Efficient and compressive sampling on the sphere

July 14, 2014

Science on the Sphere, London, UK

Y. Wiaux

Institute of Sensors, Signals, and Systems, Heriot-Watt University

basp.eps.hw.ac.uk

Collaboration with **J. D. McEwen and**

G. Puy, J.-Ph. Thiran, P. Vandergheynst, D. Van De Ville

Introduction

Motivation

Signals living on the sphere naturally arise in many fields, ranging from biomedical imaging, or geophysics, to astrophysics and others...

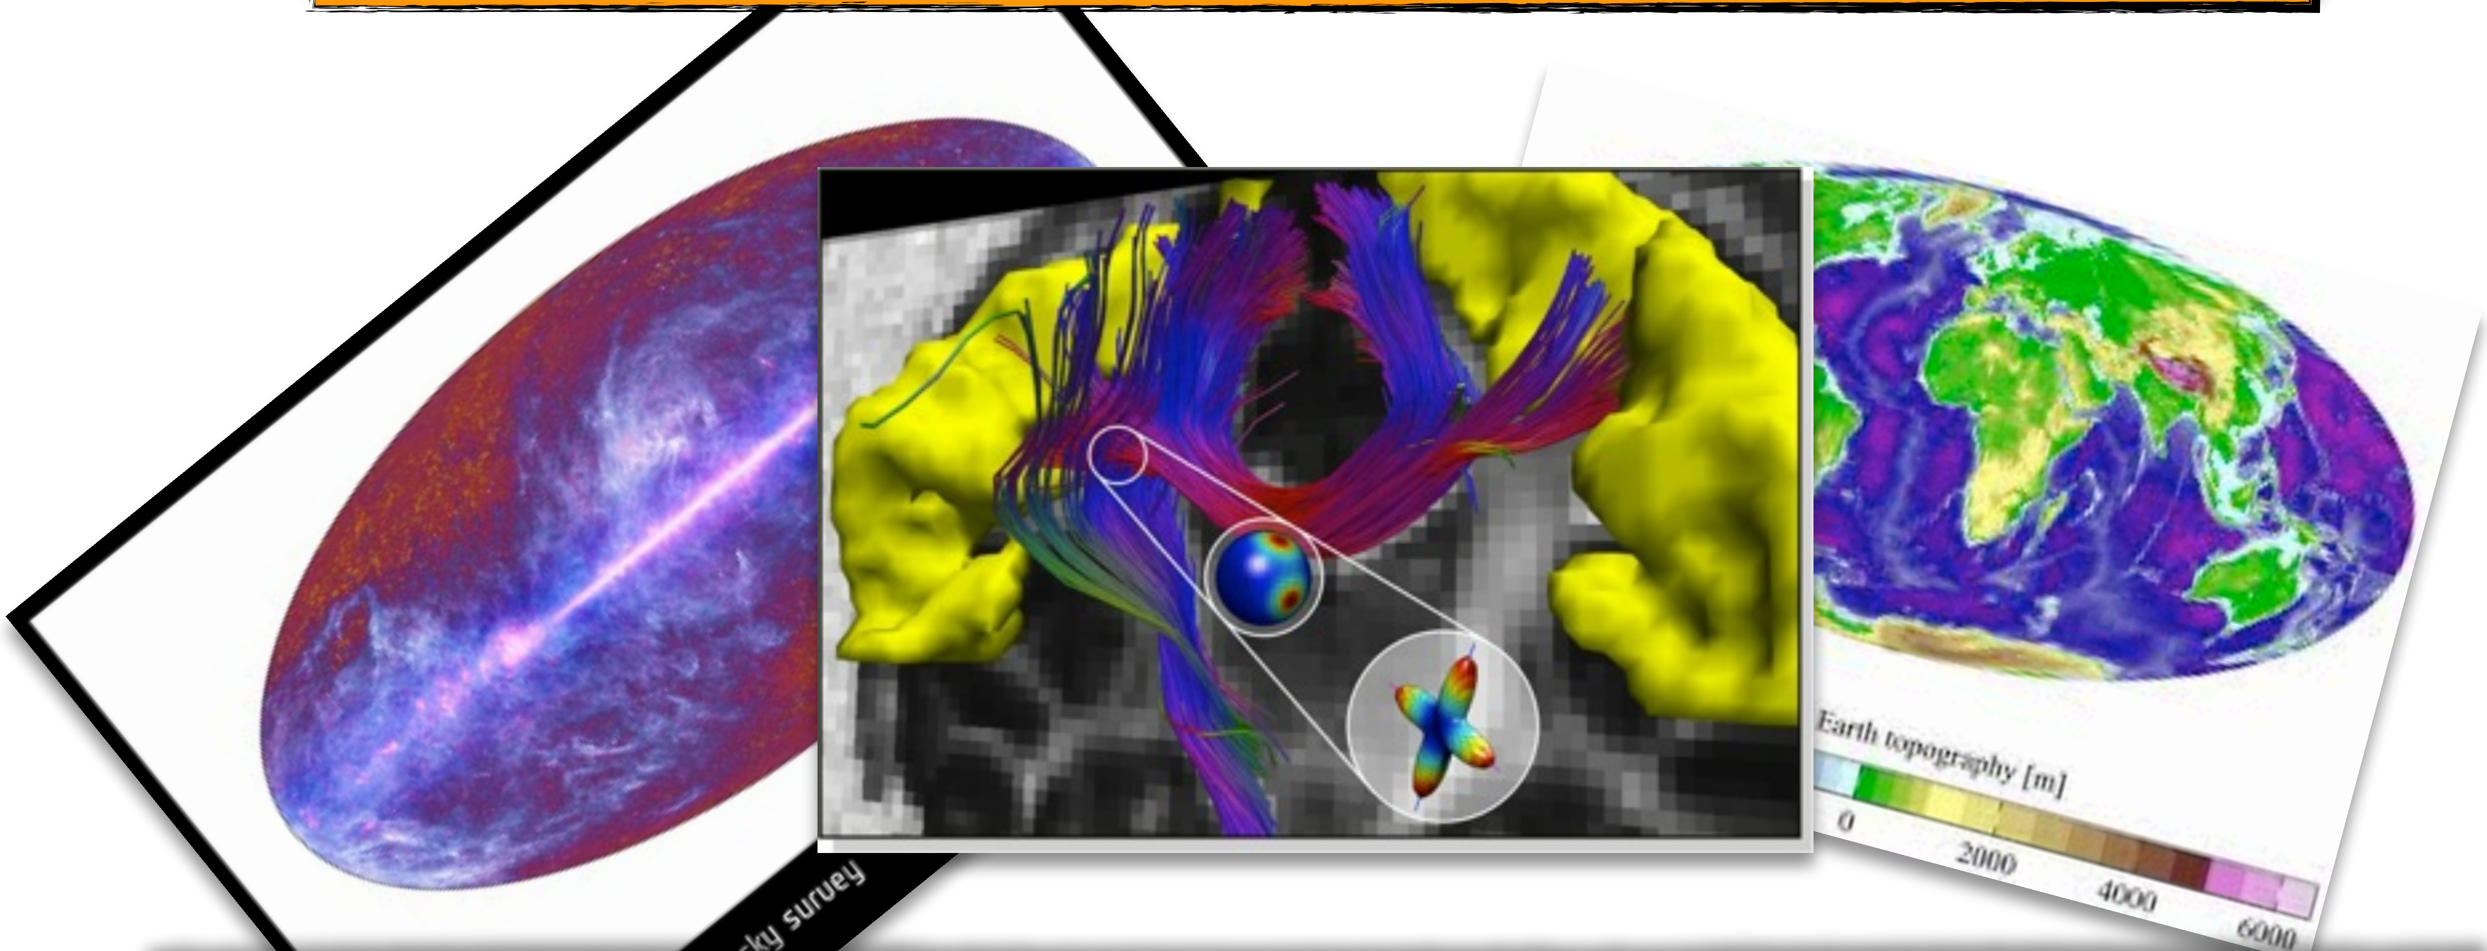

With spherical sampling issues at the core of all corresponding signal processing considerations...

Presentation summary

Overview

We discuss "Nyquist" sampling on the sphere and highlight its implications for compressed sensing.

Plan

I. A new sampling theorem for band limited signals

JD McEwen & YW
IEEE TSP
+ CODE PUBLICLY AVAILABLE

II. Sparsity and dimensionality implications for compressed sensing

JD McEwen et al.
IEEE TIP
+ CODE PUBLICLY AVAILABLE

Conclusions

I.
Sampling theorem(s)

Harmonic analysis

* Scalar and spin square integrable functions on the sphere...

Spin function $\quad {}_sf'(\theta,\varphi) = \mathrm{e}^{-is\chi}\ {}_sf(\theta,\varphi)$

for colatitude $\quad \theta \in [0,\pi]$

and longitude $\quad \varphi \in [0,2\pi)$

Scalar product $\quad \langle f,\, g\rangle = \int_{\mathrm{S}^2} \mathrm{d}\Omega(\theta,\varphi)\, f(\theta,\varphi)\, g^*(\theta,\varphi)$

Invariant measure $\quad \mathrm{d}\Omega(\theta,\varphi) = \sin\theta\, \mathrm{d}\theta\, \mathrm{d}\varphi$

Harmonic analysis

* Orthonormal spherical harmonic basis...

$$ {}_sY_{\ell m}(\theta,\varphi) = (-1)^s \sqrt{\frac{2\ell+1}{4\pi}}\, D^{\ell\,*}_{m,-s}(\varphi,\theta,0) $$

For $$ D^{\ell}_{mn}(\alpha,\beta,\gamma) = \mathrm{e}^{-im\alpha}\, d^{\ell}_{mn}(\beta)\, \mathrm{e}^{-in\gamma} $$

and with $$ \langle {}_sY_{\ell m},\ {}_sY_{\ell' m'} \rangle = \delta_{\ell\ell'}\delta_{mm'} $$

$$ \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} {}_sY_{\ell m}(\theta,\varphi)\, {}_sY^{*}_{\ell m}(\theta',\varphi') = \delta(\cos\theta - \cos\theta')\,\delta(\varphi-\varphi') $$

Harmonic analysis

* Spherical harmonic coefficients of a function...

$$ {}_s f_{\ell m} = \langle {}_s f, {}_s Y_{\ell m} \rangle \qquad \text{(Forward transform)} $$

$$ \ell \in \mathbb{N} $$

$$ m \in \mathbb{Z}, \ |m| \leq \ell $$

$$ {}_s f(\theta, \varphi) = \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} {}_s f_{\ell m} \, {}_s Y_{\ell m}(\theta, \varphi) \qquad \text{(Inverse transform)} $$

Harmonic analysis

* Band limitation: for a band limit L , the continuous signal is defined by exactly L^2 spherical harmonic coefficients... a sampling theorem is about: how many samples are needed on the sphere to recover the signal.

$$ {}_sf_{lm} = 0, \forall \ell \geq L $$

Existing exact sampling theorems

* Required numbers of sampling points [restrict to iso-latitude pixelisations for separations of variables, leading to state-of-art complexity $\mathcal{O}(L^3)$]...

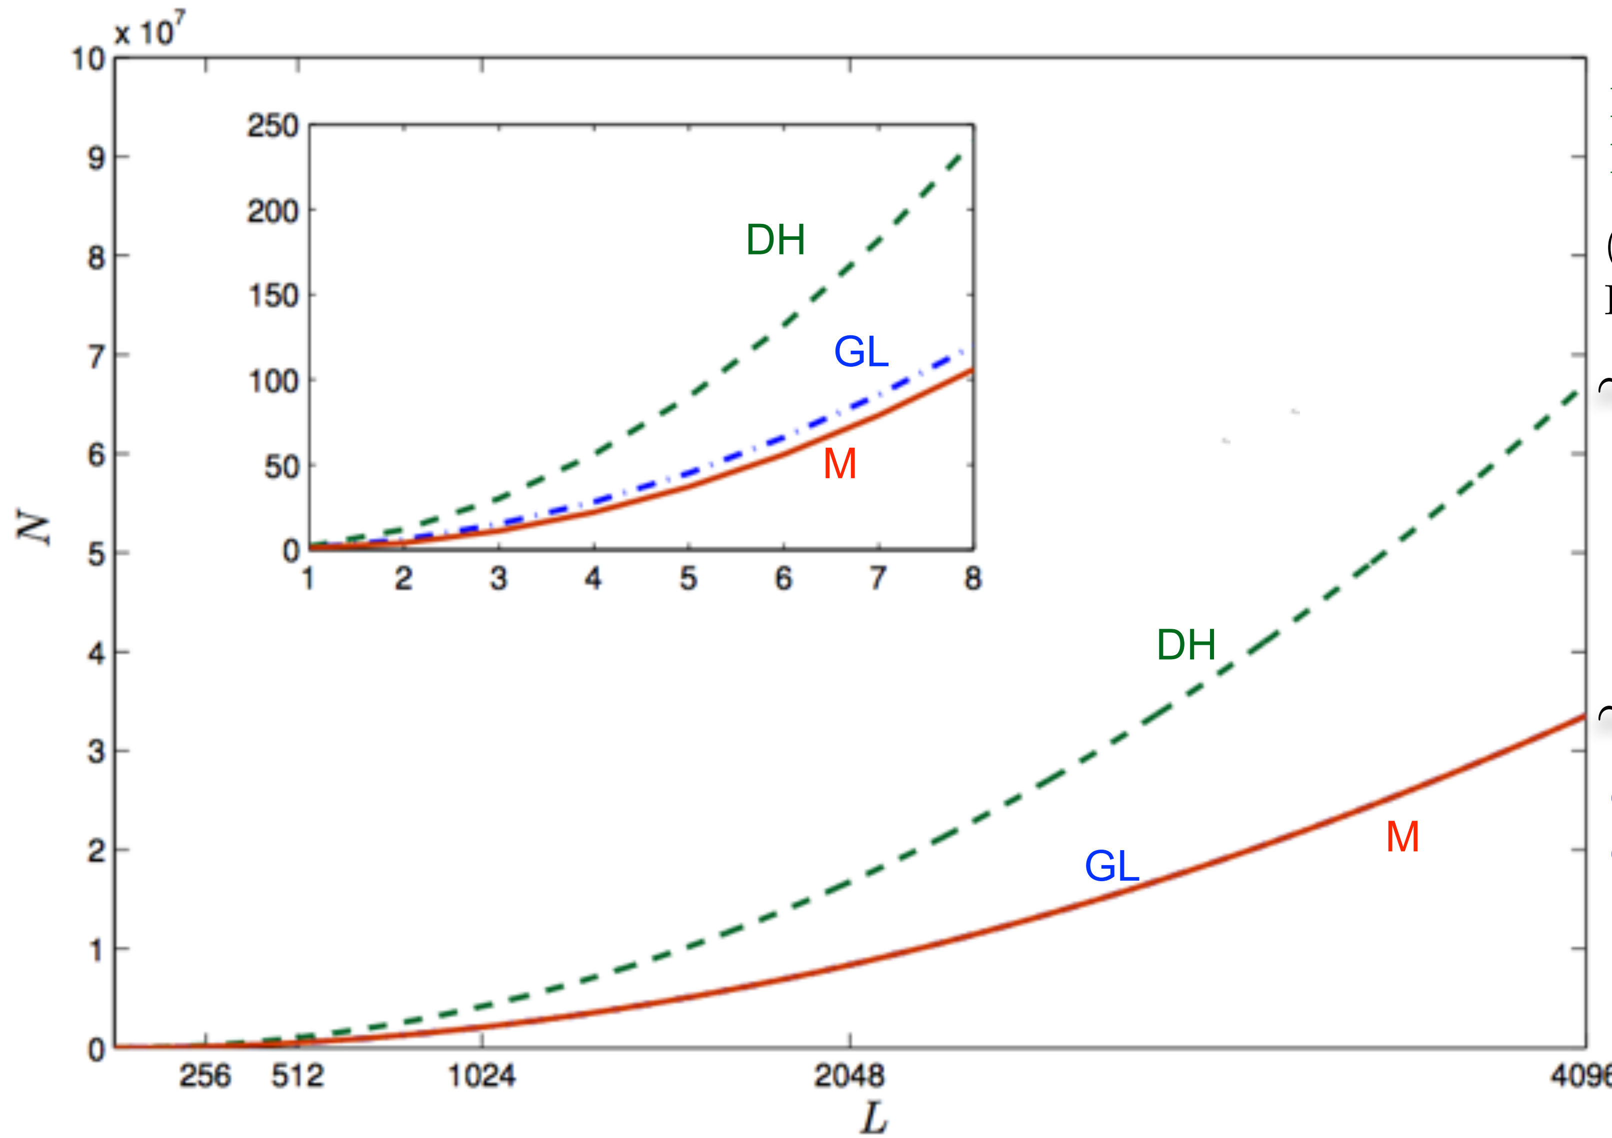

DH: Equiangular
Driscoll & Healy sampling

(Not shown) HW: Equiangular
Huffenberger & Wandelt sampling

$\sim 4L^2$

$\sim 2L^2$

GL: Non equiangular
Gauss-Legendre sampling

M: [Unstable] Equiangular
McEwen Sampling

Existing sampling theorems

* Sampling distributions (same color code)...

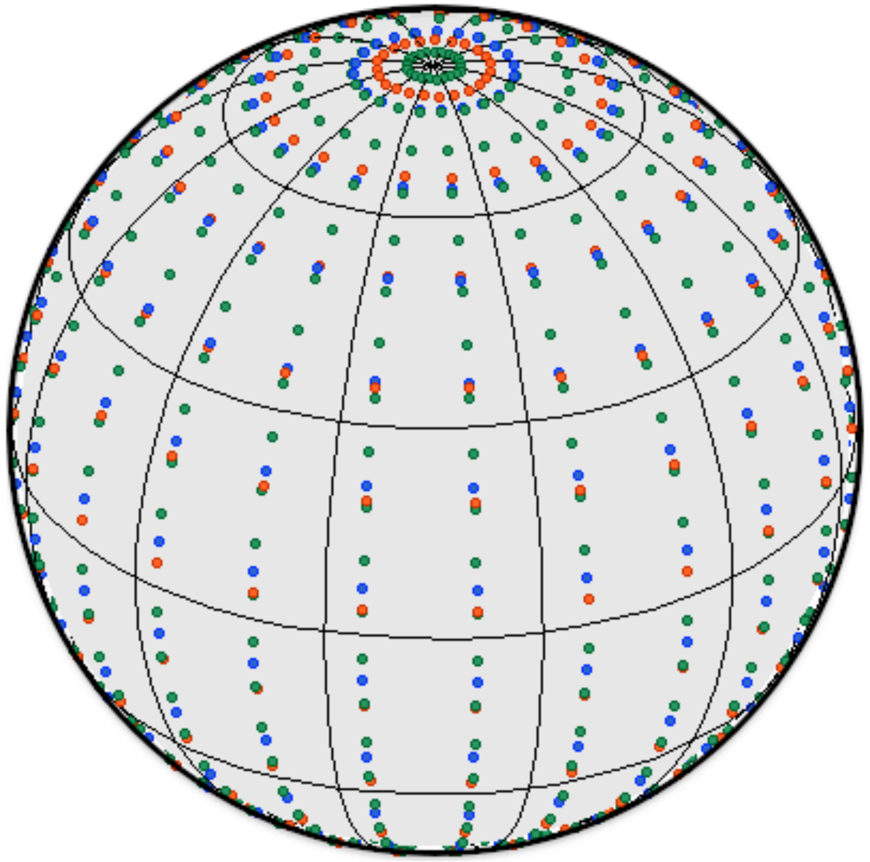

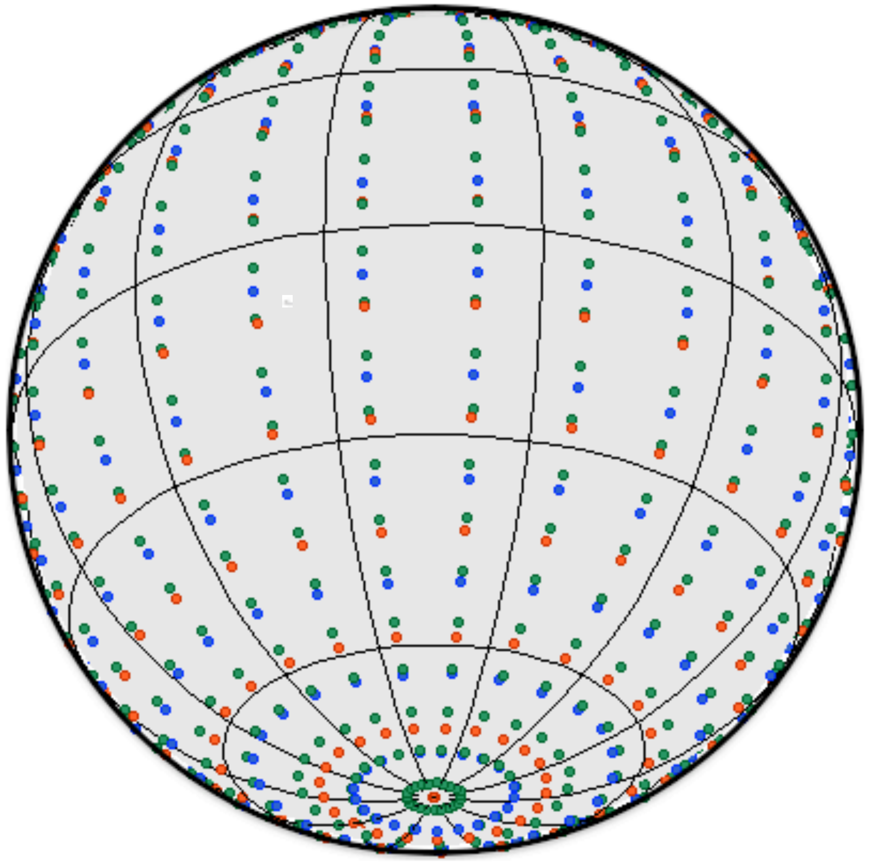

Novel sampling theorem

* The factorization of rotations implies a Fourier expansion of Wigner-d functions, so that spherical harmonic coefficients can be obtained from Fourier coefficients...

$$d^{\ell}_{mn}(\beta) = \mathrm{i}^{n-m} \sum_{m'=-\ell}^{\ell} \Delta^{\ell}_{m'm}\, \Delta^{\ell}_{m'n}\, \mathrm{e}^{\mathrm{i}m'\beta}$$

Hence $${}_sf_{\ell m} = (-1)^s \mathrm{i}^{m+s} \sqrt{\frac{2\ell+1}{4\pi}} \sum_{m'=-(L-1)}^{L-1} \Delta^{\ell}_{m'm} \Delta^{\ell}_{m',-s}\, {}_sG_{mm'} \quad \Big|\, \mathcal{O}(L^3)$$

For $${}_sG_{mm'} = \int_0^{\pi} \mathrm{d}\theta \sin\theta\, {}_sG_m(\theta)\, \mathrm{e}^{-\mathrm{i}m'\theta}$$

and $${}_sG_m(\theta) = \int_0^{2\pi} \mathrm{d}\varphi\, {}_sf(\theta,\varphi)\, \mathrm{e}^{-\mathrm{i}m\varphi}$$

Novel sampling theorem

* Extending the functions by symmetry to the torus does it all...

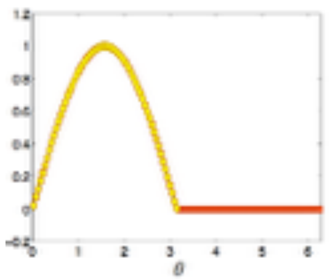

$$ {}_s\tilde{G}_m(\theta) = \begin{cases} {}_sG_m(\theta)\,, & \theta \in [0,\pi] \\ (-1)^{m+s}\, {}_sG_m(2\pi-\theta)\,, & \theta \in (\pi, 2\pi) \end{cases} $$

$$ {}_sG_{mm'} = \int_0^{\pi} \mathrm{d}\theta \sin\theta \, {}_s\tilde{G}_m(\theta)\, \mathrm{e}^{-\mathrm{i}m'\theta} $$

$$ = 2\pi \sum_{m''=-(L-1)}^{L-1} {}_sF_{mm''}\, w(m''-m') $$

$$ w(m') = \int_0^{\pi} \mathrm{d}\theta \sin\theta \, \mathrm{e}^{\mathrm{i}m'\theta} = \begin{cases} \pm \mathrm{i}\pi/2, & m' = \pm 1 \\ 0, & m' \text{ odd},\ m' \neq \pm 1 \\ 2/(1-m'^2), & m' \text{ even} \end{cases} $$

2D Fourier coefficients of ${}_sf$

Forward harmonic transform

* The algorithm provides an exact implicit quadrature rule for the harmonic coefficients:

procedure FORWARD TRANSFORM(${}_sf$)
    compute the Fourier transform of ${}_sf$ in φ
    extend the resultant function to 2π in θ
    upsample the resultant function in θ
    multiply by the inverse Fourier transform of the reflected weights and take the Fourier transform in θ to give the coefficients ${}_sG_{mm'}$
    compute the spherical harmonic coefficients ${}_sf_{\ell m}$ from ${}_sG_{mm'}$
    return ${}_sf_{\ell m}$
end procedure

BASP

HERIOT WATT UNIVERSITY

Existing sampling theorems

* Required numbers of sampling points...

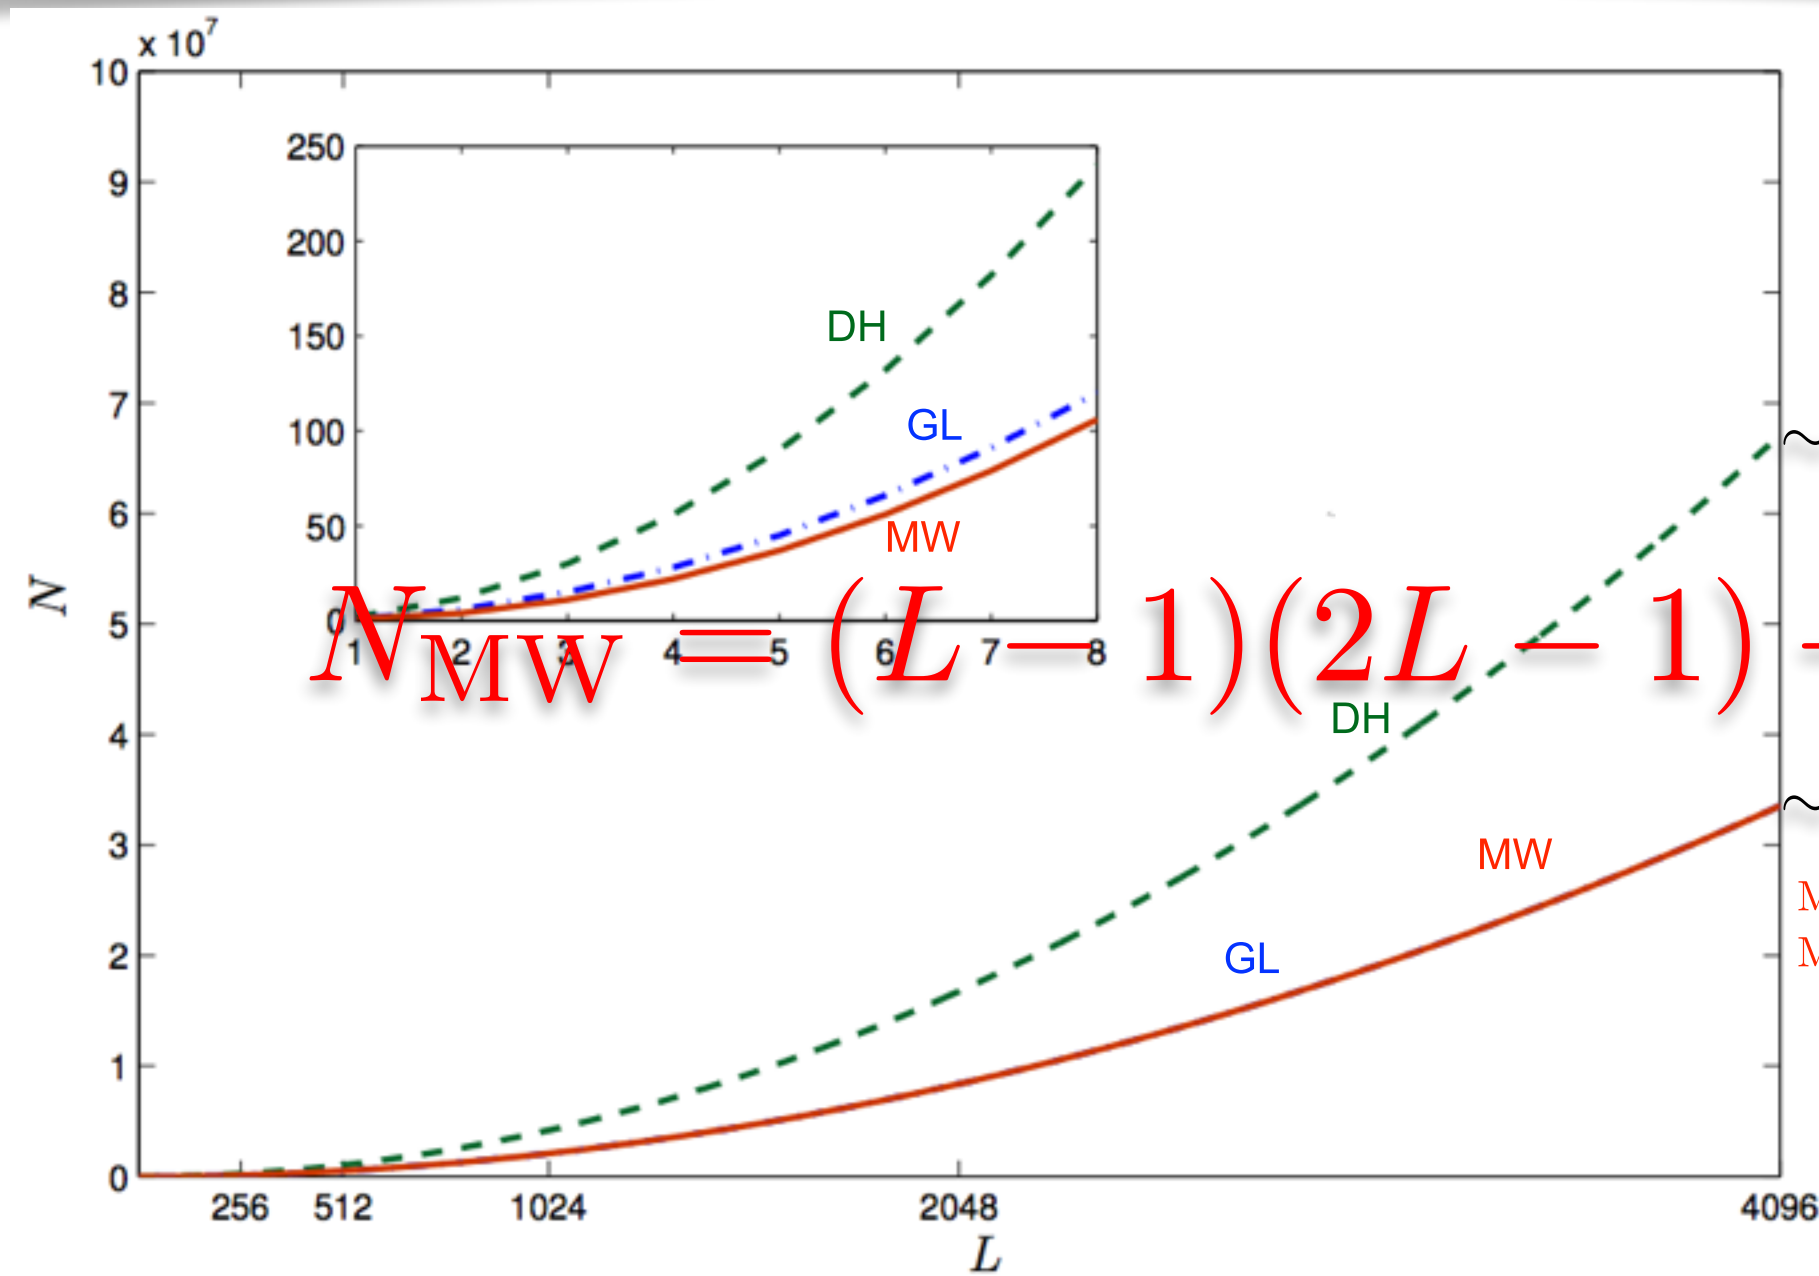

$$N_{\mathrm{MW}} = (L-1)(2L-1)+1$$

MW: [Stable] Equiangular
McEwen & Wiaux Sampling

Inverse harmonic transform

* The inverse transform algorithm proceeds from the same symmetrization arguments:

```
procedure INVERSE TRANSFORM( sf_ℓm )
    compute the Fourier coefficients sF_mm' from sf_ℓm
    compute the function samples on the extended
     domain by an inverse Fourier transform
    construct sf by discarding samples computed in the
     θ domain (π, 2π)
    return sf
end procedure
```

procedure INVERSE TRANSFORM(${}_sf_{\ell m}$)

compute the Fourier coefficients ${}_sF_{mm'}$ from ${}_sf_{\ell m}$

compute the function samples on the extended domain by an inverse Fourier transform

construct ${}_sf$ by discarding samples computed in the θ domain $(\pi, 2\pi)$

return ${}_sf$

end procedure

Algorithm exactness

* The algorithm is stable up to at least up to 4096...

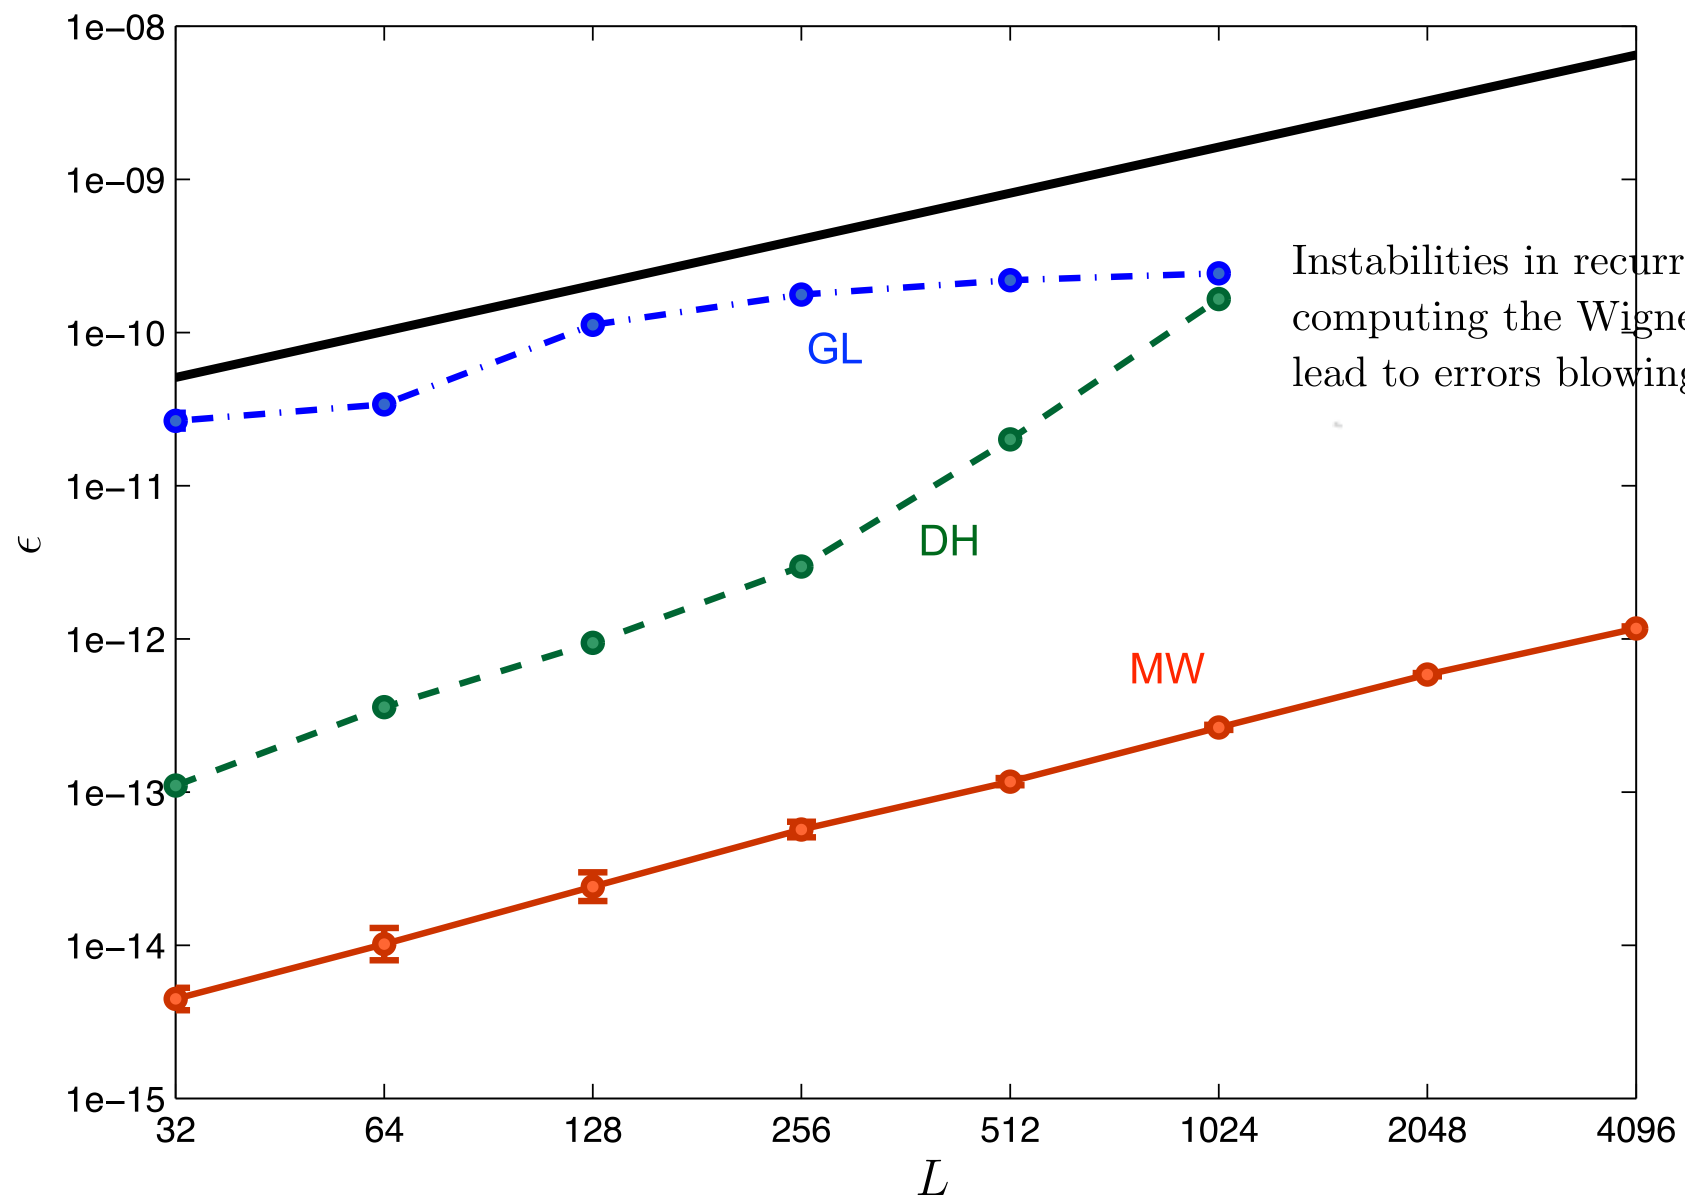

Instabilities in recurrence relations computing the Wigner-d functions lead to errors blowing up.

Algorithm complexity and speed

* The continual use of FFTs makes the algorithm really fast...

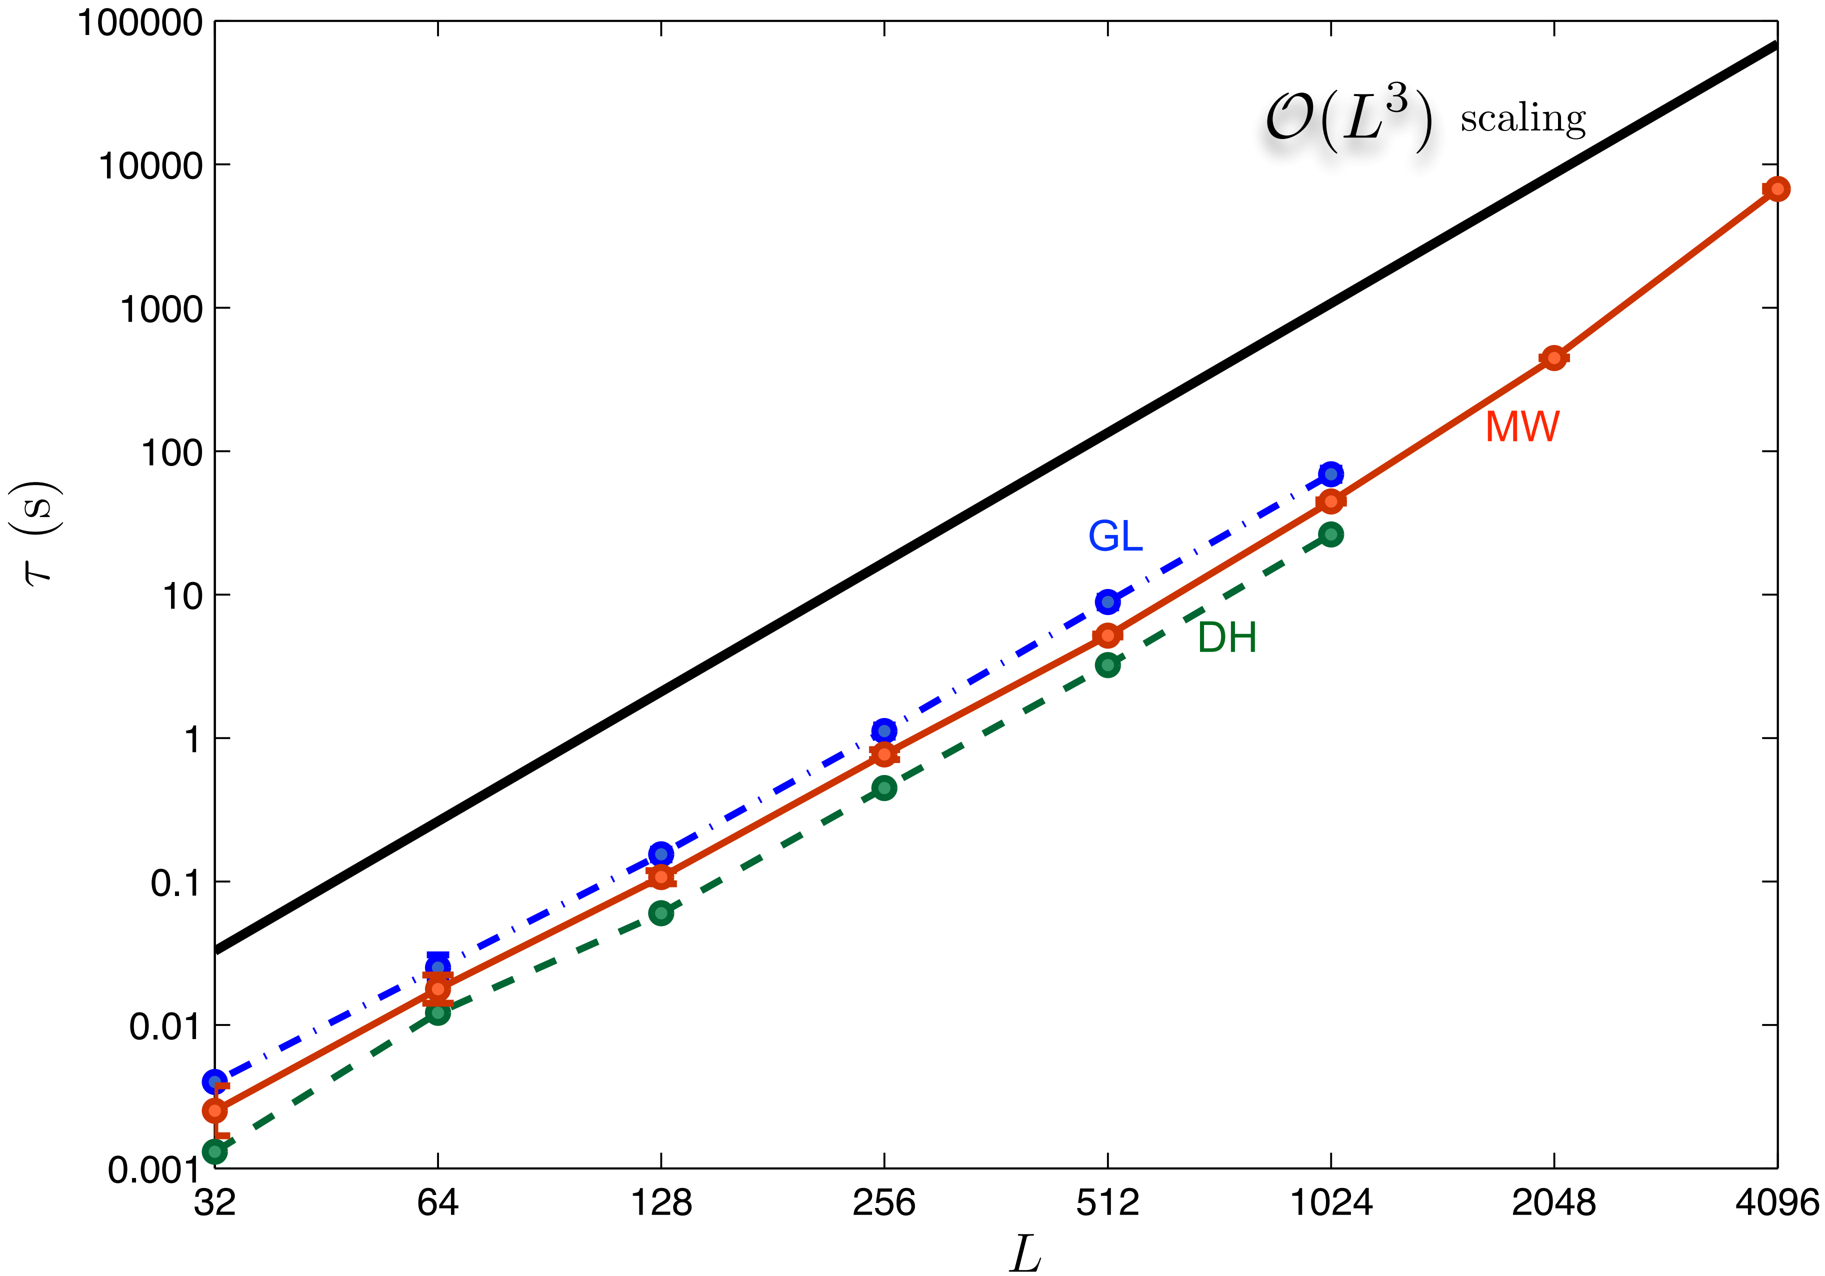

HERIOT WATT UNIVERSITY

Quadrature rule

* The algorithm is also shipped with an explicit quadrature rule for integration from $\sim L^2$ points ...

$$I = \int_{S^2} \mathrm{d}\Omega(\theta,\varphi)\, {}_sf(\theta,\varphi) = \sum_{t=0}^{L-1} \sum_{p=0}^{L-1} {}_sf(\theta_t,\varphi'_p)\, q(\theta_t)$$

with $$q(\theta_t) = \frac{2\pi}{L}\left[v(\theta_t) + (1-\delta_{t,L-1})\,(-1)^s\, v(\theta_{2L-2-t})\right]$$

and $$v(\theta_t) = \frac{1}{2L-1} \sum_{m'=-(L-1)}^{L-1} w(-m')\, \mathrm{e}^{\mathrm{i}m'\theta_t}$$

II.
Compressed sensing

TV inpainting illustration

* Earth topography map at a band limit $L = 32$ ...

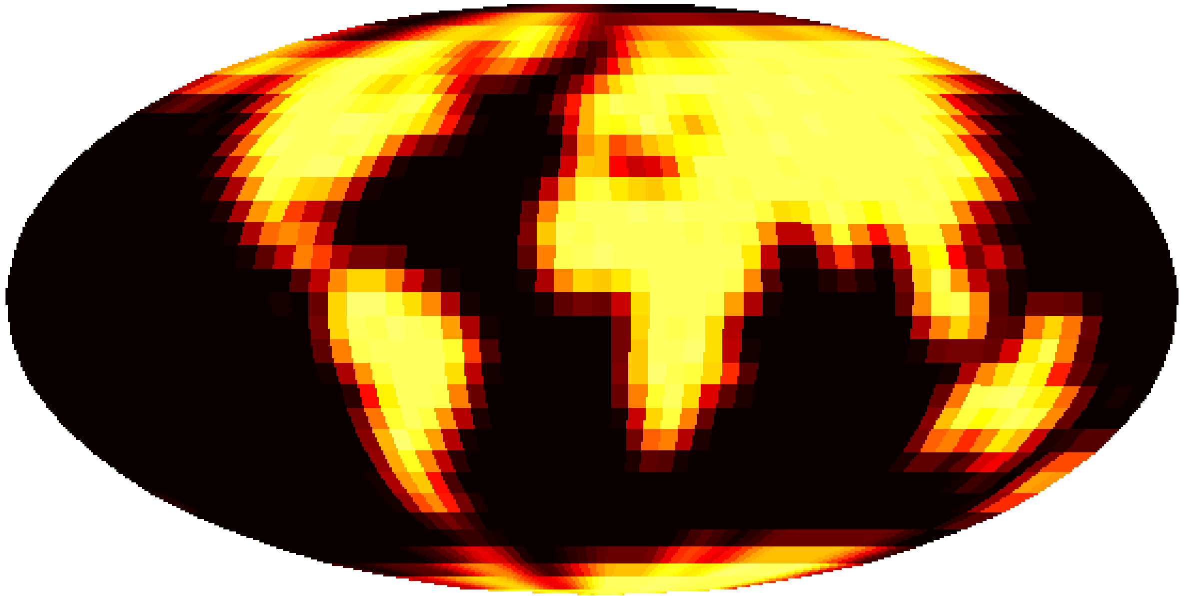

Original map

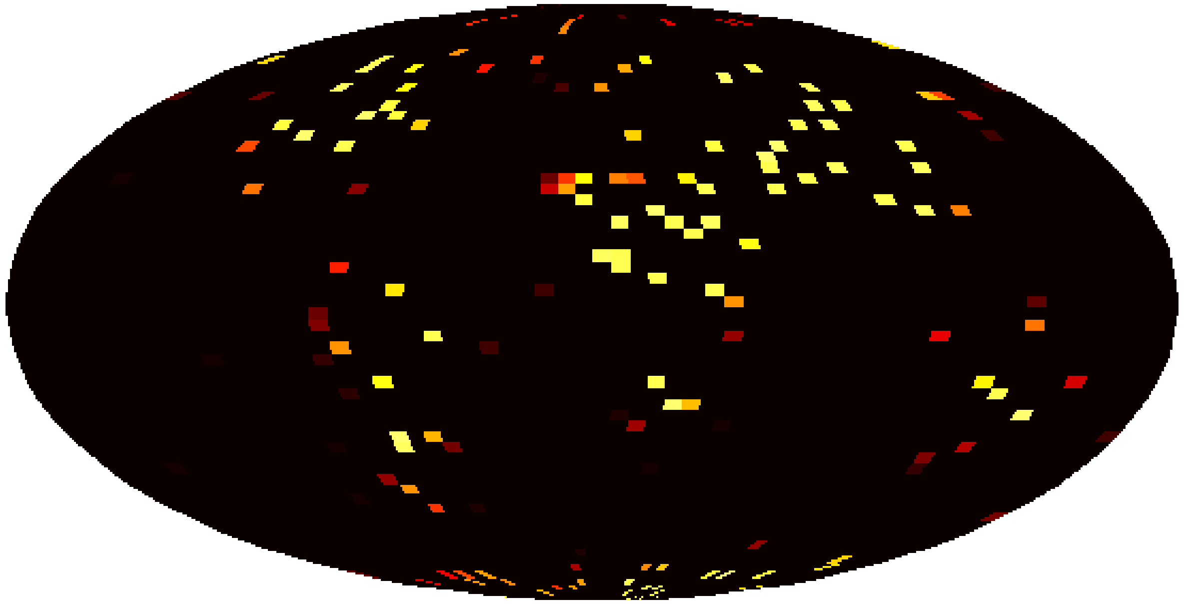

Incomplete measurements taken at random: compressed sensing!

The inverse problem

* The measurement operator Φ simply consists in a selection operator...

$$\boldsymbol{y} = \Phi \boldsymbol{x} + \boldsymbol{n}$$

For $\boldsymbol{x} \in \mathbb{R}^N$

$\boldsymbol{y} \in \mathbb{R}^M$ with $M < N_{\mathrm{MW}}$

$\Phi \in \mathbb{R}^{M \times N}$ (simple selection matrix)

$\boldsymbol{n} \in \mathbb{R}^M$

Sparsity

* The signal is sparse in the magnitude of its spherical gradient. The TV norm is defined as a continuous norm on the sphere... Also, sparsity K will be minimized on grids with minimum number of samples!

$$\|x\|_{\mathrm{TV}} \equiv \int_{\mathrm{S}^2} \mathrm{d}\Omega \, |\nabla x|$$

for the spherical gradient $$|\nabla x| \simeq \sqrt{(\delta_\theta \boldsymbol{x})^2 + \frac{1}{\sin^2\theta_t}(\delta_\varphi \boldsymbol{x})^2}$$

From the quadrature rule, the TV norm reads as a weighted l1-norm of the gradient:

$$\|x\|_{\mathrm{TV}} \simeq \sum_{t=0}^{N_\theta - 1} \sum_{p=0}^{N_\varphi - 1} |\nabla x| \, q(\theta_t)$$

Dimensionality

* The minimization problem can be posed in the spatial domain for dimensionality $N > L^2$, or in the harmonic domain for improved dimensionality L^2 !

Spatial setting:

$$\boldsymbol{x}^\star = \arg\min_{\boldsymbol{x}} \|\boldsymbol{x}\|_{\mathrm{TV}} \text{ such that } \|\boldsymbol{y} - \Phi \boldsymbol{x}\|_2 \leq \epsilon$$

Harmonic setting:

$$\hat{\boldsymbol{x}}^\star = \arg\min_{\hat{\boldsymbol{x}}} \|\Lambda \hat{\boldsymbol{x}}\|_{\mathrm{TV}} \text{ such that } \|\boldsymbol{y} - \Phi \Lambda \hat{\boldsymbol{x}}\|_2 \leq \epsilon$$

with $\Lambda \in \mathbb{C}^{N \times L^2}$

and $\hat{\boldsymbol{x}} \in \mathbb{C}^{L^2}$ identifying the vector of spherical harmonic coefficients.

Reconstruction results

* In a compressed sensing approach the under-sampling rate scales with sparsity, $M/N \propto K$, hence favoring a setting with both lower dimensionality and sparsity, as confirmed by simulation results:

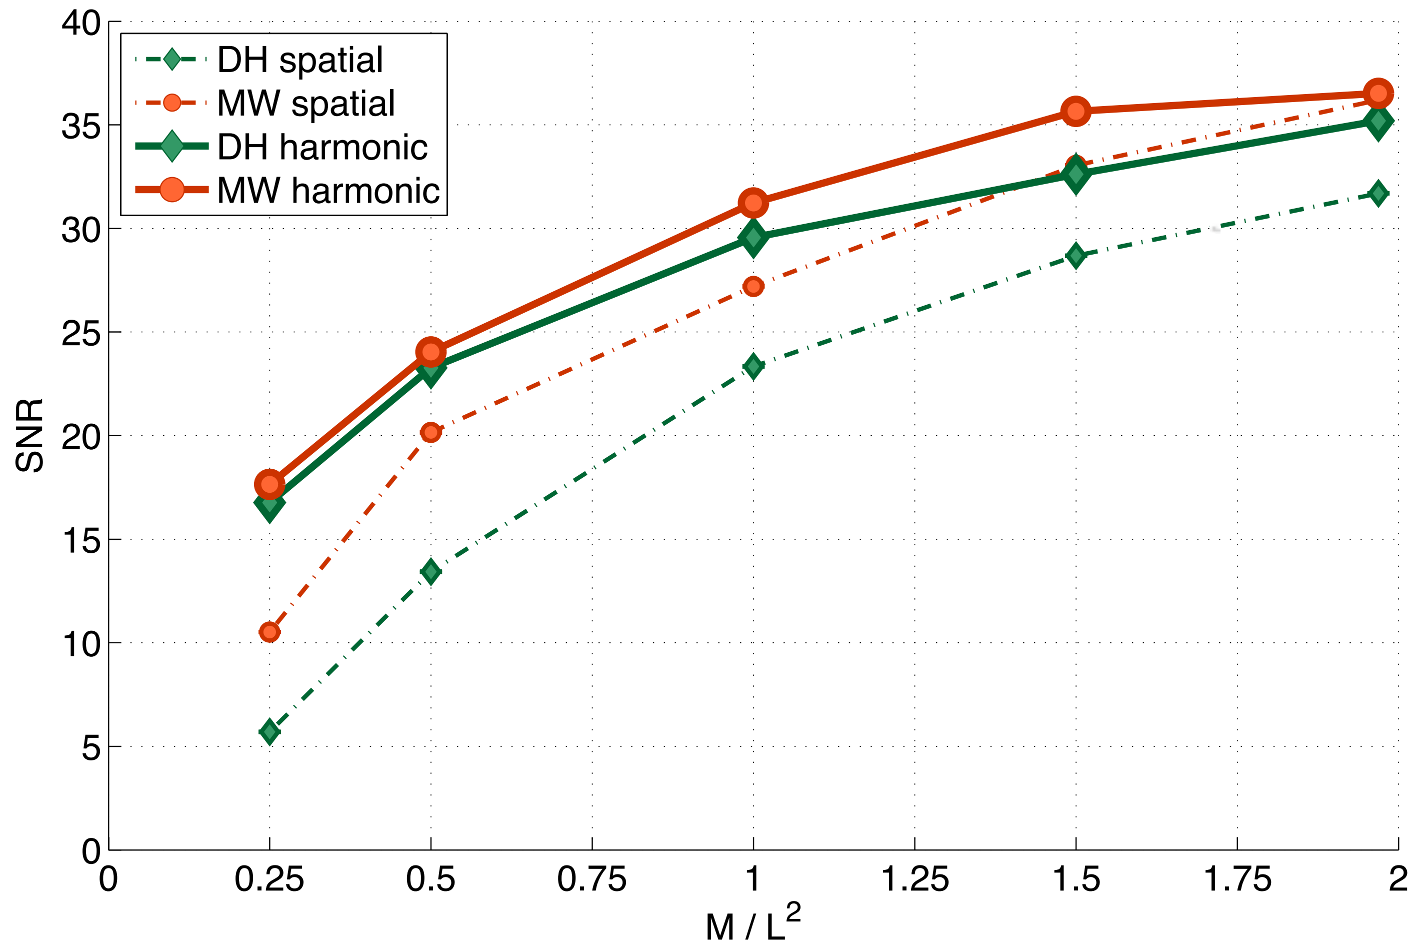

Reconstruction results

* Illustration:

$$M/L^2 = 1/4$$

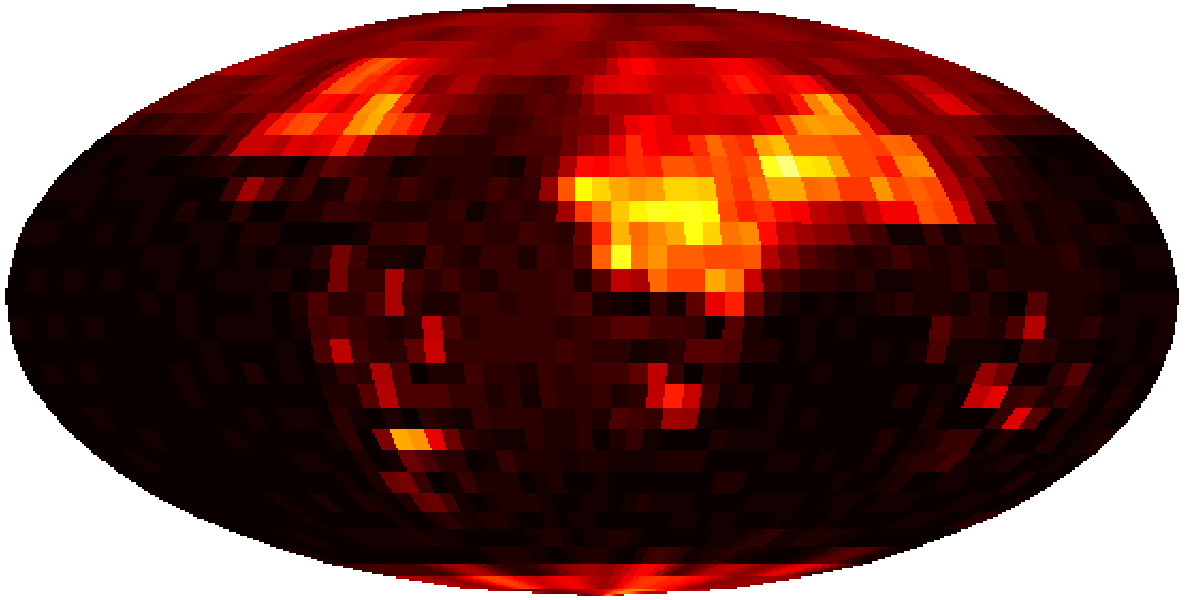

DH, spatial setting

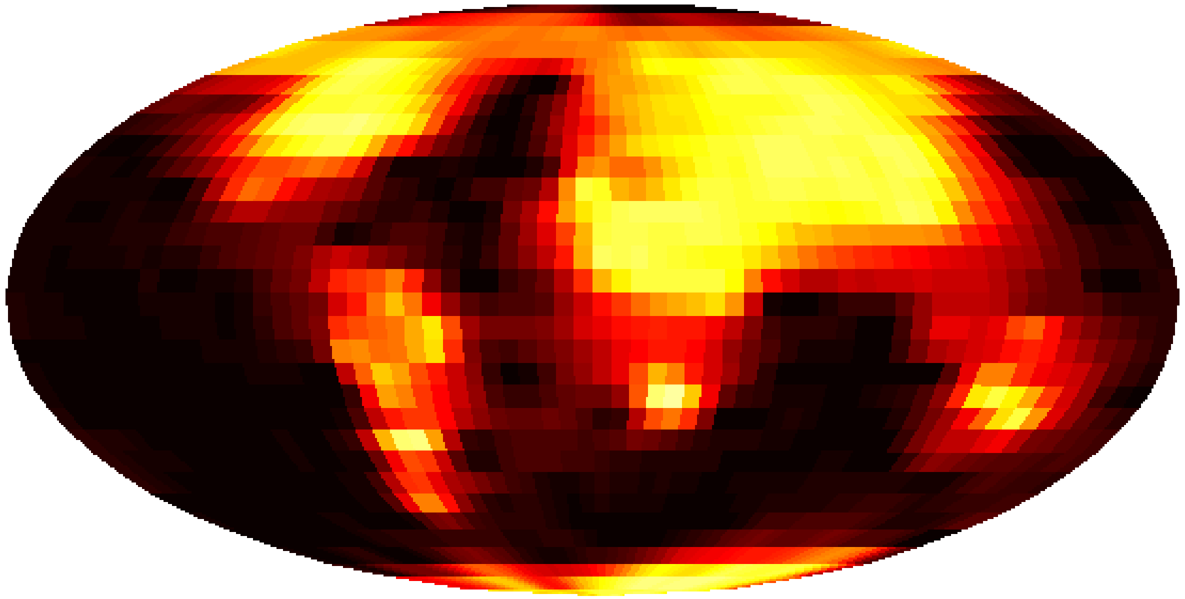

DH, harmonic setting

Reconstruction results

* Illustration:

$M/L^2 = 1/4$

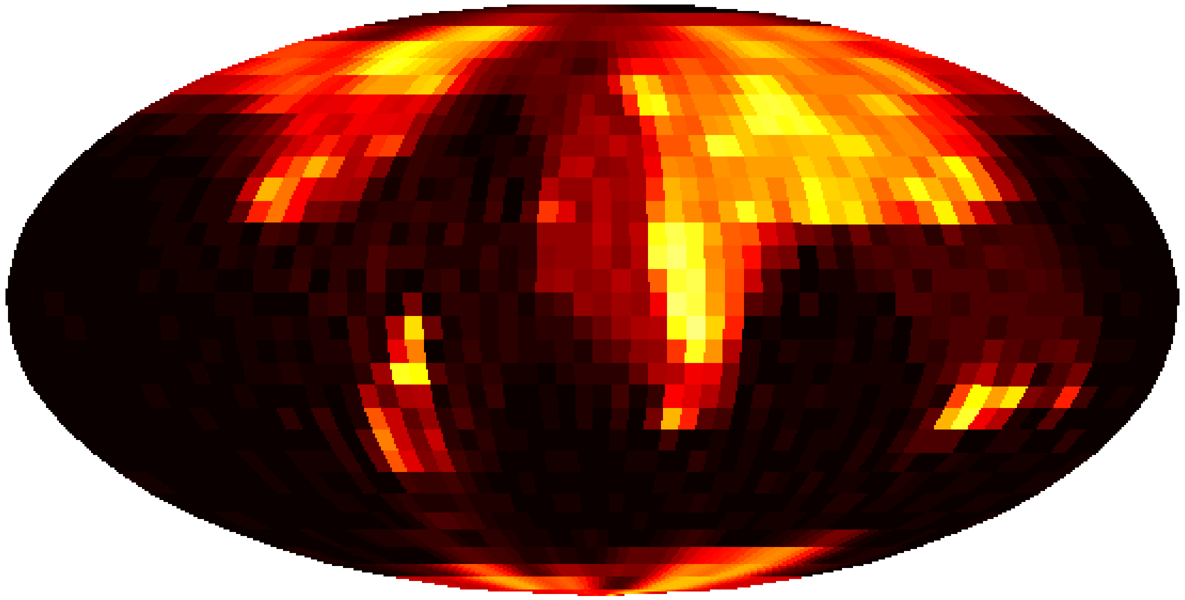

MW, spatial setting

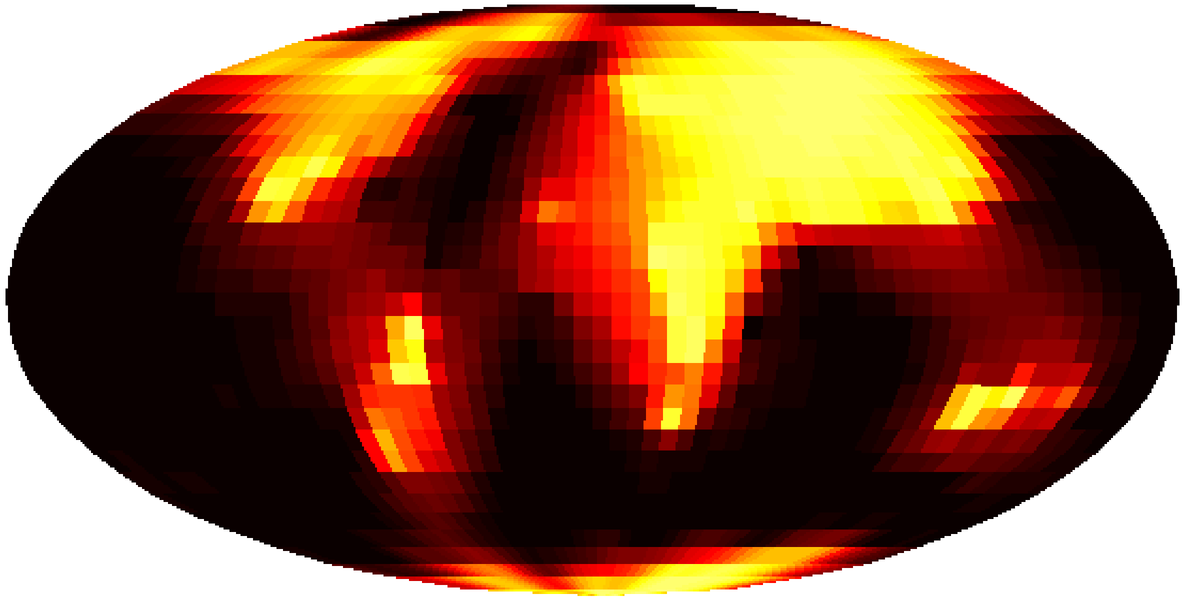

MW, harmonic setting

Reconstruction results

* Illustration:

$$M/L^2 = 1/2$$

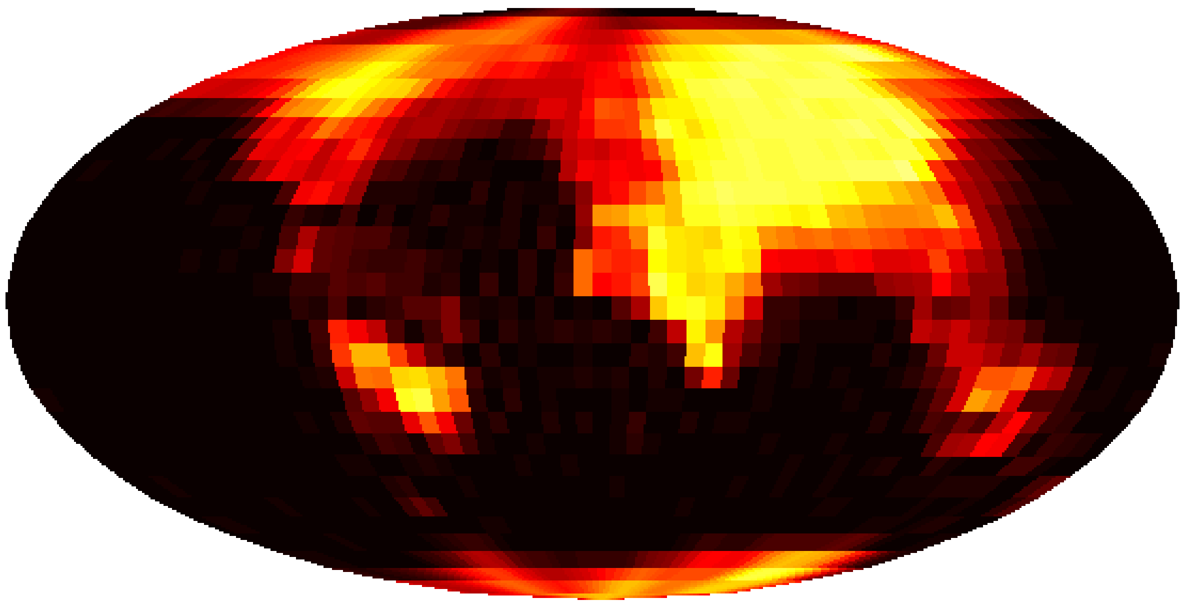
DH, spatial setting

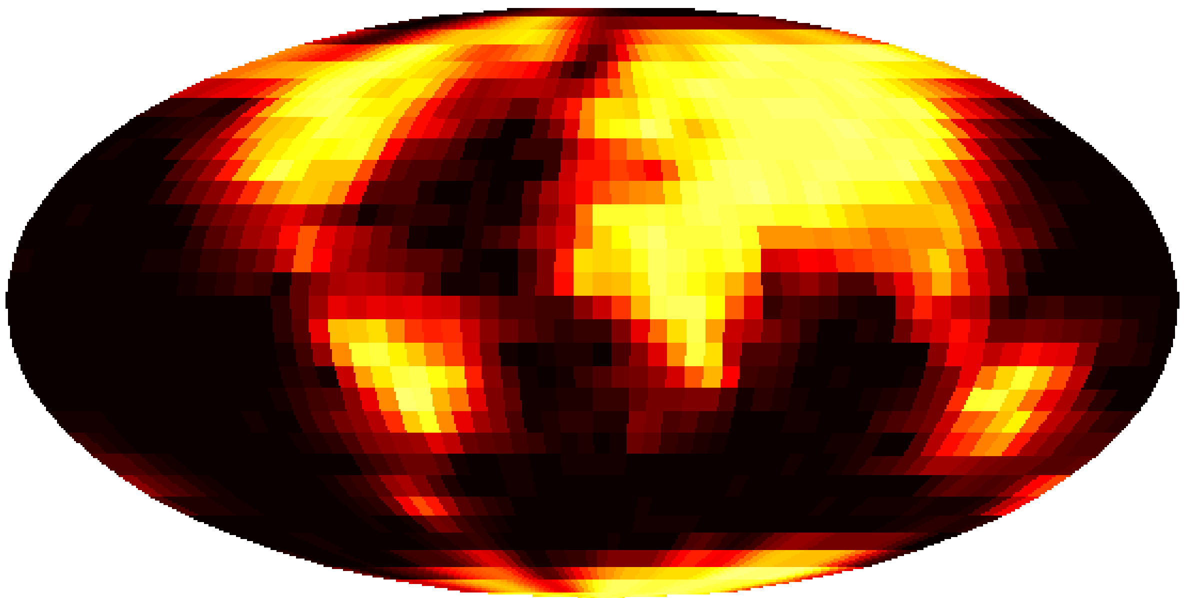
DH, harmonic setting

Reconstruction results

* Illustration:

$M/L^2 = 1/2$

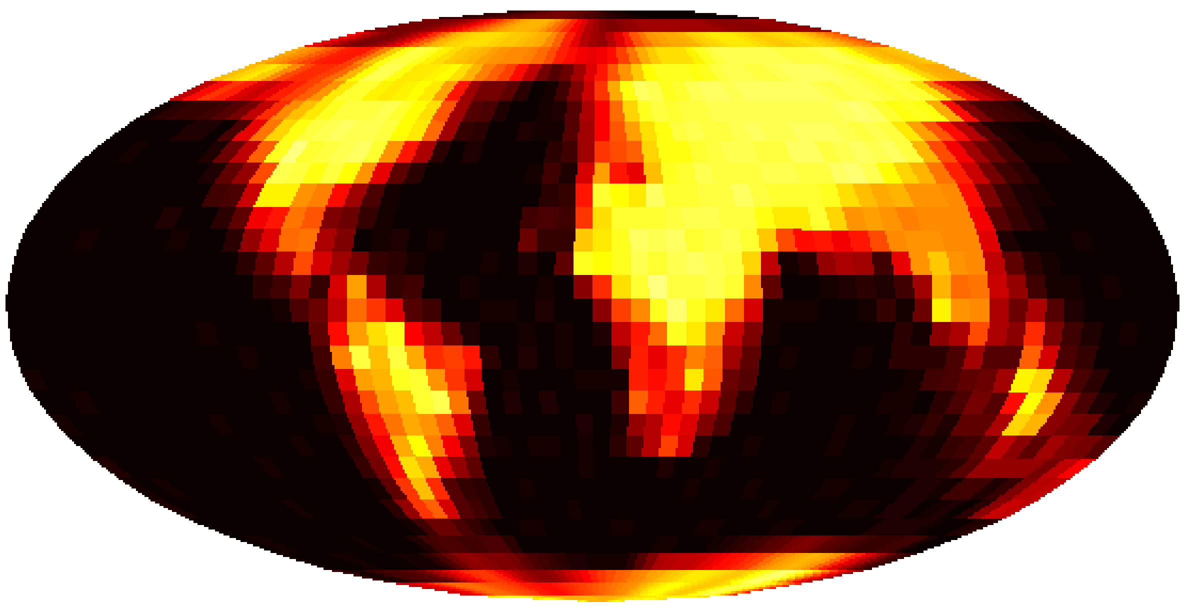

MW, spatial setting

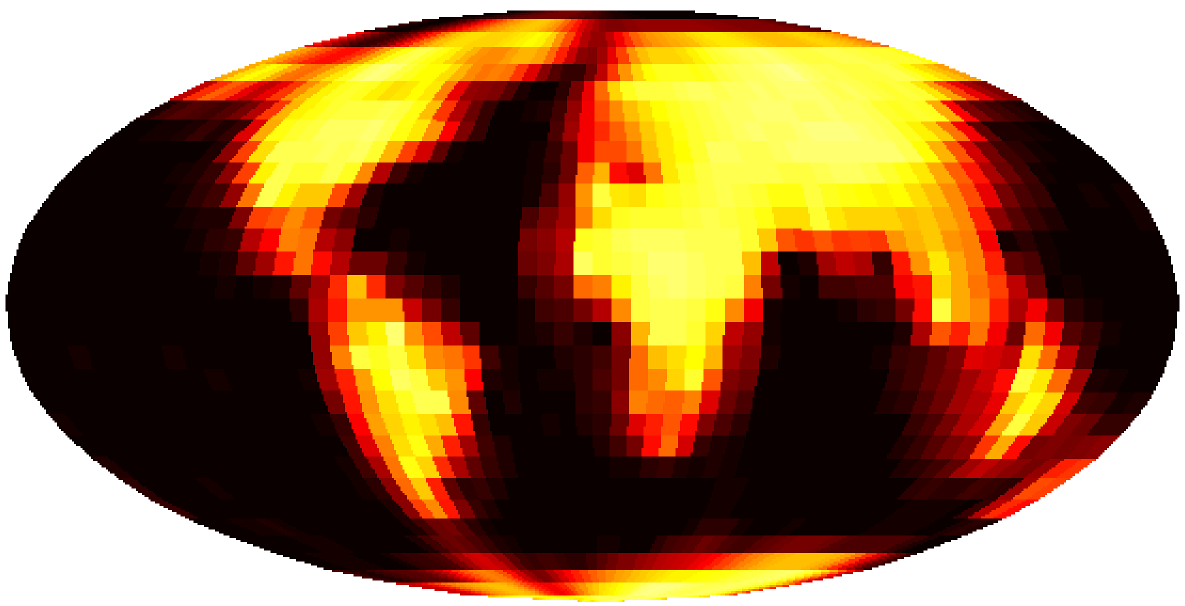

MW, harmonic setting

Reconstruction results

* Illustration:

$$M/L^2 = 1$$

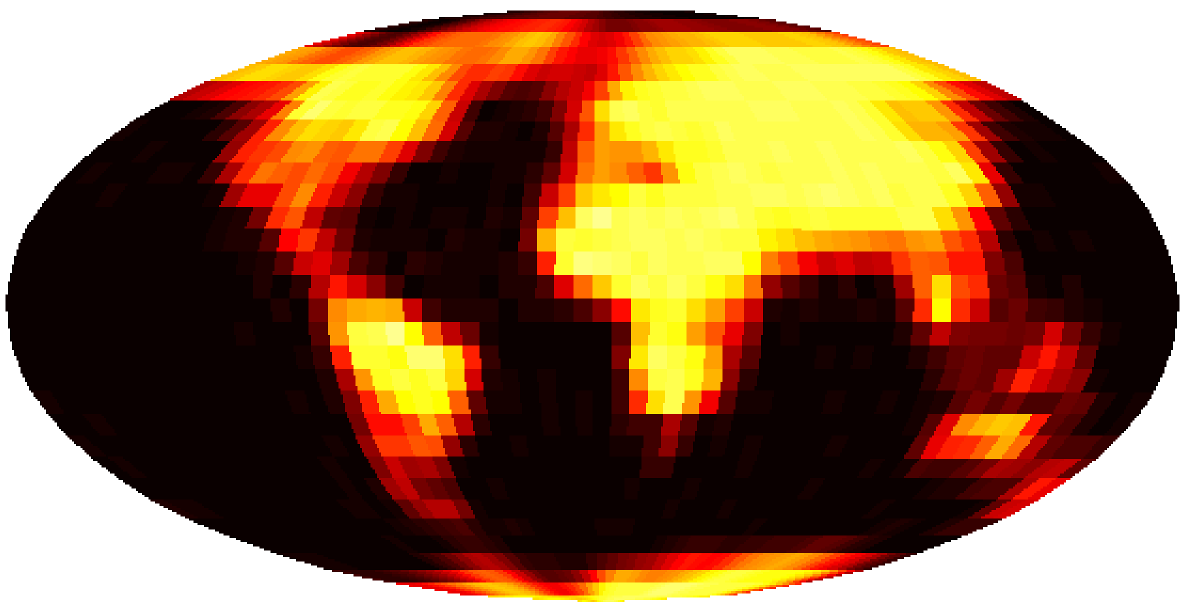

DH, spatial setting

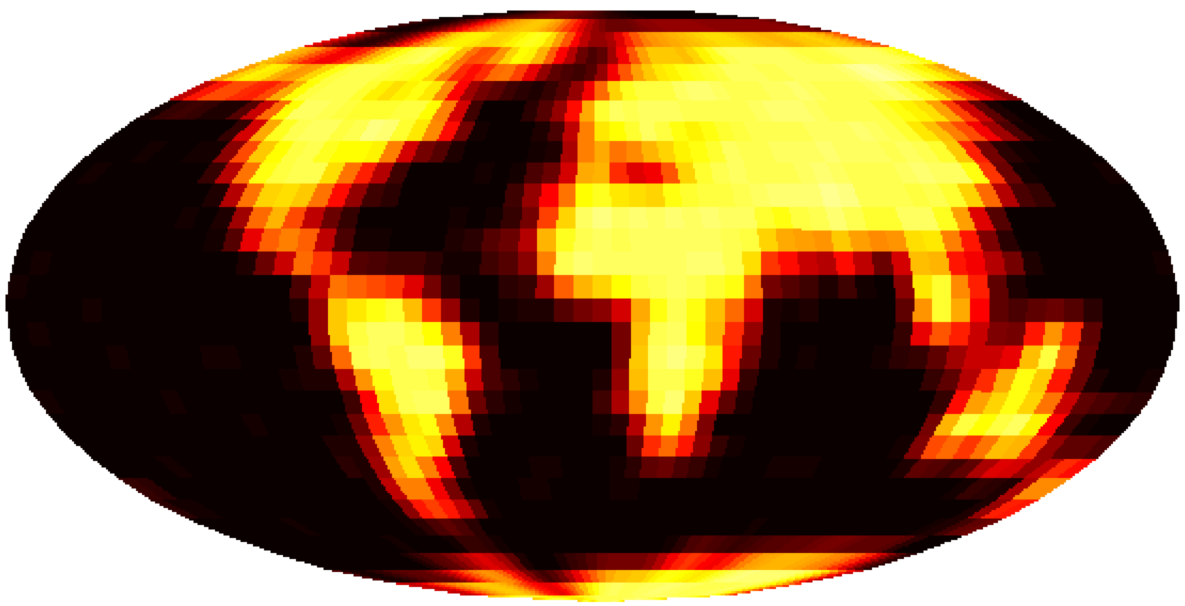

DH, harmonic setting

Reconstruction results

* Illustration:

$M/L^2 = 1$

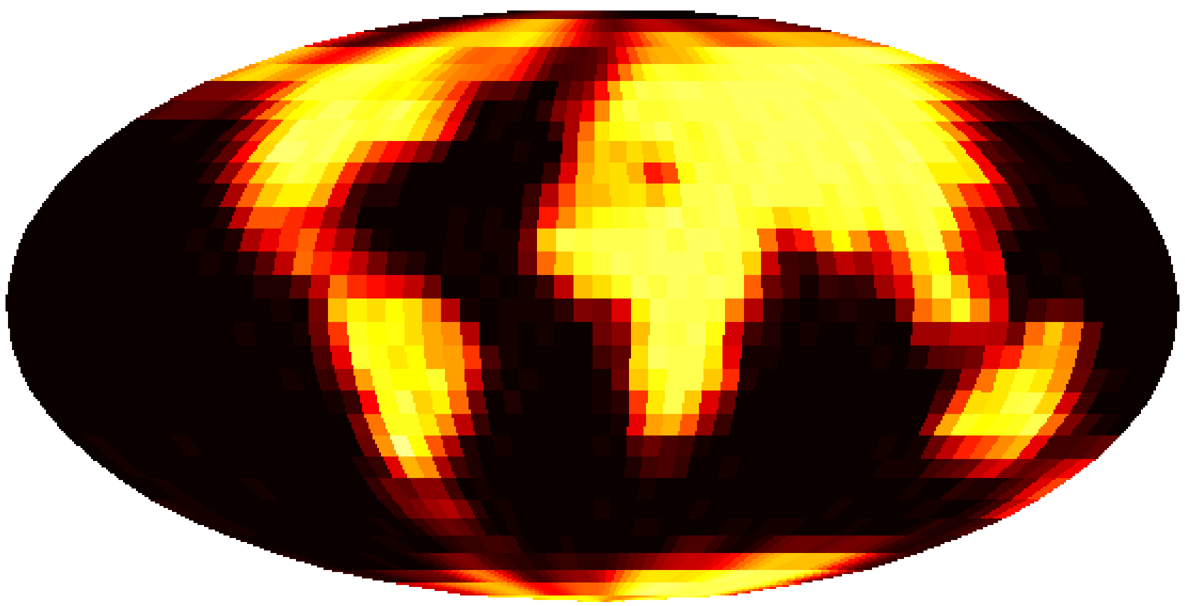

MW, spatial setting

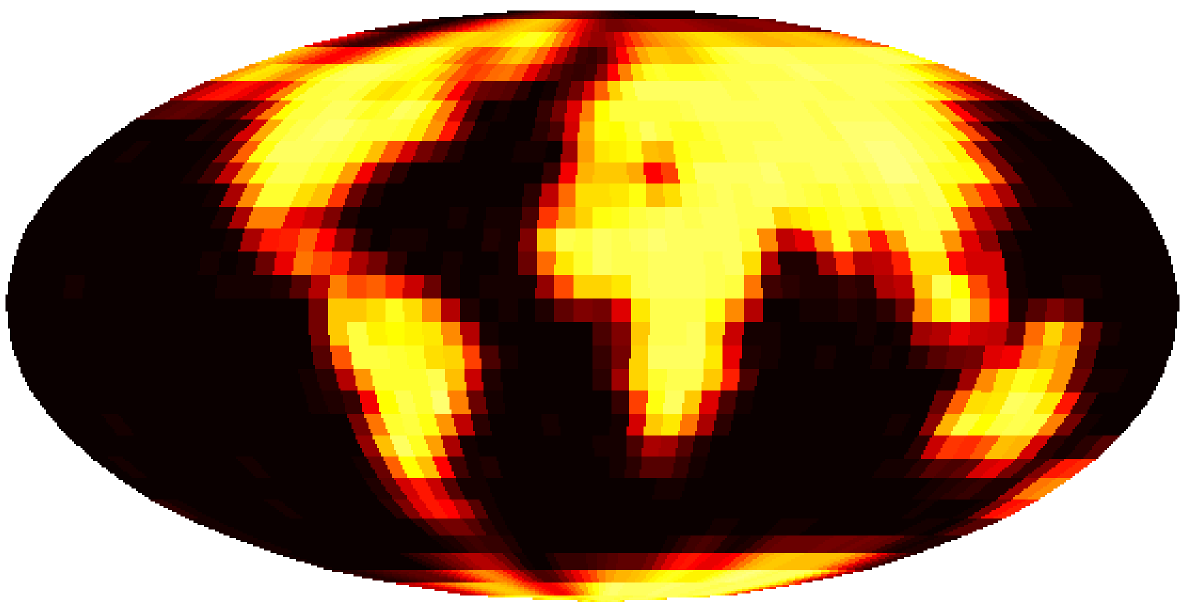

MW, harmonic setting

Reconstruction results

* Illustration:

$$M/L^2 = N_{\rm MW}/L^2 \sim 2$$

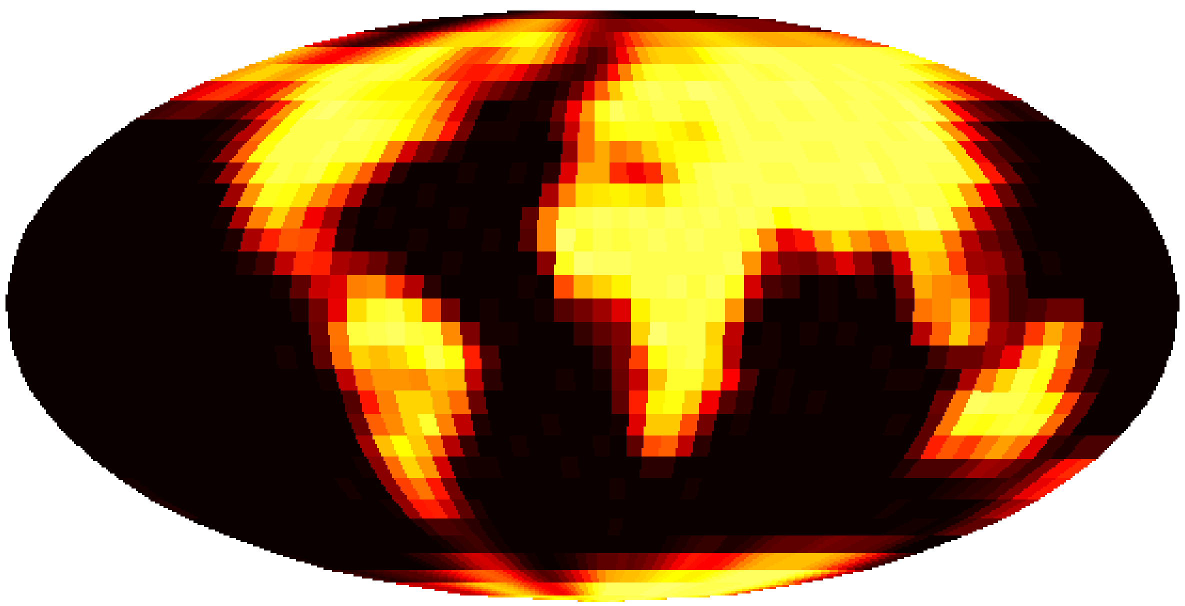

DH, spatial setting

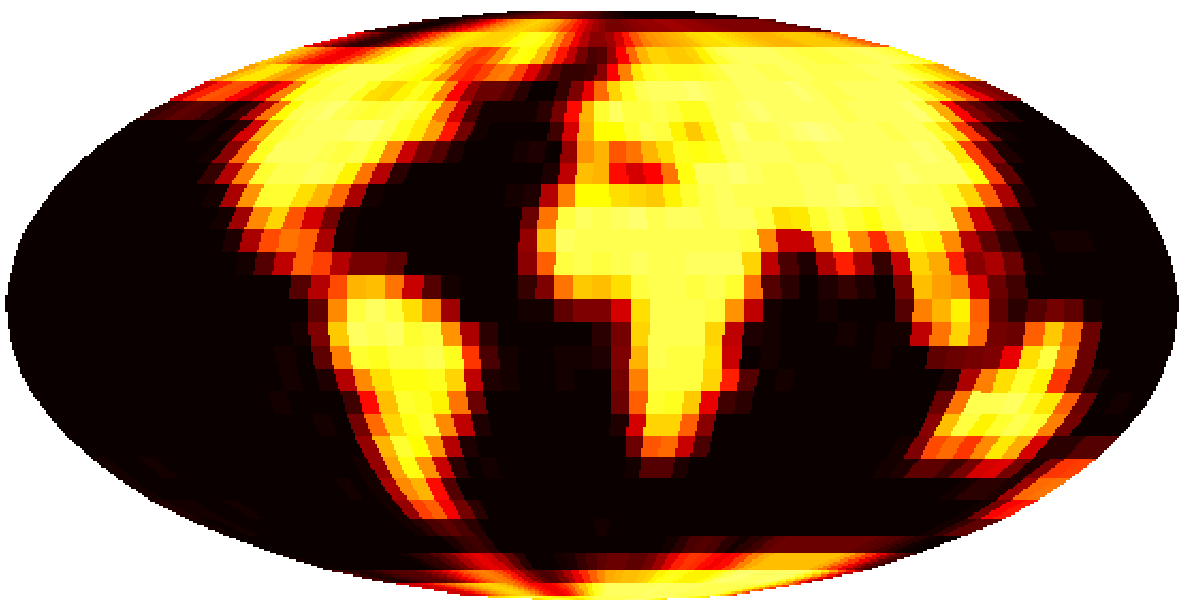

DH, harmonic setting

Reconstruction results

* Illustration:

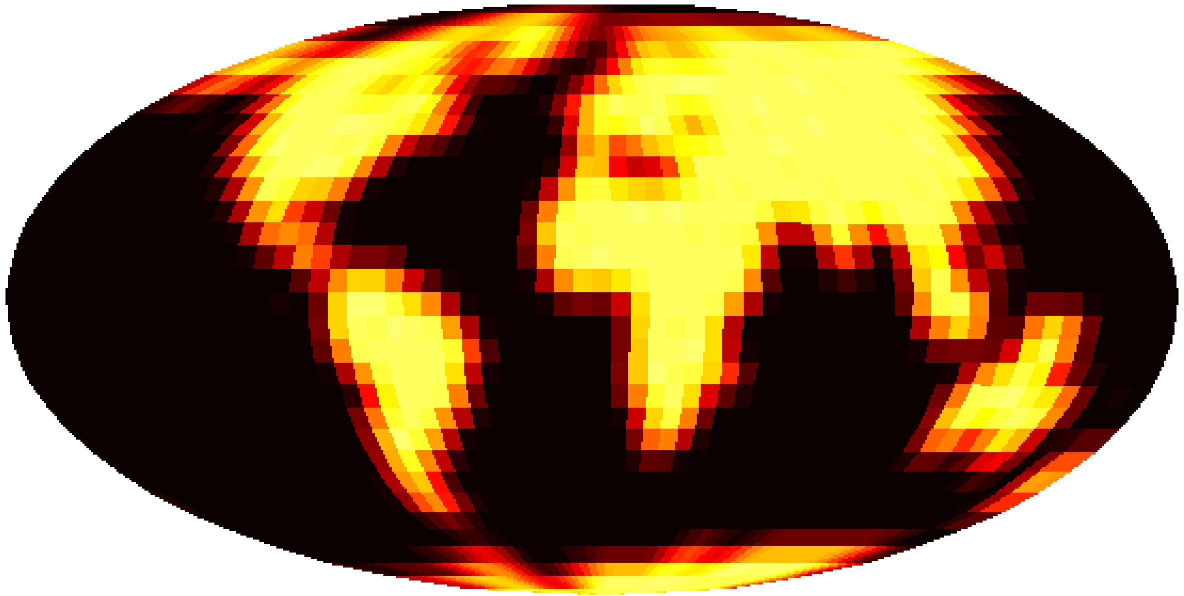

Original map

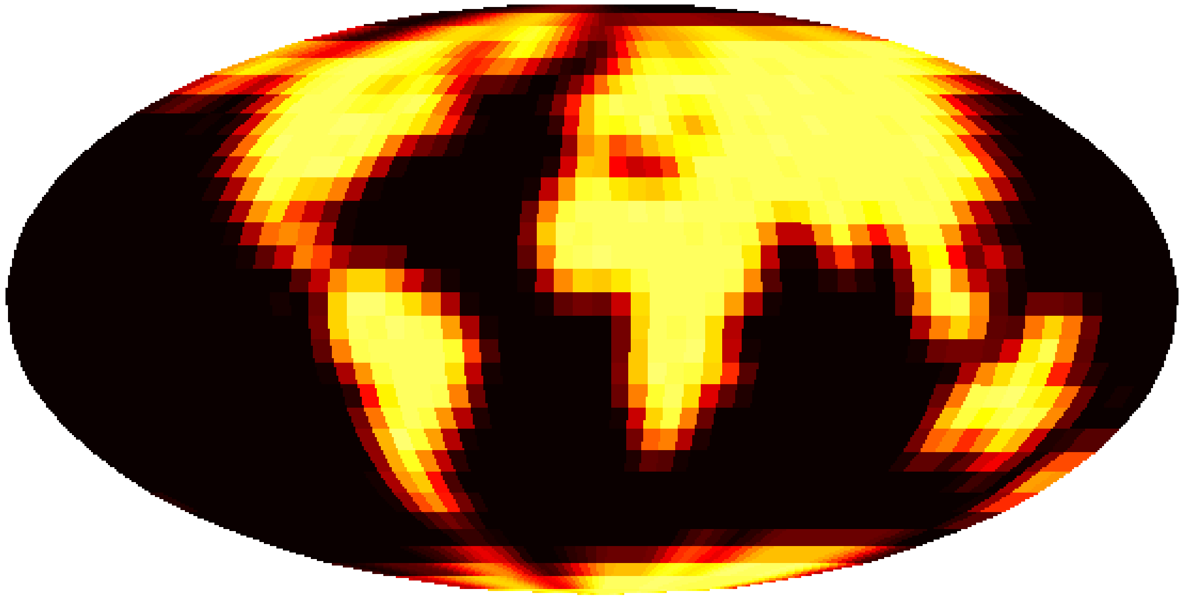

Original map

Reconstruction results

* Illustration:

$$M/L^2 = N_{\mathrm{MW}}/L^2 \sim 2$$

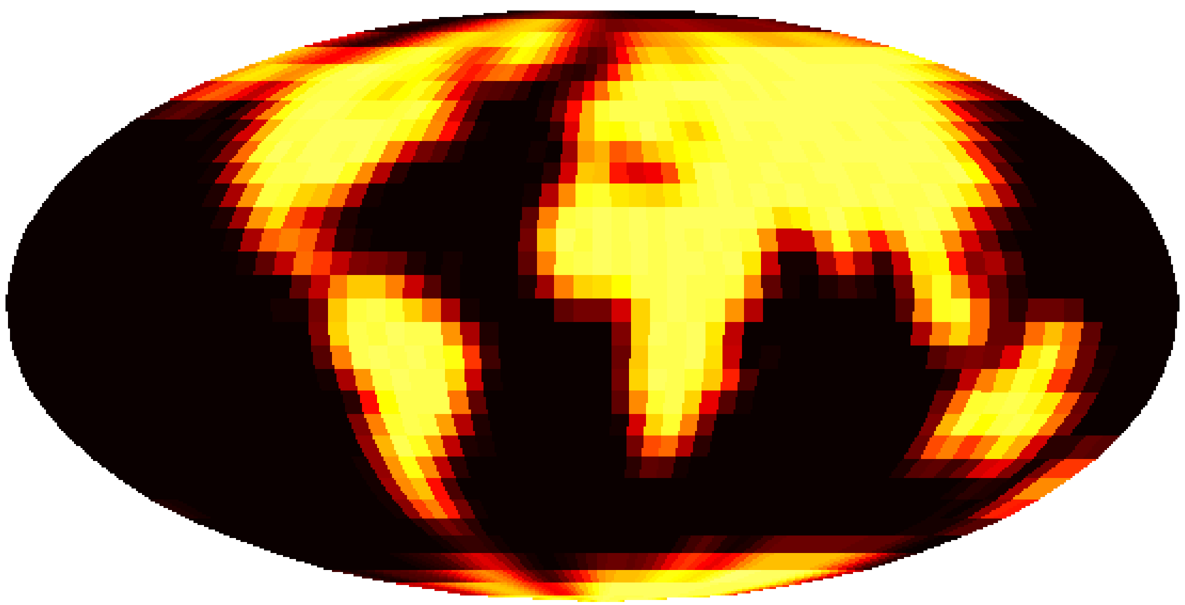

MW, spatial setting

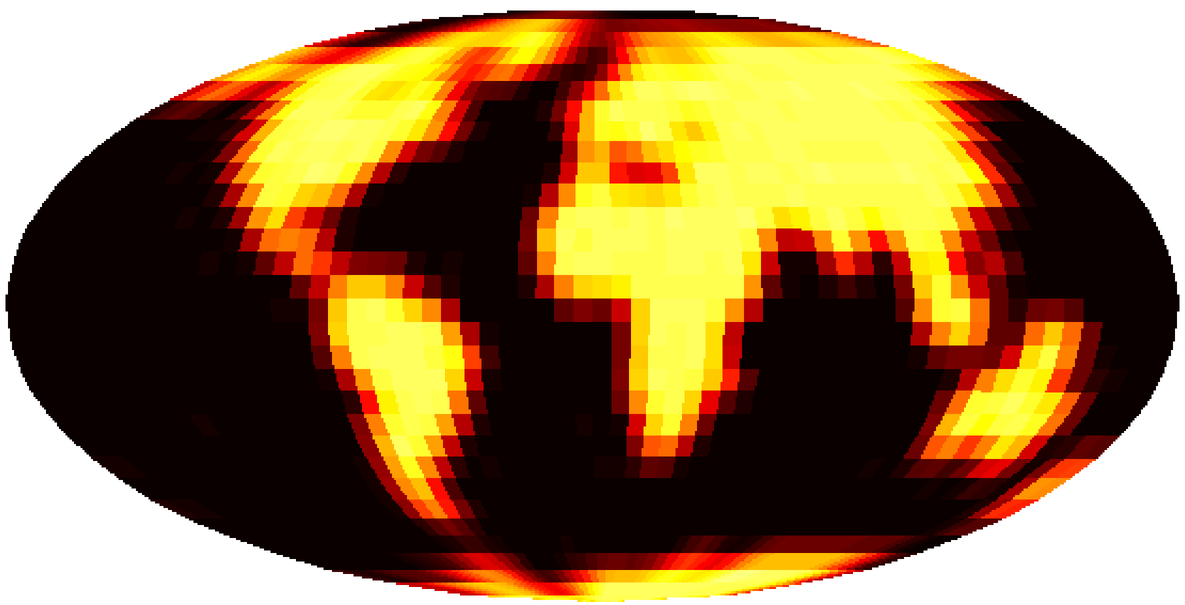

MW, harmonic setting

Reconstruction results

* Illustration:

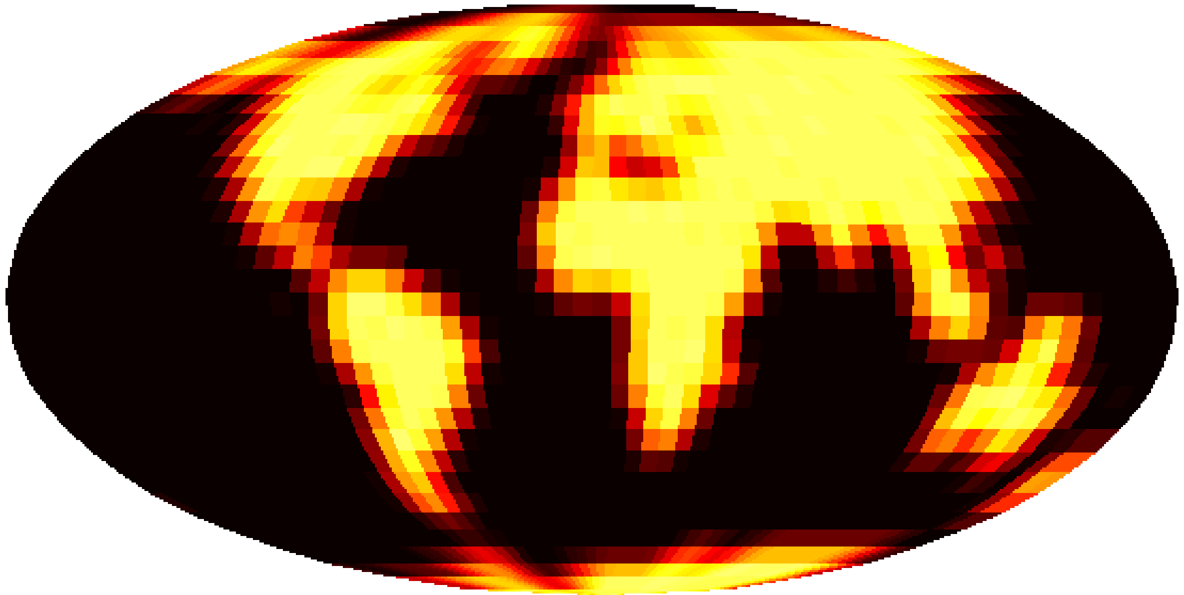

Original map

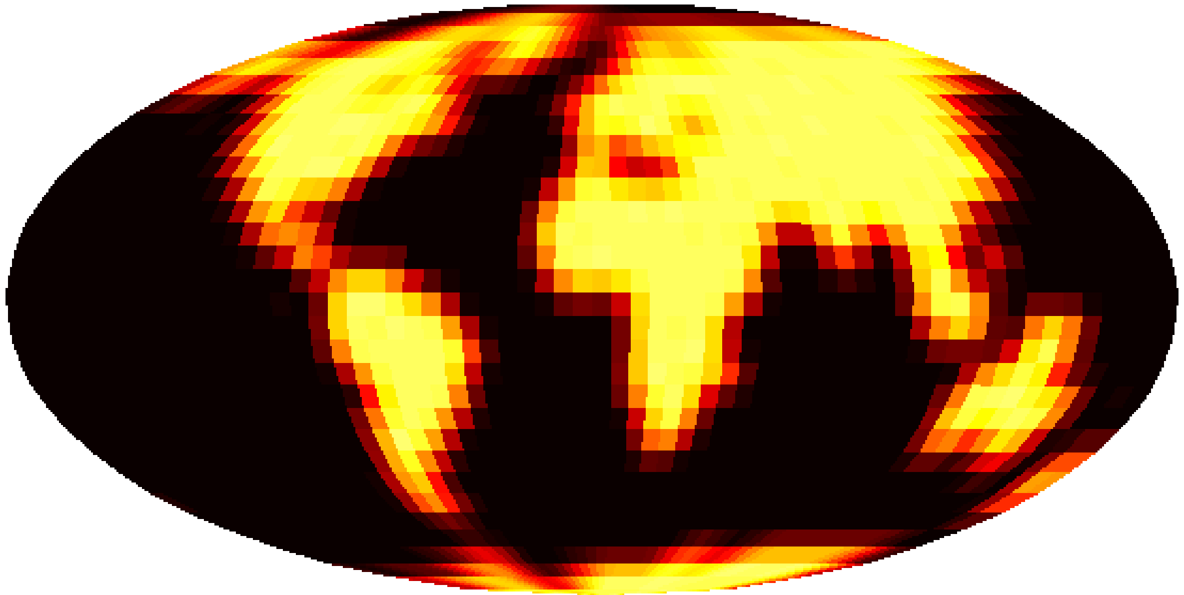

Original map

Conclusion

Take-home messages

We have introduced a novel sampling theorem on equiangular grids on the sphere requiring only $\sim 2L^2$ points, shipped with fast ($\mathcal{O}(L^3)$) and exact spherical harmonic transforms, improving the state-of-the-art.

Application: e.g. diffusion MRI, CMB ...

In a compressed sensing perspective, improving the “Nyquist” rate has important implications for dimensionality, when the signal is recovered in the spatial domain, for and sparsity for a class of priors defined in the spatial domain.

Application: e.g. radio interferometry, CMB inpainting ...