Nuclear Engineering and Technology 57 (2025) 103571

Contents lists available at ScienceDirect

Nuclear Engineering and Technology

journal homepage: www.elsevier.com/locate/net

Original Article

Simulating nuclear fuel inspections: Enhancing reliability through synthetic data

J. Knotek [a,b], J. Blažek [a,b,*], M. Kopeć [a]

[a] Research Centre Řež, Hlavní 130, Husinec, 25068, Czechia
[b] Institute of Information Theory and Automation of the CAS, Pod Vodárenskou Věží, Prague, 18208, Czechia

ARTICLE INFO

Keywords:
Synthetic data
Nuclear fuel inspection
Visual inspection
Fuel assembly inspection
Photogrammetry
Inspection algorithms
Automated inspection
Synthetic inspection videos
Nuclear fuel assembly simulation
Data anonymization
Training and validation data
Inspection cost reduction
Fuel cladding inspection
3D rendering in nuclear inspection
Operator training

ABSTRACT

Visual inspection of nuclear fuel assemblies is critical for assessing fuel reliability and ensuring safe operation. However, the sensitivity of real inspection data, along with its inflexibility and high collection costs, limits its use for research and development (R&D) tasks. These challenges hinder the ability to test and validate new inspection methodologies, making innovation slow and expensive. To address these limitations, we propose the development of synthetic nuclear fuel datasets that simulate fuel assembly inspections. These data sets replicate various defects and degradations in fuel assemblies, providing a controlled environment for hypothesis testing, operator training, and the evaluation of automated inspection techniques. Unlike real-world data, synthetic data offers the advantage of known ground-truth parameters, allowing for rigorous testing and validation. This approach enables the continuous development of inspection technologies, regardless of hardware availability and operational outages in nuclear facilities. By reducing the reliance on costly real-world experiments, synthetic data offers a scalable and flexible solution for the advancement of nuclear fuel inspection methods.

1. Introduction

Visual inspection of nuclear fuel assemblies is a key factor in long-term fuel management programs. Despite the geometric changes, usually measured by means other than the camera, that may still be visible in the images, the behavior of the fuel component surfaces plays a critical role in assessing fuel reliability. The fuel cladding is the first nonuranium safety barrier from the point of view of safe operation, and visual inspections provide an insight into its behavior during normal operation of the unit. Post-Irradiation Examination (PIE) programs take the approach of periodically inspecting at least part of the entire reactor core [1].

The primary output of nuclear fuel inspection is a video recording of the FA. Although this video directly captures the first safety barrier of the fuel, it contains a large amount of valuable data, including information on the chemical and mechanical behavior of fuel rod surfaces, structural components, geometric changes, and indirect effects such as burnup, leakages, and interactions within the core. These inspection videos are used both during inspections to evaluate the current state of the FA and after inspections to document findings, compare them with previous results, and generate reports. The combination of processed data, raw video files, and sometimes additional measurements, such as ultrasonic (UT) geometry data, creates a comprehensive database detailing how a specific fuel design behaves in particular environmental and operational conditions. However, due to the standardization of processes of a facility and the nature of the data, there is little flexibility for experimentation, driving the need for synthetic datasets that can provide a controlled environment for experimentation.

In a standard approach, the R&D actions addressing the fuel inspection programmes include a full-scale fuel mock-up, the experimental facility needs significant space, manipulators and handling equipment require materials, design, and construction. Each test also demands time for preparation and execution, leading to increased costs and workload over time. If the initial test setup is not optimal, these costs quickly escalate. A tool that can reduce these burdens during the testing or validation phase can greatly aid in developing new methods and technologies, enhancing the depth and quality of information obtained from standard fuel inspections. Such a tool would still be valuable during final tests with real fuel mock-ups (verification)

* Corresponding author.
E-mail address: blazek@utia.cas.cz (J. Blažek).

https://doi.org/10.1016/j.net.2025.103571
Received 5 November 2024; Accepted 4 March 2025
Available online 15 March 2025

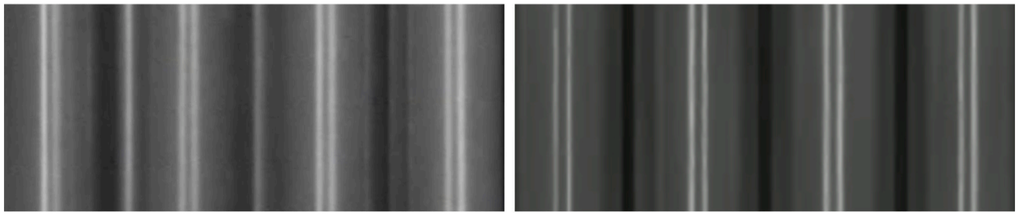

Fig. 1. Comparison of frames from real (left) and synthetic (right) nuclear fuel inspection videos. The synthetic sample does not contain any camera distortion therefore the image can be post-processed to further resemble the real-world images.

without adding extra costs, as it would help identify dead ends and suboptimal arrangements during the validation phase, allowing for necessary improvements before final testing.

In this paper, we demonstrate the process of creating synthetic inspection videos for nuclear fuel assemblies (see Fig. 1). Synthetic data sets not only offer a cost-effective and scalable alternative to real-world inspections but also provide access to known ground-truth data, which is critical for the development of precise inspection algorithms. We outline methods for generating visuals of oxidized rod surfaces, defining rod bows and fuel assembly bow at various levels, and integrating scanning protocols into the video synthesis process. To achieve realistic results, we select the photorealistic renderer Mitsuba,3 [2], which provides advanced rendering capabilities that are suited to our needs. All operations and demonstrations are conducted using a publicly available Westinghouse FA model [3], which showcases the applicability and effectiveness of our approach.

1.1. Related work

Synthetic, or procedurally generated, data has been used extensively in the entertainment industry and recently it has also been used in machine learning to enlarge or even create datasets. [4,5] demonstrate the power of synthetic data in machine learning, where procedurally generated data sets improve training and performance. Our work extends this principle to nuclear fuel inspection, providing a first step towards creating a digital twin-like environment for testing and validation. With a rough relation, the scope of the use of synthetic data in nuclear power production can be seen in the concept of Digital Twin [6]. Since synthetic fuel is not a direct concept of Digital Twin for fuels, it constitutes a first step to train the operators and optimize their work beyond the power plant itself.

Related work in the domain of nuclear fuel inspection is scarce mainly because of the reluctance to publish to avoid disclosing any manufacturer details. However, there are some exceptions. The thesis of [7] identified the benefits of constant camera speed for data processing in nuclear inspections. Our synthetic data sets build on this by introducing known ground-truth parameters, which allow for an even more efficient analysis of fuel bow and surface degradation.

Another work in this domain was published by [8] using standardized videos that allow the application of photogrammetry to extract twist and bow from three single-face videos of hexagonal FA.

This paper introduces a dataset of various FA recorded one side at a time with a fixed camera speed. The method of using a photorealistic renderer to synthesize videos is described in Section 2. The paper then follows with a description of the data sets in 3 and a discussion in Section 4.

1.2. Synthetized fuel assembly geometry

The geometric parameters of the synthetic FA were specified based on the experience of the fuel experts and the available information addressing the post-irradiation examination of the nuclear fuel as described in [1]. During the performance inside the power unit, the FA undergoes several changes, including accumulation of the surface layer (crud, oxides, etc.) and changes in geometry. The study described in [1] is a comprehensive insight into the fuel's behavior over the course of almost 10 years of operation. In this way, it gives relevant information about the overall fuel design operation. Among the parameters that are important for fuel safe operation and the scope of fuel inspections are the fuel assembly bow & twist and the single fuel rod fuel rods (FR) behavior. In favor of this article, the parameters that can directly interfere with the final scope of this study were selected, e.g., surface change implementation was limited to the basic level that represents only the burn up of fuel that corresponds with the projected FA and FR bow & twist. According to the study cited, most of the inspected FAs revealed the maximal bow of below 20 mm absolute, while the table value for the Zr-4 cladding type of fuel suggests the maximal bow on the level below 16 mm over the whole burn-up period.

According to Westinghouse data [9] on new fuel for WWER-1000 units, the nominal rod-rod gap should be at the level of $3,6$ mm in the nominal state. Closing the gap leads to a decrease in cooling abilities and a disturbance in the moderator distribution (local power changes). If exceeded, this phenomenon can lead to cladding overheating and disintegration. Good practice and the experience of the fuel experts indicate that the rod-rod gap closure of $1,5$ mm and less requires special attention during the fuel inspections. The range specified in 2 reflects this with the approximation of $1,6$ mm.

The growth of the rod is not a crucial parameter for the scope of this work but is directly related to the geometrical changes of the entire FA. Depending on the design of the fuel and its burnup, the typical rod growth changes. This phenomenon is driven by the creep rate of the cladding material and the in-rod parameters of the inner pressure. With the burnup, the gaseous fission products accumulate, and the nominal pressure increases. This imposes an axial force over the fuel rod and prompts the growth. According to the [10] report, the rod growth can reach the level of $0,83 - 1,10\%$ of nominal fuel rod length for high burnups. As the geometrical changes of the growth and bowing of the fuel in the full FA scope are combined, the reasonably and recognizable deformation of the FA will appear even in half of this scope, and the detected changes will be sufficient for this work. Furthermore, the changes in half of the level described by IAEA are more typical to the ones widely seen during the fuel inspections according to the experience of the fuel experts. This allows us to specify the maximal rod growth for this work on the level of 20 mm absolute.

2. Methods

The process of generating synthetic data follows the standard approach for creation any 3D rendered scene. It involves three key steps: modeling the geometry, assigning materials, and setting up the scene's illumination. Once these elements are properly configured, each frame is rendered individually and composed into a video. In this study, two synchronized videos were produced by simulating two cameras recording a fuel assembly (FA) from different angles at the same time. Below, the key steps are detailed. The code is available at [11].

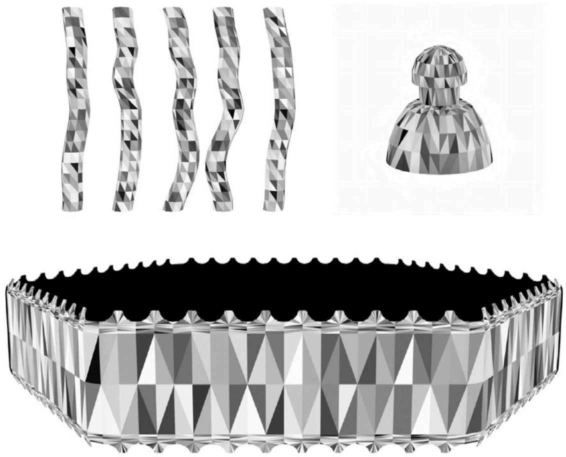

Fig. 2. The mesh of the fuel rods (left) is composed of segments that are displaced in the XY plane based on the generated rod divergence curves. The right image displays the mesh of a fuel rod tip. The bottom image shows the mesh of a spacer grid. Note that in all the mesh images, each face is assigned a random color to highlight the 3D structure.

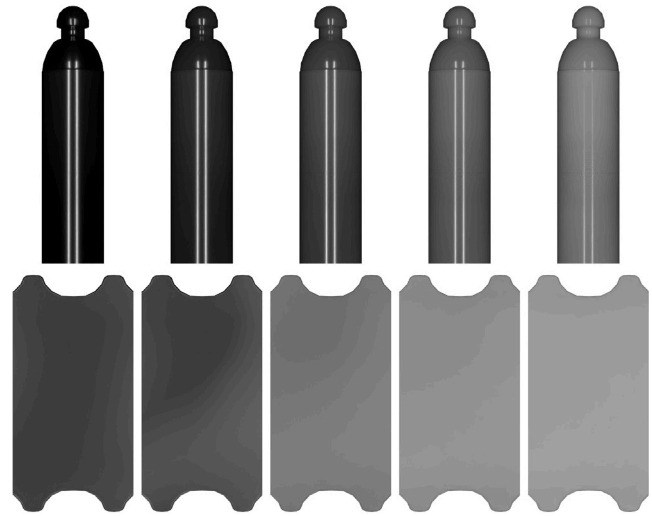

Fig. 3. Fuel rod cladding made of zirconium alloy (top) and spacer grid material made of Inconel (bottom) are shown with varying parameters to illustrate surface changes in aging fuel, ranging from the newest (left) to the oldest (right).

2.1. Modeling objects and parametrization

The generation of synthetic nuclear fuel assembly data begins with the accurate modeling of key components such as FR and spacer grids spacer grid (SG). These models are based on public blueprints and are designed to replicate the geometric and structural characteristics of real-world fuel assemblies, particularly Westinghouse [3]. The geometry of the spacer grid and fuel rod tips was created manually by replicating these blueprints.

The fuel rods are represented as a cylinder, each subdivided into smaller segments to facilitate realistic bending along their length (see Fig. 2). Segmentation allows for precise control over the deformation of the fuel rod, simulating the bowing that occurs in real assemblies due to operational stresses. The curvature of each fuel rod FR_i is defined by the bow curve $b(t)$, which models the collective bending/twisting behavior of the assembly. In addition, random deviations $d(t)$ are introduced between spacer grids to create realistic, non-uniform bow patterns for each rod. Both curves are applied when 3D FR id generated by:

$$c_i(t) = b(t) + d_i(t)$$

Additionally, fuel rod tips are shifted along the Z-axis to simulate rod growth over time, which is a known phenomenon during operation.

The spacer grid is created by first modeling a single tooth, which is a repeating geometric feature within the grid. This tooth is then duplicated and mirrored along the X-axis to form a complete grid structure for one side of the fuel assembly. This process is repeated for each side of the assembly, allowing for a seamless creation of the entire grid. Using a modular approach, this model can be easily adapted to different grid types, supporting various fuel configurations such as square or hexagonal assemblies. The modular approach is beneficial for tasks such as semantic segmentation or debris detection, and procedural generation supports different fuel rod configurations, whether square or hexagonal, without requiring significant additional effort.

2.2. Materials and oxidation

To accurately simulate the appearance of nuclear fuel assemblies, it is essential to model the material properties of key components, particularly the metallic surfaces of fuel rods and spacer grids. In our simulation, we focus on replicating the behavior of zirconium alloys and Inconel, which are commonly used materials in fuel rod cladding and spacer grids.

Incorporating realistic material properties is critical for generating high-fidelity synthetic videos. Using measured optical properties from zirconium and Inconel, we simulate the gradual oxidation of fuel rods over time, adjusting parameters such as reflectivity, roughness, and color to match real-world observations. This approach allows us to produce visuals that accurately represent different stages of fuel rod aging, providing valuable data for inspection training and algorithm development.

Over time, the surfaces of nuclear fuel rods undergo oxidation, leading to changes in their optical properties. Initially, the zirconium alloy used for cladding appears dark, smooth, and highly reflective. However, as the oxidation process progresses, the surface becomes lighter and rougher (as seen in Fig. 3, significantly altering its visual appearance. To simulate this transformation, we employ a linear interpolation between two materials: a reflective zirconium alloy [12] or Inconel [13] and a diffuse white material that mimics heavy oxidation. By applying a texture map across the model, we create an uneven surface that realistically represents the wear and aging patterns observed in real-world assemblies.

By accurately modeling surface materials and their aging processes, we ensure that the synthetic videos can serve as a realistic proxy for real-world inspections, allowing for more reliable testing and development of inspection technologies.

2.3. Illumination

Lighting plays a crucial role in the realism of synthetic fuel inspection videos. In real-world inspections, lighting conditions are often suboptimal, with variations in intensity and directionality that can obscure key features. The lighting configuration significantly affects the appearance of metallic surfaces, as demonstrated in 4. To address this, we fine-tuned the lighting parameters to achieve maximum possible realism.

The synthetic environment is illuminated using two point lights positioned symmetrically on either side of the camera. These lights are configured to simulate typical lighting conditions encountered during real-world inspections, where cameras mounted on manipulators are equipped with built-in light sources. By adjusting the intensity and color temperature of these lights, we can mimic the variability

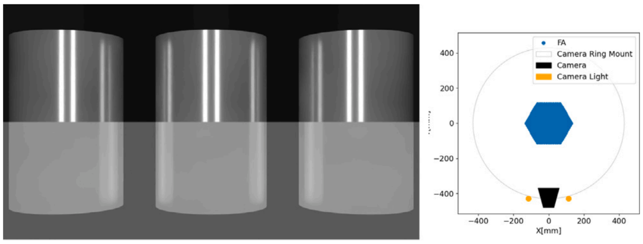

Fig. 4. The difference between local (top) and global (bottom) illumination is depicted. The dark background represents the surface of an inspection chamber, appearing nearly black with local illumination and gray with global illumination. The image highlights the significant impact that proper lighting has on the material's appearance. The local illumination setup resembles the configuration of real-world camera mounts (right) used during inspections.

Table 1
Frame Generation Time by Video Resolution for GPU NVIDIA GeForce RTX 2080 Ti.

Video Resolution	Rendering Time
750 × 600	0.8 - 1.2 sec
1920 × 1080	4 - 5 sec

found in actual inspection scenarios, ensuring that the synthetic data is representative of the lighting challenges that operators face.

To further enhance the realism of the synthetic dataset, we experimented with both local and global illumination setups. Local illumination, resembling real-world camera-mounted lights, are optimized for uniform illumination on the fuel assembly, however, creating high-contrast shadows and reflections. Global illumination, on the other hand, diffuses light more evenly across the scene, producing softer and more readable images. The comparison between these two approaches, shown in 4, highlights the critical role of proper lighting in real inspection environments.

Following these steps, the synthetic video accurately mimics real fuel assembly inspections, providing valuable data for various inspection tasks. In this place, we discuss the possibility to test light source positioning in the virtual space to achieve high-quality scans in the real world. However, for this purpose, our model is too simple, mainly due to metallic surfaces present everywhere in the fuel pool. Metallic surfaces produce a huge amount of secondary reflections that are impossible to simulate without precise scene creation. Therefore, this application is still open for further research.

2.4. Computational challenges, optimization

The computational challenges associated with this project are significant, especially in terms of optimization. Each standard video frame requires approximately 0.8 to 1.2 s of rendering time using NVIDIA GeForce RTX 2080 Ti GPU, based on preliminary observations (see Table 1). Given that the average video comprises about 3,500 frames, rendering a single video would take between 46 and 70 min. With six videos, each corresponding to a different side of the object under observation, the total rendering time reaches around 6 h. Additionally, a high-definition version of each video was rendered, which increased the rendering time to approximately 4 to 5 s per frame due to the greater pixel density and detail. This substantially increased computational demands. Fortunately, it is feasible to parallelize the rendering process across these videos, thereby considerably reducing the overall computation time.

2.5. How synthetic data can be integrated with existing inspection processes

One of the biggest issues while working on nuclear fuel inspection is the relative lack of data. The main data holder is usually the utility or the fuel vendor, and the inspection team is not always a part of either of them. This naturally limits the experience of working with the data or during the fuel inspection in general — the issues concerning fuel inspection cannot be regarded as open R&D activities so no direct educational work can be implemented or used. However, even fuel-based companies (on the same level as fuel service companies) must maintain the experience of the new generation of workers. Even in this case, the on-site crew is not prepared for everything that can be seen during the inspections, especially when new types of fuel are regarded.

The use of synthetic fuels is the answer to the market's needs. The real fuel data are a treasure and a trade secret, so those data are highly vulnerable. In the first aspect, the synthetic fuel provides total anonymization of the data and no direct relation to the special design, performance or/and unit parameters. In this way, all created data can be fully shareable and publishable. It unlocks the potential for more robust data processing and experiment preparation in favor of advanced inspection procedures. It enables collaboration with universities and also between industrial companies, so that fuel inspection issues, which were hidden beneath the fuel service, can be discussed in a wider context and problems solved around the world. Of course, the synthetic fuel will constitute here an information bearer, and only the data designer can decide on the levels of detail one will see in the data.

As fuel inspection data are limited and not fully accessible, even a prosaic task such as training new fuel inspection operators is sometimes problematic. Due to the data group and its homogeneity, the operator is oriented mainly on a FA design with a small relation to other, older types. It also reveals in a special scope of potential anomalies, so anything beyond it might be problematic for the operator. The current solution to this problem relies on long training of the operators and their work altogether with highly experienced senior employee. The knowledge and experience passing can mostly be done during the inspections themselves so the learning window is limited and narrow. However, this approach is not optimal from the point of view of the demand of the fuel inspection market. Having a tool that will help in continuous training and skilling of the personnel in the form of assisted training allows having more experts of higher expertise who will be prepared for non-standard findings because they have already seen something similar.

Finally, the main scope of this work is the perfect argument in favor of the development of synthetic fuels. Implementing any improvement in the fuel inspection policy requires R&D actions that consume manpower and money. The new approach must be tested, verified, and validated in close-to-real conditions which sometimes require building up a vast facility of infrastructure — the FA is a reasonably big component to cope with it. Supposing the first outlay of the experiment was not perfect and did not cover all required parameters, there must be implemented new changes and upgrades in the infrastructure that pull the loads even more. By changing the approach and using synthetic fuel as a model, close to the digital twin of what is happening during fuel inspection, the costs and loads can be optimized. Of course, there is still the need for a new approach validation, but the majority blind ends and errors can be solved within the modeling part of the R&D activities. This approach can also accelerate the development of new inspection technologies for the limited workload of researchers.

3. Results

In the Results section, we provide a comprehensive overview of the generated synthetic data, detailing the details of the inputs, the videos produced, the duration of their creation, and the real-world conditions they simulate. We also discuss the limitations of our generator, as revealed through various experiments conducted with the generated data set. In addition, we describe two experiments performed using this synthetic dataset, which would be exceptionally challenging to execute in real-world scenarios. Each aspect is explored in greater detail in the following subsections.

Table 2
Parameters and their respective distributions from which the generator draws a set before rendering a video. The distributions do not represent real-world measurements.

variable	distribution	what affects
Max FA Bow [mm]	Uniform(−20, 20)	FA and FR geometry
Max Fuel Rod Divergence [mm]	Uniform(−2, 2)	FR geometry
Max Fuel Rod Growth [mm]	Uniform(−10, 10)	FR length
Amount of Surface Oxidation [%]	Uniform(0.3, 0.9)	FR optical properties

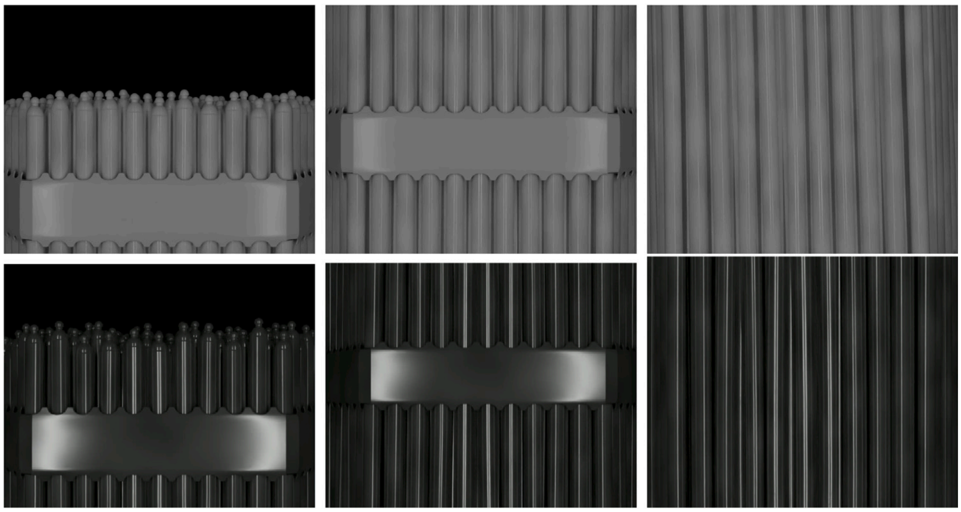

Fig. 5. Example frames from fuel with high (top) and low (bottom) levels of oxidation.

3.1. Overview of generated synthetic data

The synthetic data generator was designed to replicate a range of fuel assembly geometries and surface optical properties, allowing the simulation of diverse inspection scenarios. These simulations capture key variables such as fuel rod bow, fuel assembly bow, surface oxidation, and fuel rod growth, all of which are critical to real-world fuel assembly inspections (see Fig. 5). The generator [11] is in our data set [14] used to create seven videos per six face FA, i.e. six videos for each fuel face and auxiliary (see Section 3.2) from the top. The top camera is stationary and located 10 m above the FA top. As the video progresses to the end the distance between camera and the top is approaching 6.2 m (as the fuel is 3.2 m tall). Table 2 outlines how the geometry parameters are handled.

In the generated data set [14], the videos are synthesized using varying parameters for each FA, with each parameter sampled from its own distribution (see Table 2). The distributions for FA bow, FR bow, and FR growth are based on Section 1.2 and related papers. The percentage of oxidation then fits the visual properties of real-world videos.

To ensure the realism and accuracy of the generated videos, the synthetic data was compared against real inspection footage. A panel of experts from the nuclear industry reviewed the synthetic videos, confirming that the visual properties and geometric deformations were consistent with real-world observations. This validation step ensures that the synthetic data is suitable for both algorithm development and operator training.

The synthetic data set is designed to include information about the dimensions measured from specific experiments. In our case, this includes the bow of the fuel assembly as a continuous function, the swinging of the FA during scanning as a function, and rod distances (also represented as a function of time). Using functions allows us to calculate precise values for each pixel. All ground-truth values are shown in Fig. 6.

Additionally, the model parameters are not the only important aspect. The generator parameters also encompass camera and light source positions. Adjusting the acquisition time t, we can model synchronization issues within the scanning system.

3.2. Use case: Proof of concept prior to experimentation

In this use case, we demonstrate how synthetic data can be used to test and validate an extended dataset containing not only FA face videos but also top view. This experiment simulates real-world inspection scenarios, in which fuel assemblies are suspended on ropes and experience motion during visual inspections. The goal of a new approach is to use Digital Image Processing (DIP) algorithm for swinging trajectory extraction from the top view camera and to use it to stabilize the videos of faces (side views). From the currently used camera systems working with pulled (hanged) FA it is impossible to measure the bow or twist of FA, because such cameras are typically missing in Power plant (PP) facilities. We would like to demonstrate that an additional synchronized camera can significantly improve the outcomes of visual inspections, such as the photogrammetry of fuel bows. In comparison with other methods for bow measurement, this approach is cheaper and faster.

Conducting the same experiment in the real world poses several challenges related to additional camera placement. First, it requires special permission from the facility. Second, the scene is more complex, featuring elements such as cranes, ropes, and other objects that can obstruct the view. Third, creating a proof of concept that demonstrates the achieved quality will be complicated due to the intricate behavior of the scene, which includes varying lighting conditions and the optical properties of the fuel's top surface. Finally, the absence of ground truth further complicates the evaluation of results.

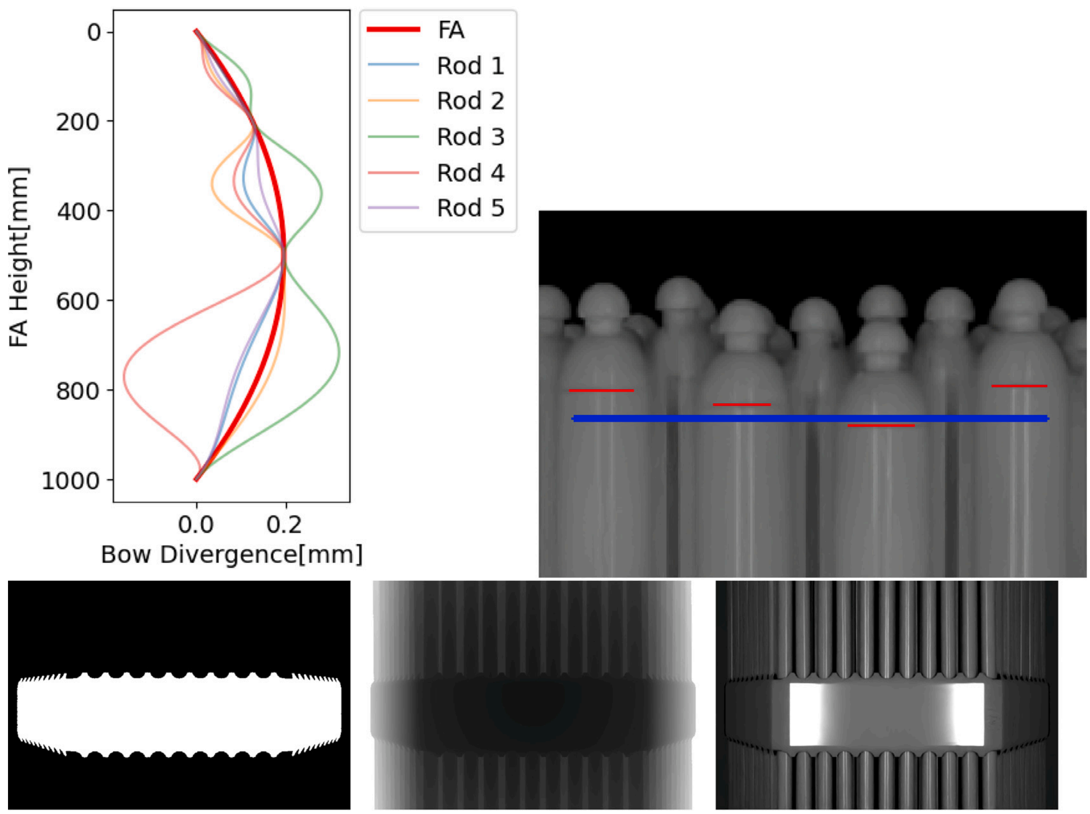

Fig. 6. Ground truth for various aspects of the FA includes: (a) Fuel bow and rod divergence curves; (b) Growth of individual rods (only five are shown); (c) Spacer grid mask (used for semantic segmentation) derived from frame (d); (d) Depth map corresponding to the same spacer grid.

Our application of synthetic data is to model an experiment focused on calculating the swinging trajectory of FA hanging on a rope. Input data is in a form presented in Fig. 7. The side view video is a pass along a FA side (one video per each pass). The top view is a video made from a static camera and moving FA (e.g., the top is either progressively closer).

The swinging trajectory can be understood as a moving center of the hexagon tracked across the video and normalized by the distance to the camera. The algorithm developed to extract angle of swinging works as follows.

- For each frame i, fit a circle with center c_i and radius r_i to the hexagonal outline of the FA's top.
- Calculate trajectory offsets $o_i = c_i - c_m$ where c_m is the mean center.
- Normalize offsets to account for changing distance from the camera. Assuming that the ratio between the radius of FA rf (in mm), the trajectory offsets t_i can be converted from the image domain to the spatial using formula $t_i = \frac{o_i}{r_i} * r_f$.

The last step is to convert the trajectory offsets to angular units $a_i = tan^{-1}(t_i/d)$ where d represents the maximum distance to FA. The algorithm was tested on synthetic top-view videos generated with known ground-truth parameters for the swinging trajectory. The results show that the algorithm successfully tracked the trajectory with a mean error of less than $0.002°$ compared to the ground-truth swing, as shown in Fig. 8. This demonstrates that a synchronized additional top-view camera can, under optimal conditions, enhance an already acquired dataset, particularly when FA scanning is performed while pulling the FA on a rope. This approach enables valuable photogrammetry to be conducted on the dataset.

It needs to be said that developing an algorithm against the synthetic data does not warrant success in using the real one. In this case, it is a tool how to properly model an experiment in the real world. Using synthetic data, we were able to test the approach in a controlled environment, adjusting variables such as camera position, lighting conditions, and the fuel assembly's swinging patterns. This flexibility allowed us to fine-tune the algorithm before conducting real-world experiments, saving both time and costs. In real-world scenarios, conducting this type of experiment would require special permissions and be subject to significant operational constraints.

This experiment highlights the potential of synthetic data as a tool for validating algorithms in the early stages of development. By providing a known ground truth and a controlled testing environment, synthetic data enables more efficient and accurate experimentation, reducing the need for costly and complex real-world trials:

- It provides a better understanding of the parameters of the experiment.
- The generated video serves as a valuable tool for conducting the experiment.
- A prototype algorithm for processing the video can be developed.
- Thanks to the existence of ground truth in synthetic dataset, we are able to evaluate precision of a new approach in advance.

The practical impact of this use-case is a possible upgrade of a facility scanning process by demonstrating the principle on a synthetic dataset. The use-case helped to identify the potential problems, and this way allows us to better estimate risks and costs of real-world trial.

3.3. Use case: Evaluation of the photogrammetry algorithm

Synthetic data generated for this purpose is shared with the public [14]. The code used for the generation of these data is available at [11].

In this use case, we evaluated a previously developed algorithm [8] for measuring the bow of fuel assemblies using photogrammetry. In

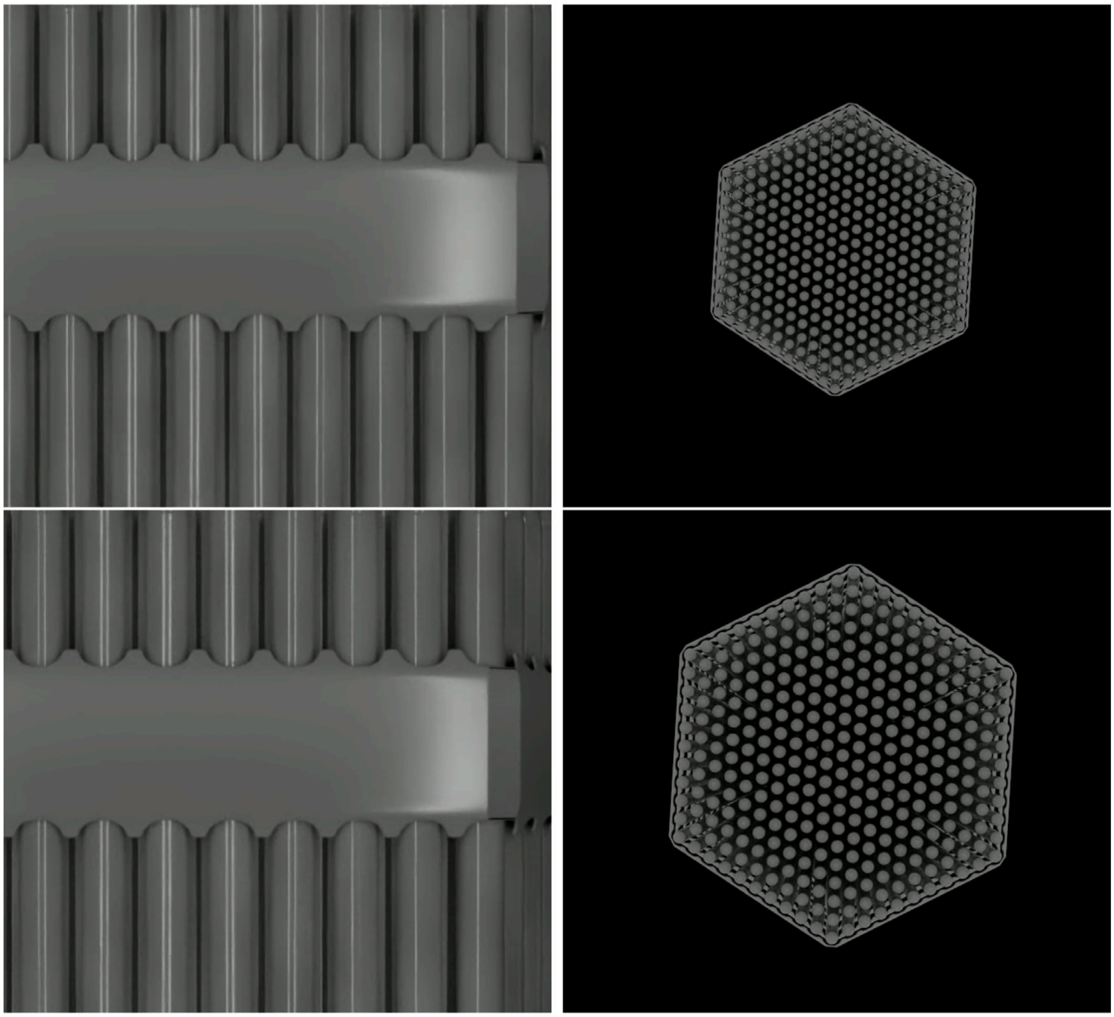

Fig. 7. The figure shows frames from the two types of auxiliary videos. The left column shows two frames capturing the same spacer grid where the top comes from the main video and the bottom is auxiliary. There is a noticeable shift in X axis that hints sideways movement. The right column shows two frames from the vertical type of auxiliary videos — one from the beginning and the other from the end (both in close up). As the fuel moves upwards, the hexagonal top approaches the camera. Notice that neither fuel head nor any rope was rendered.

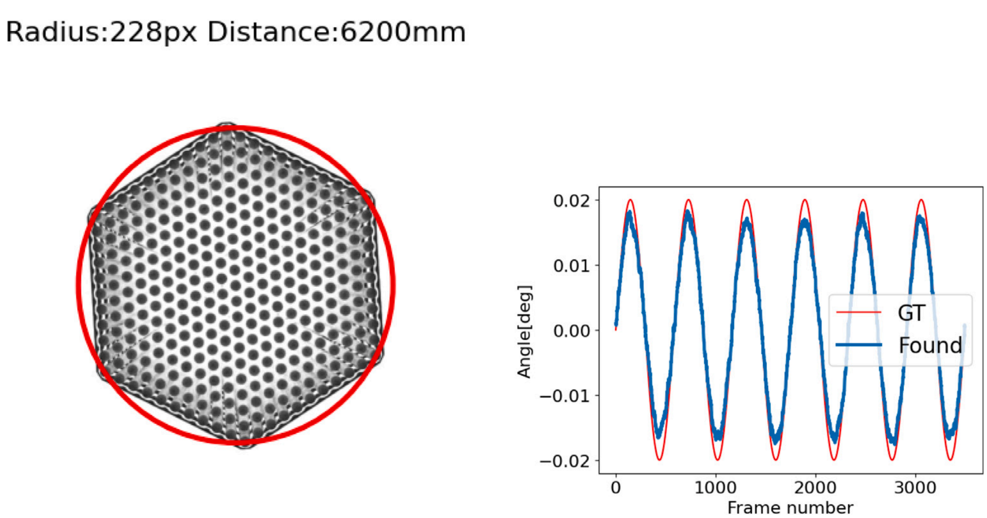

Fig. 8. The figure shows circle around FA's top (left) and swinging angle compared to ground-truth (right). The close match between the two demonstrates the accuracy of the synthetic dataset in replicating real-world fuel assembly movements.

contrary to previous use-case (Section 3.2) the algorithm works with fuel fixed during scanning (so there is no swinging and rotation of the FA during acquisition). However, the absence of ground truth data in real-world inspection videos made it challenging to assess the

Table 3
The results of bow measurement on synthetic dataset. Every column corresponds to one FA modeled. Blur is the best blur used in Devernay edge detector. Max abs err (mm) row show the maximal absolute error of a frame in millimeters from the ground truth. Mean and standard deviation are computed over all FA sides. Last three rows compare the absolute values above with the bow of the FA side (100%).

	new	semiold	old
blur	1.1	1.0	0.9
max abs err (mm)	0.20	0.61	1.37
mean abs err (mm)	0.12	0.38	0.75
std abs err (mm)	0.05	0.23	0.31
max (%) err	38.87	19.66	25.60
mean (%) err	25.65	16.96	16.41
std (%) err	9.73	1.77	6.44

algorithm's accuracy and identify sources of error. To address this, we now used synthetic data with known ground-truth parameters for fuel bow, providing a reliable baseline for comparison.

The photogrammetry algorithm works as follows:

1. apply edge detector [15] with blur parameter
2. every edge having more than 20 points is approximated by a line (least squared error method) and calculate its direction (angle) and square error of approximated points.
3. use weighted average for an estimate of frame rotation. Weight is computed as follows:

$$w(P) = \frac{|P|}{\sum_{p \in P}(p_x - \hat{p}_x)^2 + (p_y - \hat{p}_y)^2},$$

   where w is weight of a line approximated by points $p \in P$. p_x, p_y are real coordinates of the point p. And $\hat{p}$ is the closest point to p laying on the line.

FA bow is then computed as a x-shift for every frame from estimated rotations:

$$f_x(n) = \sum_{i \in \{0, \cdot, n\}} \arctan(\alpha_i - \overline{\alpha}),$$

where $f_x(n)$ denotes x-shift of the nth frame, α_i is angle of ith frame and $\overline{\alpha}$ is mean angle of the whole frame sequence (typically corresponding to a camera lean).

The use of synthetic data addresses the ground-truth issue. It also allows us to modulate input quality, allowing us to better define the limits of the algorithm and technology. Furthermore, synthetic data are ideal for testing both the algorithms and the operators.

The algorithm testing process is as follows:

1. Generate synthetic data within specific parameter ranges.
2. Photogrammetry is performed on the generated video.
3. The measured results are compared with the ground-truth.

We evaluate 3 fuel assemblies new one, semi old and old one. Results of the bow measurement are summarized in Table 3

So far, we have tested our algorithm for FA bow estimation, as published in [8]. This experiment demonstrates that the complexity of the synthetic data is relatively low and, compared to real-world data, is easier to process (e.g., there are no issues with threshold computations for edge detectors). Despite this, the experiment confirmed validity our approach to FA bow estimation and identify limits for our initial assumptions.

The current generator is capable of creating an effective platform for testing algorithms for fuel bow measurement and their future improvements. Practically, we are now able to improve algorithms without the need of operators, their manual evaluations, and new data from a PP. This improves inspection accuracy and speeds up the development of operator software tools.

4. Discussion

The R&D activities in favor of visual inspection of nuclear fuel face significant challenges due to the lack of large and diverse datasets. This is largely because data acquisition in these environments is inherently difficult, sensitive, and cannot be in raw format shared between facilities. As a result, each facility is forced to rely on its own limited data, which is often slow and costly to obtain. In addition, experiments are constrained by the need to avoid disrupting operations during power plant outages. Synthetic data offers a solution by replicating the complexity and variability of real-world inspection videos, providing a more efficient and scalable alternative to traditional data collection methods.

Our goal was to develop a synthetic video generator that simulates visual inspection data by focusing on several key aspects. We address:

- material optical properties of zirconium alloy and Inconel,
- changes in optical properties due to surface oxidation,
- geometric variations such as rod bow and fuel bow.

. Additionally, we account for different scanning protocols, including fixed fuel with a moving camera, moving fuel with a stationary camera, and both in motion. This approach generates sufficient variability for creating a dataset that is valuable for e.g. verifying geometry reconstruction algorithms and also enables deepening collaboration among experts without worrying about disclosing any details and compromising security.

4.1. Limitations of synthetic data

All demonstrated use-cases are valid with current state of the generator. But we see a huge potential in synthetic videos on many new levels, e.g. in usage for:

- Artificial Intelligence (AI) training (segmentation, anomaly detection, preprocessing)
- Operator training
- Designing of new scanners
- Evaluation of acquisition processes
- Simulation of fuel handling

All of these steps lead to faster acquisitions, reproducible measurements, redundancy, and robustness. Application of better fuel inspections then increases safety and prolongs the fuel lifespan.

However, for this very wide range of applications, the current model requires more context. That is, the current model does not contain all scene-specific objects (such as pool walls, ropes, handling crane parts, water, other FAs and pool structures interacting with the light rays). In addition, the generator does not cover various fuel designs, aging beyond oxidation, physical issues like missing parts or cracks, or other experiments-specific requirements. The context needs to be built according to the particular use-case/experiment requirements, and there is a potential to do so.

In the follow-up work we would like to improve the FR surface modeling and oxidation patterns for the use in AI training. This highlights the need for more advanced techniques to accurately capture and replicate the diverse and intricate oxidation patterns found in actual fuel assemblies.

High fidelity outcomes have also negative impact on performance. At this moment, it takes approximately one second[1] to render a single video frame as described within this paper.

[1] Using NVIDIA GeForce RTX 2080 Ti.

4.2. Future work

This is an initial approach that paves the way for more effective experimentation, with the potential to incorporate more real-world conditions in future work. Looking ahead, we envision a future with fast, photorealistic rendering software and a shared database of materials, textures, and models to speed up the groundwork that must be done to prepare any experiment. Future research will focus on expanding the diversity of fuel designs represented in the synthetic dataset, incorporating more complex aging phenomena such as cracking and missing parts. In addition, we plan to integrate more advanced machine learning algorithms to further unravel the synthetic and real fuel inspection videos without affecting the data sensitivity.

Synthetic data promise to overcome critical issues with the sharing of data related to nuclear fuel. Our synthetic data generator represents the initial effort to facilitate discussions on processing FA fuel inspection videos and to aid automating their analysis or updating the scanning protocols.

Abbreviations

The following abbreviations are used in this manuscript:

AI Artificial Intelligence

CVR Centre of Research Řež

FA nuclear fuel assembly

FR fuel rods

SG spacer grid

OIO One image overview

DIP Digital Image Processing

BSDF Bidirectional Scattering Distribution Function

PIE Post Irradiation Examinations

PP Power plant

UT ultrasonic

CRediT authorship contribution statement

J. Knotek: Conceptualization, Data curation, Methodology, Software, Validation, Writing – original draft. **J. Blažek:** Supervision, Validation, Writing – review & editing. **M. Kopeć:** Data curation, Funding acquisition, Project administration, Resources, Writing – original draft.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Jan Blazek reports financial support was provided by Czech Academy of Sciences. Marcin Kopec reports financial support was provided by Ministry of Industry and Trade of the Czech Republic. Jan Blazek and Marcin Kopec has patent A method of measuring fuel assembly deformation using digital image processing from fuel inspection videos licensed to Centrum výzkumu Řež s.r.o., Hlavní 130, 250 68 Husinec, Řež, Česká republika and Ústav teorie informace a automatizace AV ČR, v.v.i., Pod vodárenskou věží 1143/4, 182 00 Praha 8, Libeň, Česká republika. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The presented work has been carried out with the institutional support of the Ministry of Industry and Trade of the Czech Republic. We are also grateful for national support provided by the Czech Academy of Sciences through the program Strategy AV21.

Data availability

All data generated for the purpose of the described scenarios is available in our HuggingFace repo [14] under the GPLv3 license. The synthetic FA generator is also available on Github [11] under the GPLv3 license.

References

[1] D. Ernst, L. Milisdörfer, 10 years of experience with westinghouse fuel at NPP Temelín, in: International Conference VVER-2010, Prague, 2010.

[2] W. Jakob, S. Speierer, N. Roussel, M. Nimier-David, D. Vicini, T. Zeltner, B. Nicolet, M. Crespo, V. Leroy, Z. Zhang, Mitsuba 3 renderer, 2022, https://mitsuba-renderer.org.

[3] H.S. M. Dye, Enhanced Westinghouse WWER-1000 Fuel Design for Ukraine Reactors, Westinghouse Electric Company LLC, Hopkins, Columbia, USA, 2016.

[4] S.R. Richter, V. Vineet, S. Roth, V. Koltun, Playing for data: Ground truth from computer games, 2016.

[5] E. Wood, T. Baltrušaitis, C. Hewitt, S. Dziadzio, M. Johnson, V. Estellers, T.J. Cashman, J. Shotton, Fake it till you make it: Face analysis in the wild using synthetic data alone, 2021.

[6] A.K. Sleiti, J.S. Kapat, L. Vesely, Digital twin in energy industry: Proposed robust digital twin for power plant and other complex capital-intensive large engineering systems, Energy Rep. 8 (4) (2022) 3704–3726.

[7] J. Knotek, Automation of Nuclear Fuel Visual Inspection, (Master's thesis), Charles University, 2020.

[8] D. Plašienka, J. Knotek, M. Kopeć, M. Malá, J. Blažek, Measurement of nuclear fuel assembly's bow from visual inspection's video record, Nucl. Eng. Technol. 55 (4) (2023) 1485–1494.

[9] W. LLC, VVER-1000 Fuel Products, 2022.

[10] V. Krasnorutsky, Progress report: Fuel rod performance evaluation of CE 16×16 operated at steady state using TRANSURANUS and PAD codes, in: Investigation of PWR and VVER Fuel Rod Performances under High Burnup Using FEMAXI & PAD Codes, no. 15370, 2009.

[11] J. Knotek, Synthetic - nuclear fuel assembly - generator, 2024, https://github.com/jaroslavknotek/synth-nfa-generator. (Online Accessed 5 November 2024).

[12] M.N. Polyanskiy, Refractive index database, 2024, https://refractiveindex.info/?shelf=main&book=Zr&page=Palm. (Online Accessed 5 November 2024).

[13] W.V. Goodell, J.K. Coulter, P.B. Johnson, Optical constants of inconel alloy films, J. Opt. Soc. Am. 63 (2) (1973) 185–188.

[14] J. Knotek, Nuclear fuel inspection - synthetic data, 2024, https://huggingface.co/datasets/CVRez/nfa-inspection-synth. (Online Accessed 5 November 2024).

[15] R. Grompone von Gioi, G. Randall, A sub-pixel edge detector: an implementation of the Canny/Devernay algorithm, Image Process. Line 7 (2017) 347–372, http://dx.doi.org/10.5201/ipol.2017.216.