Image Reconstruction in Low-dose Cone Beam Computed Tomography

Alexander Meaney, MSc
Prof. Samuli Siltanen, PhD

Computational Inverse Problems Research Group
Department of Mathematics and Statistics
University of Helsinki, Finland
alexander.meaney@helsinki.fi

WORKSHOP:
Wave Physics and Imaging Applications
University of Helsinki
20 May 2022

Cone-beam Computed Tomography (CBCT)

Image: endocare.ca

Cone-beam Computed Tomography (CBCT)

The CT measurement at detector pixel i is defined as

$$m_i = -\log \frac{I_i}{I_0} = \int_{\text{ray path}} f(x, y, z) ds.$$

By discretizing the distribution of attenuation coefficients $f(x, y, z)$ as $f \in \mathbb{R}^n$ we can make a linear model for the X-ray measurements:

$$m = Af + \varepsilon.$$

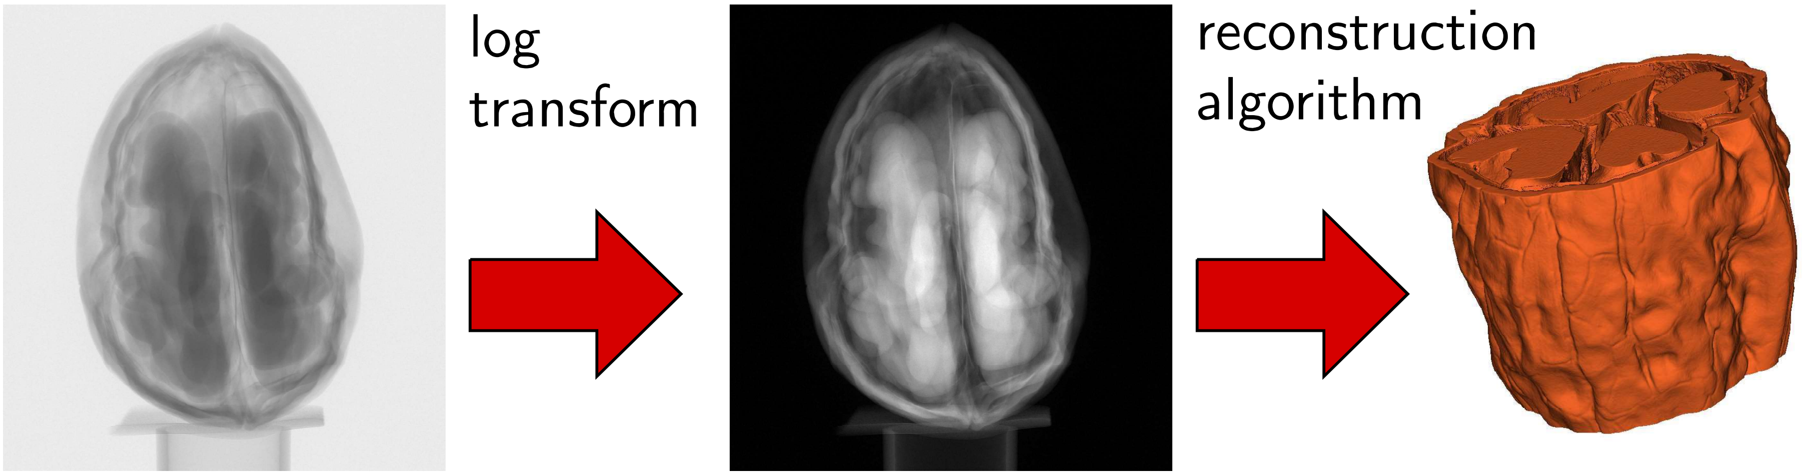

Cone-beam CT is a true 3D imaging modality (FDK reconstruction, isosurface)

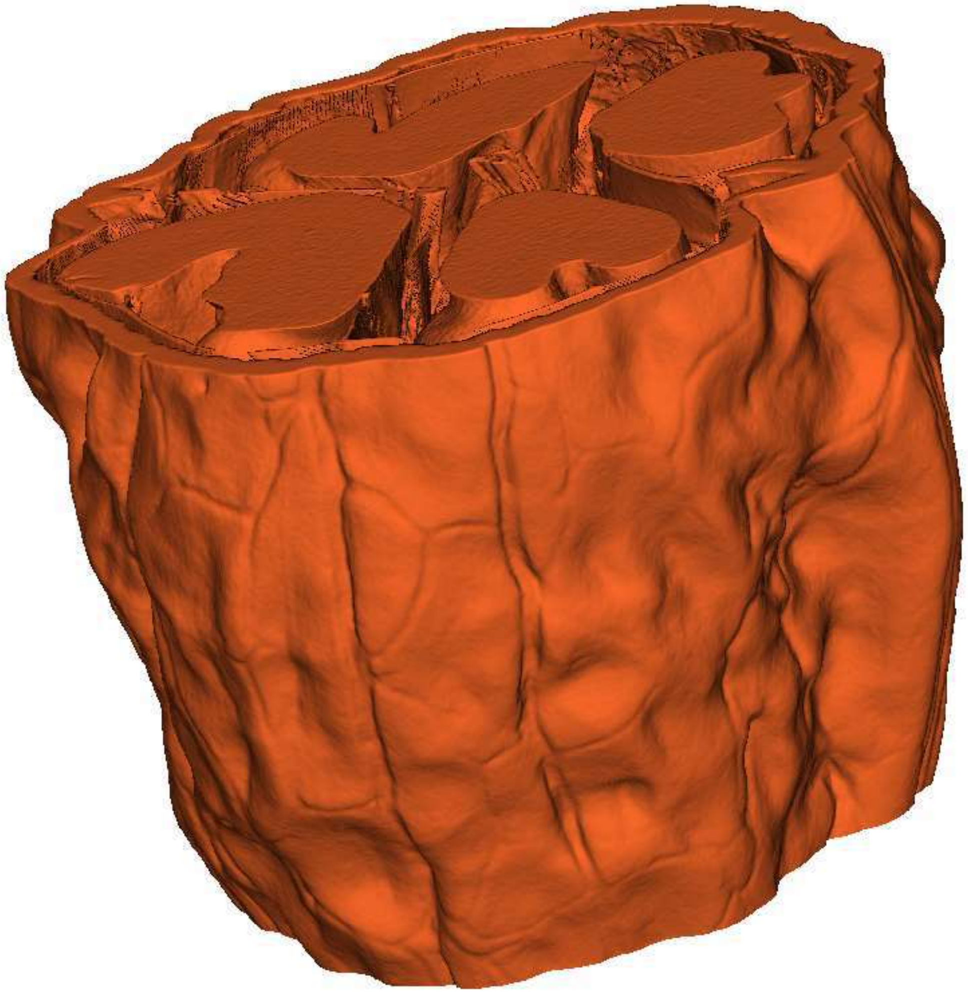

Cone-beam CT is a true 3D imaging modality (FDK reconstruction, 3D slices)

Test case: simulated human patient X-ray data

CBCT measurements of human head obtained with XCAT software

720 X-ray projections at 0.5° intervals.

100 kV X-ray tube (W target).

10 different dose levels on relative scale from 100% to 0.1%.

Projection size 320 × 320 pixels.

Cone-beam Computed Tomography: Analytical Reconstruction

Based on filtering + backprojection.

Most frequently used method is the algorithm proposed in 1984 by Feldkamp, Davis, and Kress (FDK) algorithm.

Pros:

+ Fast
+ Well understood
+ Approximately linear

Cons:

- Performs poorly with noisy data
- Performs poorly with undersampled data
- Suffers from cone-beam geometry artifacts

Dose and noise considerations in FDK

In linear reconstruction algorithms, we have

$$\text{dose} \propto N$$

and

$$SNR \propto \sqrt{N},$$

where N is the number of photons used and SNR is signal-to-noise ratio.

We therefore have

$$SNR \propto \sqrt{\text{dose}}.$$

FDK vs. reference reconstruction, rel. dose 100%

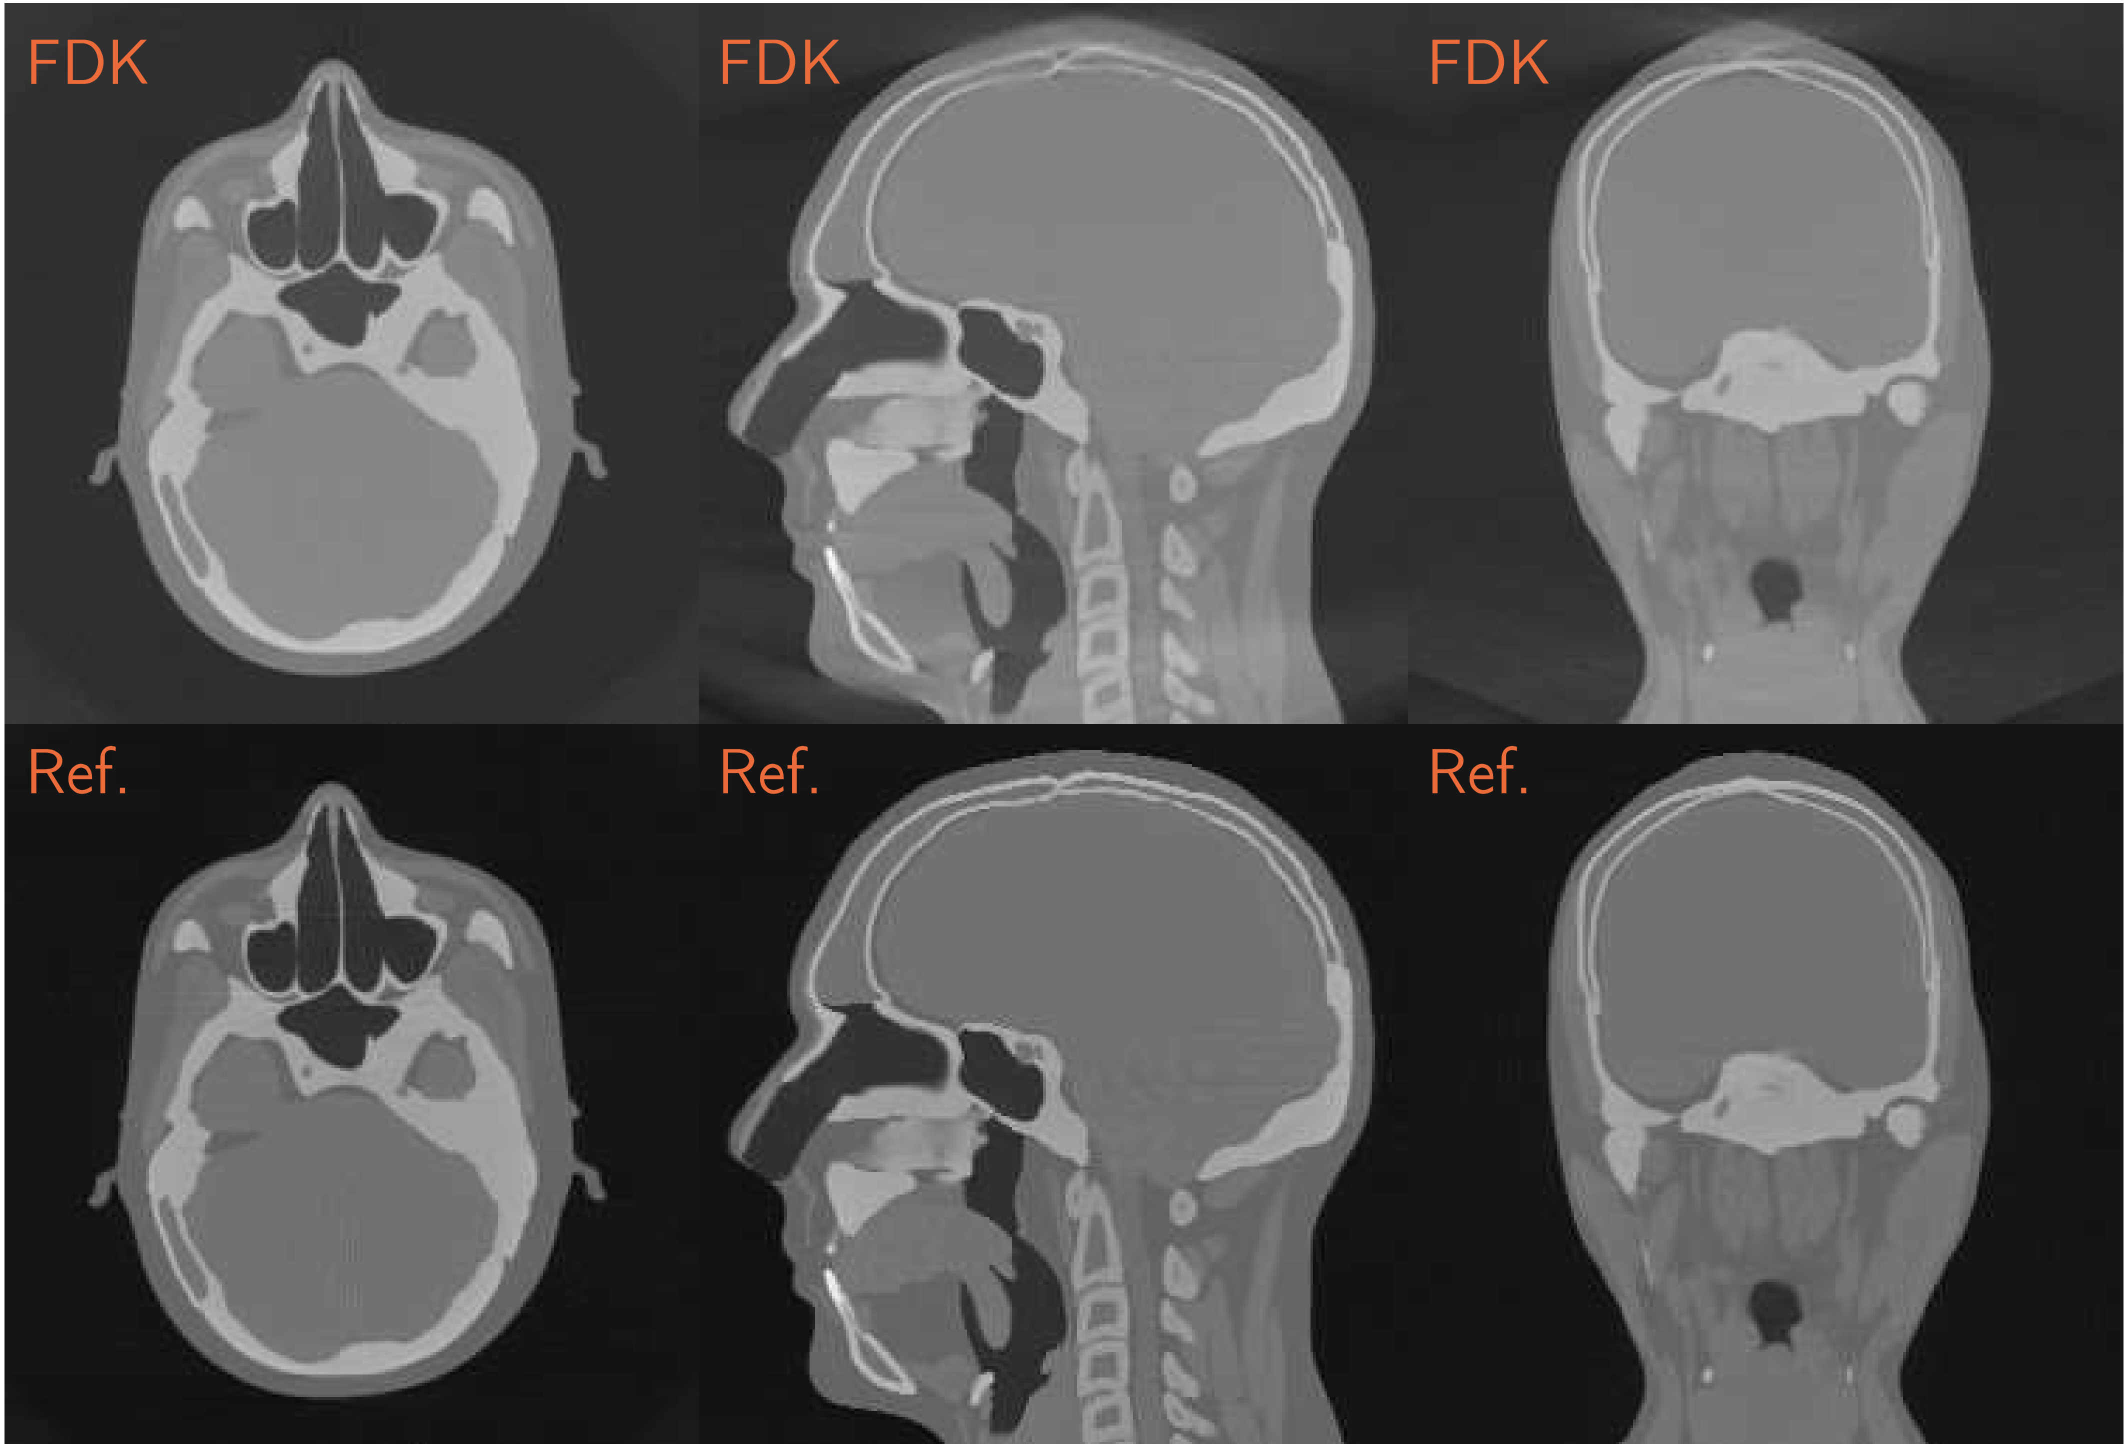

FDK vs. reference reconstruction, rel. dose 50%

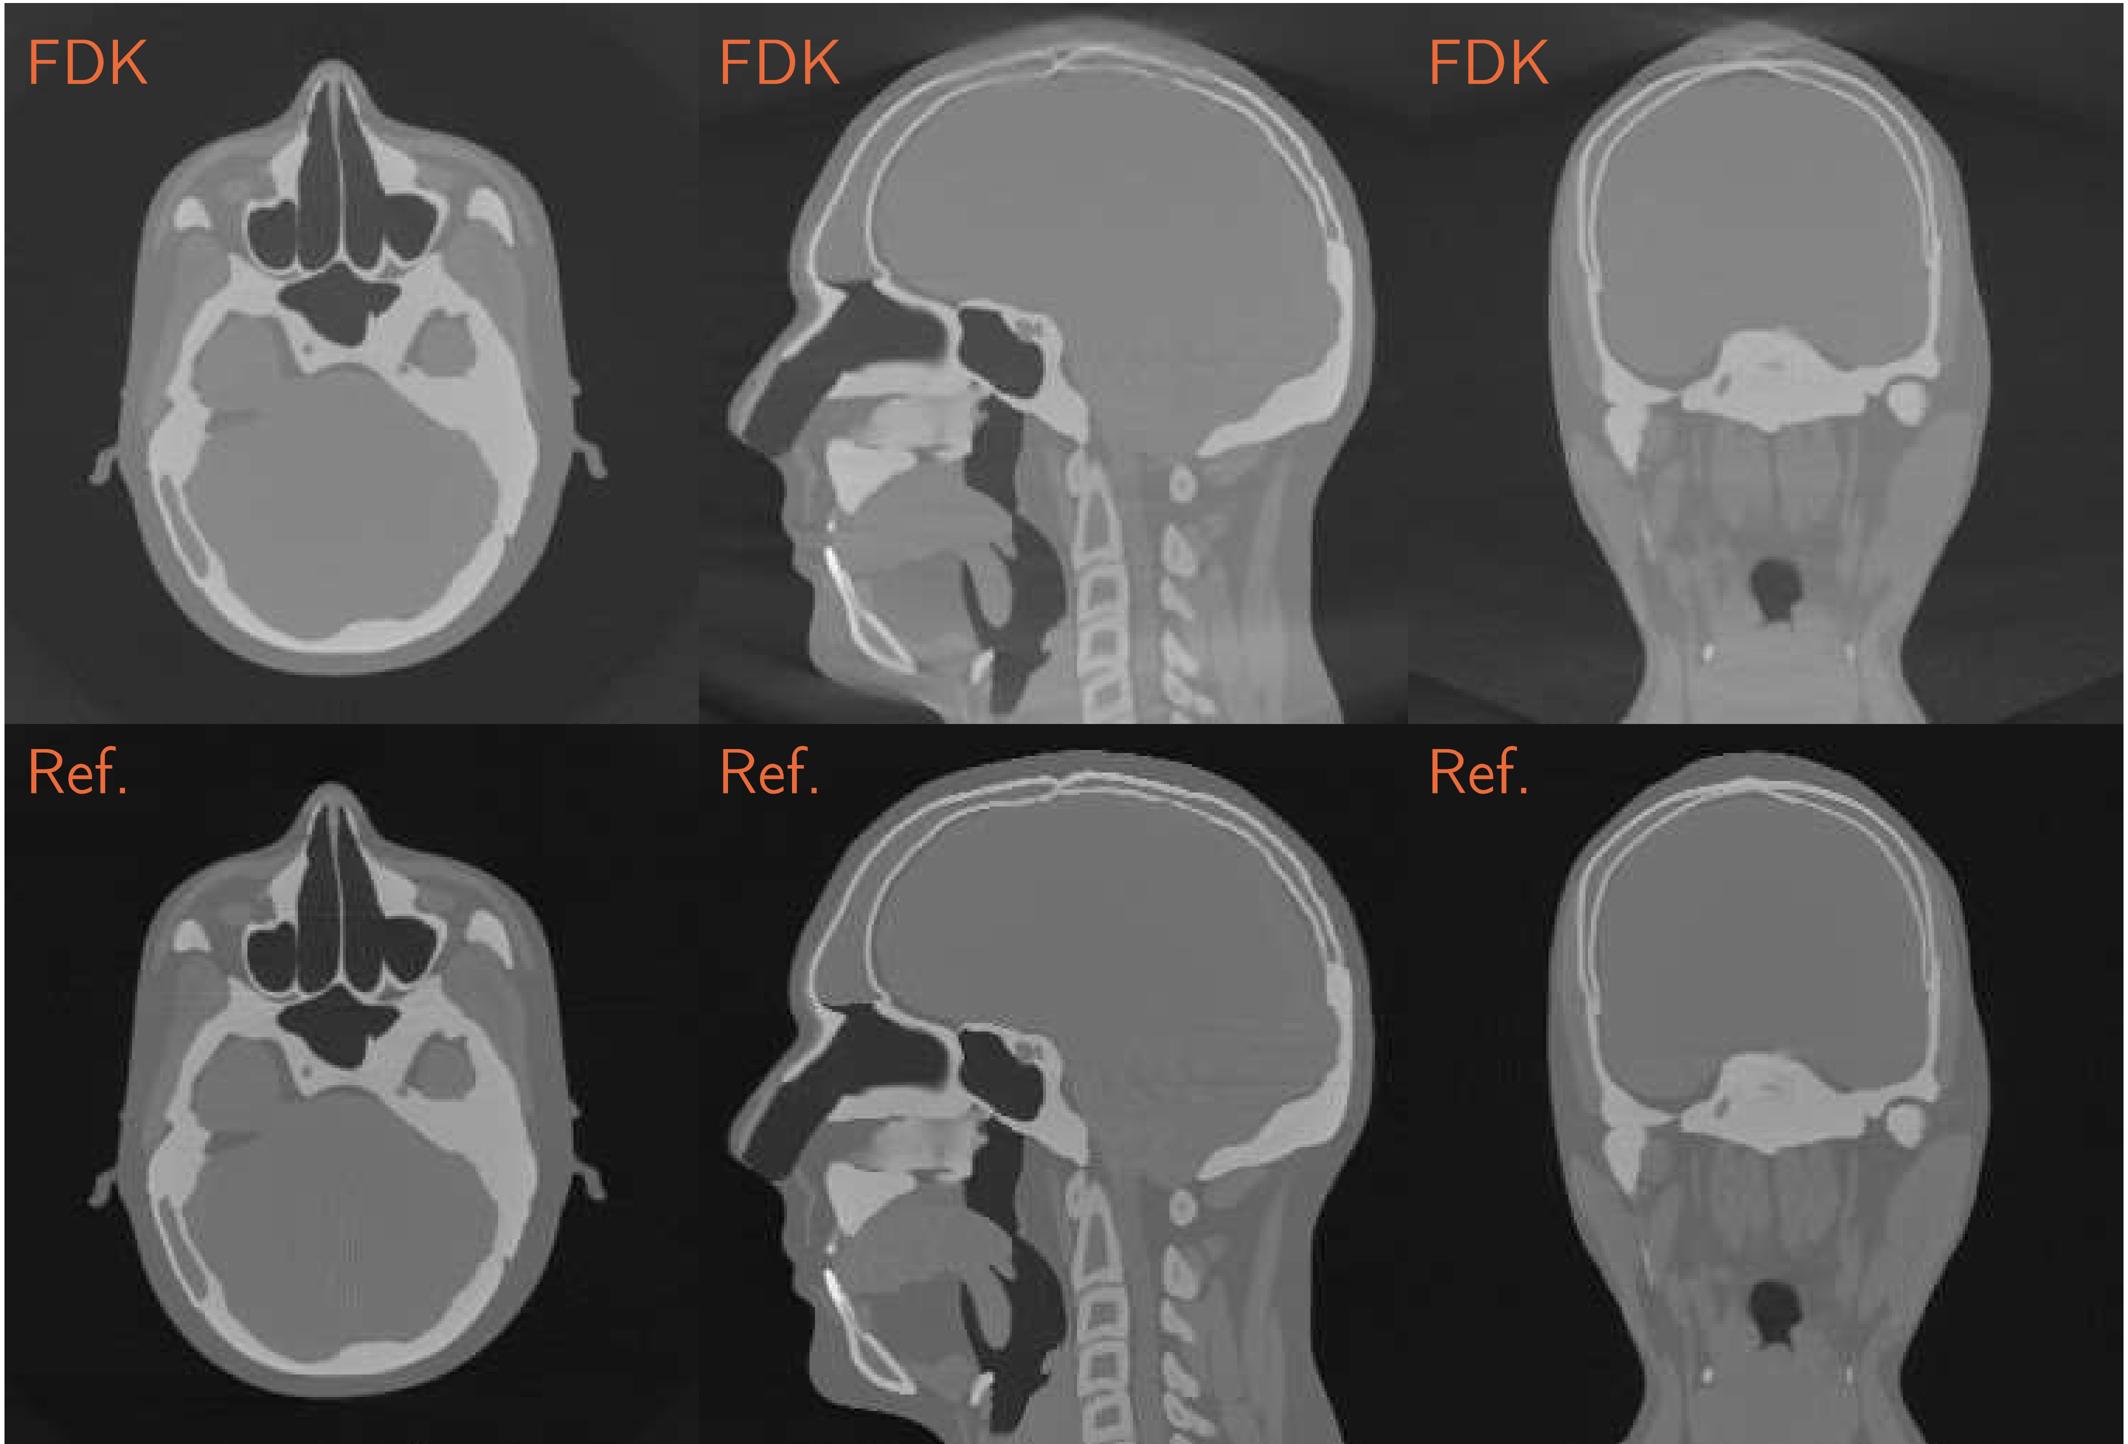

FDK vs. reference reconstruction, rel. dose 20%

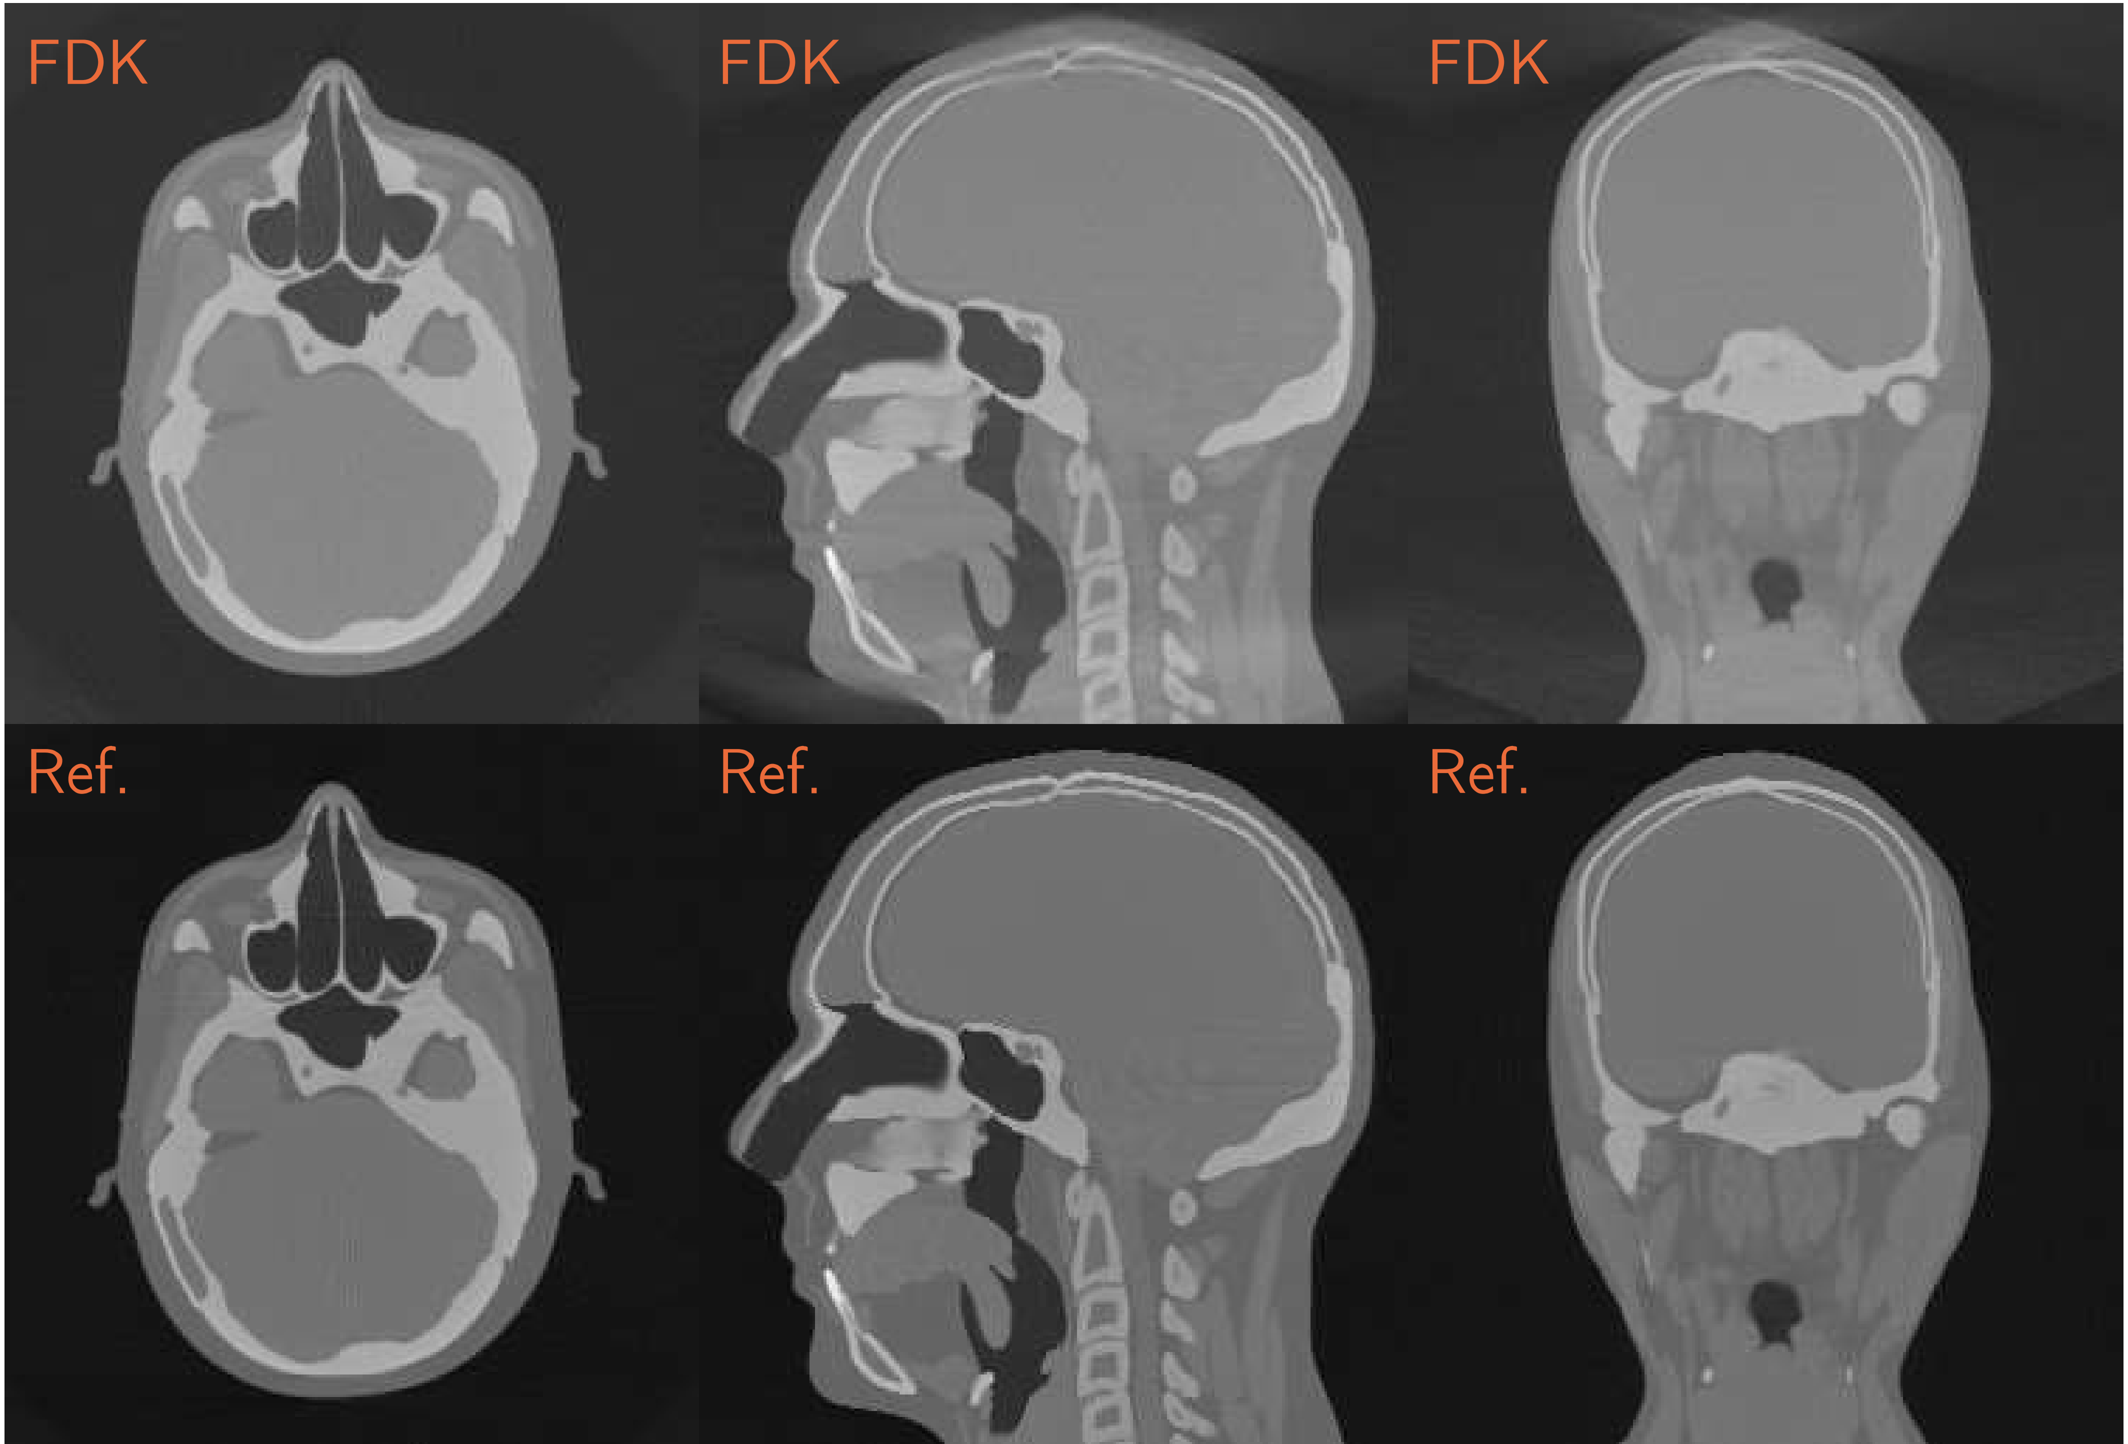

FDK vs. reference reconstruction, rel. dose 10%

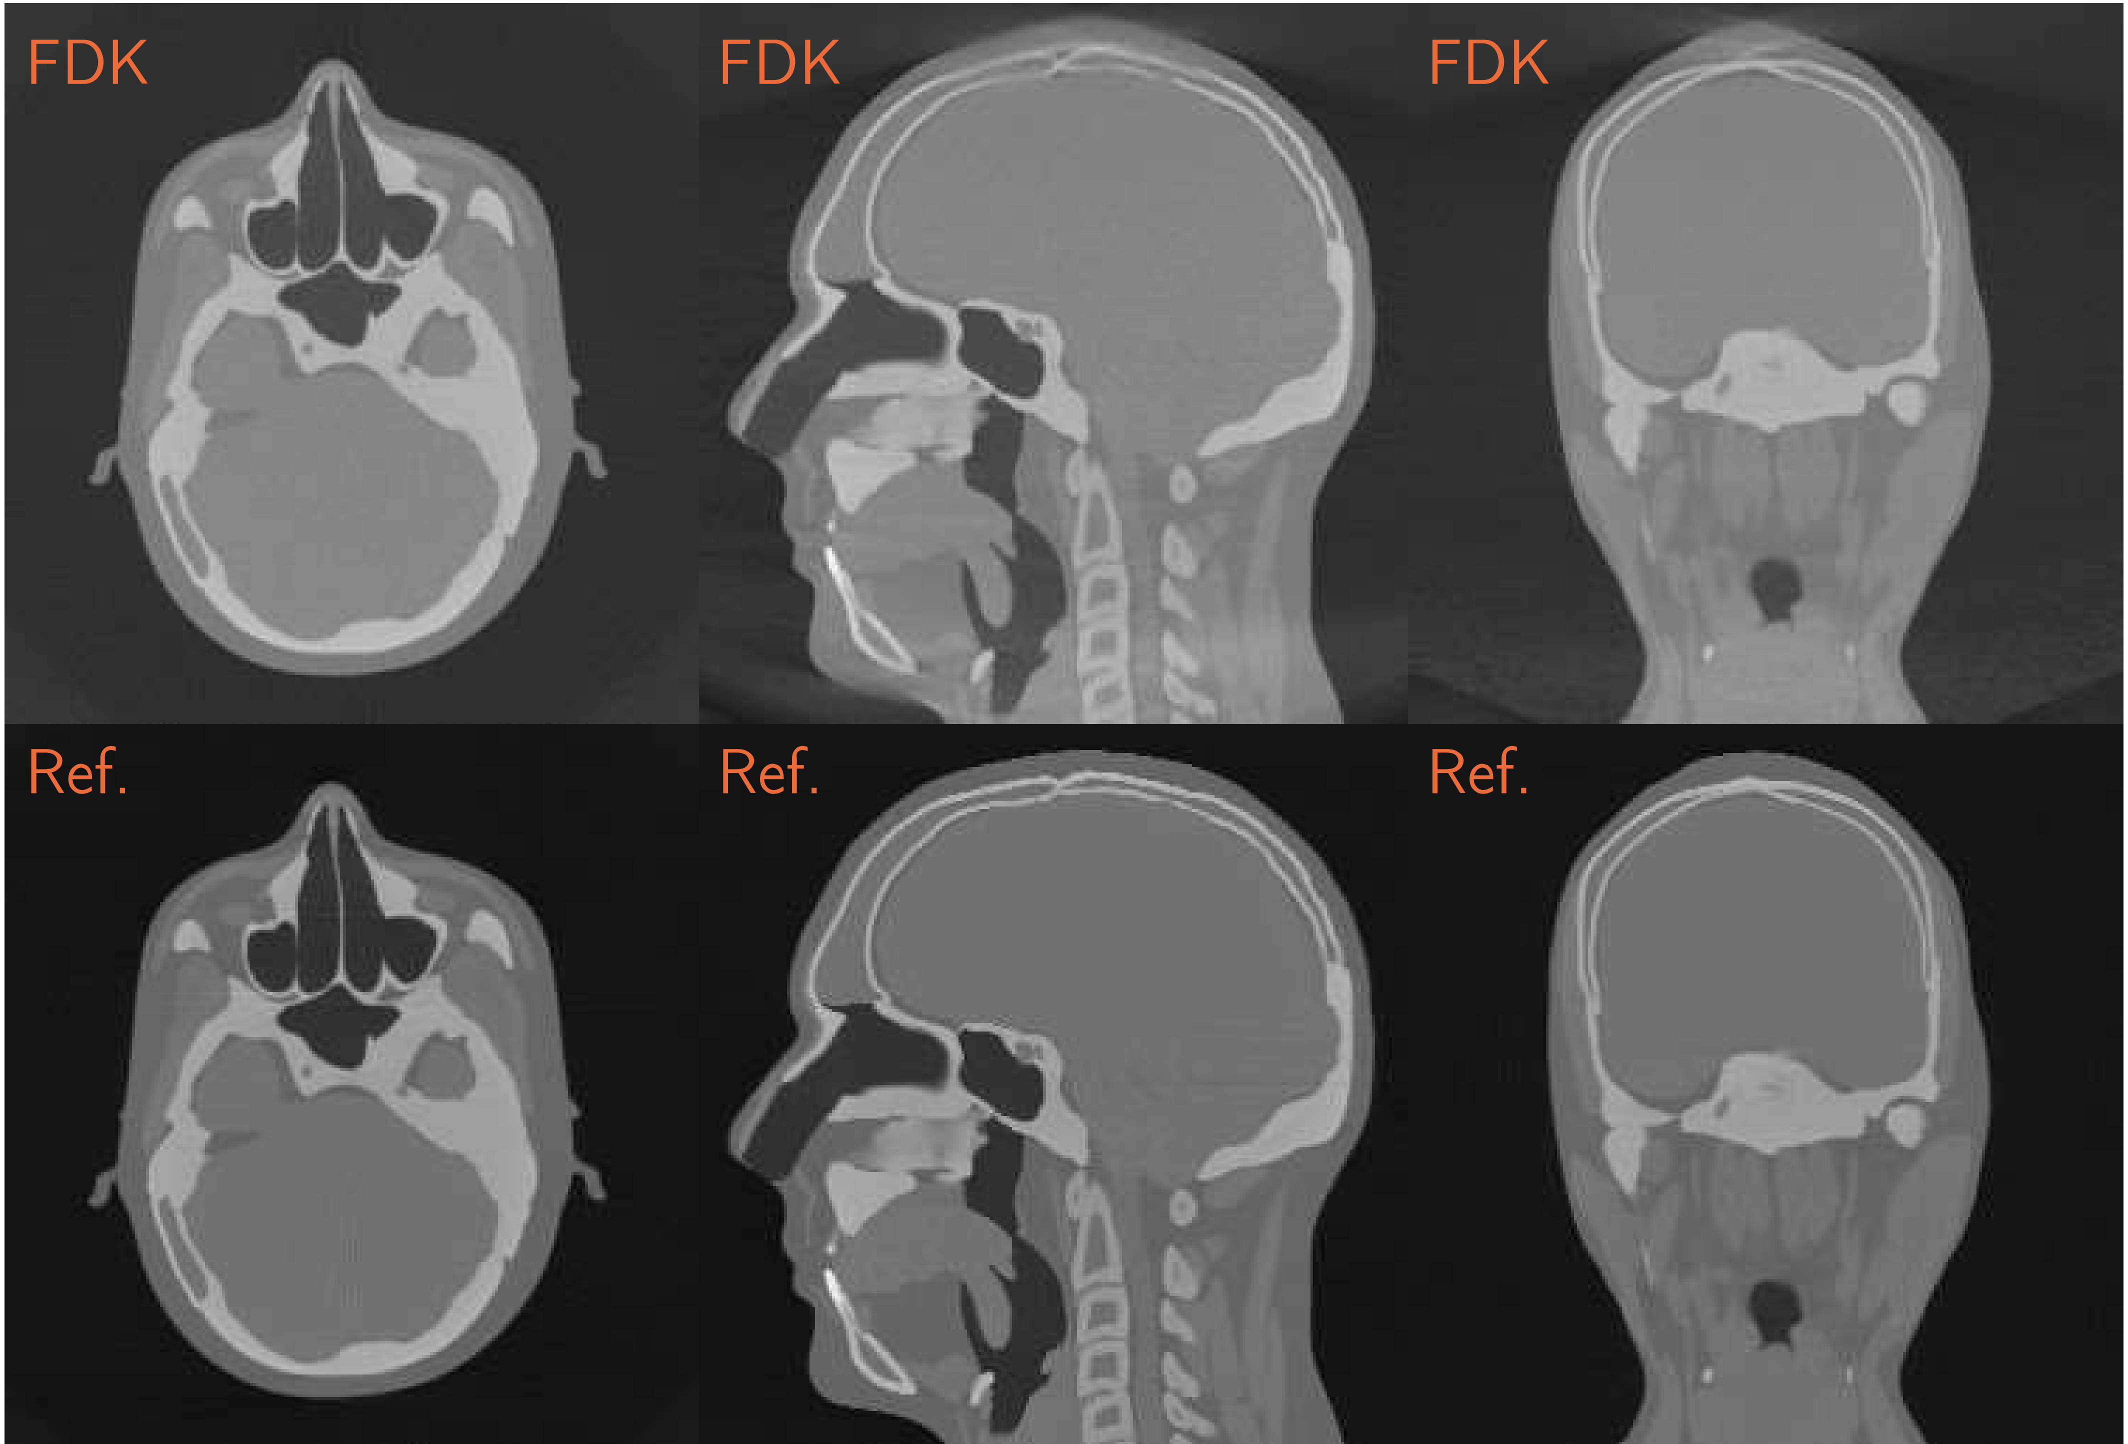

FDK vs. reference reconstruction, rel. dose 5%

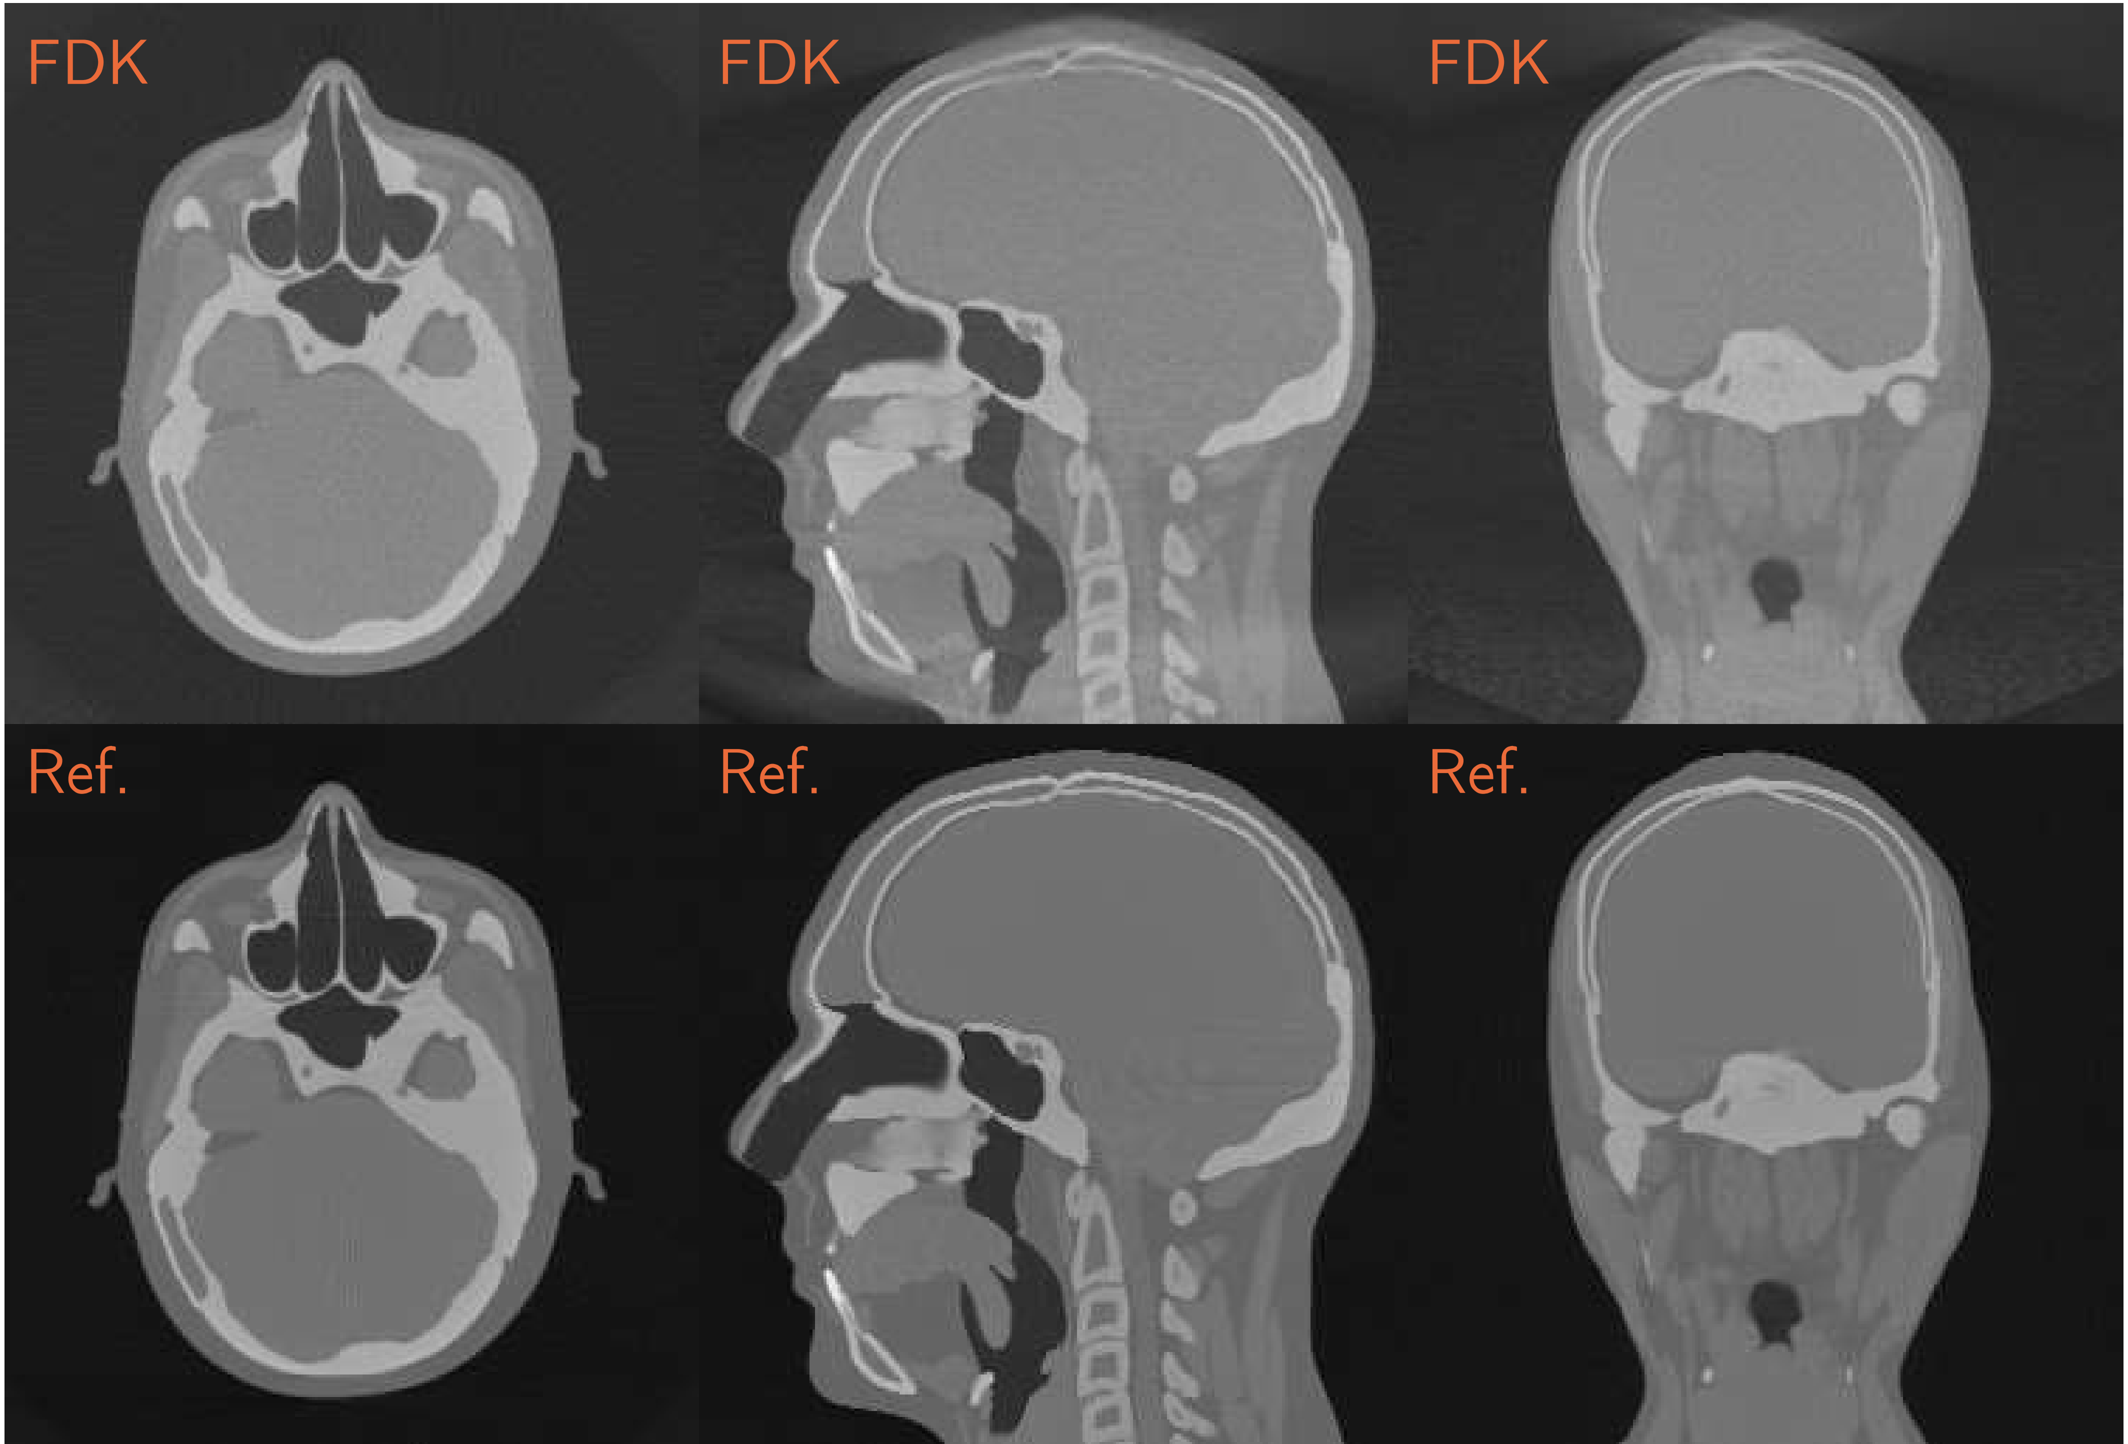

FDK vs. reference reconstruction, rel. dose 1%

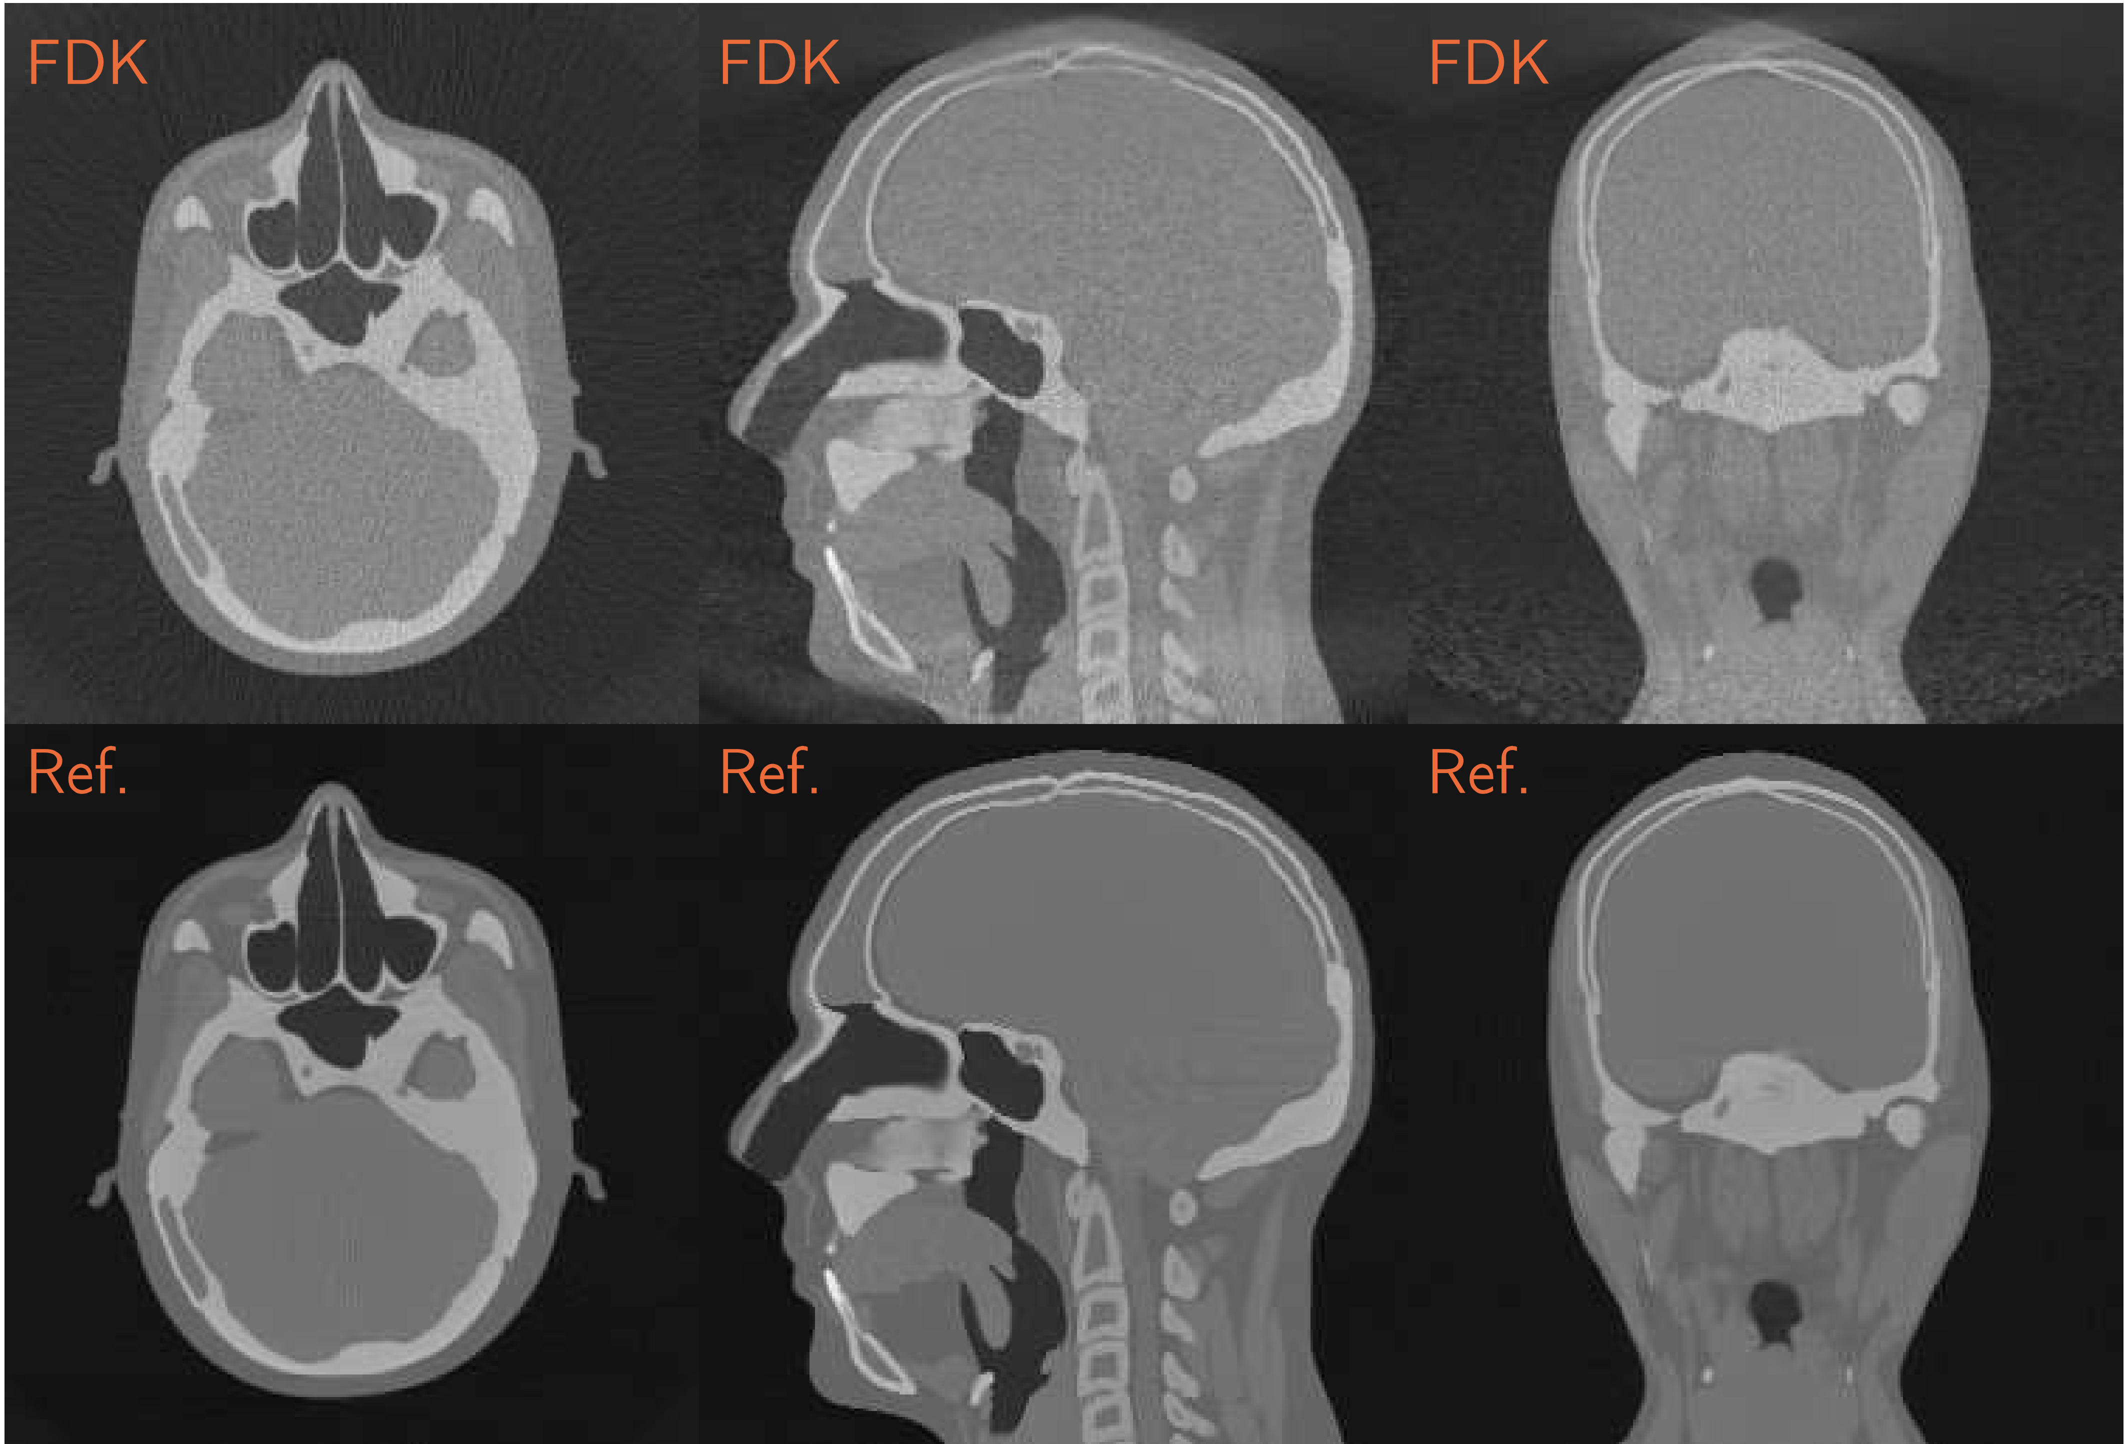

FDK vs. reference reconstruction, rel. dose 0.5%

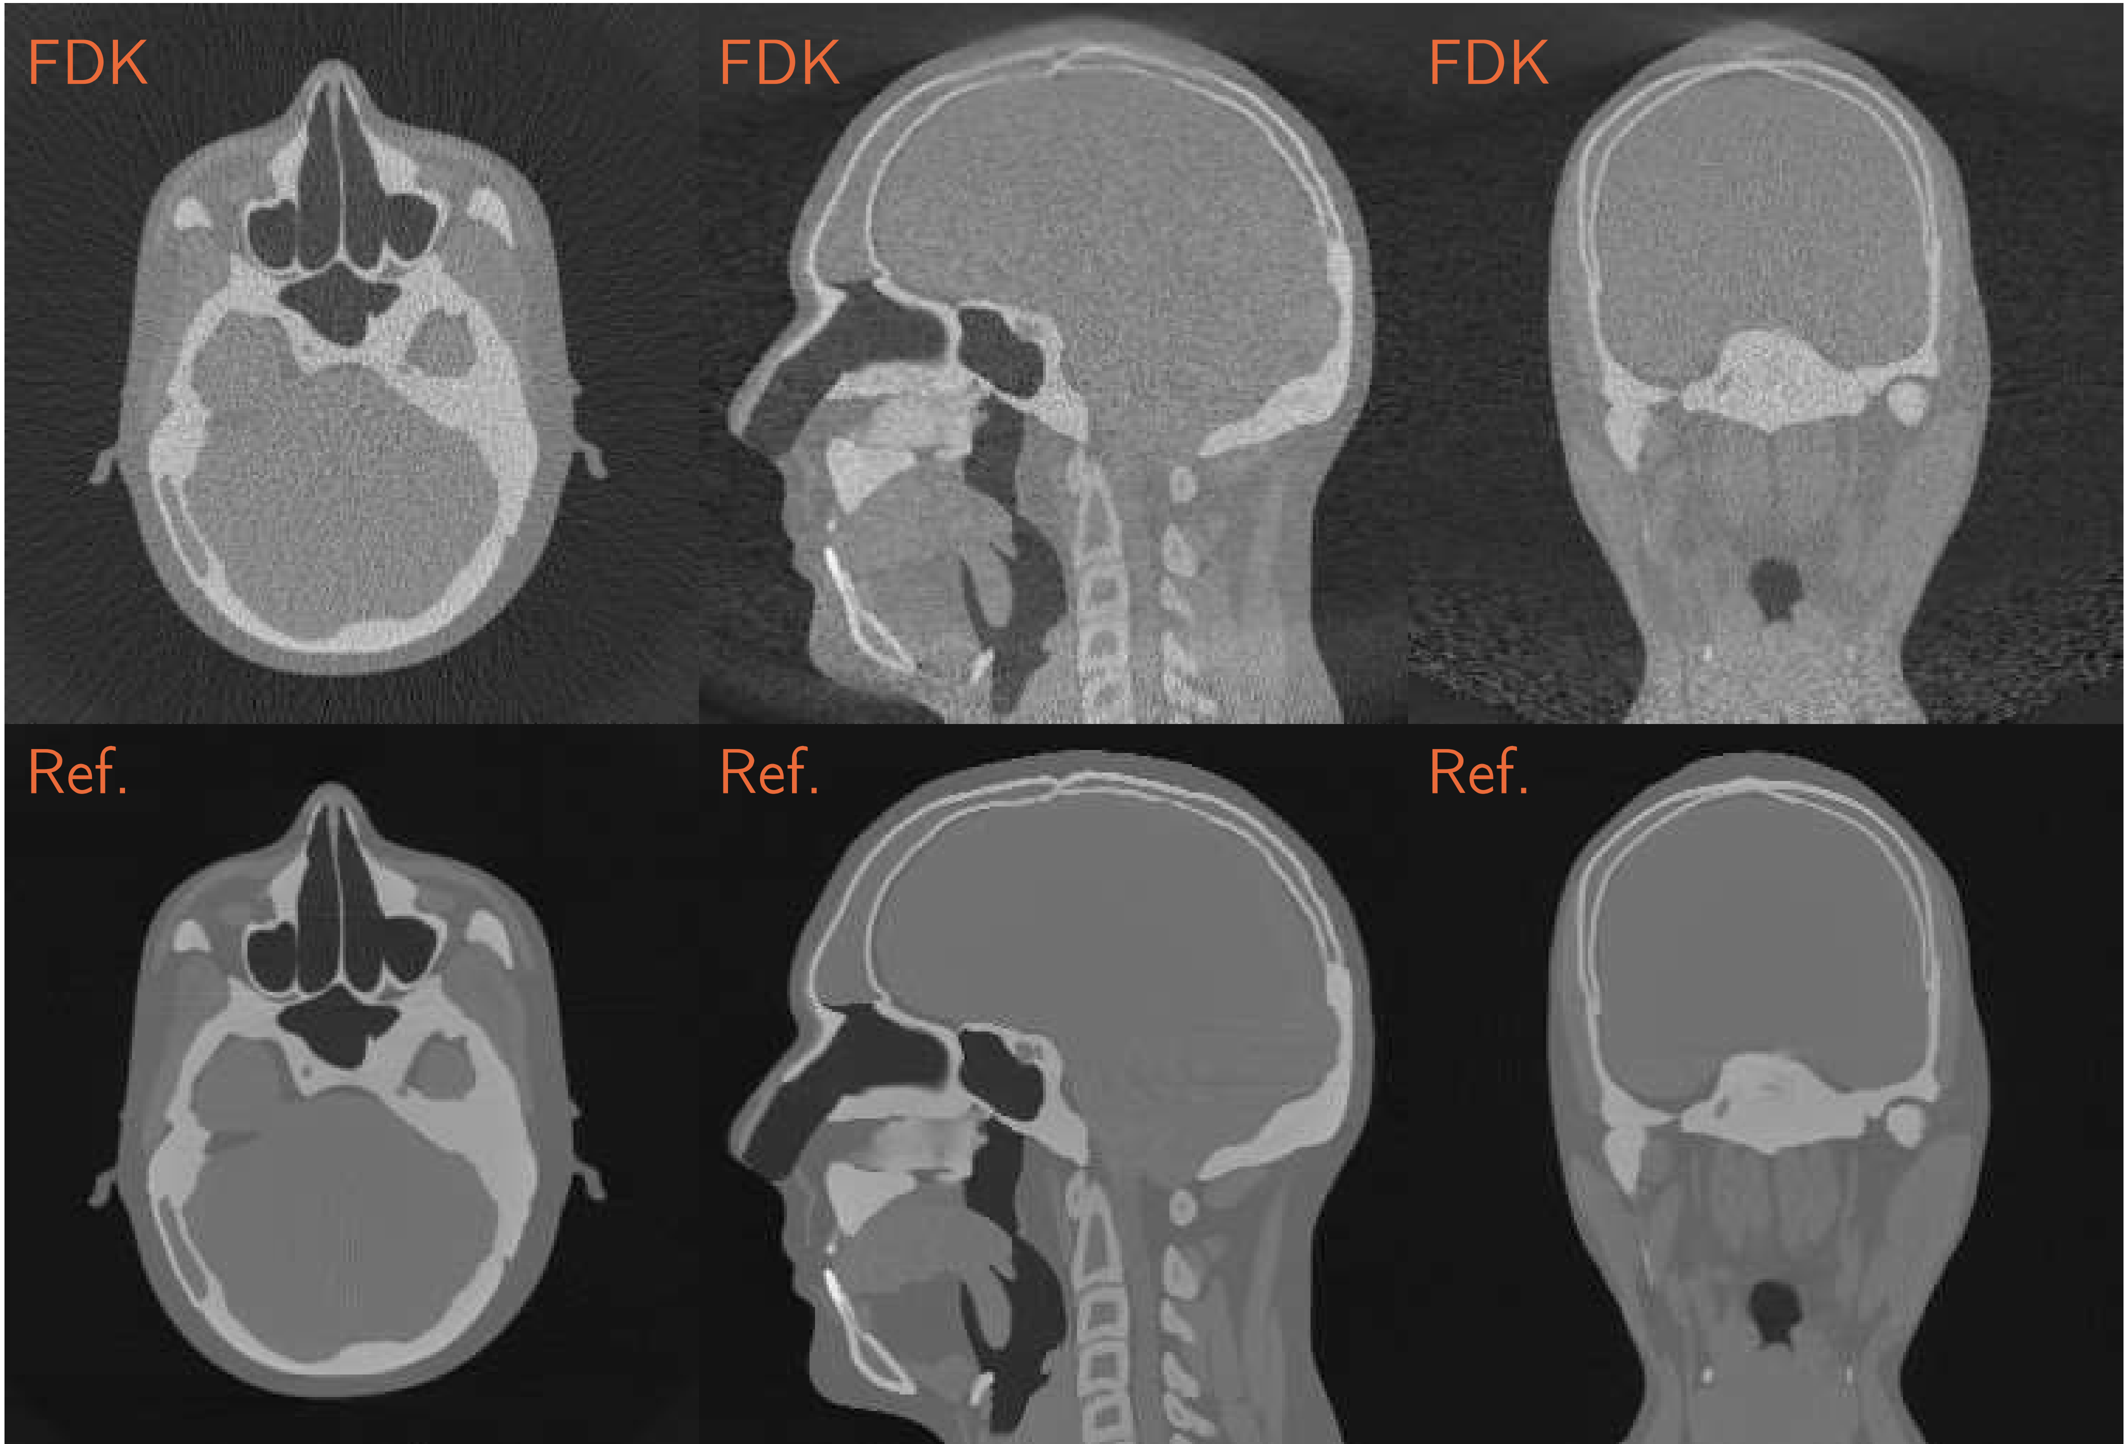

FDK vs. reference reconstruction, rel. dose 0.2%

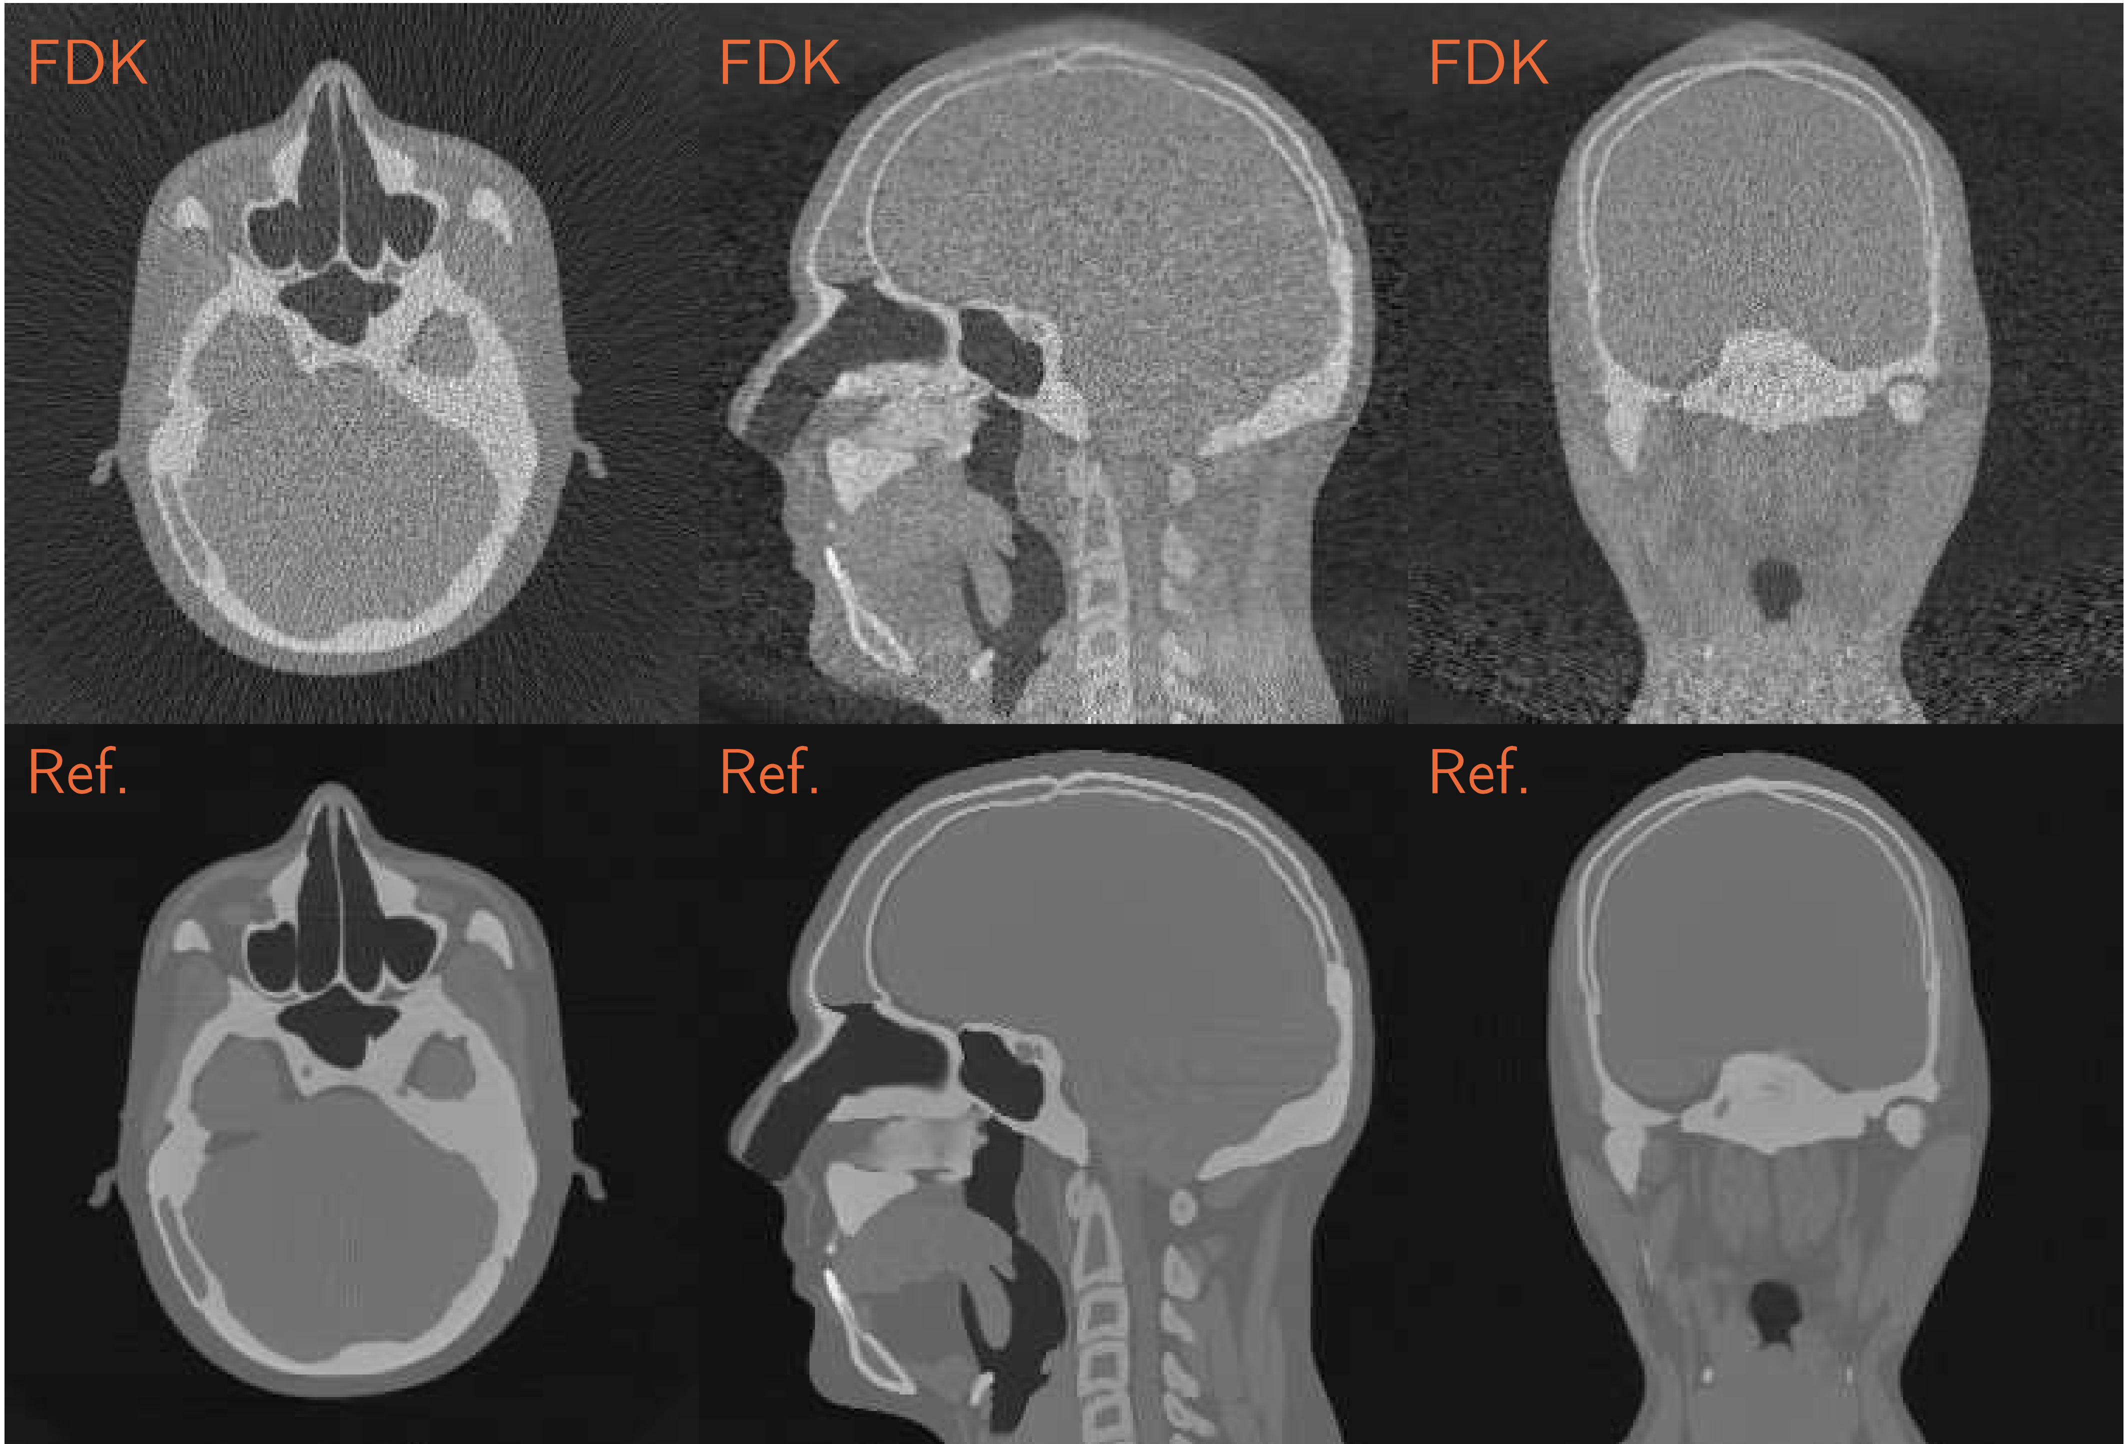

FDK vs. reference reconstruction, rel. dose 0.1%

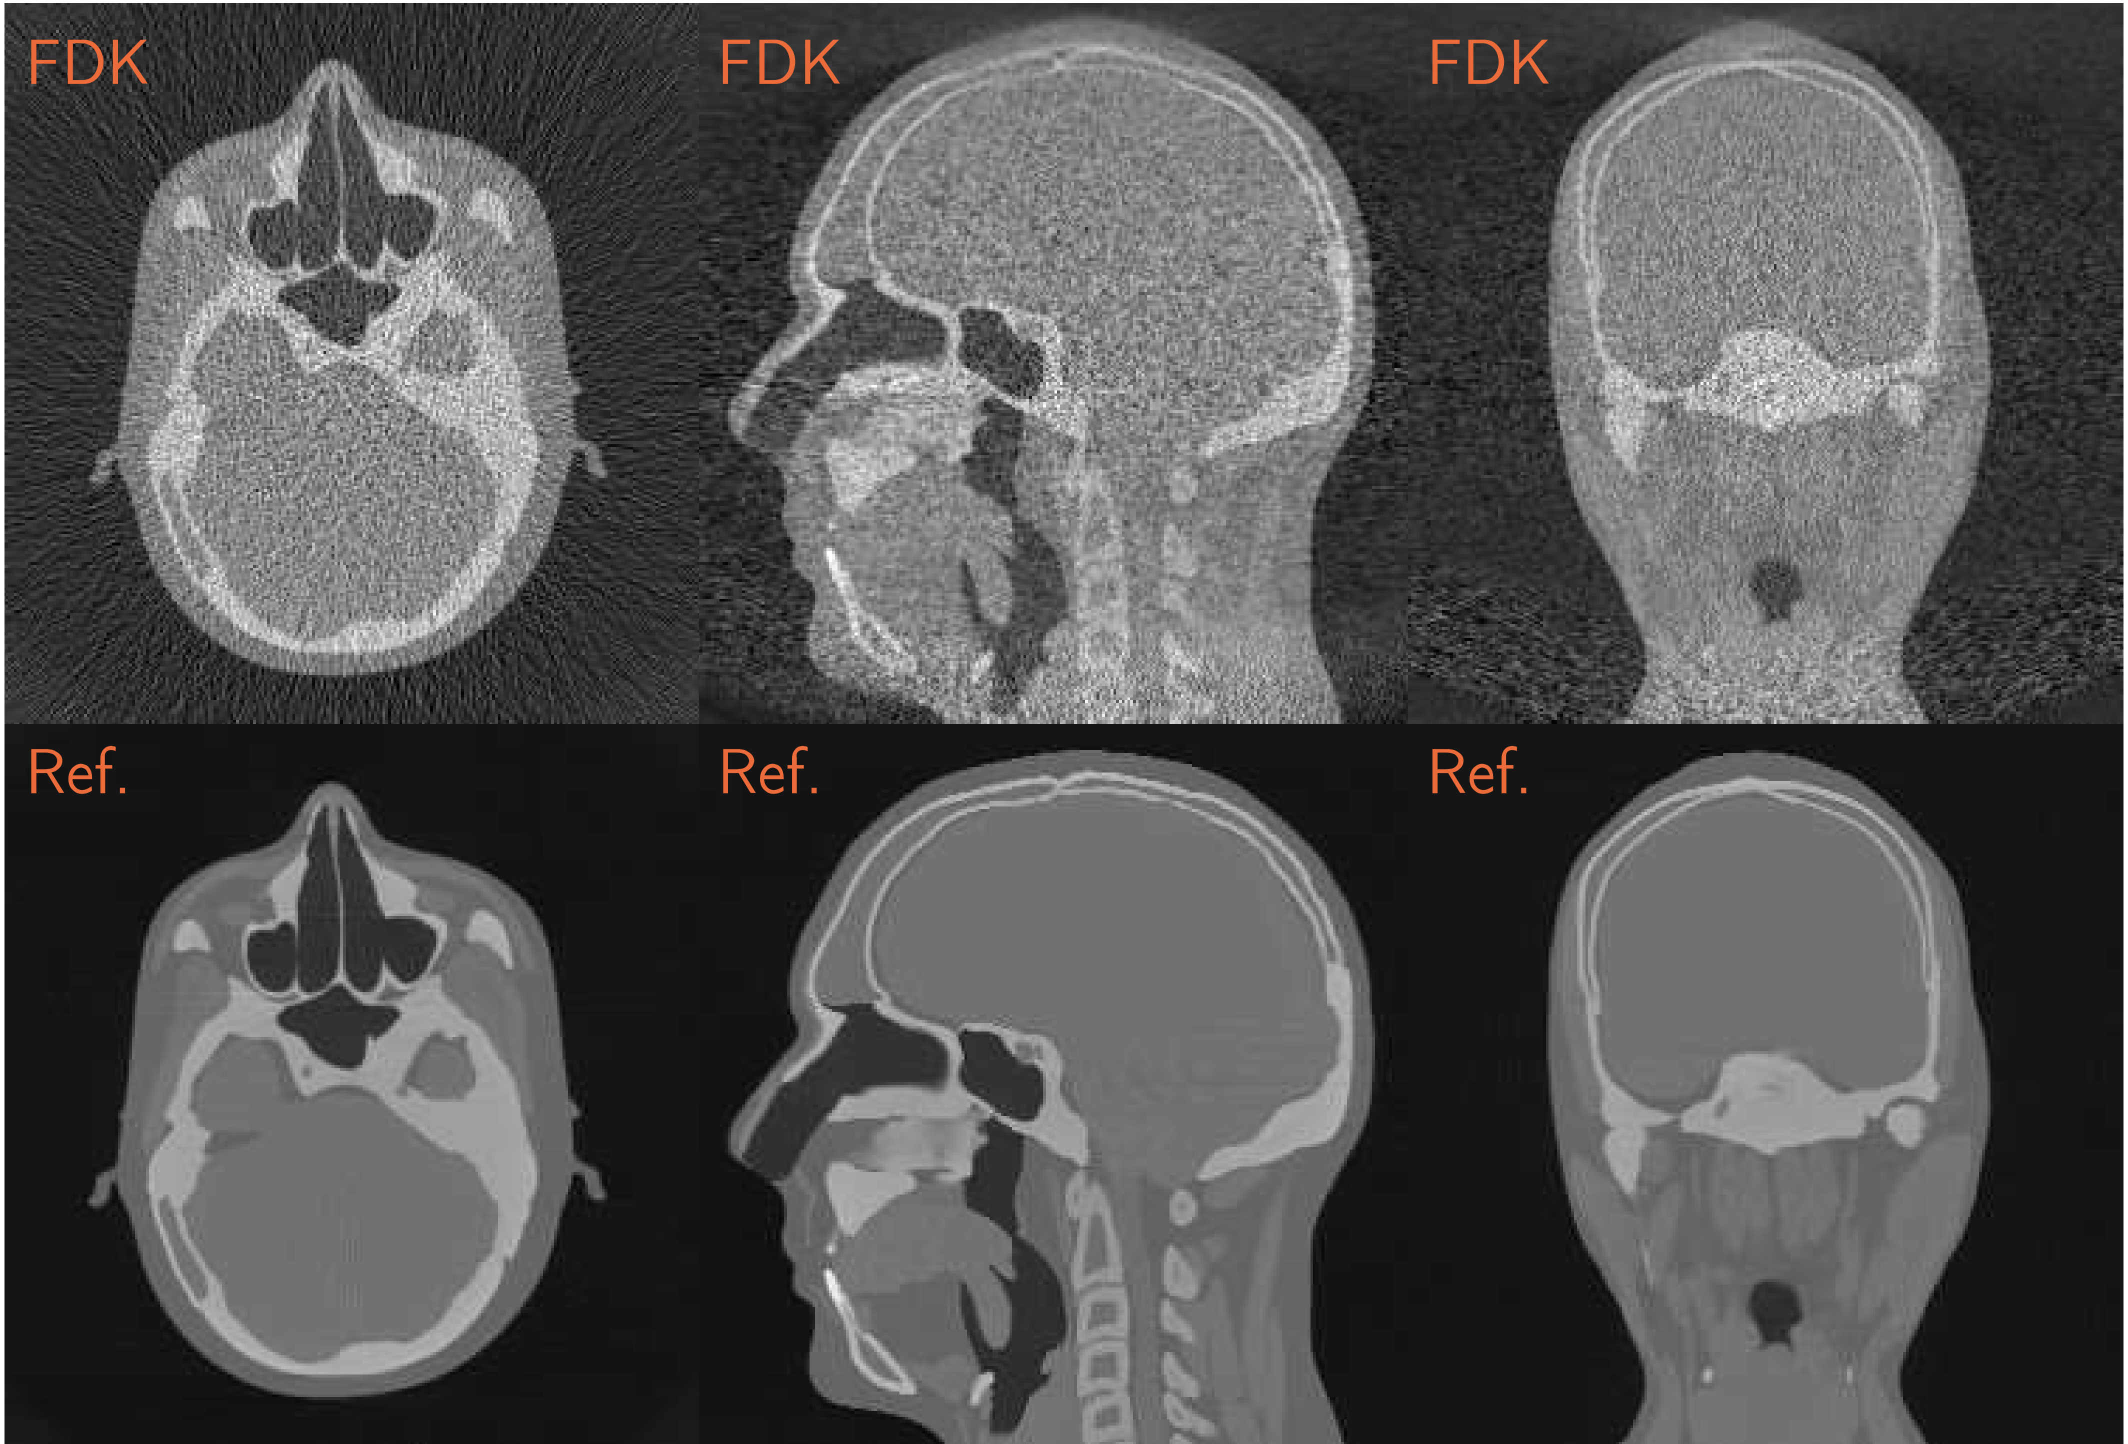

Cone-beam Computed Tomography: Iterative Reconstruction

Usually formulated as a regularized optimization problem:

$$\min_{f \in \mathbb{R}^n} \frac{1}{2}||Af - m||^2 + \mu R(f).$$

Pros

+ (Potentially) better performance with noisy and/or undersampled data
+ Allows incorporating physics modelling into the reconstruction problem

Cons:

- SLOW
- Highly sensitive to choice of the regularization parameter μ
- Choice of regularizer $R(f)$ strongly affects reconstruction results

Approach 1: Haar-CT

We formulate the reconstruction problem as

$$\min_{f \in \mathbb{R}^n_+} \frac{1}{2}||Af - m||^2 + \mu||Bf||_1,$$

where B is the Haar transform of f.

Hypothesis: Enforcing sparsity in a Haar wavelet basis will result in improvements in reconstruction quality in low-dose CBCT.

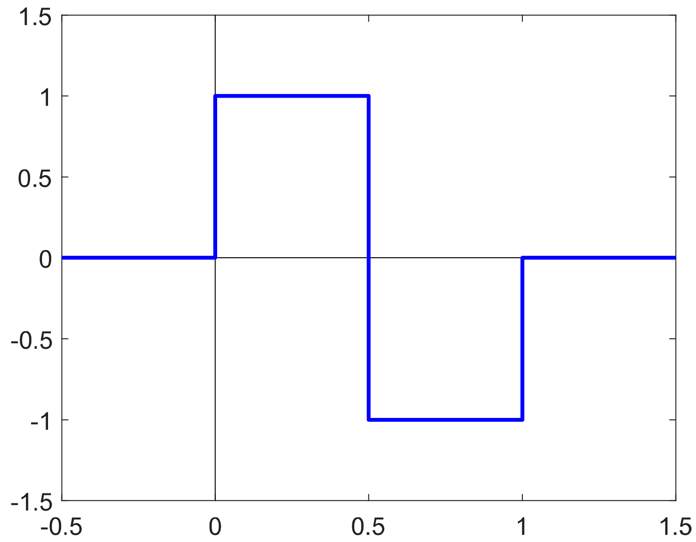

Approach 2: Anisotropic Total Variation

We formulate the reconstruction problem as

$$\min_{f \in \mathbb{R}^n_+} \frac{1}{2}||Af - m||^2 + \mu\big(||\partial_x f||_1 + ||\partial_y f||_1 + ||\partial_z f||_1\big),$$

where $\partial_x f$, $\partial_y f$, and $\partial_z f$ are the discrete derivatives of f.

Hypothesis: Enforcing sparsity in the components of the gradient will result in improvements in reconstruction quality in low-dose CBCT.

The primal-dual fixed point (PDFP) algorithm [Chen, Huang & Zhang, 2016]

Consider the following minimization problem:

$$\min_{x\in\mathbb{R}^n} f_1(x) + (f_2 \circ B)(x) + f_3(x),$$

where f_1, f_2, and f_3 are proper lower semi-continuous convex functions, f_1 is differentiable on $\mathbb{R}^n$ with a $1/\beta$-Lipschitz continuous gradient, and $B : \mathbb{R}^n \to \mathbb{R}^m$ is a linear transformation.

This can be solved using the algorithm

$$(\mathrm{PDFP}) \begin{cases} y^{k+1} &= \mathrm{prox}_{\gamma f_3}(x^k - \gamma\nabla f_1(x^k) - \lambda B^T v^k), \\ v^{k+1} &= (I - \mathrm{prox}_{\frac{\gamma}{\lambda} f_2})(By^{k+1} + v^k), \\ x^{k+1} &= \mathrm{prox}_{\gamma f_3}(x^k - \gamma\nabla f_1(x^k) - \lambda B^T v^{k+1}), \end{cases}$$

where $0 < \lambda < 1/\lambda_{\max}(BB^T)$, $0 < \gamma < 2\beta$.

The primal-dual fixed point (PDFP) algorithm [Chen, Huang & Zhang, 2016]

Our formulation of the problem can be stated as:

$$\min_{f \in \mathbb{R}^n_+} \frac{1}{2}||Af - m||_2^2 + \mu R(f),$$

where

$\frac{1}{2}||Af - m||_2^2$ is the data fidelity term,

$R(x) = ||[\partial_x f, \partial_y f, \partial_z f]^T||_1$ OR

$R(x) = ||Wf||_1$, and

μ is the regularization parameter.

The primal-dual fixed point (PDFP) algorithm [Chen, Huang & Zhang, 2016]

The optimization algorithm adapted to our problem is

$$
\text{(PDFP)}\ \begin{cases} y^{k+1} &= \mathrm{proj}_C(x^k - \gamma A^T(Ax^k - m) - \lambda B^T v^k), \\ v^{k+1} &= (I - S_{\mu\frac{\gamma}{\lambda}})(By^{k+1} + v^k), \\ x^{k+1} &= \mathrm{proj}_C(x^k - \gamma A^T(Ax^k - m) - \lambda B^T v^{k+1}), \end{cases}
$$

where

proj_C is the projection operator to the non-negative orthant of $\mathbb{R}^n$, and

S_α is the soft thresholding operator.

Soft thresholding
(This is where the interesting stuff happens)

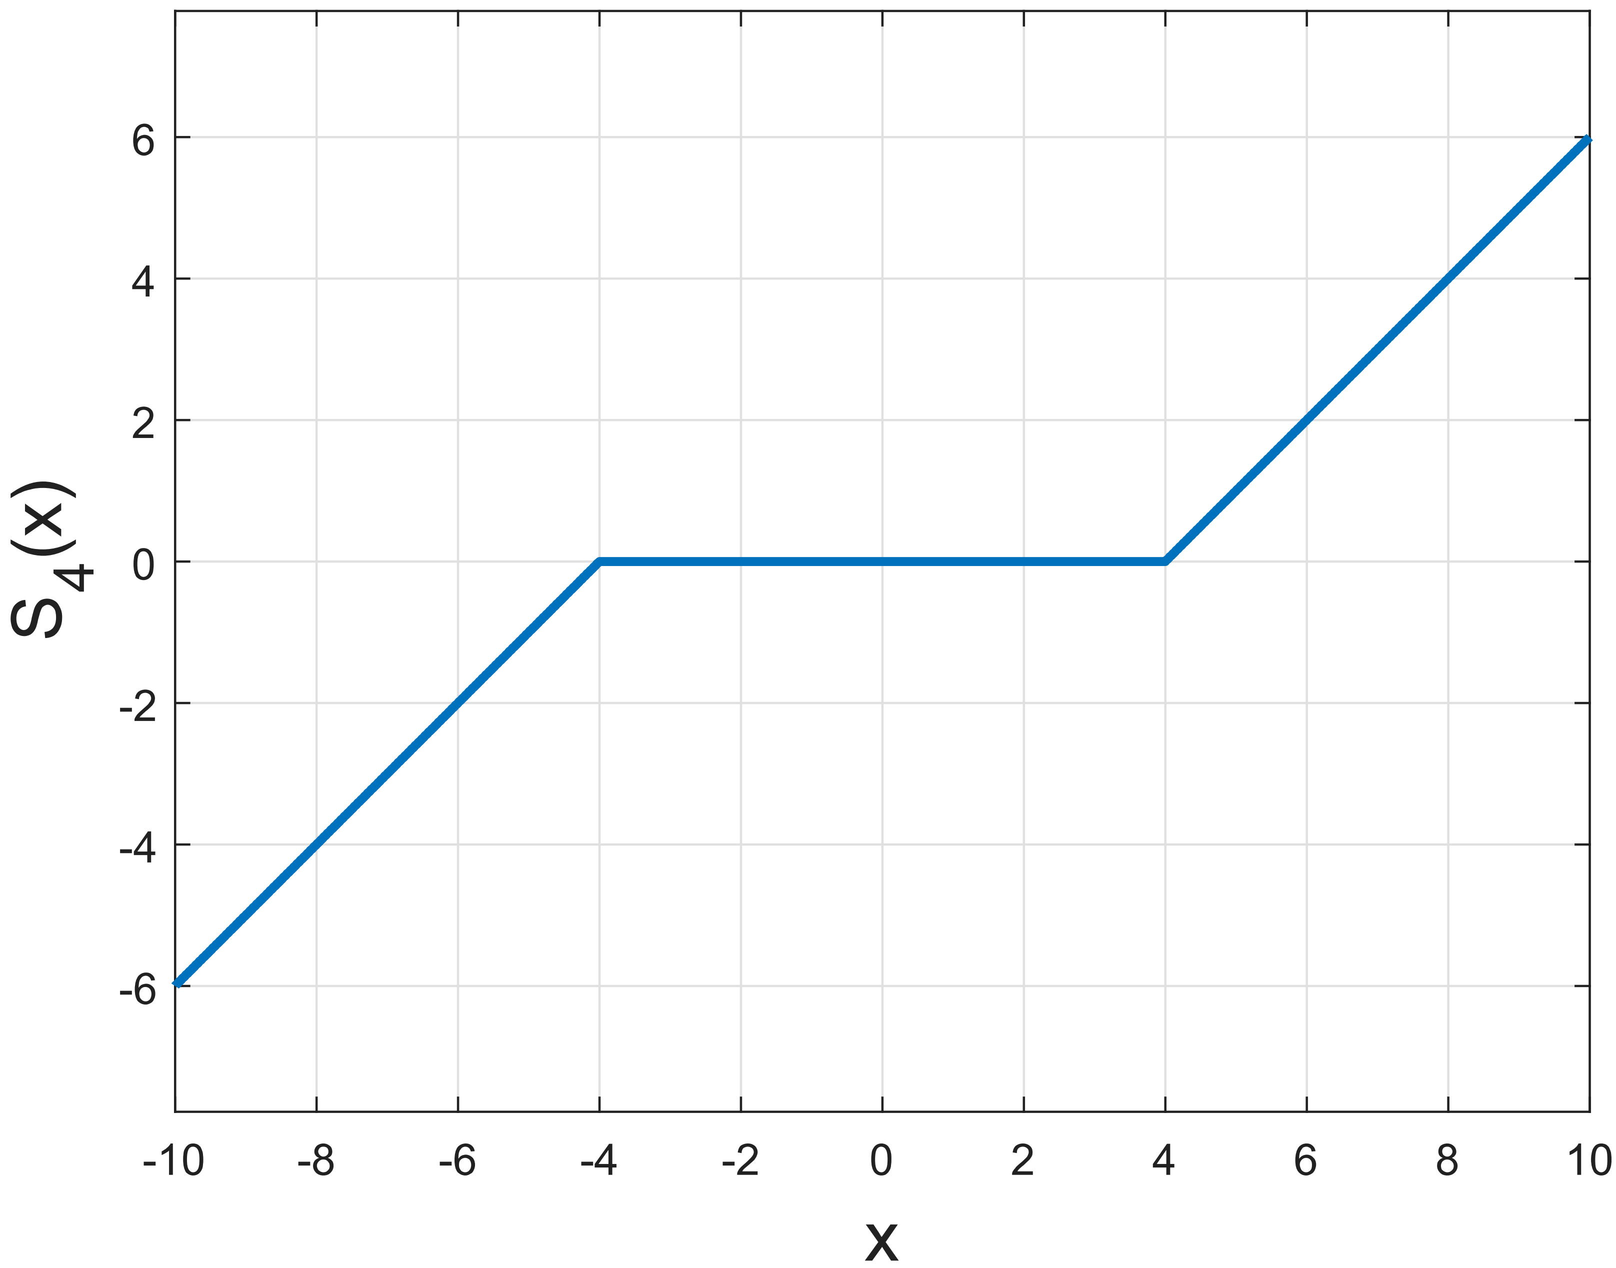

Reconstruction settings

Reconstruction size: $256 \times 256 \times 256$ pixels.

Voxel size in reconstruction: 1 mm.

Iteration stopping conditions: $||f_i - f_{i-1}|| < 10^{-3}$ or $n_{\mathrm{iter}} > 200$.

Single precision floating-point numbers.

FDK recontruction used as f_0 (convergence acceleration).

Reconstructions: Haar-CT

What does Haar wavelet sparsity look like?

How does choice of μ affect the reconstruction?

Haar wavelets, dose 100%, μ too large ($5 \cdot 10^{-5}$)

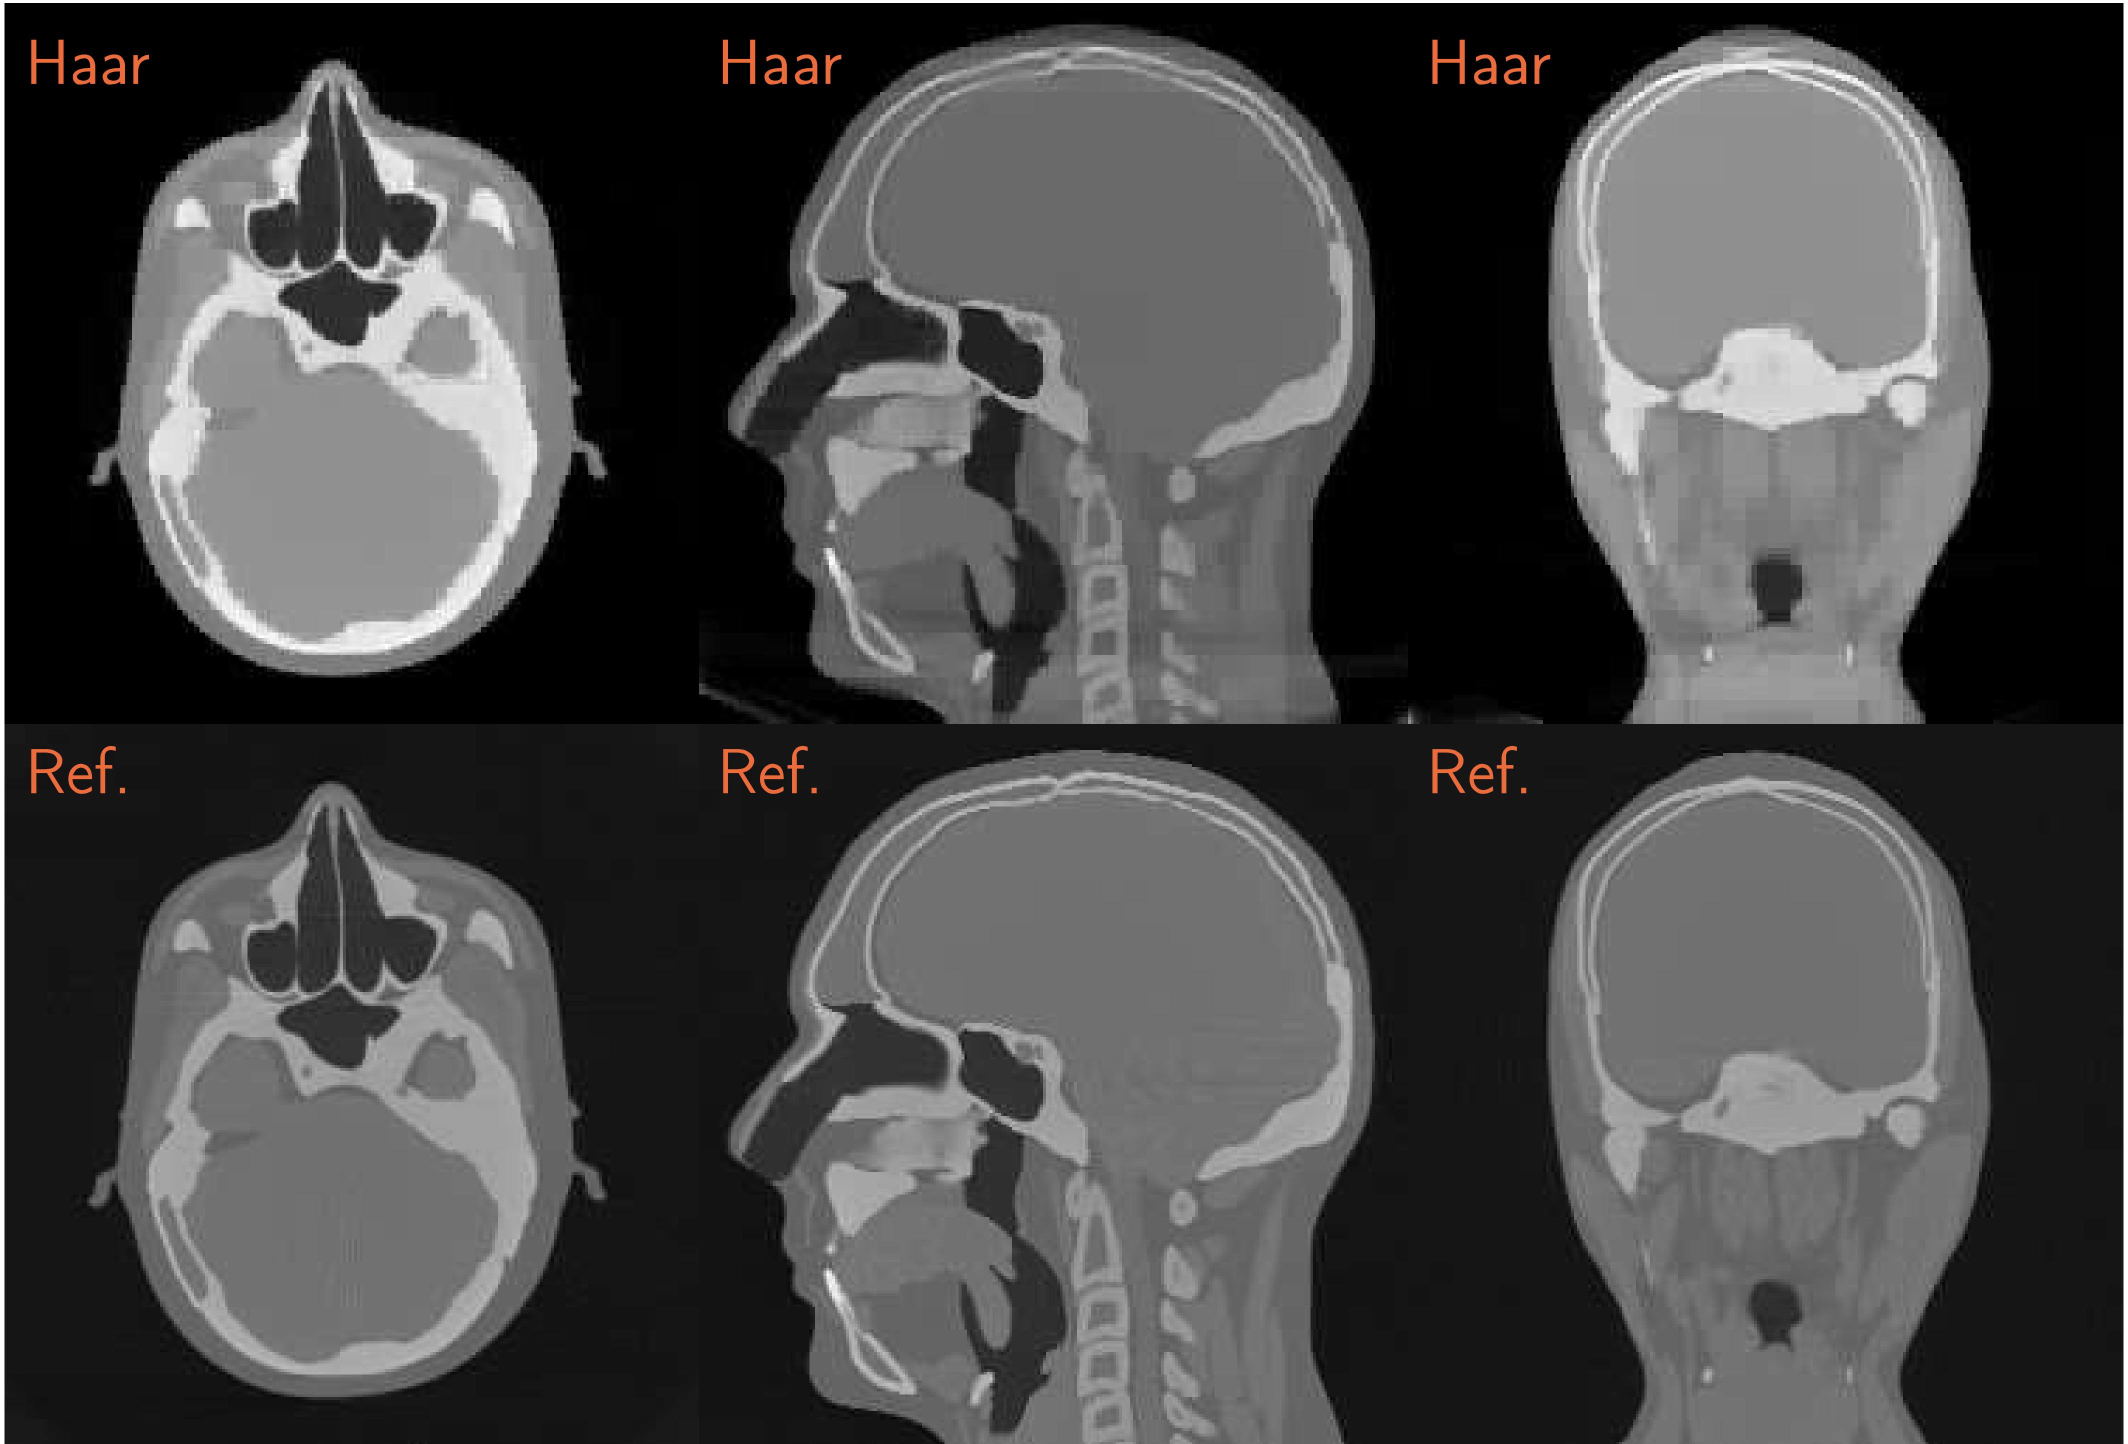

Haar wavelets, dose 100%, μ suitable ($2.5 \cdot 10^{-6}$)

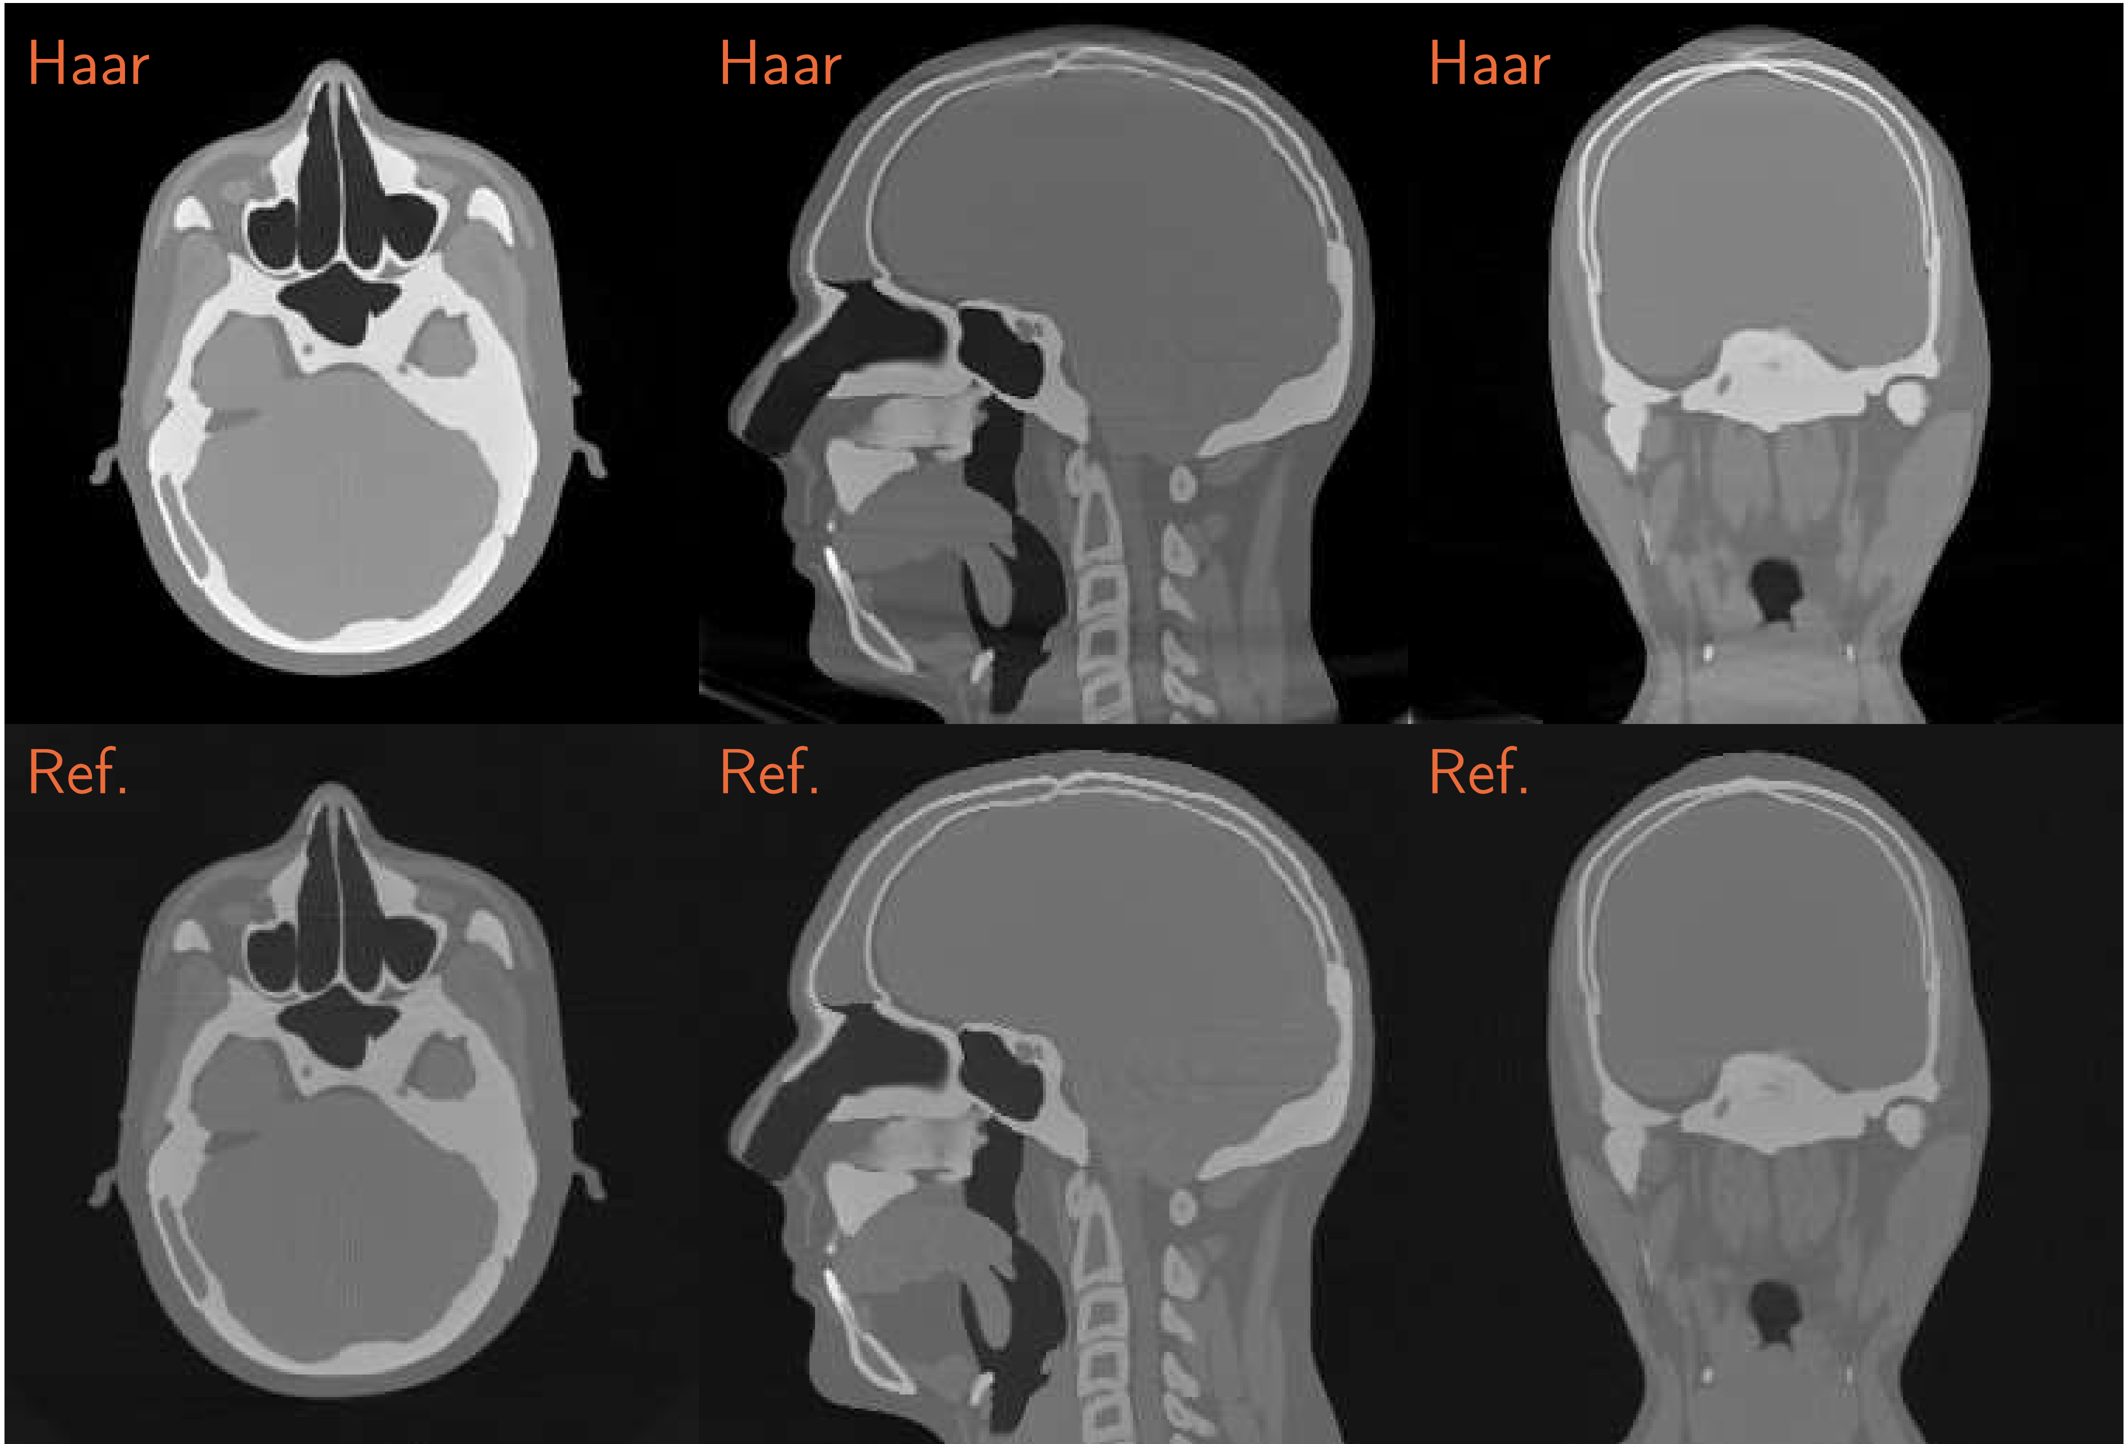

Haar wavelets, dose 100%, μ too small ($1 \cdot 10^{-7}$)

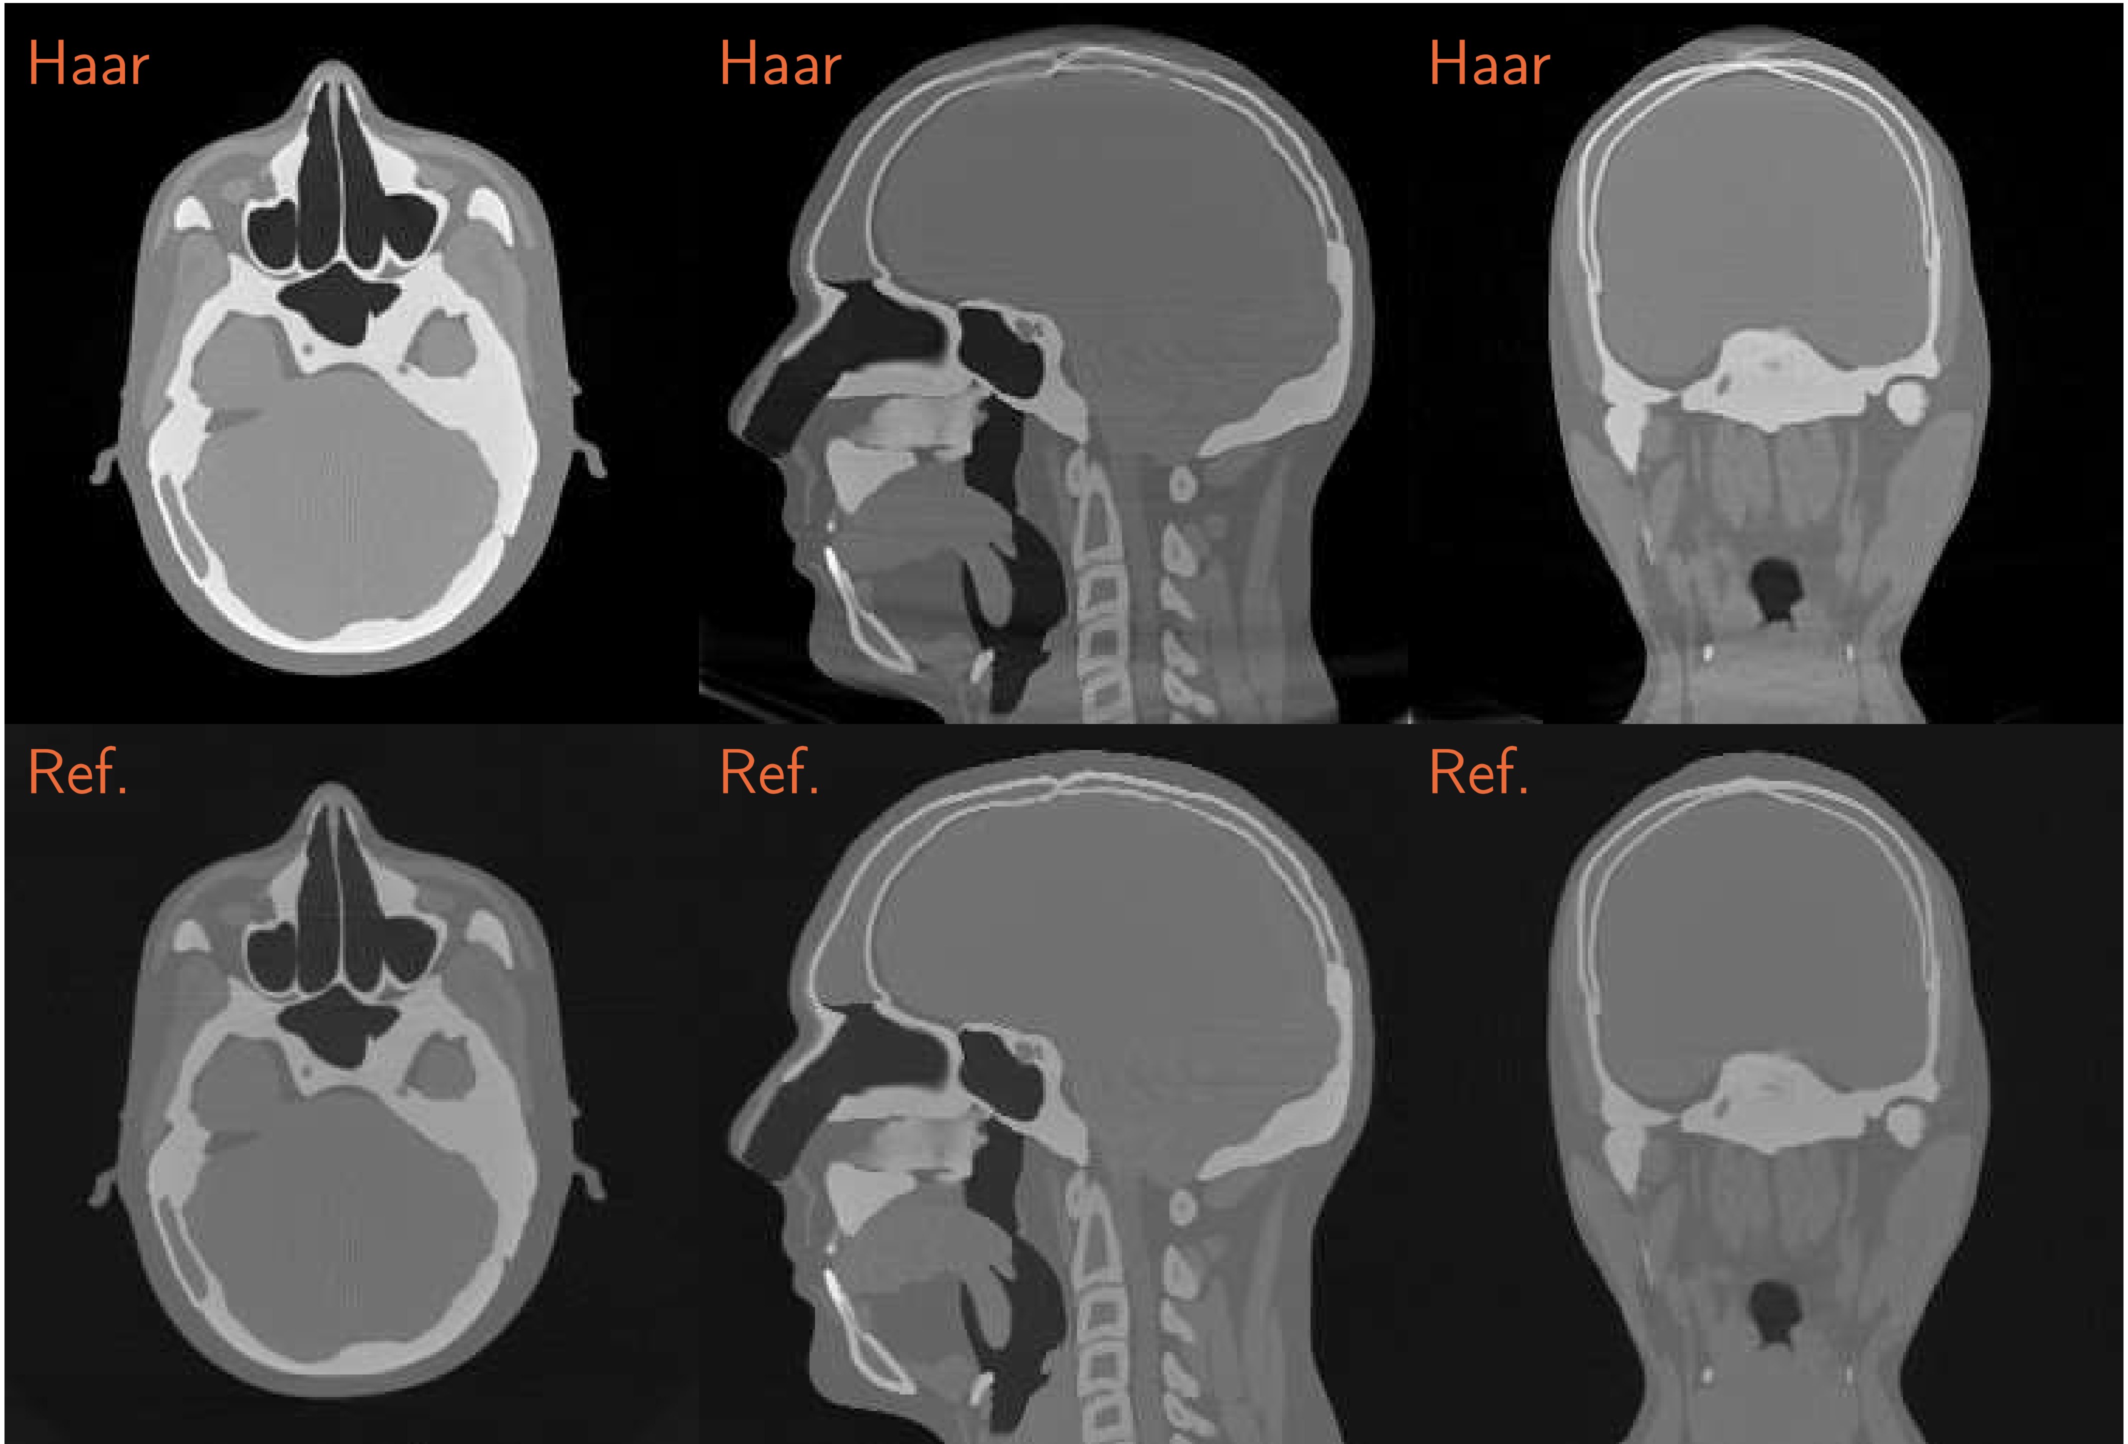

Haar wavelets, dose 10%, μ too large ($5 \cdot 10^{-5}$)

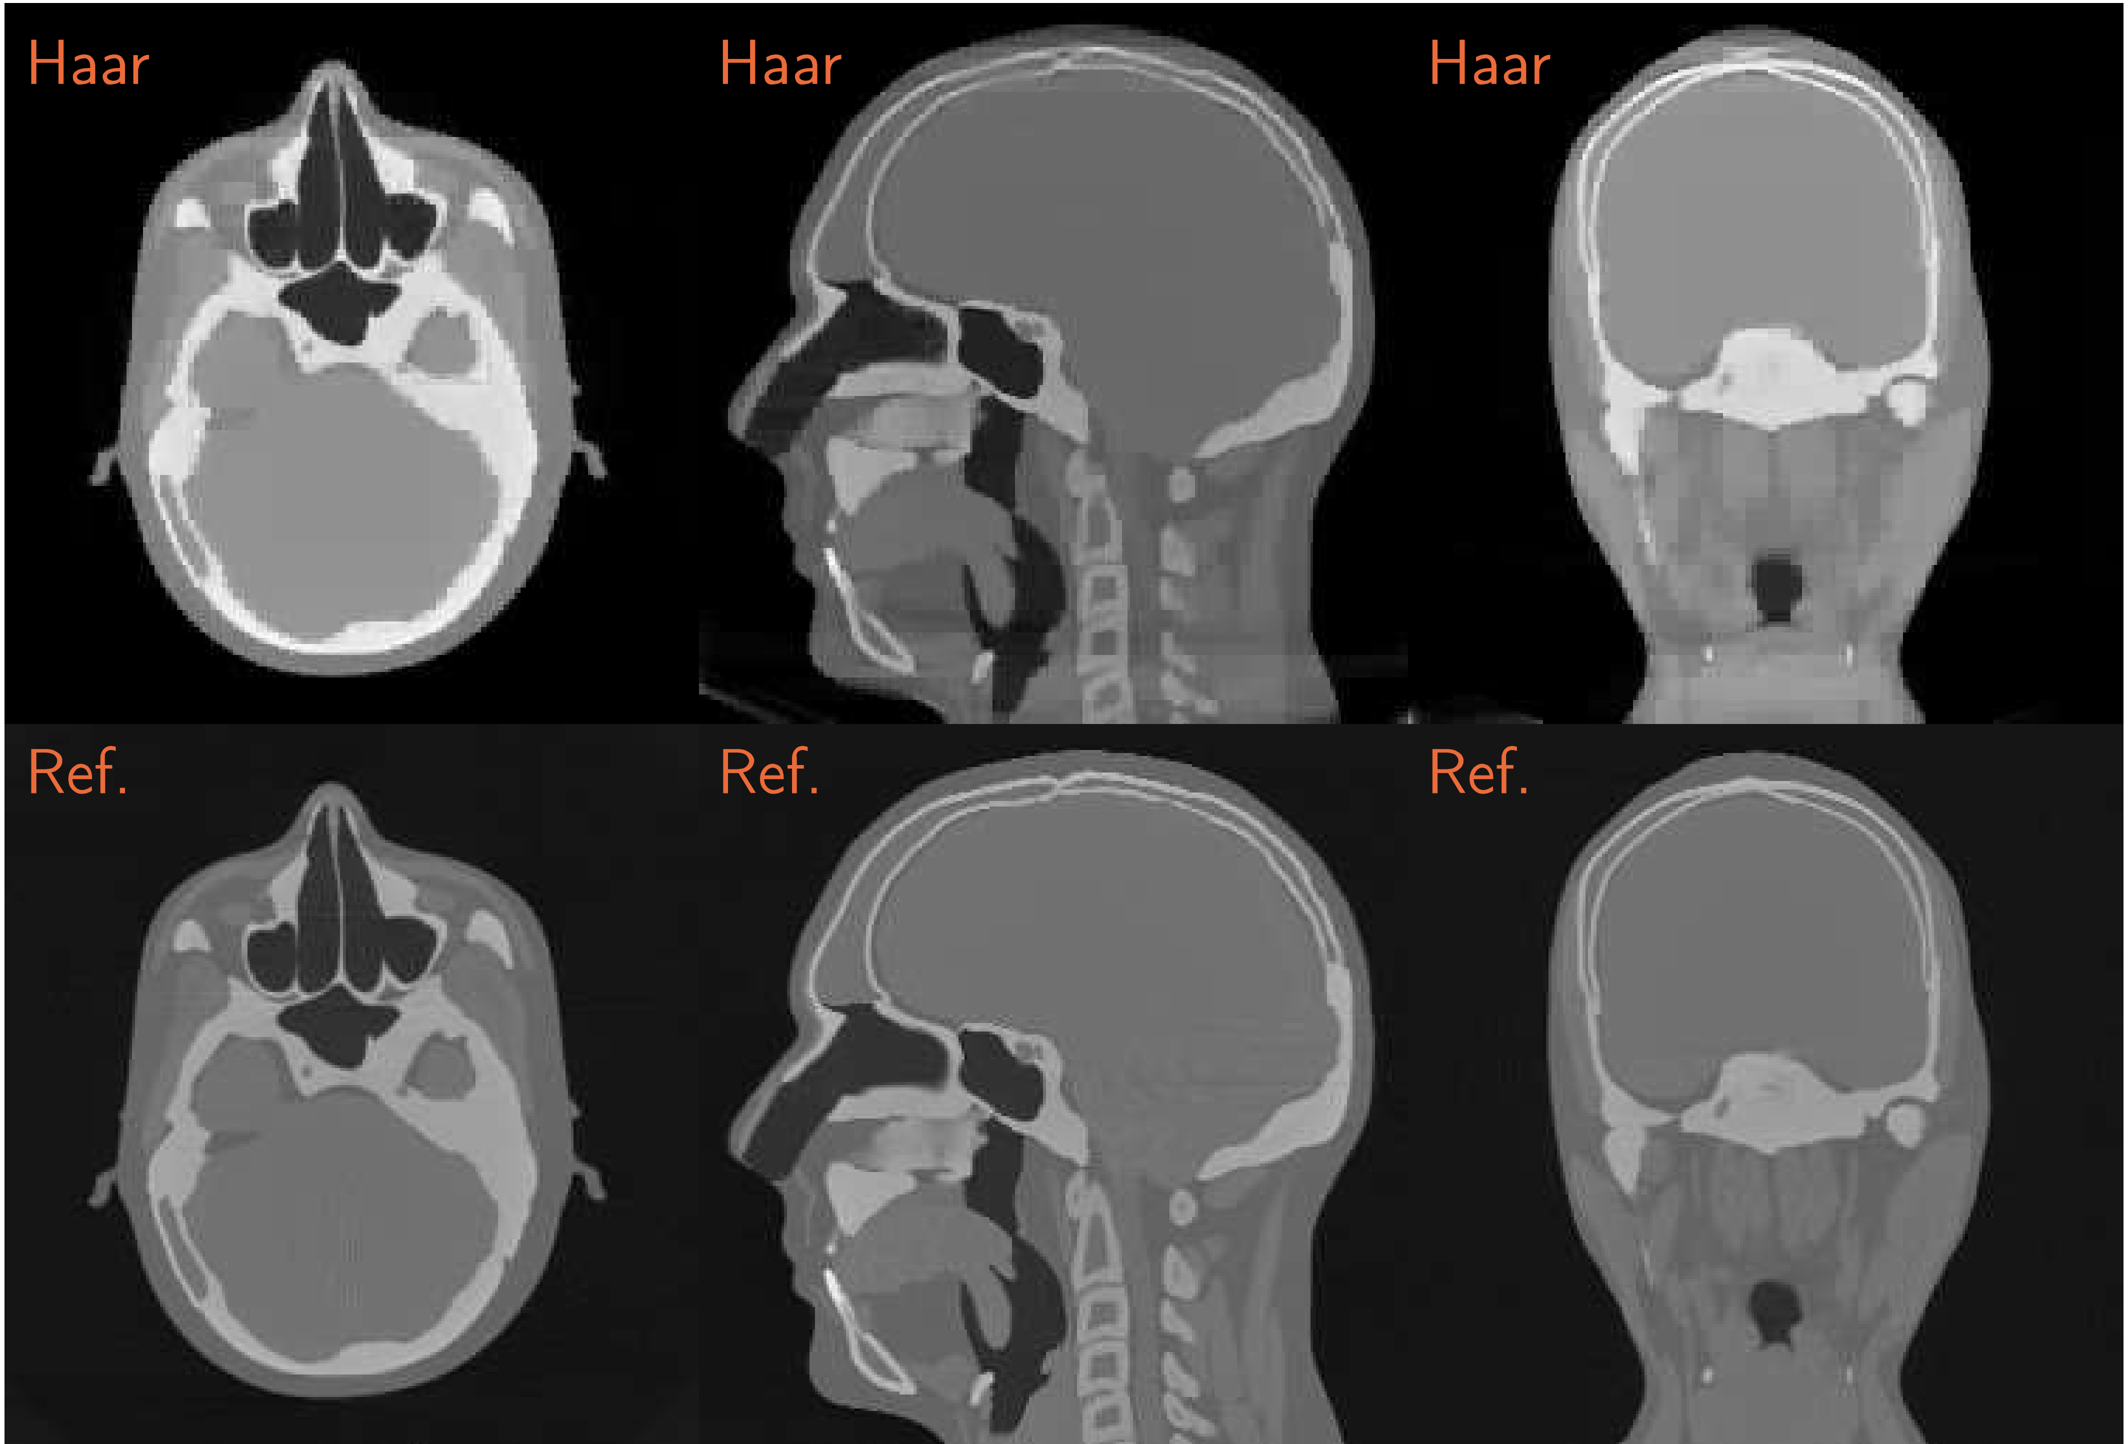

Haar wavelets, dose 10%, μ suitable $(5 \cdot 10^{-6})$

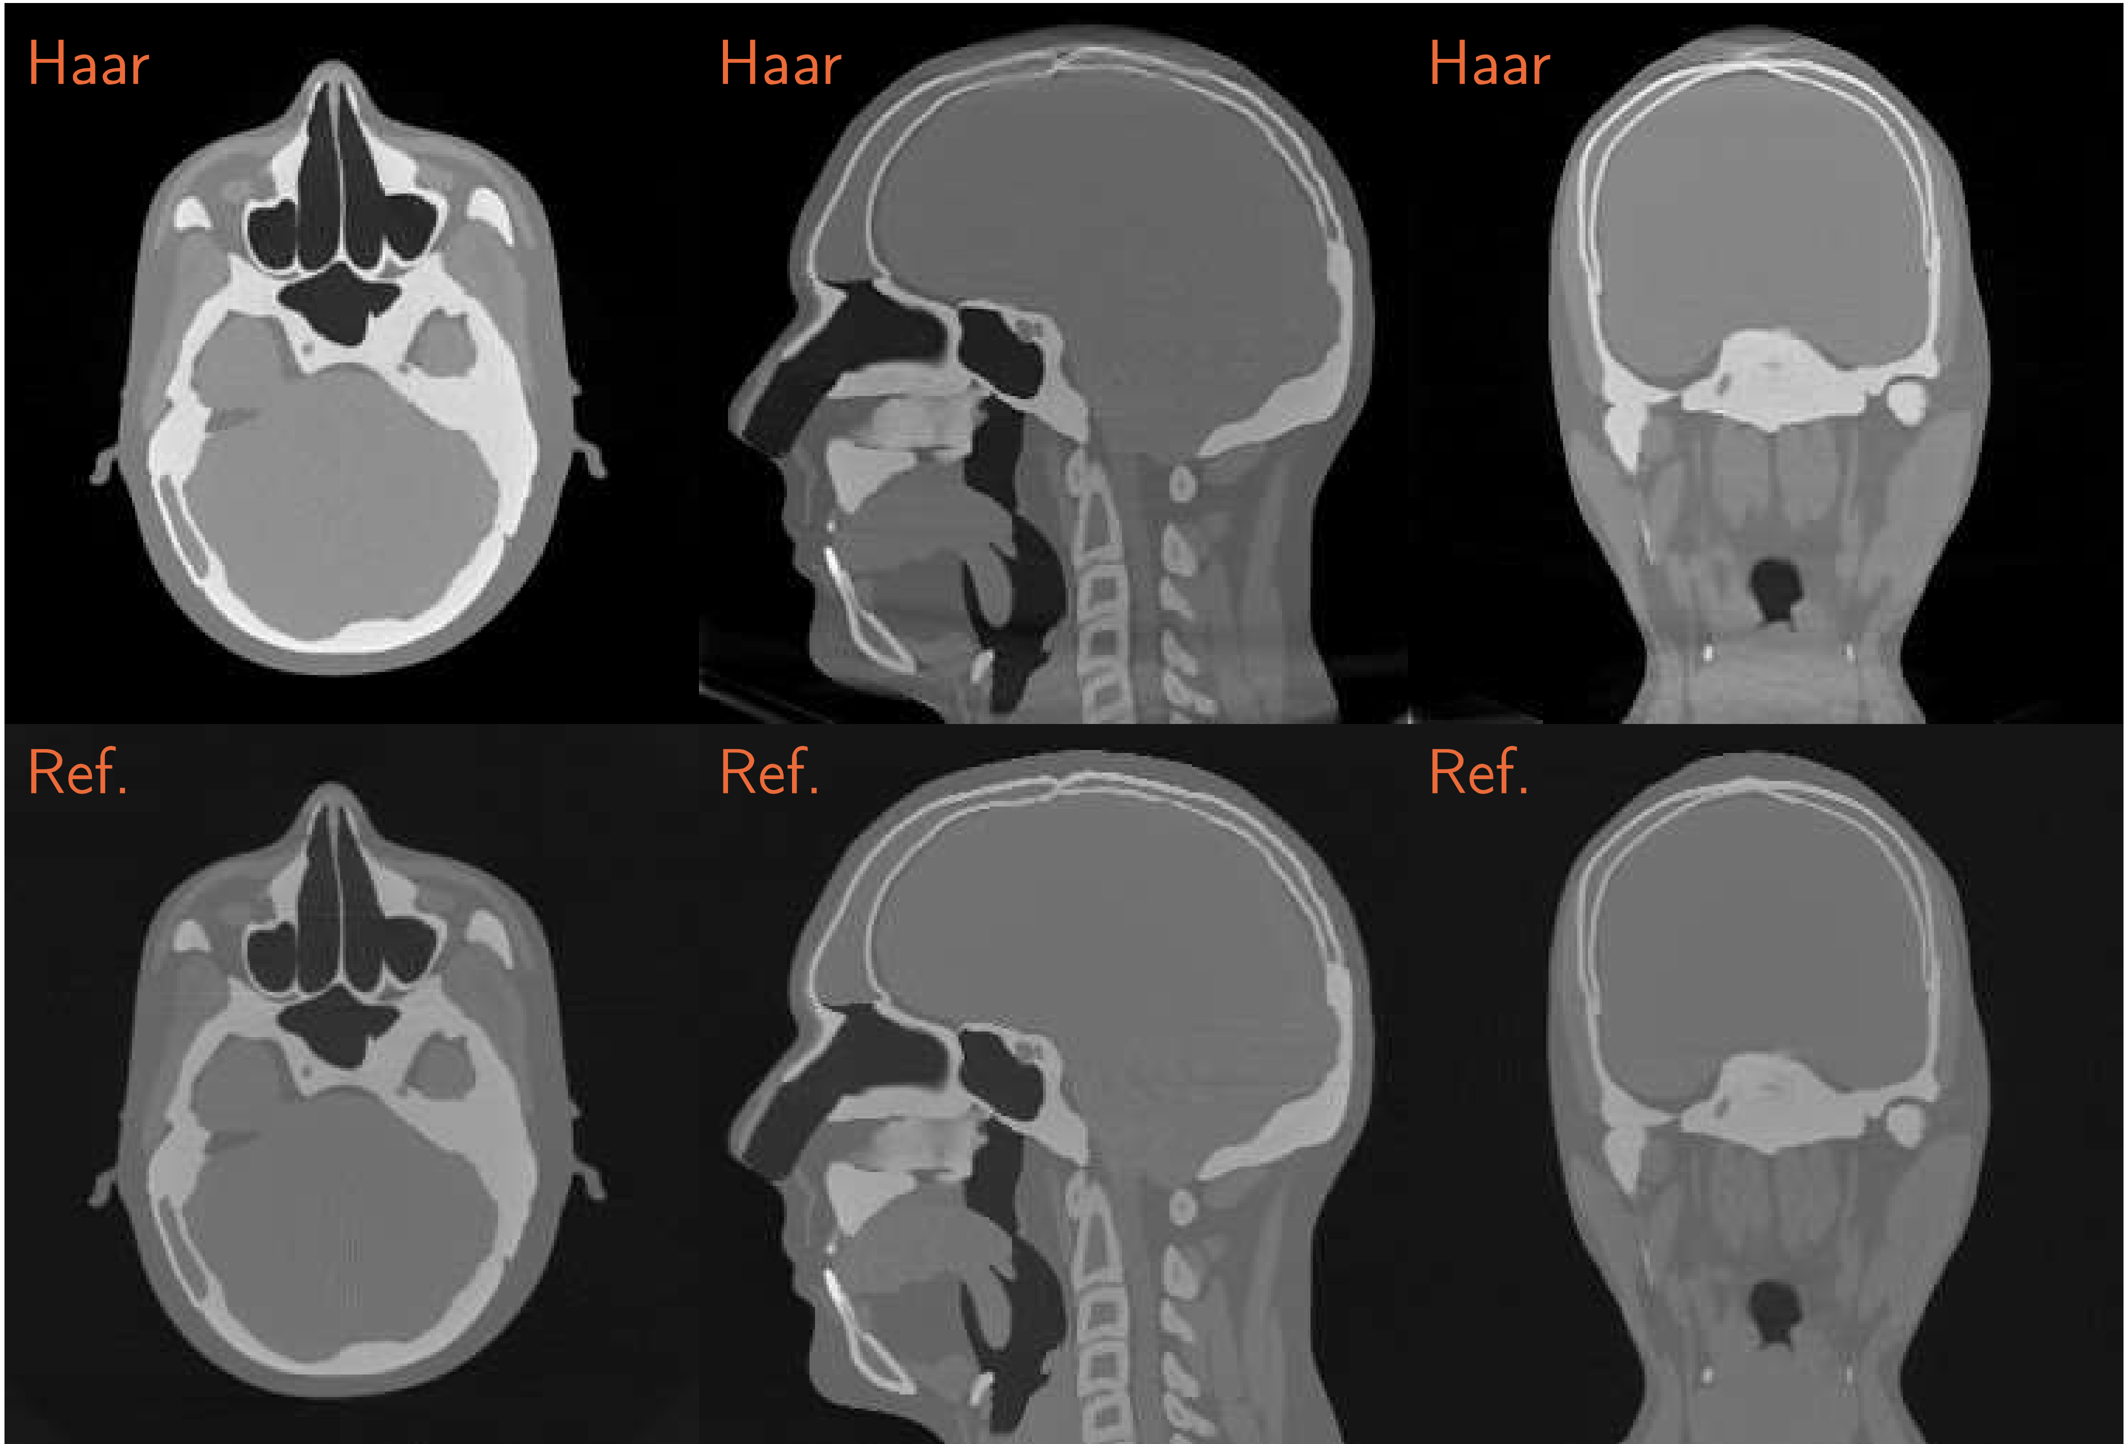

Haar wavelets, dose 10%, μ too small ($5 \cdot 10^{-7}$)

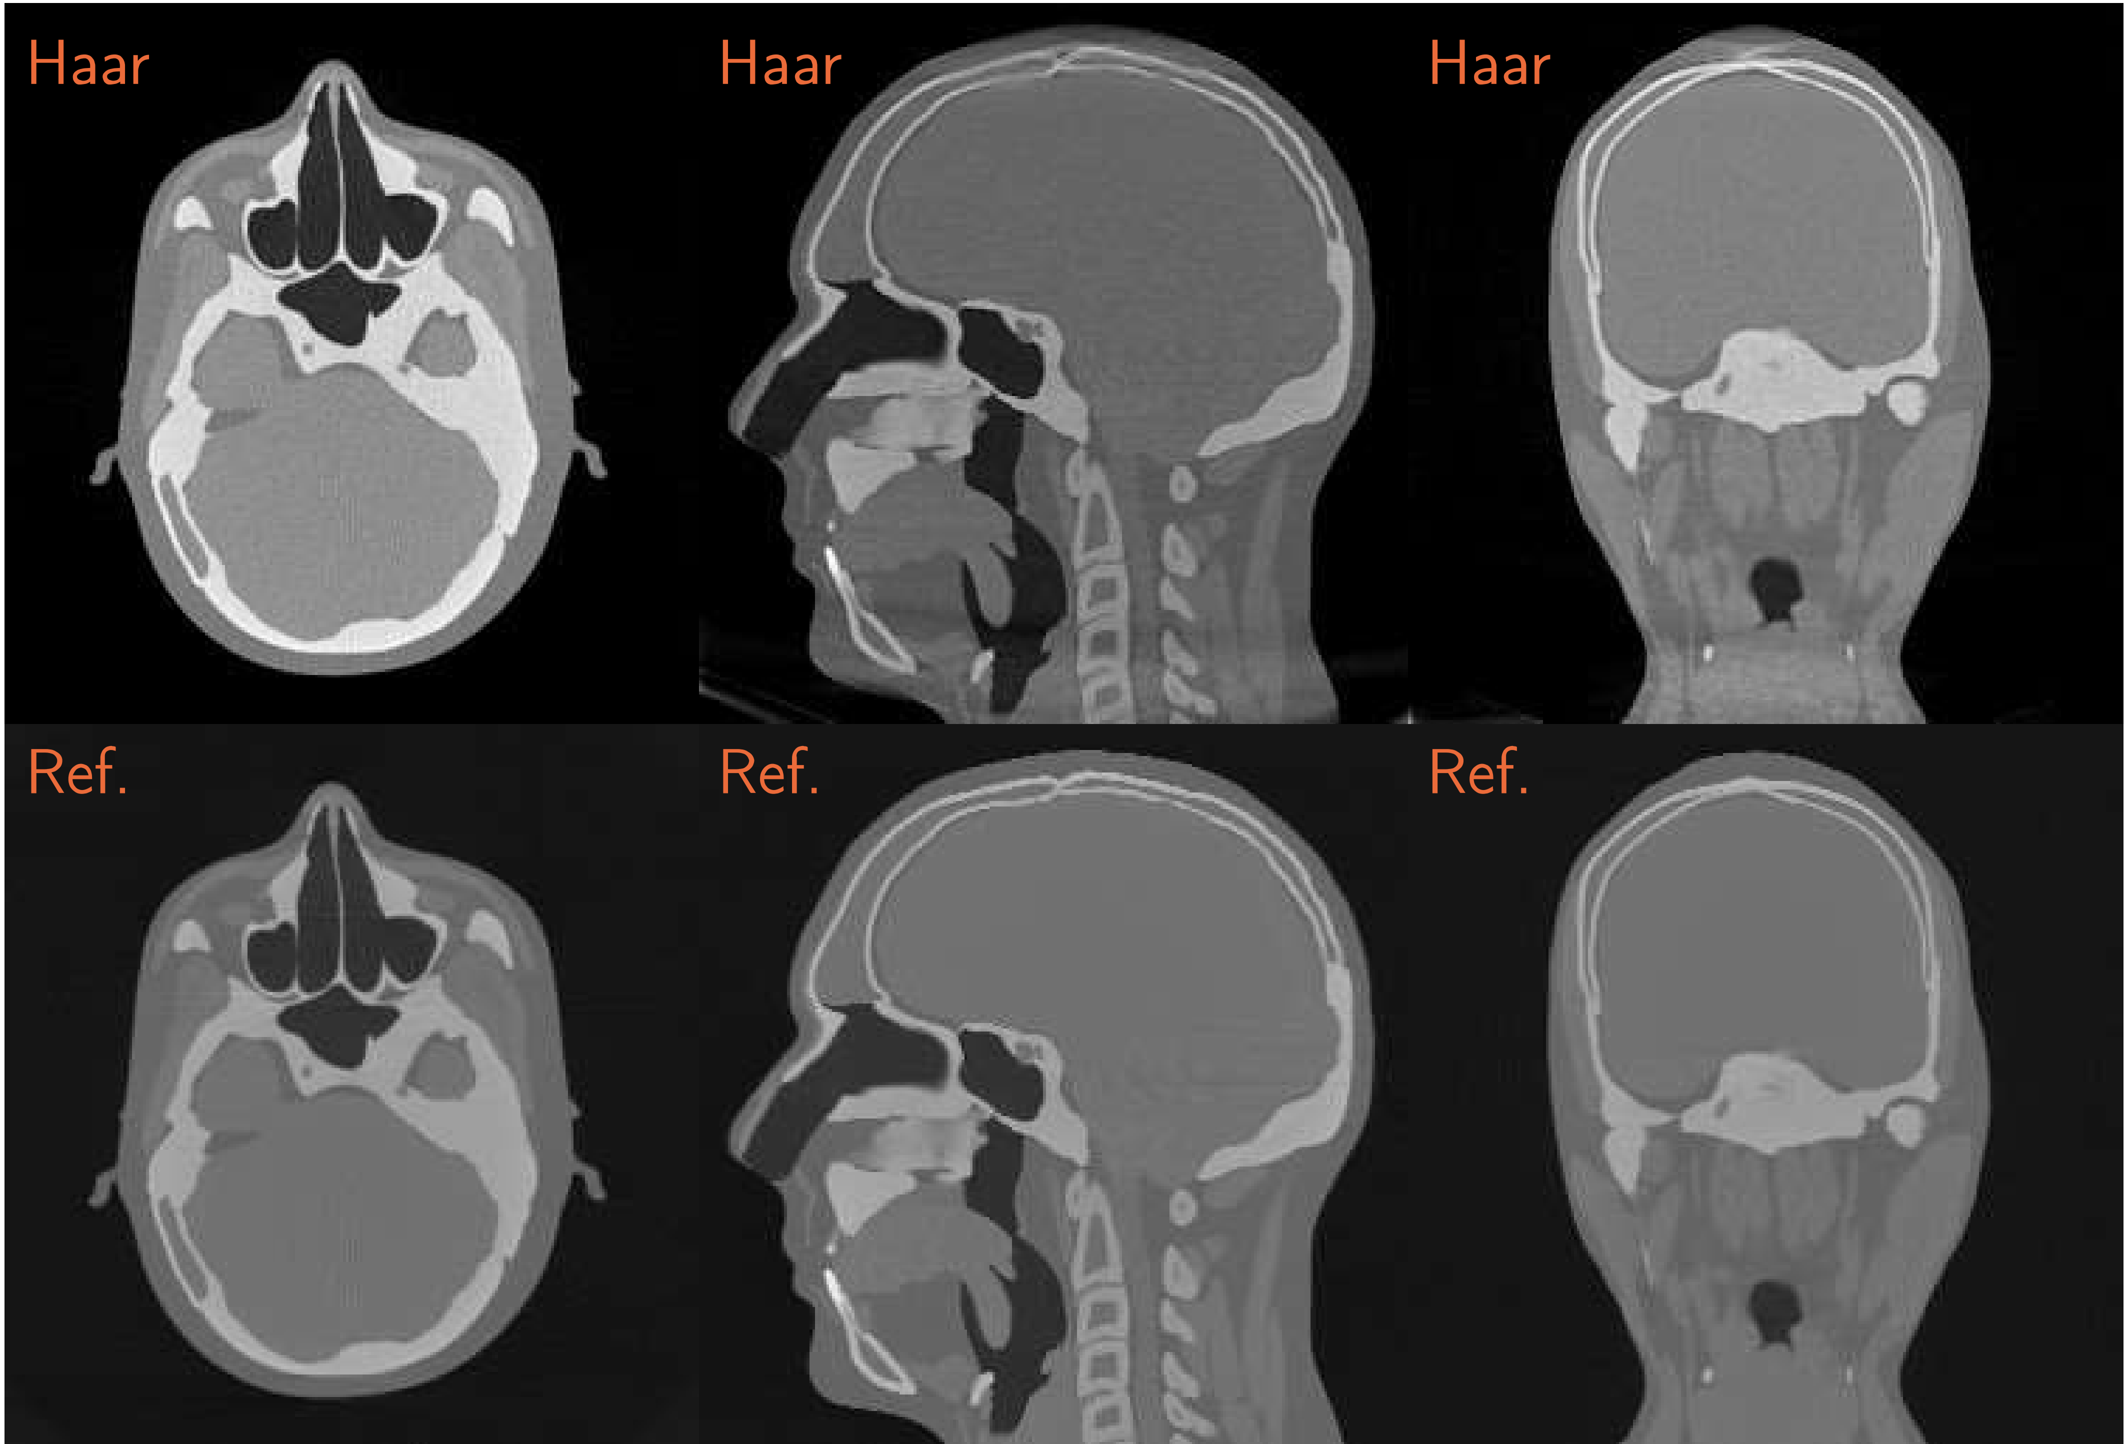

Haar wavelets, dose 1%, μ too large ($5 \cdot 10^{-5}$)

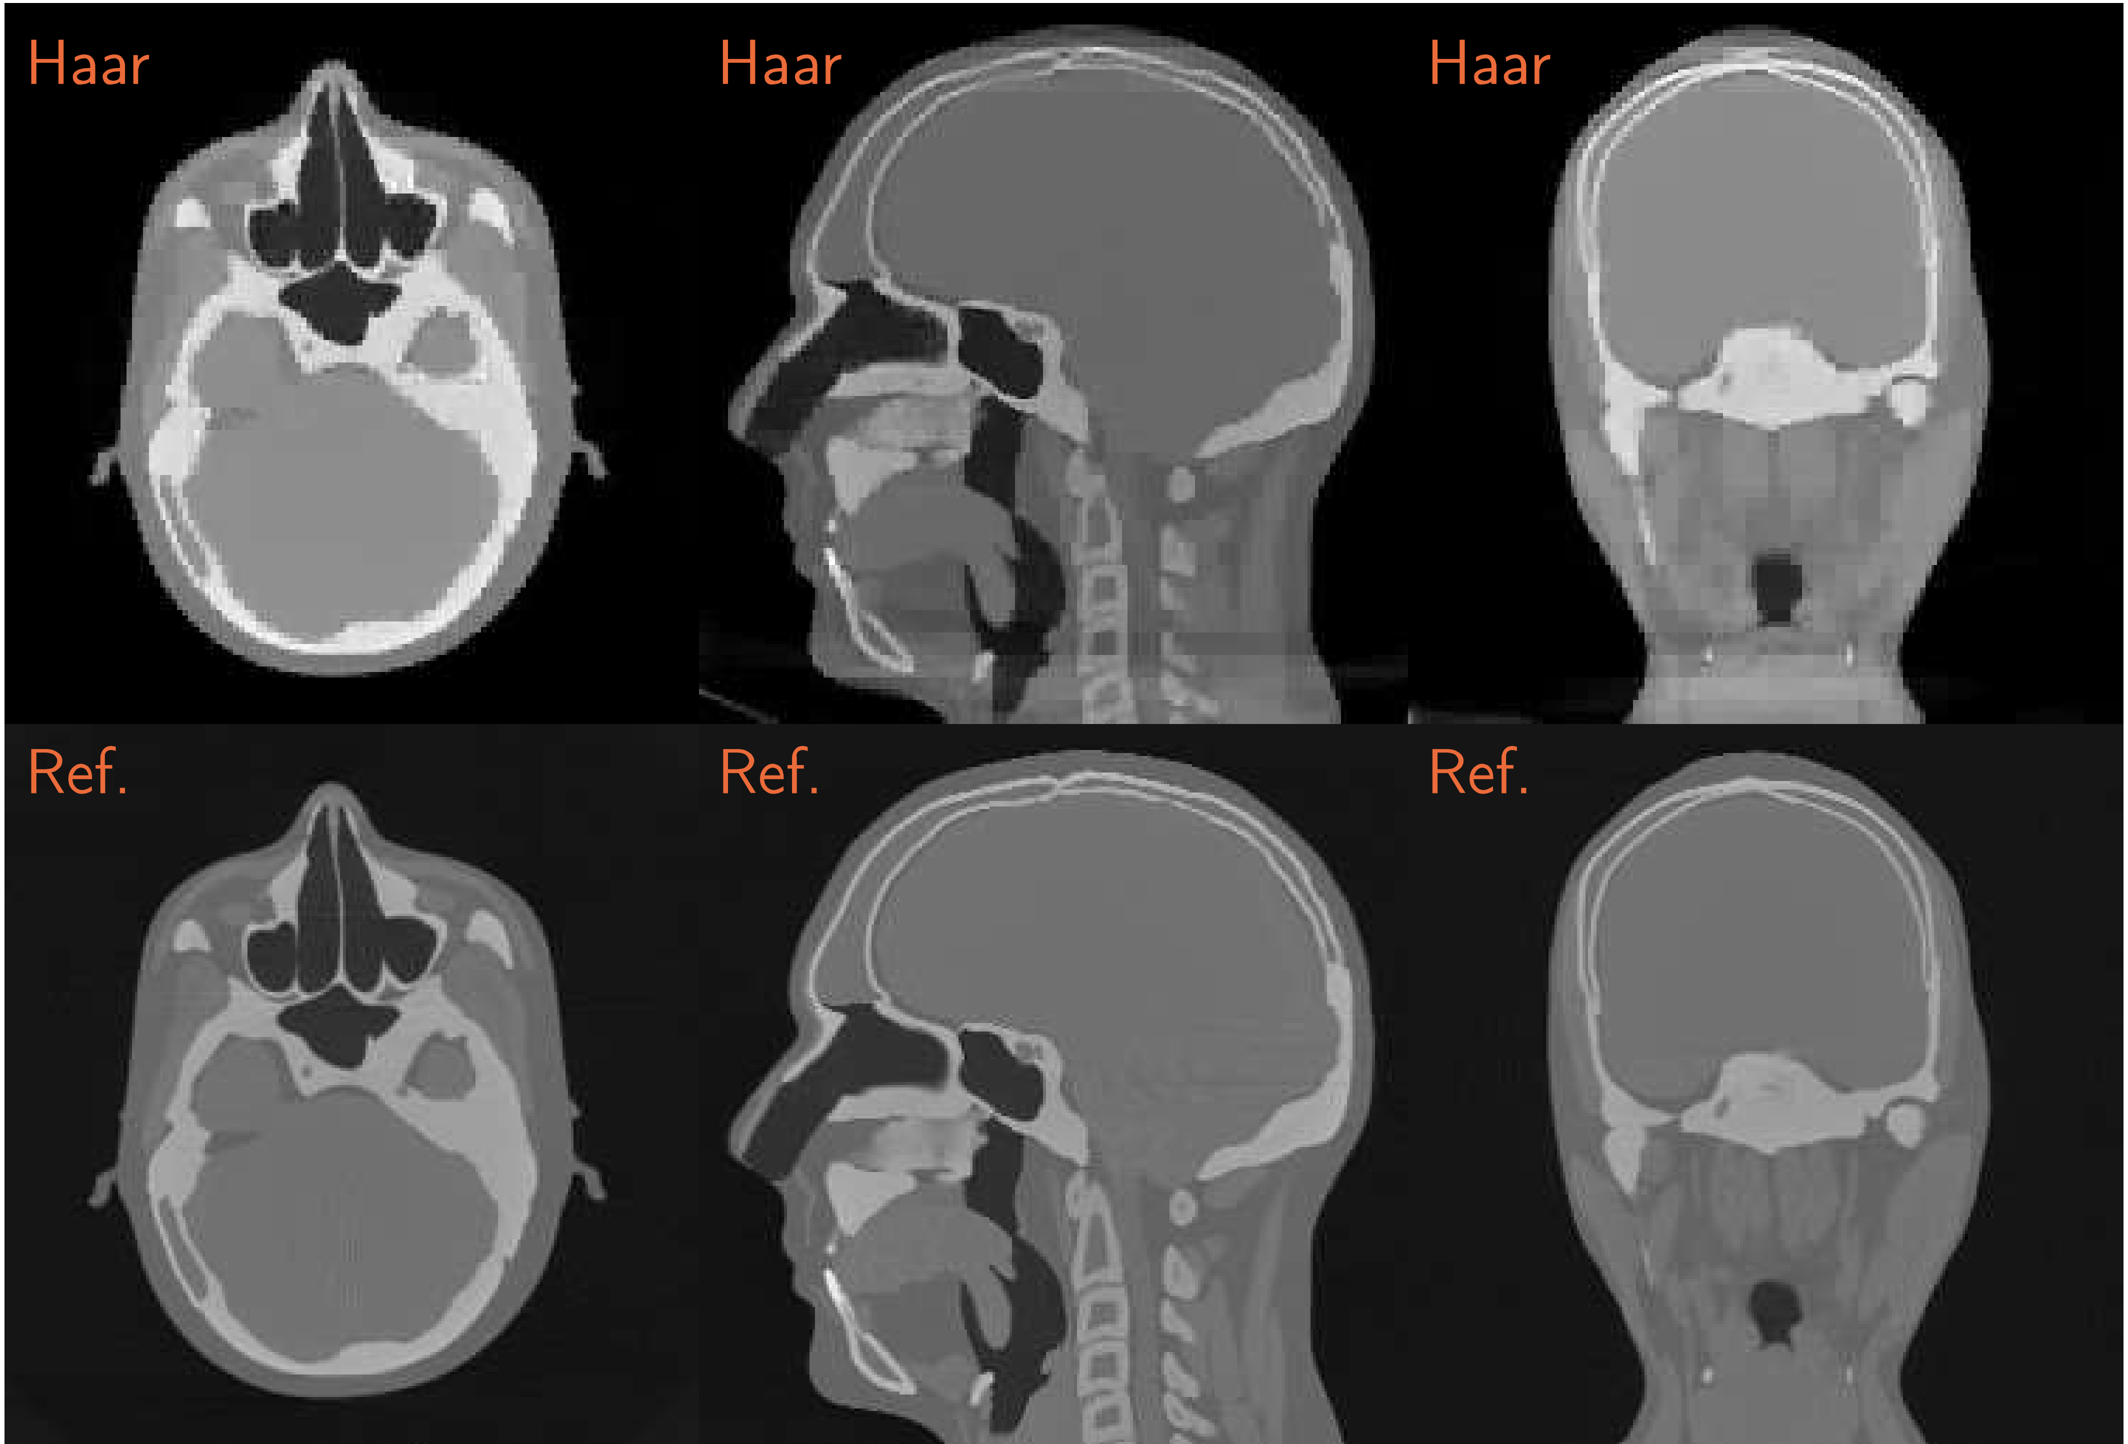

Haar wavelets, dose 1%, μ suitable $(1 \cdot 10^{-5})$

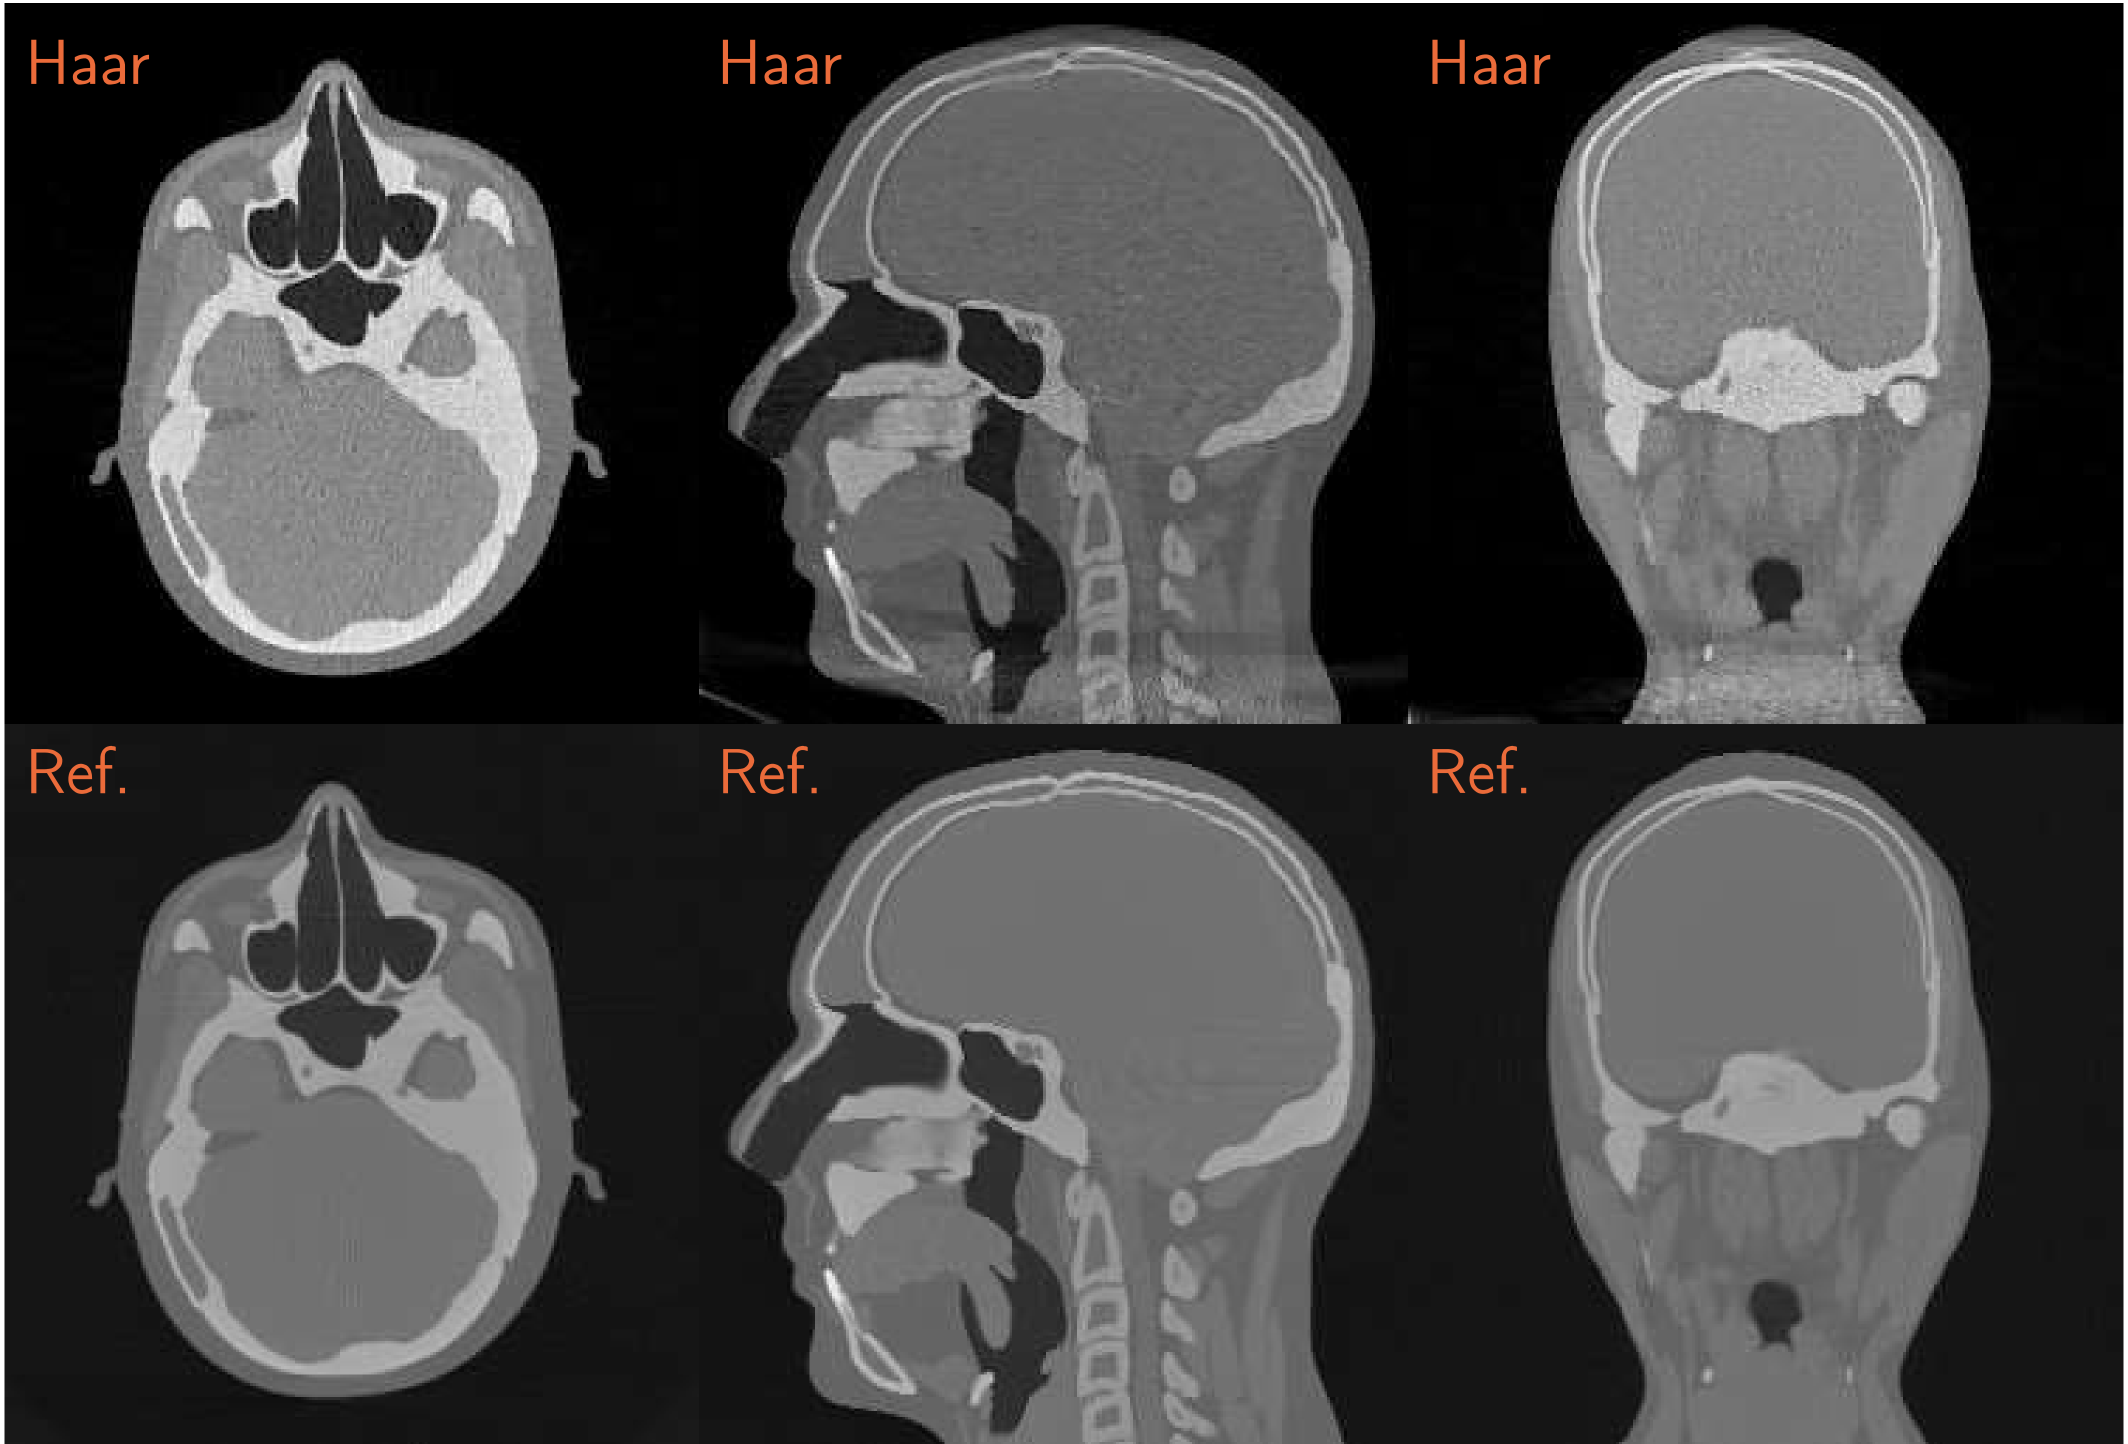

Haar wavelets, dose 1%, μ too small $(1 \cdot 10^{-6})$

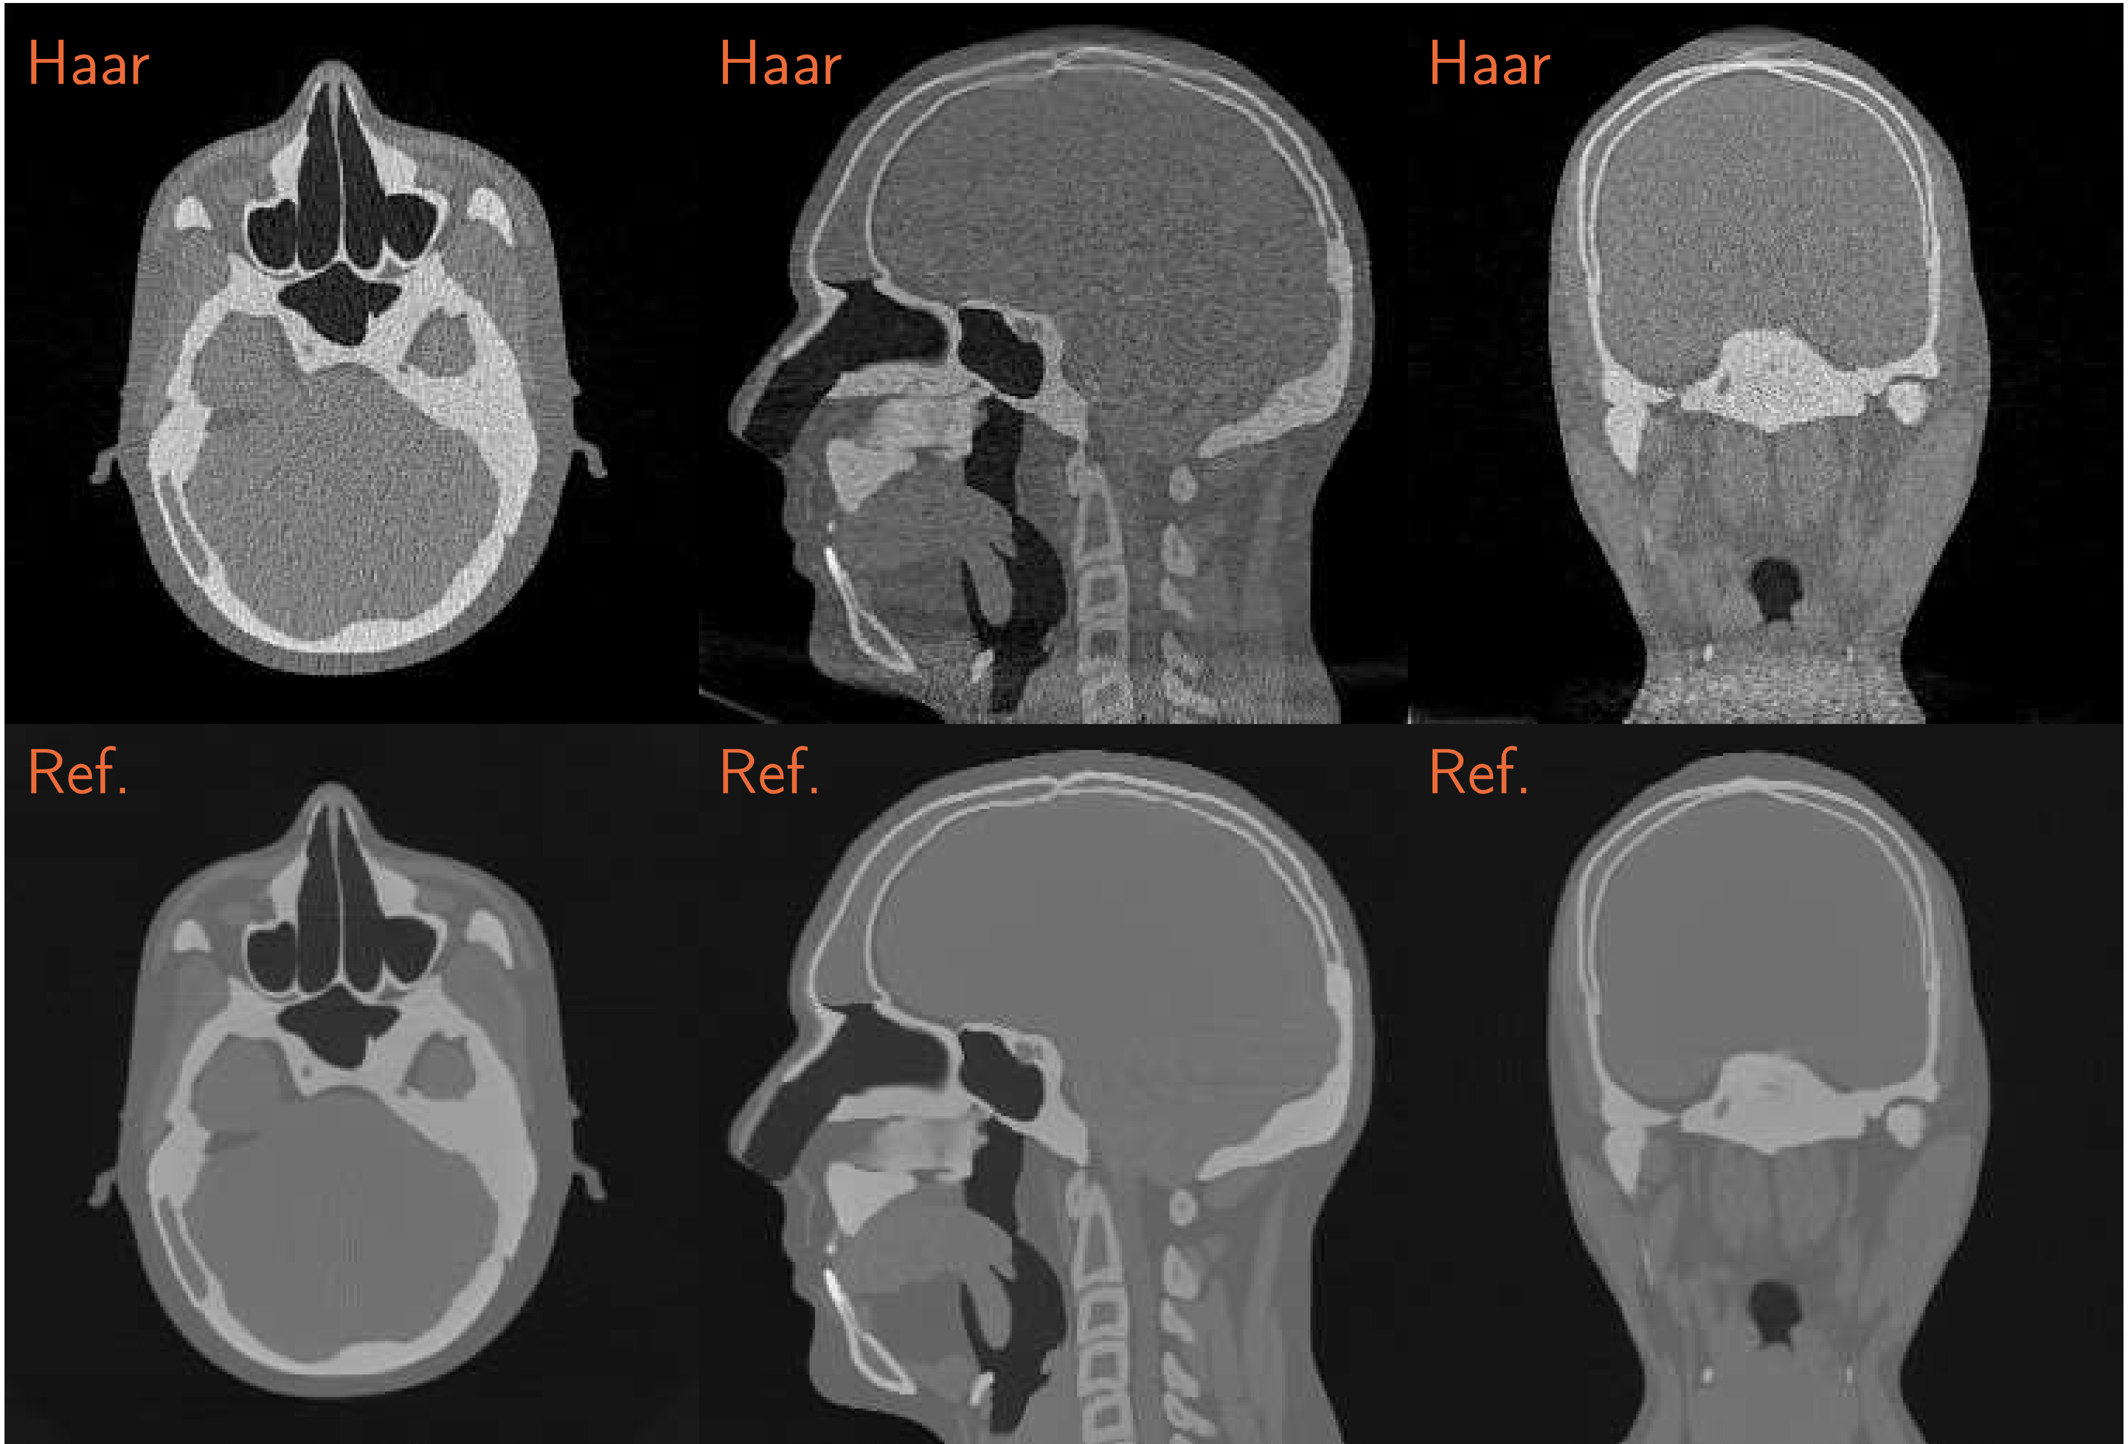

Haar wavelets, dose 0.1%, μ too large ($1 \cdot 10^{-4}$)

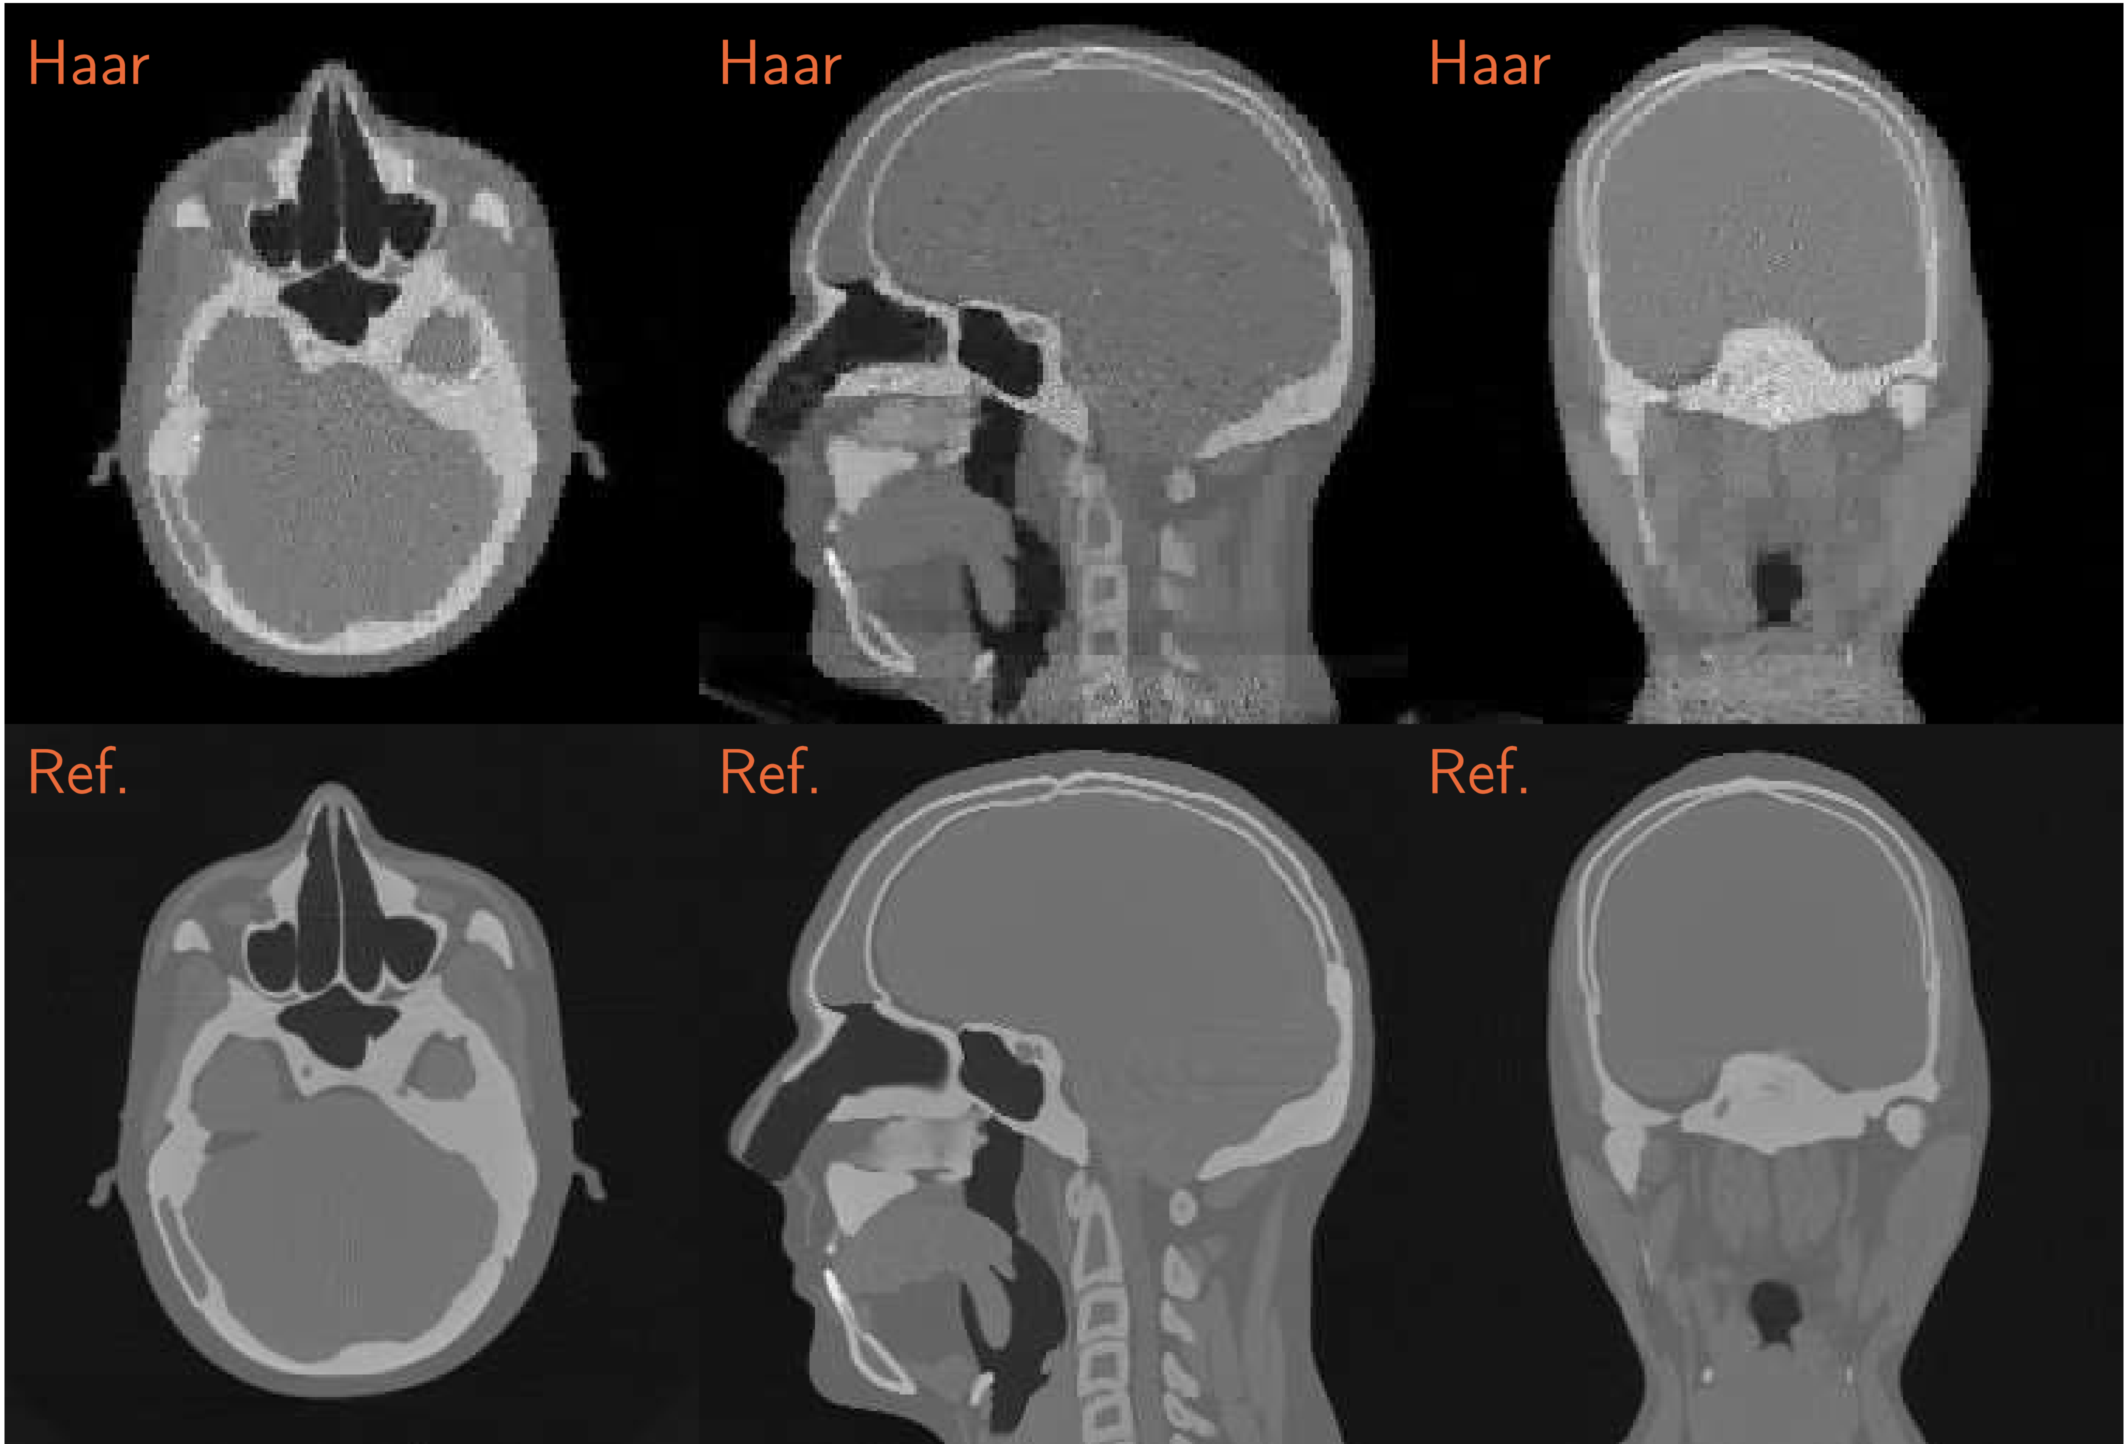

Haar wavelets, dose 0.1%, μ suitable (?) $(5 \cdot 10^{-5})$

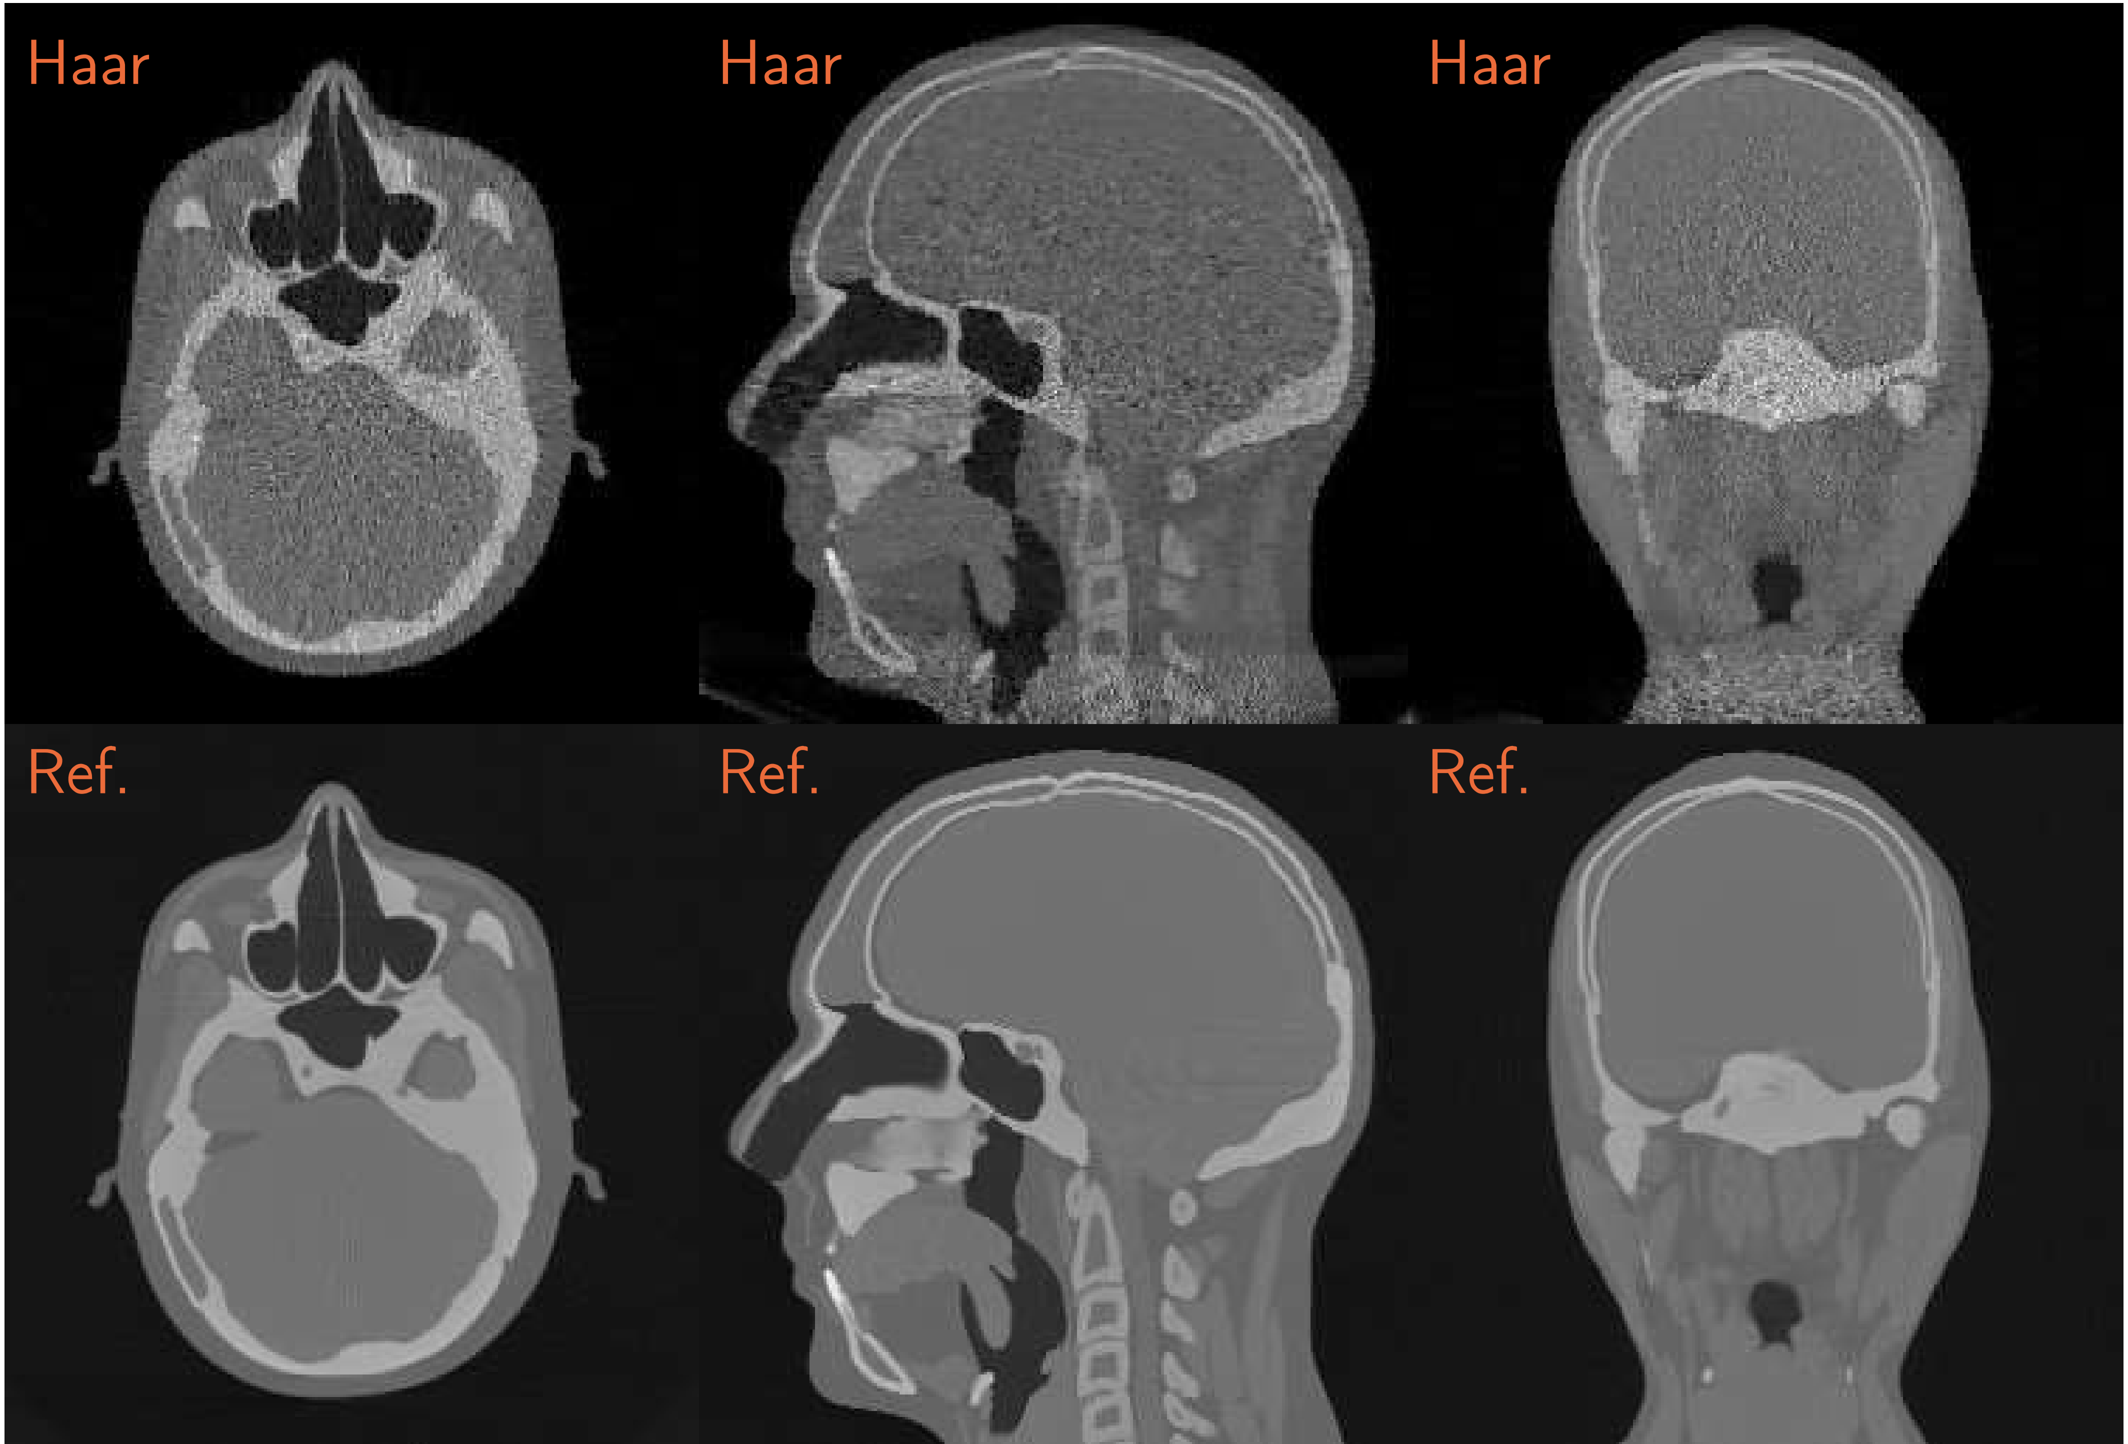

Haar wavelets, dose 0.1%, μ too small ($1 \cdot 10^{-5}$)

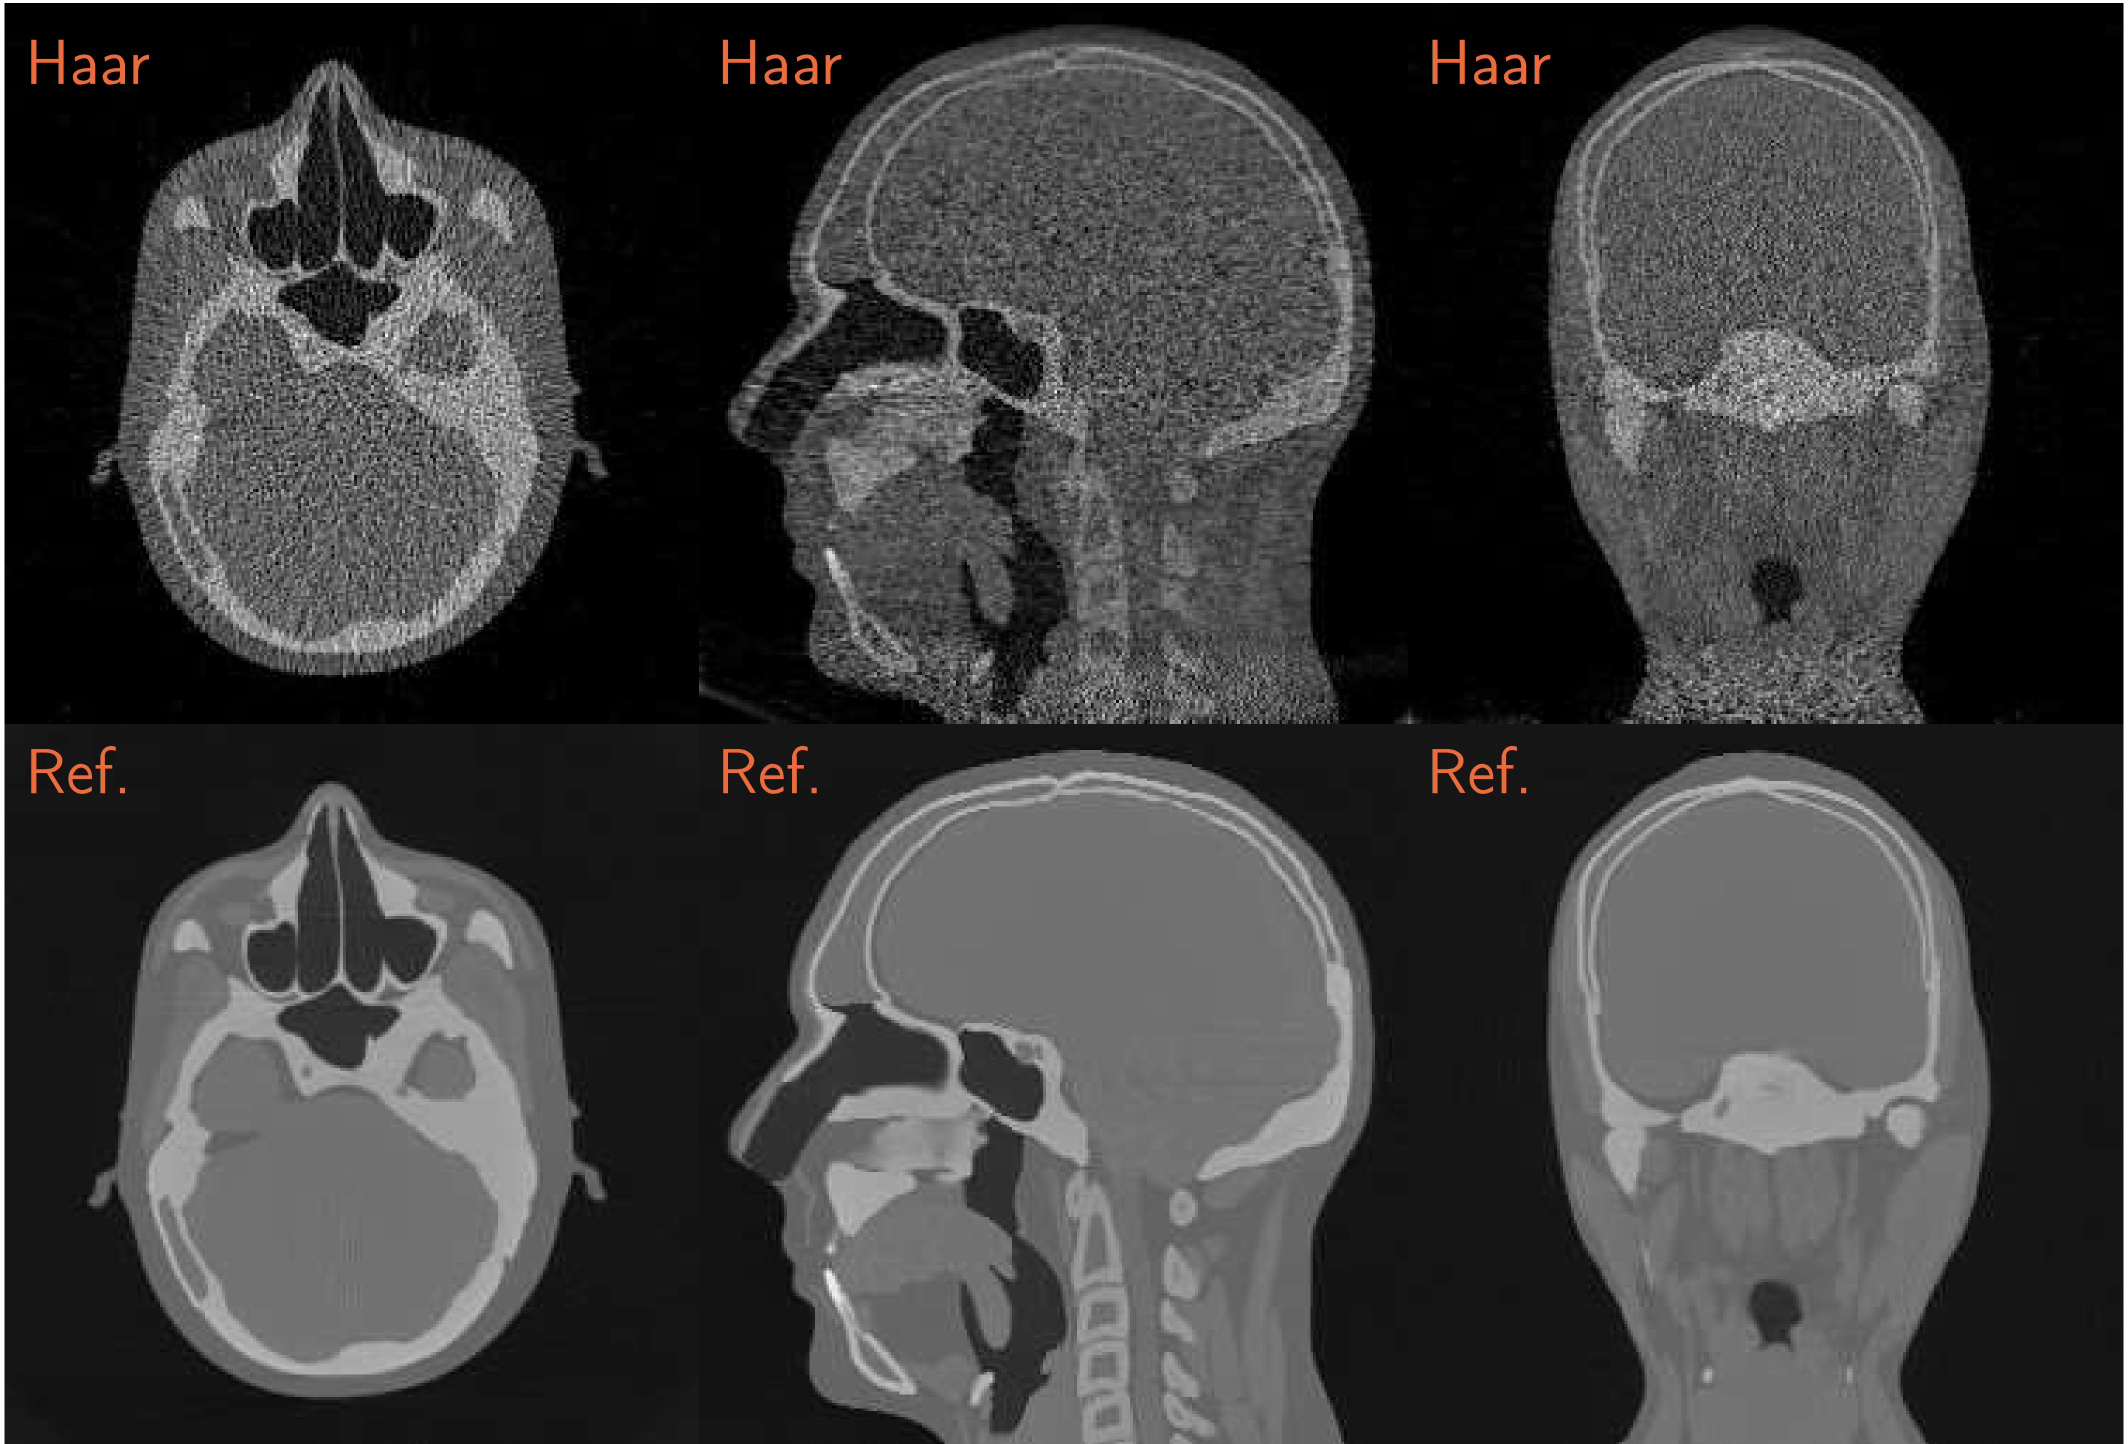

Haar Wavelet regularization

The true test:

How does Haar wavelet regularization compare to FDK?

Haar-CT vs. FDK, dose 100% ($\mu = 2.5 \cdot 10^{-6}$)

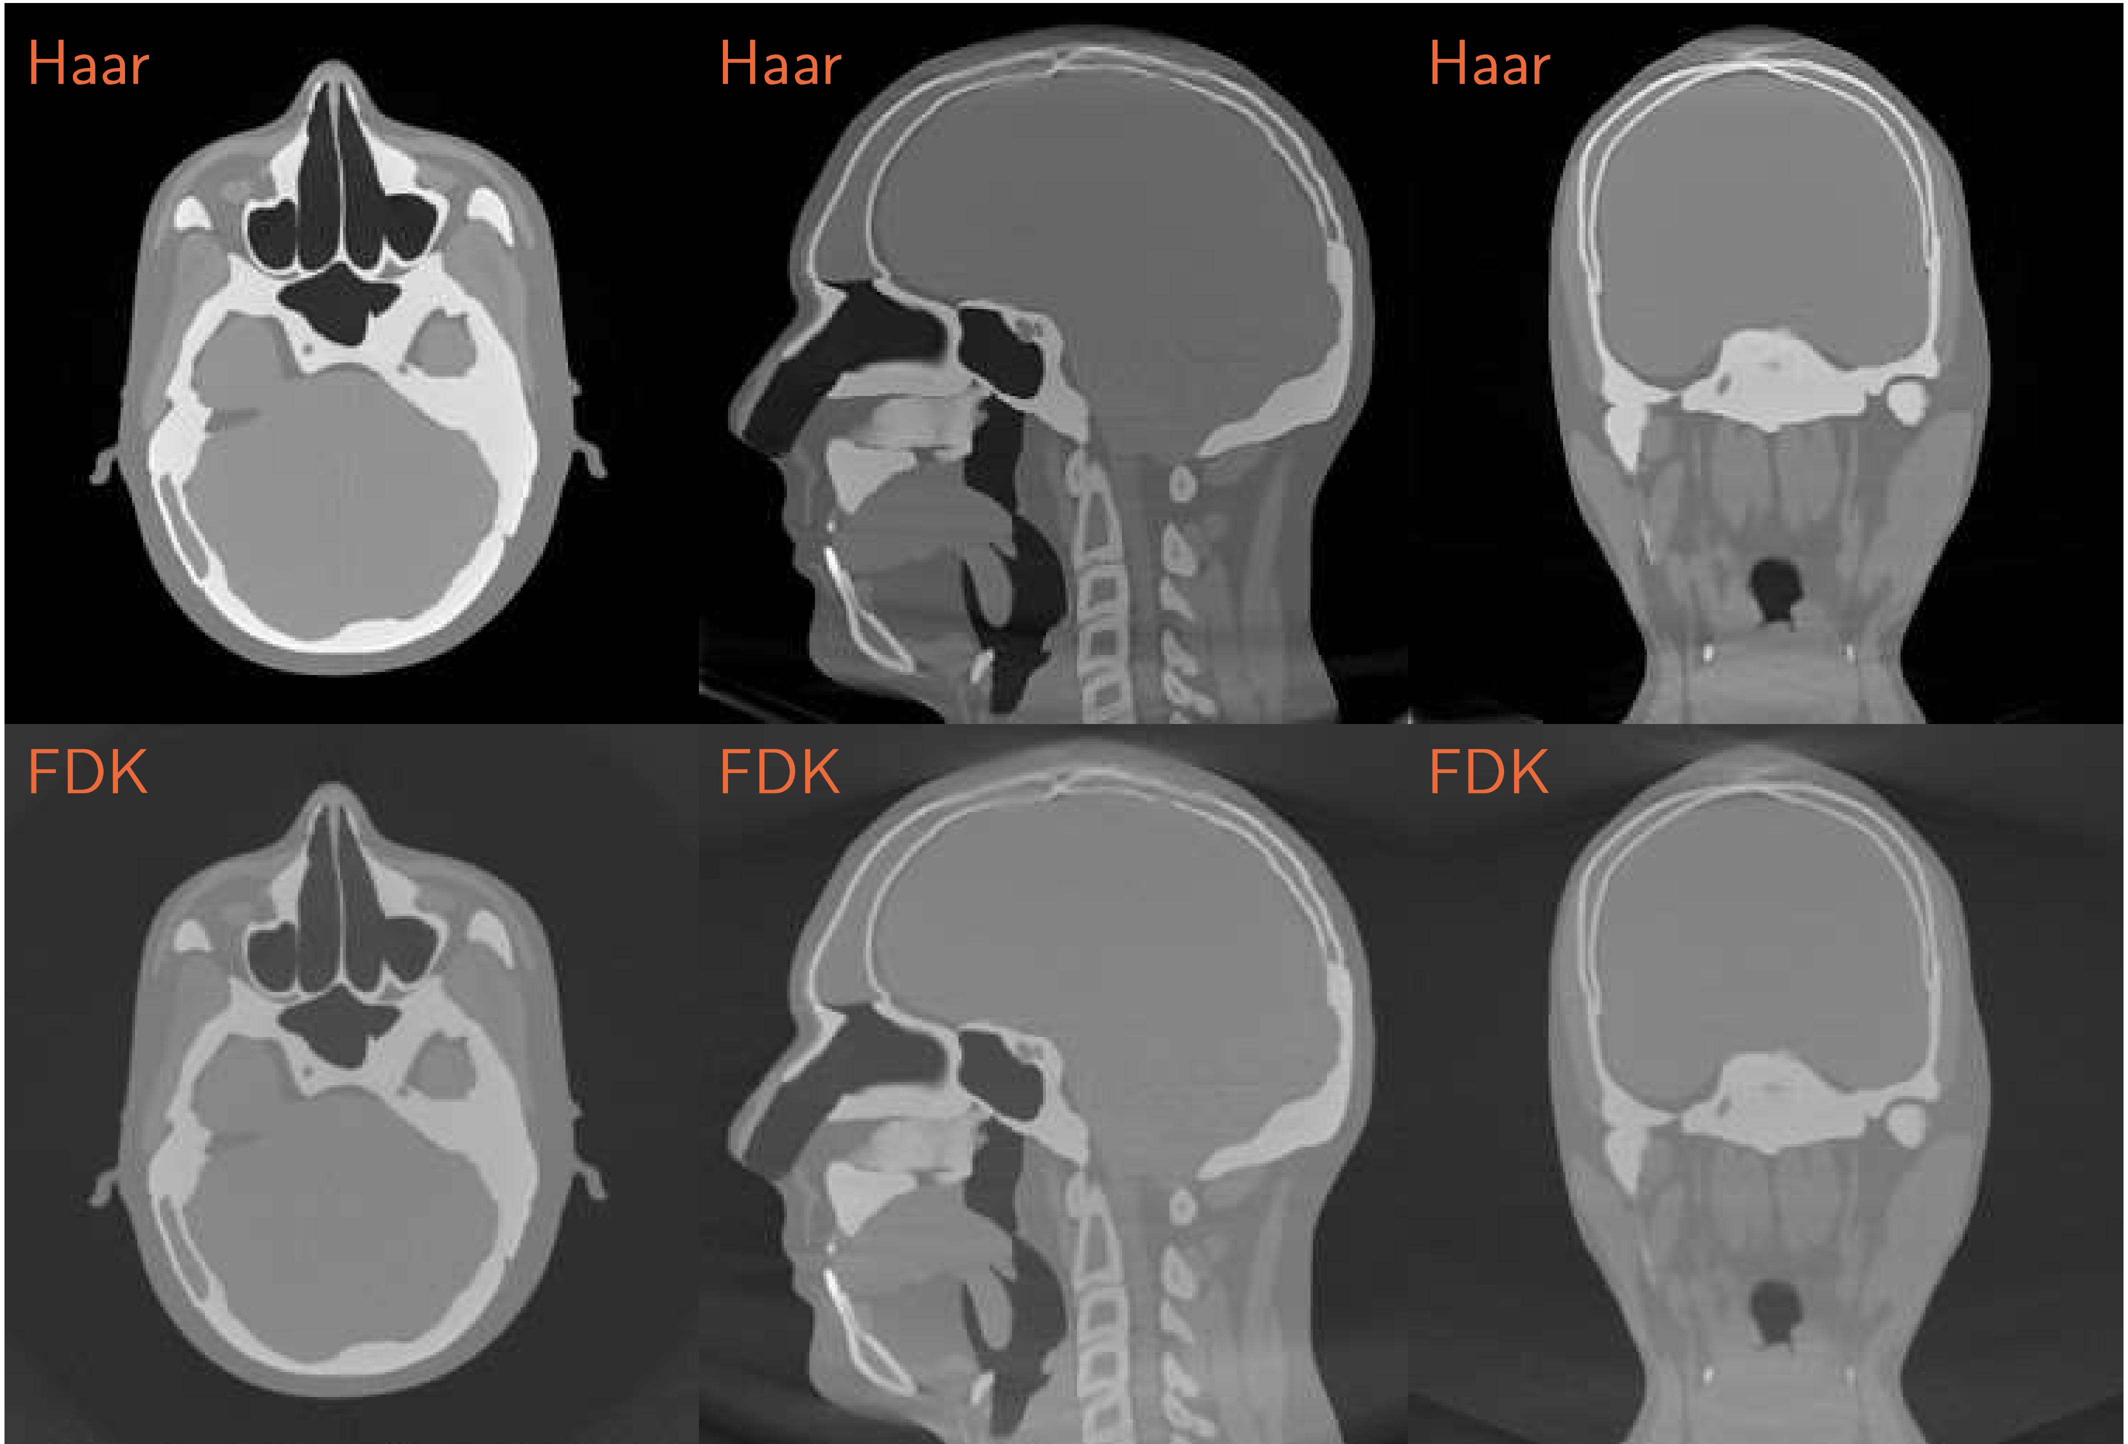

Haar-CT vs. FDK, dose 10% ($\mu = 5 \cdot 10^{-6}$)

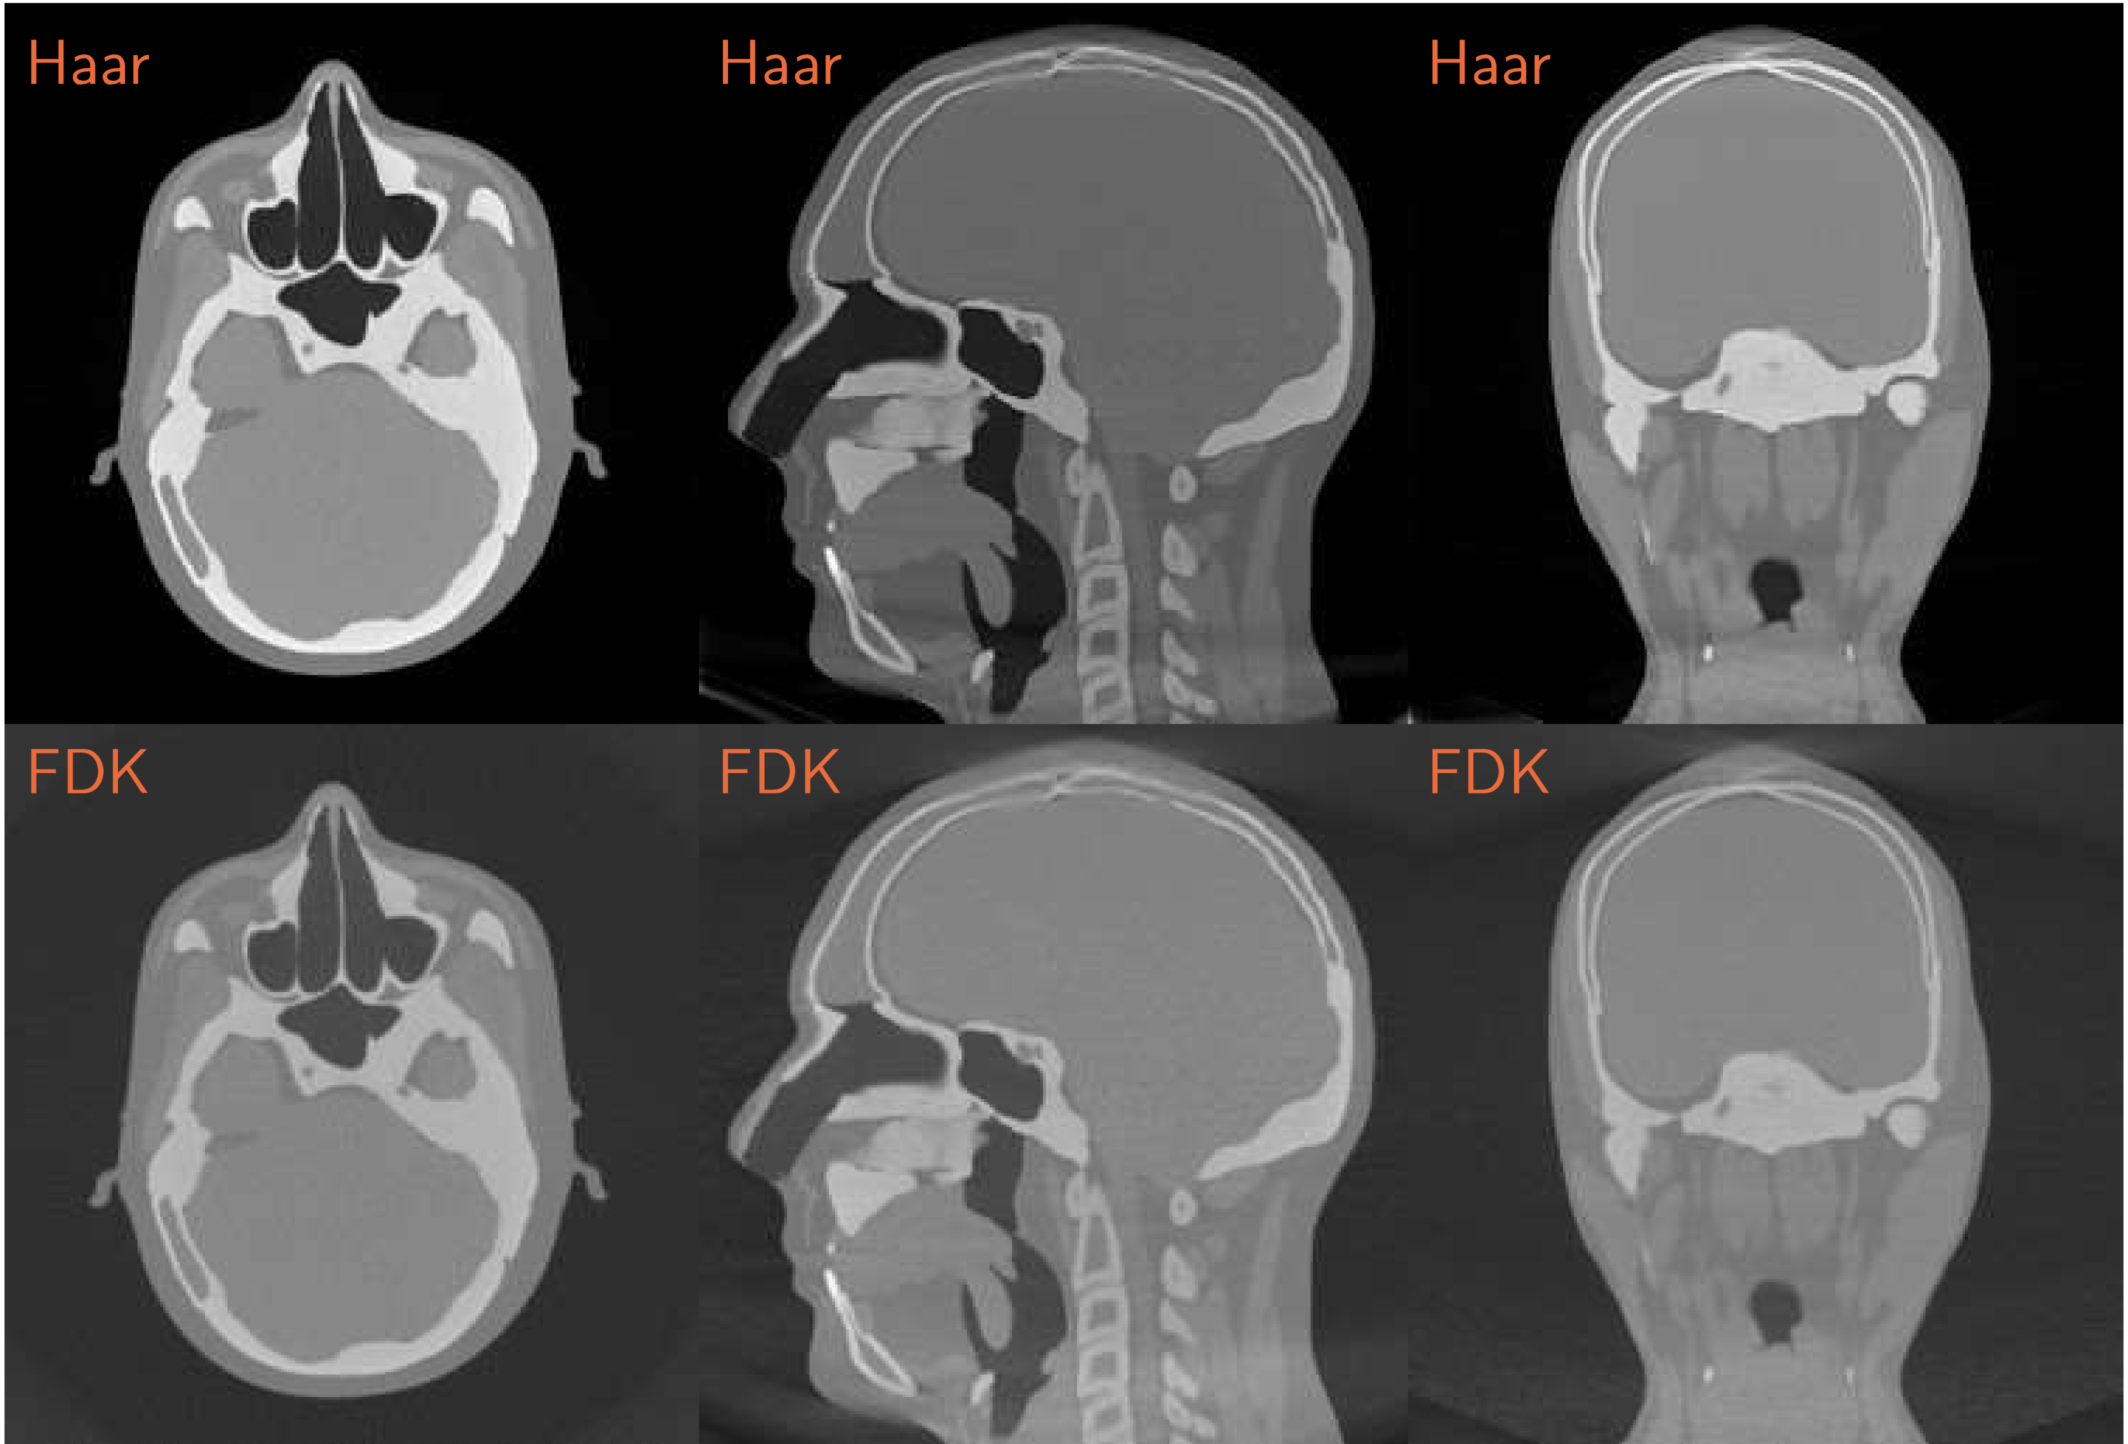

Haar-CT vs. FDK, dose 1% ($\mu = 1 \cdot 10^{-5}$)

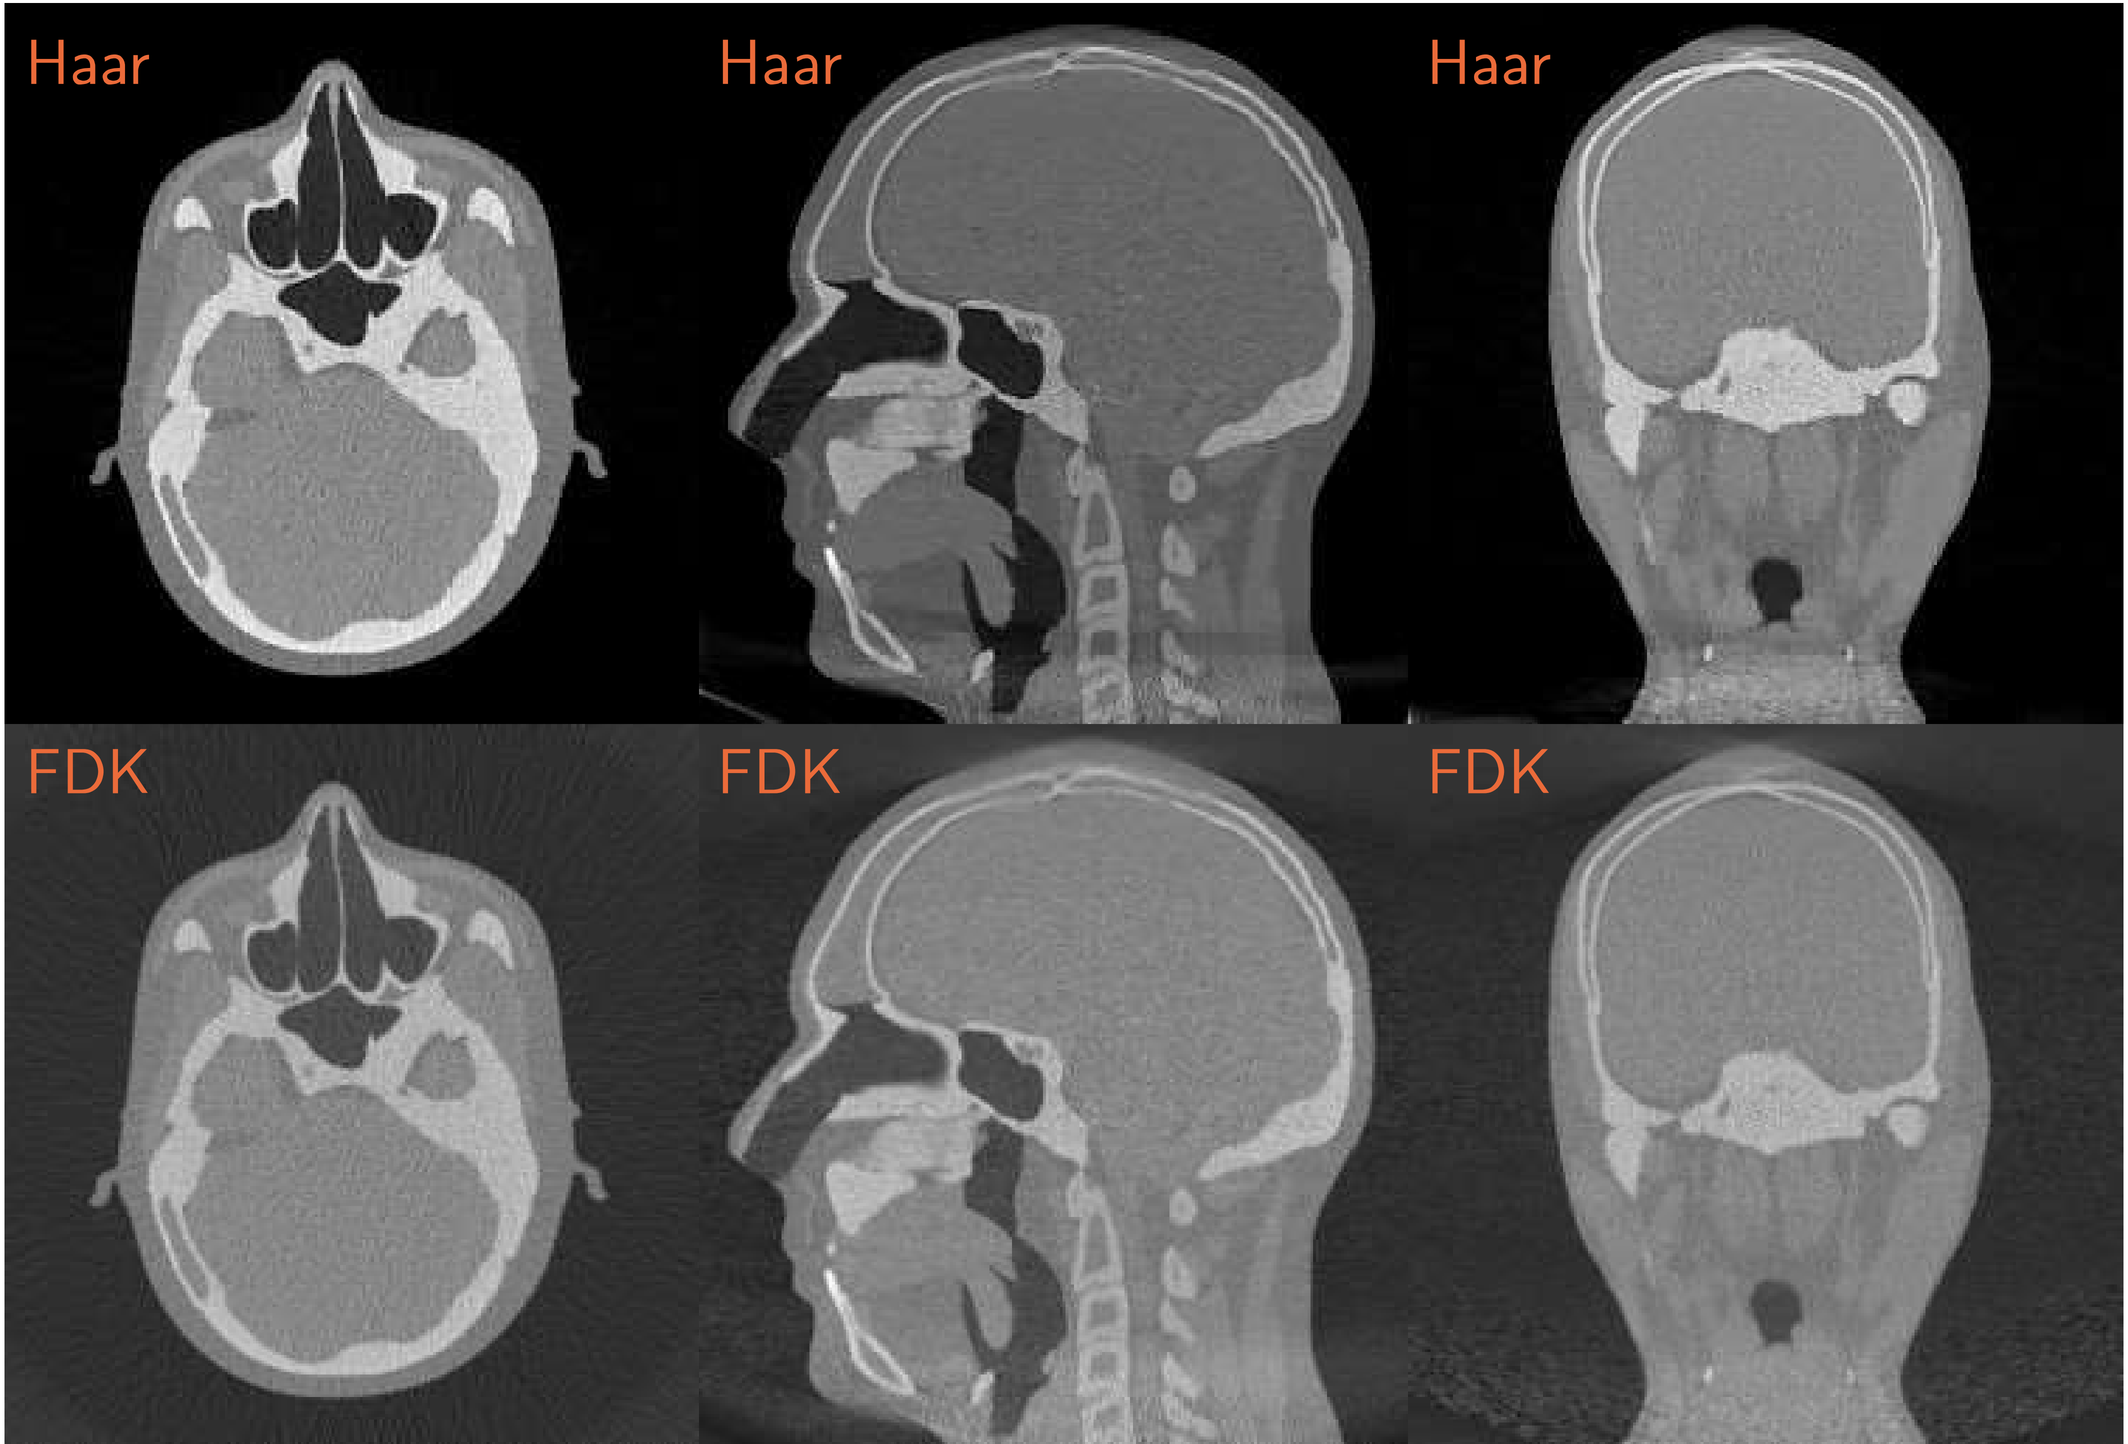

Haar-CT vs. FDK, dose 0.5% ($\mu = 2 \cdot 10^{-5}$)

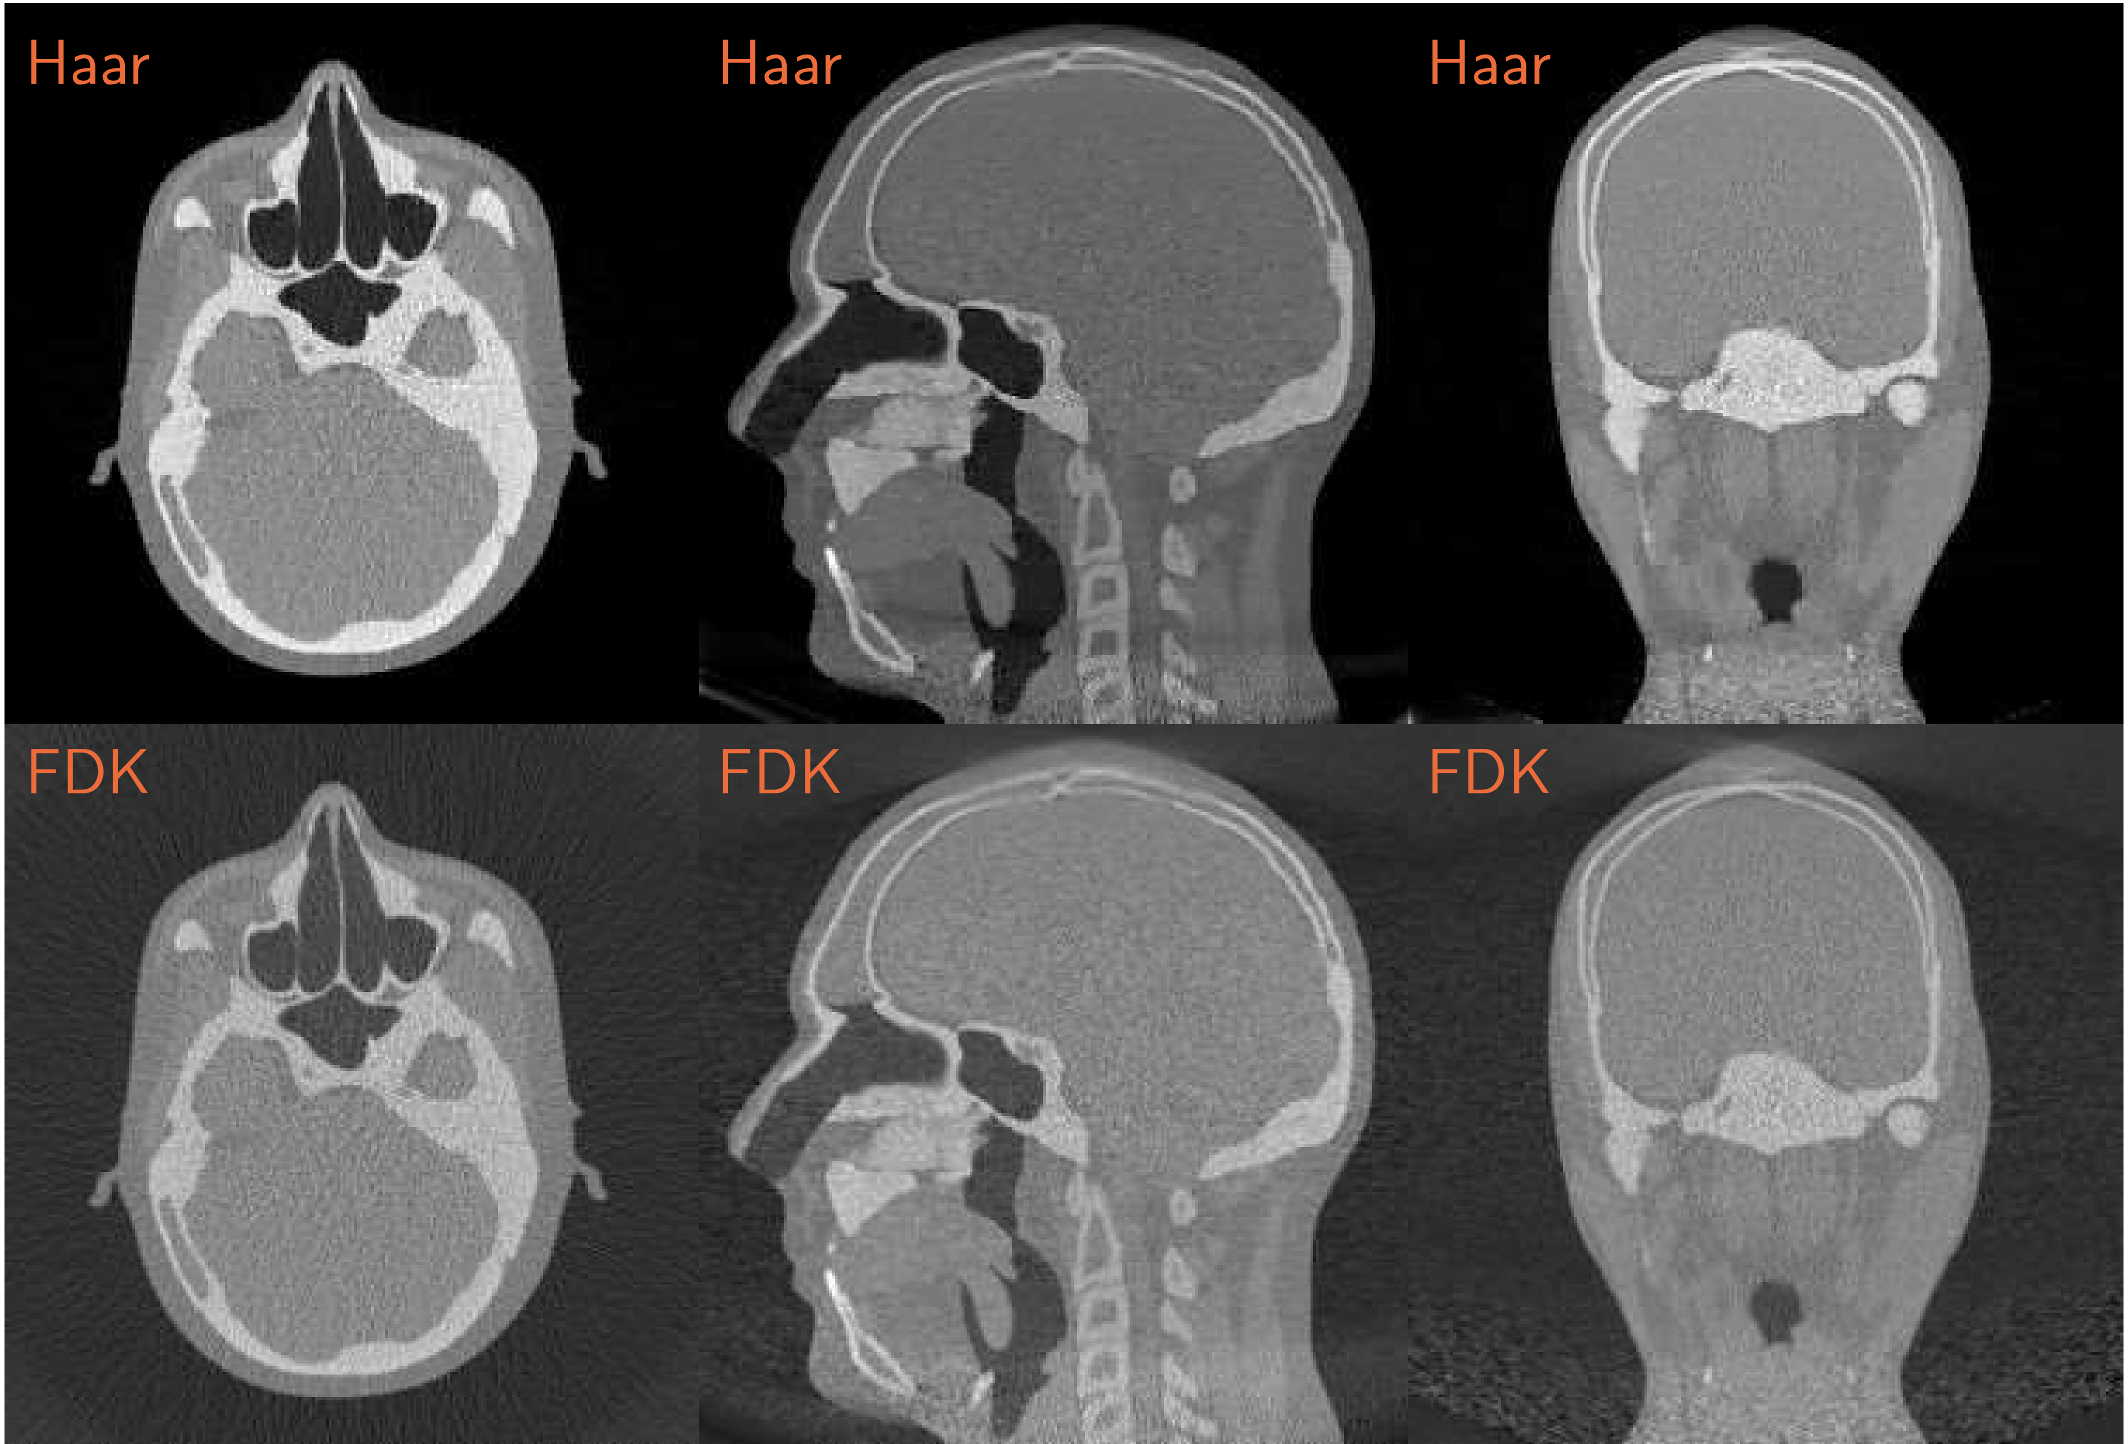

Haar-CT vs. FDK, dose 0.2% ($\mu = 5 \cdot 10^{-5}$)

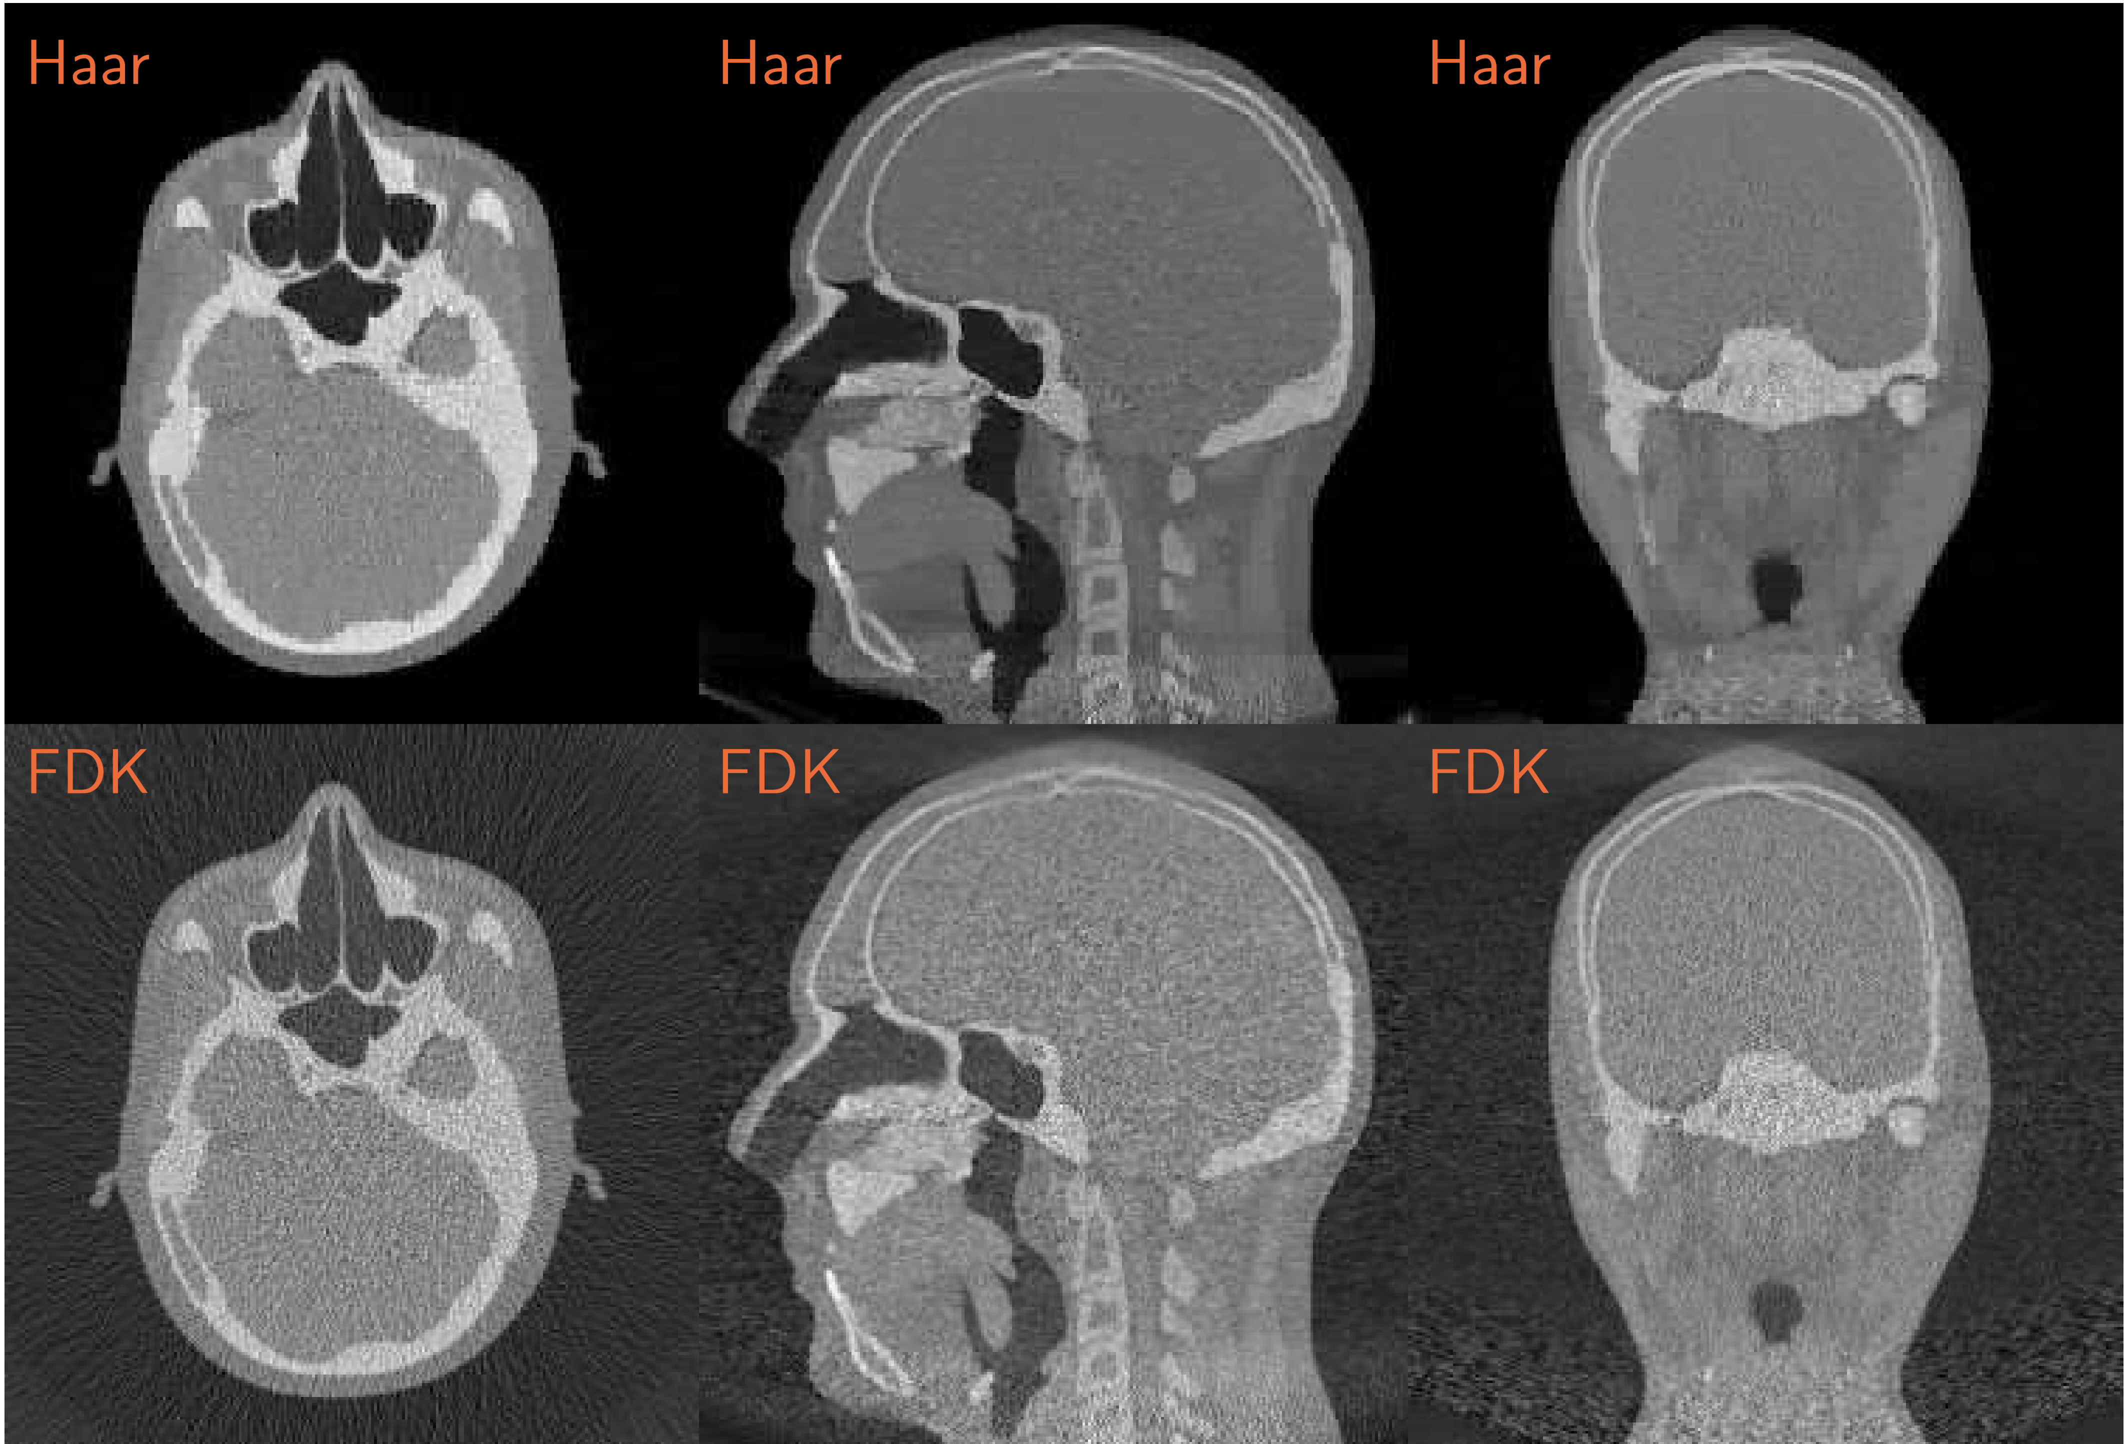

Haar-CT vs. FDK, dose 0.1% ($\mu = 5 \cdot 10^{-5}$)

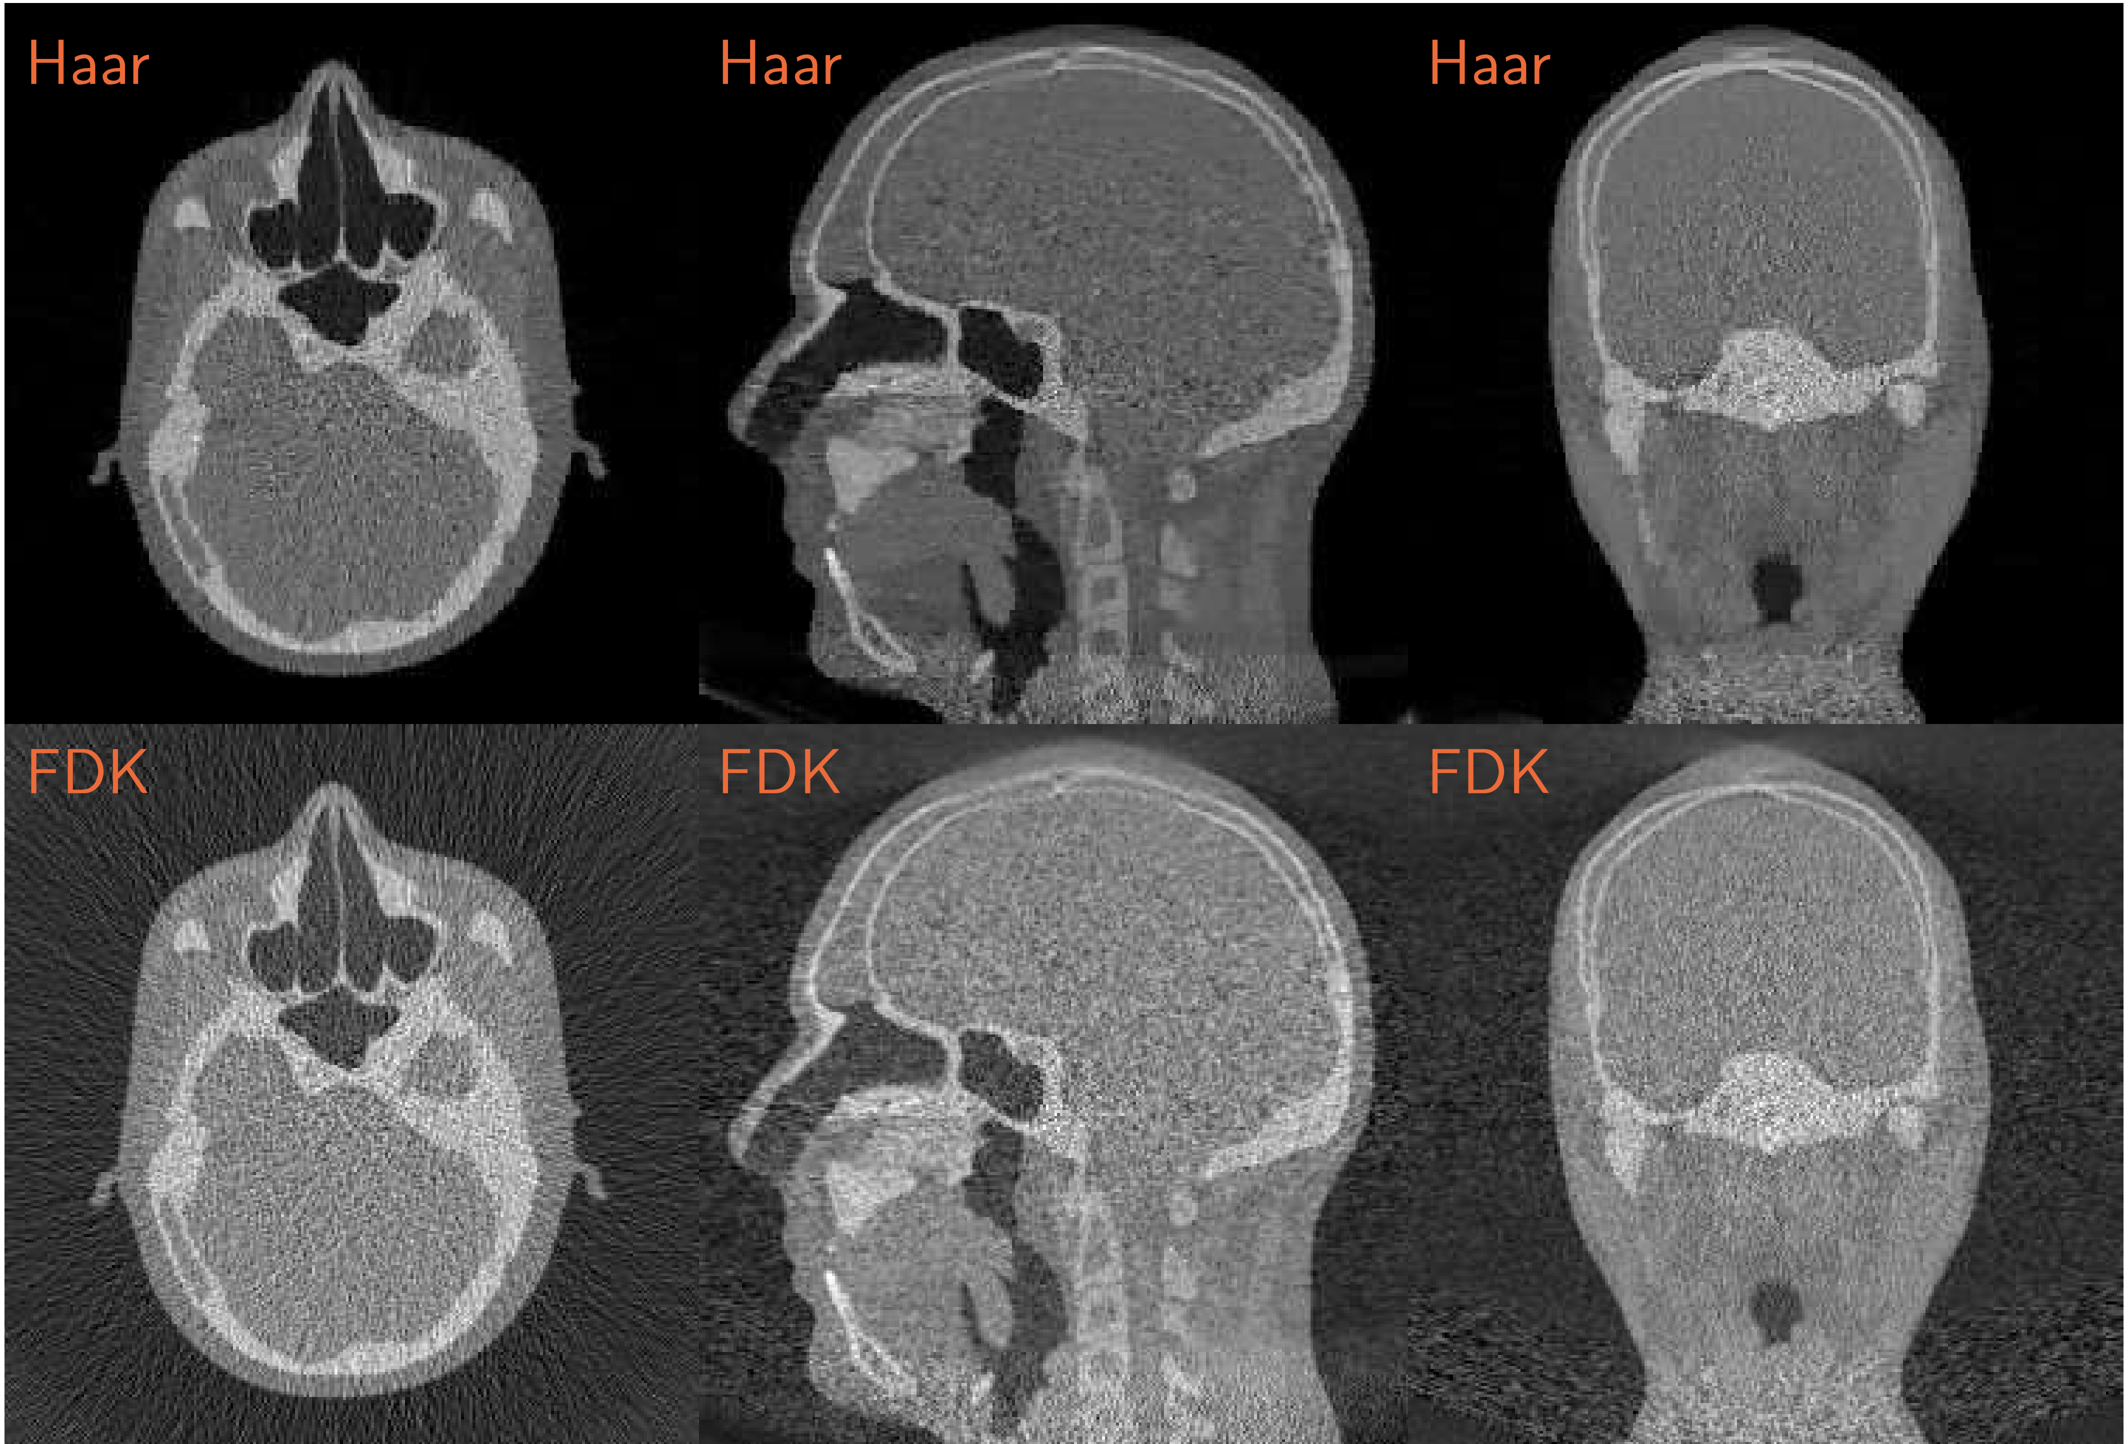

Reconstructions: Anisotropic Total Variation

What does anisotropic total variation look like?

How does choice of μ affect the reconstruction?

Anisotropic TV, dose 100%, μ too large ($5 \cdot 10^{-5}$)

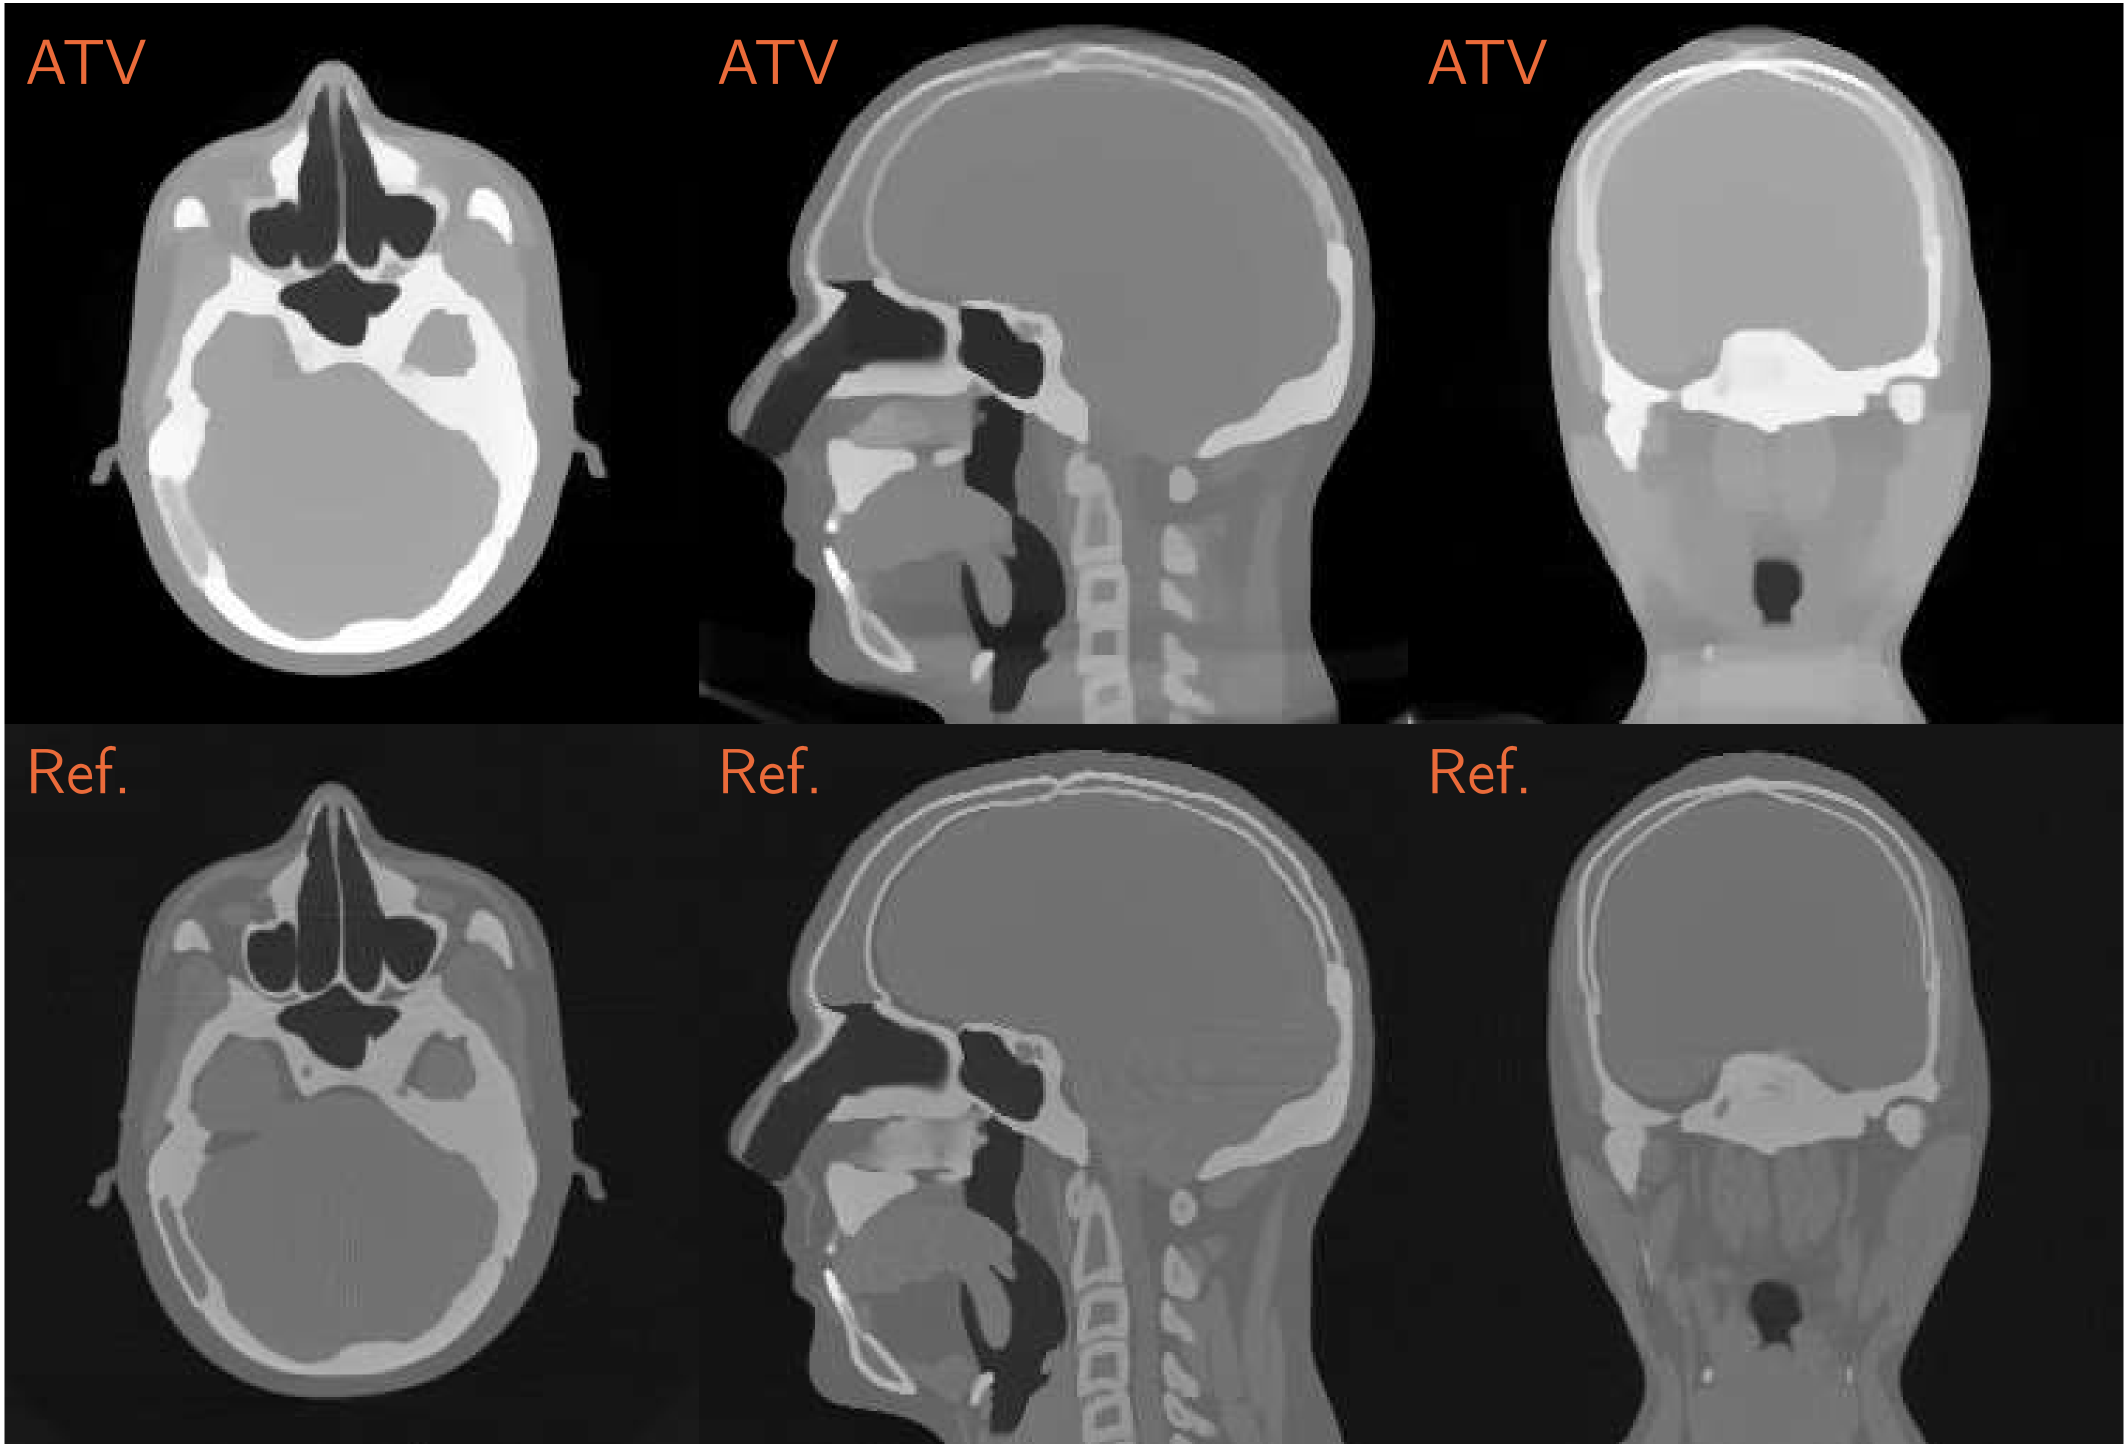

Anisotropic TV, dose 100%, μ suitable $(5 \cdot 10^{-7})$

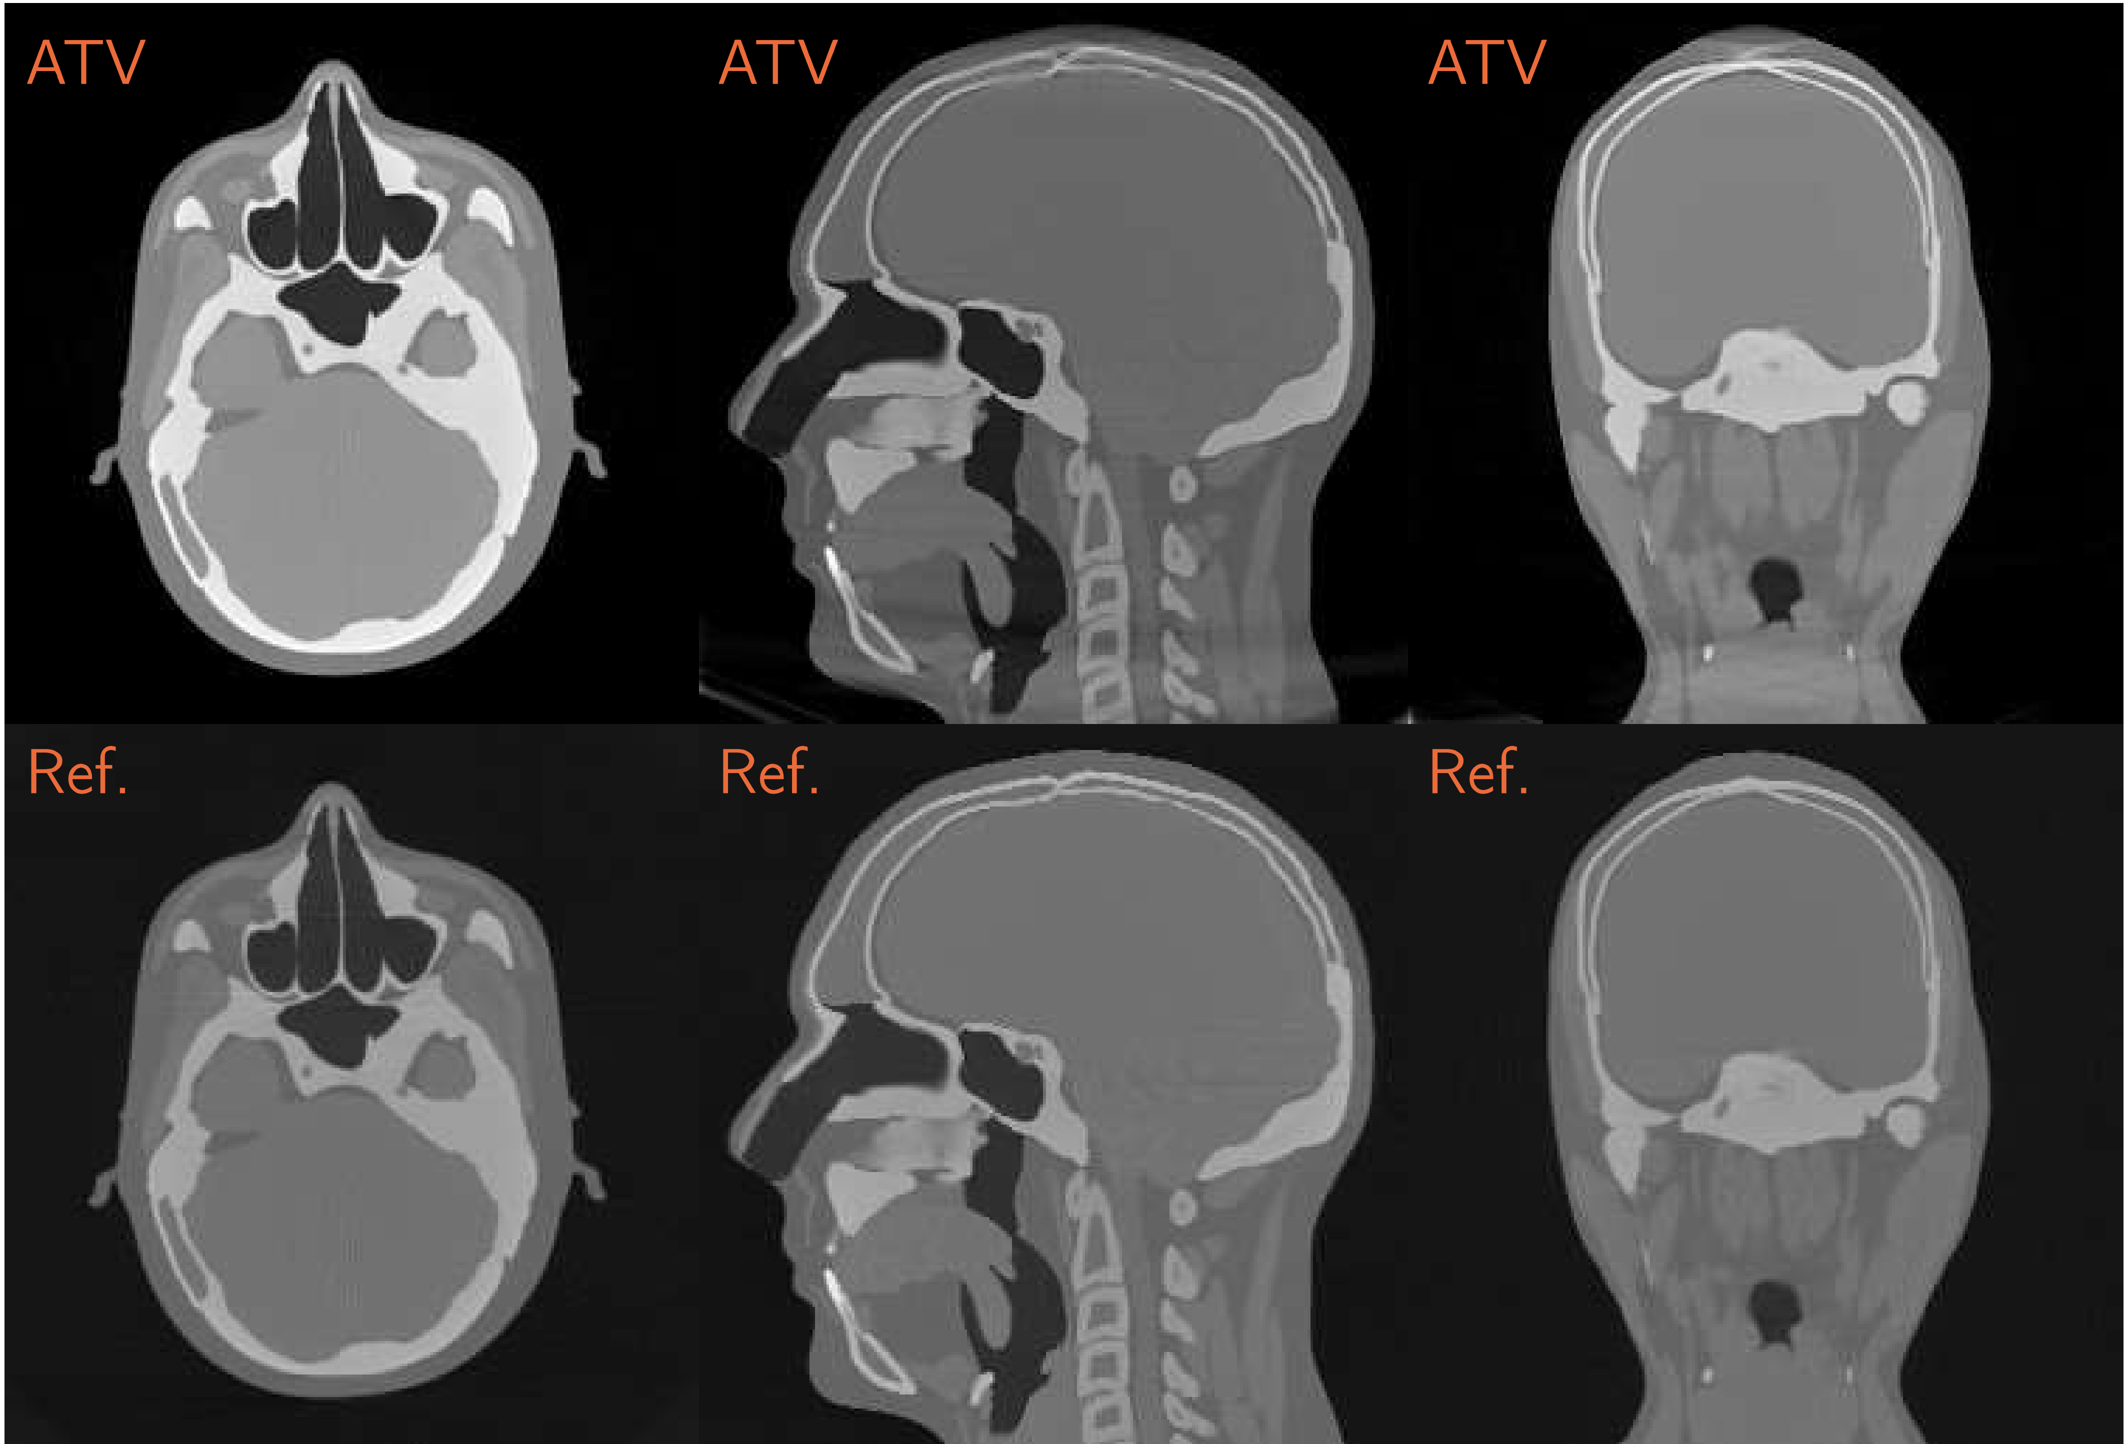

Anisotropic TV, dose 100%, μ too small $(1 \cdot 10^{-7})$

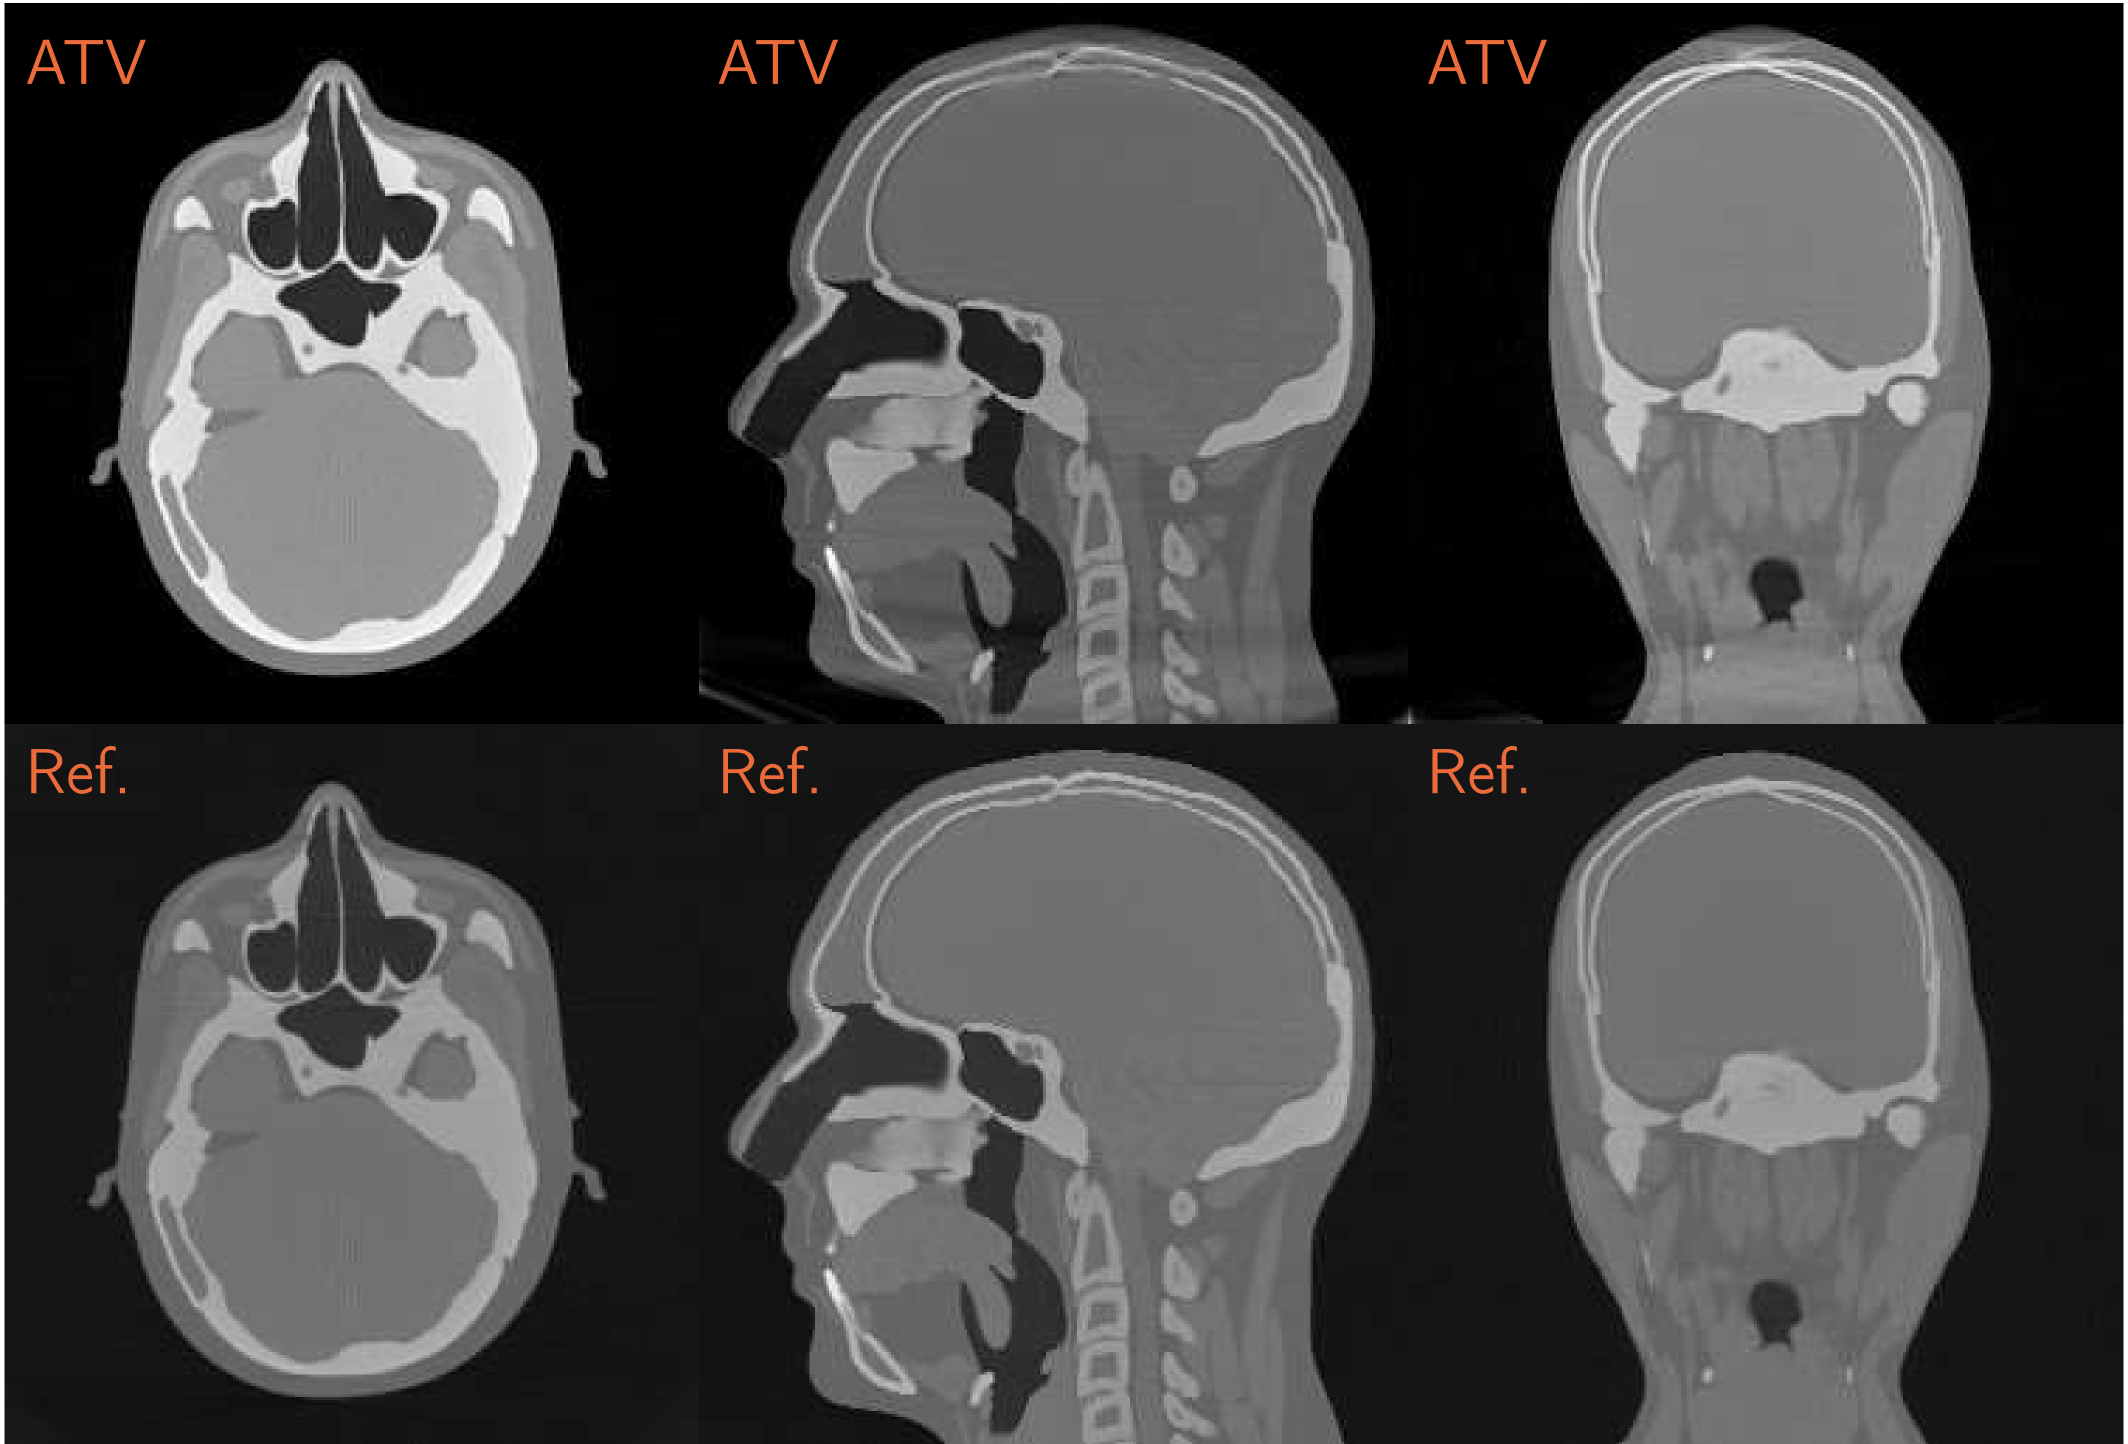

Anisotropic TV, dose 10%, μ too large ($5 \cdot 10^{-5}$)

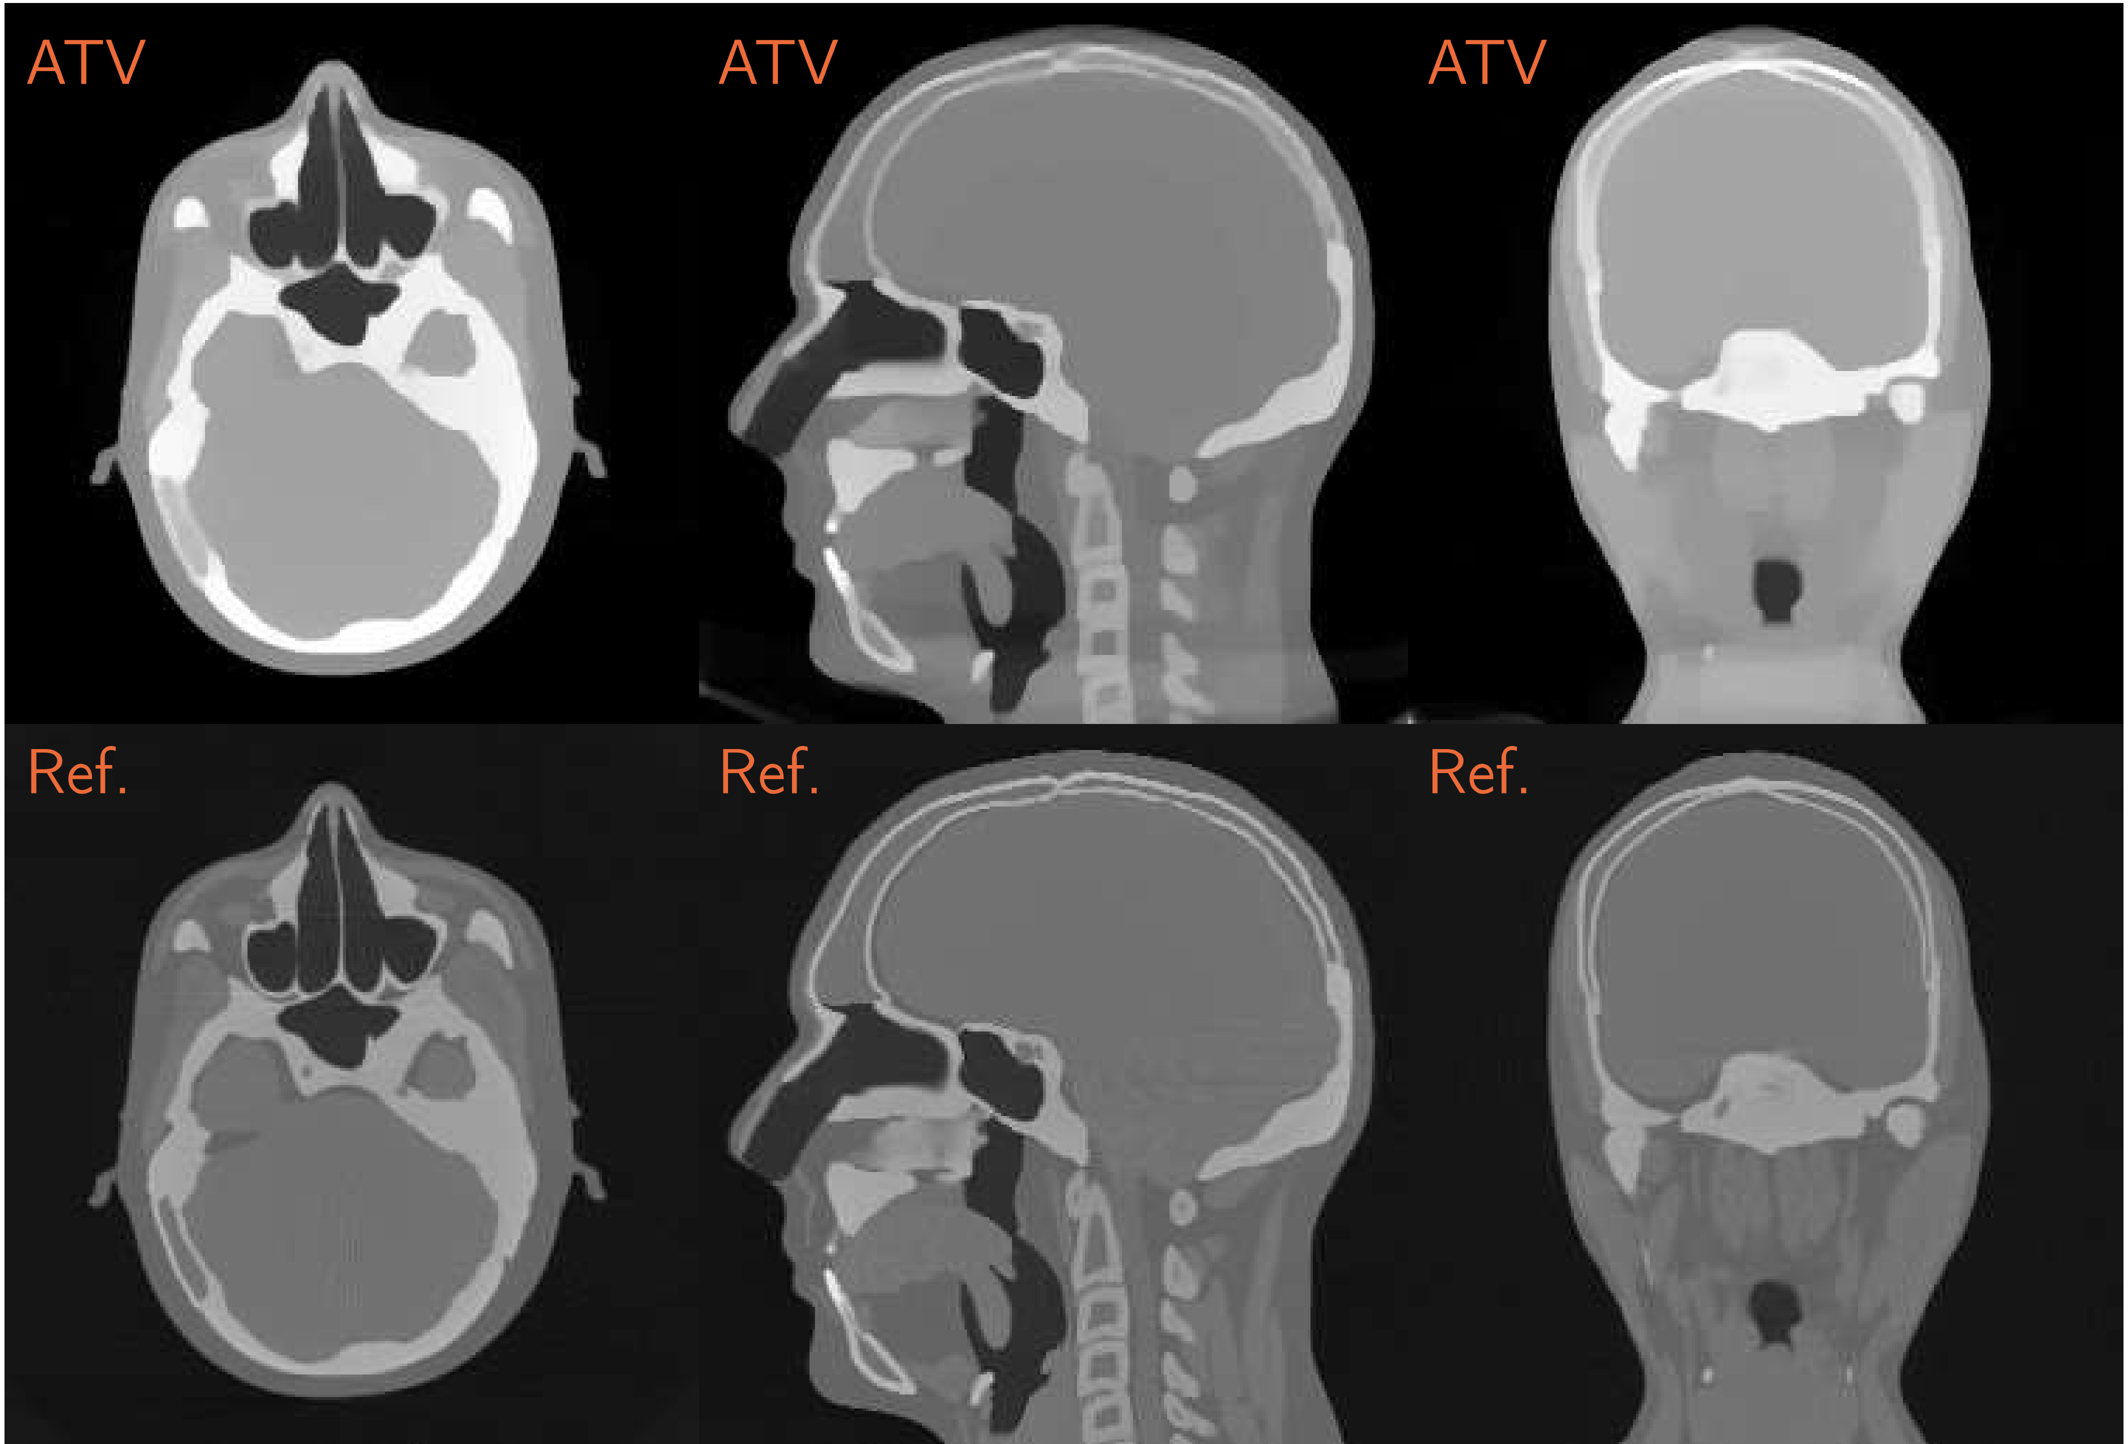

Anisotropic TV, dose 10%, μ suitable $(1 \cdot 10^{-6})$

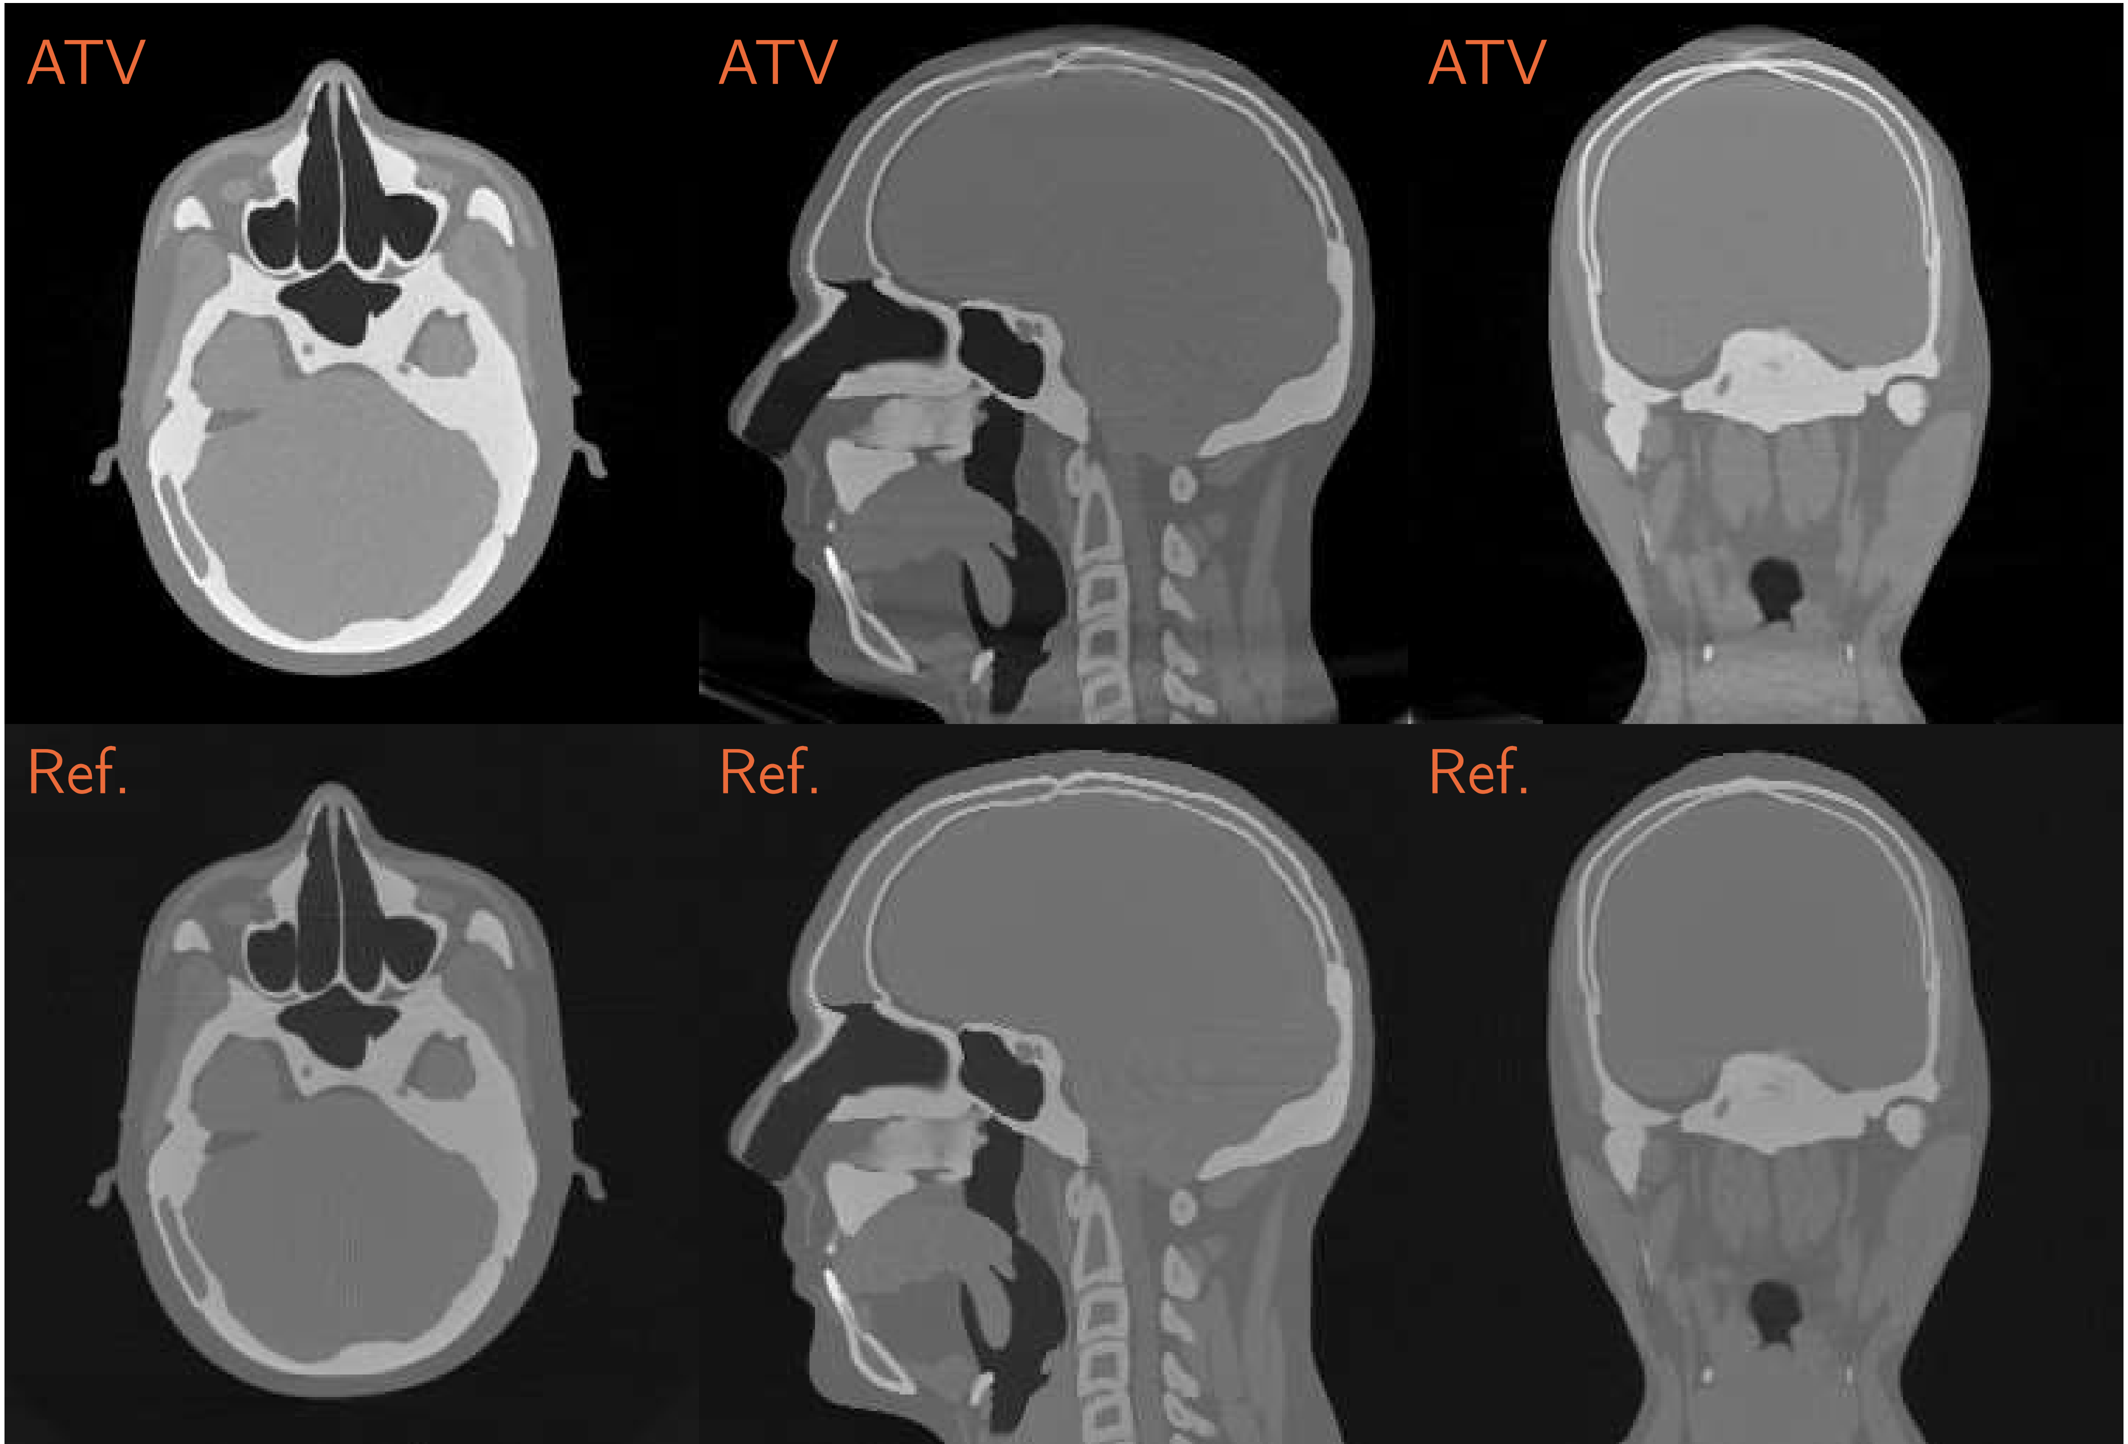

Anisotropic TV, dose 10%, μ too small $(1 \cdot 10^{-7})$

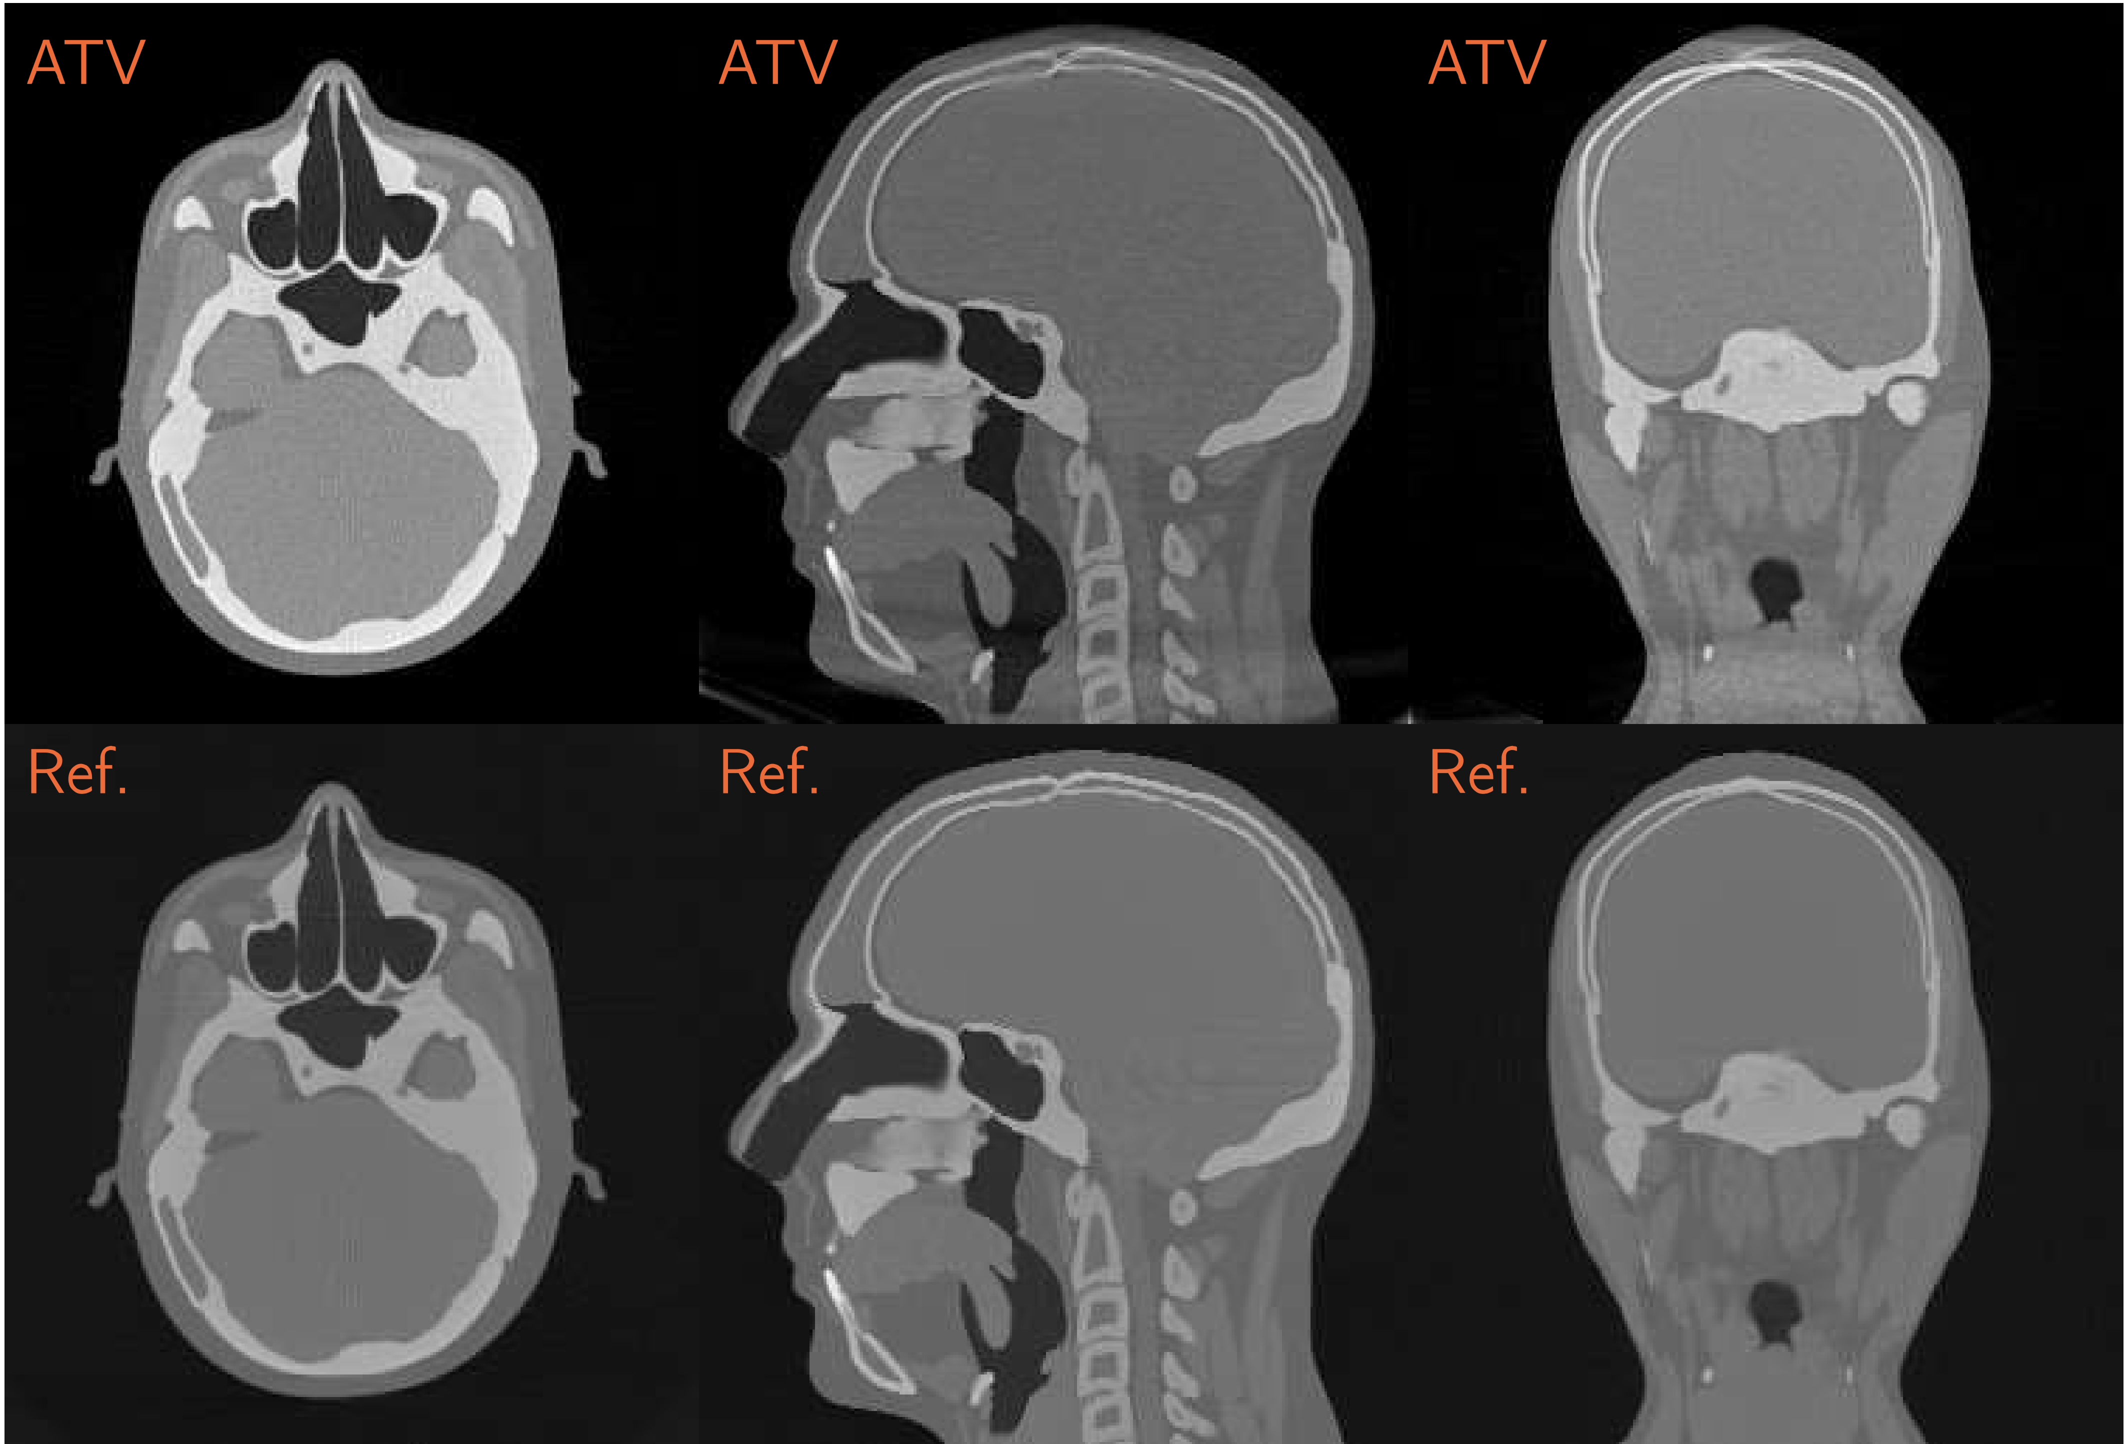

Anisotropic TV, dose 1%, μ too large ($5 \cdot 10^{-5}$)

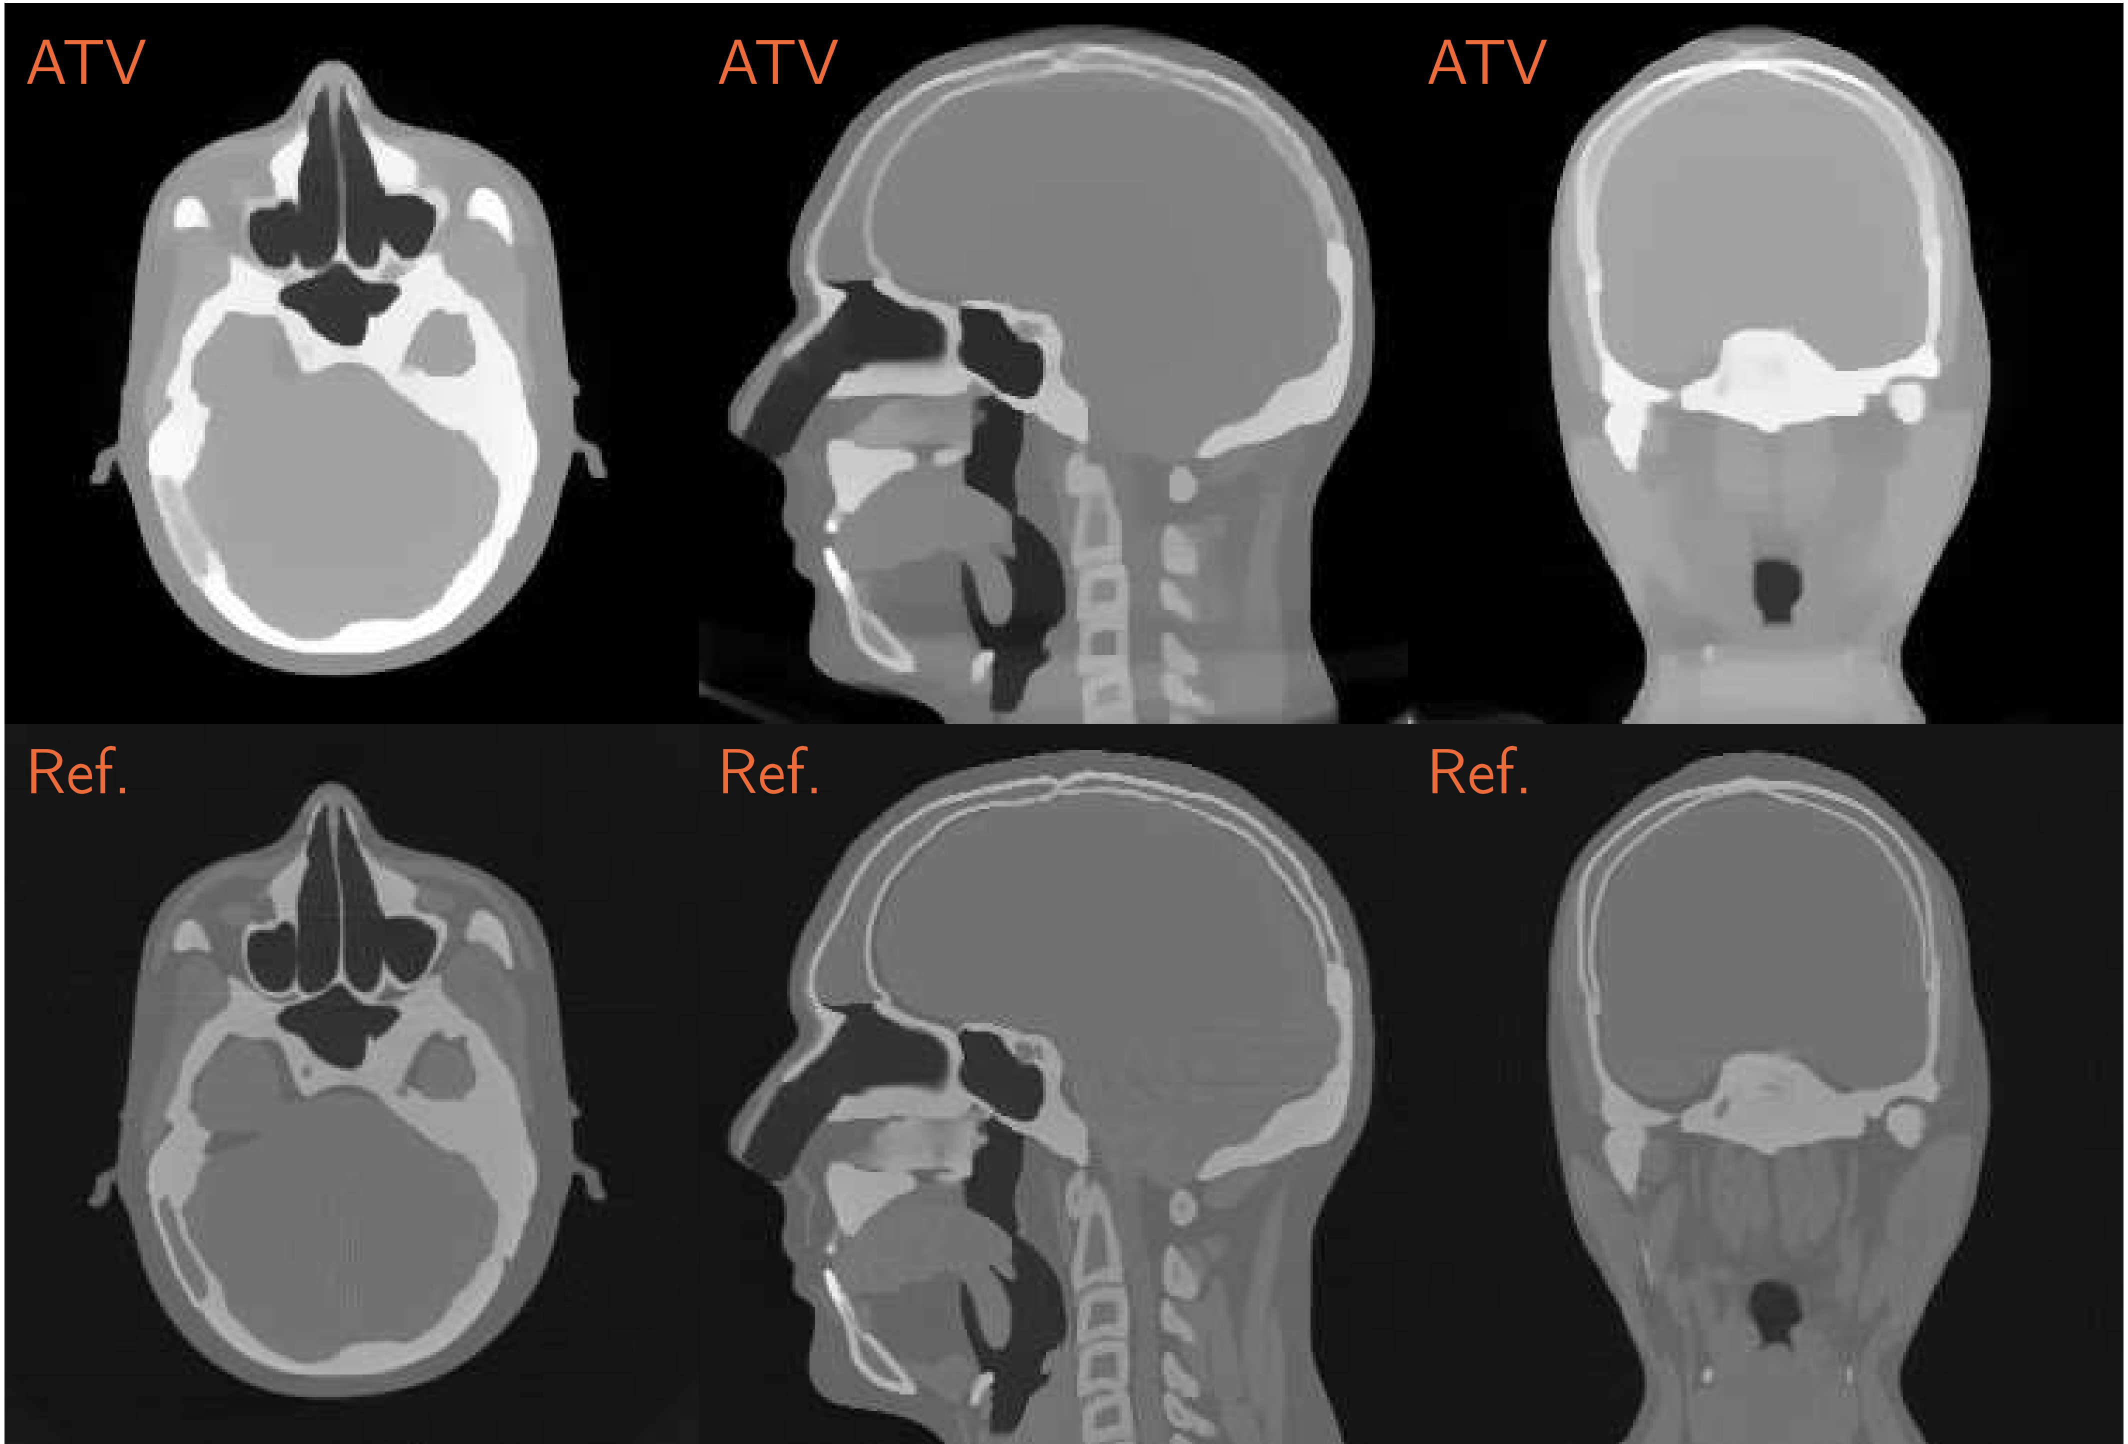

Anisotropic TV, dose 1%, μ suitable $(7.5 \cdot 10^{-6})$

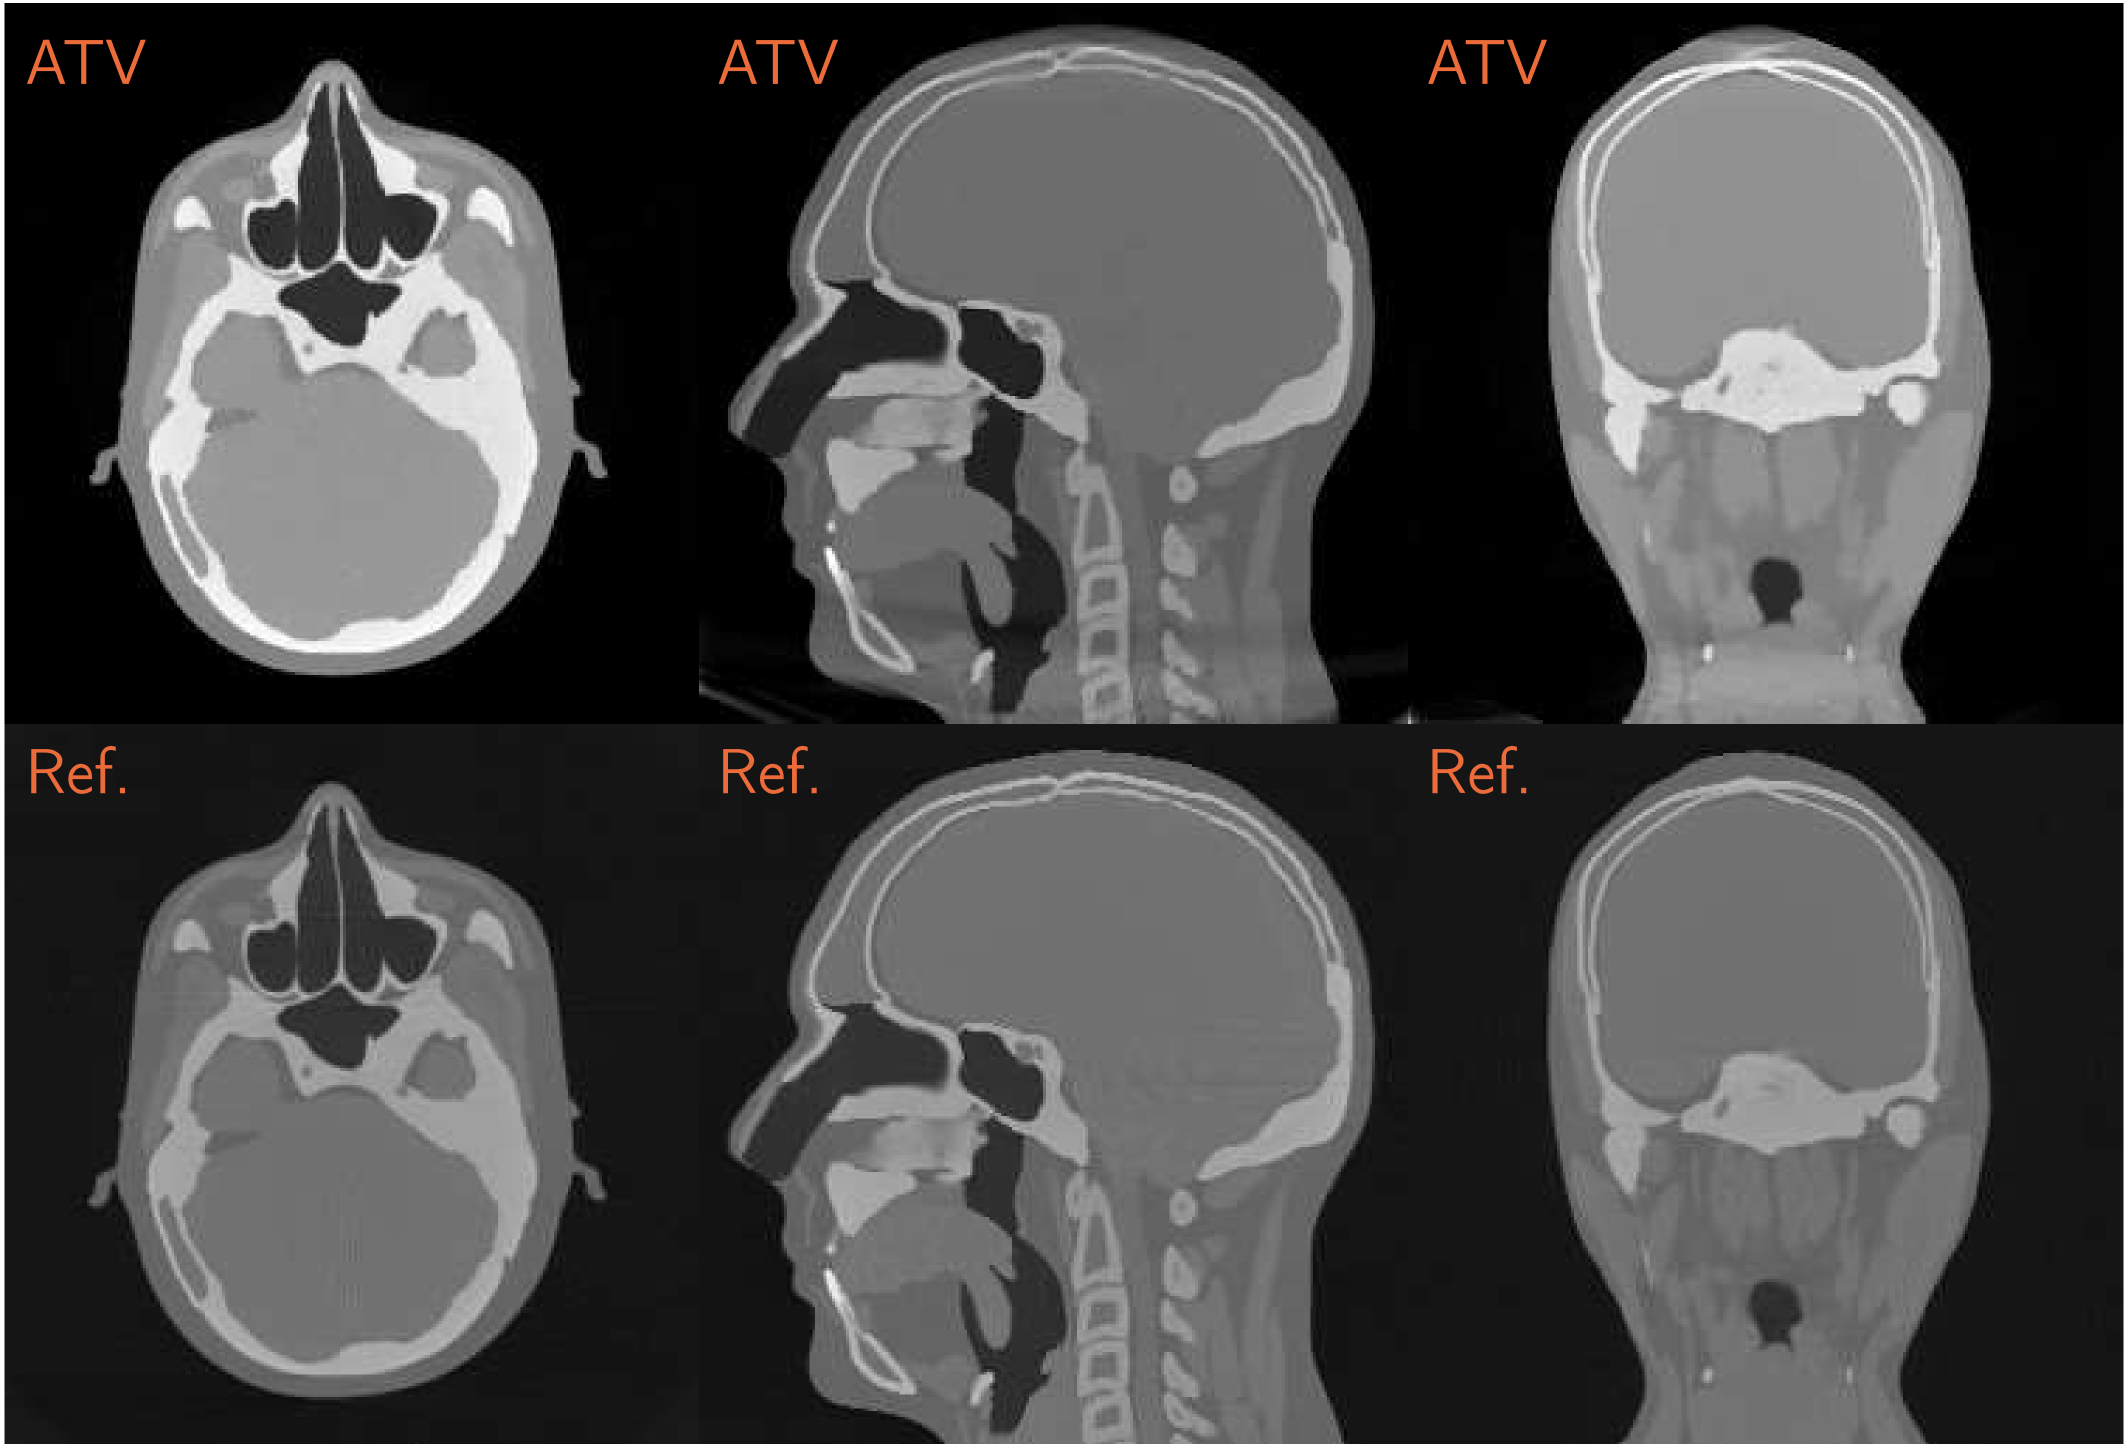

Anisotropic TV, dose 1%, μ too small ($1 \cdot 10^{-6}$)

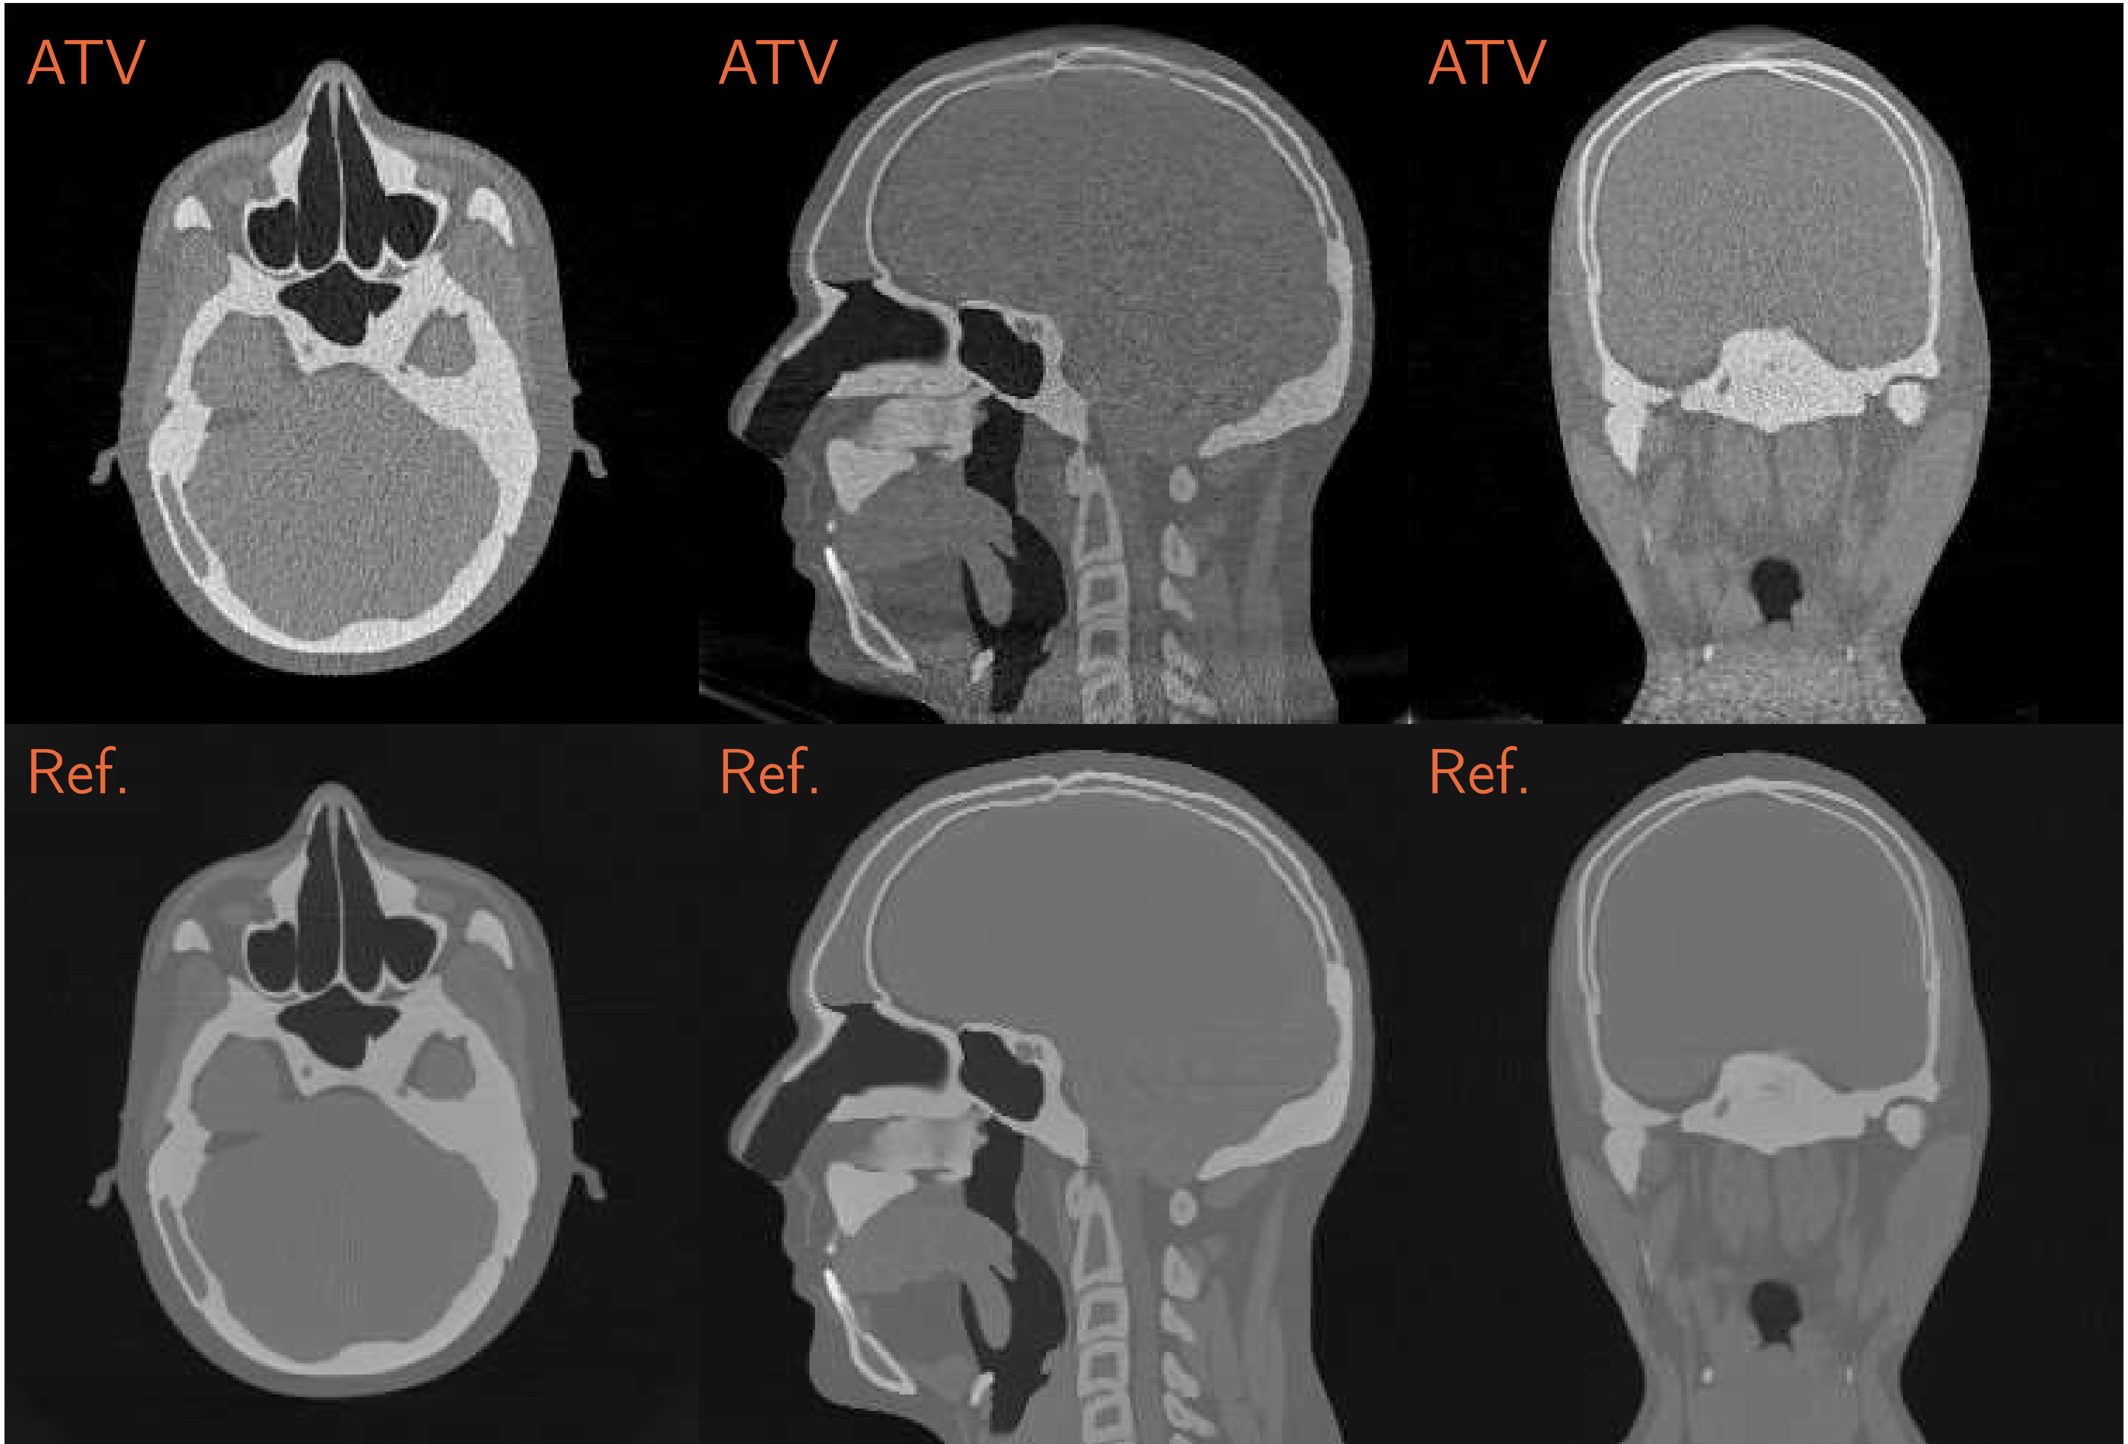

Anisotropic TV, dose 0.1%, μ too large $(1 \cdot 10^{-4})$

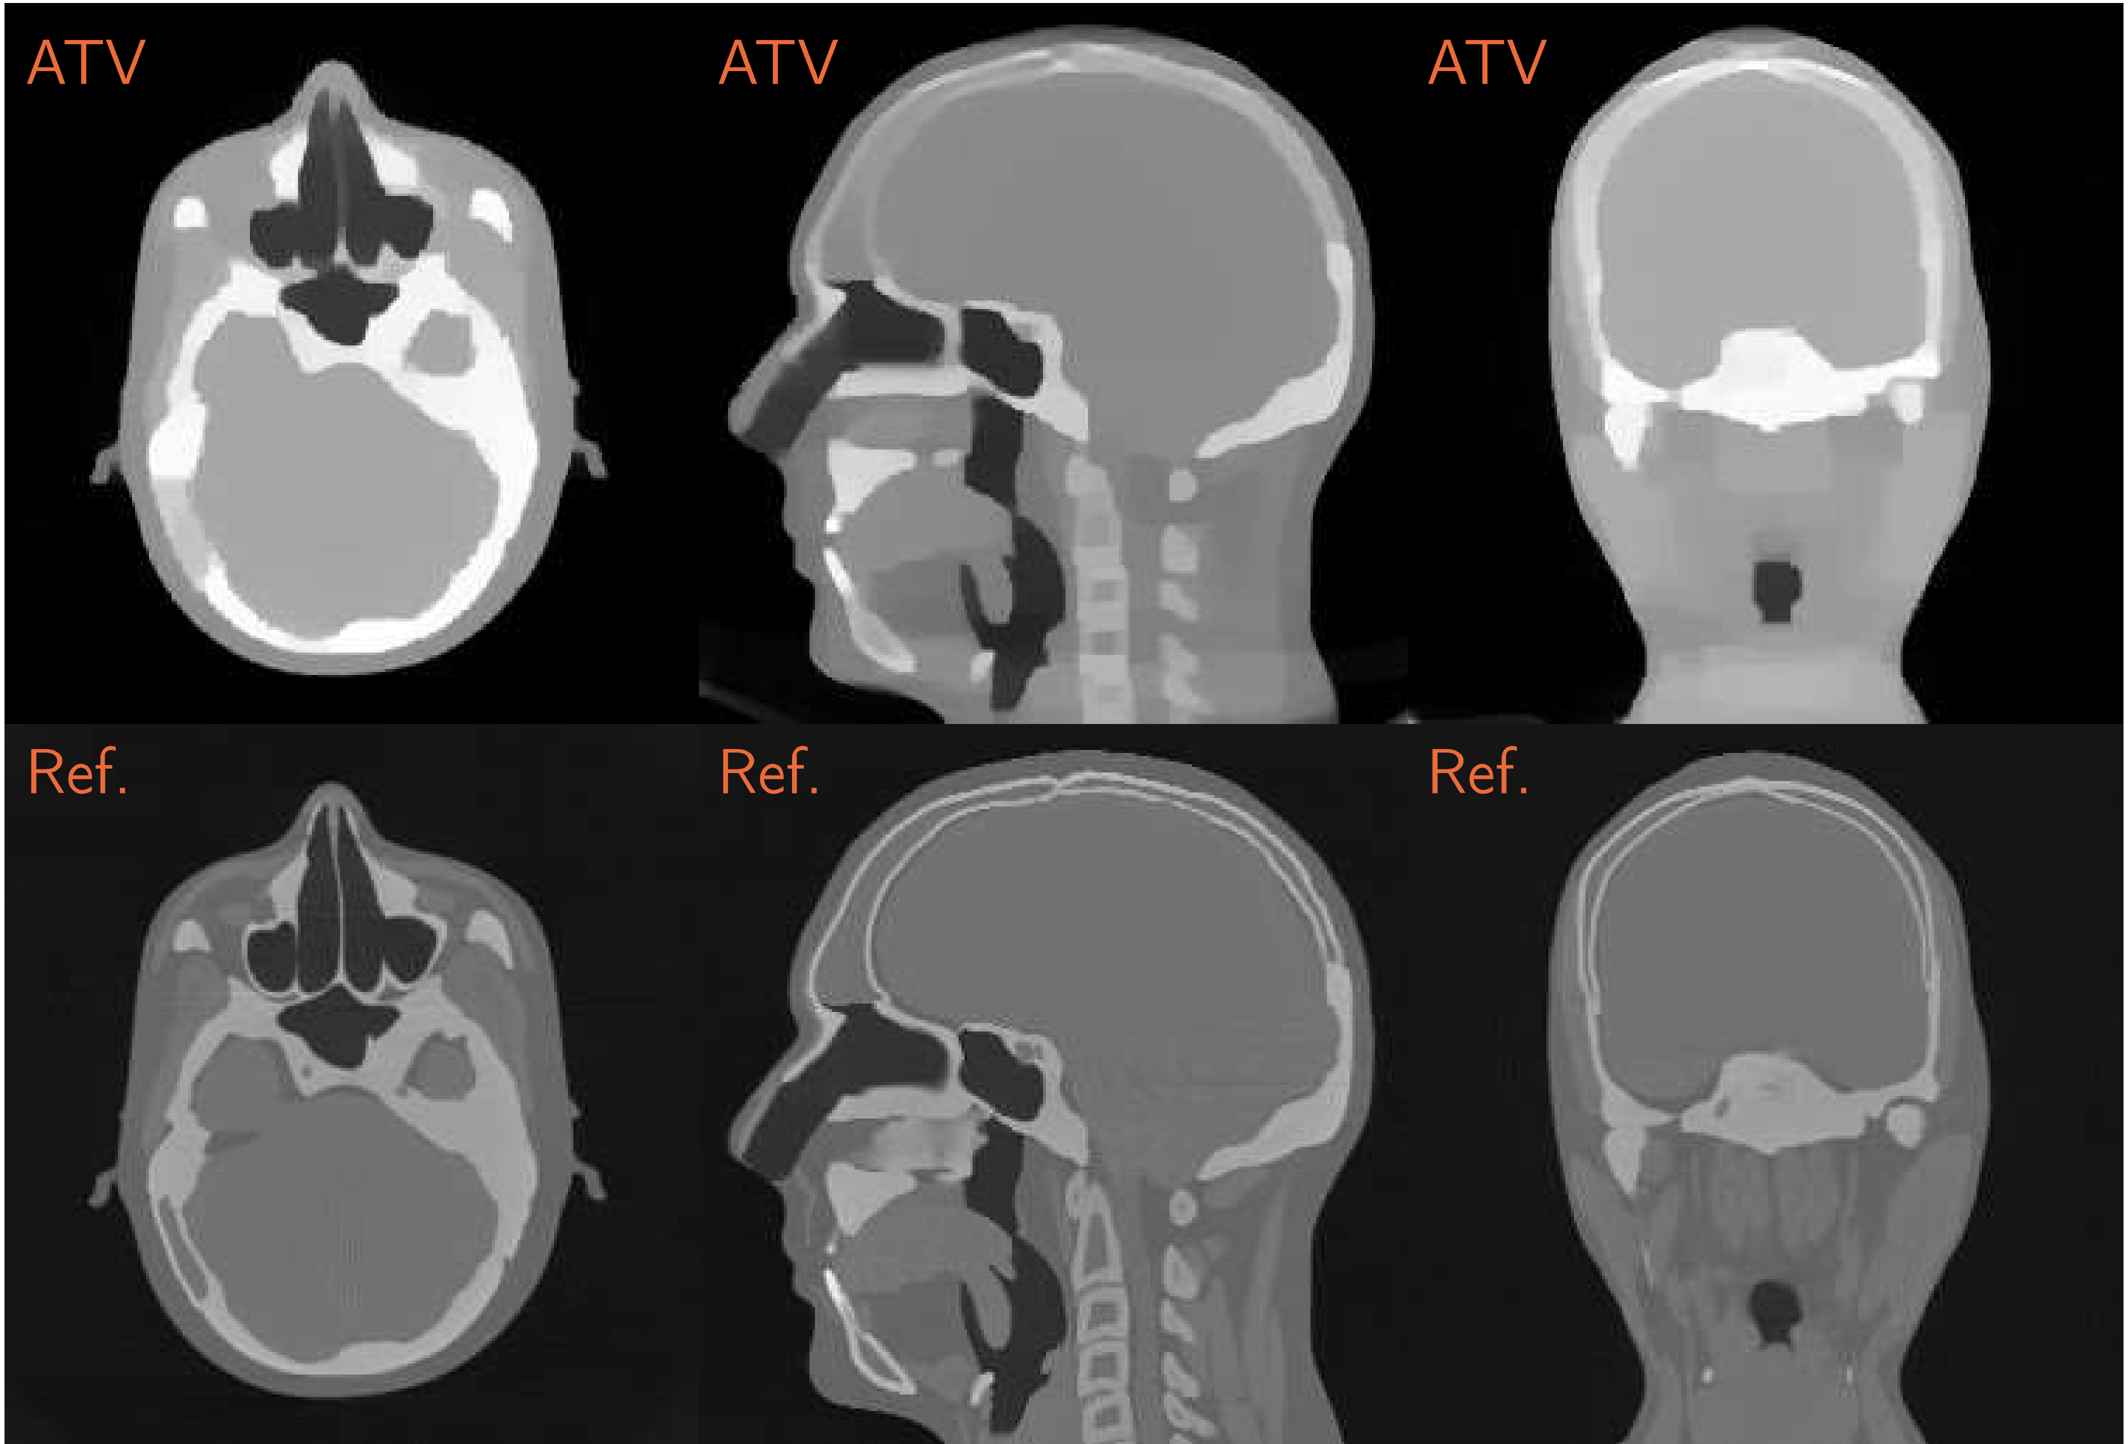

Anisotropic TV, dose 0.1%, μ suitable ($2.8 \cdot 10^{-5}$)

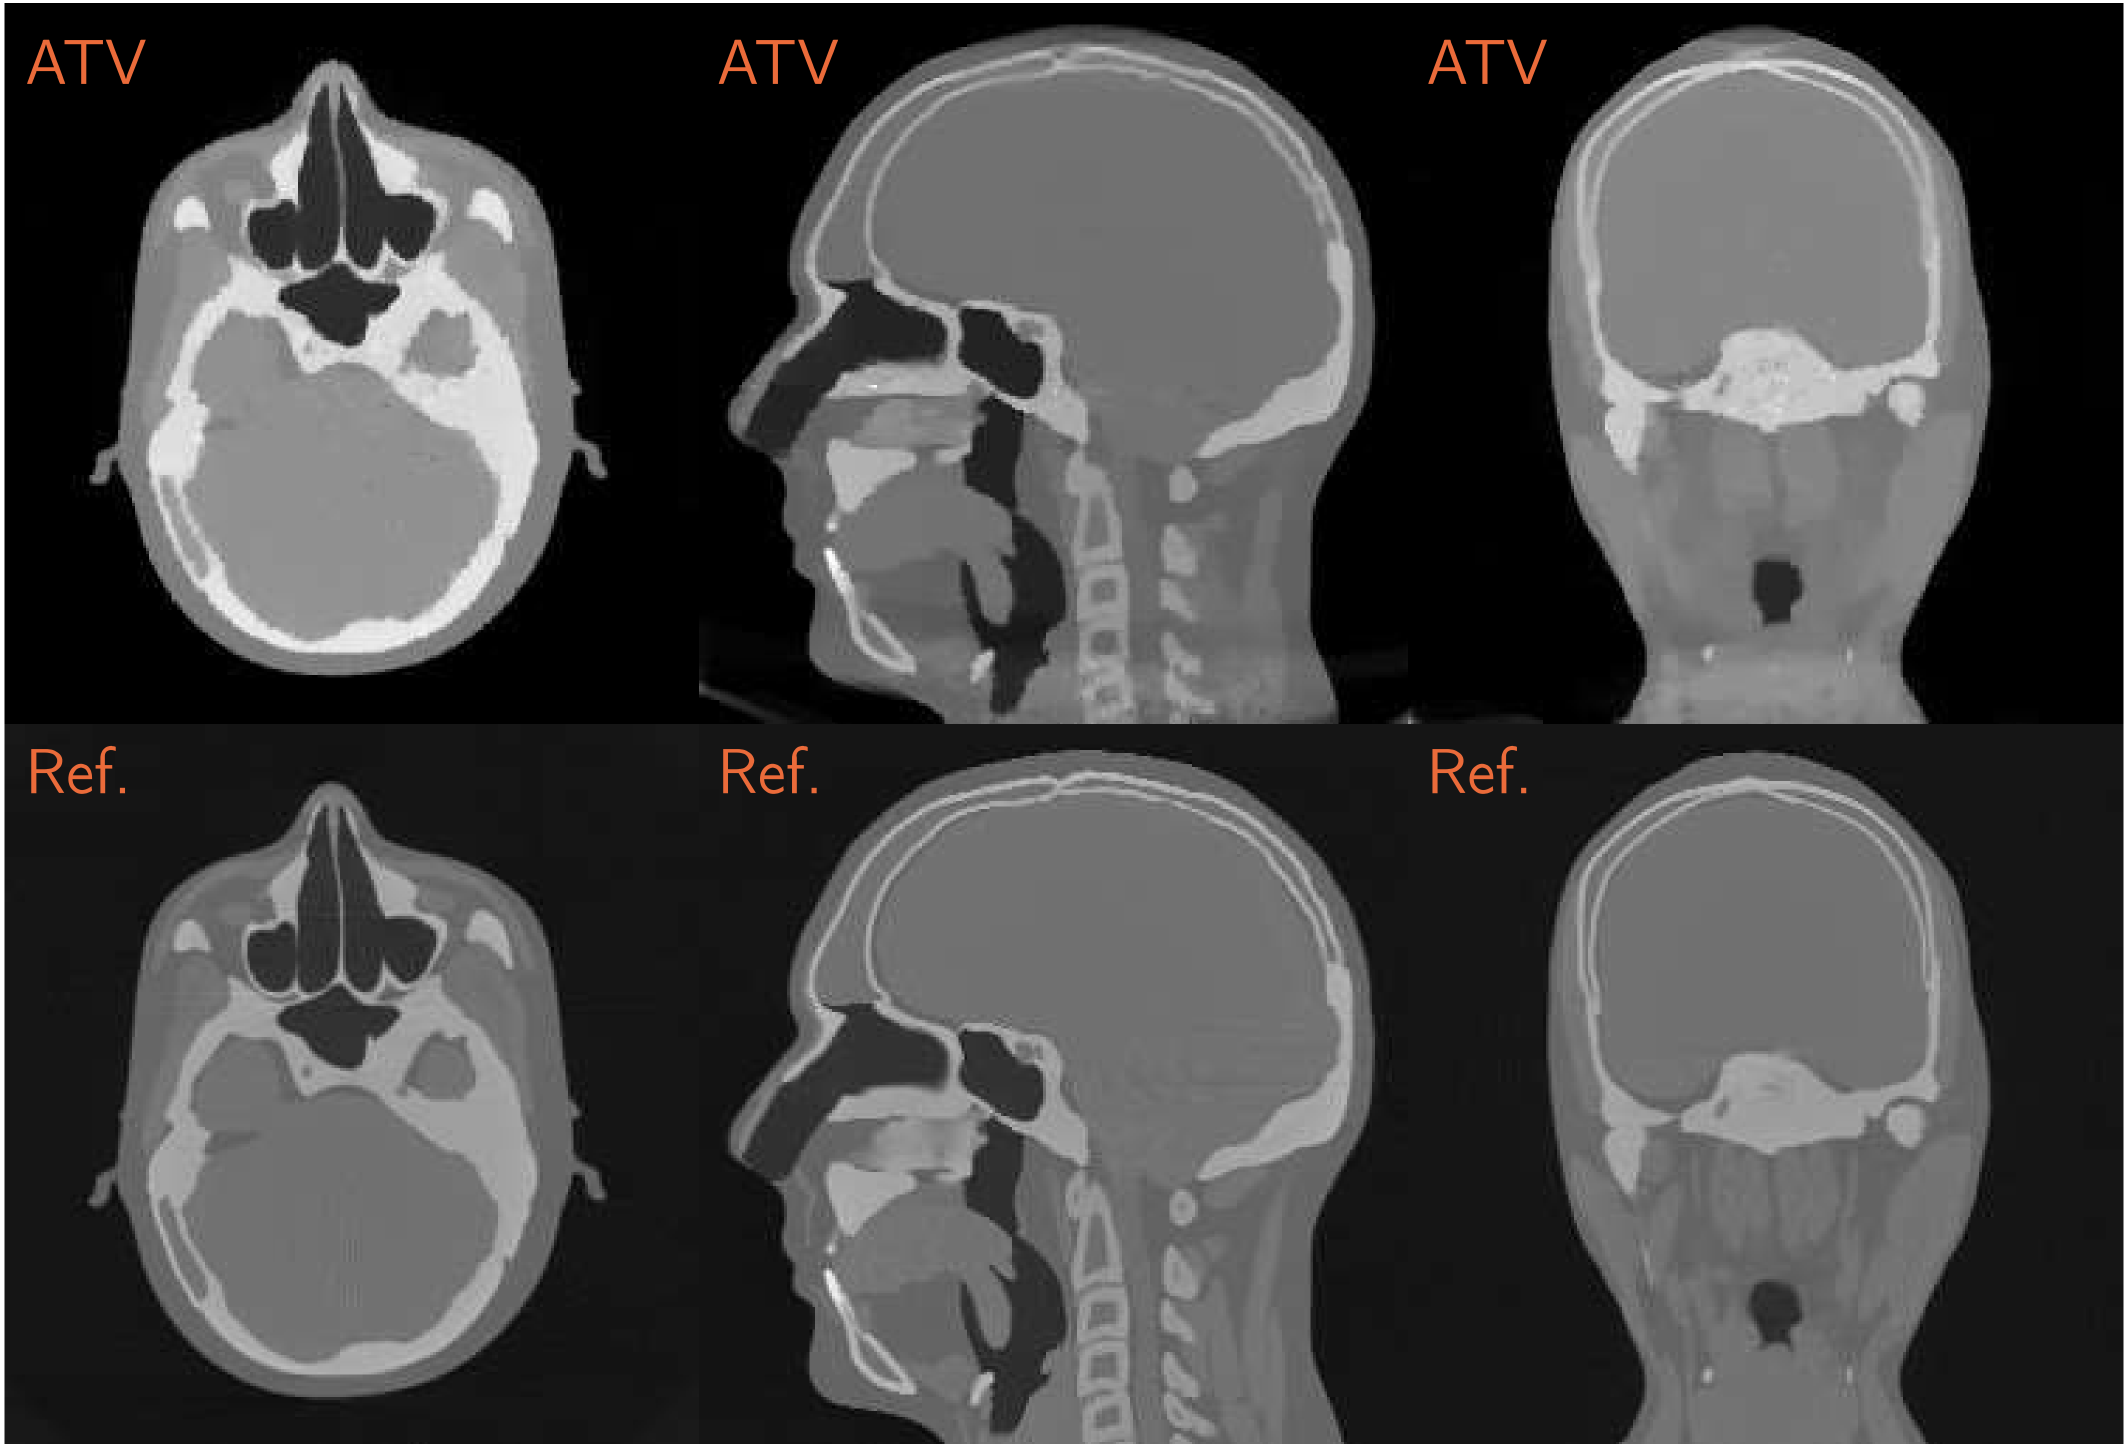

Anisotropic TV, dose 0.1%, μ too small ($7.5 \cdot 10^{-6}$)

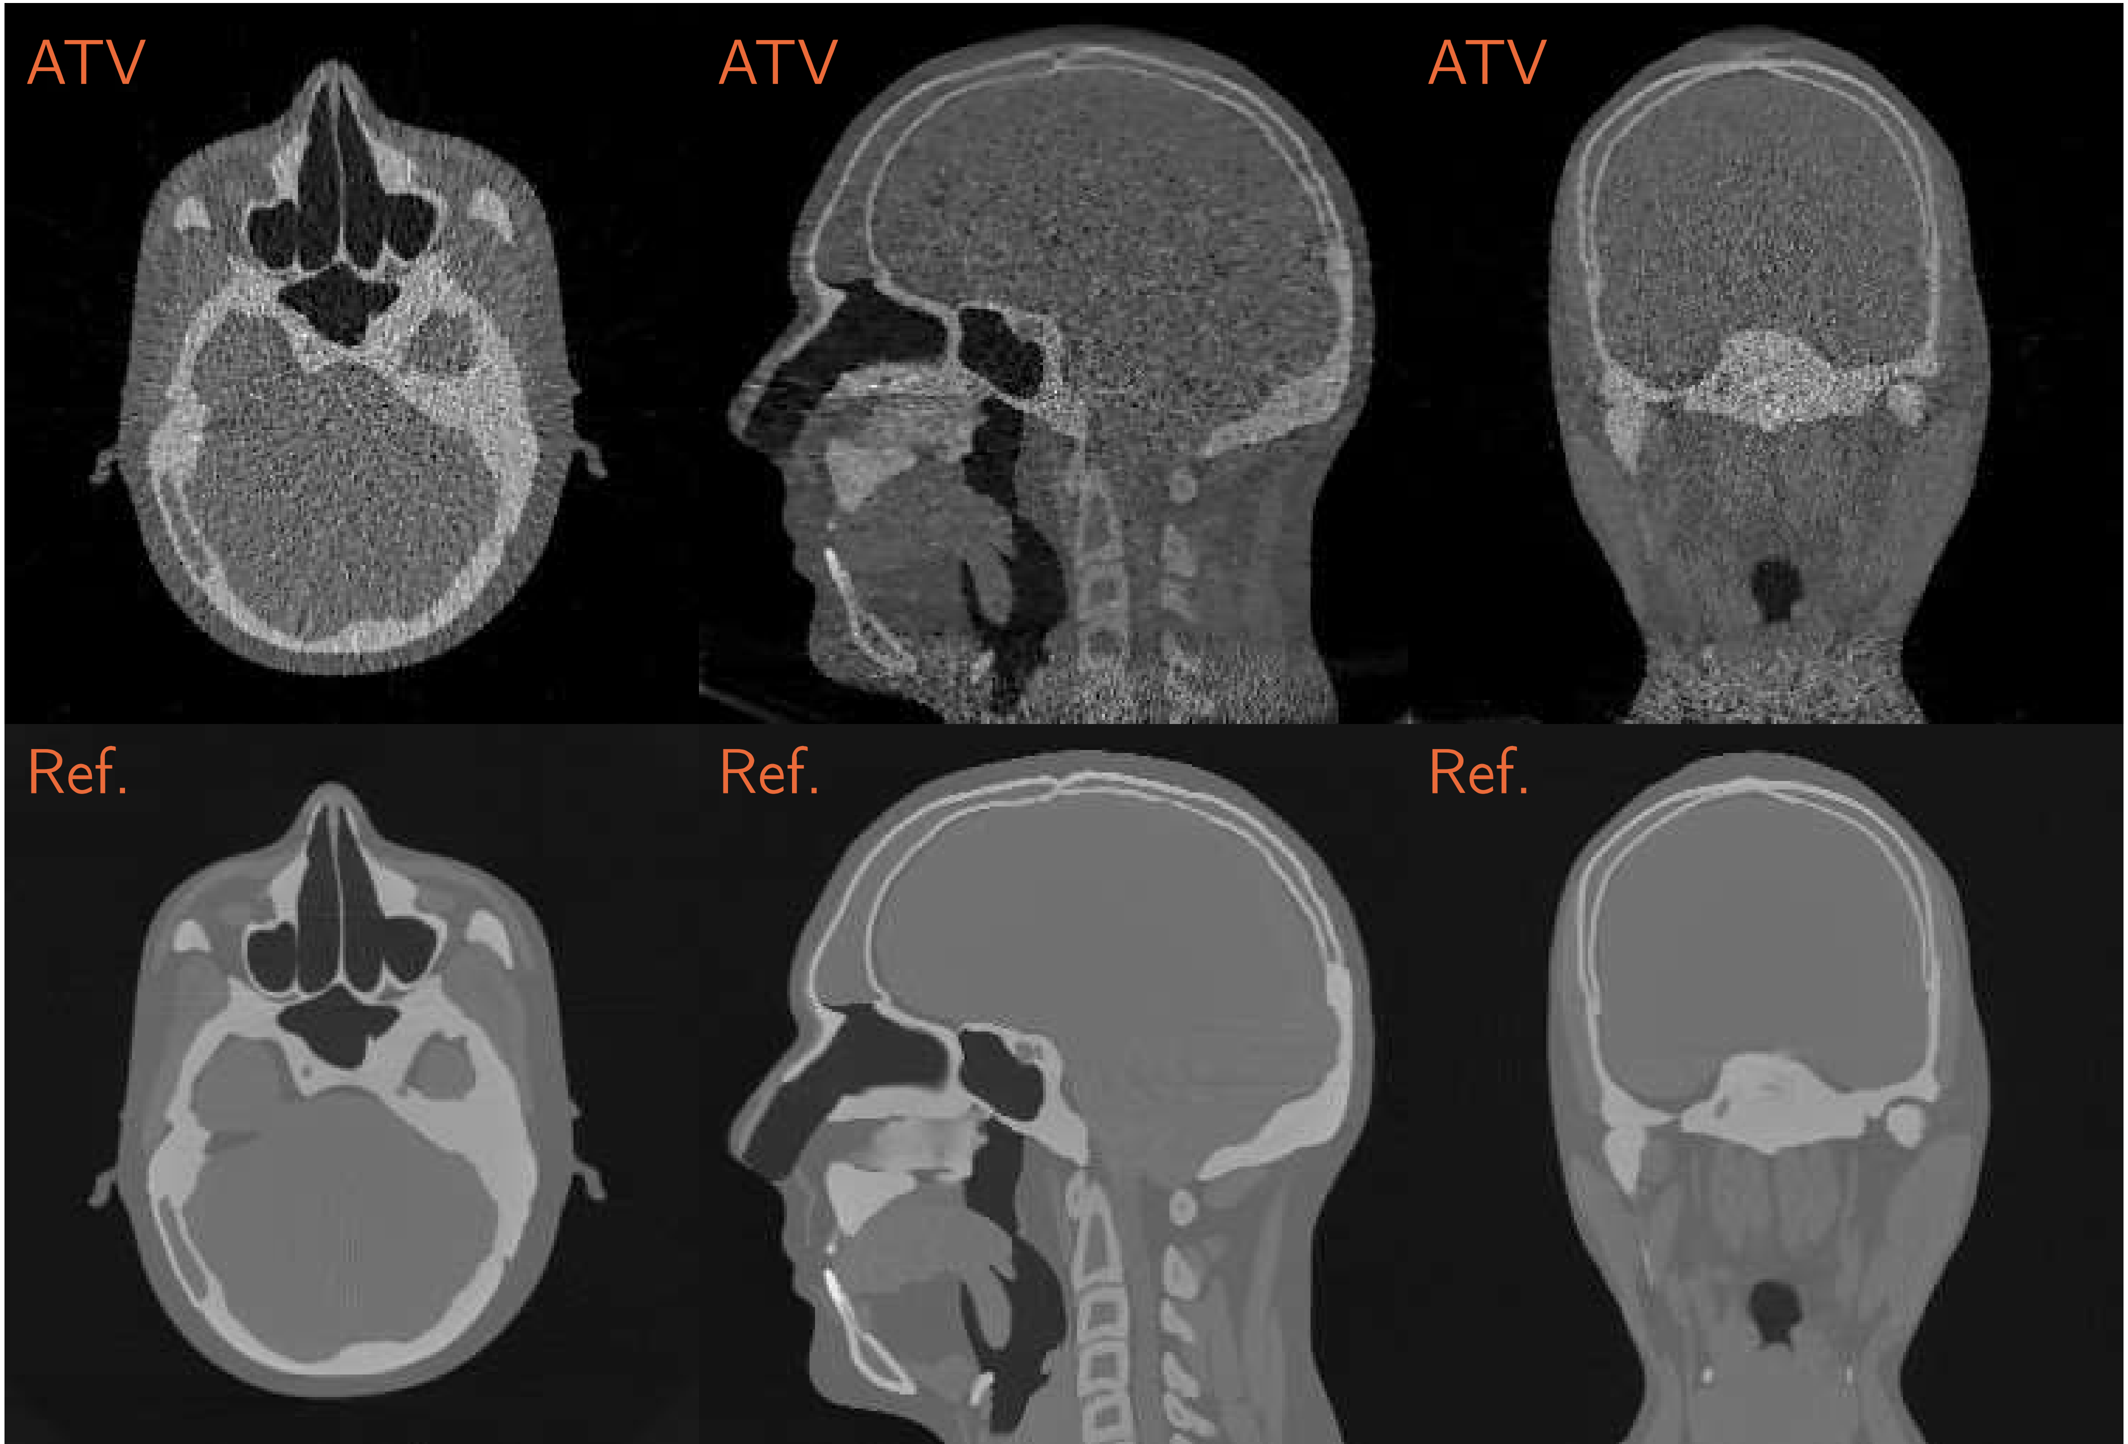

Anisotropic total variation regularization

The true test:

How does anisotropic total variation regularization compare to FDK?

Anisotropic TV vs. FDK, dose 100% ($\mu = 5 \cdot 10^{-7}$)

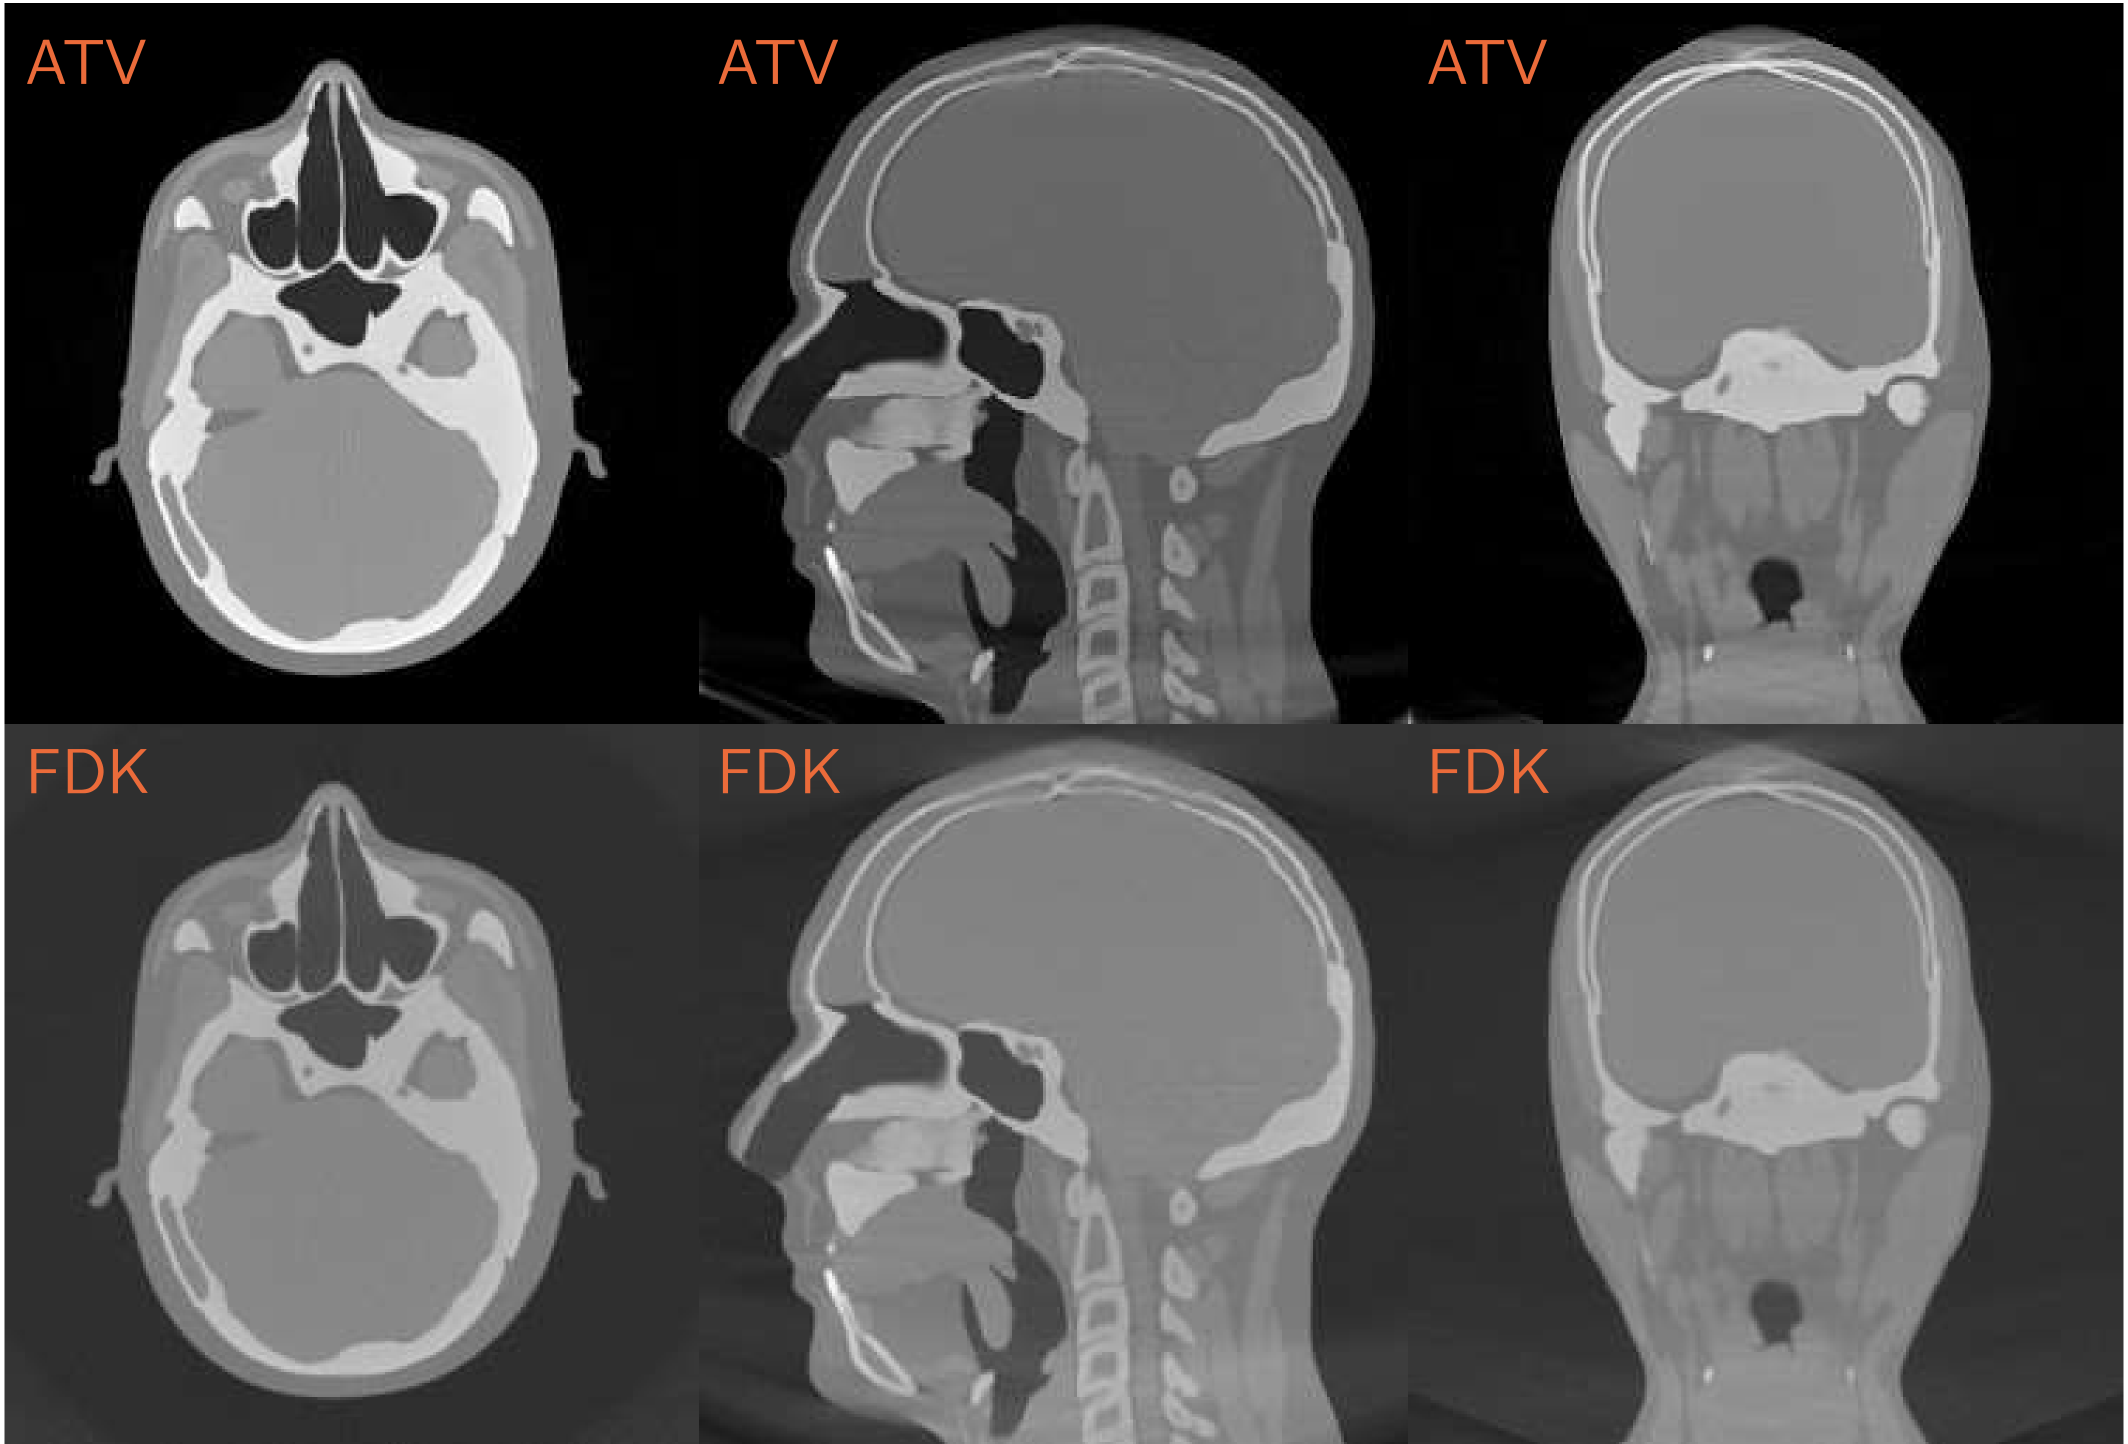

Anisotropic TV vs. FDK, dose 10% ($\mu = 2.5 \cdot 10^{-6}$)

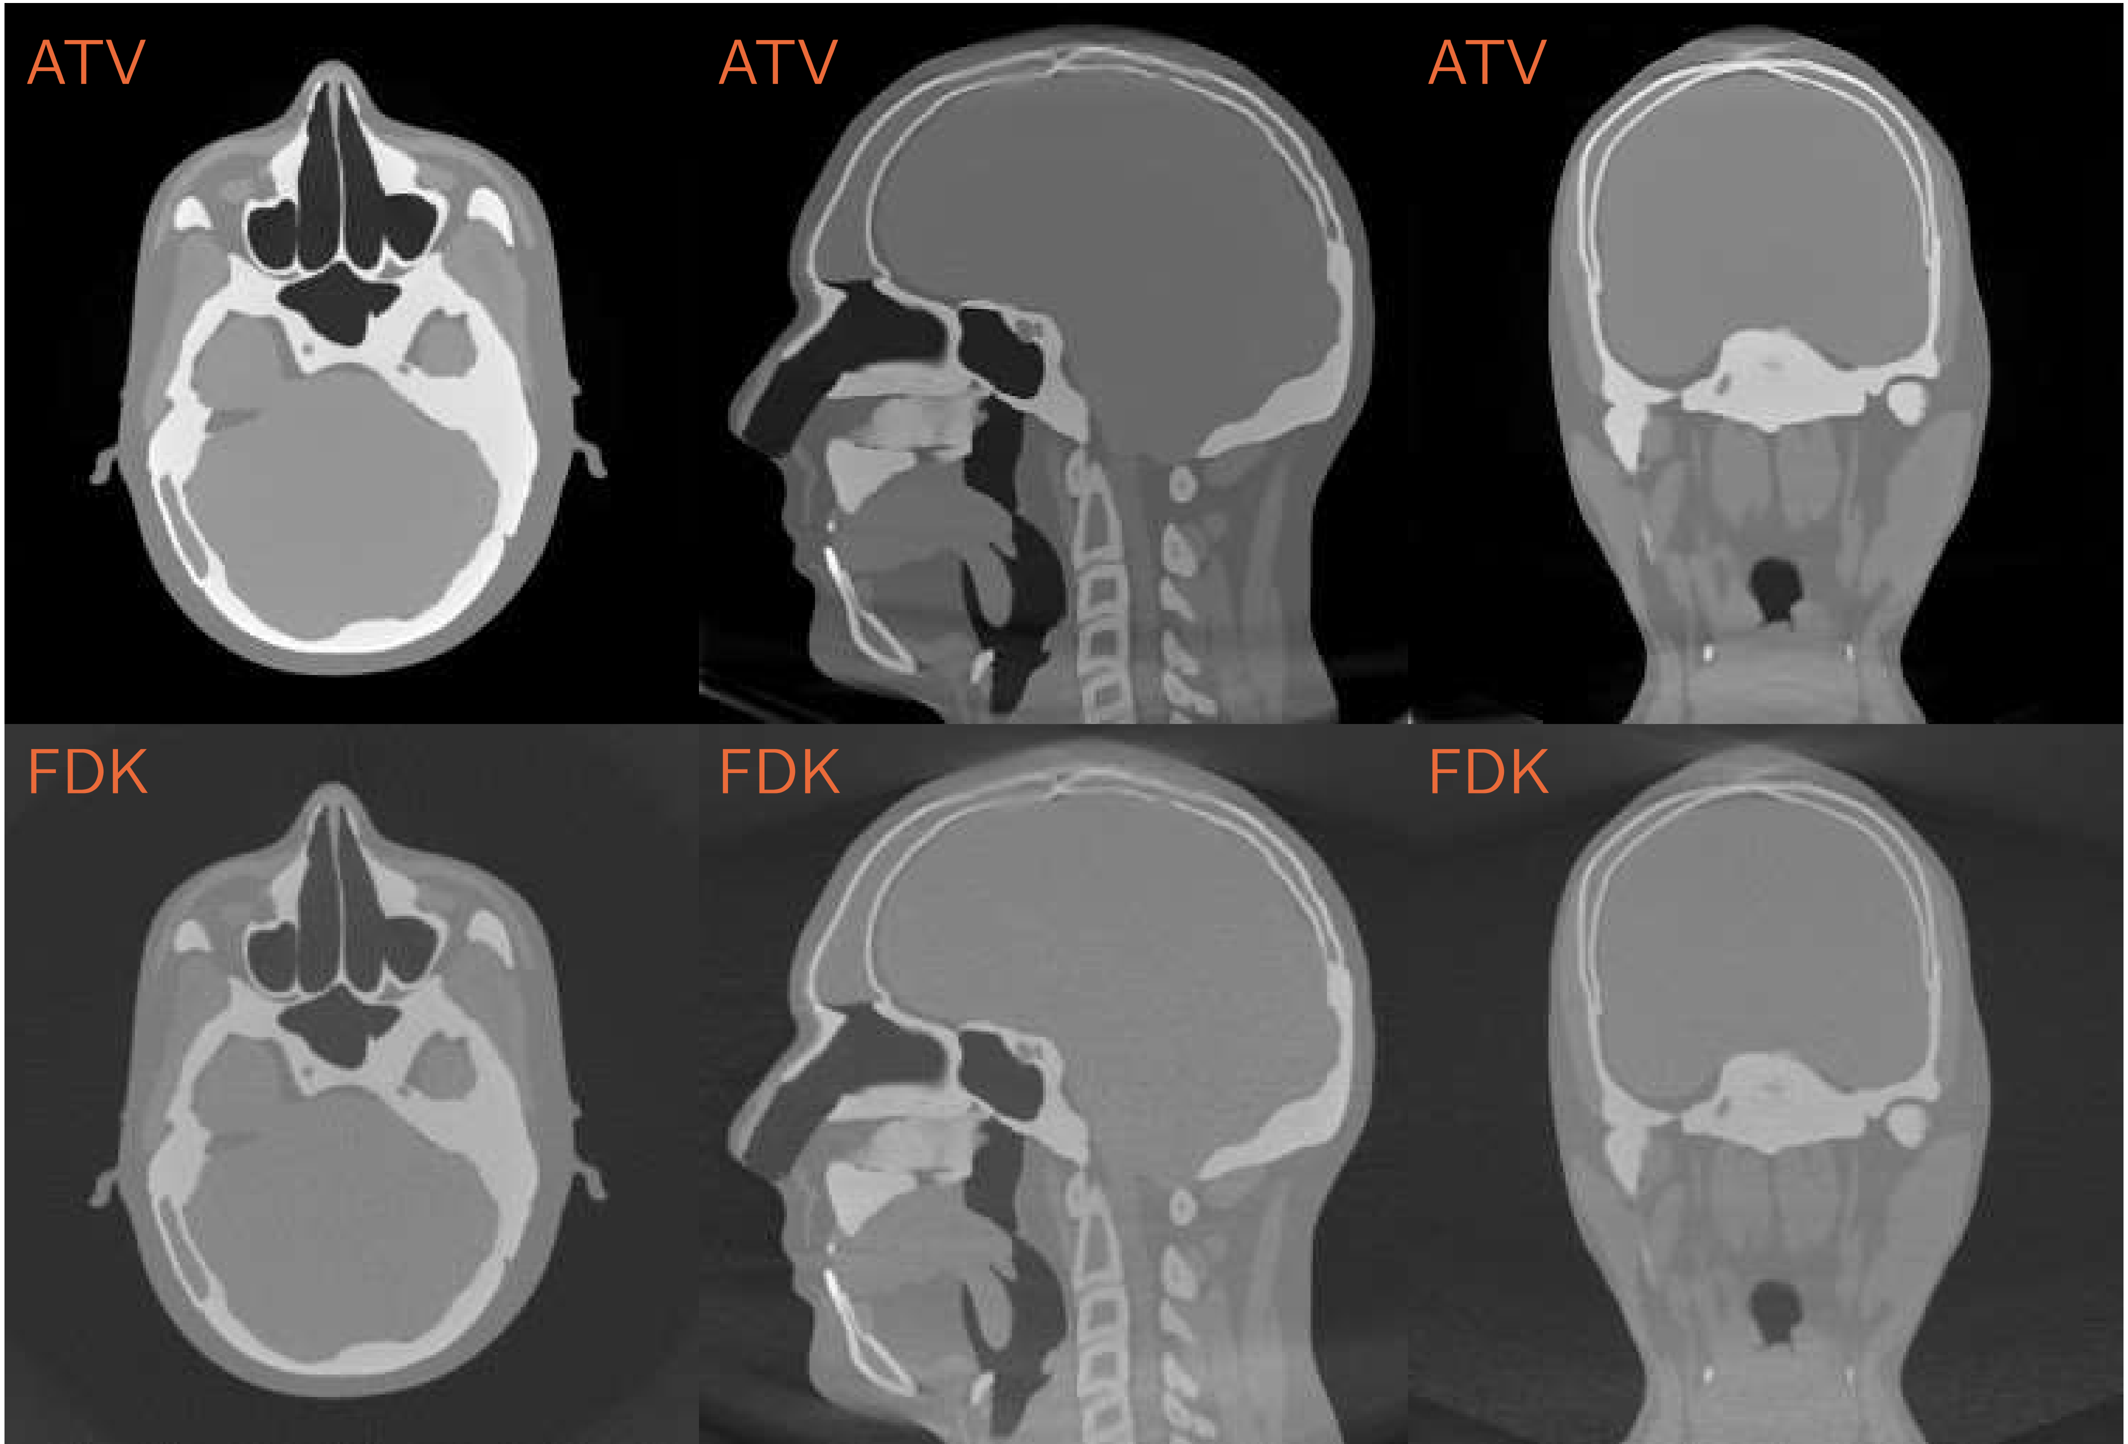

Anisotropic TV vs. FDK, dose 1% ($\mu = 7.5 \cdot 10^{-6}$)

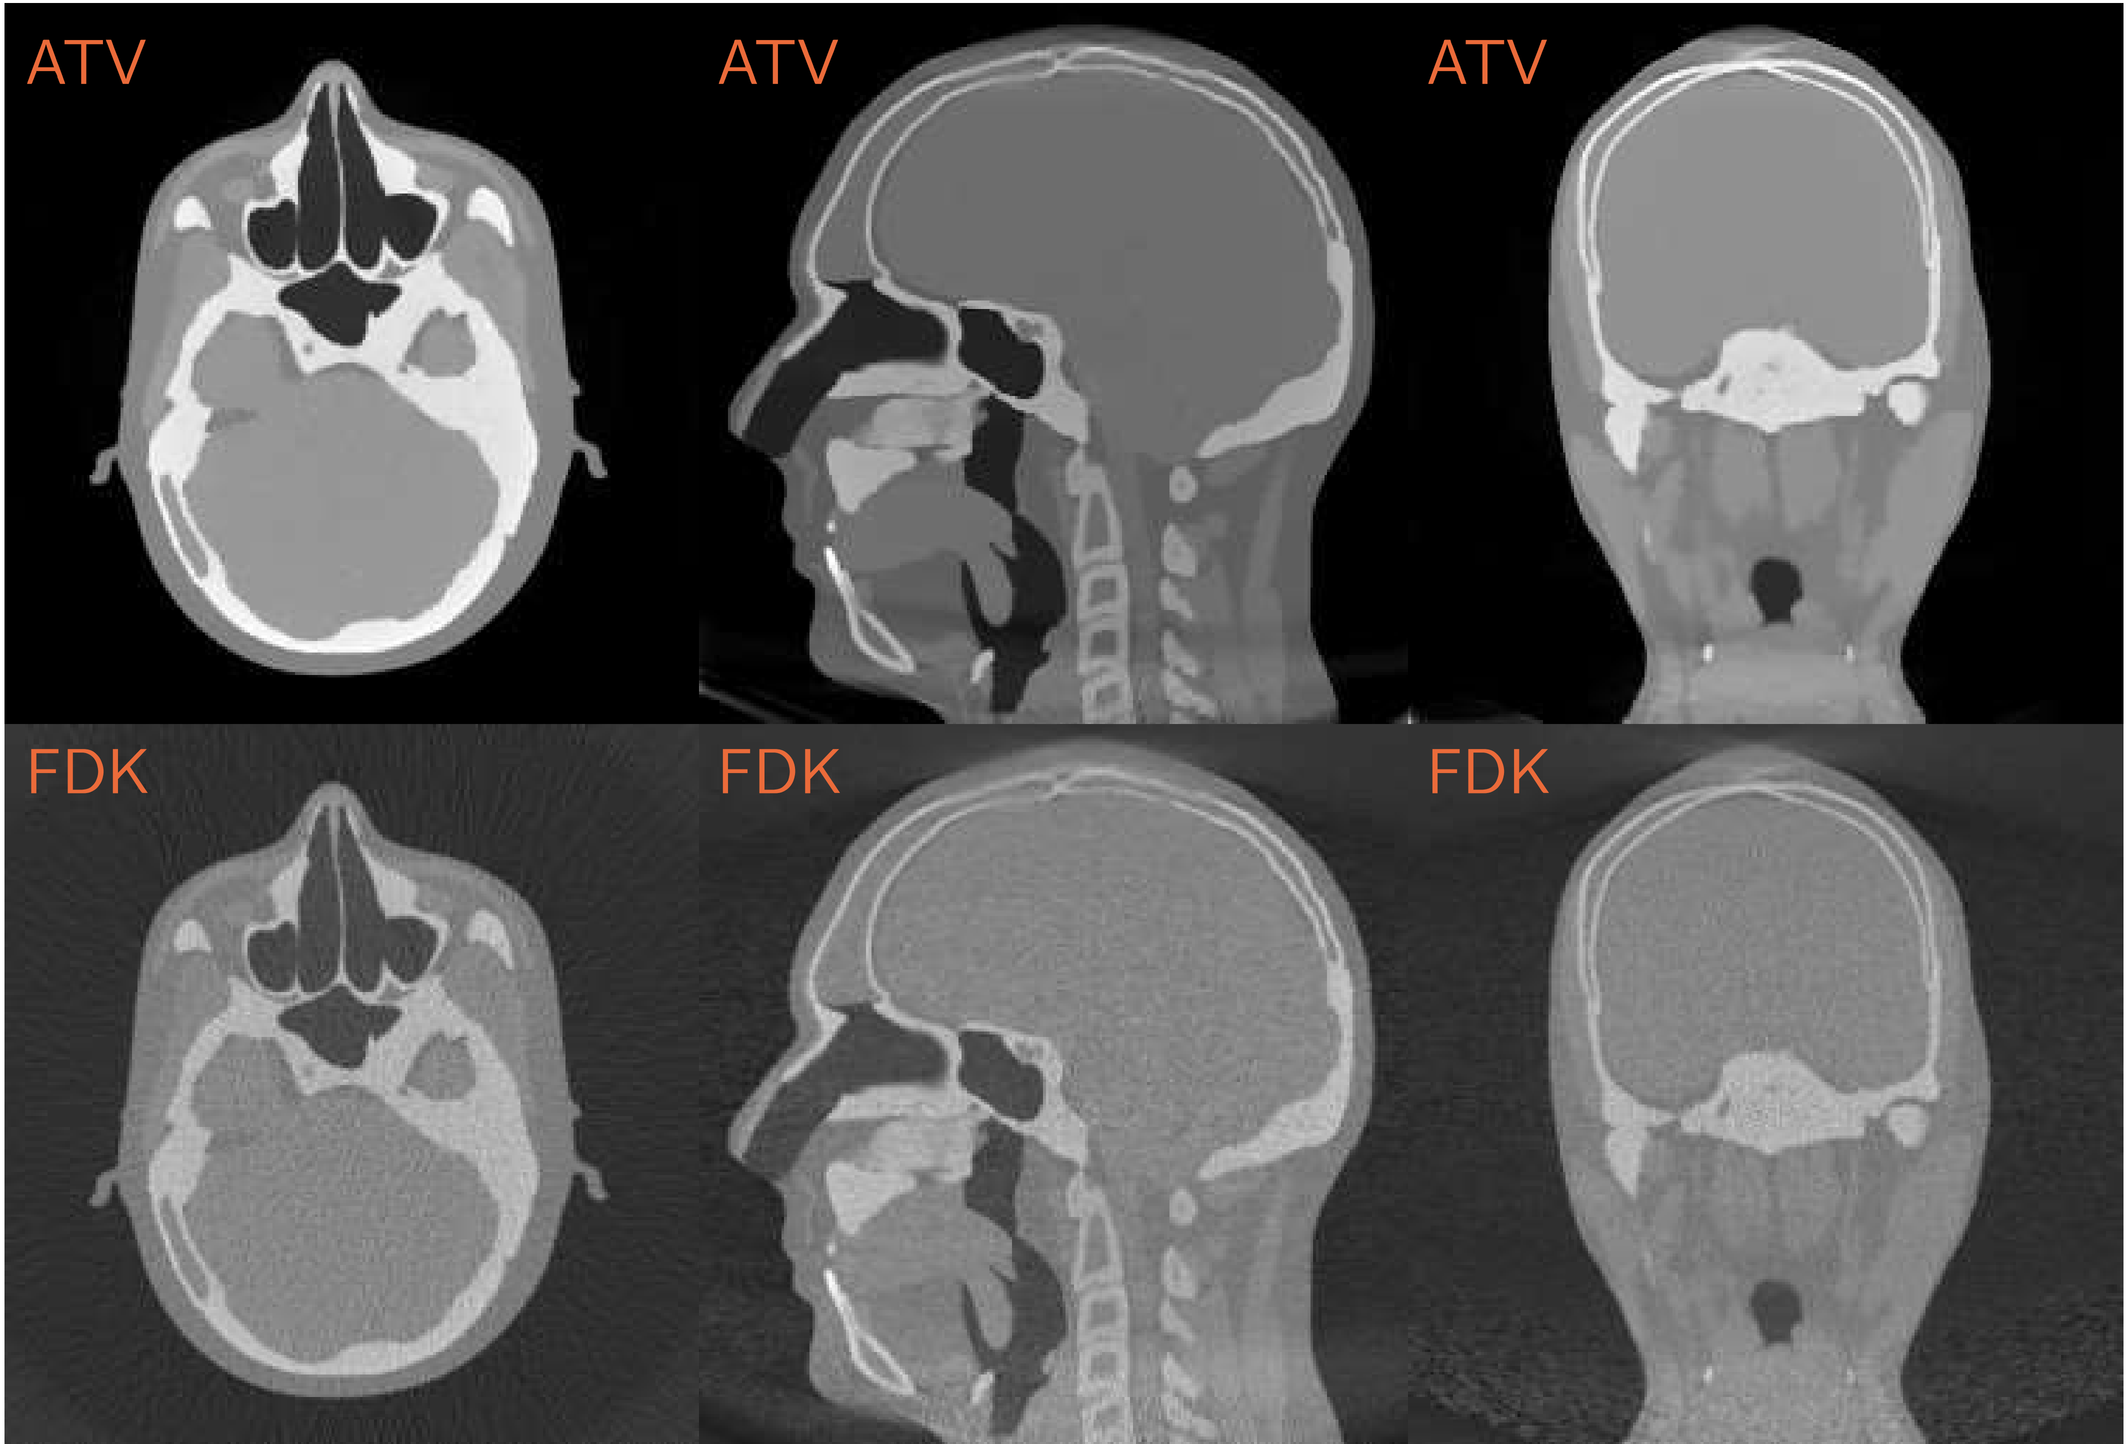

Anisotropic TV vs. FDK, dose 0.5% ($\mu = 1 \cdot 10^{-5}$)

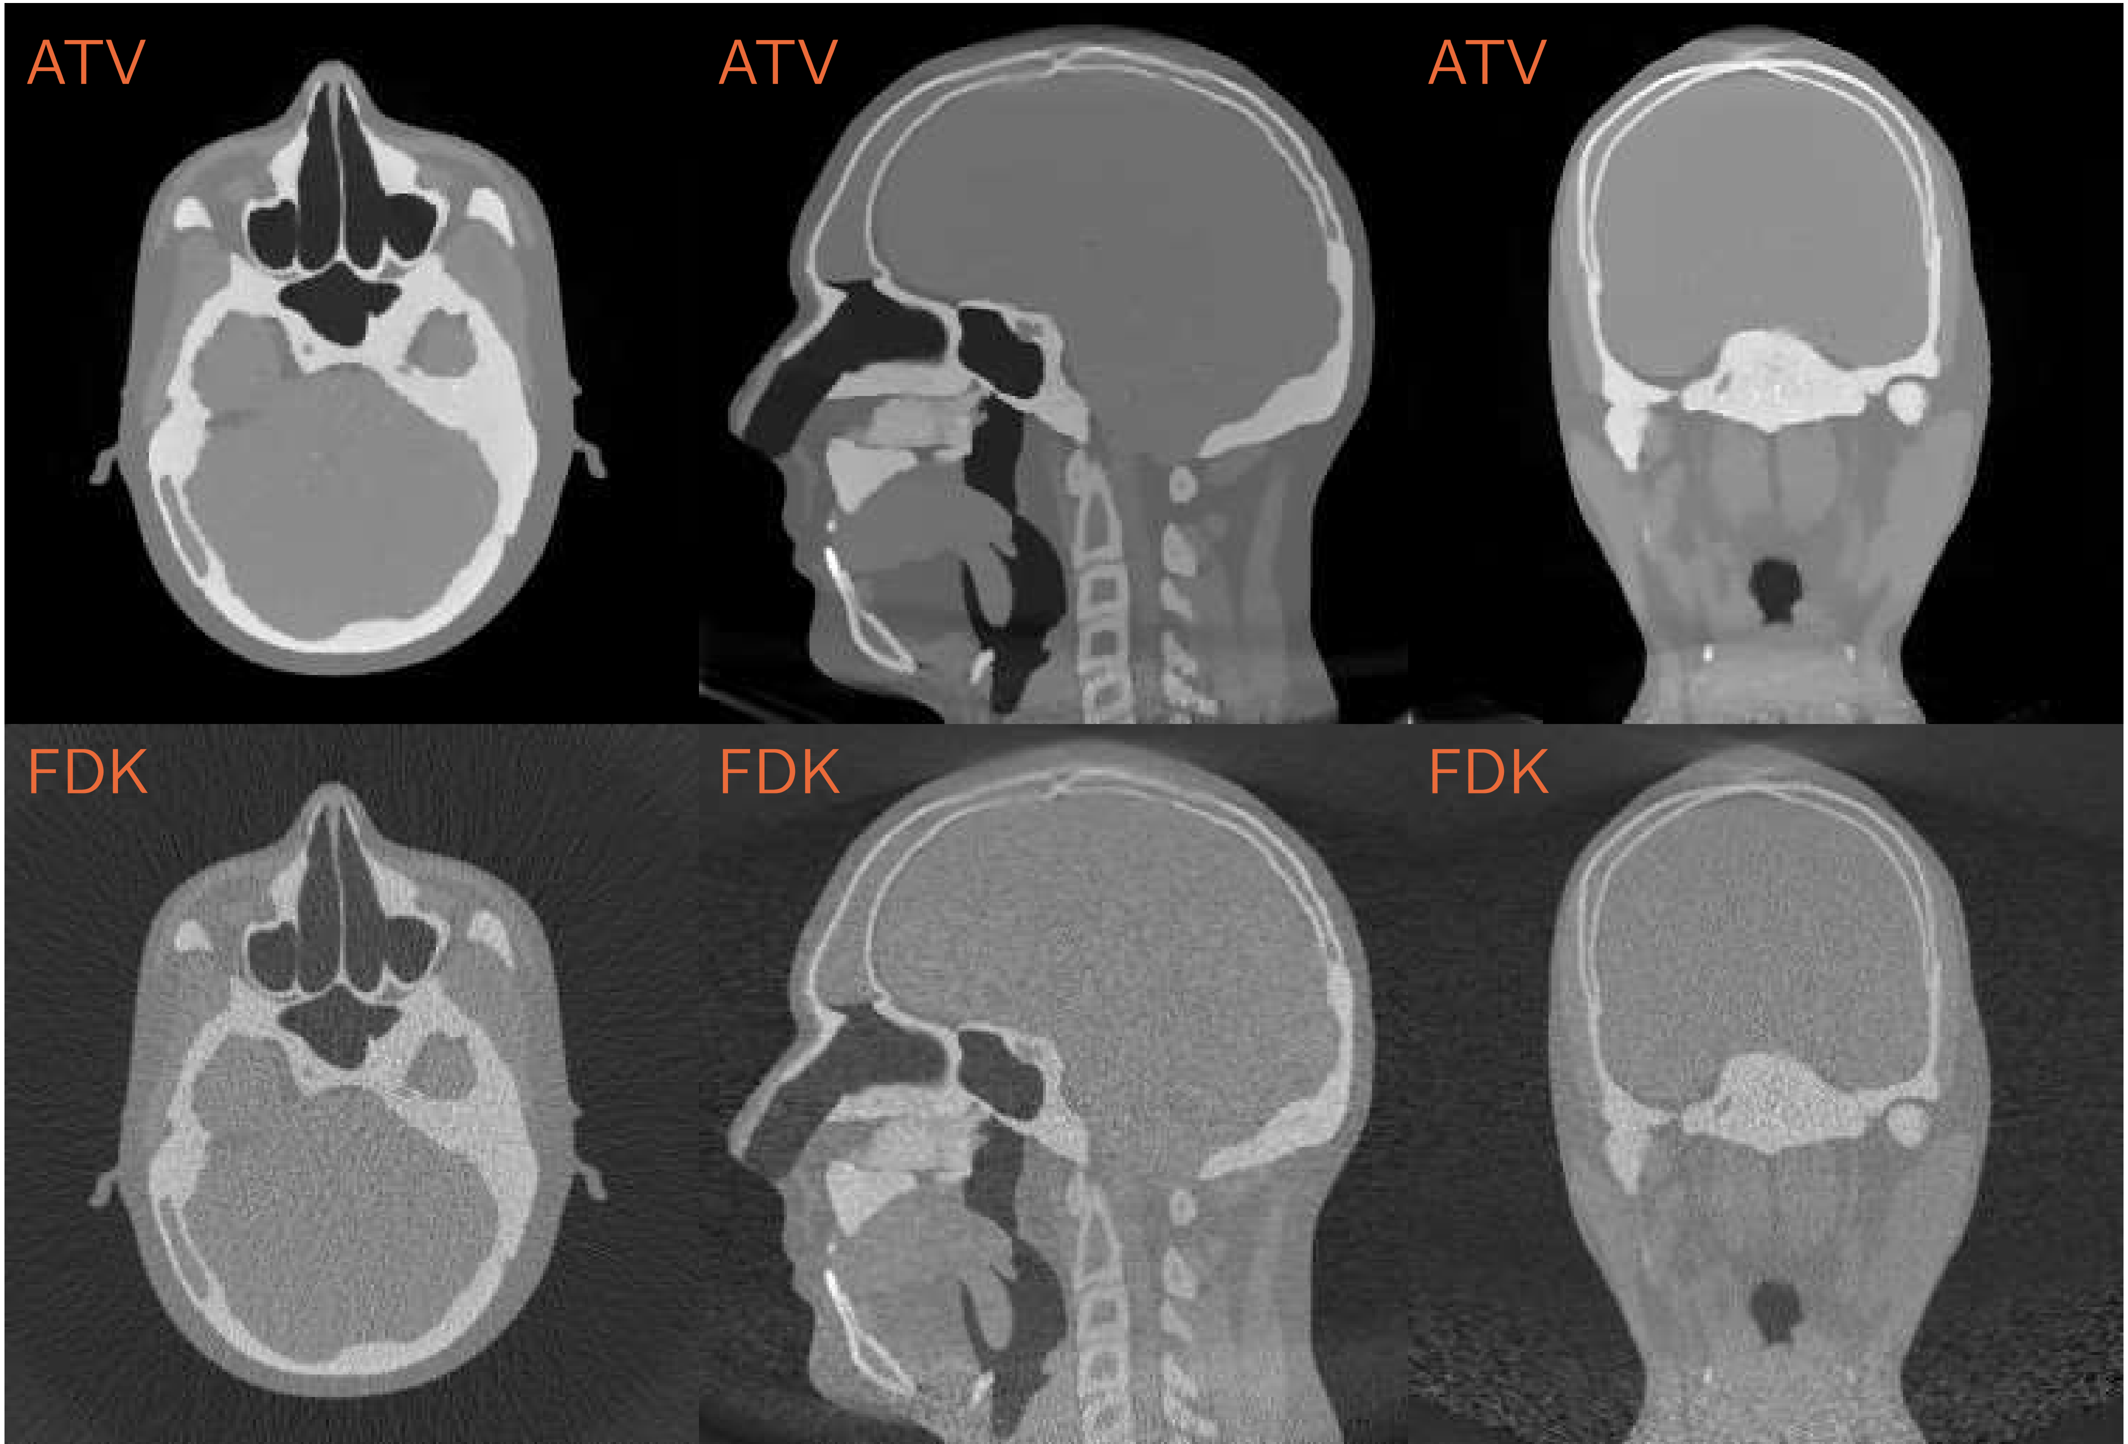

Anisotropic TV vs. FDK, dose 0.2% ($\mu = 2 \cdot 10^{-5}$)

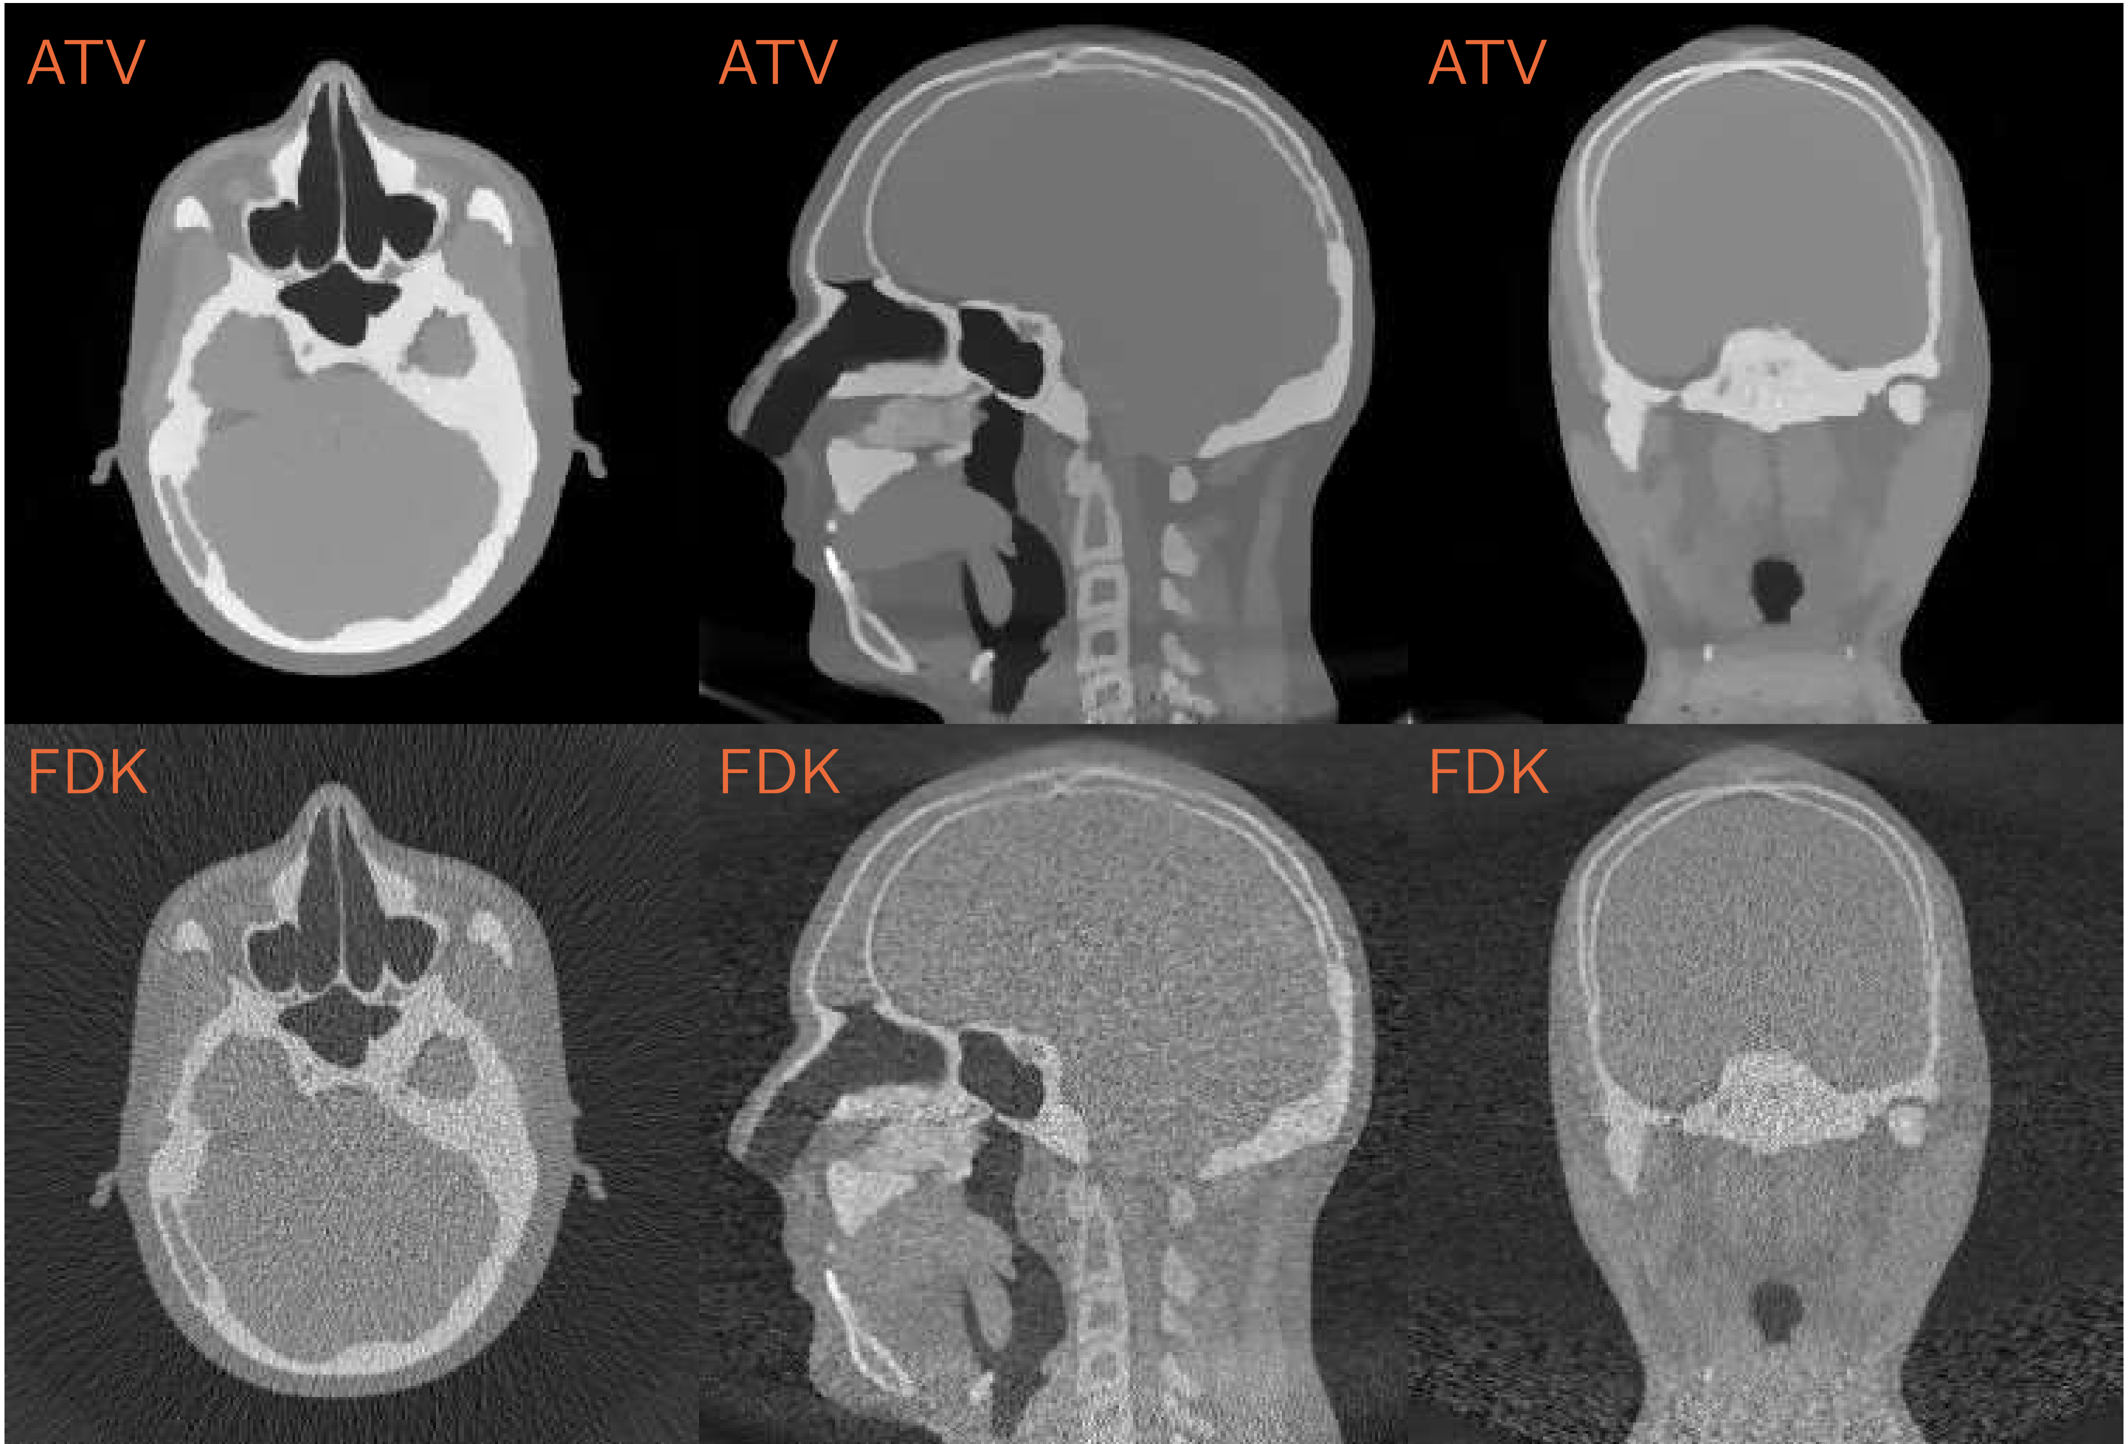

Anisotropic TV vs. FDK, dose 0.1% ($\mu = 2.8 \cdot 10^{-5}$)

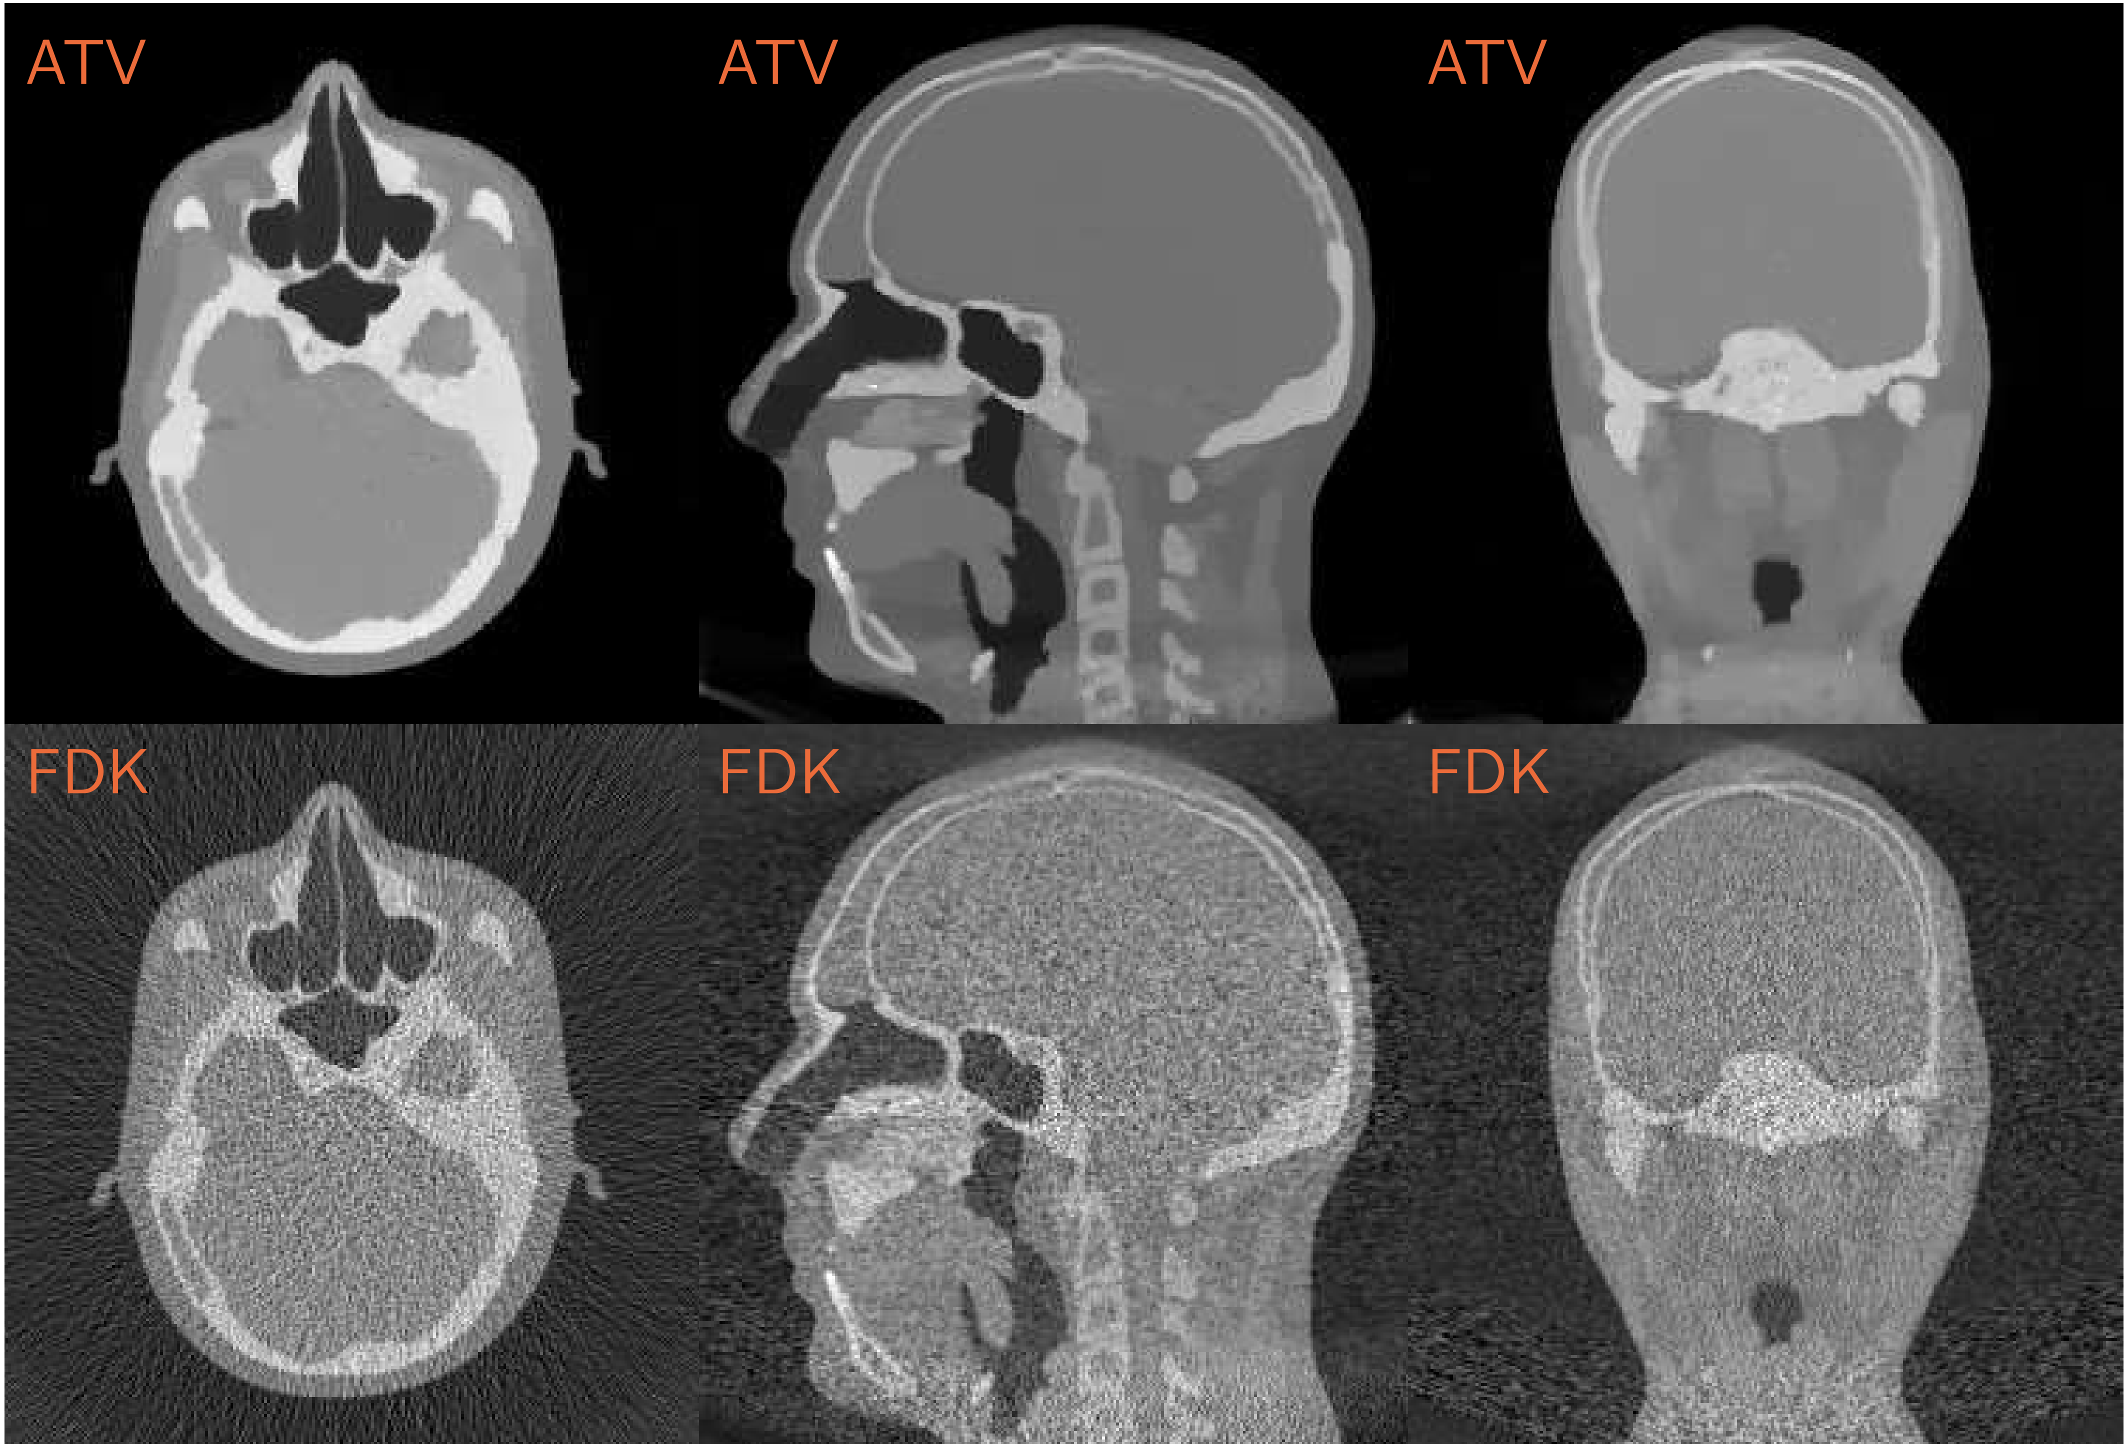

Convergence behaviour, dose 1%, $\mu = 1 \cdot 10^{-5}$

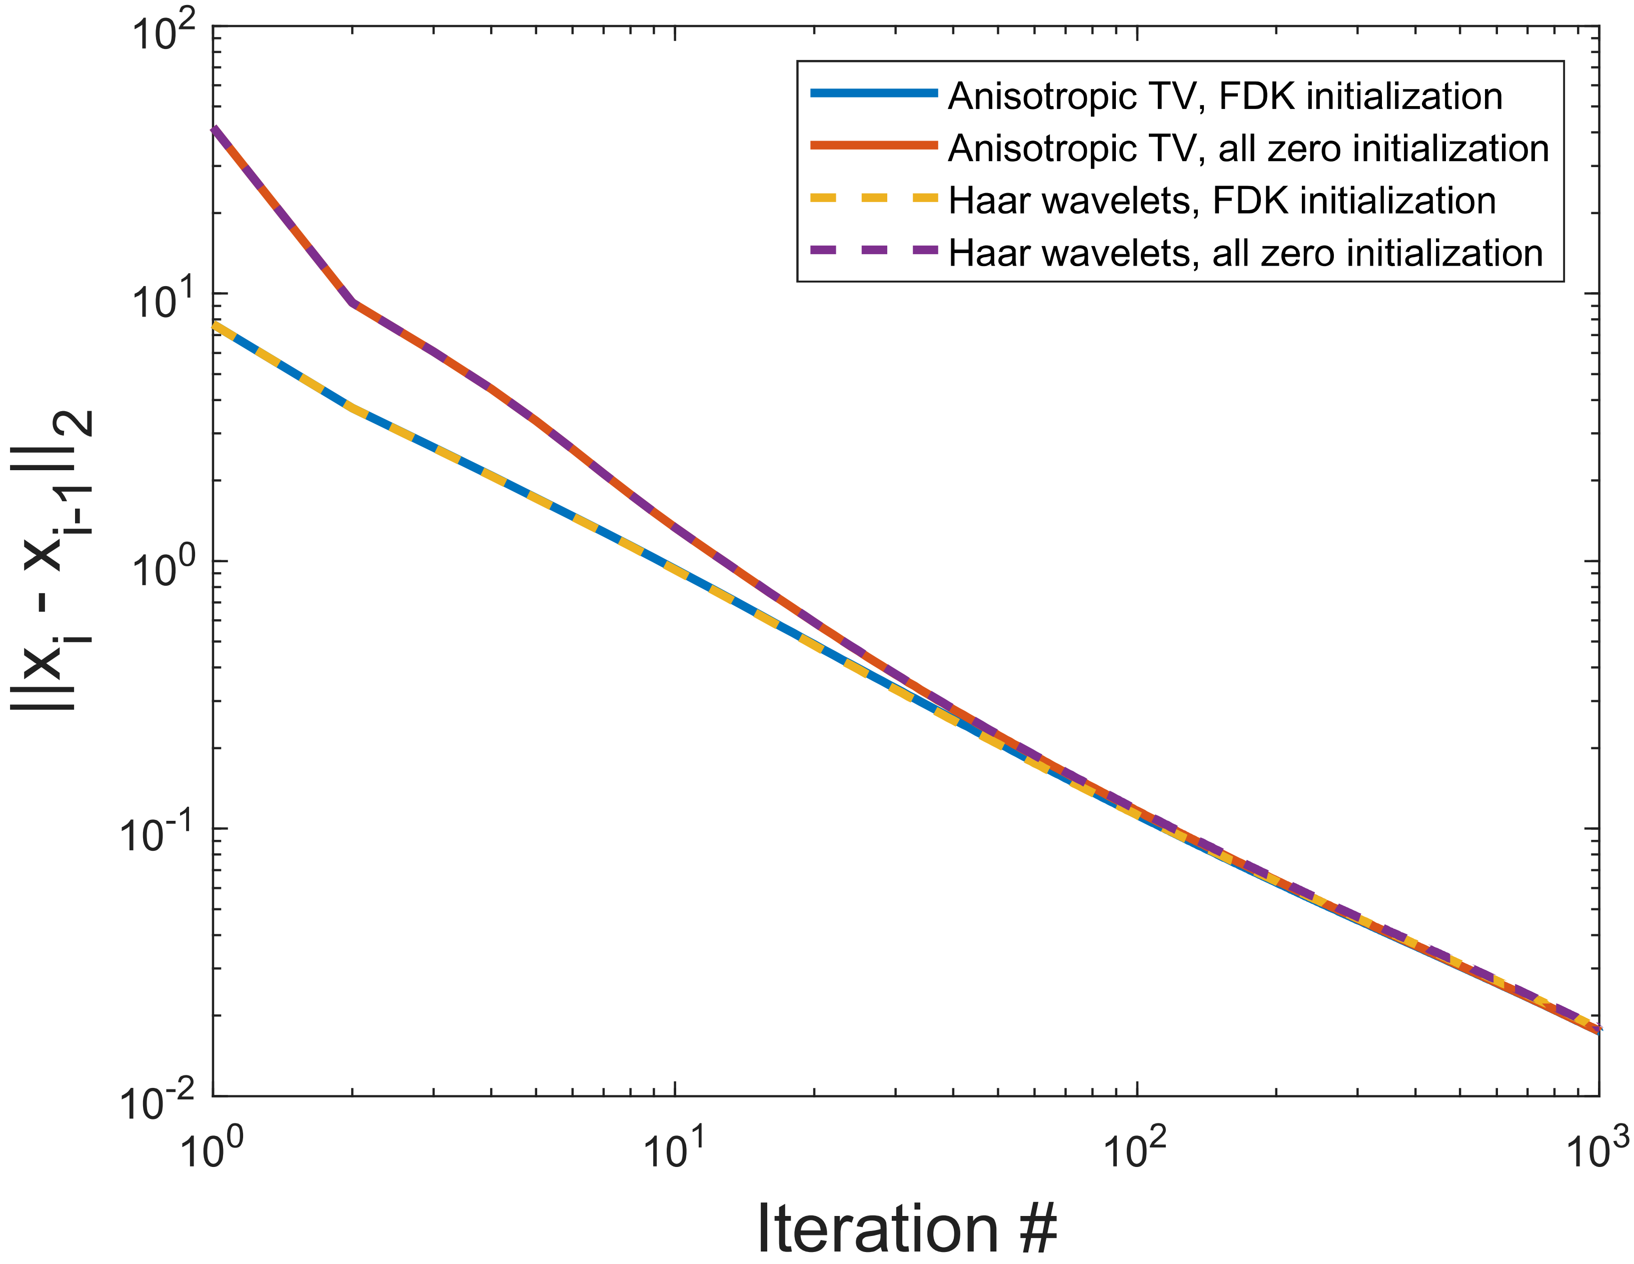

Conclusions and open questions

Iterative reconstruction techniques can improve soft tissue contrast in low dose CBCT.

Reconstruction quality is highly sensitive to choice of μ.

Long reconstruction times require semi-automated choice of regularization parameter.

Haar wavelet regularization begins to crumble at very low dose levels.

Sparsity of Haar and/or ATV components requires quantitative investigation.


THANK YOU