

Integrating Atmospheric Dispersion Modeling Priors into Cuboid Splatting for Spatiotemporal Reconstruction of Airborne Radioiodine After Nuclear Accidents

Mareike Böckel^{1,2}, Stephan Doerfel², Kathrin Meisenberg³, Oliver Meisenberg³, Max Friedrich¹ and Mattis Hartwig^{1,4}

¹singularIT GmbH

²Kiel University of Applied Sciences

³Federal Office for Radiation Protection, Germany

⁴German Research Center for Artificial Intelligence

{mareike.boeckel, mattis.hartwig, max.friedrich}@singular-it.de, stephan.doerfel@haw-kiel.de,
{kmeisenberg, omeisenberg}@bfs.de

Abstract

In nuclear power plant accidents, airborne radioiodine poses major health risks, making reliable reconstructions of its spatiotemporal distribution crucial for emergency management. Current state-of-the-art prognosis systems use atmospheric dispersion modeling but ignore posterior evidence from emergency care centers, comprising movement profiles and thyroid measurements of affected individuals. A first study showed that the AI method Cuboid Splatting can reconstruct iodine air concentrations from such data but it ignores simulations from established prognosis systems. Our multidisciplinary team extends Cuboid Splatting by incorporating these simulations as priors and subsequently correcting them using movement and thyroid data. Several ways to translate and correct priors are developed. The best-performing approaches are combined into a novel Cuboid Splatting-with-prior mechanism, which we evaluate using constructed prior scenarios representing different error types and intensities. Using Cuboid Splatting-with-prior yields more accurate reconstructions than (i) the used dispersion simulations alone and (ii) plain Cuboid Splatting without prior. Across reconstructions, the mean scenario error is 19.61%, improving on (i) by 28.02pp and on (ii) by 89.86pp, the latter with particularly large gains at high spatial resolution. These results demonstrate that combining simulation-based priors with measurement-based posterior inference can substantially improve the reconstruction of iodine air activity concentrations in nuclear emergencies.

1 Introduction

Globally, nuclear power production is on the rise to feed increasing electricity demands, among others due to growing use of server centers and generative AI models [International Energy Agency, 2025]. Discussion on EU level recently led

to the classification of nuclear energy production as a sustainable energy source [European Commission, 2022], despite its risks in cases of accidents at nuclear power plants (e.g. Fukushima, Chernobyl). Beyond accidents, the exposure and potential damage of nuclear power plants during armed conflicts pose a severe risk (e.g. Zaporizhzhia, [IAEA, 2024]).

In case of an accident at a nuclear power plant, radioactive isotopes, particularly iodine isotopes, can be released to the air in the surrounding area. This creates not only an ecological problem but also poses severe health risks for the people in the affected area, notably an increased risk for thyroid cancer. Adequate management of such a radiological emergency therefore is of immense importance for public health. Part of the management is a good situation overview of air concentrations of radioiodine isotopes not only during but also after the emergency for retrospective assessments, deriving follow-up measures, information of the public, reconstruction of exposure of people not measured directly after the accident.

In order to assess the situation in the event of a radiological emergency, decision support systems are used internationally in many countries. These generally include an atmospheric dispersion model using the Lagrange particle model approach, which can be used to estimate environmental contamination and thus the radiation exposure of the affected population based on weather and source term information. Available systems comprise among others ARGOS (PDC-ARGOS, Denmark, [Hoe *et al.*, 2009]), FLEXPART (University of Vienna, Austria, and others, [Stohl, 2006]), HYSPLIT (NOAA, USA, [Stein *et al.*, 2015]) and RODOS ("Realtime Online Decision Support System"; KIT, Germany, [Raskob *et al.*, 2011; Raskob *et al.*, 2016]). The state-of-the-art system in Germany, RODOS, models atmospheric dispersion and deposition of radioactive particles in a nuclear accident and estimates the spatiotemporal iodine concentrations in the air using a Lagrangian particle dispersion model.

However, these simulations do not incorporate an important source of empirical data generated during an emergency. In such events, people in affected areas are directed to emergency care centers, where individual thyroid exposure measurements are performed and evacuees are asked to report their movement profiles during the critical exposure period.

A sufficient number of such reported paths including the measurements provides substantial posterior evidence on the spatiotemporal distribution of radioiodine air concentrations. Recently, a first study has demonstrated that this type of individual exposure and movement data can indeed be used to extract meaningful information on the contamination distribution [Friedrich *et al.*, 2026].

In that study, Cuboid Splatting was introduced and shown to be the best performing approach. It has, however, a major downside: It cannot incorporate the information generated by state-of-the-art atmospheric dispersion modeling approaches. As a consequence, an authority tasked with estimating the spatiotemporal distribution of radioactive contamination would be forced to choose between two complementary sources of information. Either it relies on dispersion-based models, thereby ignoring posterior evidence like individual exposure data, or it uses solely the latter, thereby forfeiting the physically grounded insights provided by atmospheric dispersion modeling. This separation results in an avoidable loss of information and limits the overall quality of reconstructions of radioiodine air activity concentrations.

To address this challenge, here we present research from a *multidisciplinary collaboration* between AI experts and domain experts in radiation protection, in which the two approaches – atmospheric dispersion modeling and individual exposure data – are combined. The core idea is to treat the outputs of established atmospheric dispersion models as informative priors, which are subsequently updated using path-based iodine exposure measurements as posterior evidence, thereby unifying both approaches. The data science experts are associated with the private software company singularIT GmbH, the Kiel University of Applied Sciences (HAW Kiel) and the German Research Center for Artificial Intelligence (DFKI). Their roles encompassed method development for AI models, model evaluation and Python programming. The domain experts are professionals from the German Federal Office for Radiation Protection (BfS). Their role was to guide the method development and evaluation process to fit the addressed real world problem in the application domain and to provide dispersion simulation data.

Our research links *AI, public health, and environmental sustainability*. Enabling improved reconstructions of radioiodine air concentrations after nuclear accidents contributes to:

- SDG 3 (Good Health and Well-Being) by providing reliable exposure estimates that guide immediate health protection, targeted long-term medical follow-up, and clear risk communication, helping to reduce anxiety and uncertainty in affected populations.
- SDG 15 (Life on Land) by identifying highly contaminated areas and informing evidence-based decontamination, land-use restrictions, agricultural management, and long-term ecological risk assessment, thereby supporting the sustainable management of affected ecosystems.
- SDG 17 (Partnerships for the Goals), as robust reconstructions offer a shared scientific reference for national authorities and international organizations, improving cross-border coordination as well as enhancing transparency, and trust in institutions.

More broadly, this work advances responsible risk governance and intergenerational equity by enabling societies to better quantify, manage, and mitigate the long-term environmental and health impacts of nuclear accidents, thereby strengthening resilience to radiological crises.

Key Contributions: 1) Extension of the Cuboid Splatting reconstruction method to incorporate prior scenarios by integrating two key data sources for reconstructing radioactive air contamination: Emergency care center data (movement profiles, thyroid measurements) and RODOS simulations (dispersion modeling, meteorology, estimated release). 2) Full experimental evaluation of the final Cuboid Splatting-with-prior mechanism. 3) Development of multiple approaches for translating and correcting prior scenarios via Cuboid Splatting in pilot studies. 4) Enabling better reconstructions in the event of a nuclear emergency. Concisely, we developed a method that unifies simulation-based priors with measurement-based posterior inference for reconstructing radioiodine air concentrations, solving a methodological problem that had not been solved before.

2 Related Work and Background

The spatiotemporal iodine distribution is denoted as a three dimensional tensor X whose index set is defining the 3D-spatiotemporal grid G . When we refer to iodine distributions later, we use the term *scenario* for brevity. A model (or reconstructed) distribution $\hat{X}$ is good if the distance between X and $\hat{X}$ is minimal.

When comparing two spatiotemporal iodine distributions or scenarios (e.g. X and $\hat{X}$), the *scenario error* in percent is calculated in the following way: The point-wise absolute differences between scenario 1 (e.g. the reconstruction) and scenario 2 (e.g. the original scenario) are summed and divided by the sum of iodine concentrations in scenario 2. This approach of calculating the relative error is necessary as the scenarios contain many 0 and close to 0 values, which impedes point-wise calculation of the relative error.

Since X is unknown during optimization, we cannot minimize the distance directly. The state-of-the-art system in Germany to model atmospheric dispersion of radionuclides in case of a nuclear accident is the RODOS system [Raskob *et al.*, 2016], which estimates the spatiotemporal iodine concentrations in the air using a Lagrangian particle dispersion model F taking into account several parameters Θ comprising current weather data and information on the composition and quantity of the radioactive substances released

$$\hat{X} = F(\Theta). \quad (1)$$

Recent research assessed whether it is possible to reconstruct a scenario (the spatiotemporal iodine air concentrations) from data collected at emergency care centers spontaneously set up in a nuclear emergency [Friedrich *et al.*, 2026].

Persons from the affected area are requested to go to these centers for decontamination, to disclose their movement profile from the last hours and to undergo thyroid activity measurement. The collected data is described by the path data set D , containing a path tuple (p, g_p, A_p, y_p) for each person p , where $g_p \subseteq G$ contains the visited grid cells, A_p contains the

person specific absorption factors over time and y_p contains the thyroid activity measurement, which is modeled as

$$y_p = \sum_{i \in g_p} T A_{p,i} X_i, \quad (2)$$

where T is the timestep size.

The expert-guided path data synthesis approach, published in Friedrich et al. [2026] and used in this project as well, was developed because no real dataset of this kind exists from a nuclear emergency. It models a broad set of parameters based on scientific literature: movement types, group constellations, age-dependent breathing and iodine uptake, protective factors (e.g., respiratory masks, thyroid-blocking iodide tablets, indoor/outdoor), timing of tablet intake, and radioiodine retention curves.

Scenarios cover areas with approx. 400 to 800 km side length. Scenarios and paths are discretized to a specific resolution (spatial 10x10, 30x30, 100x100, time: 1h steps).

For reconstructing the iodine air concentrations from this indirect data source two reconstruction methods, both originating in 3D image reconstruction, were transferred to the use case in radiation protection [Friedrich et al., 2026]: Memory-based NeRFs based on Plenoxels [Fridovich-Keil et al., 2022] and Cuboid Splatting based on Gaussian Splatting [Kerbl et al., 2023], out of which Cuboid Splatting overall outperformed memory-based NeRFs. The indirect reconstruction follows the minimization of the measured exposure y_p versus predicted exposure $\hat{y}_p$ of individuals which is also referred to as the *path error*. Reconstructing $\hat{X}$ based on path data follows

$$\hat{X} = \arg \min_{\hat{X}} \sum_{p \in \pi_1(D)} (y_p - \hat{y}_p)^2, \quad (3)$$

where π is the projection function.

The state-of-the-art of the method Cuboid Splatting consists of a random initialization of a set of cuboids in spacetime and subsequently training their parameters (center, length, width, height, and iodine value) based on the path data (movement profiles and thyroid measurements) to reconstruct the 3D spatiotemporal distribution of iodine concentrations following the accident. Cuboid parameters are optimized by comparison of actual and reconstructed thyroid measurements at the end of a person’s path based on the current state of 3D iodine air concentration reconstruction. An adaptive density control mechanism adds cuboids along under-reconstructed paths and deletes cuboids that became obsolete during training. The idea of representing the scene with trainable geometric objects and dynamically adapting their density is inherited from Gaussian Splatting [Kerbl et al., 2023]. While the latter accumulates color and density along camera rays and optimizes with respect to image reconstruction error, Cuboid Splatting [Friedrich et al., 2026] accumulates iodine concentration along person trajectories and optimizes with respect to thyroid activity error.

The *aim of this paper* is to develop an approach that can exploit the result of the dispersion model $F(\Theta)$ as an informative prior for the indirect path based optimization.

For the evaluation of the reconstruction quality, we use the two introduced errors: Scenario error and path error. The expectation towards the method is to be able to improve the prior scenario and reach good reconstruction quality for spacetime cells, where training path data are available. Therefore, in the calculation of the reconstruction quality, only spacetime cells are included, which are visited by at least one training path. The path error can be computed as MSE, MAE or in a relative manner analogously to the scenario error, also it has to be differentiated into training error and test error depending on the data subset used for the calculation.

3 Methodological Advancement

Method development was conducted in two phases: (i) pilot studies developing different approaches for translation and correction of prior scenarios with Cuboid Splatting, (ii) further development of most promising approaches into a final mechanism of Cuboid Splatting-with-prior and in-depth evaluation. This section contains a detailed presentation of the final Cuboid Splatting-with-prior mechanism and a short depiction of the other approaches developed in the pilot studies, that did not become part of the final mechanism.




Figure 1: Cuboids in initialization, prior and reconstruction

3.1 Final Mechanism: Cuboid Splatting-with-Prior

The final mechanism consists of an initialization, a supervised pretraining using the prior and a full finetuning. Figure 1 shows how the cuboids representing the spatiotemporal iodine distribution change during the process.

Initialization from Prior Scenario. Algorithm 1 depicts the process for Cuboid Initialization. We start from a 3D prior scenario $\hat{X}$ of iodine concentrations on a spatiotemporal grid. Cells with negligible iodine concentrations ($\hat{x}_i \leq \tau$) are discarded. From the remaining cells, we select a fixed number N of cells using k-means++ initialization algorithm [Arthur and Vassilvitskii, 2007] over the 3D coordinates, yielding a well-distributed set of locations. For each selected spatiotemporal grid cell $m \in M$, we instantiate one cuboid with center m , iodine value $\hat{x}_m$, and isotropic size (mean distance to the $k = 3$ nearest centers). This approach builds on initialization in Gaussian Splatting [Kerbl et al., 2023], where Gaussians are initialized on an SfM-point cloud [Snavely et al., 2006; Schönberger and Frahm, 2016] and scaled based on neighbor

Algorithm 1 Cuboid Initialization

Input: Spatiotemporal prior distribution $\hat{X}$, iodine threshold τ , number of cuboids N , number of neighbors k
Output: Cuboid Initialization C

```

1:  $S \leftarrow \{i \mid \hat{x}_i > \tau\}$  {Filter by iodine threshold in  $\hat{X}$ }
2:  $M \leftarrow \text{k-means++}(S, N)$ 
3:  $C \leftarrow \emptyset$ 
4: for all  $m \in M$  do
5:    $U \leftarrow k$  nearest centers to  $m$ 
6:    $d \leftarrow \text{mean\_distance}(m, U)$ 
7:    $s \leftarrow (d, d, d)$  {Isotropic initial size}
8:    $C \leftarrow C \cup \{(m, s, \hat{x}_m)\}$ 
9: end for
10: return  $C$ 

```

distances. The resulting set of cuboids constitutes the Cuboid Initialization.

Pretraining on Prior Scenario. Algorithm 2 depicts the process for obtaining the Cuboid Prior. We optimize the initialized cuboids C to reconstruct the prior scenario $\hat{X}$. Each cuboid has trainable parameters: center, iodine value and spatial extent (length, width, height). Using Adam [Kingma and Ba, 2015], we minimize the mean squared error (MSE) between the reconstructed prior scenario represented by the current cuboid configuration $\hat{X}$ and the original prior scenario $\hat{X}$. Cuboids are rendered into a 3D scenario based on the overlap volumes between the cuboids and the space-time cells (*RenderScenario*). Every $K = 20$ epochs, we perform adaptive density control (adc):

- *RemoveObsoleteCuboids*: Deletion of cuboids that have become negligible (low iodine, minimal volume, located outside relevant region)
- *AddCuboidsInHighErrorRegions*: Addition of new cuboids in under-reconstructed regions of the prior, identified via local error (*LocalErrorMap*)

Unlike in the simplified pseudocode representation, the removal and addition of cuboids are performed in separate epochs to limit the disruptive effects of the adc mechanism in training. This pretraining adapts previous Cuboid Splatting [Friedrich *et al.*, 2026] by training on the prior scenario instead of path data. The result is the pretrained Cuboid Prior.

Finetuning on Path Data. Algorithm 3 depicts the process for obtaining the Cuboid Reconstruction. We load the pretrained cuboids C and optimize all cuboid parameters jointly as a full finetuning process using the path data D . Each training path consists of a discretized spatiotemporal trajectory, associated protection factors and a measured thyroid activity. Given a batch of paths $B \subset D$, the model accumulates iodine exposure from the current state of the reconstruction $\hat{X}$ (rendered scenario from current cuboid configuration) along each path through spacetime, taking protection factors into account, to obtain reconstructed thyroid activities (*CalculatePathExposures*). We use MSE loss and Adam, now minimizing the discrepancy between measured (Y) and reconstructed ($\hat{Y}$) thyroid activity. We apply adap-

tive density control every $K = 20$ epochs, deleting obsolete cuboids (*RemoveObsoleteCuboids*) and adding new ones (*AddCuboidsAlongHighErrorPaths*). The latter adds Cuboids during finetuning by approximating underreconstructed areas. It selects the ten training paths with the largest deficit in reconstructed thyroid exposure (*PathErrorMap*) and adds Cuboids centered on five random spacetime cells along these paths. To prevent catastrophic forgetting [Kirkpatrick *et al.*, 2017] of the prior, we employ conservative hyperparameters (low learning rate, large batch size) [Li *et al.*, 2020] and learning-rate warm-up [Kalra and Barkeshli, 2024]. The result is the Cuboid Reconstruction.

3.2 Pilot Studies (Other Approaches Tested)

Next to the above described final mechanism of Cuboid Splatting-with-prior several other initialization and finetuning approaches were implemented. The approaches constituting the final mechanism were chosen because they performed best in initial benchmarks.

Two other initialization mechanisms were implemented: *random initialization*, which distributes initial cuboids evenly in spacetime with fixed shape and low iodine values, and *unsupervised pretraining*, which places cuboids at locations of clusters identified in the prior data. Both approaches performed worse than the described *pointcloud initialization*.

The other developed finetuning approaches were based on the paradigm of Parameter Efficient Fine-Tuning (PEFT) [Xin *et al.*, 2025; Han *et al.*, 2024], promising to prevent catastrophic forgetting [Kirkpatrick *et al.*, 2017] but avoid rigidity. All approaches finetune on path data and all except Side-Tuning start by loading the Cuboid Prior from pretraining. In *Parameter Groupwise Fine-Tuning* individual parameter groups (shapes, iodine values, positions) of the cuboids are finetuned (approach inspired by [Ben Zaken *et al.*, 2022]). In *Blockwise Fine-Tuning* groups of cuboids in 3D blocks of spacetime are finetuned iteratively (approach inspired by [Barakat and Huang, 2023]). In *Adapter Fine-Tuning* all cuboid parameters are frozen and instead, adapter-parameters to correct systematic errors in the prior get trained. E.g. a factor applied to all iodine values or a summand on the time-coordinate of all cuboids (approach in-

Algorithm 2 Pretraining on Prior Scenario

Input: Spatiotemporal prior distribution $\hat{X}$, cuboid initialization C , no. of epochs T , adc interval K
Output: Cuboid Prior C

```

1: for epoch = 1 to  $T$  do
2:    $\hat{X} \leftarrow \text{RenderScenario}(C)$  {Cuboids to scenario}
3:    $L \leftarrow \text{MSE}(\hat{X}, \hat{X})$ 
4:   Update  $C$ 's parameters by gradient descent on  $L$ 
5:   if epoch mod  $K = 0$  then
6:      $C \leftarrow \text{RemoveObsoleteCuboids}(C)$ 
7:      $E \leftarrow \text{LocalErrorMap}(\hat{X}, \hat{X})$ 
8:      $C \leftarrow \text{AddCuboidsInHighErrorRegions}(C, E)$ 
9:   end if
10: end for
11: return  $C$ 

```

Algorithm 3 Finetuning on Path Data

Input: Cuboid Prior $\hat{C}$, path dataset D , no. of epochs T , adc interval K

Output: Cuboid Reconstruction C

```

1:  $C \leftarrow \hat{C}$ 
2: for epoch = 1 to  $T$  do
3:   for all batch  $B \subset D$  do
4:      $Y \leftarrow \pi_4(D)$ 
5:      $\hat{X} \leftarrow \text{RenderScenario}(C)$  {Cuboids to scenario}
6:      $\hat{Y} \leftarrow \text{CalculatePathExposures}(\hat{X}, B)$ 
7:      $L \leftarrow \text{MSE}(\hat{Y}, Y)$ 
8:     Update  $C$ 's parameters by gradient descent on  $L$ 
9:   end for
10:  if epoch mod  $K = 0$  then
11:     $C \leftarrow \text{RemoveObsoleteCuboids}(C)$ 
12:     $E \leftarrow \text{PathErrorMap}(C, D)$ 
13:     $C \leftarrow \text{AddCuboidsAlongHighErrorPaths}(C, E)$ 
14:  end if
15: end for
16: return  $C$ 

```

spired by [Rebuffi *et al.*, 2017]). In *Side-Tuning* the prior scenario is loaded and frozen in its original form (3D-array). Cuboids are trained as an additional model to correct the prior scenario. For this, cuboids can have positive and negative iodine values (approach inspired by [Zhang *et al.*, 2020]).

Benchmark findings: Parameter Groupwise Fine-Tuning of only the shape parameters yielded positive results but was outperformed by Full Finetuning. Blockwise Fine-Tuning had an infeasibly high time consumption, without promising performances. Adapter Fine-Tuning yielded partly positive results but was outperformed by Full Finetuning. Current adapters were found to be under-complex for the nature of systematic errors in prior scenarios. Side-Tuning produced poor results and preventing negative values in reconstructions without impairing gradient-based optimization was difficult.

4 Experimental Setup

This section presents the scenarios, evaluation and metrics.

4.1 Original Scenarios and Prior Scenarios

The scenarios and some prior scenarios used in this research stem from the RODOS system of the German Federal Office for Radiation Protection (BfS). Two RODOS simulations (Emsland and Brokdorf) are used as the “original” scenarios to be reconstructed. These two RODOS simulations were performed with the highest quality information available. For Emsland scenario, a prior scenario exists, where the weather data and the dispersion model used in the simulation were varied (PD). For Brokdorf scenario, three prior scenarios exist: one with varied weather data (PA), one with varied dispersion model (PB) and one with a five hour earlier start of release (PC). See the supplementary materials for details on the scenarios.

Typical Errors in Atmospheric Dispersion Modeling. The highest uncertainty in atmospheric dispersion modeling

of radioactive releases (and therefore in the RODOS simulations used in this study) lies in the source term, so the question of how much radioactive material was released and when. It is very challenging to estimate even the order of magnitude of releases in cases of accidents at nuclear power plants. Next to that, the release process over time is difficult to predict too. The second highest uncertainty in atmospheric dispersion modeling originates in the intrinsic uncertainty of weather prediction data. The RODOS simulations used for this study were computed using weather data from the German numerical weather prediction model ICON at its highest resolution (ICON-D2 of the German Weather Service DWD [Reinert *et al.*, 2023]). There are limitations e.g. in a weather forecast models’ ability to accurately simulate atmospheric processes especially for small-scale turbulence beyond the minimum resolution of the model, or to set optimal initial conditions for the model runs. A third source of uncertainty in dispersion modeling is due to intrinsic assumptions in physical parameterizations and numerical approximations. All these uncertainties are interconnected; for instance, inaccuracies in the temporal variation of the source term can amplify the effects of meteorological uncertainties. [Leadbetter *et al.*, 2020]

Because of the high uncertainties, we made moderate adjustments in the parameters when simulating the priors (PA, PB and PD). These adjustments allow for a plausible improvement evaluation on RODOS simulation quality, since parameters would need to be estimated in real use too.

Prior Construction. Several types of prior scenarios have been constructed for the evaluation of Cuboid Splatting-with-prior. They cover the following types of constructed errors:

- Systematic overestimation of the scenario (P7, P9)
- Systematic underestimation of the scenario (P1, P8, P10)
- Over- and underestimation of some timesteps (P11, P12)
- Local errors in the weather data (local turbulence) (P13)
- Systematic error in the weather data (windspeed) (P14)
- Random noise in different intensities (P2, P4, P5)
- Perfect prior (original scenario) (P3)
- Completely wrong prior (interchanged scenarios) (P6)

4.2 Evaluation Plan

We analyze the two structural components of our method (prior-to-Cuboid translation; finetuning on path data), including sensitivity to different prior errors. For the two “original” RODOS scenarios to be reconstructed based on the prior scenarios and the path data, experiments have been carried out using 25,000 paths, three spatial resolutions of reconstructions (10x10, 30x30, 100x100) and all of the above mentioned RODOS and constructed prior scenarios.

Hardware & runtime: Experiments were run on an NVIDIA Ampere A40 (48GB). On this hardware, end-to-end reconstructions (initialization, pretraining, finetuning) complete within 2–3h for resolutions 10x10 and 30x30 and 12–20h for 100x100.




Figure 2: Cuboid pipeline as reference points for metrics

Our code is available under https://github.com/OpenBfS/iodine-reconstruction_cuboid-splatting-with-prior. Implementation was done with PyTorch [Paszke *et al.*, 2019] and all functions contain docstrings. Reproducibility of experiments is ensured by configurable random seeds.

Hyperparameter Optimization and Cross Validation. A hyperparameter optimization for the translation step and for the correction step were carried out to determine the hyperparameters for the final evaluation. Key hyperparameters: *Initialization*: `nr_of_cuboids=150`, `k_neighbors=3`; *Pretraining*: `initial_lr=0.1`, `patience=8`, `num_epochs=2000`, `new_half_size=1.1`; *Finetuning*: `initial_lr=0.01`, `patience=6`, `num_epochs=1000`, `adc=True`, `lr_warmup=20`. The code repository contains csv files of hyperparameter settings. A cross validation on four different 25.000 paths training splits has been carried out to create awareness of potential performance fluctuations. While reconstructions in 10x10 and 30x30 showed relatively stable reconstruction qualities, in 100x100 some fluctuations depending on the training split showed.

4.3 Quality Metrics and Baselines

Along the pipeline of Cuboid Splatting-with-prior (Figure 2), different quality metrics can be assessed. (a)-(c) constitutes translation quality and (d1)-(e) reconstruction quality of Cuboid Splatting-with-prior. (a)-(e) is baseline 1: prior scenario quality and (d2)-(e) is baseline 2: reconstruction quality of plain Cuboid Splatting. The two baselines are domain-specific state-of-the-art methods; no other established approaches exist for this task. On baseline 1: The prior scenario is a RODOS simulation, state-of-the-art in radiological protection (see Introduction). Comparison shows how Cuboid Splatting-with-prior can correct errors from RODOS simulations using path data. On baseline 2: Cuboid Splatting is the novel and currently best-performing method for reconstructing spatiotemporal iodine concentrations from path data [Friedrich *et al.*, 2026]. Comparison shows the improvement of Cuboid Splatting-with-prior over plain Cuboid Splatting. The Cuboid pipeline (Fig. 2) allows for more metrics (e.g., initialization error, Cuboid changes, per-phase progress) implemented in code but beyond this paper’s scope.

5 Results

This section presents the main result and sensitivity studies.

5.1 Main Result

When using a prior scenario that already contains useful information but also errors (PA, PB, PD, P1, P2), the reconstruction qualities in Table 1 are achieved. Overall mean reconstruction error is 19.61%. P1 are PA and PD with additional 20% source term underestimation and P2 are priors with medium random noise on original scenarios. In summary, they show per spatial resolution:

Resolution	Prior	Emsland			Brokdorf		
		CSwp	Prior	CS	CSwp	Prior	CS
10x10	PA	-	-	-	8.79	19.14	47.87
	PB	-	-	-	11.6	32.11	47.87
	PD	9.42	79.81	17.74	-	-	-
	P1	10.79	83.85	17.74	8.94	31.64	47.87
	P2	5.88	22.21	17.74	8.94	30.07	47.87
30x30	PA	-	-	-	8.91	24.12	122.25
	PB	-	-	-	12.04	38.63	122.25
	PD	33.38	80.37	40.14	-	-	-
	P1	34.42	84.3	40.14	13.81	33.2	122.25
	P2	13.13	20.27	40.14	12.88	31.46	122.25
100x100	PA	-	-	-	20.01	47.06	252.21
	PB	-	-	-	24.1	62.31	252.21
	PD	59.85	91.95	145.31	-	-	-
	P1	59.95	93.48	145.31	19.73	42.46	252.21
	P2	18.54	19.82	145.31	16.71	31.94	252.21

Table 1: Reconstruction error [%] of Cuboid Splatting-with-prior (CSwp), Prior Scenario (Prior) and plain Cuboid Splatting (CS)

- 10x10: very good reconstruction qualities of on average 9.19 % (between 5.88% and 11.6%)
- 30x30: good to very good reconstruction qualities of on average 18.37% reconstruction error (8.91% to 34.42%)
- 100x100: good or medium reconstruction qualities of on average 31.27% reconstruction error (16.71% to 59.95%)

In all experiments, the reconstructions with Cuboid Splatting-with-prior (CSwp) lead to improvements over both state-of-the-art methods (reported in percentage points, pp):

- reconstruction quality of the used RODOS simulations as prior scenarios (Prior): CSwp on average 28.02pp (between 1.28pp and 73.06pp) better
- reconstruction quality of plain Cuboid Splatting (CS): CSwp on average 89.86pp (5.72pp to 235.5pp) better

The improvement w.r.t. plain Cuboid Splatting is very high in highest resolution (100x100) and for the more difficult to reconstruct Brokdorf scenario also in (30x30), which had >100% scenario error with plain Cuboid Splatting. Generally, the better the prior scenario quality the better the improved reconstruction quality of Cuboid Splatting-with-prior. The underlying trends of improvements were also visible when using the full reconstruction for error calculation instead of filtering it to the visited cells.

5.2 Sensitivity Studies

Several dimensions were explicitly studied for sensitivity. The average *translation error* (prior scenario to Cuboid Prior) was 0.27% in spatial resolution 10x10, 1.02% in 30x30, and 3.98% in 100x100. This and the min-max range of translation errors is displayed in Table 2. Next to resolution, translation error varies depending on characteristics of the prior scenario, mainly range of iodine values and shape.

Table 3 shows sensitivity study results for the spatial resolution (30x30); results for other resolutions and further analyses are provided in the Supplementary Materials. In case of using a *perfect* (0% scenario error, P3) or *almost perfect prior*

Resolution	Emsland		Brokdorf	
	Mean	[Min; Max]	Mean	[Min; Max]
10x10	0.27	[0.05; 0.69]	0.44	[0.15; 1.02]
30x30	1.02	[0.18; 2.65]	1.26	[0.43; 3.78]
100x100	3.98	[0.85; 7.31]	2.86	[1.50; 4.61]

Table 2: Translation error Prior Scenario to Cuboid Prior in per cent

Category	Prior	Emsland			Brokdorf		
		CSwp	Prior	CS	CSwp	Prior	CS
Perfect	P3	7.86	0.0	40.14	4.86	0.0	122.25
	P4	11.71	2.03	40.14	3.22	3.15	122.25
Wrong	P5	23.12	105.96	40.14	8.16	97.1	122.25
	P6	63.86	2044.68	40.14	89.2	96.06	122.25
Source term	P7	15.86	100.0	40.14	9.42	100.0	122.25
	P8	24.25	50.0	40.14	17.55	50.0	122.25
	P9	57.98	900.0	40.14	32.03	900.0	122.25
	P10	36.85	90.0	40.14	61.11	90.0	122.25
	P11	17.21	22.39	40.14	5.77	17.95	122.25
	P12	36.28	161.04	40.14	9.56	155.13	122.25
	PC	-	-	-	60.98	72.25	122.25
	PA	-	-	-	8.91	24.12	122.25
Weather	P13	20.81	19.57	40.14	9.38	22.08	122.25
	P14	24.83	42.24	40.14	13.79	39.55	122.25

Table 3: Reconstruction quality of Cuboid Splatting-with-prior starting from specific constructed prior scenarios in resolution 30x30

scenario (2-3%, P4), applying the method Cuboid Splatting-without-prior leads to a small but clear degradation of reconstruction quality compared to prior quality. In 30x30 on average -5.62pp. Analyzing the error sources shows that a smaller part stems from translation error, while most is caused by negative training "progress". Deactivating the adc mechanism can lower the resulting error but does not prevent it.

Using a *completely wrong prior scenario* has to be subdivided: A prior scenario with approx. 100% scenario error (strong random noise, P5) can still lead to good reconstruction results. In 30x30, on average 15.64% error and significantly better performance than with plain Cuboid Splatting. This happens if the prior - despite having a 100% scenario error - still contains some useful information like a useful shape or iodine values roughly in the right range. Another completely wrong prior scenario, modeled by switching Emsland and Brokdorf as prior scenarios (P6), can lead to a much higher prior error. E.g. Brokdorf scenario encompassing much higher iodine concentrations than Emsland scenario results in a prior scenario error of >2000%. In this case, a very high improvement w.r.t. prior scenario quality is visible (+1980.82pp) when applying Cuboid Splatting-with-prior but resulting reconstruction quality can be worse than with plain Cuboid Splatting or proper reconstruction can be impeded.

Also, *source term errors* in priors can be corrected by Cuboid Splatting-with-prior. The correction of systematically over- or underestimated prior scenarios generally works, seen by improvements over prior quality. Over- and underestimations in the same order of magnitude (doubling (P7) or halving (P8) iodine concentrations) as a starting point for reconstructions still lead to very good or good reconstruction qualities (on average 16.77% error). Over- and underestima-

tions to a different order of magnitude (times 10 (P9) or a 10th (P10) of iodine concentrations) get corrected to the right order of magnitude but then systematic over- or underestimation leads to decreased reconstruction qualities (on average 46.99% error). If only some timesteps of a scenario are over- and underestimated the correction works as well. Here for over- and underestimations in the same order of magnitude (double/half, P11) the improvement over prior quality is generally lower (on average +8.68pp) as prior quality is already high. For over- and underestimations of timesteps to a different order of magnitude (P12) the correction of these timesteps' iodine concentrations to the right order of magnitude works too, improvements over prior quality are high (on average + 135.17pp), but resulting reconstruction quality maybe lower due to the high initial errors of the prior. When the prior scenario with a systematic error along time axis was used for Brokdorf scenario reconstruction (5h earlier start of release, PC), Cuboid Splatting-with-prior achieved a partial correction of this. Examination of the change in cuboid positions during finetuning showed that the model does not "systematically" correct this error by moving cuboids +5h in time, but rather corrects it through local corrections of the cuboids. Additional analyses in the Supplementary Materials.

Weather data errors can be corrected by Cuboid Splatting-with-prior as well. For Brokdorf scenario, a RODOS simulation where only the input weather data was altered (PA) was available. This prior led to a good reconstruction (8.91% error) with clear improvement over prior quality. Correcting prior scenarios with constructed local turbulence (P13) partly worked with improvement of 12.7pp to prior quality for Brokdorf scenario but -1.24pp for Emsland scenario. This might show difficulty for Cuboid Splatting-with-prior to correct this error or could be due to already high prior scenario quality at start. Correcting prior scenarios with constructed systematic wind speed error (P14) worked well. The improvement from prior quality is on average 21.59pp. Systematic error appears easier to correct than local error.

6 Discussion

Cuboid Splatting-with-prior achieves methodological advancement, experimentally supported by the achieved results.

Main Result. As the reconstructions with Cuboid Splatting-with-prior generally achieve good reconstruction quality which provide a clear benefit over solely using the RODOS simulation or solely reconstructing with plain Cuboid Splatting, this method should be used for spacetime reconstructions of airborne iodine isotopes after nuclear accidents. Especially reconstructions in high resolution (100x100), infeasible with plain Cuboid Splatting (scenario error >100%) [Friedrich *et al.*, 2026], now become possible.

Sensitivity Studies. The *translation* of a prior scenario to Cuboid Prior works well. It is easiest in lower resolution (10x10, 30x30), slightly more difficult in higher resolution (100x100). More complex characteristics of prior scenarios increase difficulty of translation. Overall, the translation mechanism successfully retains the information from the prior for the next step. With a (*nearly*) *perfect prior*,

Cuboid Splatting-with-prior yields slightly worse reconstructions than directly using the RODOS simulation. Thus, for very high-quality priors, the prior should be used as the reconstruction. If uncertain, Cuboid Splatting-with-prior should be applied and its reconstruction quality compared to the prior. In real application, this can be done using the relative path error as described below. A *completely wrong prior* with very high scenario error can yield worse reconstructions than using no prior. Detection of poor priors as well as very good priors is feasible using the relative path error (in per cent), defined as the sum of absolute differences between actual and reconstructed thyroid measurements (calculated from the prior scenario) over all paths, divided by the sum of actual measurements. A high path error indicates a poor prior. In such case, two Cuboid Splatting reconstructions, with and without the prior, should be performed and the better one selected. If both results are poor, reconstruction should be done at lower resolution to improve robustness. A very low path error indicates a high-quality prior. Reconstruction quality can be estimated using relative path error on a hold-out test path set analogously to prior quality. The method Cuboid Splatting-with-prior showed capability to correct *source term errors* and some *weather data errors*, especially more systematic ones, from prior scenarios by finetuning on path data.

Validation Gap. The main limitation of our study is the missing real world validation arising from the evaluation entirely on synthetic data—necessitated by the lack of real data—which constrains our findings by the underlying model assumptions, chosen parameterizations, and the discretization of path data and scenarios into space-time cells. The same limitation applies to the original publication on Cuboid Splatting without prior [Friedrich *et al.*, 2026].

Real World Applicability. *Reconstruction quality* is highly relevant in real-world nuclear emergency response, because thyroid screening capacity at emergency care centers is limited: Only a subpopulation can be measured, and only during the first days after an accident due to radioactive decay. Reliable spatiotemporal reconstruction enables exposure estimation for people who were missed or could not reach testing centers. Identifying elevated exposure has direct consequences for long-term medical follow-up, including regular preventive examinations due to the increased thyroid cancer risk. Thus, improving reconstruction quality is not merely a technical gain: It increases the number of people for whom elevated risk can be recognized and adequate follow-up care initiated, which can ultimately save lives. As reconstructions are primarily used for retrospective assessment and identifying high-risk individuals, not for immediate protective actions, *timeliness* of reconstruction results is not as critical as one might think at first. Faster, low-resolution reconstructions can be published after a few hours, followed by higher-resolution ones. In case of hardware restrictions, the number of cuboids in the model should be limited to still allow reconstruction. Reconstructions are updated as more path data becomes available within the days after the accident. *Scaling* e.g. to larger areas, populations or different temporal resolution, essentially affects path overlap in space time cells. Higher overlap improves reconstruction quality. If overlap is

too low, reducing reconstruction resolution is advisable.

7 Conclusion

The developed mechanism of Cuboid Splatting-with-prior successfully brings together the two data sources: RODOS simulations and path data (movement profiles and thyroid measurements). It achieves this by first translating the prior scenario into a distribution of cuboids in spacetime and secondly correcting this Cuboid Prior based on the path data. The reconstruction results achieved have good quality and lead to improvements over using solely the RODOS simulations (state-of-the-art system in German radiological protection, baseline 1) as well as reconstructing the iodine air concentrations solely from the path data with the state-of-the-art of the method Cuboid Splatting ([Friedrich *et al.*, 2026], baseline 2). Cuboid Splatting-with-prior can successfully correct different types of errors from the prior scenario. Before application of the method, the quality of the prior scenario should be estimated based on path error. Applying Cuboid Splatting-with-prior as a unified approach, considering two possibly inconsistent independent sources of information, allows the best possible estimation of the radiation exposure for the affected persons. In this framework, region-wide atmospheric dispersion estimates are complemented by individual movement-based reconstructions, enabling both detailed exposure assessment in visited locations and plausible estimates in unobserved parts of the affected area. Directions for further research are to test more original scenarios with different characteristics, to integrate further data sources like local air measurements and developing an interpolation mechanism for spacetime cells not covered by path data.

A Supplementary Materials

The supplementary materials are available here: https://github.com/OpenBfS/iodine-reconstruction_cuboid-splatting-with-prior.

1. *Code.* Python implementations of Initialization, Pre-training and Finetuning for Cuboid Splatting-with-prior.
2. *Extended Experimental Results.* Additional results tables from the evaluation experiments (PDF).
3. *Scenario Descriptions.* Details on the original and prior scenarios employed in the evaluation (PDF).

References

- [Arthur and Vassilvitskii, 2007] David Arthur and Sergei Vassilvitskii. k-means++: the advantages of careful seeding. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '07, page 1027–1035, USA, 2007. Society for Industrial and Applied Mathematics.
- [Barakat and Huang, 2023] Basel Barakat and Qiang Huang. Improving reliability of fine-tuning with block-wise optimisation. In *22nd International Conference on Machine Learning and Applications*, Florida, USA, 2023.
- [Ben Zaken *et al.*, 2022] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. BitFit: Simple parameter-efficient

- fine-tuning for transformer-based masked language-models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–9, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- [European Commission, 2022] European Commission. Commission delegated regulation (eu) 2022/1214 of 9 march 2022 amending delegated regulation (eu) 2021/2139 as regards economic activities in certain energy sectors and delegated regulation (eu) 2021/2178 as regards specific public disclosures for those economic activities, 2022. Text with EEA relevance.
- [Fridovich-Keil *et al.*, 2022] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5491–5500. IEEE, June 2022.
- [Friedrich *et al.*, 2026] Max Friedrich, Mareike Böckel, Oliver Meisenberg, Kathrin Meisenberg, and Mattis Hartwig. Learning from 3d image reconstruction: Reconstructing spatiotemporal distributions of radioactive iodine after nuclear accidents. *Science of The Total Environment*, 1011:181159, January 2026.
- [Han *et al.*, 2024] Zeyu Han, Chao Gao, Jinyang Liu, Jeff Zhang, and Sai Qian Zhang. Parameter-efficient fine-tuning for large models: A comprehensive survey. *Transactions on Machine Learning Research*, 2024.
- [Hoe *et al.*, 2009] Steen Hoe, Paul McGinnity, Tom Charnock, Florian Gering, Lars Henrik Schou Jacobsen, Jens Havskov Sørensen, Kasper Grann Andersson, and Poul Astrup. Argos decision support system for emergency management. In *12th International Congress of the International Radiation Protection Association*. Argentine Radiation Protection Society, 2009.
- [IAEA, 2024] International Atomic Energy Agency IAEA. Two years of iaea continued presence at the zaporizhzhya nuclear power plant. Technical report, 2024.
- [International Energy Agency, 2025] International Energy Agency. Energy and ai. Tech. rep., International Energy Agency, Paris, 2025.
- [Kalra and Barkeshli, 2024] Dayal Singh Kalra and Maisam Barkeshli. Why warmup the learning rate? underlying mechanisms and improvements. In Amir Globerson, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub Tomczak, and Cheng Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37, pages 111760–111801. Curran Associates, Inc., 2024.
- [Kerbl *et al.*, 2023] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), July 2023.
- [Kingma and Ba, 2015] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [Kirkpatrick *et al.*, 2017] James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A. Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, Demis Hassabis, Claudia Clopath, Dharshan Kumaran, and Raia Hadsell. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.
- [Leadbetter *et al.*, 2020] Susan J. Leadbetter, Spyros Andronopoulos, Peter Bedwell, Karine Chevalier-Jabet, Gertie Geertsema, Florian Gering, Andrew R. Jones, Heiko Klein, Irène Korsakissok, Anne Mathieu, Tamás Pázmándi, Raphaël Périllat, Csilla Rudas, Andrey Sogachev, Péter Szántó, Jasper M. Tomas, Chris Twenhöfel, Hans de Vries, and Joseph Wellings. Ranking uncertainties in atmospheric dispersion modelling following the accidental release of radioactive material. *Radioprotection*, 55:S51–S55, 2020.
- [Li *et al.*, 2020] Hao Li, Pratik Chaudhari, Hao Yang, Michael Lam, Avinash Ravichandran, Rahul Bhotika, and Stefano Soatto. Rethinking the hyperparameters for fine-tuning. In *International Conference on Learning Representations*, 2020.
- [Paszke *et al.*, 2019] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. *PyTorch: an imperative style, high-performance deep learning library*. Curran Associates Inc., Red Hook, NY, USA, 2019.
- [Raskob *et al.*, 2011] Wolfgang Raskob, Dmytro Trybushnyi, Ievgen Ievdin, and Mark Zheleznyak. Jrodos: Platform for improved long term countermeasures modelling and management. *Radioprotection*, 46(6):S731–S736, 2011.
- [Raskob *et al.*, 2016] Wolfgang Raskob, Claudia Landman, and Dmytro Trybushnyi. Functions of decision support systems (jrodos as an example): overview and new features and products. *Radioprotection*, 51(HS1):S9–S11, 2016.
- [Rebuffi *et al.*, 2017] Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. Learning multiple visual domains with residual adapters. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 506–516, Red Hook, NY, USA, 2017. Curran Associates Inc.
- [Reinert *et al.*, 2023] Daniel Reinert, Florian Prill, Helmut Frank, Michael Denhard, Michael Baldauf, Christoph Schraff, Christoph Gebhardt, Chiara Marsigli, Jochen

Förstner, Günther Zängl, Linda Schlemmer, Ulrich Blahak, and Christian Welzbacher. Dwd database reference for the global and regional icon and icon-eps forecasting system. Technischer bericht, Offenbach, Deutschland, 2023. Research and Development at DWD.

- [Schönberger and Frahm, 2016] Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4104–4113, 2016.
- [Snavely *et al.*, 2006] Noah Snavely, Steven M. Seitz, and Richard Szeliski. Photo tourism: exploring photo collections in 3d. *ACM Trans. Graph.*, 25(3):835–846, July 2006.
- [Stein *et al.*, 2015] Ariel F Stein, Roland R Draxler, Glenn D Rolph, Barbara JB Stunder, Mark D Cohen, and Fong Ngan. Noaa’s hysplit atmospheric transport and dispersion modeling system. *Bulletin of the American Meteorological Society*, 96(12):2059–2077, 2015.
- [Stohl, 2006] Andreas Stohl. Characteristics of atmospheric transport into the arctic troposphere. *Journal of Geophysical Research: Atmospheres*, 111(D11), 2006.
- [Xin *et al.*, 2025] Yi Xin, Jianjiang Yang, Siqi Luo, Haodi Zhou, Junlong Du, Xiaohong Liu, Yue Fan, Qing Li, and Yuntao Du. Parameter-efficient fine-tuning for pre-trained vision models: A survey. 2025.
- [Zhang *et al.*, 2020] Jeffrey O. Zhang, Alexander Sax, Amir Zamir, Leonidas Guibas, and Jitendra Malik. Side-tuning: A baseline for network adaptation via additive side networks. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision – ECCV 2020*, pages 698–714, Cham, 2020. Springer International Publishing.