

# Theoretical Principles of AI with a Focus on CT Dose/Dosimetry

Marc Kachelrieß


German Cancer Research Center (DKFZ)

Heidelberg, Germany

[www.dkfz.de/ct](http://www.dkfz.de/ct)

# The Big Goal: Patient-Specific Dosimetry

- Ideally consists of
  - A segmentation of the patient's anatomy
  - An estimation of the patient dose distribution
  - A risk measure that combines
    - » Organ doses
    - » Patient habitus
    - » other patient parameters



# BUILDING BLOCKS

# **TotalSegmentator: Robust Segmentation of 104 Anatomical Structures in CT images**

Jakob Wasserthal PhD<sup>1</sup>, Hanns-Christian Breit MD<sup>1</sup>, Manfred T. Meyer MD<sup>1</sup>, Maurice Pradella MD<sup>1</sup>, Daniel Hinck<sup>1</sup>, Alexander W. Sauter MD<sup>1</sup>, Tobias Heye MD<sup>1</sup>, Daniel Boll MD<sup>1</sup>, Joshy Cyriac MSc<sup>1</sup>, Shan Yang PhD<sup>1</sup>, Michael Bach PhD<sup>1</sup>, Martin Segeroth MD<sup>1</sup>

<sup>1</sup> Clinic of Radiology and Nuclear Medicine, University Hospital Basel, Basel, Switzerland

Correspondence to [jakob.wasserthal@usb.ch](mailto:jakob.wasserthal@usb.ch)

## **Abstract**

### **Purpose**


To present a deep learning segmentation model that can automatically and robustly segment all major anatomical structures in body CT images.

### **Materials and Methods**

In this retrospective study, 1204 CT examinations (from the years 2012, 2016, and 2020) were used to segment 104 anatomical structures (27 organs, 59 bones, 10 muscles, 8 vessels) relevant for use cases such as organ volumetry, disease characterization, and surgical or radiotherapy planning. The CT images were randomly sampled from routine clinical studies and thus represent a real-world dataset (different ages, pathologies, scanners, body parts, sequences, and sites). The authors trained an nnU-Net segmentation algorithm on this dataset and calculated Dice similarity coefficients (Dice) to evaluate the

# TotalSegmentator

- Neural network based fully automated segmentation of up to 104 anatomic structures in CT images
- Two models to choose from: 1.5 mm and 3 mm
- 1.5 mm higher resolution model
- 3 mm fast model with lower resolution





# A novel methodology to train and deploy a machine learning model for personalized dose assessment in head CT

Eleftherios Tzanis<sup>1</sup> · John Damilakis<sup>1</sup>

Received: 8 December 2021 / Revised: 28 February 2022 / Accepted: 19 March 2022 / Published online: 6 April 2022  
© The Author(s), under exclusive licence to European Society of Radiology 2022

## Abstract

**Objectives** To propose a machine learning-based methodology for the creation of radiation dose maps and the prediction of patient-specific organ/tissue doses associated with head CT examinations.

**Methods** CT data were collected retrospectively for 343 patients who underwent standard head CT examinations. Patient-specific Monte Carlo (MC) simulations were performed to determine the radiation dose distribution to patients' organs/tissues. The collected CT images and the MC-produced dose maps were processed and used for the training of the deep neural network (DNN) model. For the training and validation processes, data from 231 and 112 head CT examinations, respectively, were used. Furthermore, a software tool was developed to produce dose maps from head CT images using the trained DNN model and to automatically calculate the dose to the brain and cranial bones.

**Results** The mean (range) percentage differences between the doses predicted from the DNN model and those provided by MC simulations for the brain, eye lenses, and cranial bones were 4.5% (0–17.7%), 5.7% (0.2–19.0%), and 5.2% (0.1–18.9%), respectively. The graphical user interface of the software offers a user-friendly way for radiation dose/risk assessment. The implementation of the DNN allowed for a 97% reduction in the computational time needed for the dose estimations.

**Conclusions** A novel methodology that allows users to develop a DNN model for patient-specific CT dose prediction was developed and implemented. The approach demonstrated herein allows accurate and fast radiation dose estimation for the brain, eye lenses, and cranial bones of patients who undergo head CT examinations and can be used in everyday clinical practice.


## Key Points

- The methodology presented herein allows fast and accurate radiation dose estimation for the brain, eye lenses, and cranial bones of patients who undergo head CT examinations and can be implemented in everyday clinical practice.
- The scripts developed in the current study will allow users to train models for the acquisition protocols of their CT scanners, generate dose maps, estimate the doses to the brain and cranial bones, and estimate the lifetime attributable risk of radiation-induced brain cancer.

# Tzanis 2022

- **Dose with Monte Carlo (Impact MC)**
  - Based on images (training data)
  - No scan trajectory
- **Uses “training array” with 9 columns and as many rows as voxels × patients.**
  - Scan trajectory not contained in training data (only pitch value)
  - Spatial neighborhood of voxels is not contained in the “training array”. This means two voxels with identical CT value will receive identical dose prediction.
- **Neural net with 3 hidden layers to provide brain, eye, cranial bone organ doses**
- **Dose as a regression from only 8 features.**
  - CT value, scan mode, scanner model, z-pos., tube current, rotation time, collimation, pitch, dose

mAs product ( = only 1 feature, not 2)



COMPUTED TOMOGRAPHY

# A machine learning-based pipeline for multi-organ/tissue patient-specific radiation dosimetry in CT



Eleftherios Tzanis<sup>1</sup> and John Damilakis<sup>1\*</sup> 

## Abstract

**Objectives** To develop a machine learning-based pipeline for multi-organ/tissue personalized radiation dosimetry in CT.


**Materials and methods** For the study, 95 chest CT scans and 85 abdominal CT scans were collected retrospectively. For each CT scan, a personalized Monte Carlo (MC) simulation was carried out. The produced 3D dose distributions and the respective CT examinations were utilized for the development of organ/tissue-specific dose prediction deep neural networks (DNNs). A pipeline that integrates a robust open-source organ segmentation tool with the dose prediction DNNs was developed for the automatic estimation of radiation doses for 30 organs/tissues including sub-volumes of the heart and lungs. The accuracy and time efficiency of the presented methodology was assessed. Statistical analysis (*t*-tests) was conducted to determine if the differences between the ground truth organ/tissue radiation dose estimates and the respective dose predictions were significant.

**Results** The lowest median percentage differences between MC-derived organ/tissue doses and DNN dose predictions were observed for the lung vessels (4.3%), small bowel (4.7%), pulmonary artery (4.7%), and colon (5.2%), while the highest differences were observed for the right lung's upper lobe (13.3%), spleen (13.1%), pancreas (12.1%), and stomach (11.6%). Statistical analysis showed that the differences were not significant (*p*-value > 0.18). Furthermore, the mean inference time, regarding the validation cohort, of the developed methodology was  $77.0 \pm 11.0$  s.

**Conclusion** The proposed workflow enables fast and accurate organ/tissue radiation dose estimations. The developed algorithms and dose prediction DNNs are publicly available (<https://github.com/eltzanis/multi-structure-CT-dosimetry>).

# Tzanis 2025

- Dose with Monte Carlo
  - Based on images (training data)
  - No scan trajectory
- Organ segmentation with TotalSegmentator
- For each organ uses “training array with six columns” and as many rows as slices.
- NN (not specified which) learns from the training array.
- No voxel-wise dose distribution





Contents lists available at ScienceDirect

Physica Medica

journal homepage: [www.elsevier.com/locate/ejmp](http://www.elsevier.com/locate/ejmp)

Original paper

## Is deep learning-enabled real-time personalized CT dosimetry feasible using only patient images as input?

Theocharis Berris<sup>a</sup>, Marios Myronakis<sup>a</sup>, John Stratakis<sup>b</sup>, Kostas Perisinakis<sup>a</sup>,  
Apostolos Karantanas<sup>c</sup>, John Damilakis<sup>a,\*</sup>

<sup>a</sup> Department of Medical Physics, School of Medicine, University of Crete, P.O. Box 2208, 71003 Iraklion, Crete, Greece

<sup>b</sup> Department of Medical Physics, University Hospital of Iraklion, 71110 Iraklion, Crete, Greece

<sup>c</sup> Department of Radiology, School of Medicine, University of Crete, P.O. Box 2208, 71003 Iraklion, Crete, Greece



### ARTICLE INFO

#### Keywords:

Conditional generative adversarial networks  
(cGANs)  
CT  
AI  
Synthetic image generation  
Individualized medical dosimetry

### ABSTRACT

**Purpose:** To propose a novel deep-learning based dosimetry method that allows quick and accurate estimation of organ doses for individual patients, using only their computed tomography (CT) images as input.


**Methods:** Despite recent advances in medical dosimetry, personalized CT dosimetry remains a labour-intensive process. Current state-of-the-art methods utilize time-consuming Monte Carlo (MC) based simulations for individual organ dose estimation in CT. The proposed method uses conditional generative adversarial networks (cGANs) to substitute MC simulations with fast dose image generation, based on image-to-image translation. The pix2pix architecture in conjunction with a regression model was utilized for the generation of the synthetic dose images. The lungs, heart, breast, bone and skin were manually segmented to estimate and compare organ doses calculated using both the original and synthetic dose images, respectively.

**Results:** The average organ dose estimation error for the proposed method was 8.3% and did not exceed 20% for any of the organs considered. The performance of the method in the clinical environment was also assessed. Using segmentation tools developed in-house, an automatic organ dose calculation pipeline was set up. Calculation of organ doses for heart and lung for each CT slice took about 2 s.

**Conclusions:** This work shows that deep learning-enabled personalized CT dosimetry is feasible in real-time, using only patient CT images as input.

# Berries 2024

- **Dose with Monte Carlo (Impact MC)**
  - Based on images (training data)
  - No scan trajectory
- **Pix2pix image-to-image translation cGAN**
  - Converts CT image into relative dose distribution (no absolute scale)
- **Extra-trees regression to calculate missing scale**
  - Uses mA, table height, recon diameter, slice location, WED, patient widths



RESEARCH ARTICLE

MEDICAL PHYSICS

# Real-time estimation of patient-specific dose distributions for medical CT using the deep dose estimation

Joscha Maier<sup>1</sup> | Laura Klein<sup>1,2</sup> | Elias Eulig<sup>1,2</sup> | Stefan Sawall<sup>1,2</sup> |  
Marc Kachelrieß<sup>1,2</sup>

<sup>1</sup> German Cancer Research Center (DKFZ),  
Heidelberg, Germany

<sup>2</sup> Ruprecht-Karls-University, Heidelberg,  
Germany

**Correspondence**

Joscha Maier, German Cancer Research  
Center (DKFZ), Heidelberg, Germany.  
Email: [joscha.maier@dkfz.de](mailto:joscha.maier@dkfz.de)

**Funding information**

Deutsche Forschungsgemeinschaft,  
Grant/Award Number: KA 1678/24


**Abstract**

**Purpose:** With the rising number of computed tomography (CT) examinations and the trend toward personalized medicine, patient-specific dose estimates are becoming more and more important in CT imaging. However, current approaches are often too slow or too inaccurate to be applied routinely. Therefore, we propose the so-called deep dose estimation (DDE) to provide highly accurate patient dose distributions in real time

**Methods:** To combine accuracy and computational performance, the DDE algorithm uses a deep convolutional neural network to predict patient dose distributions. To do so, a U-net like architecture is trained to reproduce Monte Carlo simulations from a two-channel input consisting of a CT reconstruction and a

# Deep Dose Estimation (DDE)

- Combine fast and accurate CT dose estimation using a deep convolutional neural network.
- Train the network to reproduce MC dose estimates given the CT image and a first-order dose estimate.



# First-Order Dose Estimate

- DDE needs information about the tube position, tube current, tube voltage, shaped filters, collimation, scan trajectory etc.
- This information is encoded in the first-order dose estimate.
- First order dose-estimate in a voxel at position  $\mathbf{r}$  with mass  $m$ :

$$D_{1st}(\mathbf{r}) = \frac{1}{m} \int \frac{d^2 N}{d\Omega dE} \sum_i P_{int,i}(\mathbf{r}, E, \alpha) E_{int,i}(E) d\Omega dE$$

- Dose distribution as if there was no scatter.
- Simple backprojection of the photon energy.
- Encodes all (!) information about the CT scan.



$$P_{int,CS}(\mathbf{r}, E, \alpha) = \mu_{CS}(\mathbf{r}, E) \cdot e^{-\int_0^r \mu(r', E) dr'} \quad E_{int,PE} = E \quad E_{int,CS} = \int \frac{d\sigma}{d\Omega}(E) \Delta E_{CS}(\theta) d\Omega$$

- Our first order dose estimate assumes the object to be water-equivalent (because it turned out to work this way).


# Comments for Practical Use

- **DDE needs to use the actual**
  - x-ray spectrum (not available in DICOM)
  - scan trajectory (not available in DICOM)
- **The patient volume**
  - must laterally show the full patient cross-section
  - should longitudinally show the scanned range plus, say, 20 cm at each end.


# Results

Thorax, tube A, 120 kV, no bowtie

CT volume



First order dose




|            | MC   | DDE    |
|------------|------|--------|
| 48 slices  | 1 h  | 0.25 s |
| whole body | 20 h | 5 s    |


MC uses 16 CPU kernels  
DDE uses one Nvidia Quadro P600 GPU

DDE training took 74 h for 300 epochs,  
1440 samples, 48 slices per sample


MC ground truth



DDE



Relative error




C = 0%  
W = 40%


# Results

Thorax, tube A, 120 kV, with bowtie

CT volume



First order dose




|            | MC   | DDE    |
|------------|------|--------|
| 48 slices  | 1 h  | 0.25 s |
| whole body | 20 h | 5 s    |


MC uses 16 CPU kernels  
DDE uses one Nvidia Quadro P600 GPU

DDE training took 74 h for 300 epochs,  
1440 samples, 48 slices per sample


MC ground truth



DDE



Relative error




C = 0%  
W = 40%


# Results

Abdomen, tube A, 120 kV, no bowtie

CT volume



First order dose




|            | MC   | DDE    |
|------------|------|--------|
| 48 slices  | 1 h  | 0.25 s |
| whole body | 20 h | 5 s    |


MC uses 16 CPU kernels  
DDE uses one Nvidia Quadro P600 GPU

DDE training took 74 h for 300 epochs,  
1440 samples, 48 slices per sample


MC ground truth



DDE



Relative error




C = 0%  
W = 40%


# Results

Abdomen, tube B, 120 kV, no bowtie

CT volume



First order dose




|            | MC   | DDE    |
|------------|------|--------|
| 48 slices  | 1 h  | 0.25 s |
| whole body | 20 h | 5 s    |


MC uses 16 CPU kernels  
DDE uses one Nvidia Quadro P600 GPU

DDE training took 74 h for 300 epochs,  
1440 samples, 48 slices per sample


MC ground truth



DDE



Relative error



C = 0%  
W = 40%

# MAPE values for DDE's Organ Dose and $D_{\text{eff}}$

| Organ and ICRP weight |      | 80 kV | 100 kV | 120 kV |
|-----------------------|------|-------|--------|--------|
| Bone marrow           | 0.12 | 5.2%  | 6.7%   | 7.1%   |
| Bone surface          | 0.01 | 5.7%  | 7.0%   | 7.2%   |
| Brain                 | 0.01 | 5.1%  | 4.9%   | 5.3%   |
| Breast                | 0.12 | 1.0%  | 1.4%   | 2.1%   |
| Colon                 | 0.12 | 0.9%  | 1.7%   | 1.9%   |
| Esophagus             | 0.04 | 1.3%  | 2.4%   | 2.3%   |
| Gonads                | 0.08 | 3.2%  | 2.7%   | 2.2%   |
| Liver                 | 0.04 | 2.9%  | 1.1%   | 0.8%   |
| Lung                  | 0.12 | 1.7%  | 3.5%   | 4.0%   |
| Remainder             | 0.12 | 0.9%  | 1.9%   | 2.3%   |
| Salivary glands       | 0.01 | 4.9%  | 5.1%   | 5.3%   |
| Skin                  | 0.01 | 2.8%  | 3.3%   | 4.2%   |
| Stomach               | 0.12 | 2.3%  | 1.1%   | 0.8%   |
| Thyroid gland         | 0.04 | 3.1%  | 3.0%   | 2.3%   |
| Urinary bladder       | 0.04 | 1.7%  | 1.7%   | 1.3%   |
| $D_{\text{eff}}$      |      | 1.2%  | 2.5%   | 2.7%   |

Weighting factors and mean absolute percentage error of the DDE organ dose values with respect to the ground truth Monte Carlo organ dose values.




**FIGURE 5** Sagittal and coronal view of the dose distribution of a 100 kV whole-body spiral computed tomography (CT) scan including a bowtie filter and an angular tube current modulation. Here, the two left columns show the ground truth, the middle columns show the deep dose estimation (DDE) prediction and the two right columns the corresponding percentage error. Note that dose to air is neglected for computational reasons, and therefore, displayed as zero

# Conclusions on DDE

- **DDE provides accurate dose predictions**
  - for circle scans
  - for sequence scans
  - for partial scans (less than  $360^\circ$ )
  - for limited angle scans (less than  $180^\circ$ )
  - for spiral scans
  - for different tube voltages
  - for scans with and without bowtie filtration
  - for scans with tube current modulation
  - for DSCT scanners, i.e. with large (A) and small (B) detector
- **In practice it may therefore be not necessary to perform separate training runs for these cases.**
- **Thus, accurate real-time patient dose estimation may become feasible with DDE.**

# PUTTING THINGS TOGETHER

# CTDI Is Not Patient Dose



- CTDIvol and DLP describe the scanner output in a reference phantom, not the dose to a patient.
- Organ doses require the patient anatomy, the x-ray spectrum, bowtie, tube current modulation, and scan trajectory.
- Gold standard: patient-specific Monte Carlo simulation. Question of this talk: where can AI help, and how much physics must stay in?

# Principles That Matter for AI-Based Dosimetry

## 1 Inductive bias

- Architecture encodes assumptions.
- 3D context is needed for depth and shielding.

## 2 Physics priors

- Put what we know into the input or the problem decomposition.
- The network then only learns the residual (e.g. scatter).
- More robust and more data-efficient.


## 3 Generalization

- Networks interpolate well. They do not extrapolate well.
- Domain shift: new scanner, x-ray spectrum, trajectory, children, obesity, implants, contrast.
- Errors grow silently outside the training data.

## 4 Uncertainty


- A dose without uncertainty is not a measurement result.
- Detect out-of-distribution cases.

# How Much Physics Is Inside?



Physics in the ground truth does not count: all approaches are trained on Monte Carlo.

# Deep Dose Estimation DDE: Physics In, Dose Distribution Out




Training data: 45 whole-body patients × 3 voltages × circle/spiral × bowtie on/off × TCM on/off; test: 8 independent whole-body patients (VISCERAL).

## Two-channel 3D U-Net

- Channel 1: CT volume (anatomy)
- Channel 2: first-order dose from ray tracing (Beer–Lambert, photo + Compton)
- Encodes spectrum, bowtie, TCM, trajectory, which the CT image cannot contain
- Network learns the higher-order (scatter) part
- Works on top of any existing method to provide dose distributions (typically MC code)

# Deep Dose Estimation: Summary



**$\leq 6.3\%$**

voxel-wise MAPE vs. MC

**$\geq 91\%$**

gamma passing rate (4%, 4 mm)

**3.1%**

mean organ dose error  
(max. 7.2%, bone surface)

**1.5 s**

whole-body CT on one GPU  
(in 2021)

Maier, Kachelrieß et al. Med. Phys. 49(4):2259-2269, 2022.

# Generalization Is Bounded by the Training Data

| DDE trained on → DDE tested on    | MAPE   | Gamma pass |
|-----------------------------------|--------|------------|
| Full training data (all settings) | ≤ 6.3% | ≥ 91%      |
| 80 kV → 120 kV                    | 10.3%  | 64%        |
| Pelvis → thorax                   | 12.5%  | 52%        |
| Circle scan → spiral scan         | 15.7%  | 50%        |
| TCM off → TCM on                  | 6.7%   | 92%        |

## Lesson

- Accurate inside the training distribution
- Degrades noticeably outside
- The model does not know that it is outside
- TCM transfers well: the physics prior carries the information

Out-of-distribution behaviour (implants, rare anatomy, children) needs to be analyzed/detected.

# Other Voxel-Wise Approaches

## Salimi et al. 2023

- Density + MC dose maps of one source position in a large water cylinder, rotated+scaled
- Dose per source position, summed with tube current
- One scanner, 120 kV (90 kV by fine-tuning)

Prior: explicit, patient-independent

## Kuhlmann et al. 2026 (PTB)

- Electron density + air-kerma map of a 32 cm water cylinder, rotated and shifted
- Real + synthetic (XCAT) training data
- Uncertainty: ensembles + heteroscedastic loss

Prior: explicit, patient-independent

## Berris et al. 2024

- 2D pix2pix cGAN, single CT slice as image input
- Dose scale lost by normalization, restored by a separate regression (WED, mA, patient width, ...)
- Dose error 8.3% (max. 19.8%)

Prior: none (scale via scalars)

For comparison, DDE prior: explicit and patient-specific (first-order dose through the patient).

Salimi et al. Eur. Radiol. 33(12):9411-9424, 2023. Kuhlmann, Martin, and Pojtinger. arXiv:2601.09235, 2026. Berris, Damilakis et al. EJMP 122:103381, 2024.


# How These Approaches Differ from DDE

|                                | <b>DDE</b><br><b>Maier, Kachelrieß 2022</b>      | <b>Salimi 2023</b>                  | <b>Kuhlmann 2026</b>                  | <b>Berris 2024</b>                     |
|--------------------------------|--------------------------------------------------|-------------------------------------|---------------------------------------|----------------------------------------|
| <b>Second input channel</b>    | first-order dose through the patient             | MC dose in 71.5 cm water cylinder   | air-kerma map of 32 cm water cylinder | none (image only); scale by regression |
| <b>Prior sees the patient?</b> | yes                                              | no (distance only)                  | no                                    | —                                      |
| <b>Spatial context</b>         | full 3D                                          | 3D (slab)                           | 3D (32 <sup>3</sup> patches)          | 2D (single slice)                      |
| <b>Output</b>                  | dose of full scan or of single views             | dose per source position, summed    | 3D dose map                           | 2D dose slices, rescaled per slice     |
| <b>Network</b>                 | 3D U-Net                                         | ResNet, no pooling, dilated kernels | 3D U-Net                              | pix2pix cGAN + extra-trees regression  |
| <b>Ground truth</b>            | in-house MC (vs. Geant4)                         | MCNPX                               | EGSnrc, validated CT source           | ImpactMC                               |
| <b>Protocols in training</b>   | 80 to 120 kV, circle/spiral, bowtie y/n, TCM y/n | 120 kV, TCM (90 kV fine-tuned)      | one scanner                           | one scanner, chest CT                  |
| <b>Cites DDE</b>               | —                                                | yes                                 | yes                                   | yes                                    |

**Per-view dose is not unique to Salimi: DDE provides it too and riskTCM is built on it. The key difference is whether the physics prior sees the patient.**

Maier, Kachelrieß et al. Med. Phys. 49(4):2259-2269, 2022. Klein, Kachelrieß et al. Med. Phys. 49(7):4391-4403, 2022.

# When Spatial Information Is Lost Before the Network



- The network cannot see where an organ lies in the body, how deep it is, or what shields it.
- Such a model is closer to an extended SSDE regression than to a dose transport surrogate.
- Non-physical inputs (z-coordinate, scanner model as a category) invite shortcut learning and prevent transfer.



# Feature-Based Models: What They Achieve

|                                | Tzanis 2022 (head)          | Berris 2024 (chest)                                  | Tzanis 2025 (30 organs)           |
|--------------------------------|-----------------------------|------------------------------------------------------|-----------------------------------|
| <b>Input</b>                   | per voxel:<br>HU + protocol | 2D CT slice (normalized);<br>dose scale from scalars | per slice and organ:<br>5 scalars |
| <b>Model</b>                   | MLP                         | pix2pix cGAN +<br>extra-trees regression             | MLP per organ                     |
| <b>Spatial context</b>         | none                        | 2D only, no z-context                                | none (organ mask only)            |
| <b>Patients (train / test)</b> | 231 / 112                   | 77 / 18                                              | 141 / 39                          |
| <b>Scanners, kV</b>            | 2 (GE, Siemens), 120 kV     | 1 (GE), kV not reported                              | 1 (GE), 120 kV                    |
| <b>Error vs. MC</b>            | mean 4.5–5.7 %              | mean 8.3 %<br>(per organ 7.2–9.3 %)                  | median 4.3–13.3 %                 |
| <b>Maximum error</b>           | up to 19 %                  | up to 19.8 % (bone)                                  | up to 35 %                        |
| <b>Strengths</b>               | measured air kerma, GUI     | no scan parameters<br>needed for the image part      | open source,<br>TotalSegmentator  |

Tzanis and Damilakis. Eur. Radiol. 32(9):6418-6426, 2022. Berris et al. Phys. Med. 122:103381, 2024. Tzanis and Damilakis. Eur. Radiol. 35(2):919-928, 2025.



# From Dosimetry to Optimization: riskTCM<sup>1</sup>



$$D_{\text{eff}} = \int d\alpha I(\alpha) D_{\text{eff}}(\alpha)$$

<sup>1</sup>L. Klein, C. Liu, J. Steidel, L. Enzmann, M. Knaup, S. Sawall, A. Maier, M. Lell, J. Maier, and M. Kachelrieß. Patient-specific radiation risk-based tube current modulation for diagnostic CT. Med. Phys. 49(7):4391-4403, July 2022. This paper received the Sylvia&Moses Greenfield Award for the best scientific paper on imaging in Medical Physics in 2022.


# Reported Errors Are Not Comparable

| Study                         | Ground truth             | Metric                        | vs. measurement |
|-------------------------------|--------------------------|-------------------------------|-----------------|
| Maier, Kachelrieß 2022 (DDE)  | in-house MC (vs. Geant4) | mean voxel MAPE, <b>Gamma</b> | no              |
| Salimi 2023                   | MC                       | voxel and organ RE / RAE      | no              |
| Kuhlmann 2026                 | EGSnrc, validated source | summed organ dose, U (k = 2)  | source only     |
| Tzanis, Damilakis 2022 / 2025 | ImpactMC + air kerma     | mean / median organ PD        | air kerma only  |
| Berris 2024                   | ImpactMC                 | mean organ error              | no              |
| Scout-Dose-TCM 2025           | MC                       | mean abs. rel. organ error    | no              |

**Different codes, different metrics.**

**Validation against dosimeters in phantoms not necessary: Needs to be done for the labels.**

# AI in the Metrological Chain



each link adds an uncertainty component; the AI link is usually not quantified

## What is needed

- Reference data sets with MC and measured doses
- Benchmark phantoms incl. children, obesity, implants
- Common metrics and uncertainty reporting

## Who could provide it

- IAEA, national metrology institutes, EURAMET
- Building on TRS-457 and AAPM TG-268 (RECORDS)
- Plus AI reporting checklists (CLAIM)

# Good Practice for AI Dosimetry Papers

## Physics

State what physics enters the model, and keep the 3D context.

## Baselines

Compare with with MC. If real scanners are involved: calibrate MC (e.g. to CTDI)

## Errors

Report maximum errors and per-organ errors, not only means.

## Statistics

“Not significantly different” is not equivalence: use equivalence tests or Bland–Altman with predefined limits.

## Domain


Test on other scanners, protocols, sizes; report the domain of validity.

## Uncertainty

Give an uncertainty per prediction and flag out-of-distribution cases.

# Conclusions on AI in CT Dosimetry

- Should provide
  - a reliable segmentation of radiation risk-relevant organs or tissues
  - a fast surrogate of Monte Carlo, not a replacement of physics.
- Physics priors and 3D context make models robust; reducing the anatomy to a few scalars before the network discards information.
- Generalization and uncertainty are open issues.
- Metrology is needed: reference data, benchmarks, and uncertainty budgets for AI-based organ doses.
- AI can make CT dosimetry fast enough to become patient-specific and perhaps even prospective.



**How integrate AI in dose prediction and monitoring?**

**→ Learn from Ioannis Sechopoulos, today, 09:30, room M2.**

# Thank You!



## The 10<sup>th</sup> International Conference on Image Formation in X-Ray Computed Tomography

July 10 – July 14, 2028, Oberstdorf, Germany  
[www.ct-meeting.org](http://www.ct-meeting.org)



**Conference Chair**  
**Marc Kachelrieß, German Cancer Research Center (DKFZ), Heidelberg, Germany**

This presentation will soon be available at [www.dkfz.de/ct](http://www.dkfz.de/ct).

Parts of the reconstruction software were provided by RayConStruct<sup>®</sup> GmbH, Nürnberg, Germany.

Some slides had been drafted by Claude.