Generative models and simulation to assess uncertainties for tokamak infrared thermography

Alexis Juven, Marie-Hélène Aumeunier, Julien Marot

To cite this version:

Alexis Juven, Marie-Hélène Aumeunier, Julien Marot. Generative models and simulation to assess uncertainties for tokamak infrared thermography. MLSP 2024 - IEEE 34th International Workshop on Machine Learning for Signal Processing, Sep 2024, Londres, United Kingdom. ⟨10.1109/MLSP58920.2024.10734728⟩. ⟨cea-05108308⟩

HAL Id: cea-05108308
https://cea.hal.science/cea-05108308v1

Submitted on 11 Jun 2025

2024 IEEE INTERNATIONAL WORKSHOP ON MACHINE LEARNING FOR SIGNAL PROCESSING, SEPT. 22–25, 2024, LONDON, UK

GENERATIVE MODELS AND SIMULATION TO ASSESS UNCERTAINTIES FOR TOKAMAK INFRARED THERMOGRAPHY

Alexis Juven⋆ *Marie-Hélène Aumeunier*⋆ *Julien Marot*†

⋆Institute for Magnetic Fusion Research (IRFM), CEA, Cadarache, Saint-Paul-lez-Durance, France
†Aix Marseille Université, École Centrale Méditerranée, CNRS, Institut Fresnel, UMR 7249, Marseille, France

ABSTRACT

Quantitative infrared thermography, which aims to infer surface temperatures from an infrared image, is challenging in metallic environments. A local increase in the flux measured by the camera may be due to indiscriminately an increase in the temperature of the observed target or due to a change in its surface properties. This paper addresses this ill-posed inverse problem and provides a method to automatically evaluate a confidence interval on the predicted temperature. For this purpose, we use simulated data, generated with ray tracing, to train a deep generative model able to approximate an unknown probability distribution. We apply the proposed method to simulated infrared images of fusion facilities, where infrared thermography is a critical diagnostic for machine protection. We show that the model estimates the temperature with a 6% mean error, and provides a 95% confidence interval, on average 20% of the predicted value, improving the reliability of the measurement.

Index Terms— cWGAN-GP, Inverse Problem, Infrared Imaging, Ray Tracing, Nuclear Fusion

1. INTRODUCTION

Infrared imaging is a critical tool used to monitor in real-time the wall temperature of nuclear fusion devices. These systems are used in existing facilities such as WEST [1] (see Fig. 1), and for the future international fusion device ITER [2].

However, mapping the observed infrared images to the true surface temperatures exhibits significant challenges. Measurements are deteriorated by parasitic disturbances, such as a flux resulting from reflections on the metallic surfaces of variable reflectivity and emissivity. Such inaccuracies in temperature estimation are a limitation for an effective machine protection, leading to risks of either underestimating temperatures, potentially causing damage, or overestimating them, which might lead to unnecessary shutdowns.

Recent work has shown promising results by using machine learning to estimate surface temperatures from infrared images [3] . Similar to modern methods of inverse rendering, which aims to recover parameters of a scene from a captured picture [4], the approach involves generating a large dataset of simulated images, which is then used to train a convolutional neural network in a supervised manner, which would not be possible with experimental data because of the absence of a labeled ground truth. The network learns a mapping between the infrared image and the actual surface temperatures, despite not being provided with actual experimental values of emissivity and reflectivity for the surfaces.

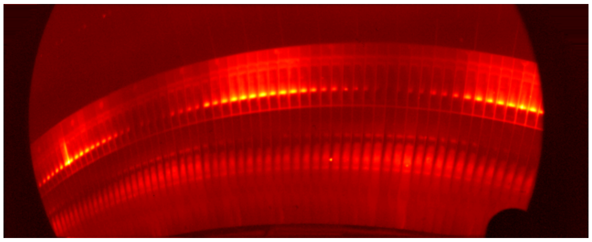

Fig. 1: Experimental image captured by one of the infrared camera focusing on the lower divertor of the WEST tokamak during plasma operation.

However, this study has shown its limitations in relying on deterministic models. Determining temperature from images is fundamentally an inverse problem, subject to ill-posedness: for a given input image, there might be a wide range of possible temperature solutions due to information loss. A deterministic model cannot account for all these possibilities, and tends towards an averaged answer over potential solutions [5]. In inverse rendering, an additional fine-tuning phase using a differentiable renderer avoids returning an average of all possible values and instead selects one plausible solution [6]. However, such an approach is not sufficient to ensure wall protection in nuclear fusion devices, where an estimation of the measurement uncertainties is required to make a decision on the subsequent operations (power decrease or shutdown).

In this study, we explore the usage of generative models providing a probability distribution of surface temperatures. This approach allows the model to account for both uncertainty as well as the potential multiplicity of solutions where the data lack information. Section 2 describes the ill-posed problem in the particular case of temperature measurement from thermography systems. The implemented model, based on the cWGAN-GP, an architecture currently used to solve inverse problems [7], [8], is described in section 3. A first evaluation of the temperature estimation model is performed in section 4, before it is applied to simulated tokamak data in section 5.

2. TEMPERATURE MEASUREMENT FROM INFRARED IMAGES: AN ILL-POSED PROBLEM

Determining surface temperatures from infrared images is a complex inverse problem involving unknown factors like scene geometry, surface emissivity, reflectivity, and camera measurement disturbances.

Infrared cameras in tokamaks capture light from observed surfaces, modeled as an outgoing radiance L^{out} leaving a surface point $\mathbf{p}$ in the direction $\boldsymbol{\omega}_o$. The outgoing radiance results from a combination of emitted radiance, which depends on the surface emissivity ε and temperature T, and parasitic reflected radiance, which depends on the surface bidirectional reflectivity ρ and incoming radiance L^{in}, integrated over all possible directions $\Omega_{\mathbf{p}}$. This study focuses on extreme scenarios where surface reflections are purely diffuse, leading to a maximal information loss [9], making inversion challenging. Under these conditions, the infrared rendering equation is expressed by eq. (1). $L^{\circ}(T)$ is the radiance emitted by a blackbody of temperature T in the camera wavelength range (in our case 3.8 to 4.0 µm [1]), an increasing function of T obtained by using Planck's law.

$$L^{\text{out}}_{\mathbf{p},\boldsymbol{\omega}_o} = \varepsilon_{\mathbf{p}} L^{\circ}(T_{\mathbf{p}}) + \rho_{\mathbf{p}} \int_{\Omega_{\mathbf{p}}} L^{\text{in}}_{\mathbf{p},\boldsymbol{\omega}_i} \cos\left[\theta_i\right] \mathrm{d}\omega_i \quad (1)$$

In this case, the captured image might not contain sufficient information to determine a precise surface temperature value. To address this challenge, our approach involves using Bayesian probability distributions to deal with uncertainty.

3. GENERATIVE MODELS FOR INFRARED INVERSE PROBLEMS

This section describes a generative model to provide probabilistic temperature outputs from a single infrared image, enabling an estimation of measurement uncertainty.

3.1. Bayesian inference

In inverse problem solving, Bayesian inference is a classical probabilistic approach to handle uncertainties and ill-posedness. Considering the random variables X (parameters to be estimated, e.g., temperature and/or emissivity), Y (measured value, e.g., image in radiance), and N (measurement noise), these variables are linked by a deterministic forward operator $\mathcal{F}$, as expressed in eq. (2).

$$Y = \mathcal{F}(X) + N \quad (2)$$

Knowing the prior probability distribution function p_X, p_Y and p_N of the three variables X, Y and N respectively, Bayes' theorem provides the posterior distribution $p_{X|Y}$ expression by scaling the prior distribution with the likelihood of the measurement, as shown in eq. (3).

$$p_{X|Y}(\mathbf{x}|\mathbf{y}) = \frac{p_N(\mathbf{y} - \mathcal{F}(\mathbf{x}))p_X(\mathbf{x})}{\int_{\mathcal{X}} p_N(\mathbf{y} - \mathcal{F}(\mathbf{x}'))p_X(\mathbf{x}')\,\mathrm{d}x'} \quad (3)$$

In our study, we aim to obtain a sampling method that accurately follows the posterior probability distribution $p_{X|Y}$, in order to estimate a range of uncertainty for the estimated parameters after a specific measurement. While the theoretical expression of this distribution can be derived, its practical application for sampling is difficult. Surrogate models approximating the desired distribution, such as deep generative models, can be used. This paper demonstrates the usage of such models to estimate uncertainties in temperature measurement from infrared images.

3.2. Deep generative models

In our study, we use neural networks, more specifically generative models, as a method for sampling from the posterior distribution $p_{X|Y}$. The core idea is to introduce an additional input, a simple random variable Z of known distribution p_Z, and to train a model $\mathcal{M}$ to approximate the complex posterior distribution $p_{X|Y}$ when being fed samples from Y and Z, as expressed in eq. (4).

$$p_{\mathcal{M}(Y,Z)|Y}(\mathbf{x}|\mathbf{y}) \simeq p_{X|Y}(\mathbf{x}|\mathbf{y}) \quad (4)$$

Various generative model architectures have been explored, like Variational Autoencoders (VAEs) [10], Generative Adversarial Networks (GANs), or diffusion models [11]. In our case, we opted for a conditional GAN variant which ensures a stable convergence and fast post-training sampling.

3.3. cWGAN-GP

The Conditional Wasserstein GAN with Gradient Penalty (cWGAN-GP) is a kind of GAN constituted by a generator $\mathcal{G}(Z, Y)$ and a discriminator $\mathcal{D}(X, Y)$ conditioned by an additional input Y. The goal is to map the latent variable $Z \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ into the posterior distribution $p_{X|Y}$, by adversarially maximizing on $\mathcal{G}$ and minimizing on $\mathcal{D}$ the loss defined by eq. (5).

$$\mathcal{L}_{\text{cWGAN}} = \mathbb{E}\left[\mathcal{D}(\mathcal{G}(Z,Y),Y)\right] - \mathbb{E}\left[\mathcal{D}(X,Y)\right] \quad (5)$$

In practice, an additional loss term $\mathcal{L}_{\text{GP}}$ is added to the discriminator's one to add a penalty when the norm of its gradient diverges from 1, as described in eq. (6), with $\widehat{X}$ being sampled by interpolating between samples produced by the generator and real X distribution samples.

$$\mathcal{L}_{\text{GP}} = \lambda_{\text{GP}} \mathbb{E}\left[\left(\left\|\nabla_{\widehat{X}} \mathcal{D}\left(\widehat{X}, Y\right)\right\|_2 - 1\right)^2\right] \quad (6)$$

This additional loss term enforces a constraint over the discriminator's search space to provide convergence guarantees: in theory, maximizing $\mathcal{L}_{\text{cWGAN}}$ for $\mathcal{G}$ and minimizing it for $\mathcal{D}$, when $\mathcal{D}$ is restricted to 1-Lipschitz functions, leads to a generator that minimizes the 1-Wasserstein distance between its generated distribution and the target distribution

[12]. Training typically involves alternating phases: the discriminator is trained to minimize $\mathcal{L}_{\text{cWGAN}} + \mathcal{L}_{\text{GP}}$ across several batches, followed by a single batch where the generator is trained to maximize $\mathcal{L}_{\text{cWGAN}}$.

Applied to infrared measurement, the cWGAN-GP learns to approximate the distribution of temperatures autonomously, given an infrared image without knowing surface materials properties (emissivity, reflectivity).

4. APPLICATION IN CONTROLLED CONDITIONS

In this initial case study, our objective is to validate the efficacy of the cWGAN-GP in providing the correct Bayesian posterior probability distribution. To achieve this, we have constructed a controlled dataset that enables the computation of the true posterior distribution, in order to compare it against the prediction.

4.1. Dataset generation

The created dataset consists of synthetic infrared images modeling a camera focusing on a small surface characterized by an unknown temperature T and emissivity ε, regrouped in a parameter $X = (T, \varepsilon)$ to estimate. In this case test, the reflected part in eq. (1) is assumed to be negligible: it is equivalent to expressing the outgoing radiance L^{out} as a direct function of T and ε. The captured image is represented as a matrix $Y \in \mathbb{R}^{3\times3}$, disturbed by a noise $N \in \mathbb{R}^{3\times3}$. The resulting noisy measurement is described by eq. (7).

$$Y = \mathcal{F}(X) + N = [\varepsilon L^{\circ}(T)]_{3\times3} + N \tag{7}$$

This simple example highlights two important factors of inverse problems: firstly, the measurement is affected by noise, and secondly, even with an ideal measurement, multiple pairs of temperature and emissivity values could yield the same radiance.

The dataset of 60,000 elements is generated by sampling the parameter $X = (T, \varepsilon)$ with a prior distribution p_X modeled as the mixture of two multivariate Gaussian distributions. Fig. 2 shows the resulting joint and marginal distributions in temperature and emissivity of the dataset. A Gaussian noise $N \sim \mathcal{N}(\mathbf{0}_{3\times3}, \sigma_N^2)$ is used to perturb the measurement.

4.2. Network architecture and learning phase

We used small multilayer perceptrons for both the generator and discriminator, each with four hidden layers of 15 neurons and LeakyReLU activation (slope 0.2). The generator processes a 13-dimensional input (a 3x3 image vector plus a 4-dimensional latent vector) and outputs temperature and emissivity values, while the discriminator uses an 11-dimensional input (9-dimensional image vector plus temperature and emissivity) to return a scalar.

The model was trained for 200 epochs on 50,000 images (batch size 64), alternating between training the discriminator on four batches and the generator on one. We set the discriminator's gradient penalty loss λ_{GP} to 10 and used an Adam optimizer (parameters: 0.5 for β_1, 0.9 for β_2, learning rate 10^{-4}).

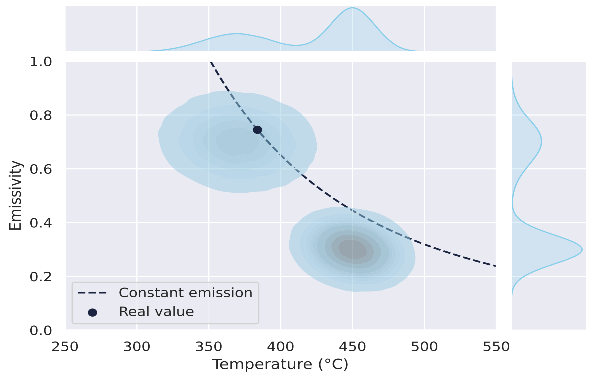

Fig. 2: Prior distribution for the generation of the controlled environment dataset: the joint probability distributions of temperature (T) and emissivity (ε), with their marginal probabilities plotted at the top and right, respectively. For a given couple (T^*, ε^*) represented by the black dot, the dashed line shows the set of possible (T, ε) combinations that lead to the same emitted radiance.

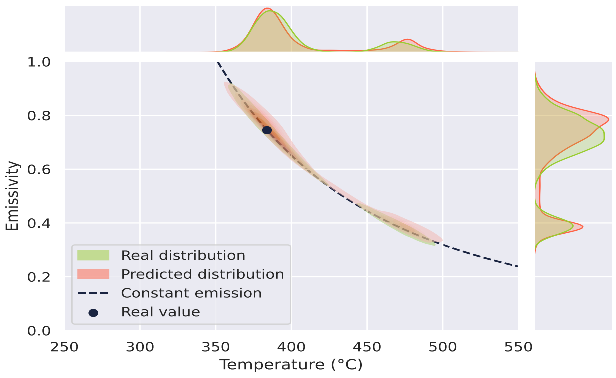

Fig. 3: True Bayesian (green) and predicted by the cWGAN (red) posterior distributions of temperature (T) and emissivity (ε) after the measurement of a noisy image generated from a true (T^*, ε^*) couple. The distributions match closely, and are both concentrated around the dashed line that represents the set of possible (T, ε) combinations that leads to the same emitted radiance.

4.3. Results

After training, the model is tested on a set of 10,000 images, generated from the distribution p_X. To evaluate the model accuracy, the predicted posterior distribution $p_{\mathcal{M}(Y,Z)|Y}$ is compared with the true posterior distribution $p_{X|Y}$. This ground truth distribution is sampled using Metropolis-Hastings algorithm [13], since p_X, p_N, and $\mathcal{F}$ can be easily evaluated.

Fig. 3 compares the true and predicted posterior distributions for a given measurement of the test set. In a first qualitative way, a good agreement is found between the prediction and the ground truth value.

For a quantitative analysis, we fitted both distributions to two-component Gaussian mixtures, comparing parameters like means (μ_T, μ_ε), standard deviations (σ_T, σ_ε), the corre-

Metric	$\left\|w^{\mathrm{P}} - w^{\mathrm{R}}\right\| / w^{\mathrm{R}}$	$\left\|\mu_T^{\mathrm{P}} - \mu_T^{\mathrm{R}}\right\|$	$\left\|\mu_\varepsilon^{\mathrm{P}} - \mu_\varepsilon^{\mathrm{R}}\right\|$
Mean	10.0%	2.7 °C	0.010
Metric	$\left\|\sigma_T^{\mathrm{P}} - \sigma_T^{\mathrm{R}}\right\| / \sigma_T^{\mathrm{R}}$	$\left\|\sigma_\varepsilon^{\mathrm{P}} - \sigma_\varepsilon^{\mathrm{R}}\right\| / \sigma_\varepsilon^{\mathrm{R}}$	$\left\|r^{\mathrm{P}} - r^{\mathrm{R}}\right\|$
Mean	22.0%	19.4%	0.017

Table 1: Mean errors on the test set, obtained by comparing the real (R) and predicted (P) Gaussian mixture fits.

lation coefficient (r), and Gaussian weights (w). The mean absolute or relative errors over a test set of 10,000 samples are listed in Table 1, indicating a high degree of accuracy: a 2.7 °C (resp. 0.01) mean temperature (resp. emissivity) error.

These promising results demonstrate the model ability to deduce automatically the correct posterior distribution, suggesting the potential for complex scenarios where exact distribution expressions may be unknown.

5. APPLICATION TO TOKAMAK SIMULATIONS

For this second experiment, we aim to get closer to the practical situations encountered while exploiting images provided by the infrared cameras of a tokamak. More complex simulated images are processed, as they result from a dataset of simulations of the WEST tokamak. Contrary to the last experiment, information is not sufficiently controlled to enable a computation of a real Bayesian posterior distribution. We expect the cWGAN to still automatically approximate this posterior distribution which will enable to attach uncertainties to temperature predictions.

5.1. Dataset description

The training dataset is the one used in [3], which contains simulations of a WEST tokamak camera focusing on the lower divertor. This component, which is exposed to a maximum heat flux, is particularly susceptible to surface changes due to erosion-deposition phenomena or damage such as melting during operation, which affect its emissivity.

The synthetic dataset is created using ray tracing to simulate infrared images from a given thermal scene. The following assumptions are made for the scenes: (a) uniform temperature for each main in-vessel components (inner and outer wall, antenna, baffle, ripples protection), (b) parameterized function applied on lower divertor to roughly reproduce the experimental strike line heat pattern [14], (c) uniform emissivity for each in-vessel components, with the exception of the spot of interest located on the lower divertor. Our study focuses on this spot by varying the temperature and/or emissivity. This allows us to generate pairs of similar images in the dataset, *i.e.* with the same measured radiance given different couples (T, ε) used as input, as shown in Fig. 4. It is worth noting that this assumes that the reflected light is diffuse. This hypothesis places us in the most restrictive case, where we cannot rely on specular reflection patterns to distinguish between different solutions.

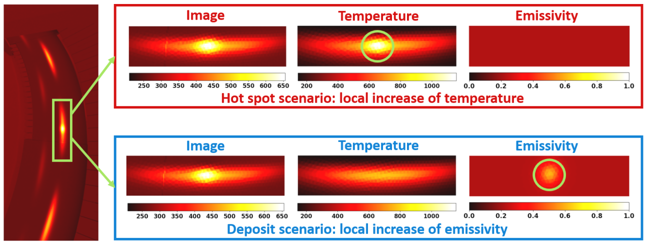

Fig. 4: Simulated image from the dataset showing a camera's view of the tokamak WEST's lower divertor. A spot on the divertor exhibits either increased temperature or emissivity, resulting in identical radiance and illustrating an ill-posedness challenge.

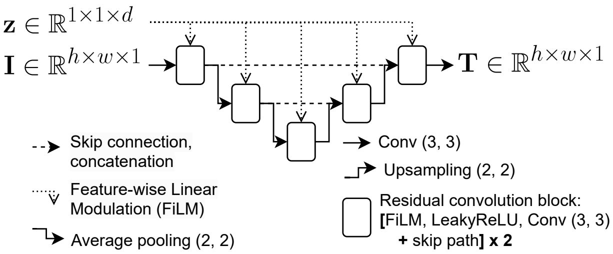

Fig. 5: Schematic representation of the conditional residual U-Net used as generator. A latent vector is injected at each convolution step by scaling the tensor features.

The dataset, generated with varied camera positions, orientations, emissivities (ranging from 0.05 to 0.5), and temperatures (50 to 500 °C, exceeding 2,000 °C for the lower divertor), consists of 128x128 pixel images. These images are cropped from 640x256 pixel originals and transformed from radiance measurements to apparent blackbody temperatures, paired with the true temperature maps.

5.2. Network architecture and learning phase

The goal is to train the model to produce temperature maps from a given input image.

The generator, inspired by [8], is a conditional U-Net, augmented from the standard U-Net to integrate a latent vector $\mathbf{z}$ of dimension 4 using Feature-wise Linear Modulation (FiLM) [15]. Fig. 5 illustrates the model. This U-Net variant incorporates FiLM layers at each stage to ensure the latent vector influences both semantic and spatial output details. The network uses residual blocks, each with a FiLM layer, a convolution layer, and an activation layer (LeakyRELU with slope 0.3), plus a skip path. Downsampling and upsampling are performed using AveragePooling and nearest-neighbor interpolation, respectively. The model starts with 32 channels per layer, doubling after every two downsampling stages, leading to a total of 11.8 million trainable weights.

The discriminator processes two-channel inputs: the image and the associated temperature map. It uses residual convolution blocks and (2, 2) max-pooling, with each block consisting of two LeakyReLU layers, a (3, 3) convolution, and a skip path with a (1, 1) convolution. Starting with 16 units

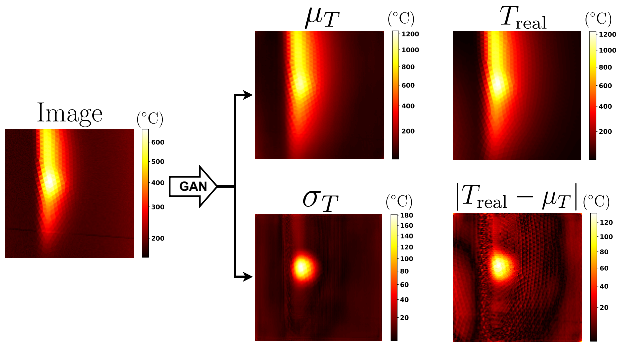

Fig. 6: Example of prediction made by the cWGAN on an image from the test set, compared with its associated true temperature map ($\mathbf{T}_{\text{real}}$). A set of 128 predictions is generated, and aggregated to compute a mean temperature map ($\boldsymbol{\mu}_T$) and standard deviation map ($\boldsymbol{\sigma}_T$). The predicted $\boldsymbol{\sigma}_T$ can identify the local spot of greater uncertainty and approximate (to some scaling factor) the absolute error made when using $\boldsymbol{\mu}_T$ as temperature prediction.

per layer, the number doubles after each one of the 7 pooling operations. The final stages use a dense layer of 128 units followed by a single unit layer for output. This leads to a total of 19.7 million trainable weights.

We trained the generator and the discriminator for 1,000 epochs using a training set of 30,000 images, with a batch size of 48, and alternating between training the discriminator on four batches and the generator on one. The discriminator's gradient penalty loss λ_{GP} is set to 10. We used an Adam optimizer to update the weights of both models, with parameters set to 0.5 for β_1 and 0.9 for β_2, with a learning rate of 10^{-4}.

5.3. Results

The model performance is assessed on a test set consisting of 10,000 elements. For each test image, the trained generator is fed 128 different latent vectors to produce a set of predictions. These predicted temperature maps are aggregated to form an average map ($\boldsymbol{\mu}_T$) and a standard deviation map ($\boldsymbol{\sigma}_T$). We demonstrate in this section that each predicted σ_T can effectively be used to construct reliable confidence intervals for the real temperature (T_{real}). An example of $\boldsymbol{\mu}_T$ and $\boldsymbol{\sigma}_T$ predictions is shown in Fig. 6 for a specific test image. The predicted $\boldsymbol{\sigma}_T$ can identify the local spot of greater uncertainty and approximate (to some scaling factor) the absolute error made when using $\boldsymbol{\mu}_T$ as a temperature prediction.

To evaluate our model, we analyze separately the region of the image where the apparent hot spot is located (referred as "ill spot" since it might be due to an increase in temperature or emissivity), and the rest of the image. This is achieved by tracking the location of the ill spot in each image and using a 15x15 pixel window centered around it to split predictions. Fig. 7 confirms that the model exhibits greater uncertainty at the ill spot compared to other areas, by showing the mean predicted relative standard deviation (μ_T/σ_T) as a function of the predicted mean temperature (μ_T). Table 2 compares the overall errors $|\mu_T - T_{\text{real}}|$ to the predicted σ_T: there is an approximation of one degree between the two, which is proof that σ_T seems to be a good estimate of the error.

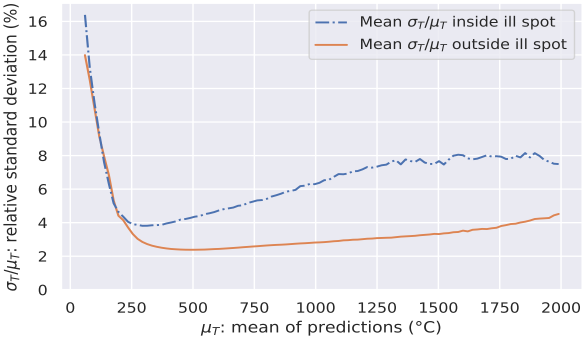

Fig. 7: Mean relative standard deviations of the cWGAN predictions in function of the predicted temperature on the test set. Predictions are separated into the one belonging to the ill spot or outside the point. The network has a greater standard deviation in its predictions for the ill spot, indicating that it is able to spot the zone of greater uncertainty.

Metric	Inside ill spot	Outside ill spot
Mean $\lvert\mu_T - T_{\text{real}}\rvert$	29.6 °C	12.0 °C
Mean predicted σ_T	30.7 °C	10.9 °C
Mean $\lvert\mu_T - T_{\text{real}}\rvert / T_{\text{real}}$	5.9%	6.8%
Mean predicted σ_T / μ_T	5.8%	6.5%

Table 2: Comparison of two error metrics on the test set with the mean predicted deviations.

Moreover, we illustrate the applicability of constructing confidence intervals for the actual temperature (T_{real}) using the predicted μ_T and σ_T. Figure 8 illustrates the probability of T_{real} to be less than $k\sigma_T$ away from μ_T, for both inside and outside ill spot predictions. The close alignment of the two curves suggests that σ_T, regardless of its magnitude, can be reliably used to create confidence intervals. Specifically for our task, the curves indicate a 95% probability that the true temperature will be within 3 σ_T from μ_T, over 80% for 2 σ_T, and more than 50% for 1 σ_T. These results demonstrate the ability of the model to generate meaningful uncertainties after training with an appropriate synthetic dataset: the derived confidence intervals, based on the predicted σ_T, are a valuable indicator for decision-making processes in terms of machine protection.

6. CONCLUSION AND PERSPECTIVES

This study addresses the challenges of temperature measurement from infrared images in the complex environment of nuclear fusion devices (tokamak) using metallic surfaces of unknown emissivity and reflectivity. This paper deals with ill-posedness while estimating the surface temperature or emissivity. Our approach is based on a probabilistic model able to reproduce the distribution of the training data and thus provide a range of uncertainty on the solution (emissivity and/or temperature). For this purpose, we developed a Conditional Wasserstein GAN with Gradient Penalty (cWGAN-GP)

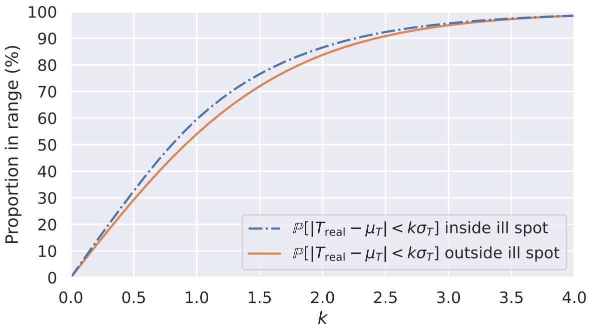

Fig. 8: Proportion of the real temperatures from the test set that are less than $k\sigma_T$ away from μ_T, as a function of k, with μ_T and σ_T being the mean and standard deviation of the predictions. Predictions are separated into the one belonging to the ill spot or outside the point. The similarity of the curves indicates that the predicted standard deviation can be used as a good estimator of confidence interval, regardless of its value. For instance, around 95% of the real values are less than 3 standard deviation away from mean predictions.

model, which was trained and evaluated on simulated infrared images. We show that such a model is powerful in generating a reliable temperature probability distribution, even in scenarios where the exact distribution is unknown, or in cases of extreme uncertainty. In the specifically studied setup, the average surface temperature is predicted with an average error of 6% on 10,000 simulated test images with surface temperatures lying between 50 and 2,000 °C. On average, the 95% confidence interval on temperature predictions represents 20% of the estimated value. These results, which provide a good confidence interval around the temperature prediction, are crucial information for decision making and improving the safety of tokamak operations.

A significant aspect of this method is its automatic adaptability to the complexity of the underlying inverse problem. While this flexibility is advantageous, it also depends on the representativeness of the training dataset, from which the posterior distribution of the solution to the inverse problem is learned. In our case, given the absence of ground truth in real-world scenarios, the realism of the simulated data used for the learning phase is crucial. This requires generating a massive training dataset spanning the possible scenarios encountered during tokamak operations, while limiting the range of variation of input parameters to avoid introducing too much uncertainty. For example, several efforts are made to characterize the material optical properties in laboratory.

In addition, a further step towards its application in tokamak should be to extend the model to handle multimodal inputs and to incorporate temporal dynamics. This could significantly improve prediction accuracy and uncertainty assessment on experimental data.

Acknowledgment

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200 — EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

References

[1] X. Courtois *et al.*, “Full coverage infrared thermography diagnostic for WEST machine protection,” *Fusion Engineering and Design*, SI:SOFT-30, vol. 146, pp. 2015–2020, Sep. 2019.

[2] R. Reichle *et al.*, “Defining the infrared systems for ITER,” *Review of Scientific Instruments*, vol. 81, no. 10, 10E135, Oct. 2010.

[3] A. Juven, M.-H. Aumeunier, and J. Marot, “U-Net for temperature estimation from simulated infrared images in tokamaks,” *Nuclear Materials and Energy*, vol. 38, p. 101 562, Mar. 2024.

[4] Z. Li, M. Shafiei, R. Ramamoorthi, K. Sunkavalli, and M. Chandraker, “Inverse rendering for complex indoor scenes: Shape, spatially-varying lighting and svbrdf from a single image,” in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 2472–2481.

[5] T. Hastie, R. Tibshirani, J. H. Friedman, and J. H. Friedman, *The elements of statistical learning: data mining, inference, and prediction*. Springer, 2009, vol. 2.

[6] S. Sengupta, J. Gu, K. Kim, G. Liu, D. Jacobs, and J. Kautz, “Neural Inverse Rendering of an Indoor Scene From a Single Image,” in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Seoul, Korea (South): IEEE, Oct. 2019, pp. 8597–8606. (visited on 01/26/2022).

[7] J. Adler and O. Öktem, *Deep Bayesian Inversion*, arXiv:1811.05910 [cs, math, stat], Nov. 2018.

[8] D. Ray, H. Ramaswamy, D. V. Patel, and A. A. Oberai, “The efficacy and generalizability of conditional GANs for posterior inference in physics-based inverse problems,” *Numerical Algebra, Control and Optimization*, Dec. 2022, Publisher: Numerical Algebra, Control and Optimization.

[9] R. Ramamoorthi and P. Hanrahan, “A signal-processing framework for inverse rendering,” in *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, ser. SIGGRAPH ’01, New York, NY, USA: Association for Computing Machinery, Aug. 2001, pp. 117–128.

[10] H. Goh, S. Sheriffdeen, J. Wittmer, and T. Bui-Thanh, “Solving Bayesian Inverse Problems via Variational Autoencoders,” in *Proceedings of the 2nd Mathematical and Scientific Machine Learning Conference*, theisssn: 2640-3498, PMLR, Apr. 2022, pp. 386–425.

[11] C. Saharia *et al.*, “Palette: Image-to-Image Diffusion Models,” in *ACM SIGGRAPH 2022 Conference Proceedings*, ser. SIGGRAPH ’22, New York, NY, USA: Association for Computing Machinery, Jul. 2022, pp. 1–10.

[12] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. C. Courville, “Improved Training of Wasserstein GANs,” in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017.

[13] S. Chib and E. Greenberg, “Understanding the Metropolis-Hastings Algorithm,” *The American Statistician*, Nov. 1995, Publisher: Taylor & Francis Group.

[14] T. Eich *et al.*, “Inter-ELM Power Decay Length for JET and ASDEX Upgrade: Measurement and Comparison with Heuristic Drift-Based Model,” *Physical Review Letters*, vol. 107, no. 21, p. 215 001, Nov. 2011.

[15] E. Perez, F. Strub, H. d. Vries, V. Dumoulin, and A. Courville, “FiLM: Visual Reasoning with a General Conditioning Layer,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, Apr. 2018, Number: 1.