

Variational Diffusion Reconstruction for Multiple Artifact Correction in CT Imaging

Daniel Sanderson, Alex Denker, Zeljko Kereta, Kris Thielemans, Simon Arridge, Monica Abella

Universidad Carlos III de Madrid
University College London

March 3, 2026

Motivation

Usage of clinical CT systems is restricted in veterinary clinics:

- Space and budget limitations
- Systems with closed gantries that limit access to the animal.

Solutions currently available in the market:


- They do not allow the acquisition of radiography, fluoroscopy and CT with diagnostic quality
- They use a bulky closed gantry.

It is of interest to have a compact multimodal system based on an open gantry that allows easy access to the animal



Context

- Academia (UC3M)/ Industry (SEDECAL) collaboration.
- Multimodal system for veterinary imaging: Radiography + Dynamic radiography + CBCT
- Limited gantry rotation for CBCT
 - Reconstruction with conventional methods results in limited-data artifacts



State of art

Model Based Inverse Reconstruction (MBIR) methods: best results obtained with an image prior (gating, radiotherapy).

- Few applications have an image prior available.

DL supervised methods¹: postprocessing methods.

- High risk of hallucinations

Hybrid methods (DL + MBIR): prior is obtained with a DL model.

- Supervised models ² ³: low robustness to domain shifts.
- Diffusion models: high resolution priors, unpaired data, inherently iterative.
 - Types
 - Gradient based / **projection based** → Faster
 - Sampling / **non sampling** → Better data consistency
 - Current limitations
 - Only address 2D reconstruction
 - Do not focus on fast inference times
 - Solve a single type of artifact ⁴

¹[Madesta et al., 2024]

²[Zhang et al., 2021]

³[Piol et al., 2024]

⁴[Webber and Reader, 2024]

Goal

To implement a CT reconstruction algorithm to obtain artifact-free images in the clinical practice.

Requisites:

- ✓ Fast reconstruction times
- ✓ Robust against multiple artifacts
- ✓ Reconstructions are consistent with acquired data (minimal appearance of hallucinations)
- ✓ Compatible with 3D reconstruction

Proposed method: ADMM-RED-Diff

- Optimization problem given by RED-Diff⁵.
 - Projection based hybrid method for faster reconstructions
 - Non sampling method that uses diffusion model as regularizer
- Solved with ADMM: simpler to decouple the data consistency and regularization terms and solve them independently

Algorithm ADMM-RED-Diff

Input: $y, A, K, M, T, S, \{\alpha_t, \sigma_t, \lambda_t\}_{t=1}^T$

Initialize: x^0

for $k = 0, \dots, K$ **do**

$t \in \{T, T - S, \dots, 1\}$

$\epsilon \sim \mathcal{N}(0, I_n)$

$x^k \leftarrow \text{CG}((A^T A + \rho I), (A^T y + \rho z^{(k-1)} - \eta^{(k-1)}), M)$

$z^k = \text{SGD}\left(\lambda R(\alpha_t z^{k-1} + \sigma_t \epsilon) + \frac{\rho}{2\lambda} \left\| z - \left(x^{(k)} + \frac{1}{\rho} \eta^{(k-1)}\right) \right\|_2^2, 1\right)$

$\eta^k = \eta^{(k-1)} + \rho(x^{(k)} - z^{(k)})$

end for

Return: x

⁵[Mardani et al., 2023]

ADMM-RED-Diff: Optimization problem

Posterior distribution $p(x_0|y)$ is approximated by a variational distribution $q(x_0|y) \sim N(\mu, \sigma^2 I)$.

Optimized by minimizing the KL divergence:

$$\min_q \int q(x_0|y) \log \left(\frac{p(y|x_0)p(x_0)}{p(y)q(x_0|y)} \right)^{-1} = \underbrace{-\mathbb{E}_{q(x_0|y)} \log p(y|x_0)}_{\text{Data consistency}} + \underbrace{KL(q(x_0|y)||p(x_0))}_{\text{Regularization}} + C$$

Reconstruction is simplified to a traditional MBIR optimization problem with a diffusion model ϵ_θ as prior:

$$\min_{\mu} ||y - Ax_0||_2^2 + R(x_0) \quad (1)$$

where $R(\mu)$ is a stochastic regularizer whose exact gradient can be recovered without having to compute the Jacobian of the diffusion model ⁶ $\rightarrow$ Fast:

$$\nabla_{x_0} R(x_0) = \mathbb{E}_{t,\epsilon} [\lambda_t (\epsilon_\theta(x_0; t) - \epsilon)] \quad (2)$$

where $\lambda_t = \frac{2\beta_t\sigma_v^2\alpha_t}{\sigma_t}$, and which is approximated as $\lambda_t = \lambda \frac{\sigma_t}{\alpha_t}$.

⁶[Song et al., 2021]

ADMM-RED-Diff: Solving with ADMM

ADMM proposes the following constrained optimization problem:

$$\min_x \underbrace{\|y - Ax\|_2^2}_{f(x)} + \underbrace{\lambda R(z)}_{g(z)} \quad \text{subject to} \quad x - z = 0$$

where λ has been added to control the weight of $R(z)$. The constraint is incorporated via the Augmented Lagrangian:

$$\min_x \|y - Ax_0\|_2^2 + \lambda R(z) + \frac{\rho}{2} \|x_0 - z\|_2^2 + \eta^T (x_0 - z)$$

This formulation allows to split the original problem:

$$x_0^k = \min_x \frac{1}{2} \|Ax_0 - y\|_2^2 + \eta^{(k-1)T} (x_0 - z^{k-1}) + \frac{\rho_k}{2} \|x_0 - z^{k-1}\|_2^2 \quad (3)$$

$$z^k = \min_z \lambda R(z) + \eta^{(k-1)T} (x_0^k - z) + \frac{\rho_k}{2} \|z - x_0^k\|_2^2 \quad (4)$$

$$\eta^k = \eta^{k-1} + \rho(x_0^k - z^k) \quad (5)$$

ADMM-RED-Diff: Solving with ADMM

z-subproblem

Solve with a single stochastic gradient descent (SGD) step given that each iteration would require re-evaluating the regularization term.

x-subproblem

Solve with conjugate Gradient (CG): the fastest for linear systems

- ✓ Fastest method for convex problems
- ✗ Only valid for Gaussian forward model

The feasibility of the solution does not depend on hyperparameter scheduling

$$c(x_k) \approx -\frac{1}{\rho}(\eta^* - \eta^k) \quad (6)$$

Evaluation

Datasets

- Clinical thorax images:
 - Publically available dataset CT-RATE: 48 patients
- Veterinary:
 - Canon Aquilon Exceed LB: 1 dog

Data simulation


- ASTRA-toolbox with polychromatic spectrum at 110 kVp
- $LSA \in [120, 360]$, $LNP \in [1, 12]$

Ablation study:

- Convergence
- Stability of hyperparameters

Comparison methods


- Baseline: RED-Diff + Adam, $T=1000$
- Sampling diffusion: DDS^a, $T=100$
- Supervised: DeepParker




^a[Chung et al., 2023]

ADMM-RED-Diff: Ablation

- ADMM convergence despite the stochasticity of the regularizer:




- Low variability in performance for the appropriate choice of hyperparameters.
- Improved performance for multiple iterations over the x-subproblem.
- Optimal balance between performance and reconstruction time achieved for 100 inference steps (≈ 30 s).




Method comparison: clinical data

The proposed method recovers details without introducing false textures




Method comparison: clinical data

The proposed method recovers details without introducing false textures




Method comparison: clinical data

The proposed method recovers anatomical information in LSA regions




Method comparison: clinical data

The proposed method recovers anatomical information in LSA regions



Method comparison: Veterinary data

The proposed method is capable of generalizing to out of domain data



It exhibits faster reconstruction times than RED-Diff solved with the Adam optimizer

DeepParker	DDS	RED-Diff T-1000	ADMM RED-Diff
0.001x	1x	4.2x	1x

Table: Reconstruction times (s) for $160^\circ/3^\circ$ and $x=30$ s

Discussion

The proposed method allows for fast reconstructions with minimal appearance of hallucinations.

- Non sampling method: better exploits the spatial consistency of MBIR algorithms.
- ADMM solving: acceleration of the reconstruction process

Limitations and future work

- Evaluation performed on noise-free data
 - Approximate Poisson noise with a weighted gaussian model
- Implemented for Fan-Beam geometry: adaptation to Cone Beam geometry for real data.
 - Accelerate model inference: distillation, neural architecture search
 - Random selection of multi-view slices, 3D prior in a latent space, patch based methods
- Evaluate on lesion detection tasks

Thank You!



Chung, H., Lee, S., and Ye, J. C. (2023).

Decomposed diffusion sampler for accelerating large-scale inverse problems.

arXiv preprint arXiv:2303.05754.



Madesta, F., Sentker, T., Gauer, T., and Werner, R. (2024).

Deep learning-based conditional inpainting for restoration of artifact-affected 4d ct images.

Medical Physics, 51(5):3437–3454.



Mardani, M., Song, J., Kautz, J., and Vahdat, A. (2023).

A variational perspective on solving inverse problems with diffusion models.

arXiv preprint arXiv:2305.04391.



Piol, A., Sanderson, D., Del Cerro, C. F., Lorente-Mur, A., Desco, M., and Abella, M. (2024).

Hybrid reconstruction approach for polychromatic computed tomography in highly limited-data scenarios.

Sensors, 24(21):6782.



Song, Y., Durkan, C., Murray, I., and Ermon, S. (2021).

Maximum likelihood training of score-based diffusion models.

In Ranzato, M., Beygelzimer, A., Dauphin, Y., Liang, P., and Vaughan, J. W., editors, *Advances in Neural Information Processing Systems*, volume 34, pages 1415–1428. Curran Associates, Inc.



Webber, G. and Reader, A. J. (2024).

Diffusion models for medical image reconstruction.

BJR— Artificial Intelligence, 1(1):ubae013.



Zhang, K., Li, Y., Zuo, W., Zhang, L., Van Gool, L., and Timofte, R. (2021).

Plug-and-play image restoration with deep denoiser prior.

IEEE Transactions on Pattern Analysis and Machine Intelligence, 44(10):6360–6376.