Singular Values & Functions

Per Christian Hansen

DTU Compute
Department of Applied Mathematics and Computer Science
Technical University of Denmark

Some Notation

	Vectors	Functions
Norm (2-norm)	$\|\boldsymbol{x}\|_2^2 = \sum_{i=1}^{n} \|x_i\|^2$ $= \boldsymbol{x} \cdot \boldsymbol{x} = \boldsymbol{x}^T \bar{\boldsymbol{x}}$	$\|f\|_2^2 = \int_a^b \|f(x)\|^2 \, \mathrm{d}x$ $= \langle f, f \rangle$
Inner prod.	$\boldsymbol{x} \cdot \boldsymbol{y} = \sum_{i=1}^{n} x_i \, \bar{y}_i = \boldsymbol{x}^T \bar{\boldsymbol{y}}$	$\langle f, g \rangle = \int_a^b f(x) \, \overline{g(x)} \, \mathrm{d}x$
Weighted ditto		$\langle f, g \rangle_w = \int_a^b f(x) \, \overline{g(x)} \, w(x) \, \mathrm{d}x$
Orthonormal	$\boldsymbol{v}_i \cdot \boldsymbol{v}_j = \boldsymbol{v}_i^T \boldsymbol{v}_j = \delta_{ij}$	$\langle v_i, v_j \rangle = \delta_{ij}$

All vectors are column vectors, the superscript "T" denotes transposition, and a bar denotes complex conjugation.

Reminder: Fourier Series of Periodic Functions

The Fourier series of a 2π-periodic function f is

$$f(x) = \sum_{n=-\infty}^{\infty} c_n[f]\, \mathrm{e}^{\mathrm{i}\, n\, x}, \qquad \mathrm{i} = \sqrt{-1},$$

with the Fourier coefficients

$$c_n[f] = \frac{1}{2\pi} \int_{-\pi}^{\pi} f(x)\, \mathrm{e}^{-\mathrm{i}\, n\, x}\, \mathrm{d}x = \langle f, \psi_n \rangle, \qquad \psi_n = \frac{1}{2\pi}\, \mathrm{e}^{\mathrm{i}\, n\, x}.$$

The functions ψ_n form an orthogonal basis for $L^2(-\pi, \pi)$, and they are a very convenient basis for analysing the behavior of periodic functions.

Wanted: More Insight

We have studied an efficient algorithm – filtered back projection (FBP) – for computing the CT reconstruction.

And we have also seen that the reconstruction is somewhat sensitive to noise in the data.

- How can we further study this sensitivity to noise?
- How can we possibly reduce the influence of the noise?
- What consequence does that have for the reconstruction?

We need a mathematical tool that lets us perform a detailed study of these aspects: the *singular value decomposition/expansion*.

But before going into these details, we will start with a simple example from signal processing, to explain the basic idea.

Motivation: Signal Restoration

Input $\Rightarrow$ **System** $\Rightarrow$ **Output**

Assume that we know the characteristics of the system, and that we have measured the noisy output signal $g(t)$. Now we want to reconstruct the input signal $f(t)$.

The mathematical (forward) model, assuming 2π-periodic signals:

$$g(x) = \mathcal{K}[f](x) = \int_{-\pi}^{\pi} h(y-x)\, f(y)\, \mathrm{d}y \qquad \text{or} \qquad g = h * f \ \text{(convolution)}.$$

Here, the function $h(t)$ (called the "impulse response") defines the system.

Deconvolution: reconstruct input f from output $g = h * f$

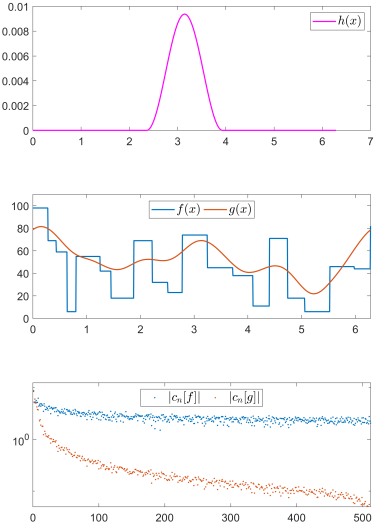

Convolution and Deconvolution in Fourier Domain

Due to the linearity, we have

$$g = h * f = h * \left(\sum_{n=-\infty}^{\infty} c_n[f]\, \psi_n \right) = \sum_{n=-\infty}^{\infty} c_n[f]\, (h * \psi_n).$$

Hence, all we need to know is the system's response $h * \psi_n$ to each basis function $\psi_n = \mathrm{e}^{\mathrm{i}\, n\, t}$.

For the periodic systems we consider here, the convolution of h with ψ_n produces a scaled version of ψ_n:

$$h * \psi_n = \mu_n\, \psi_n, \qquad \text{for all } n,$$

where $\mu_n = \langle h, \psi_n \rangle = c_n[h]$ (no proof). Hence, with $c_n[g] = \langle g, \psi_n \rangle$:

$$g = \sum_{n=-\infty}^{\infty} c_n[g]\, \psi_n = \sum_{n=-\infty}^{\infty} c_n[f]\, c_n[h]\, \psi_n = \quad \Leftrightarrow \quad \boxed{f = \sum_{n=-\infty}^{\infty} \frac{c_n[g]}{c_n[h]}\, \psi_n.}$$

Deconvolution is transformed to a simple algebraic operation: *division*.

Straightforward Reconstruction from Noisy Data

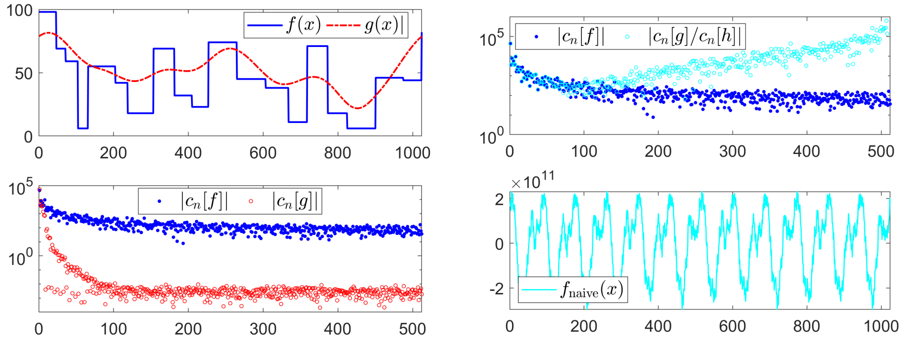

▷ Top left: one period of input $f(x)$ and noisy output $g(x)$ (noise invisible).

▷ Bottom left: corresponding Fourier coefficient; note the "noise floor."

▷ Top right: the reconstructed Fourier coefficients $c_n[g]/\mu_n$ are dominated by the noise for $n > 100$; a straightforward reconstruction is useless.

▷ Bottom right: the straightforward and useless reconstruction.

Filtered/Truncated Reconstruction from Noisy Data

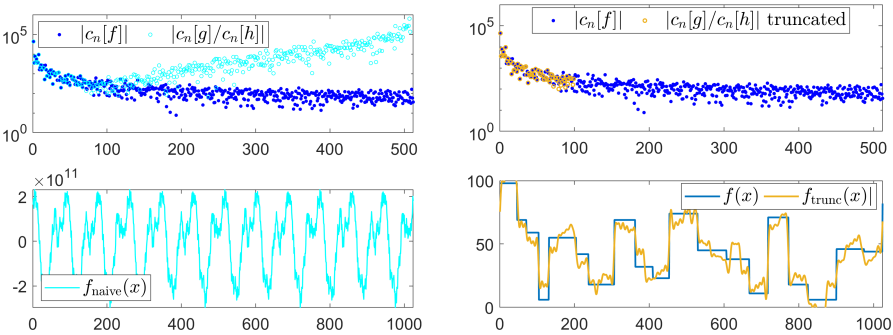

▷ Left: same as previous slide.

▷ Top right: let us keep the first ± 100 coefficients only.

▷ Bottom right: comparison of $f(t)$ and the truncated reconstruction using ± 100 terms in the Fourier expansion. It captures the general shape of $f(t)$.

What We Have Learned So Far

- With the right choice of basis functions, we can turn a complicated problem into a simper one.
- Here: the basis functions are the complex exponentials; deconvolution $\rightarrow$ division in Fourier domain.
- Inspection of the expansion coefficients reveals how and when the noise enters in the reconstruction.
- Here: the noise dominates the output's Fourier coefficients for *higher frequencies*, while the low-frequency coefficients are ok.
- We can avoid most of the noise (but not all) by means of **filtering**, at the cost of loosing some details.
- Here: we simply truncate the Fourier expansion for the reconstruction.

Let us apply the same idea to parallel-beam CT reconstruction!

The Radon Transform

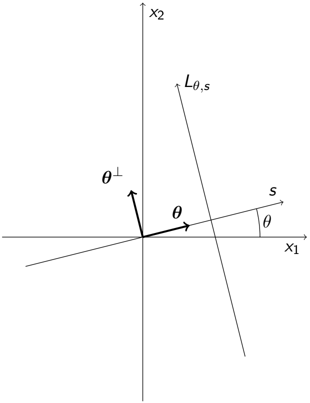

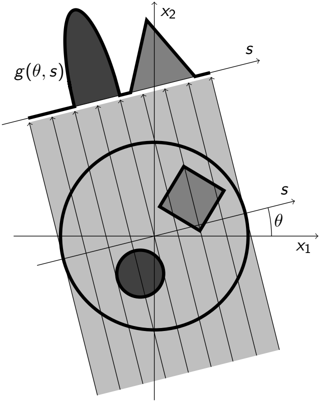

- The Radon transform $\boxed{g = \mathcal{R}\, f}$ Sometimes g is called $p_\theta(s)$.
- The image: $f(x_1, x_2)$ with $(x_1, x_2) \in \mathbb{D}$, the *disk* with radius 1.
- The sinogram (the data): $g(\theta, s)$ with $s \in [-1, 1]$ and $\theta \in [0, 2\pi]$.

Singular Values and Singular Functions

There exist unique scalars $\sigma_{n,k}$ and orthonormal functions $u_{n,k}(\theta, s)$ and $v_{n,k}(x_1, x_2)$ such that

$$\mathcal{R}\, v_{n,k} = \sigma_{n,k}\, u_{n,k}, \qquad n = 0, 1, 2, 3, \ldots \quad k = 0, 1, 2, \ldots, n.$$

The scalars are called the singular values:

$$\sigma_{n,k} = 2\sqrt{\pi/(n+1)} \qquad \text{with multiplicity } n+1.$$

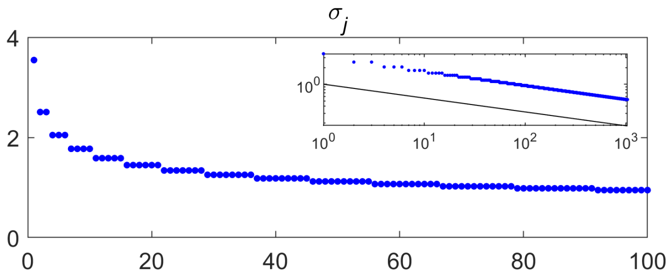

If $\sigma_{n,k} = \sigma_j$ with $j = \frac{1}{2}n(n+1) + k + 1$, then $\sigma_j \propto j^{-1/4}$ for large j.

(The word "singular" is used in the sense "special" or "unique.")

The Left Singular Functions

The left singular functions are given by

$$u_{n,k}(\theta, s) = \begin{cases} \frac{1}{\pi}\sqrt{1-s^2}\, U_n(s)\, \cos\big((n{-}2k)\,\theta\big)\ , & n{-}2k \geq 0 \\ \frac{1}{\pi}\sqrt{1-s^2}\, U_n(s)\, \sin\big((n{-}2k)\,\theta\big)\ , & n{-}2k < 0 \end{cases}$$

in which U_n are the Chebyshev polynomials of the second kind.

Note the convenient fact that the variables θ and s separate.

Some Left Singular Functions for $n = 0, 1, 2, 3, 4$

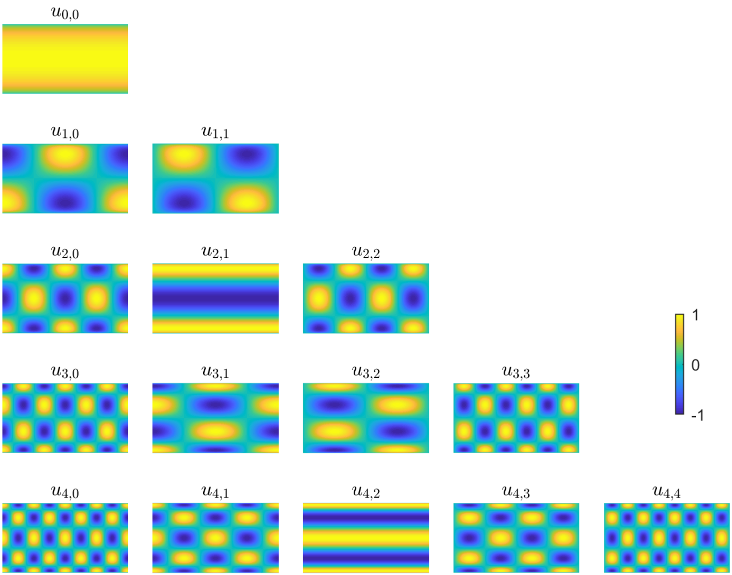

The Right Singular Functions

It is convenient to introduce polar coordinates (r, ϕ) such that $x_1 = r\cos\phi$ and $x_2 = r\sin\phi$. Then the right singular functions are given by

$$v_{n,k}(x_1, x_2) = \tilde{v}_{n,k}(r, \phi) = \sqrt{\frac{n+1}{\pi}}\, Z_{n,k}(r, \phi) \ , \qquad (1)$$

where $Z_{n,k}$ are the (real) Zernike polynomials:

$$Z_{n,k}(r, \phi) = \begin{cases} R_n^{n-2k}(r)\, \cos((n-2k)\,\phi) \ , & 2k \le n \\ R_n^{2k-n}(r)\, \sin((n-2k)\,\phi) \ , & 2k > n \end{cases} \qquad k = 0, \ldots, n \quad (2)$$

in which

$$R_n^{n-2k}(r) = (-1)^k\, r^{n-2k}\, P_k^{(n-2k,0)}(1 - 2r^2) \qquad (3)$$

and where $P_k^{(n-2k,0)}$ are the Jacobi polynomials.

Note the nice feature: in the form $\tilde{v}_{n,k}(r, \phi)$ the variables r and ϕ separate.

Some Right Singular Functions for $n = 0, 1, 2, 3, 4$

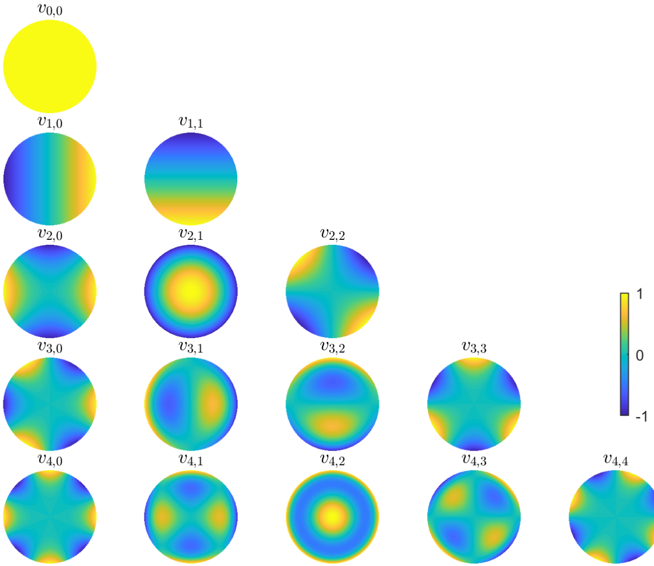

Singular Functions and Expansions

The functions $u_{n,k}$ are an orthonormal basis for $[-1,1] \times [0,2\pi]$.

The functions $v_{n,k}$ are an orthonormal basis for the unit disk $\mathbb{D}$.

The expansions of f and g take the form

$$f(x_1,x_2) = \sum_{n,k} \langle f, v_{n,k} \rangle \, v_{n,k}(x_1,x_2), \qquad g(\theta,s) = \sum_{n,k} \langle g, u_{n,k} \rangle_w \, u_{n,k}(\theta,s).$$

$$\begin{aligned} \langle f, v_{n,k} \rangle &= \int_0^{2\pi} \int_0^1 \tilde{v}_{n,k}(r,\phi)\, f(r,\phi)\, r \,\mathrm{d}r\, \mathrm{d}\phi, \\ \langle g, u_{n,k} \rangle_w &= \int_{-1}^1 \int_0^{2\pi} u_{n,k}(\theta,s)\, g(\theta,s)\, w(s)\, \mathrm{d}\theta\, \mathrm{d}s, \end{aligned}$$

where

$$w(s) = \frac{1}{\sqrt{1-s^2}} \; .$$

What We Learned

- All singular values $\sigma_{n,k}$ decay, by definition.
- Singular functions $u_{n,k}$ and $v_{n,k}$ with higher index n have higher frequencies.
- The higher the frequency, the more the damping in $\mathcal{R}\, v_{n,k} = \sigma_{n,k}\, u_{n,k}$.
- Hence the Radon transform

$$
\begin{aligned}
g &= \sum_{n,k} \langle g, u_{n,k} \rangle_w \, u_{n,k} = \mathcal{R}\, f = \mathcal{R} \sum_{n,k} \langle f, v_{n,k} \rangle \, v_{n,k} \\
&= \sum_{n,k} \langle f, v_{n,k} \rangle \, \mathcal{R}\, v_{n,k} = \sum_{n,k} \langle f, v_{n,k} \rangle \, \sigma_{n,k} \, u_{n,k}
\end{aligned}
$$

  is a "smoothing" operation
- . . . and the reverse operation $f = \mathcal{R}^{-1} g$ amplifies higher frequencies!

These are intrinsic properties of the mathematical problem itself.

The Coefficients for the Sinogram

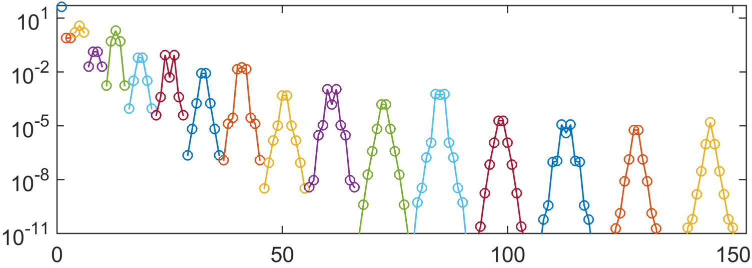

These are the coefficients $\langle g, u_{n,k} \rangle$ for the sinogram corresponding to the Shepp-Logan phantom – ordered according to increasing index n.

They decay, as expected. The specific behavior for $k = 0, \ldots, n$ is due to the symmetry of the phantom.

Left Singular Functions $u_{n,k}$ and a "Boundary Condition"

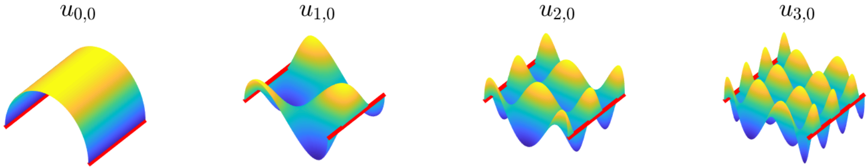

Due to the factor $\sqrt{1-s^2}$, all the left singular functions satisfy

$$u_{n,k}(\theta, s) \to 0 \quad \text{for} \quad s \to \pm 1.$$

This reflects the fact that rays through the disk $\mathbb{D}$ that almost grace the edge of the disk contribute very little to the sinogram.

This puts a restriction on sinograms $g(\theta, s)$ that admit a reconstruction:

- The sinogram $g = \mathcal{R} f$ is a sum of the singular functions $u_{n,k}$.
- Hence, the sinogram inherits the property $g(\theta, s) \to 0$ for $s \to \pm 1$.
- A perturbation Δg of g that does not have this property may not produce a bounded perturbation $\mathcal{R}^{-1} \Delta g$ of f.

When the Noise Violates the "Boundary Condition"

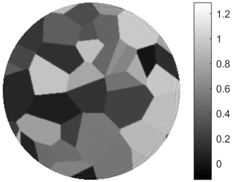

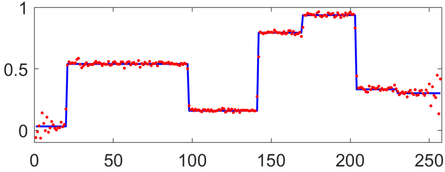

We added an increased amount of noise in $g(\theta, s)$ near $s = \pm 1$.

Top: the reconstruction computed by means of FBP. Bottom: the middle vertical column of pixels. We see severe artifacts near the edge of the disk.

The artifacts are "un-physical" – specifically, some of the pixels in the reconstruction (the attenuation coefficients) have negative values.

Let's Reconstruct

In terms of the singular values and functions, the inverse Radon transform takes the form

$$f(x_1, x_2) = \sum_{n,k} \frac{\langle g, u_{n,k} \rangle_w}{\sigma_{n,k}} \, v_{n,k}(x_1, x_2).$$

Since the image $f(x_1, x_2)$ has finite norm (finite energy), we conclude that the magnitude of the coefficient $\langle g, u_{n,k} \rangle_w / \sigma_{n,k}$ must decay "sufficiently fast."

The Picard Condition. The expansion coefficients $\langle g, u_{h,k} \rangle_w$ for $g(\theta, s)$ must decay sufficiently faster than the singular values $\sigma_{n,k}$, such that

$$\sum_{n,k} \left| \frac{\langle g, u_{n,k} \rangle_w}{\sigma_{n,k}} \right|^2 < \infty.$$

When noise is present in the measured sinogram $g(\theta, s)$, this condition is not satisfied for large n (cf. the signal restoration example from before). $\to$ This calls for some kind of filtering.

Let's Introduce Filters

A simple remedy for the noise-magnification, by the division with σ_{mk}, is to introduce filtering:

$$f(x_1, x_2) = \sum_{n,k} \varphi_{n,k} \, \frac{\langle g, u_{n,k} \rangle_w}{\sigma_{n,k}} \, v_{n,k}(x_1, x_2).$$

The filter factors $\varphi_{n,k}$ must decay fast enough that they, for large n, can counteract the growing factor $\sigma_{n,k}^{-1}$. More on this later in the course.

We can think of the filter factors as *modifiers* of the expansion coefficients $\langle g, u_{n,k} \rangle_w$ for the sinogram.

In other words, they ensure that the filtered coefficients $\varphi_{n,k} \, \langle g, u_{n,k} \rangle_w$ decay fast enough to satisfy the Picard condition from the previous slide.

The filtering inevitably dampens the higher frequencies associated with the small $\sigma_{n,k}$ and hence some details and edges in the image are lost.

Connection to Filtered Back Projection

Recall the filtered back projection (FBP) algorithm:

1. For fixed θ compute the Fourier transform $\hat{g}(\theta, \omega) = \mathcal{F}\big(g(\theta, s)\big)$.
2. Apply the ramp filter $|\omega|$ and compute the inverse Fourier transform $g_{\text{filt}}(\theta, s) = \mathcal{F}^{-1}\big(|\omega|\, \hat{g}(\theta, \omega)\big)$.
3. Do the above for all $\theta \in [0, 2\pi]$.
4. Then compute $f(x_1, x_2) = \int_0^{2\pi} g_{\text{filt}}(x_1 \cos\theta + x_2 \sin\theta, \theta)\, d\theta$.

It is the ramp filter $|\omega|$ in step 2 that magnifies the higher frequencies in the sinogram $g(\theta, s)$.

This amplification is *equivalent* to the division by the singular values $\sigma_{n,k}$ in the above analysis.

Filtered Back Projection, now with Low-Pass Filtering

How the filtered back projection algorithm (FBP) is really implemented:

1. Choose a *low-pass filter* $\varphi_{\mathrm{LP}}(\omega)$.
2. For every θ compute the Fourier transform $\hat{g}(\theta,\omega) = \mathcal{F}\big(g(\theta,s)\big)$.
3. Apply the *combined ramp & low-pass filter*, and compute the inverse Fourier transform $g_{\mathrm{filtfilt}}(\theta,s) = \mathcal{F}^{-1}\big(|\omega|\,\varphi_{\mathrm{LP}}(\omega)\,\hat{g}(\theta,\omega)\big)$.
4. Then $f_{\mathrm{rec}}(x_1,x_2) = \int_0^{2\pi} g_{\mathrm{filtfilt}}(x_1\cos\theta + x_2\sin\theta,\theta)\,d\theta$.

The low-pass filter $\varphi_{\mathrm{LP}}(\omega)$ counteracts the ramp filter $|\omega|$ for large ω. It is equivalent to the filter factors $\varphi_{n,k}$ introduced on slide 23.