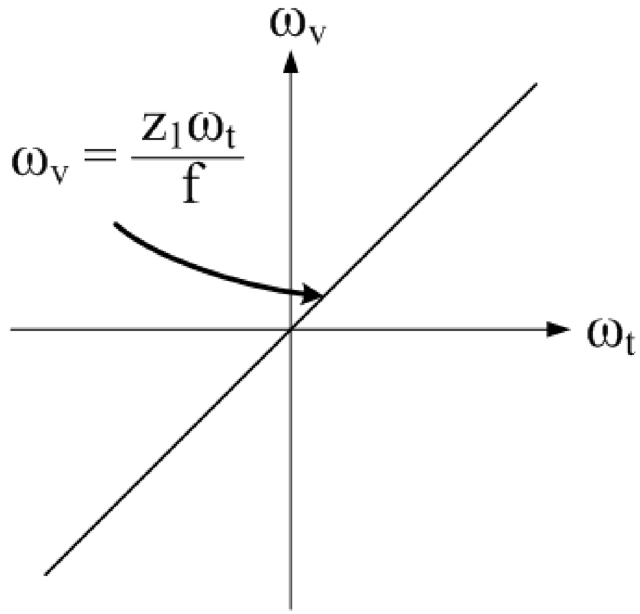

(c) Plenoptic Spectrum

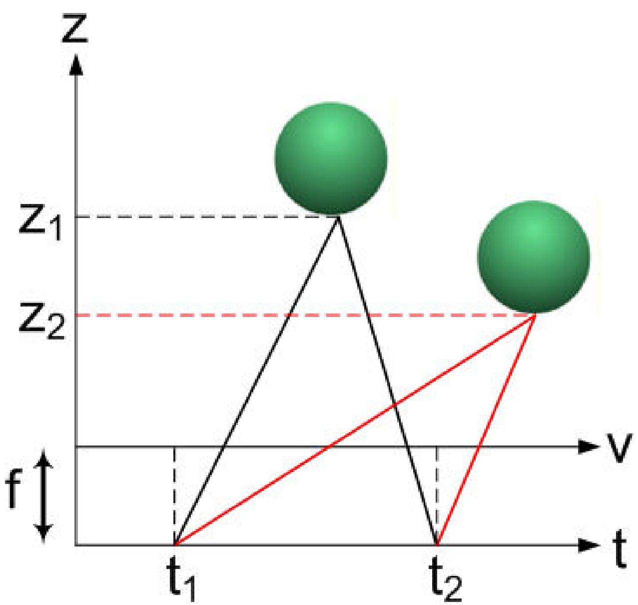

(a) Scene

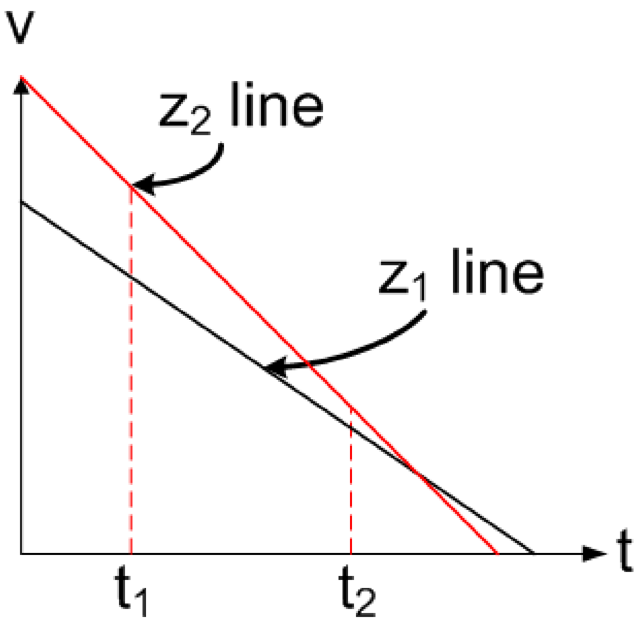

(b) EPI

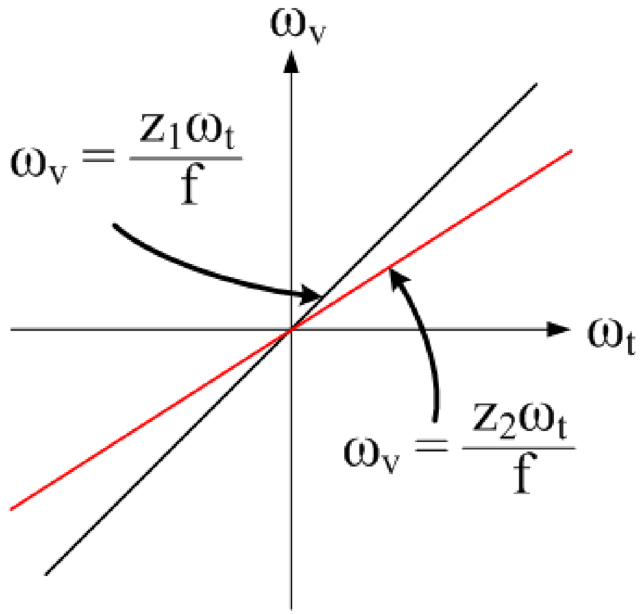

(c) Plenoptic Spectrum

Plenoptic Spectral Analysis

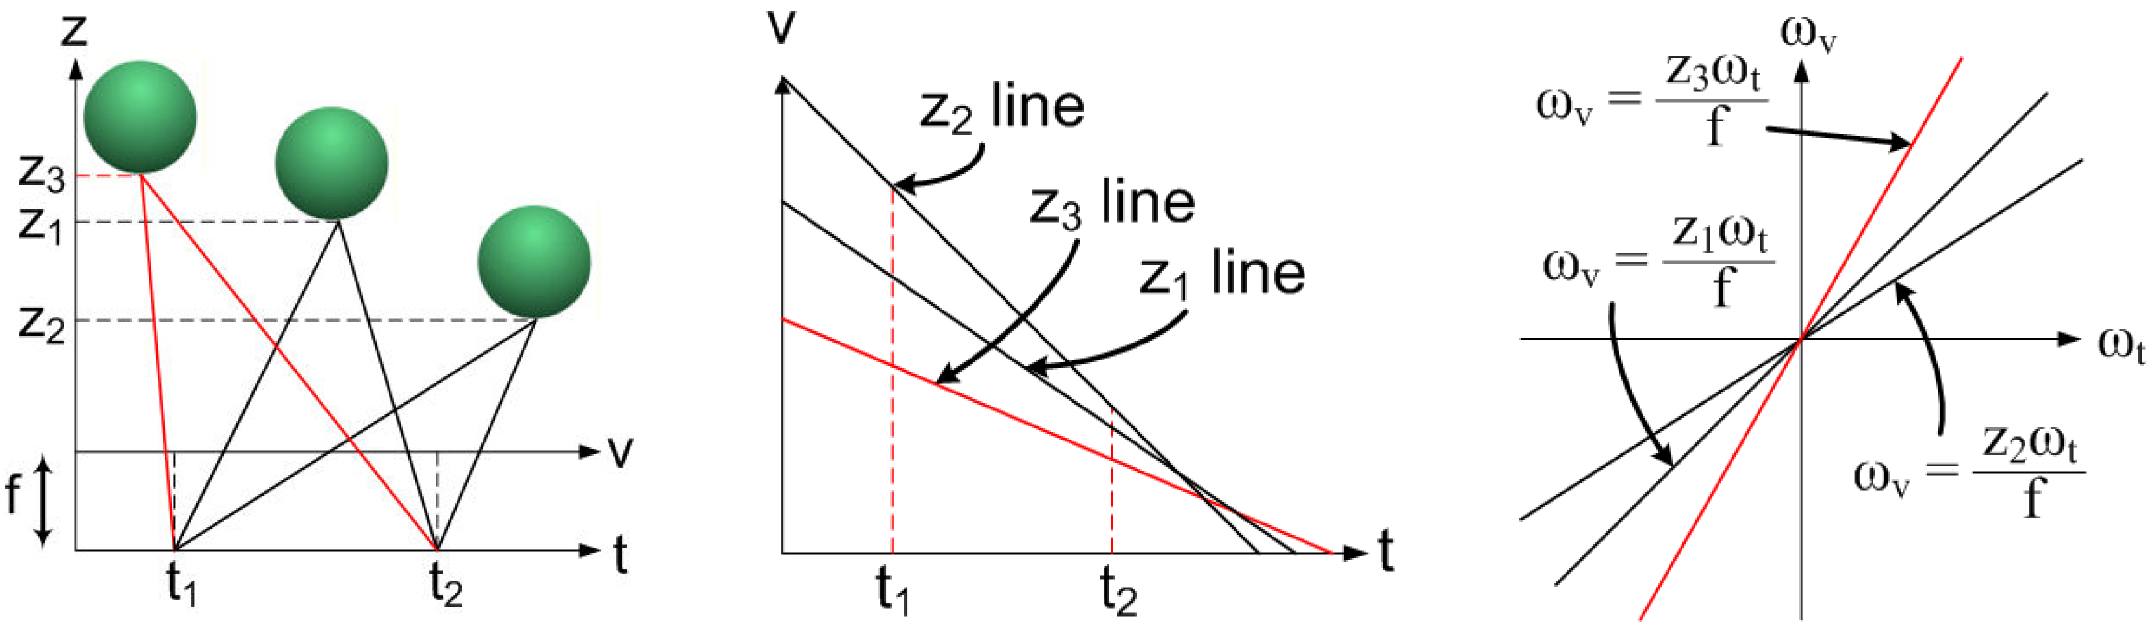

(a) Scene (b) EPI (c) Plenoptic Spectrum

Plenoptic Spectrum exactly bound within two lines relating to the minimum and maximum depths of the scene [3,4]

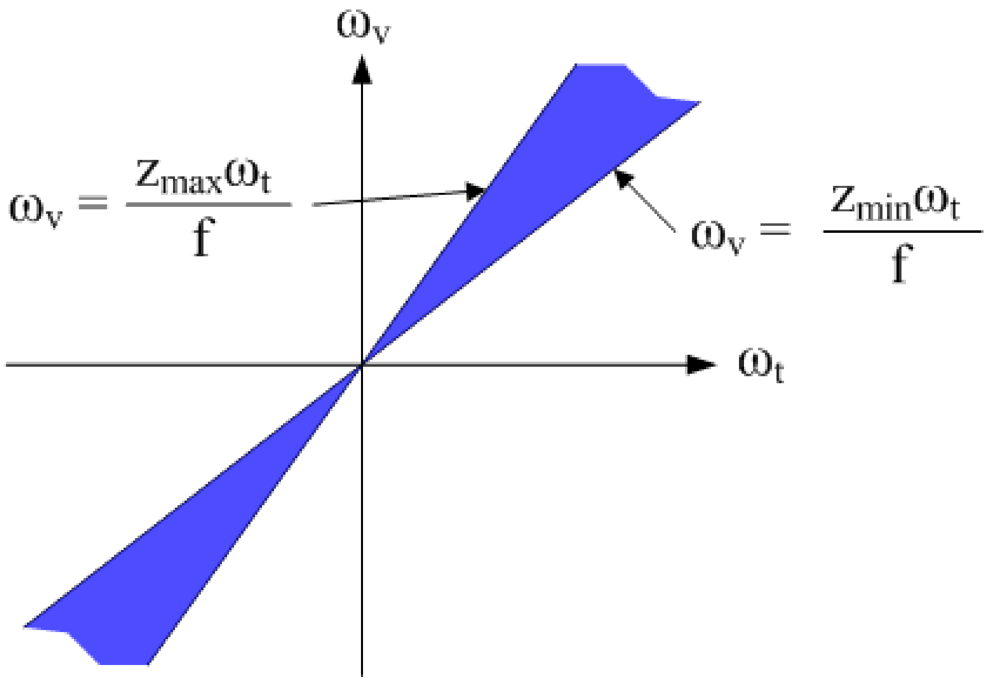

Sampled Plenoptic Spectrum

- Finite Camera Resolution $\Delta v \implies$ Enforced Lowpass Filtering in ω_v
- Sampling in t of period $\Delta t \implies$ Replicated Plenoptic Spectra
- Undersampling $\implies$ Replicated Spectra Overlap $\implies$ Aliasing

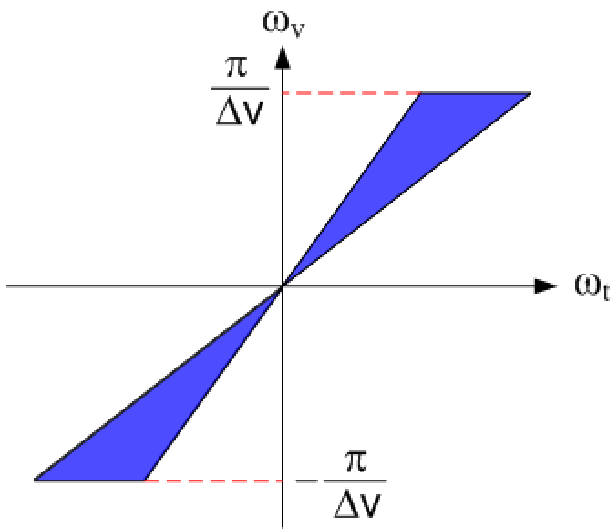

(a) Plenoptic Spectrum Sampled in v

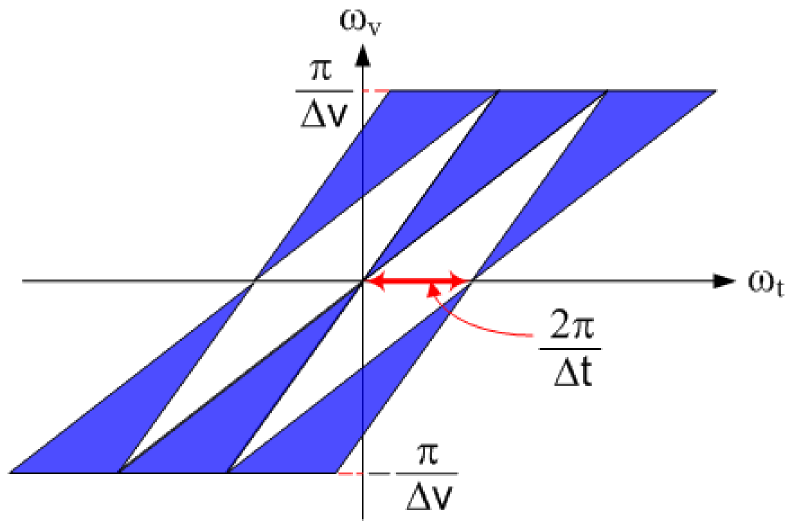

(b) Plenoptic Spectrum Sampled in t

Sampled Plenoptic Spectrum

- Finite Camera Resolution $\Delta v \Longrightarrow$ Enforced Lowpass Filtering in ω_v
- Sampling in t of period $\Delta t \Longrightarrow$ Replicated Plenoptic Spectra
- Undersampling $\Longrightarrow$ Replicated Spectra Overlap $\Longrightarrow$ Aliasing

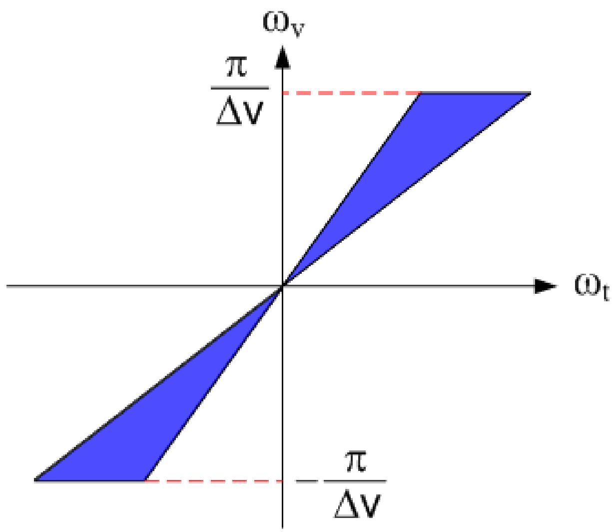

(a) Plenoptic Spectrum Sampled in v

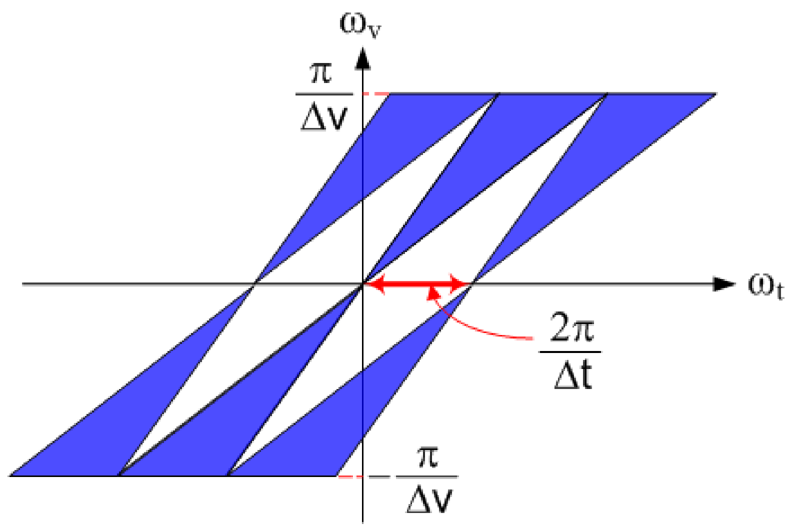

(b) Plenoptic Spectrum Sampled in t

Assumes $\Longrightarrow$ Infinite Scene Width and Infinite Field of View (FoV),

$\Longrightarrow$ Uniform Sampling in t

Slanted Plane Geometry

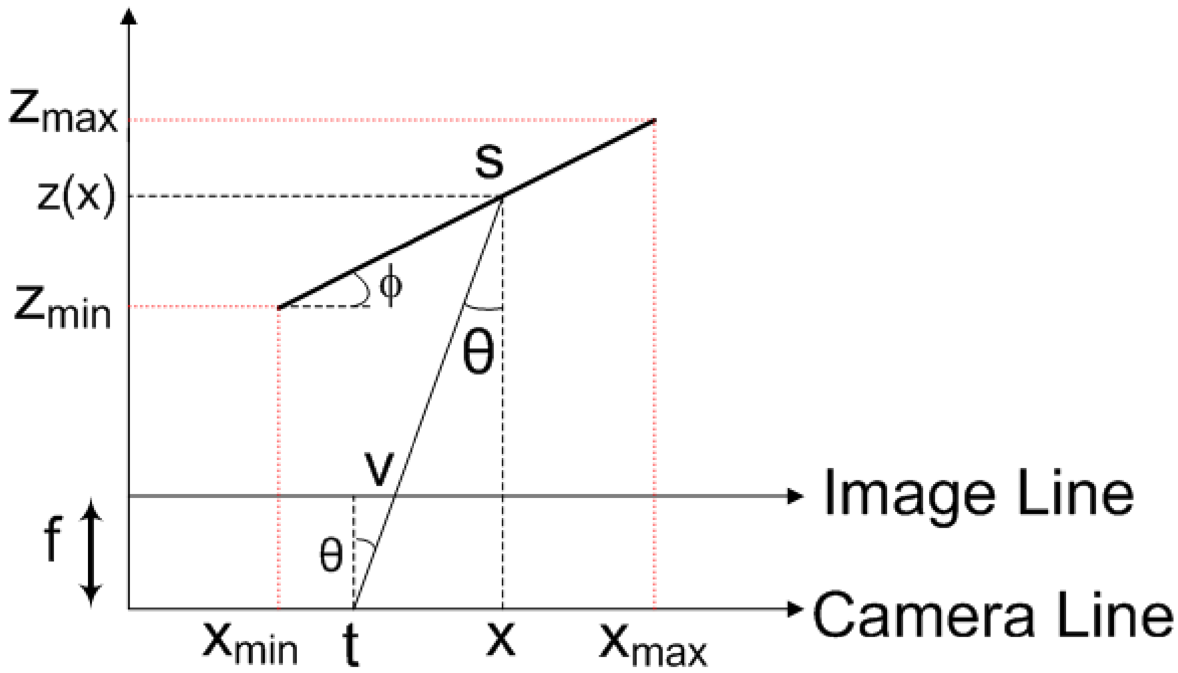

Functional Scene Model [4]:

- s is the Curvilinear Coordinate
- x is the Projection of s onto t
- ϕ is the Slant of the Plane

Sinusoidal Texture Signal Pasted to Scene Surface

Constraints:

- Finite Field of View (FoV) for the Cameras $\Longrightarrow v \in [-v_m, v_m]$
- Finite Plane Width $\Longrightarrow s \in [0, T]$
- Lambertian Scene

Effect of the Constraints on the EPI

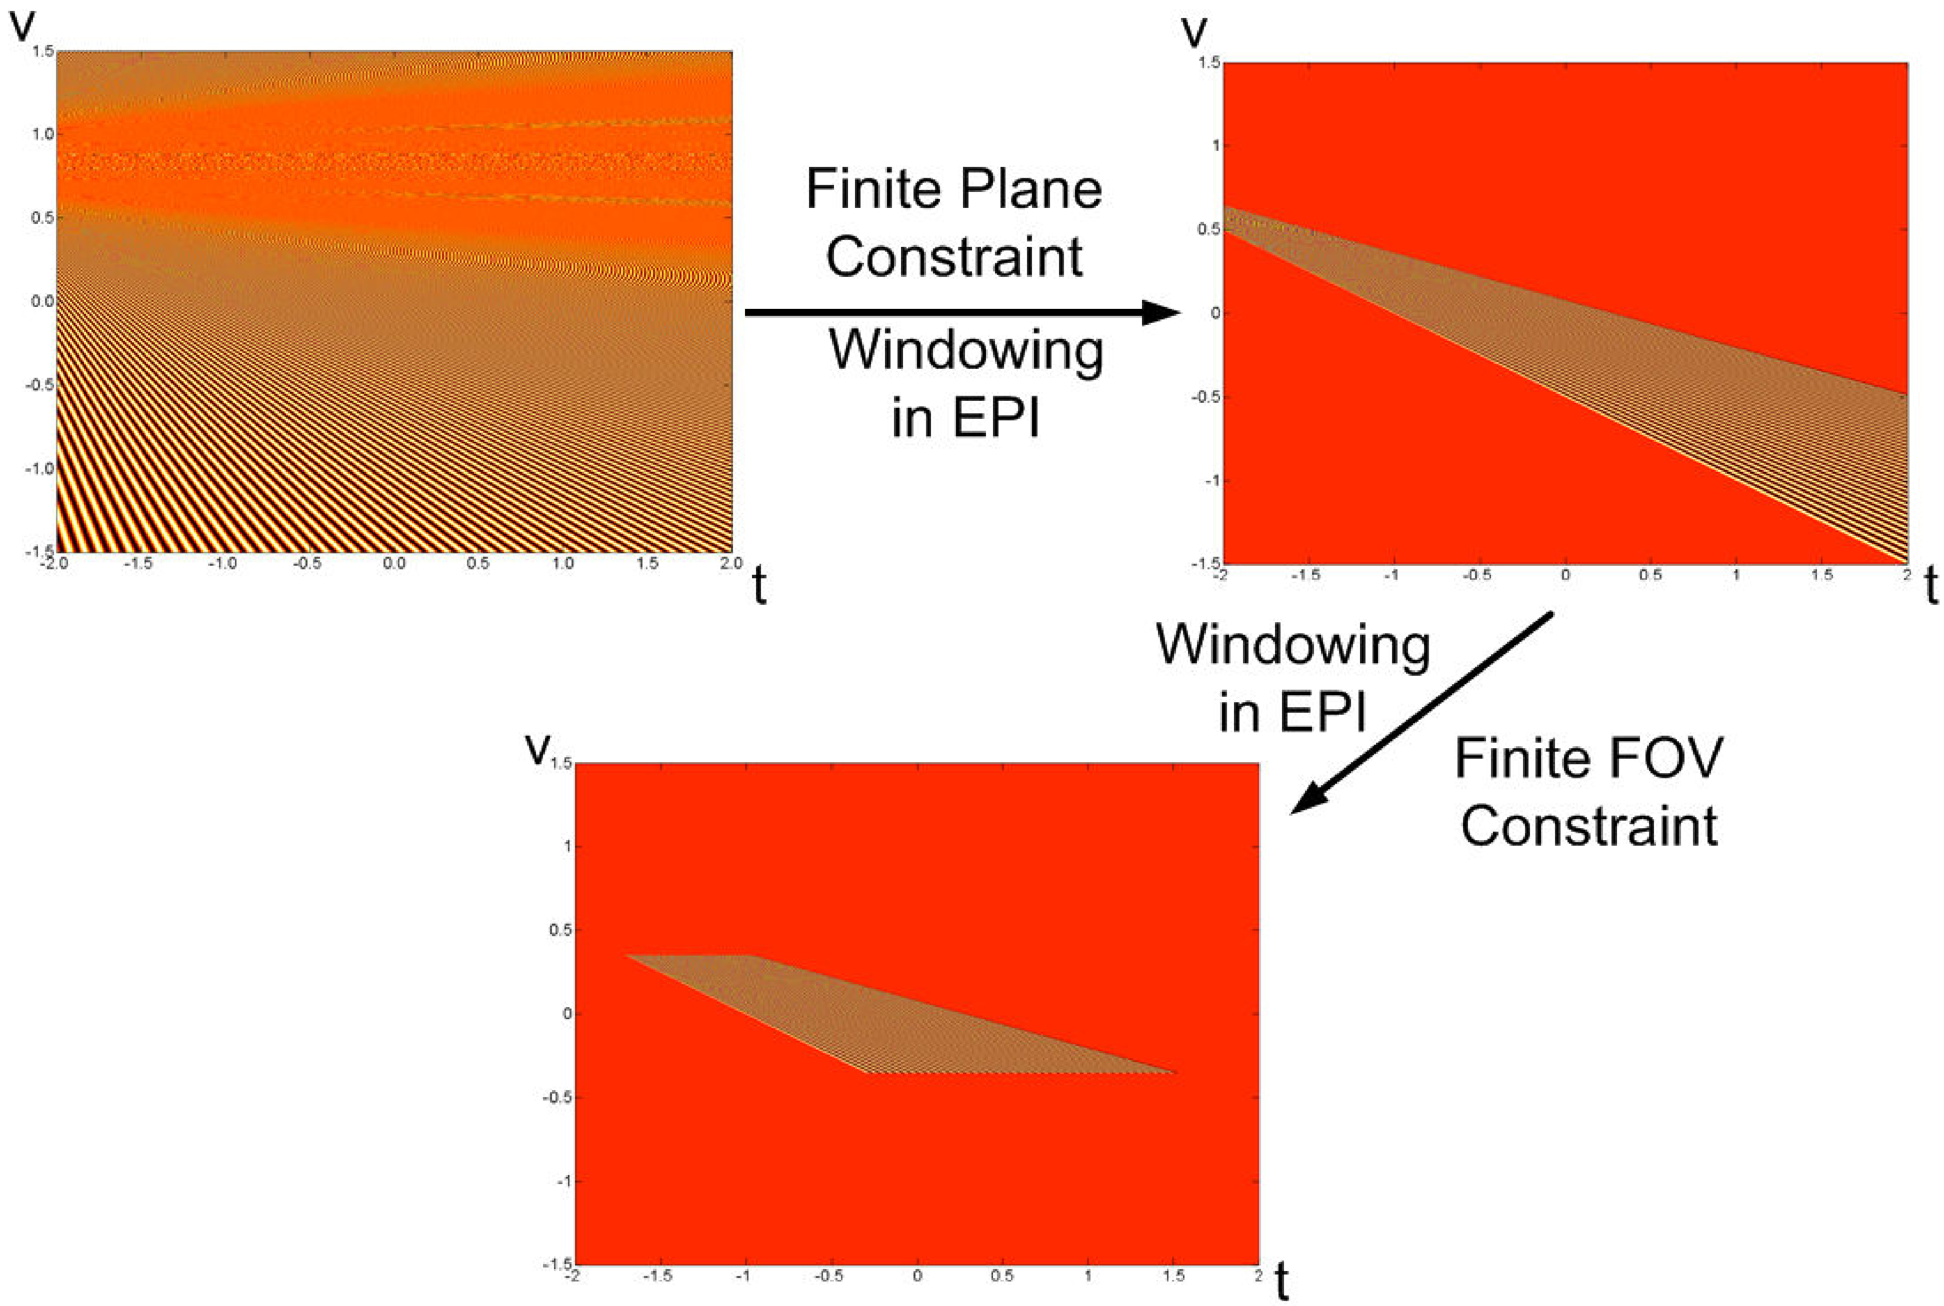

Slanted Plane Plenoptic Spectrum

Plenoptic spectrum is band-unlimited in both ω_t and ω_v

↬ Using our closed-form expression [5], characterise the plenoptic spectrum using 6 lines.

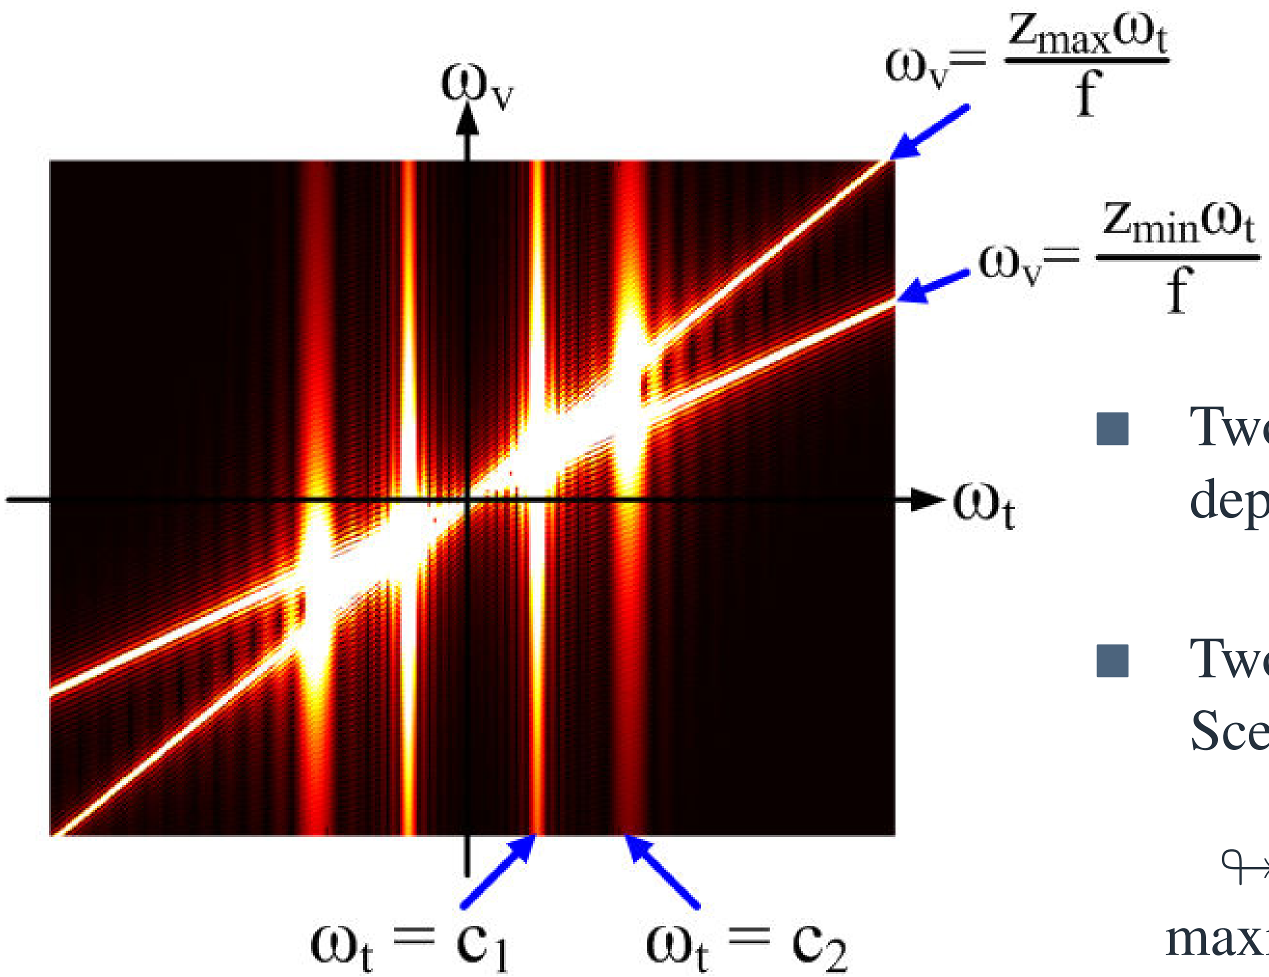

- Two lines ⇒ Maximum and minimum depths of the plane.
- Two pairs of lines ⇒ Finite (FoV) and Scene Geometry.

↬ c_1 and c_2 are the minimum and maximum modulated frequency of the texture

Essential Bandwidth for a Slanted Plane

Parametric model of the essential bandwidth, comprising 4 parameters:

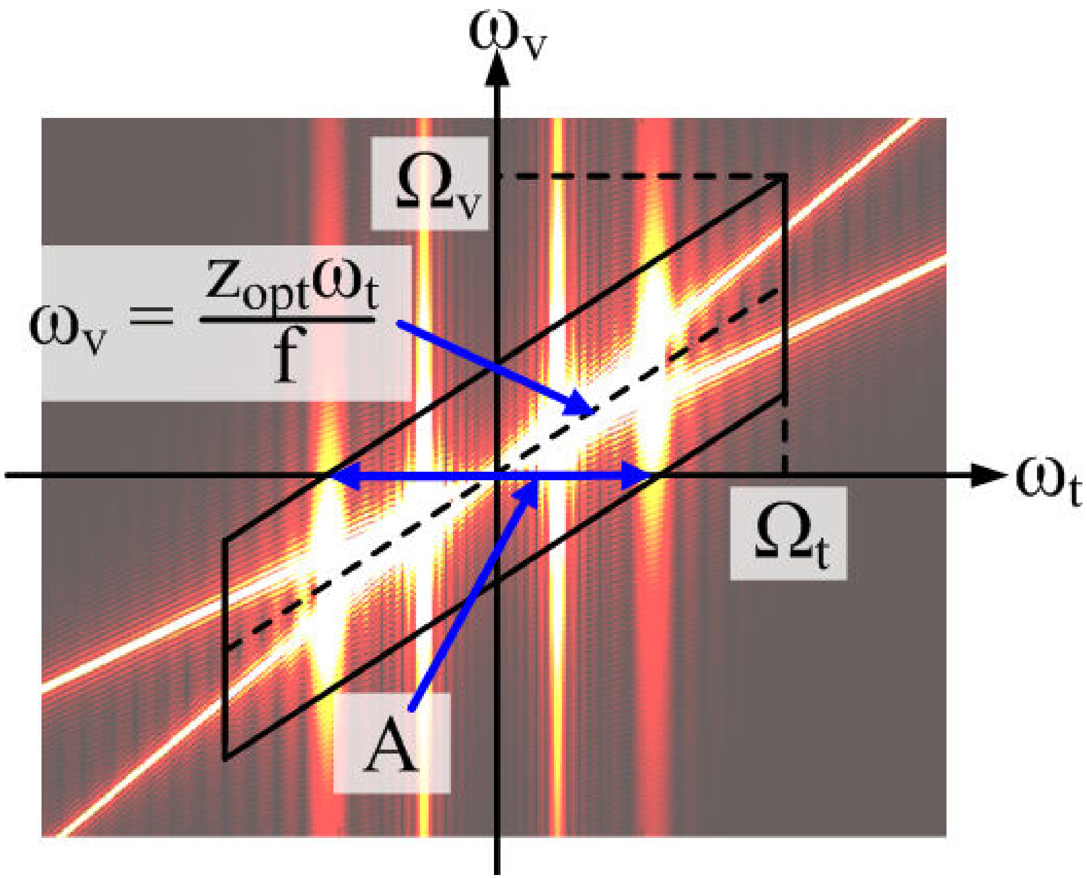

$$\Omega_t = \frac{\omega_s f}{f\cos(\phi) - v_m|\sin(\phi)|} + \frac{2\pi}{T},$$

$$\Omega_v = \Omega_t \frac{z_{max}}{f} + \frac{\pi}{v_m},$$

$$z_{opt} = \frac{z_{max} + z_{min}}{2},$$

$$A = \frac{T|\sin(\phi)|\Omega_t}{z_{opt}} + \frac{2\pi f}{z_{opt} v_m}$$

Essential Bandwidth for a Slanted Plane

Maximum spatial sampling period:

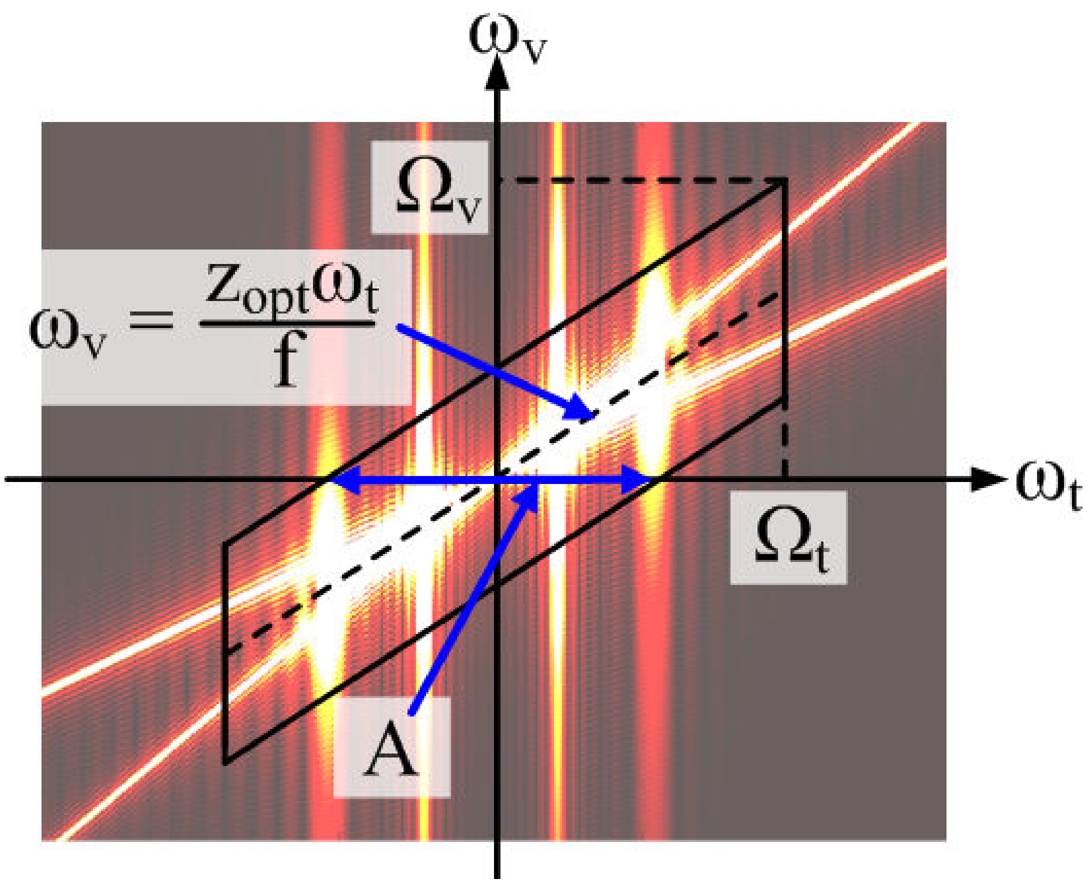

$$\Delta t = \frac{2\pi}{A},$$
$$= \frac{2\pi z_{opt} v_m}{v_m \Omega_t T \left|\sin(\phi)\right| + 2\pi f}$$

Sampling Realistic Scenes

Smoothly varying scene surface with bandlimited texture:

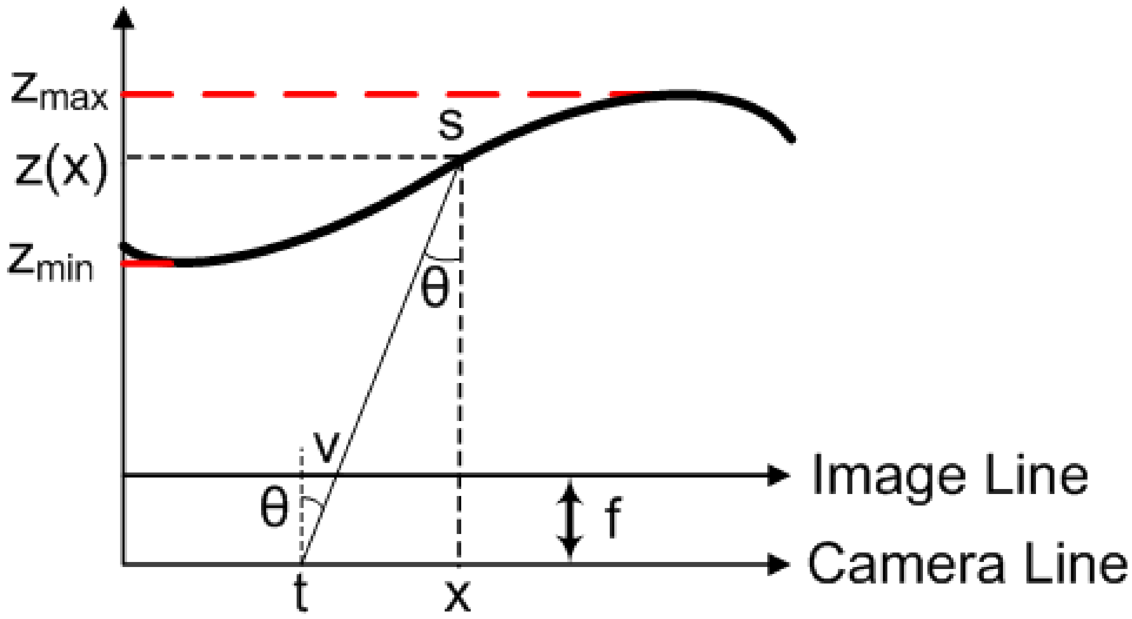

Our approach:

- Adaptively approximate the scene surface using a set of L slanted planes, given N_T samples.
- Determine a piecewise constant sample rate using the previous theory.
- Non-uniformly sample and reconstruct the plenoptic function.

Evaluating the Surface Approximation

Measure the error in the plenoptic domain, comprising two parts:

$$D(N) = \underbrace{\gamma}_{\boxed{\text{Geometric Error}}} + \underbrace{\alpha(N)}_{\text{Aliasing Error}}$$

- Error caused by approximating the scene surface with a set of slanted planes.
- Decreases the more exact the approximation.

Evaluating the Surface Approximation

Measure the error in the plenoptic domain, comprising two parts:

$$D(N) = \underbrace{\gamma}_{\text{Geometric Error}} + \underbrace{\alpha(N)}_{\boxed{\text{Aliasing Error}}}$$

- Error caused by approximating the scene surface with a set of slanted planes.
- Decreases the more exact the approximation.

- Error caused by undersampling the scene with N samples.
- Decreases as the number of samples increases.
- Approximated as twice the energy outside the reconstruction filter in the frequency domain.

Evaluating the Surface Approximation

Measure the error in the plenoptic domain, comprising two parts:

$$D(N) = \underbrace{\gamma}_{\text{Geometric Error}} + \underbrace{\alpha(N)}_{\boxed{\text{Aliasing Error}}}$$

- Error caused by approximating the scene surface with a set of slanted planes.
- Decreases the more exact the approximation.

- Error caused by undersampling the scene with N samples.
- Decreases as the number of samples increases.
- Approximated as twice the energy outside the reconstruction filter in the frequency domain.

Small number of samples $\Rightarrow$ Surface approximation controlled by Aliasing Error

Large number of samples $\Rightarrow$ Surface approximation controlled by Geometric Error

Sample Allocation per Plane

The sample allocation problem is defined in terms minimising the distortion function given N_T samples:

The problem:

$$\min \left\{ \sum_{i=0}^{L} D_i(N_i) \right\} \text{ s.t. } N_T = \sum_{i=0}^{L} N_i,$$

Solve using a Lagrange multiple λ, thus the cost function:

$$\sum_{i=1}^{L} \left(\gamma_i + K_i \left(\frac{16\Delta v}{A_i \pi} + \frac{2A_i \Delta v}{\pi} + \frac{8W_i}{\Delta v\,(N_i - 1)} \right) + \lambda N_i \right)$$

$K_i \implies$ constant for the ith plane.

$W_i \implies$ distance the plane is visible on the camera line.

Sample Allocation per Plane

The sample allocation problem is defined in terms minimising the distortion function given N_T samples:

The problem:

$$\min\left\{\sum_{i=0}^{L} D_i(N_i)\right\} \text{ s.t. } N_T = \sum_{i=0}^{L} N_i,$$

The solution:

$$N_i = \sqrt{\frac{8K_iW_i}{\Delta v \lambda}} + 1, \quad \text{where} \quad \lambda = \frac{\left[\sum_{i=1}^{L}\sqrt{8K_iW_i}\right]^2}{\Delta v\,(N_T - L)^2}$$

Assumes that $N_T > L$ and $N_i \geq 1$, $\forall i$.

$\looparrowright$ so we have an exact solution to the Lagrange multiplier

Optimising the Surface Approximation

Determining the optimum surface approximation:

- Binary-tree approach $\Longrightarrow$ Start with an initial 'fine-grain' approximation.
- Initially split the surface into 2^k equal pieces resulting in L planes.
- Determine the initial λ and sample allocation between the L planes by solving the minimisation problem.
- Merge the leaves of the tree to reduce the overall distortion.

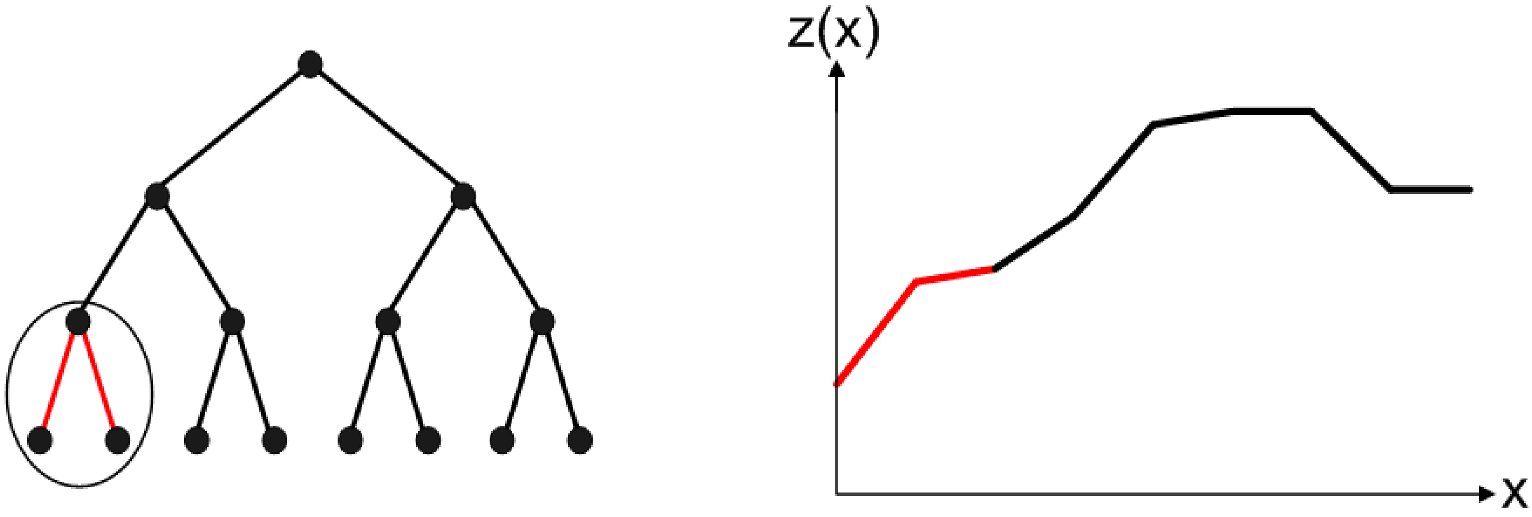

(a) Binary-tree (b) Surface Approximation

Optimising the Surface Approximation

Determining the optimum surface approximation:

- Binary-tree approach $\Longrightarrow$ Start with an initial ‘fine-grain’ approximation.
- Initially split the surface into 2^k equal pieces resulting in L planes.
- Determine the initial λ and sample allocation between the L planes by solving the minimisation problem.
- Merge the leaves of the tree to reduce the overall distortion.

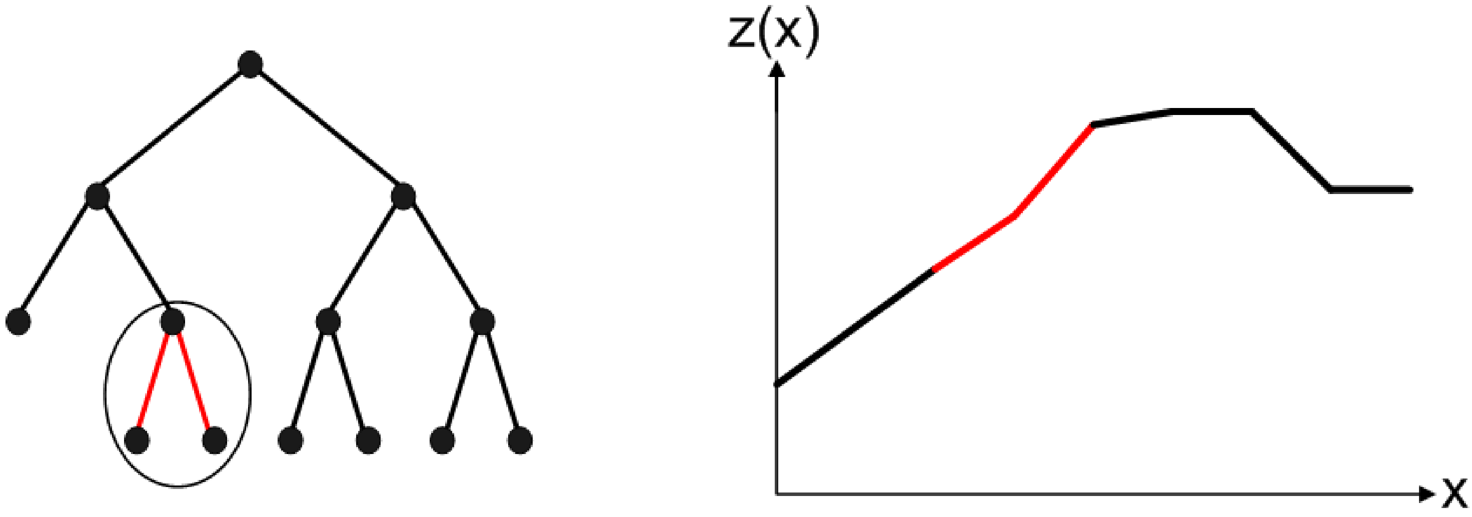

(a) Binary-tree (b) Surface Approximation

Optimising the Surface Approximation

Determining the optimum surface approximation:

- Binary-tree approach $\Longrightarrow$ Start with an initial 'fine-grain' approximation.
- Initially split the surface into 2^k equal pieces resulting in L planes.
- Determine the initial λ and sample allocation between the L planes by solving the minimisation problem.
- Merge the leaves of the tree to reduce the overall distortion.

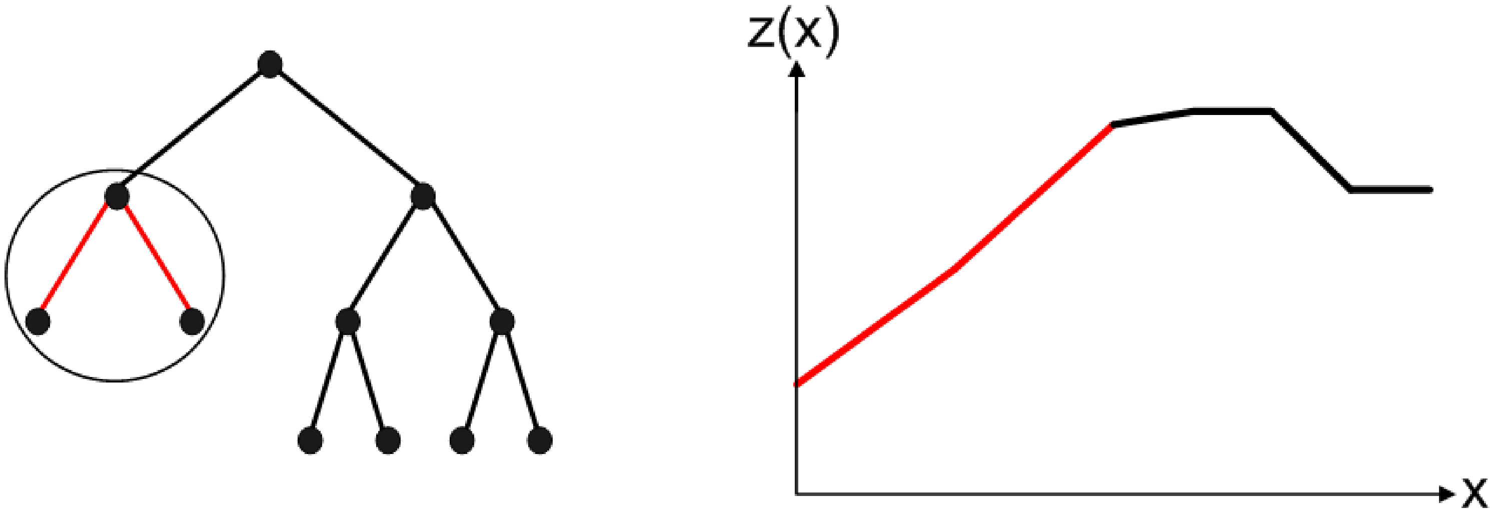

(a) Binary-tree (b) Surface Approximation

Surface Approximation Simulations

Approximation of the Piecewise Quadratic Surface using $N_T = 25$:

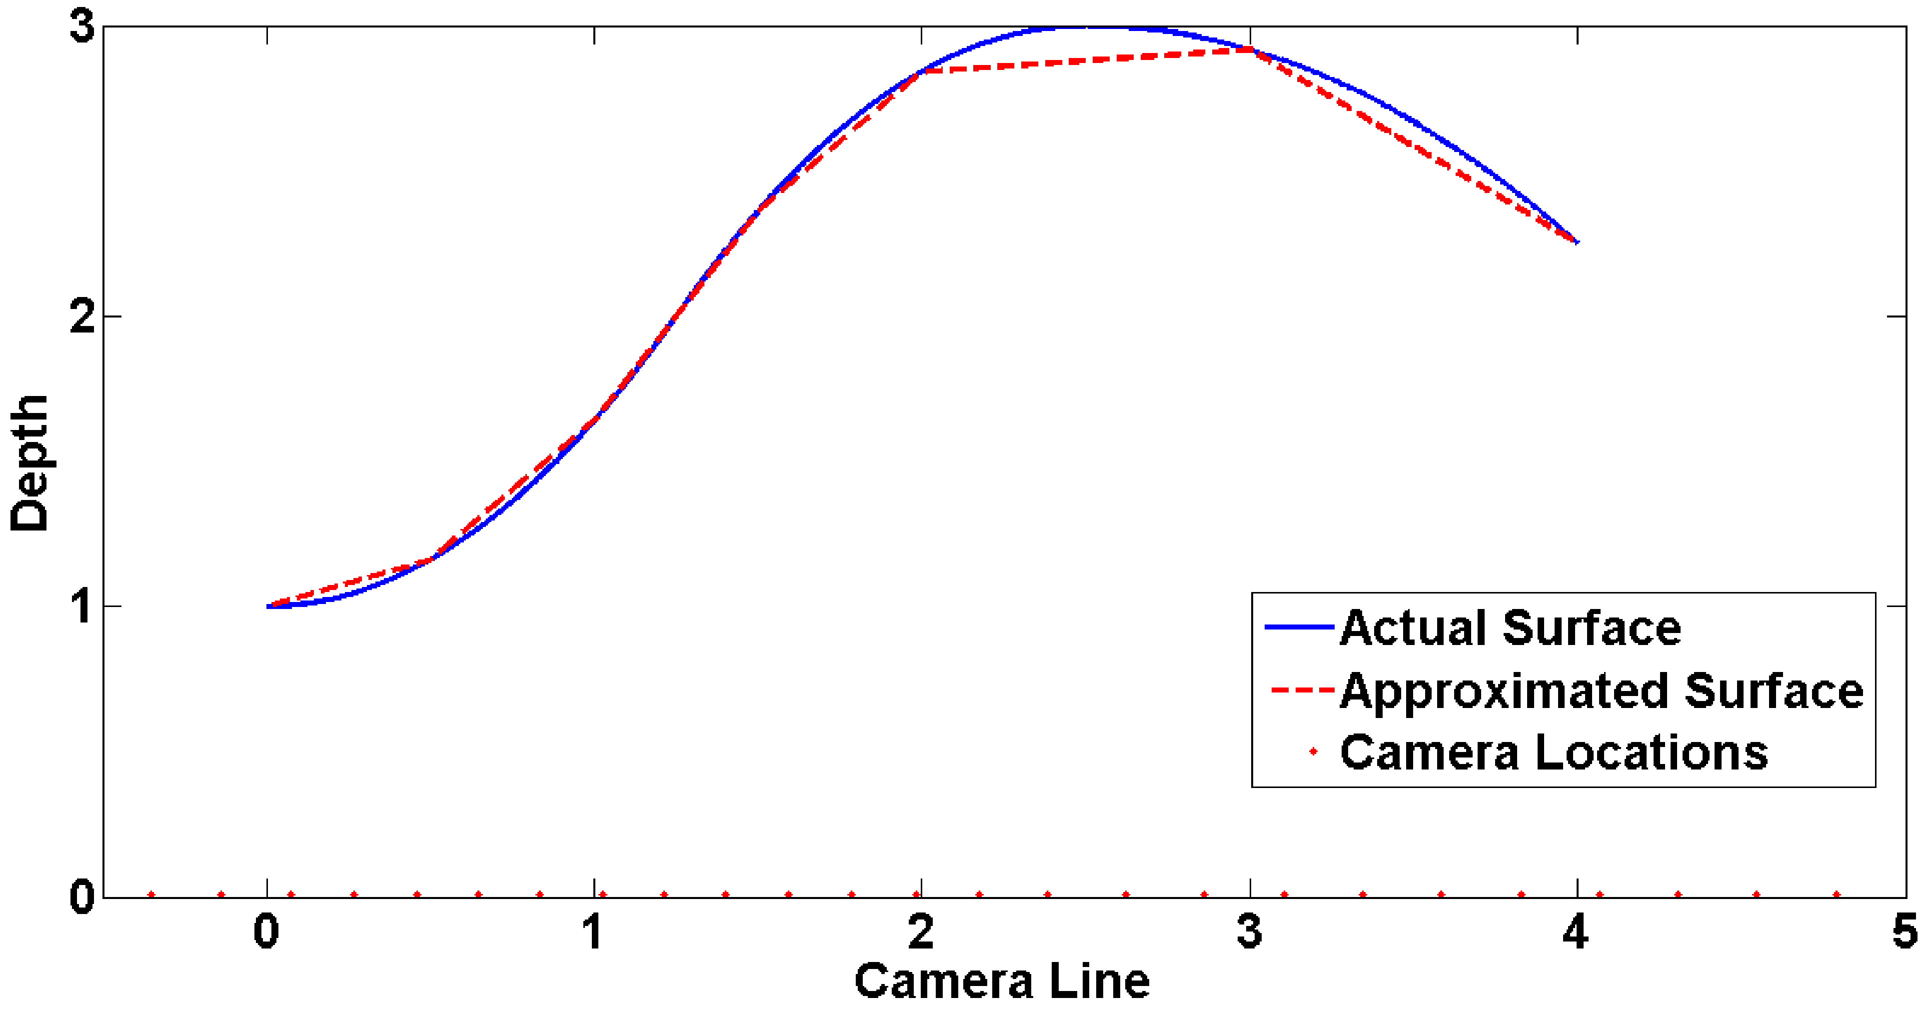

Initial Number of Planes = 16, Final Number of Planes = 6

Surface Approximation Simulations

Approximation of the Piecewise Quadratic Surface using $N_T = 130$:

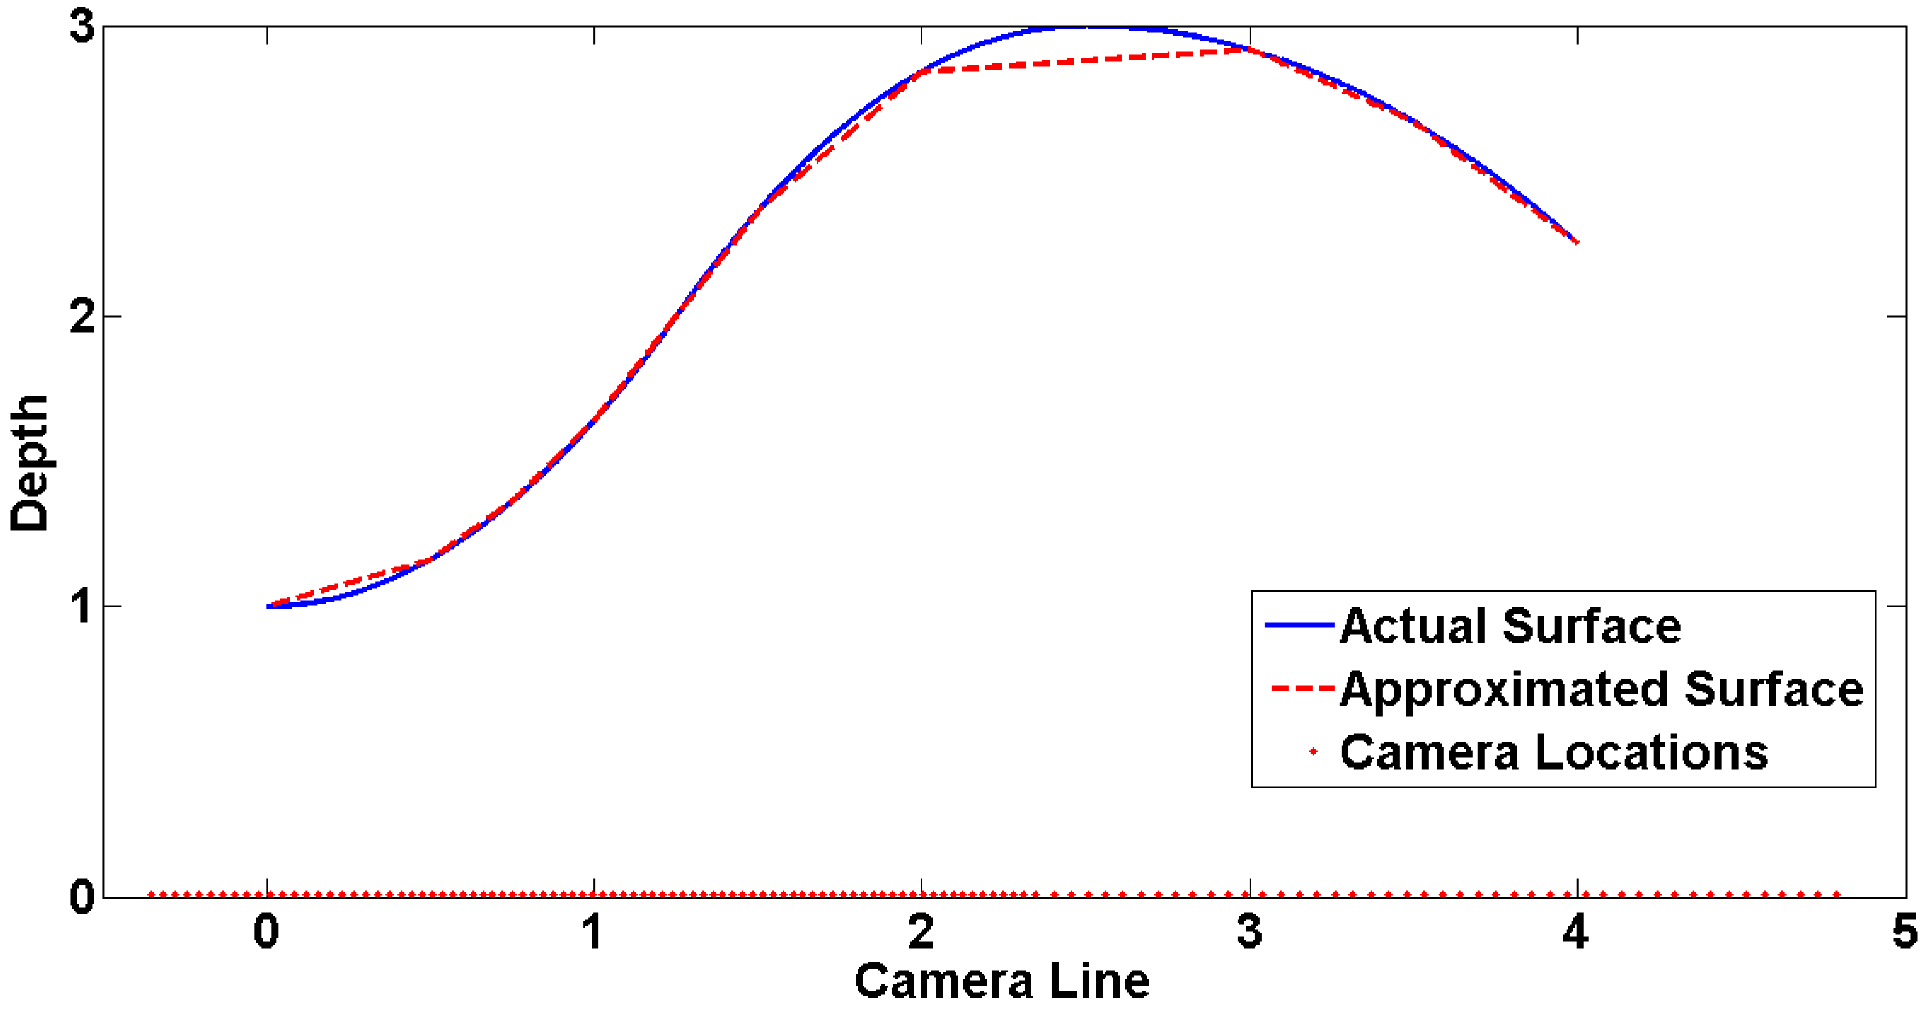

Initial Number of Planes = 16, Final Number of Planes = 8

Surface Approximation Simulations

Approximation of the Piecewise Quadratic Surface using $N_T = 275$:

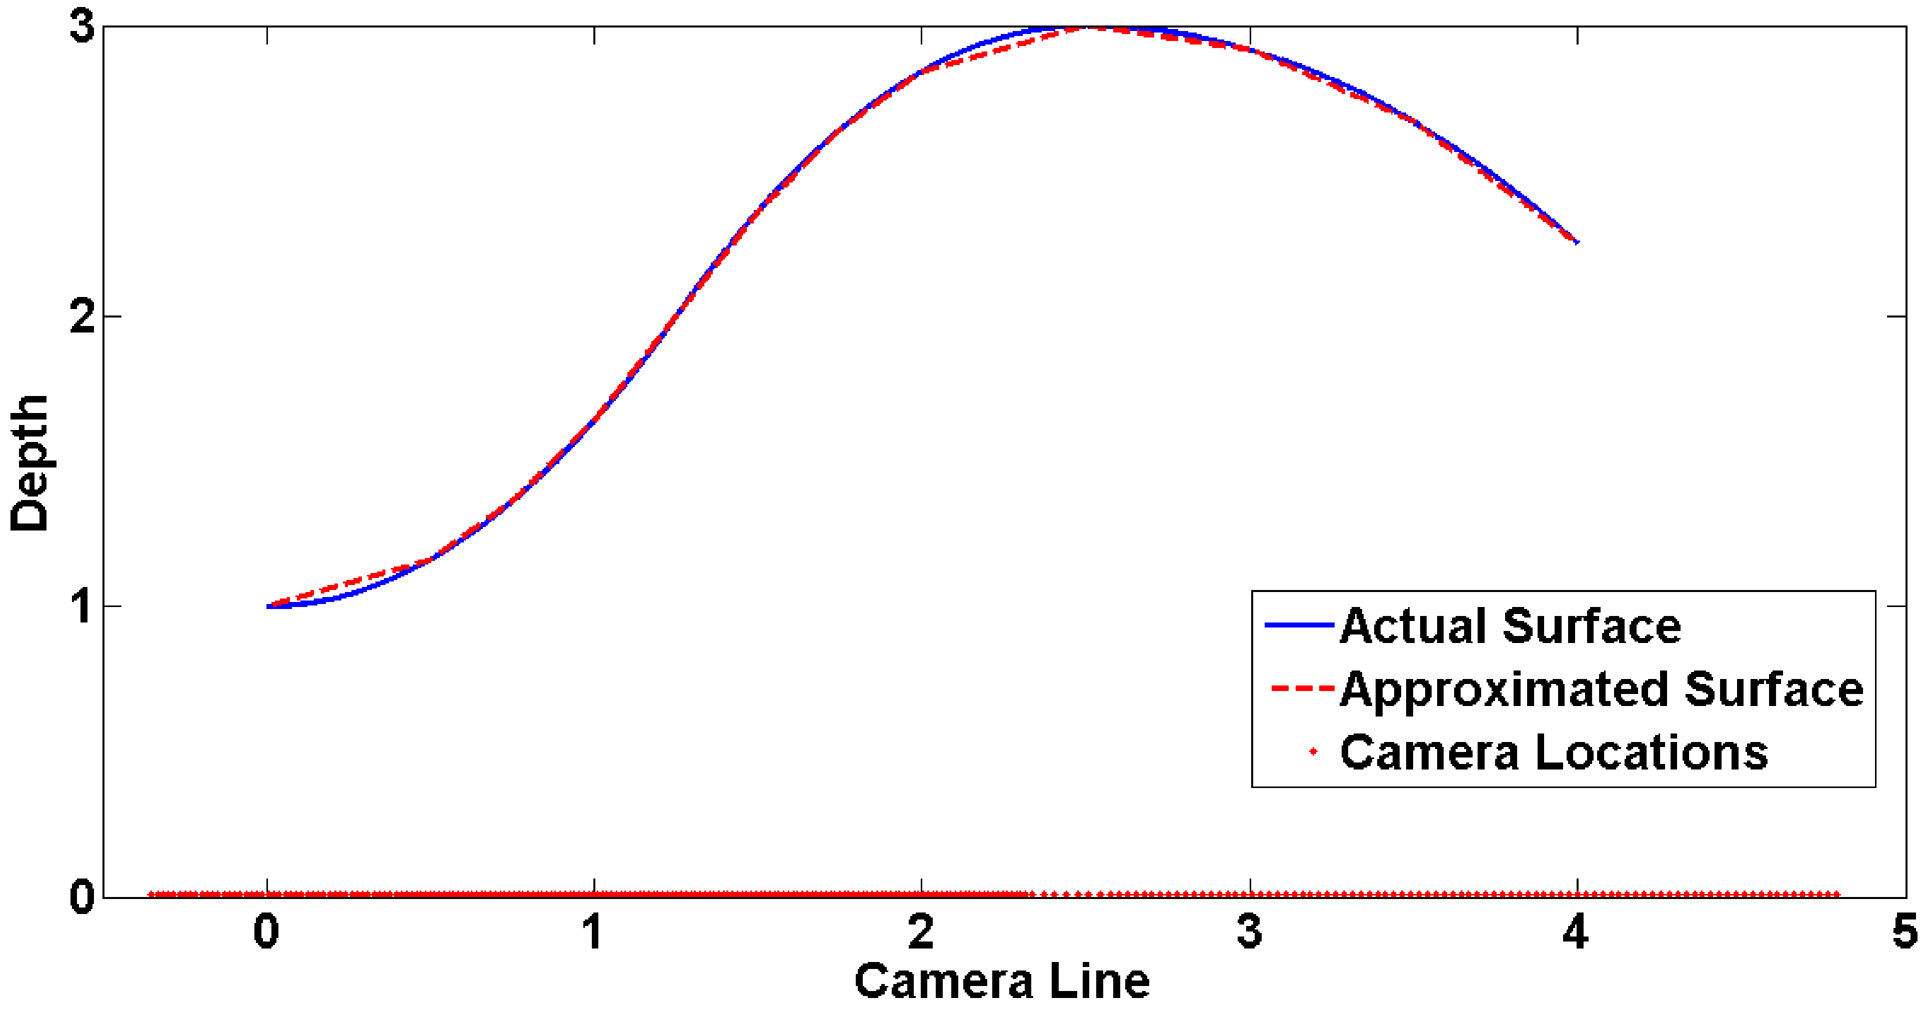

Initial Number of Planes = 16, Final Number of Planes = 10

Simulations Results

Comparison between uniform and adaptive reconstruction for the piecewise quadratic surface.

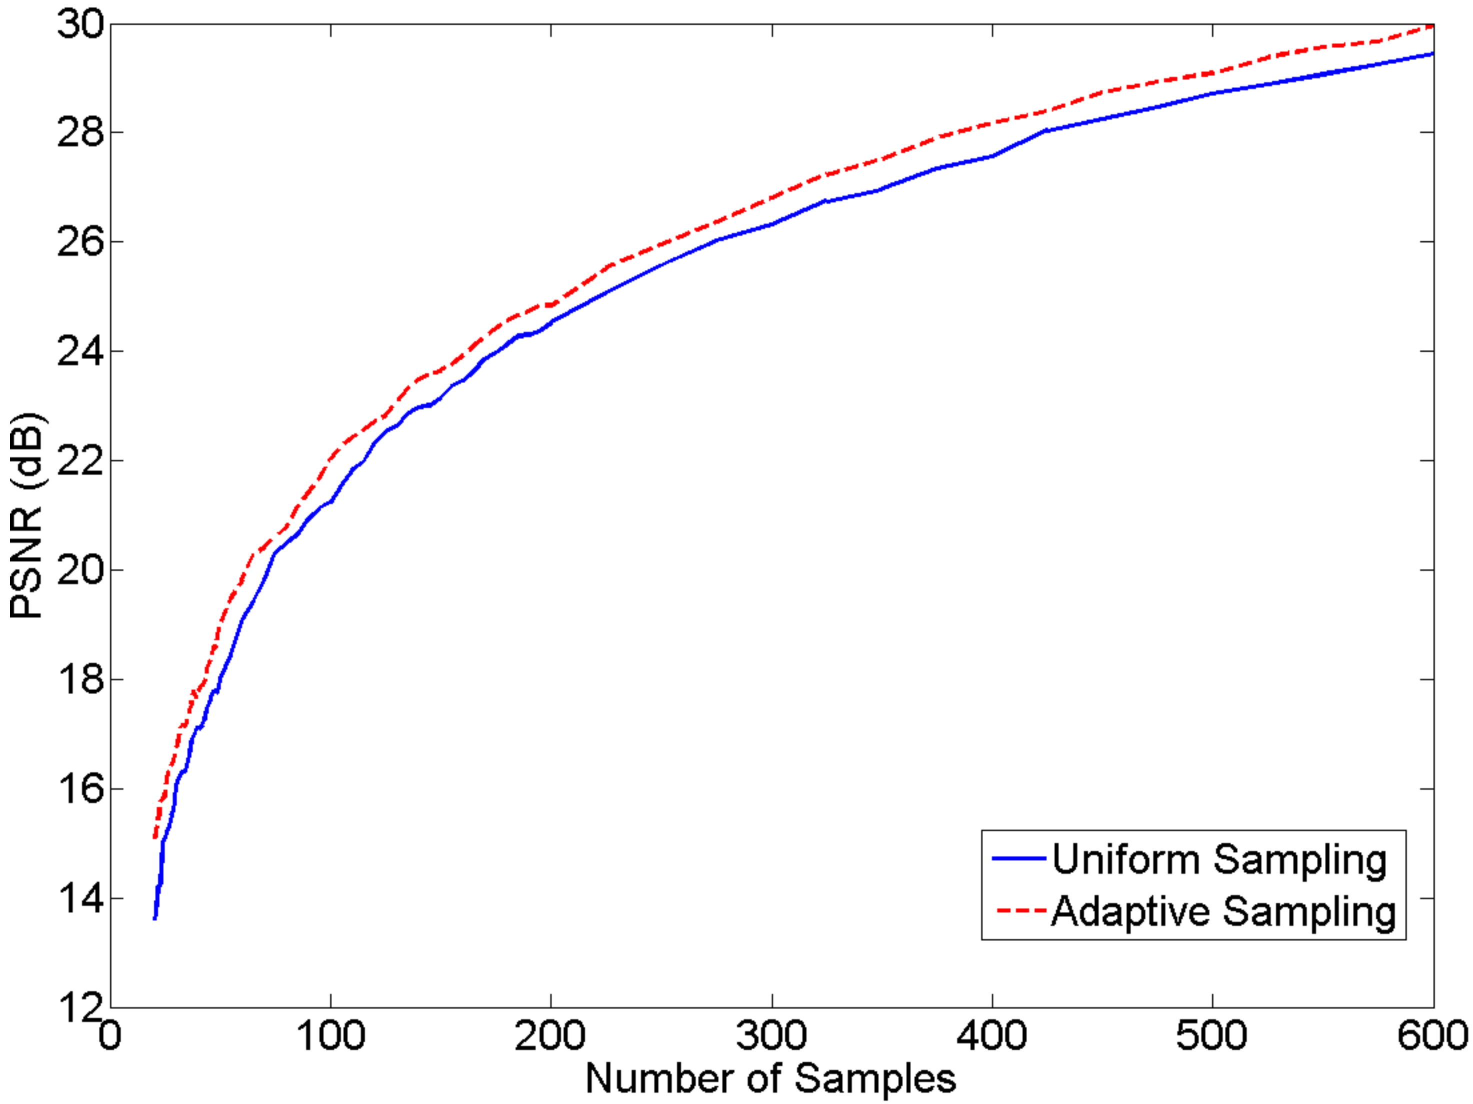

Applied to Real Images

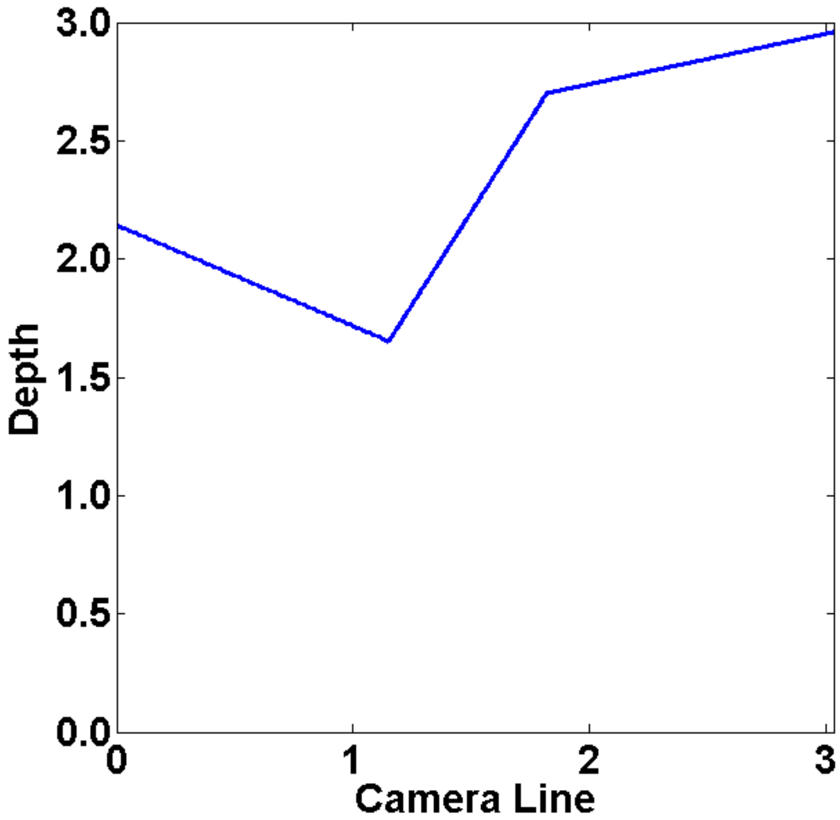

(a) Scene Geometry

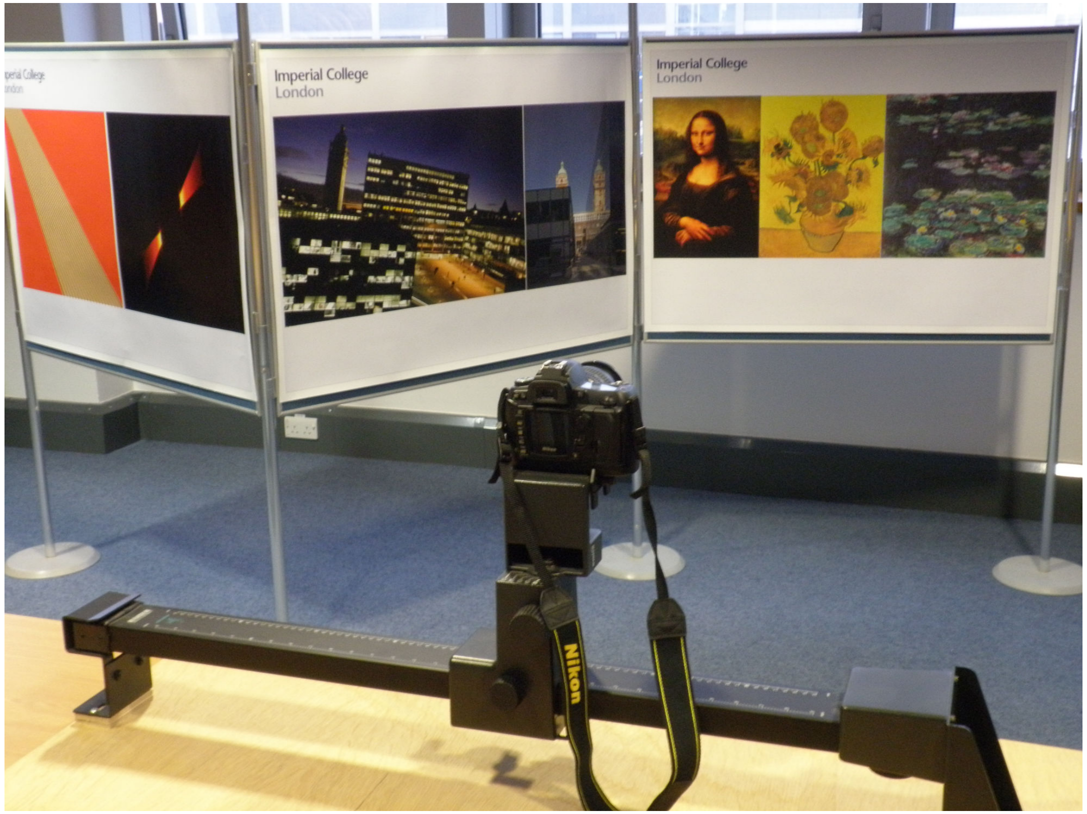

(b) Acquiring the Data

$\looparrowright$ Initial Image Set = 253 Images (1cm apart)

Applied to Real Images

Comparison, in PSNR, for the reconstruction of the plenoptic function when sampled uniformly and adaptively.

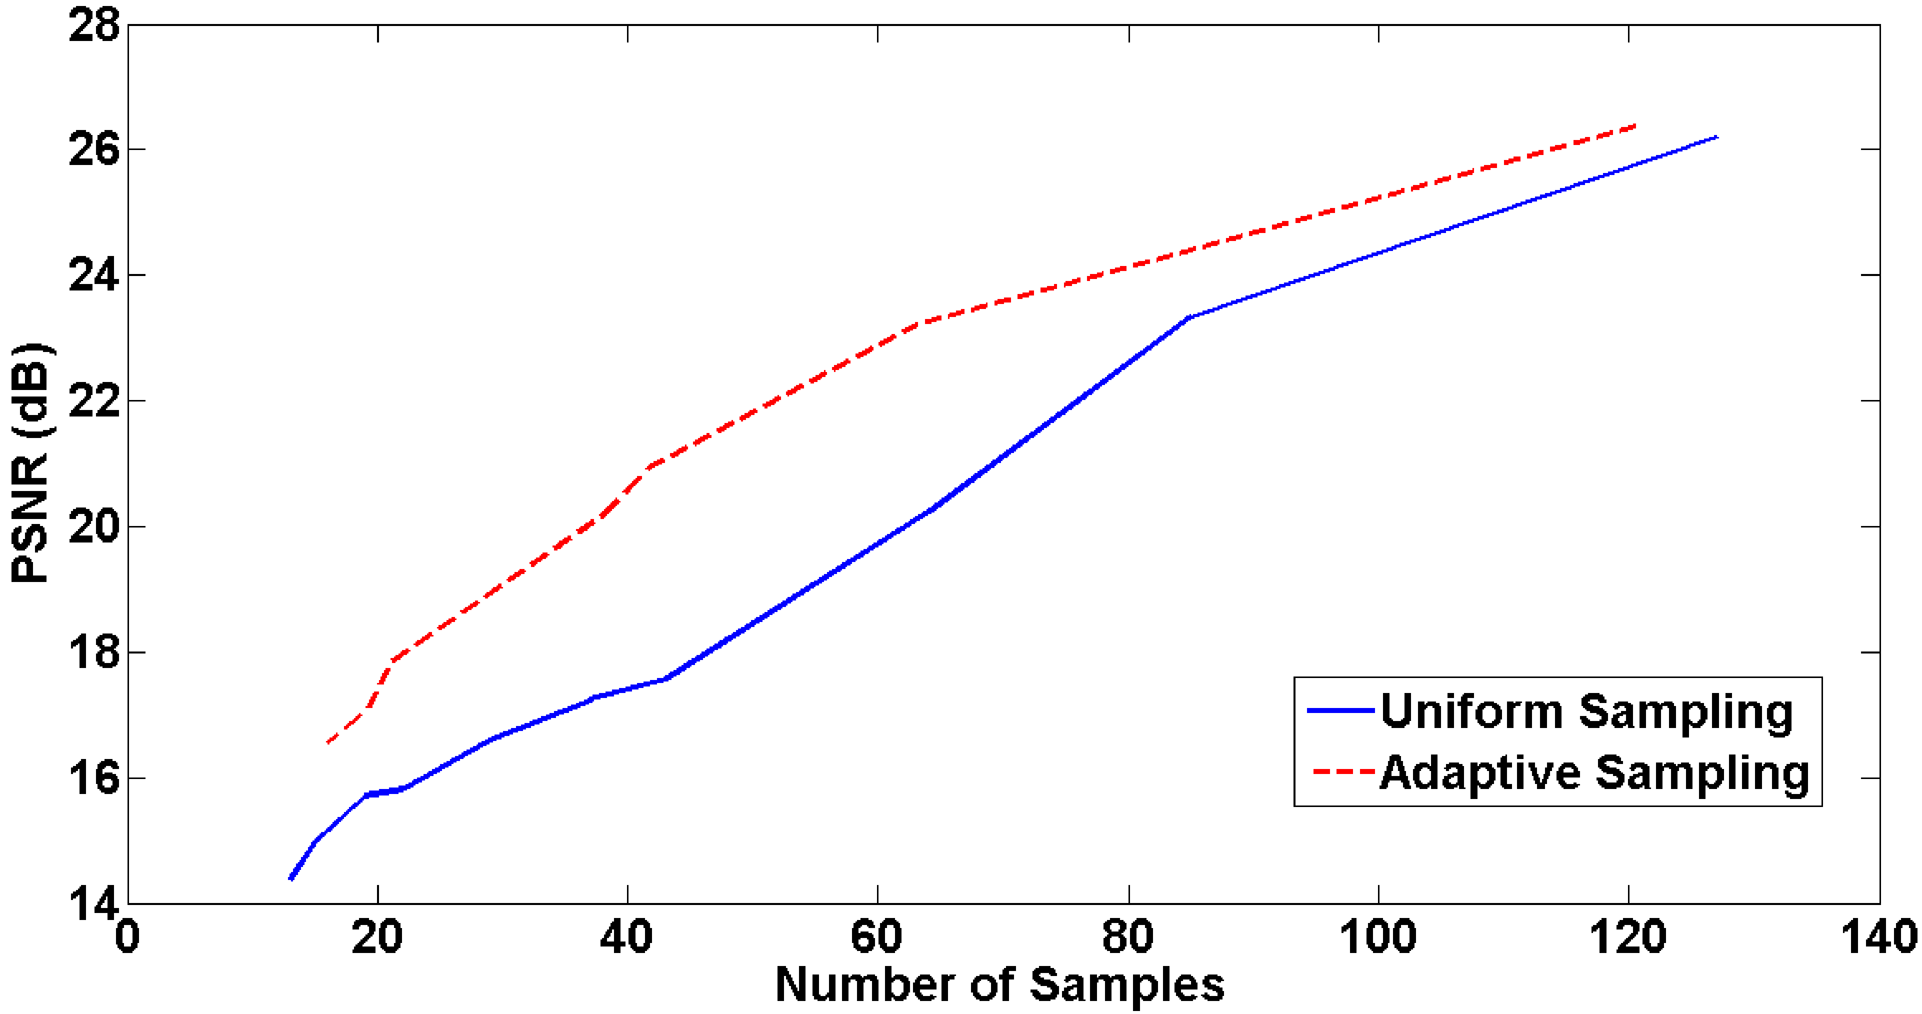

Applied to Real Images

Example of a rendered image:

↬ Uses 85 Samples

(a) Original Image

(b) Uniform Sampling

(c) Adaptive Sampling

Conclusions

- Presented a method for positioning a finite number of samples for a smooth scene based on a spectral analysis of a slanted plane.
- The smooth surface is approximated by a set of slanted planes and the samples are allocated to minimise the distortion, using a Lagrange multiplier.
- The surface approximation is optimised in a binary-tree framework and adapts given the number of samples available.
- Non-uniform sampling scheme outperforms normal uniform sampling.

References

1. J. Pearson, P.L. Dragotti and M. Brookes. Accurate non-iterative depth layer extraction algorithm for image based rendering. IEEE ICASSP 2011, pages 901-904.
2. E.H. Adelson and J.R. Bergen. The plenoptic function and the elements of early vision. In Computational Models of Visual Processing, pages 3-20. MIT Press, Cambridge, MA, 1991.
3. J.X. Chai, S.C. Chan, H.Y. Shum, and X. Tong. Plenoptic sampling. In Computer graphics (SIGGRAPH'00), pages 307-318, 2000.
4. M.N. Do, D Marchand-Maillet, and M. Vetterli. On the bandwidth of the plenoptic function. IEEE Transactions on Image Processing, 2011. Preprint.
5. C. Gilliam, P.L. Dragotti, and M. Brookes. A closed-form expression for the bandwidth of the plenoptic function under finite field of view constraints. IEEE ICIP 2010, pages 3965-3968.

Non-uniform Sampling and Reconstruction

Generate a piecewise constant sample rate profile in t using the sample allocation:

- Choose the highest sample rate in an overlap region.
- Determine the non-uniform sample locations.

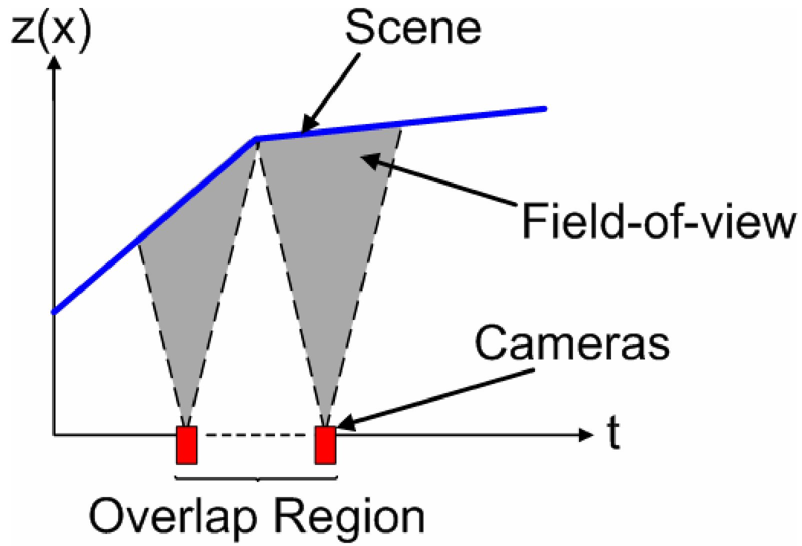

Reconstruction the plenoptic function:

- Split into regions with constant sample rate and reconstruct separately.
- Combine each region using a local interpolation to smooth the transition from one rate to another.
- Local interpolation is based on time-warp sampling theory.