



Super-resolution reconstruction of scalar fields from the pyrolysis of pulverised biomass using deep learning

A. Shamooni ^a*, R. Cheng ^a, T. Zirwes ^a, O.T. Stein ^b, A. Kronenburg ^a

^a Institute for Reactive Flows (IRST), University of Stuttgart, Pfaffenwaldring 31, 70569 Stuttgart, Germany

^b Engler-Bunte-Institut, Simulation of Reacting Thermo-Fluid Systems, Karlsruhe Institute of Technology, Engler-Bunte-Ring 7, 76131 Karlsruhe, Germany



ARTICLE INFO

Dataset link: <https://doi.org/10.18419/DARUS-5519>

Keywords:

Super-resolution
GAN
Biomass combustion
Particle
Turbulence

ABSTRACT

Recently, advanced deep-learning techniques have been successfully applied as deconvolution operators to super-resolve the low-resolution data in large-eddy simulation (LES). The super-resolution (SR) operator provides an approximate inverse to the filter operators in LES such that the under-resolved and un-resolved sub-grid information can be reconstructed from the resolved scales. In this work, a particle-aware attention-based conditional super-resolution generative adversarial network (PACASRGAN) is proposed for the fourfold SR of gas field scalars which are generated by the pyrolysis process in a hot turbulent flow laden with pulverised biomass particles. Multiple carrier-phase direct numerical simulations (DNS) of two-way coupled particle-laden flows with heat and mass transfer, that mimic the near-burner field of pulverised biomass combustion (PBC) systems, are carried out to build the training/testing datasets. The model performance is assessed in an *a priori* manner by investigating statistical quantities of interest for the modelling in LES of PBC. The results show that the proposed model can super-resolve the temperature and mixture fraction fields to a good accuracy and outperforms unconditional GAN models. Particles create localised sources/sinks via two-way coupling which sharpen scalar gradients in the subgrid. The particle mask and feature vector encode this localisation to improve the predictions of the generator. The scalar spectra, the conditional average of unresolved scalar variances, the probability density function (PDF), and the conditional average of the square of the mixture fraction gradient from the reconstructed fields follow the DNS values well. Slight deviations are observed at rich conditions in conditional statistics and at the tail of the PDFs. Nonetheless, the results demonstrate that SR is applicable to two-way coupled particle-laden flows with heat and mass transfer, providing reasonably accurate high-resolution information for both the entire gas field and particle positions.

1. Introduction

Super-resolution (SR) is a computer vision task which aims at mapping a low-resolution (LR) image to the corresponding high-resolution (HR) image [1]. In terms of scientific data science, the SR can be regarded as a deconvolution operation which reconstructs direct numerical simulation (DNS)-like fields from the filtered data available in large-eddy simulation (LES). The outcome can e.g. be used to increase the resolution of PIV measurements or to help sub-grid model developments by providing extra sub-grid information. SR has been recently shown to have a great potential for the modelling of sub-grid scale stress [1,2] and combustion [3–5]. Pure convolutional neural networks (CNNs) with pixel-wise loss may produce better pixel-wise results, however, non-physical and over-smoothed structures may be created [6]. Hence, novel CNN training designs based on generative

adversarial networks (GANs) [7] have been utilised [4,5,8,9] in scientific SR to create a balance between pixel-wise and statistical errors in the regression problem.

The focus of the current study is on SR of turbulent two-way coupled particle laden flows in the presence of heat and mass transfer. Specifically, the target is the SR of the filtered fields produced by the devolatilisation process of pulverised biomass combustion (PBC) from Euler–Lagrange (EL) numerical simulations. Euler–Lagrange LES of PBC with flamelet approaches have recently gained attention [10]. In such models, sub-grid inhomogeneities in the carrier gas due to turbulence are often neglected or simple closures for scalars such as mixture fraction are used. SR models can be utilised to enhance the sub-grid information in these models. The application of SR to the flow fields involving particles is scarce [11–14]. In [11] the particles were one-way coupled, and the SR task was therefore comparable to particle-free

* Corresponding author.

E-mail address: ali.shamooni@irst.uni-stuttgart.de (A. Shamooni).

cases. In [12–14], traditional GANs, originally designed for particle-free flows, were applied to spray combustion without explicit particle feedback to the network. In our recent work [15], we demonstrated that particle data can improve SR performance in predicting turbulence modulation in two-way coupled flows. However, the coupling was restricted to kinematic interactions, i.e. momentum exchange through drag only. In fact, a perfect SR model trained on particle-free turbulence fails to recover particle-turbulence interactions when tested on particle-laden data [15].

The present study extends the previous research by addressing pyrolysis in turbulent PBC, where particles influence the carrier gas not only via drag but also through heat absorption and volatile release. Specifically, we aim to super-resolve density, temperature and mixture fraction fields, mapping LR inputs from filtered EL-DNS of biomass pyrolysis to HR. For this purpose, a particle-aware conditional attention super-resolution GAN (PACASRGAN) is proposed, the generation of DNS datasets for unsteady devolatilisation under different carrier-gas temperatures and mass loadings is described, and the model is evaluated in *a priori* tests. Results are presented, the model is compared with traditional GANs, and the importance of particle data is investigated.

2. Modelling

2.1. Network design

The goal of the modelling is to conduct SR of density, temperature, and volatile mixture fraction fields, viz. $\xi = \{\rho, T, Z\}$, which are generated by the pyrolysis process in a hot turbulent flow laden with pulverised biomass particles. We select these quantities as local information is needed for accurate modelling of phase transfer processes that determine devolatilisation. The GAN approach [7] has been employed which is composed of two competing networks, namely the *generator* which is trained to produce fake HR data and the *discriminator* which is trained to distinguish between the real and fake HR data. When unsteady particle-gas interactions exist, the generator faces several challenges, since particles are localised and sparse, but they interact with the turbulent field. The model needs to learn which particle information modify the carrier gas. Furthermore, gas SR should work without particles in particle-free regions, i.e., the model must ignore missing particle regions and only super-resolve gas features. To fulfil these goals, we propose a particle-aware conditional attention SRGAN (PACASRGAN) model which contains a modified residual-in-residual dense block (RRDB) as its generator and a UNet discriminator [16].

As can be seen in Fig. 1 (top), the input data to the generator is passed through a first convolutional layer to extract 128 features, then it is passed to 16 attention residual-in-residual dense blocks (A-RRDB). The output is passed through a convolutional layer before receiving information from the skip connection. Then, it is passed through 2 up-sampling layers which are followed by a convolution, non-linearity (Leaky ReLU) and the final convolution to produce the 3 output channels. Each A-RRDB involves 3 residual dense blocks with an augmented convolutional block attention module (CBAM) [17]. The CBAM contains a sequence of channel attention (CA) and spatial attention (SA) blocks to fuse gas and particle feature maps selectively, ensuring that only the most important particle regions contribute to the final encoding after each block. CA highlights important feature maps across channels, while SA highlights the most important locations in the feature map. As depicted in Fig. 1, in CA first the input features are average- and max-pooled, they are passed to a shared multilayer perceptron (MLP) which is a contraction-expansion MLP (with a factor of 16), and a sigmoid function. The output is multiplied element-wise to each channel of the original feature map. In SA, the input features are average- and max-pooled. The concatenated outcome is passed through a convolution layer and a sigmoid function to output a single channel, with values between 0 and 1, that are multiplied element-wise to each channel of the input feature map. The details of the kernel

sizes, number of kernels and the strides are mentioned above each convolution layer in Fig. 1.

The generator G creates SR fields from LR inputs,

$$\hat{\xi} = \{\hat{\rho}, \hat{T}, \hat{Z}\} = G(C, \bar{v}, k_{sgs}, P), \quad (1)$$

where \bar{v} is the filtered gas viscosity, $C = \{\xi_{LR}, M_p\}$ contains the LR filtered vector of scalars, $\xi_{LR} = \{\bar{\rho}, \bar{T}, \bar{Z}\}$, and a low-resolution particle mask (M_p). M_p is a binary field defined at each LR grid point, indicating whether particles are present within the corresponding LR grid cell. M_p serves as an additional conditioning input to both the generator and the discriminator, allowing the model to account for regions influenced by particle presence, such as sites where devolatilisation mass is injected. Further, we follow [11] and condition the generator on sub-grid kinetic energy, which has shown to improve the generator's prediction of the gas fields [11]. The generator is additionally conditioned on particle information. The particle information input vector in the generator input, P , in Eq. (1) is defined as,

$$P = \{\log(1 + N_p), \log(1 + \sum_k \rho_{p,k}), \log(1 + \sum_k \dot{S}_{p,k})\}, \quad (2)$$

denoting the sum of number of particles, their mass densities, and the mass source term generated by particles at each LR grid point. These additional conditioning inputs enhance the physical consistency of the super-resolved fields, particularly in regions influenced by local mass release from pyrolysis.

The discriminator D has a UNet architecture as shown in Fig. 1 (bottom), containing spectral normalisation layers as in [16], conditioned on high frequency information of the HR data, the LR particle presence mask (M_p), and the LR gas-phase scalars fields, ξ_{LR} . It has been shown that conditioning on LR data improves the stability of the adversarial training [11]. The input LR data are upsampled by the SR factor (4x) using a custom differentiable and locally conservative upsampling layer. The layer is designed to preserve mass-bound properties of the conserved scalars. The high-frequency information of the HR data (ground truth DNS or fake SR) are extracted using stationary wavelet transform [18] that is applied to the input HR data. Although many convolution kernels already exist in the architecture of D , constraining the discriminator decision on wavelet-transformed data aims at forcing the discriminator to learn the high-frequency patterns. Haar wavelets [18] have been employed and only the three bands of the second level transformations have been used in the discriminator to limit the computational costs. The input HR, its high frequency information and upsampled LR are concatenated and passed through a series of convolution layers as shown in Fig. 1. The output of D has a similar resolution as its input with each pixel containing information about the decision of the discriminator on that pixel.

2.2. Loss functions

The total generator loss combines the pixel-wise, scalar gradient, and adversarial terms, viz. $\mathcal{L}_{\text{total}} = \lambda_{\text{pixel}} \mathcal{L}_{\text{pixel}} + \lambda_{\text{grad}} \mathcal{L}_{\text{grad}} + \lambda_{\text{adv}} \mathcal{L}_{\text{adv}}^G$. The pixel-wise loss is evaluated as L1-norm error as it tends to produce sharper edges and finer details compared to L2-norm [16]. To improve the reconstruction of sharp features and gradients, and physical consistency of the super-resolved scalar fields, a gradient loss is used. For each scalar field $\xi \in \{\rho, T, Z\}$ at grid point i , the gradient loss is computed as,

$$\mathcal{L}_{\text{grad}} = \sum_{\xi \in \{\rho, T, Z\}} \frac{1}{N} \sum_{i=1}^N \sum_{n=1}^2 \left(\left| \frac{\partial \hat{\xi}_i}{\partial x_n} - \frac{\partial \xi_i}{\partial x_n} \right| \right), \quad (3)$$

where N is total number of grid points at each slice and n is the direction. The gradients are computed using the second-order central difference method of PyTorch. The adversarial loss for the generator is based on a non-relativistic conditional discriminator D ,

$$\mathcal{L}_{\text{adv}}^G = -\mathbb{E} \left[\log D(\hat{\xi}, \text{SWT}(\hat{\xi}), C) \right], \quad (4)$$

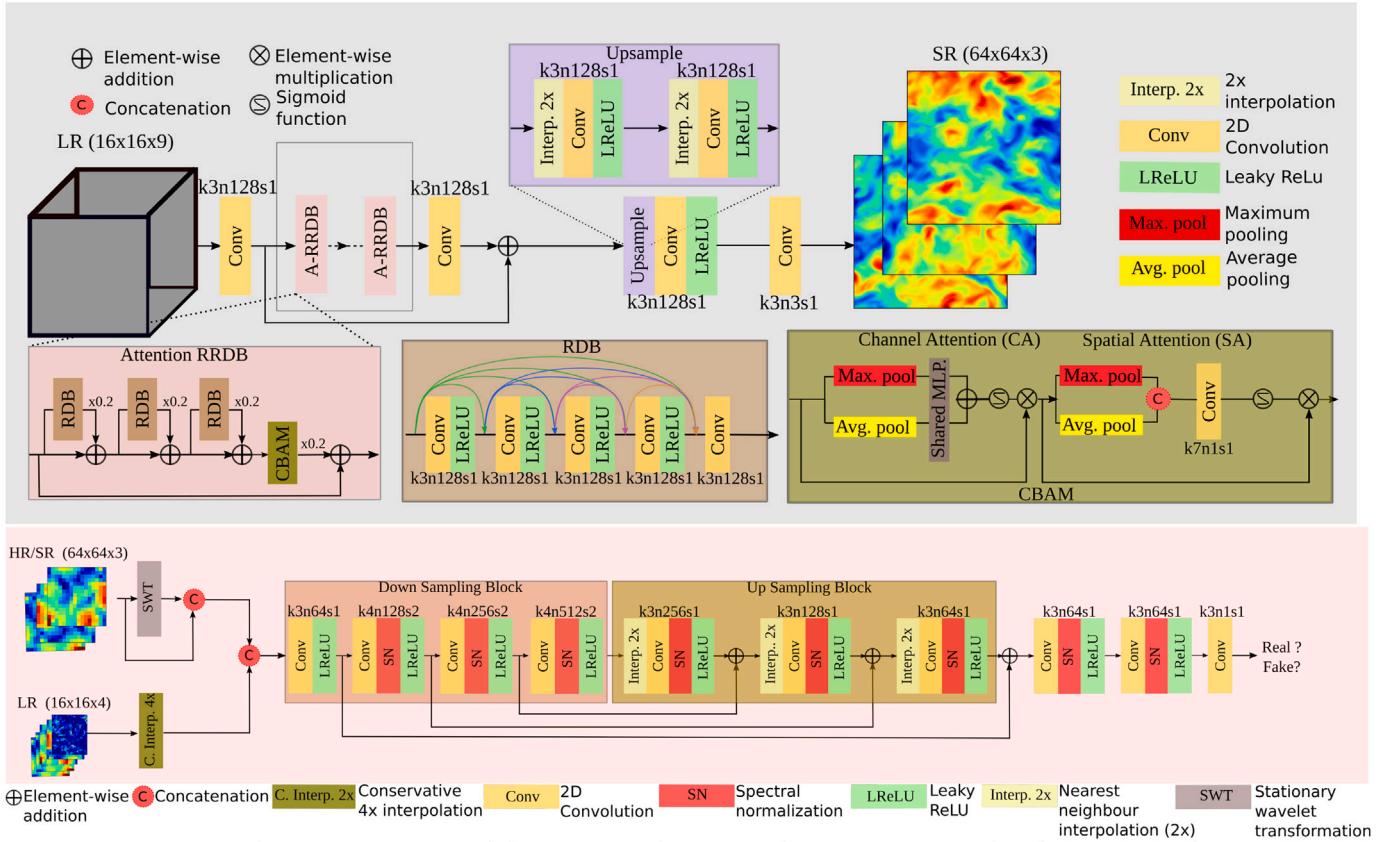


Fig. 1. The generator (top) and discriminator (bottom) architectures employed in the current study.

where $SWT(\hat{\xi}) = \{LH_2(\hat{\xi}), HL_2(\hat{\xi}), HH_2(\hat{\xi})\}$, with LH_2 , HL_2 , HH_2 the horizontal, vertical and diagonal Haar wavelet coefficients [18] at the second level of the stationary wavelet transformation. Finally the discriminator D is trained by minimising the following loss function,

$$\begin{aligned} \mathcal{L}_{\text{adv}}^D = & -\mathbb{E}[\log D(\xi, SWT(\hat{\xi}), C)] \\ & -\mathbb{E}[\log (1 - D(\hat{\xi}, SWT(\hat{\xi}), C))] \\ & -\mathbb{E}[\log (1 - D(\xi', SWT(\xi'), C))], \end{aligned}$$

where ξ' is the shuffled version of ground truth HR data (ξ) [11] in the training batch. This means that by shuffling, the HR-LR pairs are not matched any more and the discriminator is penalised whenever it identifies ξ' as ground truth although it is inconsistent with the LR conditions in C .

2.3. Training procedure

The adversarial training is carried out in one shot with no pre-training. Training employs a custom data pipeline including random batch selection (16 samples/GPU), flips/rotations, Favre filtering with a top-hat filter of width $4\Delta_{x,DNS}$ and downsampling by a 1/4 ratio, channel-wise normalisation, and paired random cropping to avoid overfitting (64^2 HR and 16^2 LR from corresponding 128^2 and 32^2 original sizes). A fixed filter type/width is used here, while a posteriori applications will require generalisability. This can be achieved by retraining with both Gaussian and top-hat filters across multiple widths to enhance robustness [11,19,20]. The downsampling factor of four ($4\Delta_{x,DNS} \approx 8\eta$), yields resolved-to-total ratios of 94% for kinetic energy and 90% for temperature, consistent with previous studies [19,21]. The lower ratio for mixture fraction (76%) reflects subgrid devolatilisation effects. Larger factors between HR/LR generally reduce the accuracy, though Gaussian filtering at fixed downsampling ratio can

alleviate such degradation [11]. The chosen patch size corresponds to approximately $1.9l_t \approx 128\eta$, ensuring several large- and small-scale eddies per sample and consistent with previous studies [19,22]. Here, l_t is the integral length scale. While larger patches increase the receptive field and may in theory improve performance, we found they reduced small-scale fidelity. The discriminator tends to prioritise large-scale features, an effect amplified in particle-laden cases. The Adam optimiser is utilised with the learning rates of 10^{-4} , 10^{-5} for generator and discriminator, respectively, that are decreased by half after 200 K, 300 K, and 400 K iterations. The hyper-parameters in the total loss function are $\lambda_{\text{pixel}} = 0.1$, $\lambda_{\text{grad}} = 0.05$, $\lambda_{\text{adv}} = 0.012$, in such a way that after the first 100 iterations the losses have comparable values within the same order of magnitudes. The code is implemented in the basicSR [16] framework using PyTorch 2.3. The training is performed on four A100 GPUs for 525000 iterations which takes approximately 48 h.

Inference benchmarks are performed on a single NVIDIA RTX 6000 Ada GPU with input tensors of shape $[1, 9, 32, 32]$ and output of shape $[1, 3, 128, 128]$. The inference time is 20.5 ms with peak allocated memory of 221 MB. The intended application of PACASRGAN is to act as a deconvolution operator in a posteriori LES of two-way coupled particle-laden flows. To provide a first estimate of the scalability to 3D, we implemented the 3D extension of the generator network with randomly initialised weights. For an input/output with previous shapes, the average latency over 20 runs is 247.4 ± 2 ms with a peak allocated memory of 5.1 GB. For multi-patch inference [22] with $[1, 9, 16, 16, 16]$ input, latency and peak memory reduce to 38.5 ± 0.05 ms, and 1.16 GB, respectively. To estimate training costs, we scale the 2D training time by the ratio of forward inference times (≈ 12), corrected by a factor of 0.25, an estimation to account for the larger voxel content per 3D sample. The numbers are feasible using current GPU clusters. Further, multi-CPU/GPU training and a posteriori LES deployment can be followed [22].

Table 1

Specifications for particle-laden forced DNS setups used for training/testing.

Case	LT1	LT2	LT3	MT1	MT2	MT3	HT1	HT2	HT3	OT1	OT2
Initial gas temperature [K]	1100	1100	1100	1300	1300	1300	1500	1500	1500	1200	1200
Particle mass loading	0.15	0.3	0.6	0.15	0.3	0.6	0.15	0.3	0.6	0.15	0.6
Averaged kinetic energy [m^2/s^2]	32	27	25	57	52	42	84	80	63	43.5	32
Averaged dissipation rate $\times 10^4$ [m^2/s^2]	2.0	1.6	1.0	4.5	3.5	2.3	9.0	7.1	4.8	2.9	1.5
Particle diameter (d_p) [μm]	29.5	29.5	29.5	27.2	27.2	27.2	25.3	25.3	25.3	28.2	28.2
Particle number density $\times 10^9$ [m^{-3}]	5.4	10.8	21.5	5.6	11.7	23.4	6.3	12.6	25.1	5.6	22.4

3. Dataset generation

3.1. Direct numerical simulation

Three-dimensional Euler–Lagrange carrier-phase DNS [23] of pulverised biomass particle pyrolysis in a forced turbulence are carried out in different environments to generate the training/testing datasets. The homogeneous reactions are suppressed to only focus on the heat and mass transfer between the two phases. The turbulence is linearly forced [24] to a target Taylor Reynolds number of a particle-free flow ($Re_\lambda = 53$). The Kolmogorov length scale (η) is 10^{-4} m and the grid resolution is 128^3 . A cloud of pulverised walnut shell biomass particles [23] with initial mass density of 650 kg/m^3 are injected into a hot turbulent environment in a box-shaped domain with a length of 0.0256 m where pyrolysis occurs.

Four groups of DNS setups with different carrier gas temperatures are considered which are denoted by LT (low temperature), MT (medium temperature), HT (high temperature), OT (out-of-sample temperature) labels in Table 1. The first three groups are used to build training/testing datasets while the last group is only used for testing. The initial conditions of the turbulent carrier gas are obtained by preliminary particle-free stationary forced turbulent simulations with a similar Taylor scale Reynolds number. Since small and large scales are set by design, the forcing parameters are changed to account for different viscosities. Particles are allowed to interact with the turbulent carrier-phase kinematically, i.e., without heat exchange, for one eddy turn over time before heat- (and thereby mass-) transfer is switched on. This relaxation time period is important for the kinematic turbulence-particle interaction to be settled and for the carrier gas to reach a new stationary state. The new state is different for each temperature group and each case within each group as particles extract kinetic energy from the carrier gas and decrease the dissipation rate in forced turbulence which can be seen in Table 1 by comparing time-averaged statistics, i.e., kinetic energy and the dissipation rates between the cases in each group. The observations are consistent with the findings in [25] for pure kinematic interactions between particles and turbulence. In the current study we have limited the parameter space to carrier gas temperature and particle mass loading variations, while the Stokes number and target (particle-free) Reynolds number are fixed. The Stokes numbers based on Kolmogorov and Taylor length scales are $St_\eta = 10$ and $St_\lambda = 1.6$, respectively. This can be achieved by varying initial particle diameter, number density, and turbulence intensity as can be seen in Table 1. Note that in the simulations the particle diameter (d_p) remains constant during devolatilisation while its mass density changes [23]. The simulations cover the whole devolatilisation process for each case and further pure mixing.

The main governing equations are mentioned in [10,23] and are only briefly described here. The mass loss of an individual particle is governed by $dm_p = -(\dot{m}_{p,vol})dt$, with the devolatilisation rate being $\dot{m}_{p,vol} = (m_{p,VM,0} - m_{p,VM})Ae^{-E/(R_u T_p)}$, where T_p , $m_{p,VM,0}$ ($= 0.76 m_{p,0}$), and $m_{p,VM}$ are the particle temperature, the initial and instantaneous mass of volatile matter, respectively. The rate coefficients ($A = 7 \cdot 10^5 \text{ s}^{-1}$, $E = 5 \cdot 10^7 \text{ J/Kmol}$) are calculated in [23]. The drag force is considered to be the only volume force on the particle and modelled by a non-spherical drag model with a shape factor of 0.75 [23]. The particle temperature is governed by [10], $dT_p/dt = (T - T_p)/\tau_{con} + \dot{Q}_{devol}/m_p c_{p,p}$,

with $\tau_{con} = \left(\Pr \rho_p d_p^2 c_{p,p} \right) / (6 \text{Nu} c_p \mu)$ the time scale of heat transfer by convection. Nu is the Nusselt number, and Pr is the Prandtl number ($= 0.7$). d_p , μ and c_p are the particle diameter, gas viscosity and heat capacity at the particle position, respectively. $c_{p,p}$ is the heat capacity of the particle, expressed according to [10] and \dot{Q}_{devol} is the heat of devolatilisation [23]. The two-way coupling between the two phases is governed by the exchange of mass from devolatilisation, momentum from drag force, and enthalpy from convection, as well as the heat of devolatilisation, which have been described in detail in [10]. The particle properties are similar to the ones in [23] with the exception that CH_4 is considered to be the sole volatile species to reduce the complexity. The volatile mixture fraction Z , that is transported for analysis purposes, is initially zero and receives a two-way coupling mass source term, $\dot{S}_p = -\frac{1}{\Delta V} \sum_k^{N_p} \left[\frac{dm_{p,k}}{dt} \right]$, from the particles, with ΔV being the volume of the grid cell where the particle resides, and N_p is the number of particles present in the local cell volume.

3.2. Training/validation/testing datasets

Eighteen 3D snapshots from each case in LT, MT, and HT groups are selected randomly and uniformly, with respect to the cumulative mass generated by the particles to account for the unsteadiness of the devolatilisation process. This is carried out by binning the cumulative devolatilisation mass into equally distanced bins from which the training data is sampled. In this way it is ensured that enough data is sampled at each devolatilisation step, in addition to the final pure mixing process. From the remaining un-selected snapshots available from the DNS, we randomly select eight snapshots from each case to build the validation/testing (50%–50%) datasets. Since the SR model is a 2D SR network, we use all 2D $x - y$ -slices at 128 z -locations of each selected DNS snapshots. This will result in 20736 two dimensional 128×128 samples for training and 3456 samples for testing.

4. Results and discussion

In Fig. 2, the reconstructed mixture fraction (\hat{Z}) and temperature (\hat{T}) are compared with the DNS and LR input data for a randomly selected slice from the test dataset. Particles have been removed from the figures for the sake of clarity, however, their effects are obvious; In the DNS, the local peaks of Z and local reduction of temperature are clear signs of the devolatilisation of individual particles. The loss of high-frequency details can be clearly observed in \hat{Z} and \hat{T} while the model has recovered small scale details in \hat{Z} and \hat{T} qualitatively well.

In Fig. 3, the one-dimensional energy density spectra of mixture fraction and temperature for the snapshots shown in Fig. 2 are depicted. The energy density spectrum of a scalar, i.e., $E_\xi^{1D}(\kappa_x \eta)$, is computed as a function of wavenumber along the x -direction, κ_x , by first applying a discrete Fourier transform along the x -direction for each y -location and further averaging along the y -direction. As can be seen in Fig. 3, in the LR fields the variances are under-resolved before the cut-off ($= \pi/(4\Delta_{x,DNS})$), shown by the vertical dashed lines, while in the SR fields both under-resolved and un-resolved variances (after the cut-off) are reconstructed well. This shows that the model captures the average contribution of each spatial scale to the scalar field variance well.

To assess the performance of the model on the whole test dataset which contains unseen data from various temperature and mass loading

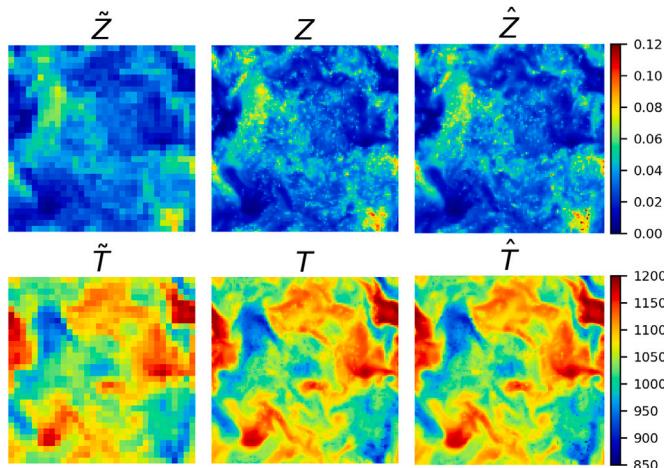


Fig. 2. Snapshots of mixture fraction (top) and temperature (bottom) of low-resolution input (left) vs. DNS (middle) vs. SR (right) from a randomly selected slice in test dataset.

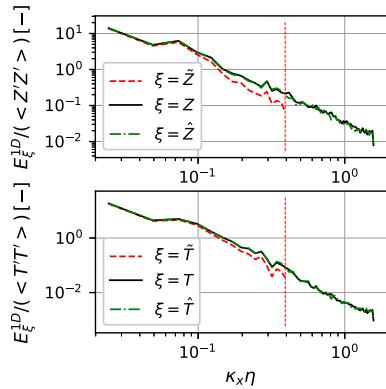


Fig. 3. Energy density spectra of mixture fraction (top) and temperature (bottom) normalised by the DNS variances. The vertical dashed lines show the LR cut-off.

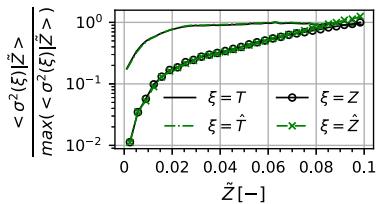


Fig. 4. Conditional average of unresolved sub-grid scalar variances evaluated on the whole test dataset.

conditions, the conditional averages of reconstructed unresolved scalar variances have been compared with the DNS counterparts. The unresolved variance of a scalar ξ at filter scale Δ , is defined by $\sigma^2(\xi; \Delta) = \xi^2 - \bar{\xi}\bar{\xi}$. The conditional averages are computed by conditioning on the filtered mixture fraction. In Fig. 4 it can be seen that the conditional mean of unresolved temperature variance from SR matches very well with the DNS. Both DNS and SR profiles show an increasing trend with respect to the mixture fraction. At low \tilde{Z} , the temperature variance is caused solely by the variance in the turbulent carrier gas. High \tilde{Z} values denote the cells where particles extract heat and release mass which causes the temperature variance to increase.

In Fig. 4, the statistics of the unresolved mixture fraction variance are shown as well. Similar to the profiles of temperature, the conditional mean of both SR and DNS increase with increasing mixture

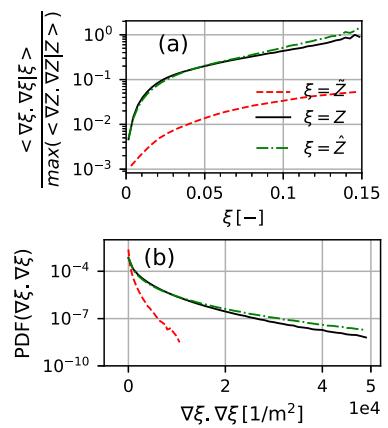


Fig. 5. Conditional average (a) and PDF (b) of the square of mixture fraction gradient evaluated on the whole test dataset.

fraction. With increasing the devolatilisation, Z increases and since the devolatilisation occurs at sub-grid levels, sub-grid variance increases. Since Z is initially zero and it is only produced by the devolatilisation process, the variance increases monotonically by increasing the injected mixture fraction from the particles. It can be seen that the statistics of the unresolved reconstructed field, \hat{Z} , matches very well with the DNS statistics. Slight deviations at high mixture fractions can be observed which is due to the slight overshoots of \hat{Z} at cells with no particle and high mixture fraction which will be explained in more detail later. Nonetheless, the predictions around stoichiometry (≈ 0.055) show good agreement with the DNS.

In Fig. 5, the PDF and conditional average of the square of the mixture fraction gradient is shown. This quantity is proportional to the scalar dissipation rate which is of interest in turbulent combustion models and determines whether a flame can establish itself in the vicinity of the particle or whether group combustion would be a more likely combustion scenario.

As can be seen in Fig. 5b, the PDF computed by LR fields is bounded by much lower peak gradients compared to the DNS. Furthermore, the number of nearly zero-gradient cells is about one order of magnitude higher than the DNS counterpart. This is due to the generated smoothed fields by the filtering in the LR. On the other hand, the distribution of the square of the gradients from the SR field follows the DNS PDF up to probabilities of around 10^{-7} with slight over-estimations for rare events at the tail of the PDF. Note that such an over-estimation of the rare high gradients has been observed in previous GAN models for the velocity gradients as well and is thought to be related to the balance coefficients in the loss functions, namely λ_{pixel} , λ_{grad} , and λ_{adv} [19]. Another reason could be the lack of training data points at rare events which can be mitigated by artificially increasing the number of training samples which have such properties. In Fig. 5a, the conditional average of dot products of the mixture fraction gradient is shown. The gradients from the LR field ($\xi = \tilde{Z}$) are lower than the ones in the DNS and SR across the whole mixture fraction space which is due to the smoothing effect of filtering on the gradients. The SR fields, however, can recover the vanishing gradients of the LR field and match well with the DNS counterparts. Slight deviations can be observed at rich conditions which highlight the point that the over-estimations observed and discussed in the PDF actually occur in rich conditions.

In Figs. 6 and 7a the performance of the model when reconstructing the scalar fields at particle positions is evaluated. The reconstruction of the particles' neighbour information is important as this data can be used to improve the Lagrangian particle heat and mass transfer modelling. The PDF is obtained by restricting the test samples to the ones with maximum Z exceeding 0.055 and further masking the data points in each slice based on the presence of particles. In this way

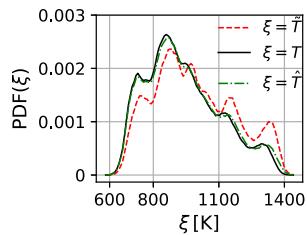


Fig. 6. PDF of temperature at the particle positions.

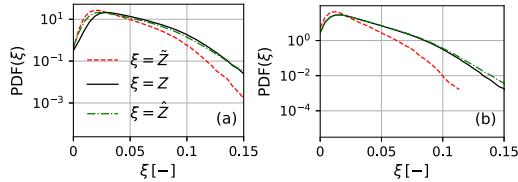


Fig. 7. PDF of mixture fraction from the samples restricted to the grid points where particles are present (a) and not present (b).

the test is restricted to particle-containing cells and later stages of the devolatilisation. In both figures it can be seen that the PDF of LR temperature has overshoots and undershoots which is the result of the filtering. On the other hand, there is a very good agreement in temperature predictions from the SR field compared to the DNS (Fig. 6). However, discrepancies can be observed in mixture fraction predictions in Fig. 7a. The generator detects LR overshoots and undershoots in lean and rich conditions, respectively, in Fig. 7a and tries to improve the reconstructed fields at particle locations by decreasing/increasing the mixture fraction in the corresponding regions. However, the network's confidence in the reconstruction is moderate only and this is why the final solution in SR lies in-between the DNS and the LR. Hence, the final distribution is biased towards low mixture fraction values when compared to the DNS in Fig. 7a. This behaviour is again attributed to λ values in the loss function. These coefficients aim at balancing the regression minimisation process between pixel-wise errors and statistical errors, e.g. PDF distributions. The deconvolution task is a one-to-many problem. A balance should be made between choosing the averaged solution among all possible solutions and choosing a single solution from a specific distribution which has the best statistical match. The choice made in this study is to balance the loss contributions such that all have comparable values in the order of magnitudes at the start of the training process. However, these hyper-parameters can be tuned via e.g. Bayesian hyper-parameter optimisation techniques which is the focus of future studies. Overall, the predictions are satisfactory considering the underdetermined nature of the deconvolution problem.

In Fig. 7b the same sample slices as in Fig. 7a have been used, however, the PDF of mixture fraction at the computational cells where no particle is present has been evaluated. It can be seen that the predictions with the SR data are matched perfectly with the DNS, except at very rich conditions. Comparing Fig. 7a and Fig. 7b shows that the overshoots observed at high mixture fraction values in Fig. 4 and Fig. 5 are caused by slight over predictions in particle-free cells.

We now compare the model performance against traditional GAN models with their generator only receiving low-resolution scalars, without conditioning and particle information input. The generator and discriminator architectures follow [2,19]. The model is pre-trained with pixel and gradient losses for 600 epochs, fine-tuned with an adversarial loss for 120 extra epochs, using the same loss weights as in [19], labelled “tradGAN”, and optimised weights ($\lambda_{adv} = 6 \cdot 10^{-2}$), labelled “tradGAN,opt” in Fig. 8. As can be seen, while the PACASRGAN (green line) recovers both under-resolved and un-resolved parts of the spectra, the other models fail, especially in recovering the unresolved part.

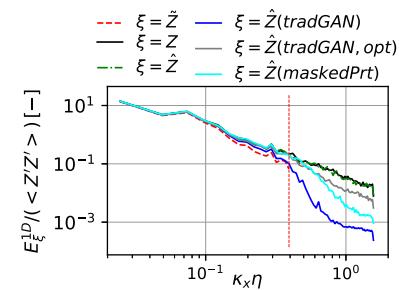


Fig. 8. Energy density spectra of mixture fraction. The test slice was randomly selected from the test dataset.

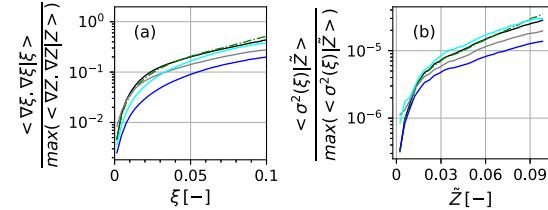


Fig. 9. Conditional averages of scalar dissipation rate (a) and sub-grid variance of mixture fraction evaluated on the whole test dataset using different models. The line labels are similar to Fig. 8.

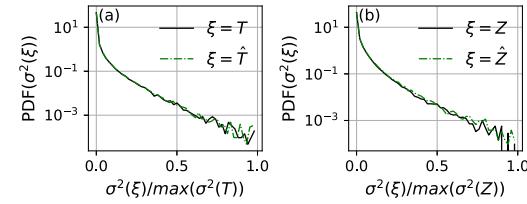


Fig. 10. PDF of subgrid variance of temperature (a) and mixture fraction (b) evaluated on OT1 and OT2 datasets.

Furthermore, it can be seen in Fig. 8 that when additional particle data, i.e. particle mask (M_p), and the vector containing particle-related information (P), are masked (zeroed out) from the PACASRGAN during the inference, the results become considerably worse (“maskedPrt”). This shows that the network effectively utilises such information. The underlying mechanism is that, at the finite Stokes numbers considered here, particles create localised sources/sinks via two-way momentum and heat/mass-transfer coupling. These processes sharpen scalar gradients and modulate subgrid variability. The particle mask and feature vector encode this localisation to improve the predictions.

In Fig. 9, the conditional averages of scalar dissipation rate and sub-grid mixture fraction variance are shown. The large under-prediction by the traditional GAN model (blue line) in Fig. 9(a) is due to the over-smoothed fields generated by this model which is also seen in Fig. 9(b). The performance can be improved by changing the loss weights (grey lines), however, it is still worse than the PACASRGAN. Moreover, the masking analysis shows that particle information is indeed used to improve both sub-grid variances (cyan line in Fig. 9(b)), and overall scalar dissipation rate (cyan line in Fig. 9(a)).

Finally, in Fig. 10, the generalisability of the model with out-of-sample tests for the devolatilisation process in an environment which is not seen during the training is assessed. The PDF of subgrid variances of both scalars computed over 1024 randomly selected samples from OT1 and OT2 DNS cases are recovered reasonably well. The spectra of mixture fraction and temperature (not shown for the sake of brevity) were found to also match reasonably well with the DNS.

5. Summary and conclusions

In this work, a particle-aware conditional attention-based super-resolution GAN (PACASRGAN) model is proposed for the super-resolution of density, temperature (T) and mixture fraction (Z) fields, generated in the devolatilisation process in a hot turbulent environment which is the initiating process of turbulent pulverised biomass combustion. Carrier-phase DNSs of pulverised biomass particles pyrolysis at various carrier-gas temperatures, and particle mass loadings, with fixed target Reynolds ($Re_\lambda = 53$) and Stokes ($St_\eta = 10$) numbers are carried out to generate the training/testing datasets. A modified RRDB generator network augmented with channel and spatial attention modules receives LR gas scalar fields, as well as particle information mapped to the LR grid. A UNet discriminator is conditioned on high frequency information of the HR input data, the LR counterparts, and the particle existence in a cell map at the LR level and decides whether the generated data by the generator are fake or real. Both networks are trained simultaneously in an adversarial manner with weighted pixel-wise, gradient, and adversarial losses. The results show that the proposed model can super-resolve the scalar fields of interest with a good accuracy and outperforms unconditional GAN models by effectively utilising particle information. Specifically, the scalar spectra, conditional averages of unresolved scalar variances, PDF and conditional average of the square of the mixture fraction gradient from the reconstructed fields follow the DNS values reasonably well. Deviations from the DNS are observed at rich conditions and for rare events at the tail of the square of Z gradient PDF. The investigation of the reconstruction of Z and T in the particles neighbourhood shows that the PDF of T predictions follows the DNS values very well. However, over-/under-estimation at lean/rich mixtures exists for the PDF of mixture fraction at particle positions. Such a behaviour is attributed to the weight of individual loss function components. These hyper-parameters can be tuned via e.g. Bayesian hyper-parameter optimisation techniques which is the focus of future studies.

Novelty and significance statement

1. Super-resolution (SR), an advanced deconvolution method based on machine-learning, is applied for the first time to turbulent two-way coupled particle-laden flows.
2. A new SR model is proposed to consider particle-flow interactions.
3. The pyrolysis in pulverised biomass combustion (PBC), an appealing option for sustainable energy transition, is considered.
4. Multiple pulverised biomass particle-laden DNSs are conducted to create the training/testing dataset, which is a valuable particle-laden dataset for future ML studies.

CRediT authorship contribution statement

A. Shamooni: Conceptualization, Methodology, Software, Validation, Investigation, Formal analysis, Writing – original draft, Writing – review & editing. **R. Cheng:** Software, Validation, Formal analysis, Writing – review & editing. **T. Zirwes:** Supervision, Writing – review & editing. **O.T. Stein:** Conceptualization, Funding acquisition, Project administration, Supervision, Writing – review & editing. **A. Kronenburg:** Conceptualization, Funding acquisition, Project administration, Supervision, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We acknowledge the financial support by the German Research Foundation, project No. 513858356, and the China Scholarship Council, Grant No. 202206020071. O.T. Stein acknowledges the financial support by the Helmholtz Association of German Research Centers (HGF), within the research field *Energy*, program *Materials and Technologies for the Energy Transition* (MTET), topic *Resource and Energy Efficiency*. The HPC resources are provided by HLRS and HoreKa.

Data availability

The code and data used in this study are available at the data repository of the University of Stuttgart (DaRUS) at <https://doi.org/10.18419/DARUS-5519> and can be downloaded via access request through the system.

References

- [1] K. Fukami, K. Fukagata, K. Taira, Super-resolution analysis via machine learning: a survey for fluid flows, *Theor. Comput. Fluid Dyn.* 37 (4) (2023) 421–444.
- [2] M. Bode, M. Gauding, Z. Lian, D. Denker, M. Davidovic, K. Kleinheinz, J. Jitsev, H. Pitsch, Using physics-informed enhanced super-resolution generative adversarial networks for subfilter modeling in turbulent reactive flows, *Proc. Combust. Inst.* 38 (2) (2021) 2617–2625.
- [3] M. Bode, M. Gauding, D. Goeb, T. Falkenstein, H. Pitsch, Applying physics-informed enhanced super-resolution generative adversarial networks to turbulent premixed combustion and engine-like flame kernel direct numerical simulation data, *Proc. Combust. Inst.* 39 (4) (2023) 5289–5298.
- [4] L. Nista, C. Schumann, T. Grenga, A. Attili, H. Pitsch, Investigation of the generalization capability of a generative adversarial network for large eddy simulation of turbulent premixed reacting flows, *Proc. Combust. Inst.* 39 (4) (2023) 5279–5288.
- [5] T. Grenga, L. Nista, C. Schumann, A.N. Karimi, G. Scialabba, A. Attili, H. Pitsch, Predictive Data-Driven Model Based on Generative Adversarial Network for Premixed Turbulence-Combustion Regimes, *Combust. Sci. Technol.* 195 (15) (2023) 3923–3946.
- [6] R. Cheng, A. Shamooni, T. Zirwes, A. Kronenburg, Improved super-resolution reconstruction of turbulent flows with spectral loss function, *Phys. Fluids* 37 (3) (2025) 035208.
- [7] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial networks, *Commun. ACM* 63 (11) (2020) 139–144.
- [8] M. Bode, Applying physics-informed enhanced super-resolution generative adversarial networks to finite-rate-chemistry flows and predicting lean premixed gas turbine combustors, 2022.
- [9] M. Bode, Applying physics-informed enhanced super-resolution generative adversarial networks to turbulent non-premixed combustion on non-uniform meshes and demonstration of an accelerated simulation workflow, 2022.
- [10] T.D. Luu, A. Shamooni, O.T. Stein, A. Kronenburg, S. Popp, H. Nicolai, H. Schneider, X. Wen, C. Hasse, Flame characterisation of gas-assisted pulverised coal combustion using FPV-LES, *Proc. Combust. Inst.* 39 (3) (2023) 3249–3258.
- [11] H. Tofighian, J.A. Denev, N. Kornev, A conditional deep learning model for super-resolution reconstruction of small-scale turbulent structures in particle-laden flows, *Phys. Fluids* 36 (11) (2024).
- [12] M. Bode, AI super-resolution: Application to turbulence and combustion, in: N. Swaminathan, A. Parente (Eds.), *Machine Learning and Its Application To Reacting Flows: ML and Combustion*, Springer International Publishing, 2023, pp. 279–305.
- [13] M. Bode, Applying Physics-Informed Enhanced Super-Resolution Generative Adversarial Networks to Large-Eddy Simulations of ECN Spray C, *SAE Int. J. Adv. Curr. Pr. Mobil.* 4 (6) (2022) 2211–2219.
- [14] M. Bode, J.H. Göbbert, Acceleration of complex high-performance computing ensemble simulations with super-resolution-based subfilter models, *Comput. & Fluids* 271 (2024) 106150.
- [15] A. Shamooni, R. Cheng, T. Zirwes, H. Tofighian, O.T. Stein, A. Kronenburg, Super-resolution of turbulent velocity fields in two-way coupled particle-laden flows, *Phys. Fluids* 37 (9) (2025) 093383.
- [16] X. Wang, L. Xie, C. Dong, Y. Shan, Real-ESRGAN: Training Real-World Blind Super-Resolution with Pure Synthetic Data, *Proc. IEEE Int. Conf. Comput. Vis.* (2021) 1905–1914.
- [17] S. Woo, J. Park, J.-Y. Lee, I.S. Kweon, CBAM: Convolutional block attention module, in: *Proceedings of the European Conference on Computer Vision, ECCV, 2018*, pp. 3–19.
- [18] D. Sundararajan, *Discrete wavelet transform: A signal processing approach*, CourseSmart Series, Wiley, 2016.

[19] L. Nista, H. Pitsch, C.D.K. Schumann, M. Bode, T. Grenga, J.F. MacArt, A. Attili, Influence of adversarial training on super-resolution turbulence reconstruction, *Phys. Rev. Fluids* 9 (6) (2024) 064601.

[20] G. Arumapperuma, N. Sorace, M. Jansen, O. Bladek, L. Nista, S. Sakhare, L. Berger, H. Pitsch, T. Grenga, A. Attili, Extrapolation Performance of Convolutional Neural Network-Based Combustion Models for Large-Eddy Simulation: Influence of Reynolds Number, Filter Kernel and Filter Size, *Flow, Turbul. Combust.* 115 (2025) 1261–1290.

[21] L. Nista, C. Schumann, G. Scialabba, T. Grenga, H. Pitsch, A. Attili, The influence of adversarial training on turbulence closure modeling, in: AIAA SCITECH 2022 Forum, 2022, pp. 1–9.

[22] L. Nista, C.D. Schumann, P. Petkov, V. Pavlov, T. Grenga, J.F. MacArt, A. Attili, S. Markov, H. Pitsch, Parallel implementation and performance of super-resolution generative adversarial network turbulence models for large-eddy simulation, *Comput. & Fluids* 288 (2025) 106498.

[23] A. Shamooni, X. Wen, P. Debiagi, A. Stagni, J.W. Gärtnner, T. Zirwes, O.T. Stein, C. Hasse, A. Kronenburg, Carrier-phase direct numerical simulation and flamelet modeling of alkali metal emissions from pulverized biomass flames, *Proc. Combust. Inst.* 40 (1) (2024) 105309.

[24] P.L. Carroll, G. Blanquart, A proposed modification to Lundgren's physical space velocity forcing method for isotropic turbulence, *Phys. Fluids* 25 (10) (2013) 105114.

[25] A.H. Abdelsamie, C. Lee, Decaying versus stationary turbulence in particle-laden isotropic turbulence: Turbulence modulation mechanism, *Phys. Fluids* 24 (1) (2012).