

ACCOUNTING FOR UNCERTAINTY IN MEDICAL DATA: A CUDA IMPLEMENTATION OF NORMALIZED CONVOLUTION

Stefan Lindholm and Joel Kronander

Scientific Visualization Group
Linköping University



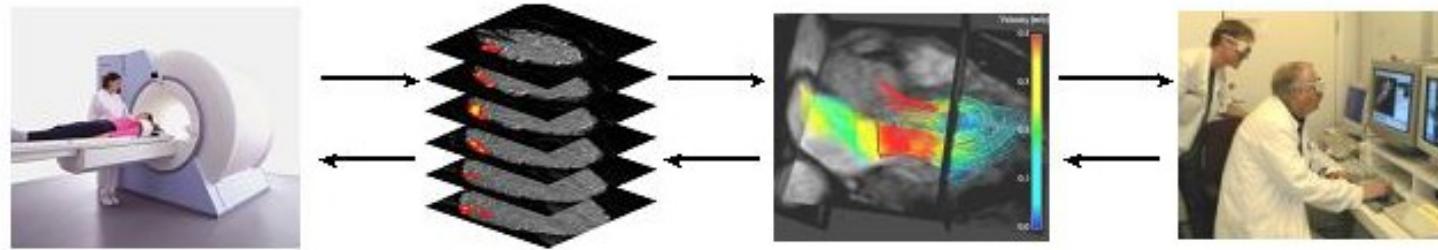
Linköping University
expanding reality



**Edge Detection
Bilateral Filtering
Transfer Functions
Gradient Estimation**

FILTERING MEDICAL IMAGES

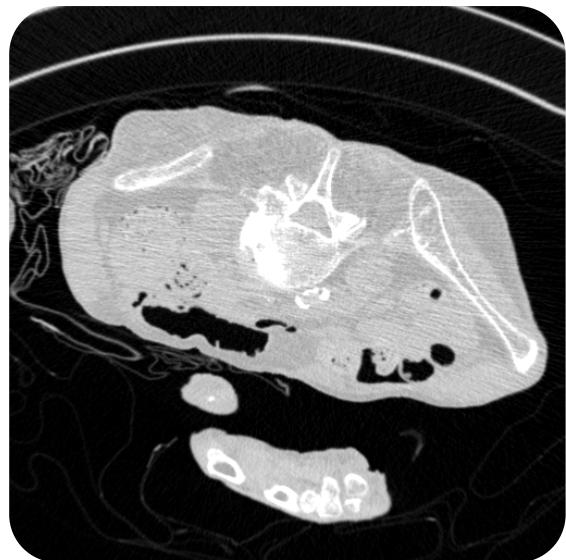




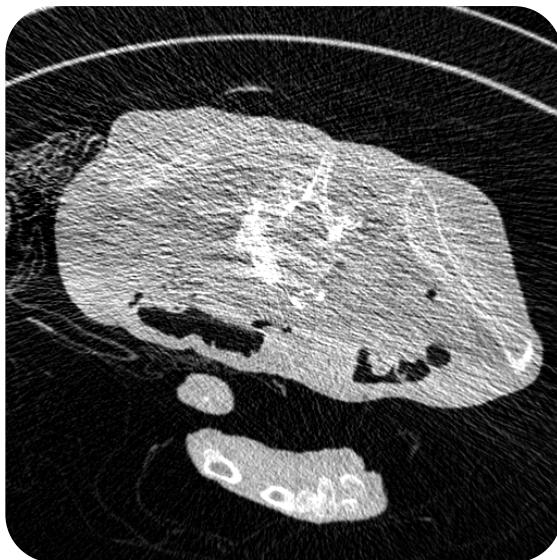
UNCERTAINTY



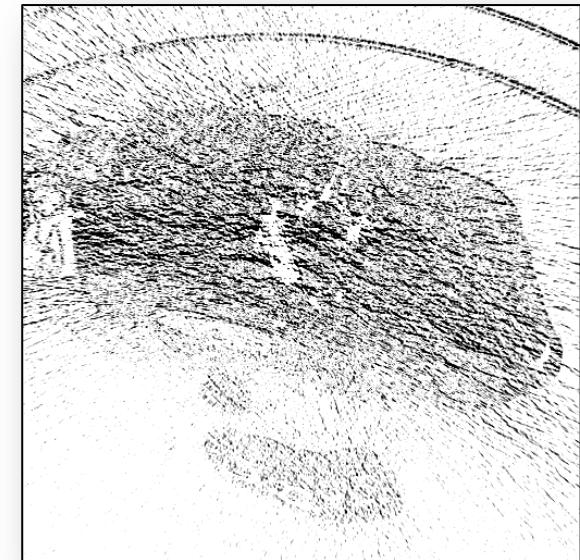
KNOWN UNCERTAINTY



180 mAs CT



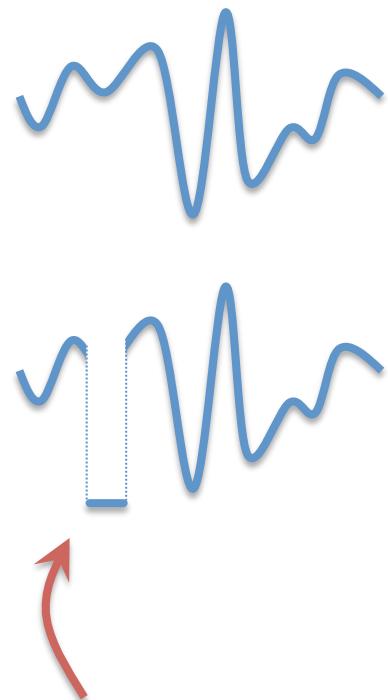
12 mAs CT



Notation: $s(k)$ $r(k)$



KNOWN UNCERTAINTY



Filtering a known signal is straightforward
..but..

How do we treat uncertainty?

Local filter kernels, adaptive to the local
signal uncertainty!

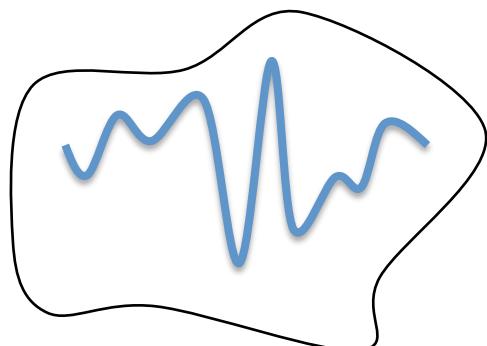
Uncertain sample



THEORY



CONVOLUTION

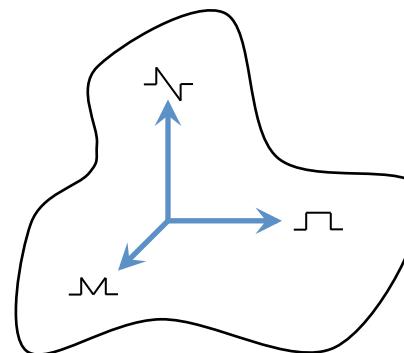


$$h = s * f \quad \otimes$$

Polynomial expansion

$$f_1(l) \quad \vdots \quad f_M(l)$$
$$h(k) = \sum_l s(k + l) \overline{f(-l)}$$

G
Interpret convolution as a projection

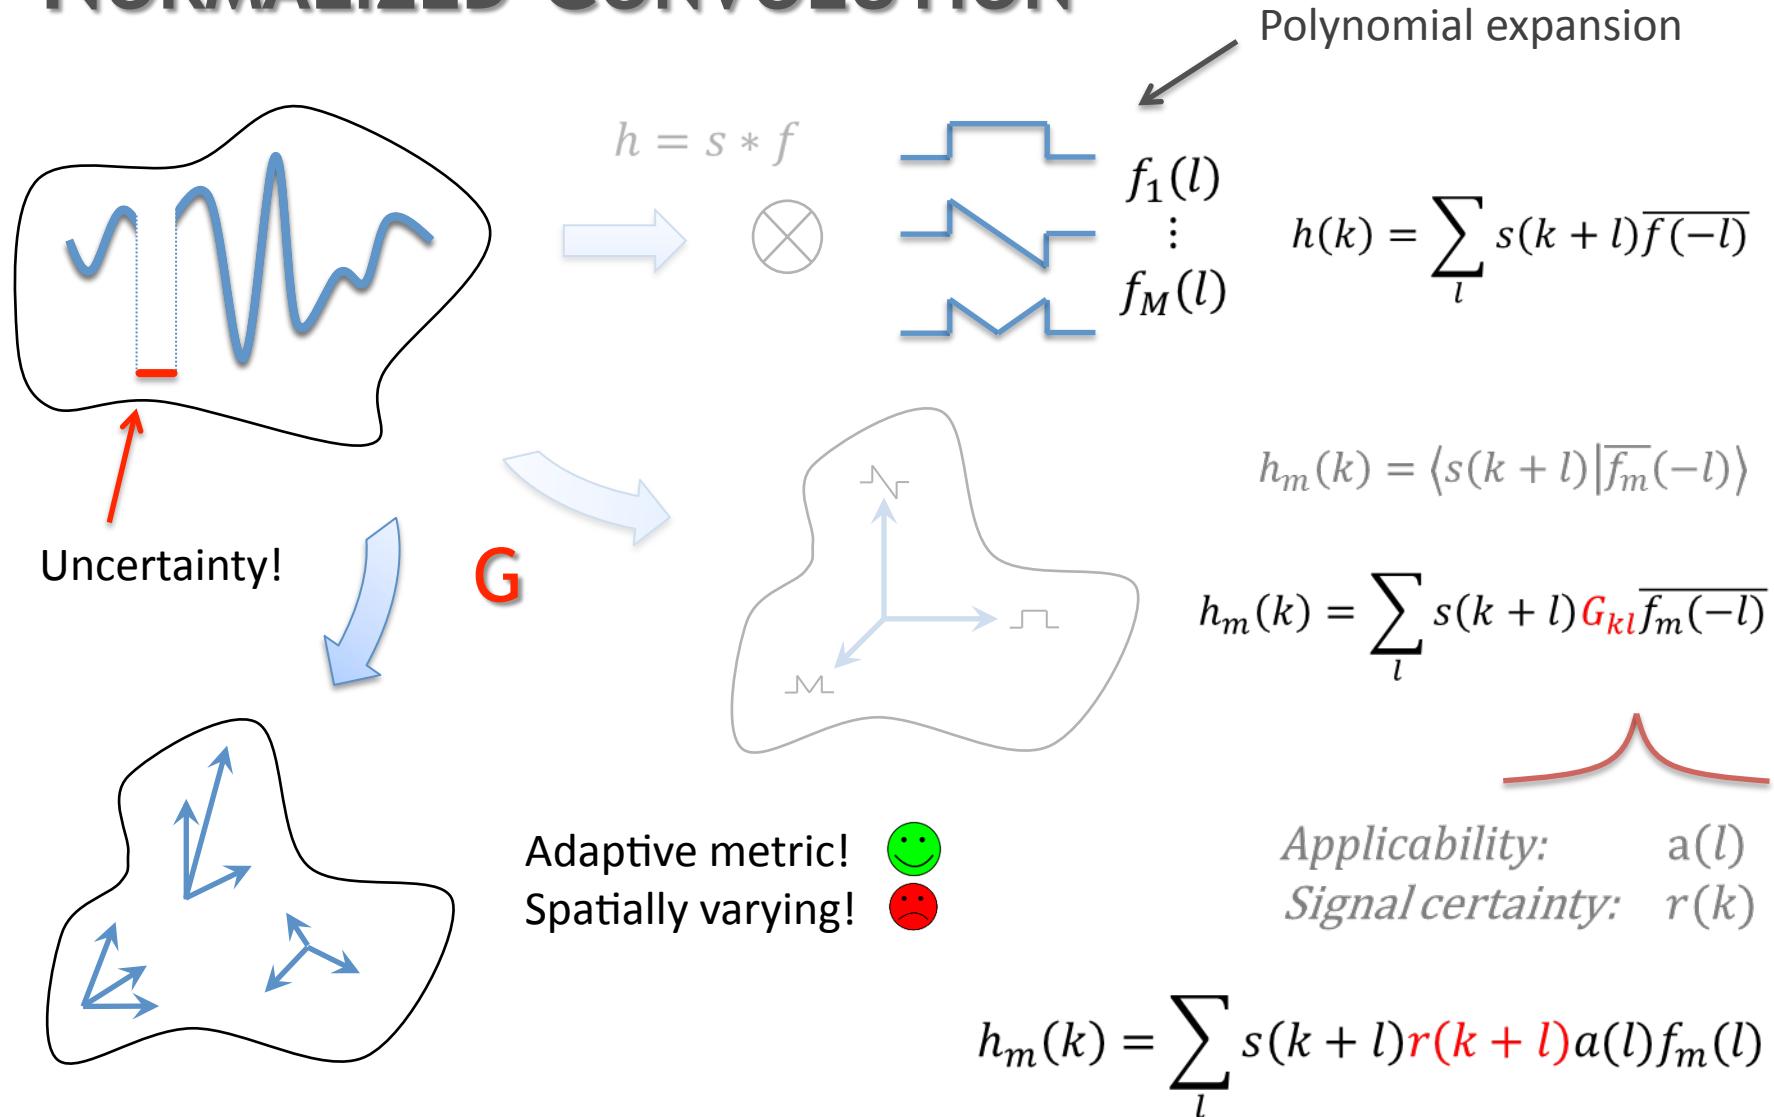


$$h_m(k) = \langle s(k + l) | \overline{f_m}(-l) \rangle$$

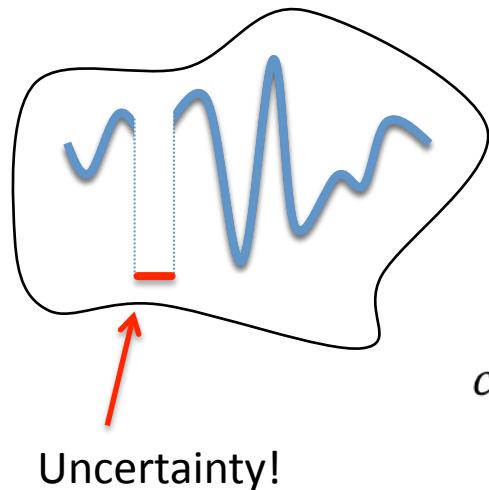
$$h_m(k) = \sum_l s(k + l) G_l \overline{f_m}(-l)$$



NORMALIZED CONVOLUTION

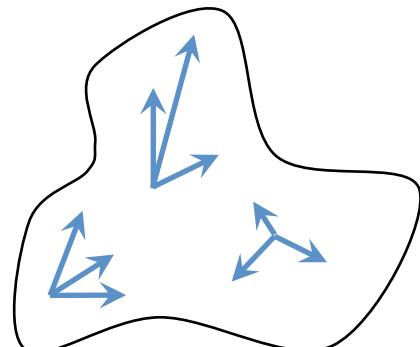


NORMALIZED CONVOLUTION



Polynomial Expansion for Orientation and Motion Estimation. [Far02] Farnebäck

$$c = \begin{pmatrix} \langle a * b_1 * \bar{b}_1 | r \rangle & \dots & \langle a * b_1 * \bar{b}_M | r \rangle \\ \vdots & \ddots & \vdots \\ \langle a * b_M * \bar{b}_1 | r \rangle & \dots & \langle a * b_M * \bar{b}_M | r \rangle \end{pmatrix}^{-1} \begin{pmatrix} \langle a * b_1 | r * s \rangle \\ \vdots \\ \langle a * b_M | r * s \rangle \end{pmatrix}$$



One matrix inversion per pixel in the image!

Adaptive metric!

Spatially varying!

Pre-computation!

G is only $M \times M$!

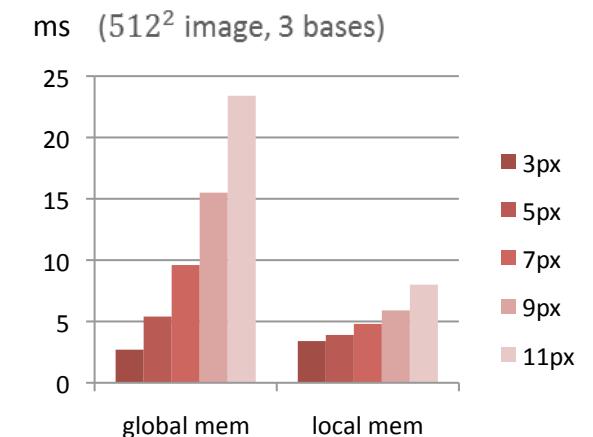
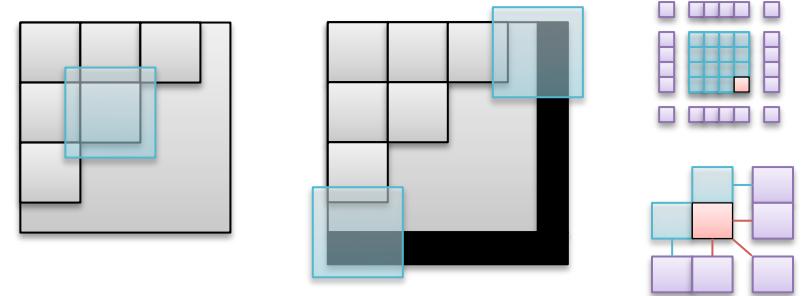


CUDA



CUDA

- Standard optimizations
 - 512 cores
 - Coalesced memory reads
 - Latency hiding
 - Loop unrolling
- Optimizations for Normalized Convolution
 - Templates for recursive matrix inversion
 - Pre-processor defines



CUDA

$$A^{-1} = \frac{1}{\det(A)} * \text{Adj}(A)$$

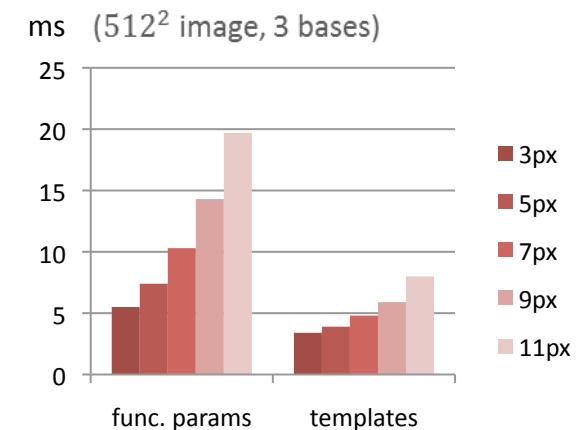
- Standard optimizations

- 512 cores
- Coalesced memory reads
- Latency hiding
- Loop unrolling

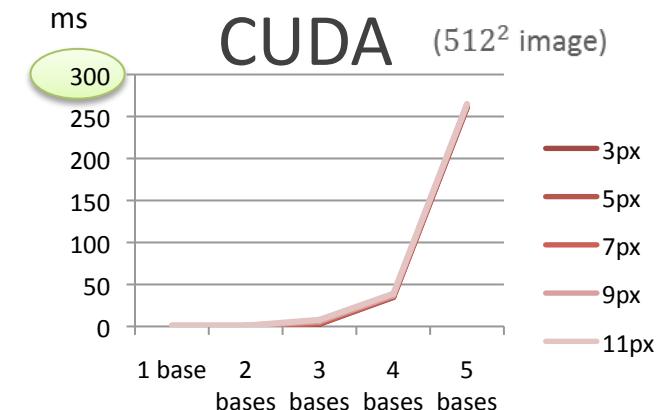
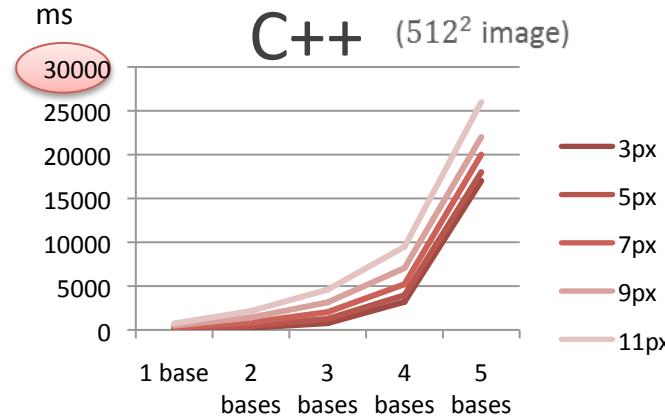
- Optimizations for Normalized Convolution

- Templates for recursive **matrix inversion**
- Pre-processor defines

```
template<int M> __device__ float Determinant()
    { ... Determinant<M-1>() ... }
template<> __device__ float Determinant<3>();
template<> __device__ float Determinant<2>();
template<> __device__ float Determinant<1>();
```



CUDA



CONCLUSIONS

G is only $M \times M$

...but even $M \times M$ is costly

...and we need precision

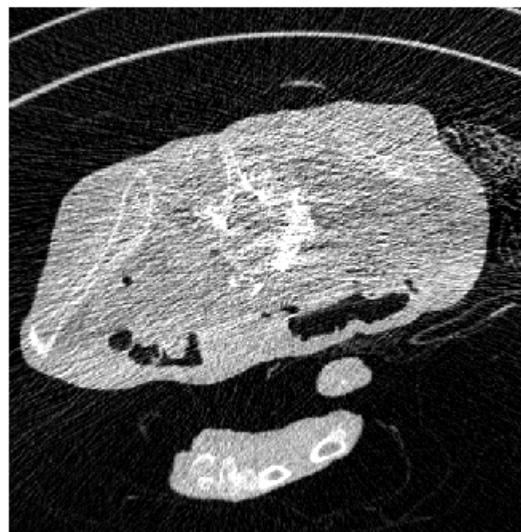
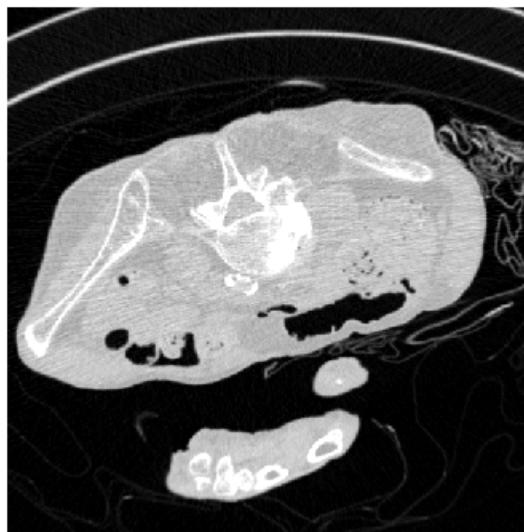
We can achieve real-time performance!



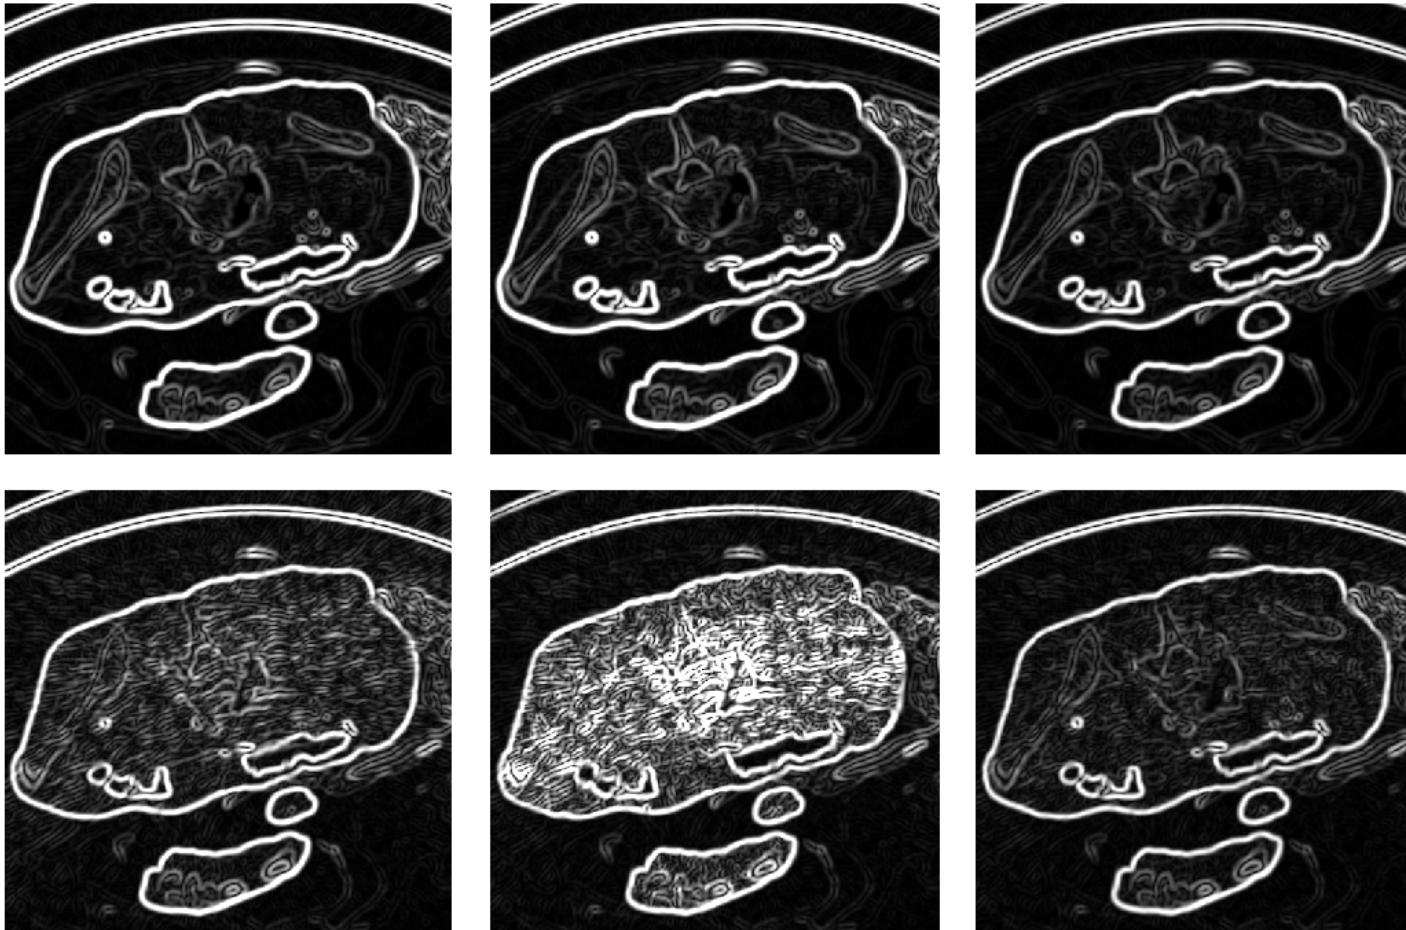
RESULTS



RESULTS



RESULTS



THANK YOU!



Linköping University
expanding reality



NORRKÖPING
VISUALIZATION CENTER



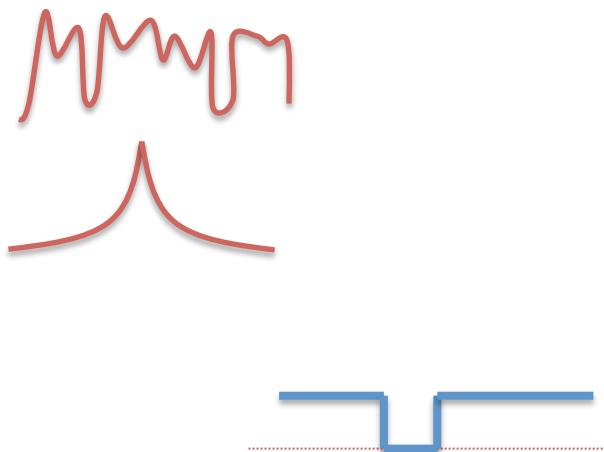
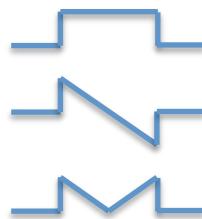
CRESEARCH

$$h_m(k) = \sum_l s(k+l) \overline{f_m(-l)}$$

$$h_m(k) = \sum_l s(k+l) G_l \overline{f_m(-l)}$$

$$h_m(k) = \sum_l s(k+l) r(k+l) a(l) f_m(l)$$

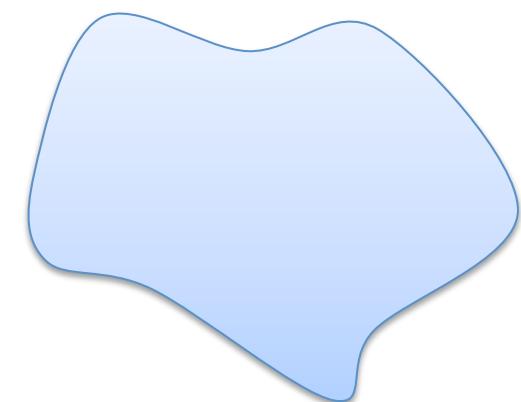




$$h(k) = \langle s(k+l) | f^*(-l) \rangle$$

$$c = G^{-1}$$

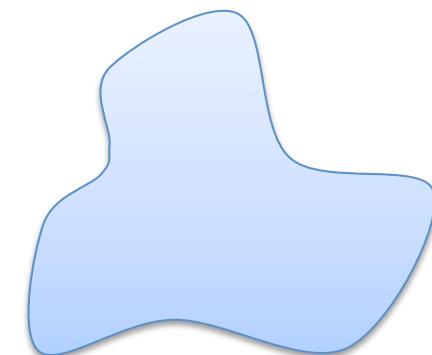
$$c = (B^* G_0 B)^{-1}$$



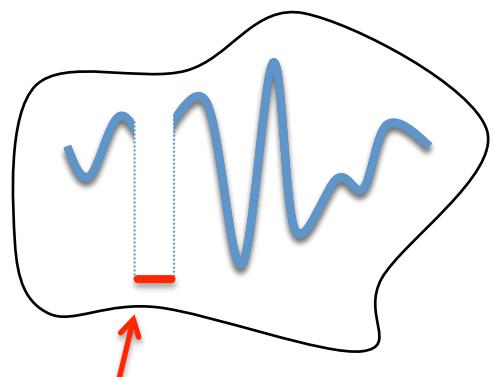
$$[111] * [111] = (1+1+1)/3 = 1$$

$$[\textcolor{red}{?}11]_{\textcolor{red}{[011]}} * [111] = \begin{matrix} \textcolor{red}{(0+1+1)/3} \\ \textcolor{red}{(1+1)/2} \end{matrix} = \begin{matrix} 0.67 \\ 1 \end{matrix}$$

$$\begin{matrix} \textcolor{blue}{|} & \textcolor{blue}{|} & \textcolor{blue}{|} \\ \textcolor{white}{|} & \textcolor{blue}{|} & \textcolor{blue}{|} \end{matrix} \otimes \begin{matrix} \textcolor{red}{|} & \textcolor{red}{|} \\ \textcolor{red}{|} & \textcolor{red}{|} \end{matrix} = \frac{\textcolor{blue}{|} + \textcolor{blue}{|} + \textcolor{blue}{|}}{3} = \begin{matrix} \textcolor{blue}{|} \\ \textcolor{white}{|} \\ \textcolor{blue}{|} & \textcolor{blue}{|} & \textcolor{blue}{|} \end{matrix}$$



CONVOLUTION

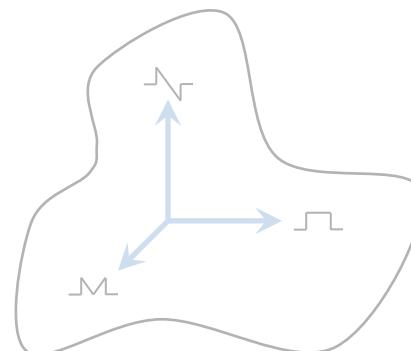
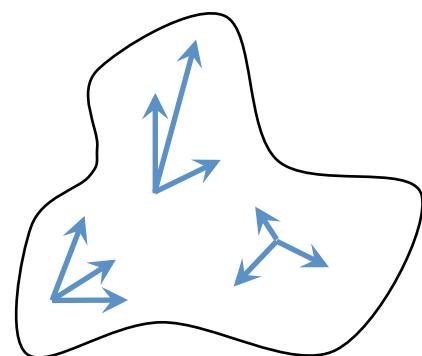


$$h = s * f \quad \otimes$$

$f_1(l) \quad \vdots \quad f_M(l)$

Polynomial expansion

$$h(k) = \sum_l s(k + l) \overline{f(-l)}$$



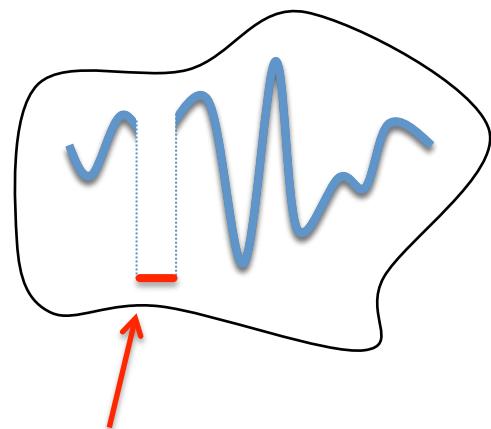
$$h_m(k) = \langle s(k + l) | \overline{f_m}(-l) \rangle$$

$$h_m(k) = \sum_l s(k + l) G_{kl} \overline{f_m}(-l)$$

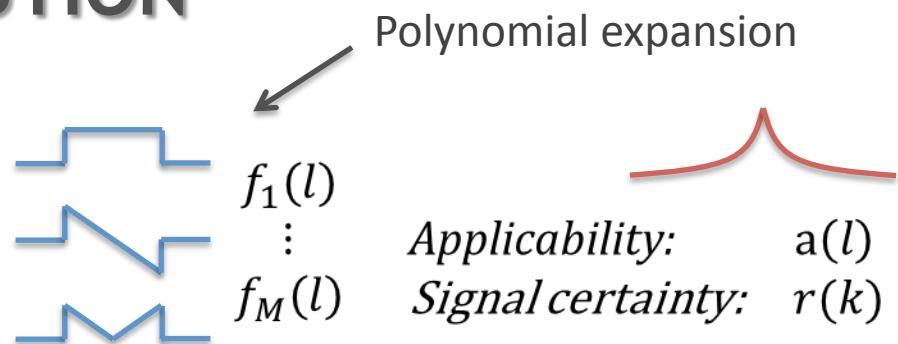
$$h_m(k) = \sum_l s(k + l) r(k + l) a(l) f_m(l)$$



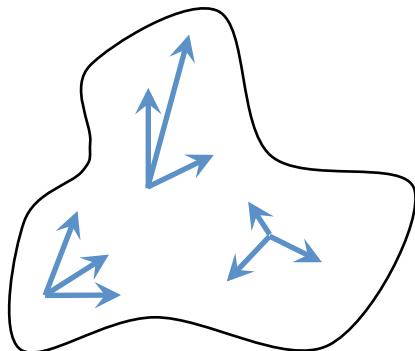
NORMALIZED CONVOLUTION



Uncertainty!



$$h_m(k) = \sum_l s(k + l)r(k + l)a(l)f_m(l)$$



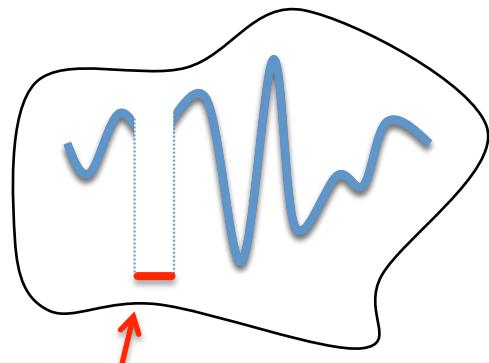
Different metric!
Spatially varying!

$$c = G^{-1}h_m$$

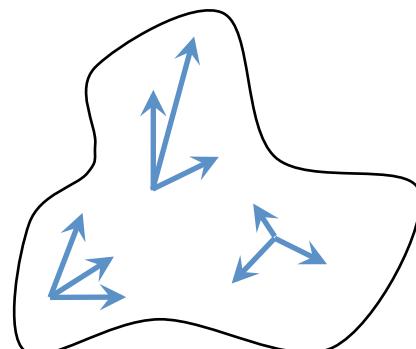
Weighted least squares



NORMALIZED CONVOLUTION



Uncertainty!



$$c = \begin{pmatrix} \langle a * b_1 * \bar{b}_1 | r \rangle & \dots & \langle a * b_1 * \bar{b}_M | r \rangle \\ \vdots & \ddots & \vdots \\ \langle a * b_M * \bar{b}_1 | r \rangle & \dots & \langle a * b_M * \bar{b}_M | r \rangle \end{pmatrix}^{-1} \begin{pmatrix} \langle a * b_1 | r * s \rangle \\ \vdots \\ \langle a * b_M | r * s \rangle \end{pmatrix}$$

Polynomial Expansion for Orientation
and Motion Estimation. [Far02] Farnebäck

$$c = G^{-1} h_m$$

- Different metric! 😞
- Spatially varying! 😞
- G is only $M \times M$! 😊
- Pre-computation! 😊

