Bona-fide Riesz projections for density estimation

Pol del Aguila Pla[1,2] and Michael Unser[2]

[1]: CIBM Center for Biomedical Imaging, Switzerland

[2]: Biomedical Imaging Group, École polytechnique fédérale de Lausanne, Lausanne, Switzerland

Motivation: Positron Emission Tomography

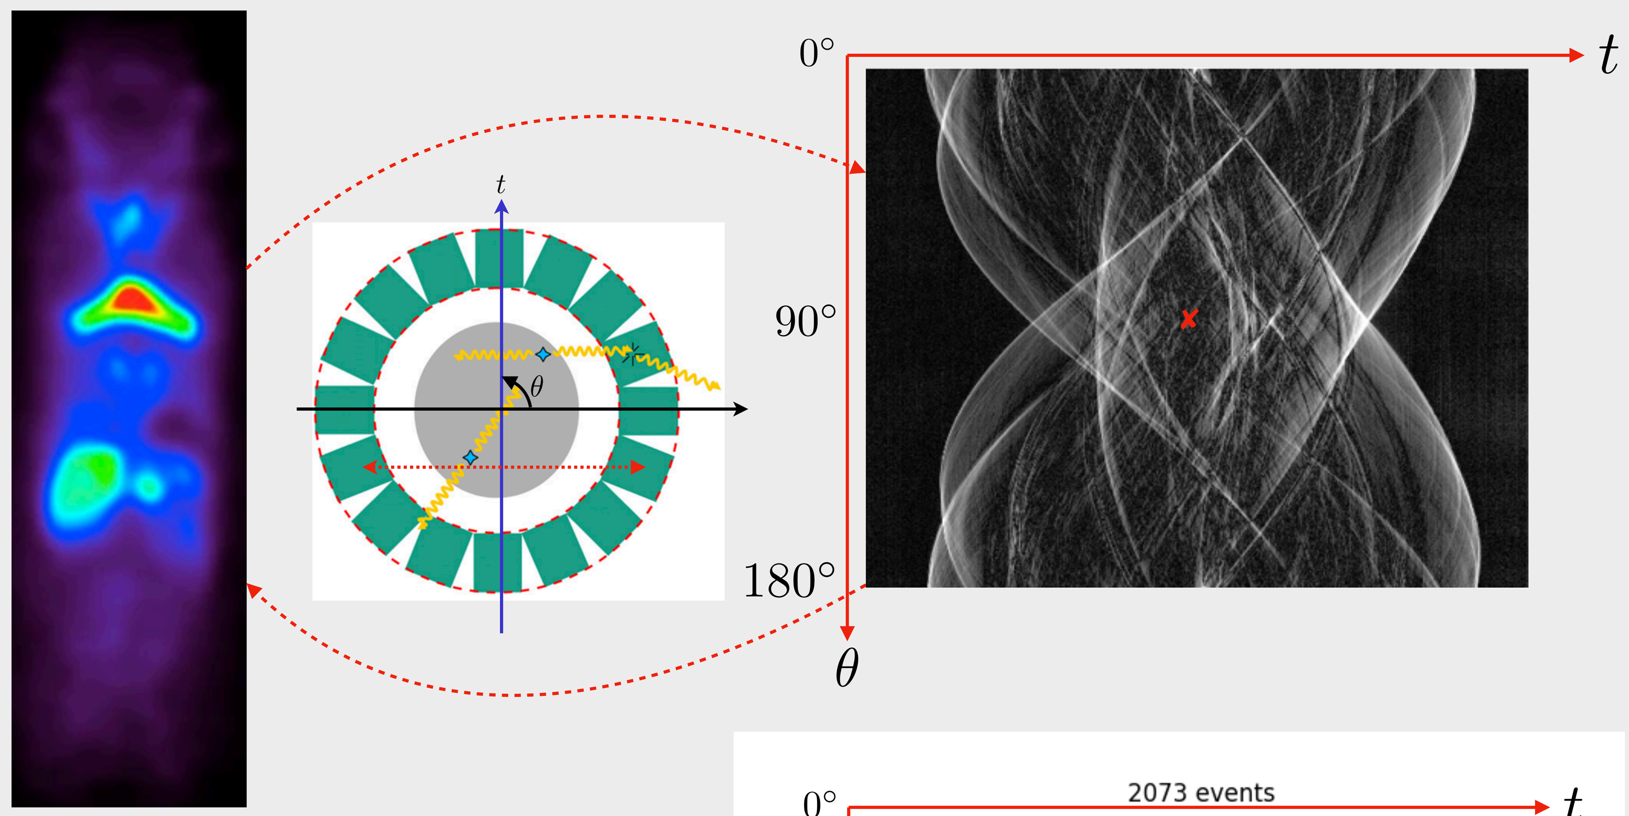

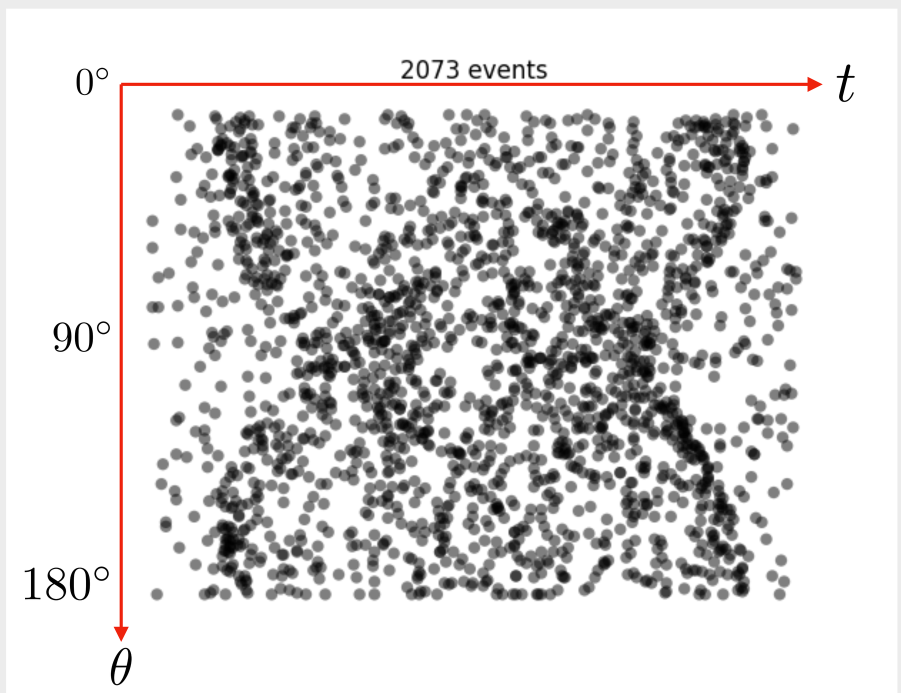

- Random sampling from a sinogram
- Scalable estimators beyond histograms

Background: Generalized Sampling

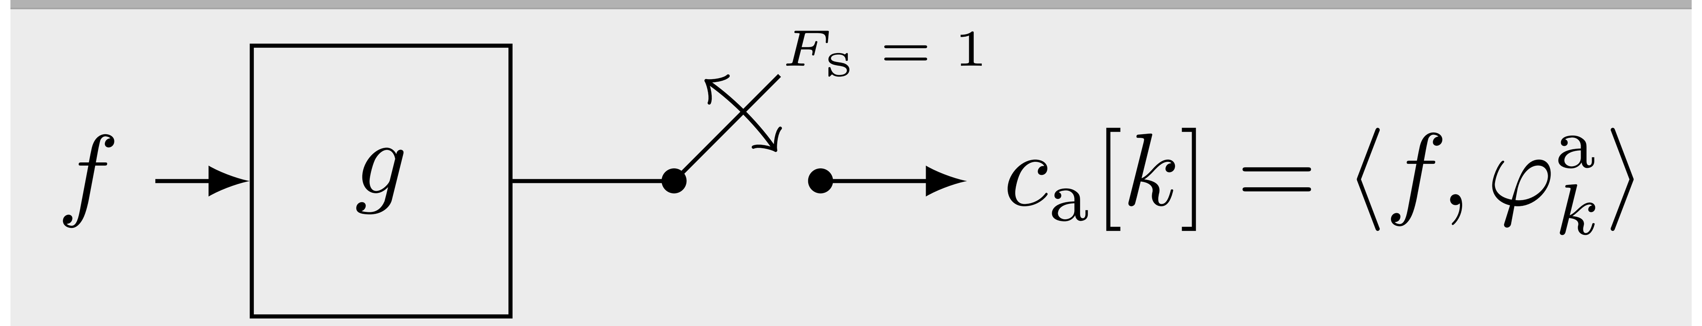

- $c_a[k] = (g * f)(k) = \langle f, \varphi_k^a \rangle$ where $\varphi_k^a(x) = g(k - x)$
- Represent f on a shift-invariant space

$$V_s = \left\{ \tilde{f} = \sum_{k\in\mathbb{Z}} c_s[k]\varphi_k^s : c_s \in \ell_2 \right\} \subset L_2(\mathbb{R})$$

Orthogonal projection (min L2)

- If $V_a = \{\breve{f} = \sum_{k\in\mathbb{Z}} c_a[k]\varphi_k^a : c_a \in \ell_2\} = V_s$

$\Rightarrow$ optimal c_s (L_2-wise) can be obtained by convolution

$\tilde{f} = \sum_{k\in\mathbb{Z}} c_s[k]\varphi_k^s$ with $c_s[k] = (q * c_a)[k], \forall k \in \mathbb{Z}$

Measurement Consistency Principle

$$\langle \tilde{f}, \varphi_l^a \rangle = \sum_{k\in\mathbb{Z}} (q * c_a)[k] \langle \varphi_k^s, \varphi_l^a \rangle$$

$$= (r_{a,s} * q * c_a)[l] = c_a[l] = \langle f, \varphi_l^a \rangle ,$$

$$\Leftrightarrow \hat{q}(\omega) = 1/\hat{r}_{a,s}(\omega), \text{ with } r_{a,s}[k] = \langle \varphi_k^a, \varphi^s \rangle .$$

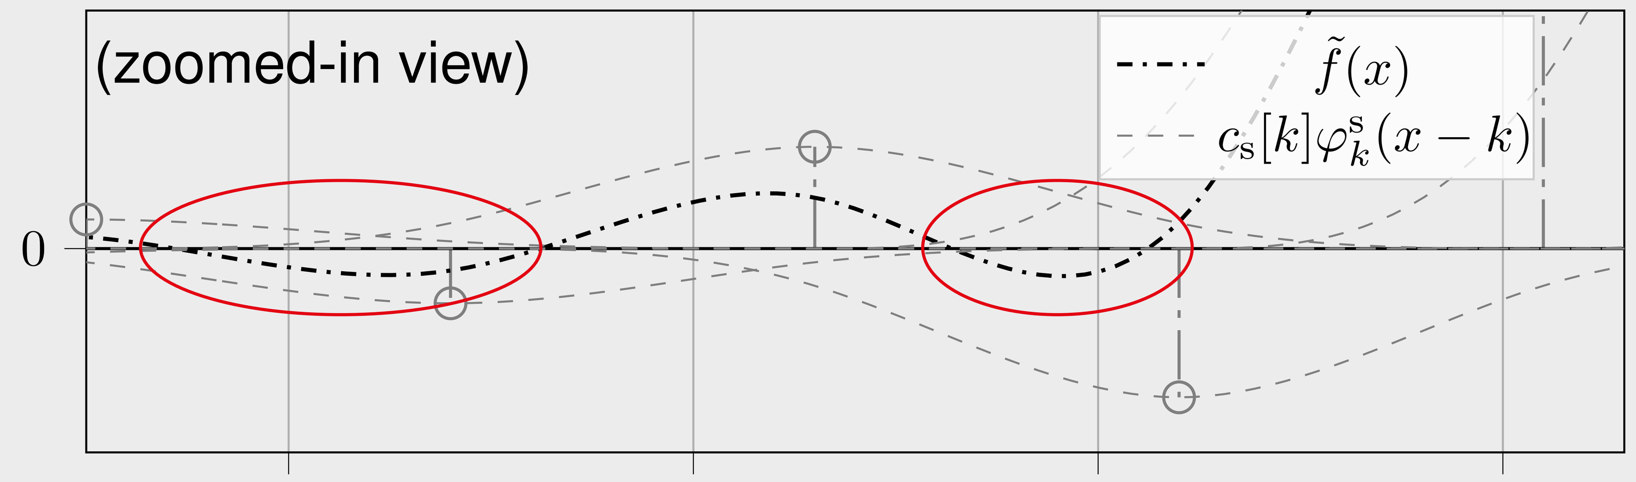

Bona-fide Relaxation of Consistency

$$\min_{\breve{f}\in V_s} \left\{ \| \langle f, \varphi_k^a \rangle - \langle \breve{f}, \varphi_k^a \rangle \|_2^2 \right\} \quad \mathcal{BF}$$

$$\text{such that } \breve{f}(x) \geq 0, \forall x \in \mathbb{R} \text{ and } \int \breve{f}(x)\mathrm{d}x = 1$$

- Approximating the non-negativity constraint with a factor M interpolation

$$\min_{c_s\in\ell_2} \left\{ \| c_a[k] - (r_{a,s} * c_s)[k] \|_2^2 \right\}$$

$$\text{such that } (c_s^{\uparrow M}[k] * \varphi^s(k/M))[q] \geq 0, \forall q \in \mathbb{Z}$$

$$\text{and } \sum_{k\in\mathbb{Z}} c_s[k] = 1$$

- Efficient, easy convex problem

Statistical Results

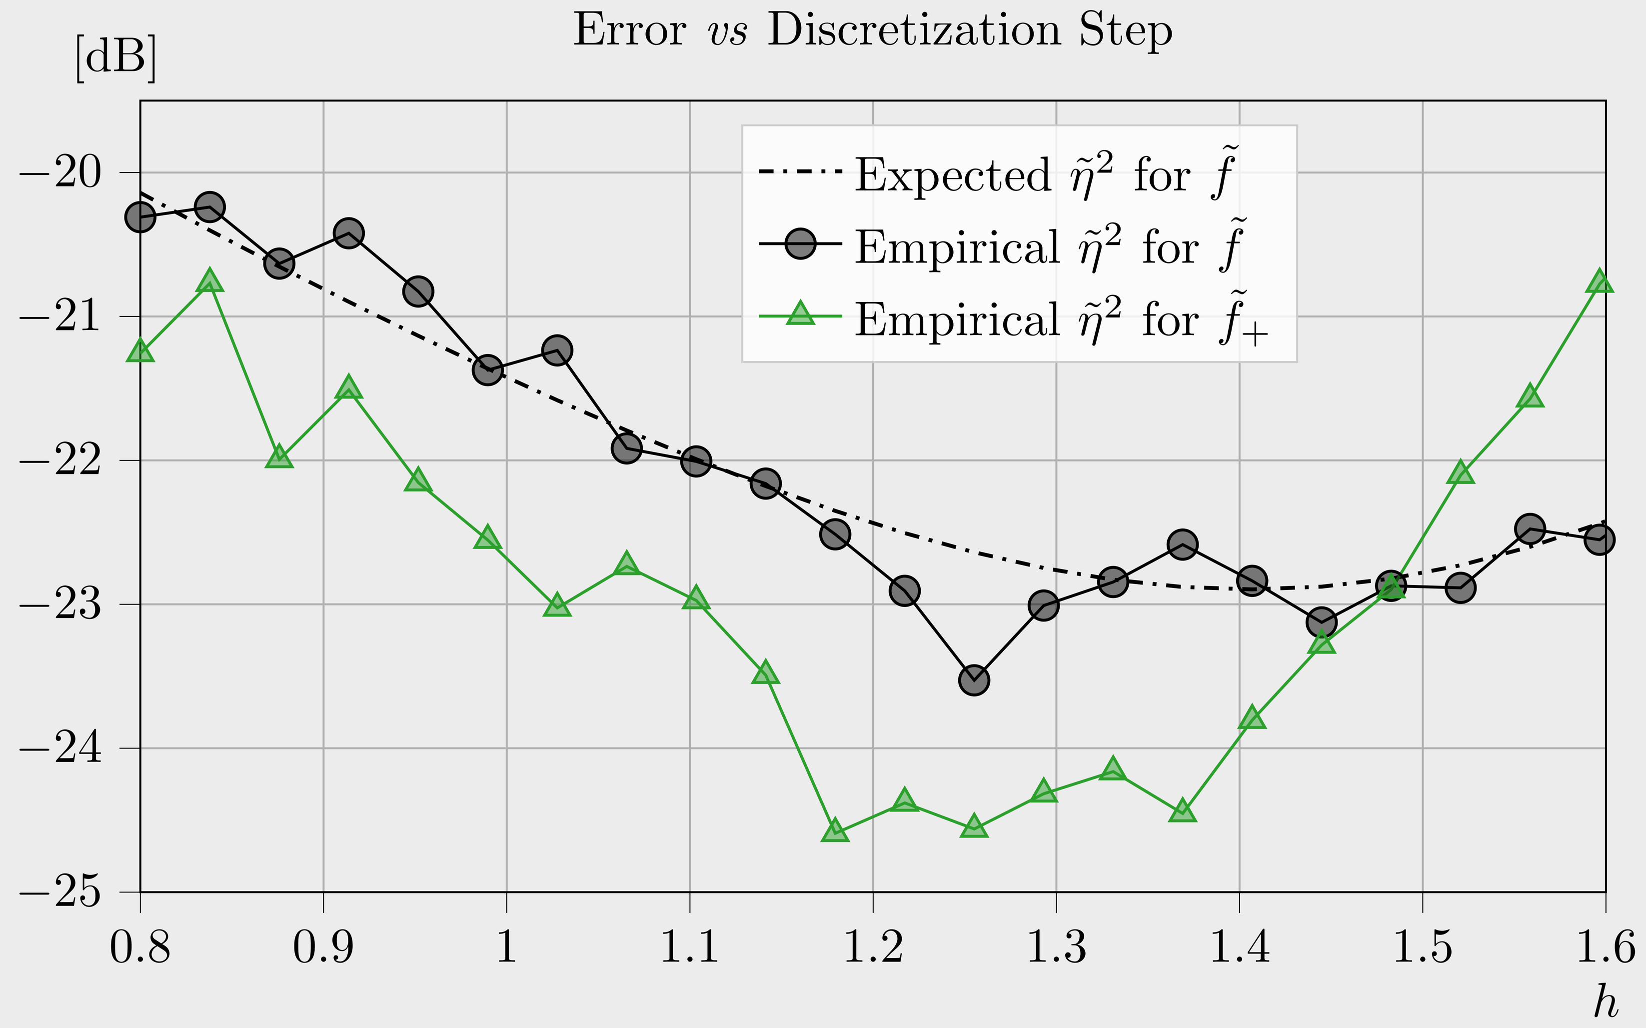

- Estimating a standard Gaussian from 10^3 samples, h from 0.8 to 1.6, averaged over 120 realizations

Example Estimate

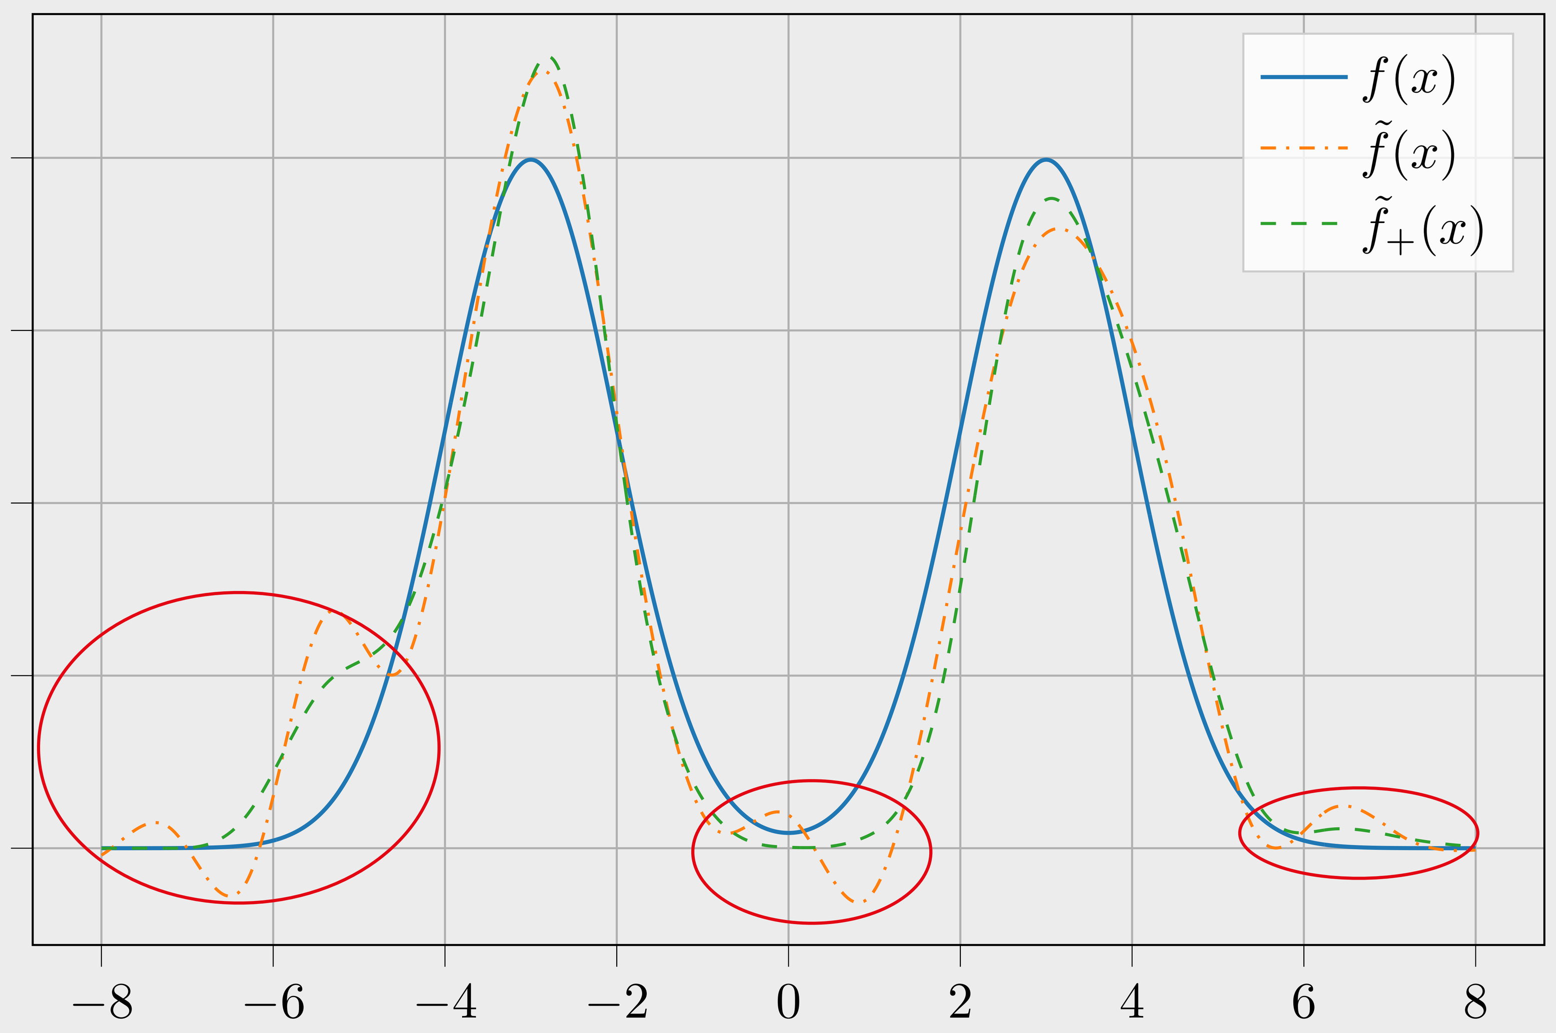

- Estimating an equal mixture between $\mathcal{N}(-3, 1)$ and $\mathcal{N}(3, 1)$ from 10^2 samples ($\varphi^a = \varphi^s = \beta_3$ and $r_{a,s}[k] = \beta_7(k)$) with $h = 0.9$