




High-precision inversion of dynamic radiography using hydrodynamic features

MALIHA HOSSAIN,^{1,*} BALASUBRAMANYA T. NADIGA,² OLEG KOROBKIN,² MARC L. KLASKY,² JENNIFER L. SCHEI,² JOSHUA W. BURBY,² MICHAEL T. MCCANN,²  TREVOR WILCOX,² SOUMI DE,² AND CHARLES A. BOUMAN¹

¹Purdue University, University, West Lafayette, IN 47907, USA

²Los Alamos National Laboratory, Los Alamos, NM 87545, USA

*mhossain@purdue.edu

Abstract: While radiography is routinely used to probe complex, evolving density fields in research areas ranging from materials science to shock physics to inertial confinement fusion and other national security applications, complications resulting from noise, scatter, complex beam dynamics, etc. prevent current methods of reconstructing density from being accurate enough to identify the underlying physics with sufficient confidence. In this work, we show that using *only* features that are robustly identifiable in radiographs and combining them with the underlying hydrodynamic equations of motion using a machine learning approach of a conditional generative adversarial network (cGAN) provides a new and effective approach to determine density fields from a dynamic sequence of radiographs. In particular, we demonstrate the ability of this method to outperform a traditional, direct radiograph to density reconstruction in the presence of scatter, even when relatively small amounts of scatter are present. Our experiments on synthetic data show that the approach can produce high quality, robust reconstructions. We also show that the distance (in feature space) between a testing radiograph and the training set can serve as a diagnostic of the accuracy of the reconstruction.

© 2022 Optica Publishing Group under the terms of the [Optica Open Access Publishing Agreement](#)

1. Introduction

In many scientific applications arising in material science, shock physics, inertial confinement fusion, and in national security applications including stockpile stewardship, radiography is routinely used to probe complex, evolving density fields and develop understanding of the underlying physics. Even as the reconstruction of the density of an object from line-integrated radiographic projections has a long history dating back to Radon [1], the noisy and complex multi-scale and multi-physics environment in many of these applications makes an accurate reconstruction of the density field a challenging task. To this end, we focus on the reconstruction of dynamic radiographic experiments at the Los Alamos National Laboratory Dual-Axis Radiographic Hydrodynamic Test (DARHT) facility. This facility serves as the nations premier facility for providing confidence in the nations nuclear arsenal. As such, it is of paramount importance to obtain the highest degree of confidence in the reconstructions obtained from the radiographic data.

Over the past several decades image reconstruction methods have evolved from simple analytical methods such filtered-back projection (FBP) methods, X-ray CT (e.g. Feldkamp-Davis-Kress or FDK methods), and the Inverse Abel Transform for axisymmetric systems [2–4]. These methods rely on relatively simple mathematical models of the imaging system. Modern approaches to solving the radiographic inversion problem for poly-energetic radiographic systems work with complex non-linear, non-convex forward models and employ iterative reconstruction techniques [5]. Iterative reconstruction algorithms are based on more sophisticated models for the imaging

system's physics and models for sensor and noise statistics. These methods are often called model based image reconstruction (MBIR) methods or statistical image reconstruction (SIR) methods [6].

Mathematically, we may express the image reconstruction problem as:

$$\hat{x} = \arg \min_x A(x) + \sum_i \alpha_i R_i(x) \quad (1)$$

where \hat{x} denotes the reconstructed density profile, $A(x)$ is the data fidelity term capturing the forward model of the imaging process and the statistical models of measurements and noise, and the $R_i(x)$ are regularizer models for the desired image. The scalar parameters α_i control the relative strength of the respective regularizers.

Figure 1 depicts the objective of the forward modeling approach. These algorithms iteratively estimate the unknown image based on the system (physical or forward) model, measurement statistical model, and assumed prior information about the inspected object. For example, minimizing penalized weighted-least squares (PWLS) cost functions has been popular in X-ray CT, and these costs include a statistically weighted quadratic data-fidelity term (capturing the imaging forward model and noise variance) and a penalty term called a regularizer that models the prior information about the object. [6]

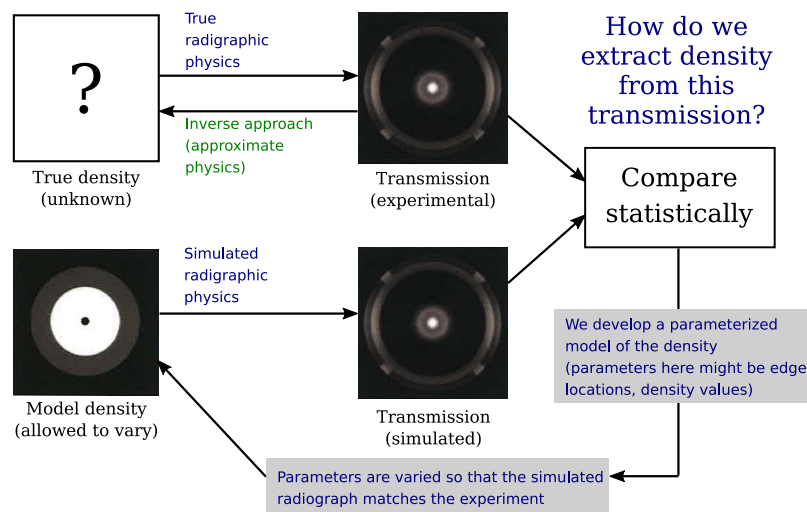


Fig. 1. Forward modeling approach to radiographic reconstruction.

The determination of the correct density distribution from 1 is contingent not only upon formulating a proper physics-based forward model, but also on imposing appropriate regularization and effective schemes to solve the optimization problem.

The uncertainties in the reconstruction process arise from our inability to exactly represent various aspects of the radiographic measurement system, such as scatter, beam spot movement, beam-target interactions, beam dynamics repeatability, and aspects of the image formation process.

Indeed, one of the major sources of error in the density reconstructions is attributed to scattered radiation. Scatter typically creates a loss of contrast and leads to image artifacts such as cupping, shading, streaks, etc. Scatter is caused by many types of photon-matter interactions [7] including Compton scatter, Rayleigh scatter, pair production, scatter involving the scene/background, etc. Many techniques have been proposed for scatter correction [8–13] (cf. [14, 15] for detailed review) in medical imaging, nondestructive testing, and other settings. The more recent techniques

[12,13] learn or fit scatter models based on training data including those generated from Monte Carlo N-particle transport code (MCNP) simulations [16].

The complex nature of scatter makes it impossible to accurately capture it even with sophisticated transport simulations, e.g. MCNP simulations. The presence of anomalous scatter fields in experimental data, such as from scene scatter exacerbates the problem. Such fields may not be accurately characterized and removed. Moreover, existing scatter (or other noise) correction techniques are not yet able to provide highly accurate quantitative density reconstructions [13] and break down in the presence of anomalous fields including scatter attributed to surrounding structures. As an illustration we examined the propagation of a shock into a non-uniform density field created by the implosion of a steel shell. A shock is a type of propagating disturbance that moves faster than the local speed of sound in the medium. In our problem the shock is generated by the convergence of the steel shell at the axis. Figure 2 in the upper left panel shows the relative levels of direct and scattered radiation, as calculated by MCNP [16] for a density field generated from a hydrodynamic simulation using CTH [17].

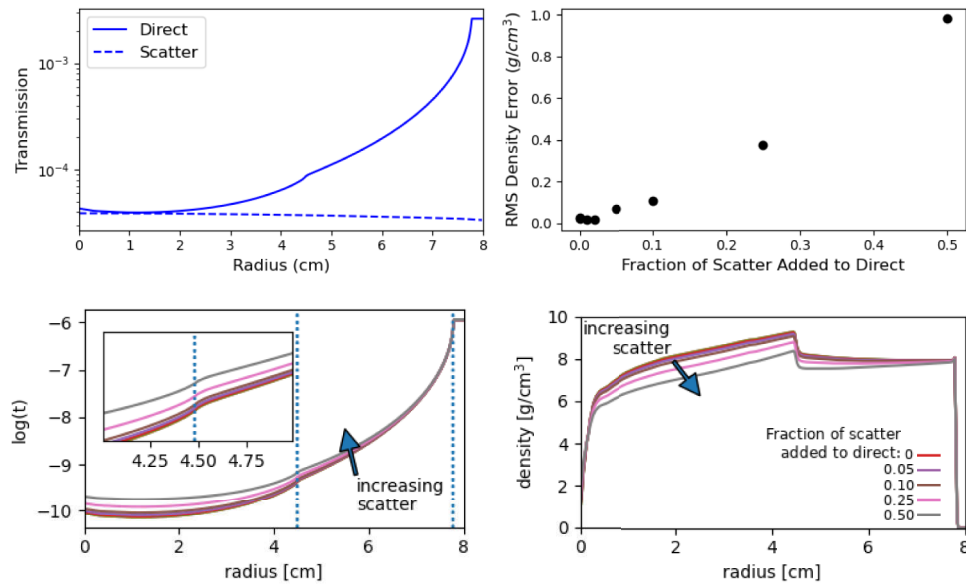


Fig. 2. Top left panel: transmission of the direct (solid) and scattered (dashed line) radiation through a sample object. Top right panel: the RMS error $\|\delta\rho\|_2$ between the descattered density and the ground truth, as a function of the fraction of scattered radiation added to the direct transmission. Bottom left: transmission for different added levels of scatter. Bottom right: descattered density profiles for the same levels of scatter. In transmission (bottom left), even in the presence of scatter, signatures of the hydrodynamic shock and edge (dotted lines) are clearly visible in the radiograph and are not displaced by the scatter (see the inset). For the density profiles (right), on the other hand, significant distortions are present after the density reconstruction, so that the RMS density error (top right) is high.

While the use of the traditional forward modeling approach to obtain quantitative density fields has been demonstrated to be sensitive to the ability of the forward model to accurately capture scatter and other radiographic imaging effects, there remain a few aspects of the radiograph that are relatively insensitive to these effects. Figure 2, bottom left panel, depicts the synthetic radiograph for a density field. Dotted lines indicate the hydrodynamic features, namely shock and edge, easily extracted using traditional edge finding techniques from synthetic radiographs, such as e.g. Canny edge detection [18]. These features are not significantly altered by the

presence of the scattered field. Figure 2 also shows the effect of varying levels of scatter on density reconstruction. The insensitivity of the location of the shock and edge to the level of scatter indicates the robustness of these features. Uncertainties in shock and edge locations arise from experimental aspects such as beam motion, magnification error, detector resolution, and image noise; however, density reconstructions based upon feature location, including errors, may provide a higher degree of accuracy compared to density reconstructions based upon a forward model or inversion of a de-scattered radiographic image.

Figure 2 supports an approach of focusing on the robust temporal sequence of shock and edge location features to perform density reconstructions. Accordingly, in this work, we explore the possibility that distinguished radiographic features, together with constraints imposed by dynamics, are sufficient by themselves to perform density reconstruction from radiographic sequences. We present a machine learning approach employing a conditional generative adversarial network (cGAN) to learn the mapping from robust radiographic features to the underlying density field.

This approach seeks to obtain density profiles that lie on a hydrodynamic manifold rather than the traditional approach of using individual radiographic time slices to perform independent density reconstructions. Namely, when employing dynamic radiography, the density being imaged evolves from one snapshot to the next, leading to a sequence of correlated radiographs. These correlations may be explained using models from continuum mechanics such as hydrodynamics or elastodynamics. Owing to these correlations, more information may be gleaned from a radiographic sequence than from the individual images regarded as independent observations.

The notion of a hydrodynamic manifold is a useful way to formalize the additional dynamical information present in a radiographic sequence of N snapshots when compared with N independent radiographs. The ideal Euler equations comprise an example continuum model governing density evolution in many realistic scenarios where dynamic radiography is applied.

Assuming that appropriate boundary conditions are known and fixed, there is a well-defined mapping $\mathcal{H} : \sigma \mapsto z$, from tuples of the form

$$\sigma \equiv (P, \rho_0, \mathbf{u}_0, e_0),$$

to solutions of the Euler equations,

$$z \equiv (\rho, \mathbf{u}, e),$$

where ρ is the mass density, \mathbf{u} is the fluid velocity, e is the specific internal energy, and $p = P(\rho, e)$ is the pressure expressed as a point-wise function of ρ and e . Here P denotes an equation of state and $(\rho_0, \mathbf{u}_0, e_0)$ denotes an initial condition for the Euler equations. Note that ρ_0, \mathbf{u}_0 , and e_0 are functions of space alone, while ρ, \mathbf{u} , and e are functions of both space and time. In full generality, the **hydrodynamic manifold** \mathcal{M} is the image of hydrodynamic solution mapping \mathcal{H} . Note that \mathcal{M} is a subset of a large function space containing arbitrary time-dependent density, velocity, and specific internal energy fields that satisfy the given boundary conditions. Among that most general class of fields, those contained in the hydrodynamic manifold are special because they solve the system of Euler equations, with *some* equation of state, and *some* initial condition.

The image of the space of physically admissible parameters Σ under the hydrodynamic solution map \mathcal{H} is the **restricted hydrodynamic manifold**, $\mathcal{M}_\Sigma \equiv \mathcal{H}(\Sigma)$. When the choice of Σ is understood from context, we will abuse terminology and simply refer to \mathcal{M}_Σ as “the” hydrodynamic manifold. In fact, the restricted hydrodynamic manifold \mathcal{M}_Σ is a smaller set than the hydrodynamic manifold \mathcal{M} , $\mathcal{M}_\Sigma \subset \mathcal{M}$. But \mathcal{M}_Σ is physically more meaningful since it excludes solutions of the Euler equations (or whatever continuum model is relevant to a specific problem) that could never occur in an experiment anyway.

In this article, our primary interest in hydrodynamic manifolds lies in their ability to distinguish arbitrary sequences of putative density profiles from sequences that are consistent with dynamics. We will say that a sequence of candidate densities $\{\rho_i\}_{i=1,\dots,N}$, separated temporally by an

increment Δt , **lies on the hydrodynamic manifold** if there is some $\sigma \in \Sigma$ and some $t_0 \in \mathbb{R}$ such that $(\rho, \mathbf{u}, e) = \mathcal{H}(\sigma)$ satisfies

$$\rho(\mathbf{x}, t_0 + [i - 1]\Delta t) = \rho_i(\mathbf{x}), \quad i = 1, \dots, N.$$

If a sequence does not lie on the hydrodynamic manifold, it is **dynamically inconsistent**. Since the hydrodynamic manifold is a reflection of prior physics knowledge, reconstructions of dynamic radiographs that lie on a hydrodynamic manifold should be preferred over reconstructions that do not.

This will enforce a physics-based constraint on the temporal evolution of the density field. To guarantee that the sequence of density fields obtained from the hydrodynamic features lie on a hydrodynamic manifold, we use parameter estimation, if a parametric form is available or potentially learned, to project our solutions onto an available hydrodynamic manifold.

In the remainder of this work we present details of our new approach for performing density reconstructions using the robust features extracted from the dynamic radiographic images. We introduce the dynamic radiography problem in Section 2. A test problem for examination of this approach is introduced in Section 3.

Section 4 presents two machine learning approaches using cGANs to allow for the determination of the density fields from a dynamic radiographic sequence from the robust hydrodynamic sequence. The results of the density reconstructions using the cGANs are presented in 5 along with projections of the cGAN predictions to the hydrodynamic manifold. Section 6 presents a discussion of the observations and conclusions drawn from the numerical experiments.

2. Dynamic radiography

A typical limited-view dynamic radiography problem may be described as follows. An experimenter conducts an experiment during which an object with density $\rho(\mathbf{x}, t)$ evolves in time. Here $\mathbf{x} \in Q$ is a point in the experimental chamber Q and $t \in [T_1, T_2]$ denotes time chosen from the interval $[T_1, T_2]$ over which the experiment takes place. At some instant $t_0 \in [T_1, T_2]$, a sequence of N X-ray pulses with temporal spacing Δt begins to impinge on the object. Each pulse passes through the object, suffering attenuation determined by the optical depth of the target material. Upon emerging from the object, and potentially scattering off obstacles peripheral to the experiment, the attenuated pulses interact with the detector, producing one in a sequence of radiographic images. While the detector positions and orientations are known, small uncertainties in propagation characteristics of X-rays through the material (e.g. absorption cross sections) may produce significant uncertainties in modeled transmission amplitudes. Moreover, details of environmental scattered radiation are difficult or even impossible to describe from first principles. Additional radiographic issues that impede the ability to perform accurate density reconstructions include beam spot movement, beam energy spectrum fluctuations, beam-target interactions, and beam dynamics repeatability. The problem then is to determine the density field at the times $t_i = t_0 + [i - 1]\Delta t$, $i = 1, \dots, N$, given the radiographic sequences. Figure 3 presents a sequence of density fields evolving in time along with the accompanying radiographic images. The focus of our investigations will be on the imaging of axisymmetric objects, hence a single view of the object is sufficient to perform the density reconstruction if a perfect forward model were known.

While there have been some attempts to incorporate dynamic constraints in the radiographic community, [19–26] most established techniques for dynamic radiographic X-Ray reconstruction fail to exploit knowledge of laws underlying the time evolution of a density field to enhance accuracy. This dynamical information that relates different images in a radiographic sequence places constraints on candidate density fields that analysis of each radiograph in isolation cannot detect. It is therefore reasonable to suspect that new, more performant reconstruction techniques may be devised by incorporating these constraints.

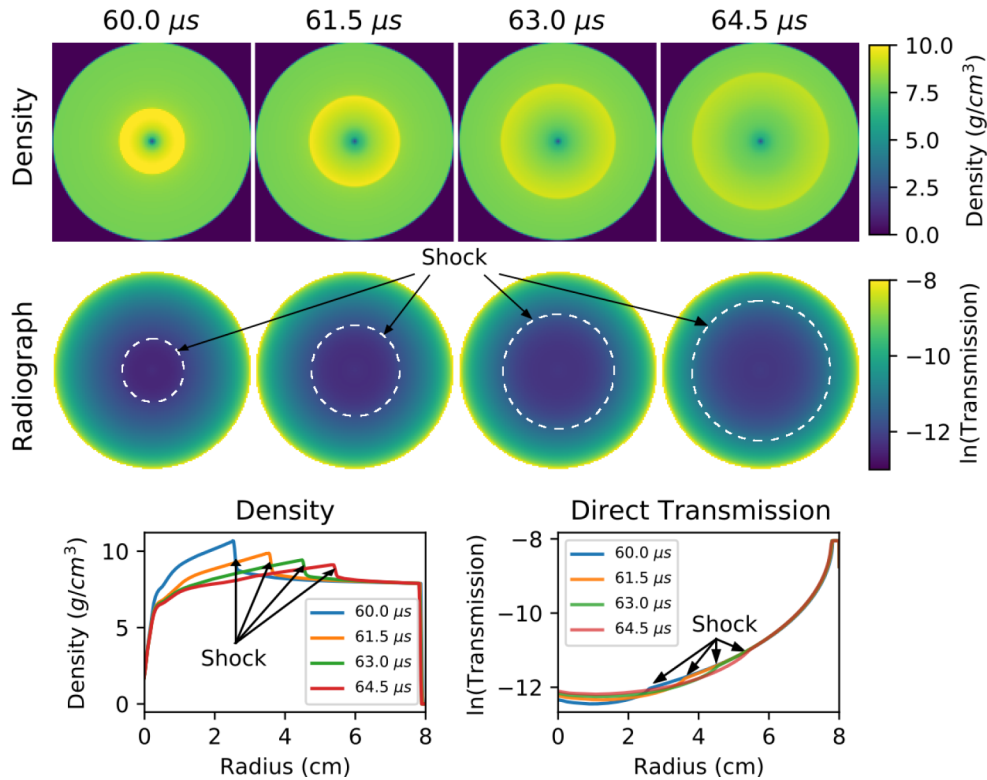


Fig. 3. Density fields and total transmission radiographs from a series of four 1D hydrodynamic simulations.

3. Test problem

To test our new approach to dynamic radiography, we studied the propagation of shocks generated by an imploding steel shell. We limited our study to 1D spherical symmetry, a Mie-Grüneisen (MG) equation of state (EOS), and a Steinberg-Guinan strength model applicable to metals at high strain rates [27]. To simulate implosion dynamics, we ran the CTH hydrodynamic code [28] on a uniform grid with 650 cells extending from the origin to a maximum radius of $R_D = 16$ cm, giving a grid cell size $\Delta x = 246.15$ μm . Initial conditions for the shell were chosen as follows: uniform density equal to that of steel under standard atmospheric conditions, inner and outer radii of $R_{\text{in}} = 8$ cm and $R_{\text{out}} = 10$ cm, respectively, and a uniform implosion velocity of $v_{\text{impl}} = -675$ m s^{-1} . (See Fig. 4.)

For each simulation, we generated synthetic radiographs at a series of times between 63 μs – 64.5 μs in 0.5 μs intervals using the Monte Carlo particle transport simulation software MCNP6 [16]. From these synthetic images, we extracted shock and outer edge locations using standard edge-detection methods. (see Fig. 3.) The object center was placed 133 cm from the source and the detector was placed 50 cm from the object, resulting in a magnification of 1.376. We irradiated the object with a 4 MeV mono-energetic X-ray source and generated a synthetic radiographic image using the radiography tally. This is a quasi-deterministic calculation that maps each particle to each pixel. Radiographs of direct and scattered radiation were generated separately and $1\text{e}3$ and $5\text{e}5$ particles were used, respectively. Additionally, the scatter simulations employed variance reduction by using the PDS card, which calculates the probability of all possible interactions and significantly reduces rare event fliers from Compton scattering. [16] Experimental contributions

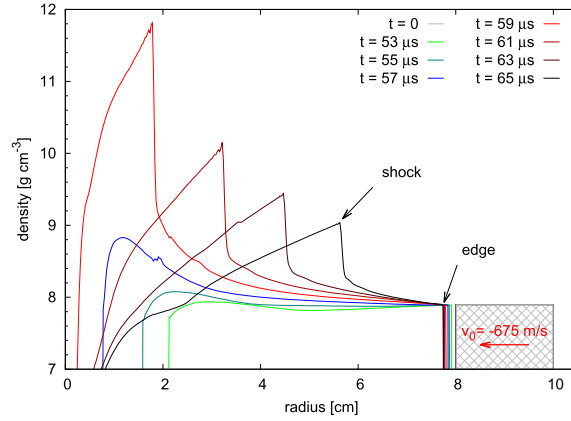


Fig. 4. Evolution of the radial density profile in a representative case. Initially, the material (steel) is uniformly distributed in a spherical shell, which is given a negative constant radial velocity, to initiate an implosion. Once the inner edge bounces, an expanding shock is formed propagating in the non-constant dynamic density background. We concentrate on the discontinuous features of the post-bounce solution, namely shock and outer edge (marked by arrows).

that degrade the image, such as source blur, energy-dependent detector efficiency, and noise, were not included in these simulations.

We generated a dataset of simulations by varying parameters characterizing the MG equation of state [17],

$$P\left(\chi := 1 - \frac{\rho_0}{\rho}, T\right) = \frac{\rho_0 c_s^2 \chi \left(1 - \frac{1}{2} \Gamma_0 \chi\right)}{(1 - s_1 \chi)^2} + \Gamma_0 \rho_0 c_V (T - T_0), \quad (2)$$

over a 5-dimensional (5D) cube. Here ρ_0 and T_0 are the reference density and temperature, respectively, c_s is the speed of sound, Γ_0 is the Grüneisen parameter at the reference state, s_1 is the slope of the linear shock Hugoniot, and c_V is the specific heat capacity at constant volume. Of the six parameters in 2, we kept the reference density ρ_0 fixed and varied the vector $\sigma = (T_0, c_s, s_1, \Gamma_0, c_V)$ around its nominal value listed in Table 1. Each component of σ was varied in 2% intervals within $\pm 10\%$ of the nominal value. Each component therefore assumed 11 discrete values, implying a total dataset comprising $11^5 \equiv 161051$ simulations.

Table 1. Nominal parameter values in the employed Mie-Grüneisen equation of state. The reference density parameter ρ_0 is fixed.

Material	ρ_0 [cm g ⁻³]	T_0 [K]	c_0 [mm/ μ s]	s_1	Γ_0	c_V [erg g ⁻¹ eV ⁻¹]
Steel	7.896	293.15	4.569	1.49	2.17	5.18×10^{10}

For testing purposes, we generated an 15 additional simulations, labeled J1-J15. Simulations J1-J10 were performed using the MG equation of state, while J11-J15 used a SESAME equation of state [29]. Parameters of these simulations are listed in Table 2.

Table 2. Test simulations J1-J15 used to evaluate the quality of density reconstruction and parameter estimation. All quantities are in terms of percentile deviations from the nominal values (see Table 1). Simulation database admits a $\pm 10\%$ variation in every parameter, so only the first four simulations (J1-J4) are within the bounds of the database. The last 5 test simulations (J11-J15) are performed with a different EOS. Finally, cases (J14-J15) also use a Preston Tonks Wallace material strength model. [30]

	EOS: Mie-Grüneisen					EOS: Sesame			
	T_0	c_s	s_1	Γ_0	c_v	material	strength model	scaling ratio	
J1	0.2	-6.1	6.0	-6.0	-10.0	J11	STEEL_HP	Steinberg	1
J2	0.2	5.1	6.0	-6.0	-10.0	J12	STEEL_HP	Steinberg	0.99
J3	0.2	5.1	5.4	-6.0	-10.0	J13	STEEL_HP	Steinberg	0.98
J4	0.2	5.1	5.4	5.1	-10.0	J14	STEEL_HP	PTW 304LSS1	0.98
J5	0.2	5.1	5.4	5.1	-31.2	J15	STEEL_HP	PTW 304LSS1	0.99
J6	-11.7	5.1	5.4	5.1	-31.2				
J7	-11.7	-16.8	5.4	5.1	-31.2				
J8	-11.7	-16.8	18.8	5.1	-31.2				
J9	-11.7	-16.8	18.8	-13.4	-31.2				
J10	-11.7	-16.8	18.8	-13.4	7.4				

4. Reconstruction using hydrodynamic features and learned dynamics: data-driven modeling

As discussed in the introduction, we are interested in reconstructing density fields ρ using only robust features of radiographs, which in the current setting consist of the shock and edge locations. As such, we formulate our learning problem as one of learning a map from sequences of shock and edge locations F to a full density field ρ and denote this process as $F \mapsto \rho$. For training data we use a set of density field sequences obtained from our simulation database and the corresponding shock and edge locations for each sequence.

A challenging aspect of these data is that the mapping from features to density can be one-to-many. Our approach to handling this aspect of the problem is to assume that the underlying (unknown) model generating the data is probabilistic. Thus, instead of learning a map $F \mapsto \rho$, we are led to learn to sample from the conditional probability $p(\rho | F)$ of encountering a density field ρ given a robust feature sequence F . In this setting, a (conditional; see later) generative model is called for. This is because at its core a generative model includes a representation of the *probability distribution* of the data, e.g., as indicated above.

While there are a handful of generative model classes ranging from generative adversarial networks (GANs) [31] to variational autoencoders [32] to continuous normalizing flows [33] to diffusion maps [34] and others, we focus attention on GANs given the extreme level of activity seen in this class of models since their introduction in 2014. GANs are a powerful class of generative models that are trained using an adversarial technique: while a discriminator continually learns to discriminate between the generator's synthetic output and true data, given a noise source, the generator progressively seeks to fool the discriminator by synthesizing yet more realistic samples. The process continues until a quasi-equilibrium is reached, wherein the discriminator's feedback to the generator does not permit further large improvements of the generator, but rather leads to stochastic variations that explore the local basin of attraction.

Indeed, in the preceding part of this section, when we have used the terminology of generative models and GANs, we have really been talking about the conditional form of the generative model(s) since that was natural for the problem on hand.

In the following section we explore two methods of generating a sequence of density fields conditioned on the shock and edge locations. Details of each method can be found in Appendices A.

and B.. In the first approach, we use a conditional WGAN (cWGAN) model, trained with shock and edge locations extracted with pixel level resolution. In the second approach, we use the subpixel feature extraction method (detailed in Appendix C.) to condition a cGAN. In both approaches, our cGANs accept as an input constraints the locations of the shock and edge at four predefined time steps, separated by $0.5 \mu\text{s}$. This imitates the scenario with DARHT, where four radiographs of the dynamical field are taken in rapid succession. Similar to [35], we incorporate an L^1 reconstruction loss term for the generator update, as well as additional discrete modules to enforce prior knowledge about the density reconstruction such as the known mass and feature locations. The examination of two different conditional GAN architectures, resolutions at which the features may be extracted from the radiographs, as well as inclusion of physics based constraints in one network architecture allows for a more thorough investigation of the properties of the solution space and behavior of our mapping from feature space to density fields.

5. Results

In this section, we describe and compare the results obtained from the two different network architectures.

5.1. cWGAN results

The left panel of Fig. 5 shows the discriminator loss as a function of training epoch. Training can be seen as consisting of three phases: an initial phase (up to about 100 epochs) when there are large changes given the random initial condition resulting both in a poor generator and the discriminator not knowing the Wasserstein metric. On incremental learning on the parts of both the generator and the discriminator, the initial phase of violent changes is followed a period of slower changes (between epochs 100 and 1000). In the final phase, there is an asymptotic approach to equilibrium characterized by extended periods of stable behavior of both the generator and discriminator. In this figure, it is further observed that when the coefficient associated with the gradient penalty λ in (5) is varied over three orders of magnitude ($0.01 \leq \lambda \leq 10$) the cWGAN trains reliably and stably without suffering any of the modes of failure such as mode-collapse that are common in the training of GANs. In the following, we show results for the smallest value of the gradient-penalty coefficient (0.01) while noting that the results for the other values of the gradient-penalty coefficient differ so slightly as to be visually indistinguishable most often for the diagnostics considered. This further confirms the robustness of the WGAN results presented and discussed.

The top set of density fields in the right panel of Fig. 5 correspond to the most frequently occurring time-sequence of shock and edge locations in the numerical simulations. When this sequence of shock and edge locations are input to the trained generator (of the cWGAN) along with a further random number, one set of four density profiles corresponding to the four shock and edge locations is obtained. The lower set of density fields show an ensemble of one hundred such predictions. This constitutes an initial verification of the ability of using a cWGAN to reconstruct a full density profile given just the shock and edge locations.

Next, we further examine the ability of the trained generator to reconstruct the density field given shock and edge locations where now the shock and edge locations come from test numerical computations, some of which use an equation of state that is significantly different from that used to generate the data that was used for training the cWGAN. We consider cases J1, J10, J11, and J15 in Table 2. Furthermore, since we have the density profiles for these cases, we additionally examine the ability of the trained discriminator to establish similarity to (or likewise, tell them apart from) the density profiles used for training.

Figure 6 compares the cWGAN reconstruction of the density profiles from the shock and edge locations for cases J1, J10, J11 and J15. Each panel in the Figure additionally indicates the root

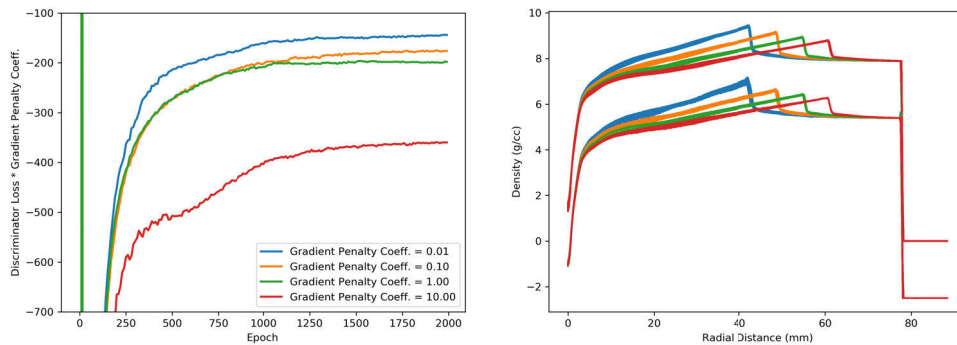


Fig. 5. The left panel shows the discriminator loss as a function of training epoch. The loss is averaged over ten epochs to reduce the number of data points. The averaging also leads to smoother loss curves. When the coefficient associated with the gradient penalty λ in (5) is varied over three orders of magnitude ($0.01 \geq \lambda \geq 10$) the cWGAN trains reliably and stably. The right panel shows two ensembles of density profiles for four successive time steps: 63 μs (blue), 63.5 μs (orange), 64 μs (green) and 64.5 μs (red). The top ensemble of density profiles represents reference computations, while the bottom ensemble of density profiles (offset by -2.5 g/cc to facilitate comparison) represents the cWGAN reconstructions. The right panel verifies the ability of the trained generator, given a sequence of shock and edge locations, to reconstruct a range of reference density profiles.

mean square error (RMSE) between the average of the ensemble of generated density profiles and the ground truth (see also Table 3).

The RMSE of the cWGAN density reconstruction in case J1 is small and the predicted profile is seen to match the reference numerical profile well. However, the RMSE is seen to be larger for cases J11 and J15 with corresponding large mismatches in the density behind the shock. Indeed, it turns out that the equation of state used in case J1 is well represented in the training data whereas that is not the case with cases J11 and J15.

This is further borne out in terms of the discriminator suggesting that while the density profile of case J1 is well within the kinds of density profiles seen in the training data that is not the case for J11 and J15. This can be seen in Fig. 7 which shows the discriminator function value for the training set of density profiles as a histogram and for each of the test cases J1 through J15 as delta functions.

In the noisy environment of a hydrotest, various effects such as scatter, beam spot asymmetry, and motion blur can affect the recovery of even the most robust of features such as the shock and edge locations. For example, in test case J1, experimentation with the level of scatter suggests that the location of the shock can be, depending upon experimental design, perturbed by as much as 300 microns whereas the location of the edge can be perturbed by as much as 100 microns (while noting that the hydro simulations here use a uniformly space grid with a mesh size of 246.15 microns).

Figure 8 shows the results of reconstructing the sequence of density profiles when the shock and edge locations are perturbed as stated above, but when the reference profiles are left unperturbed. Again a thousand perturbations were considered. On comparison with the top left panel of Fig. 10, changes in the location of the shock of individual profiles is evident; because of the large number of profiles plotted, this perturbation can only be seen as a thickening of the envelope of profiles at the location of the shocks. The result of the uncertainty of the edge location can also be seen, but is significantly smaller, and this is commensurate with the lower level of uncertainty in locating the edge due to scatter from radiographs. Small changes in the structure of the shock near the downstream end can also be seen at the different times of the sequence of ensemble-averaged

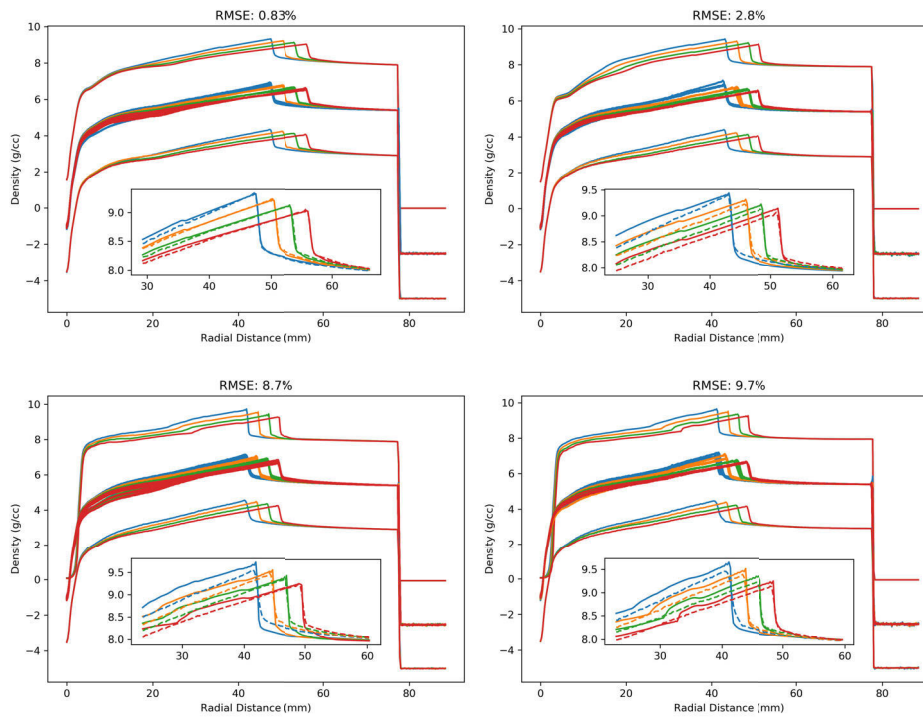


Fig. 6. Top: Cases 1 (left) and 10 (right); Bottom: Cases 11 and 15. In each of these figures, the top set of density fields come from the test numerical integrations and form the reference to judge the quality of the cWGAN predictions. The next set of density fields comprise an ensemble of a thousand predictions from the cWGAN; they are offset by -2.5 (g/cc) to facilitate comparison. The final set of density fields that are offset by a further -2.5 cm g^{-3} correspond to the mean over the thousand cWGAN predictions. The inset is a zoom of the reference (solid) and predicted (dashed) density fields at the four times considered. The colors for four successive time steps are the same as in Fig. 5. The RMSE is indicated on top, with the percentage taken with respect to the nominal density $\rho_0 = 7.896$ g/cc. Larger differences between the equations of state used in cases 11 and 15, as compared to the equations of state used for the training data lead to larger density reconstruction errors in those cases.

Table 3. Errors in density profiles for projections onto hydrodynamic manifold by parameter estimation: cWGAN (left) vs cGAN-DF (right). All quantities are in terms of percentile deviations from the nominal values (see Table 1), and $\rho_0 = 7.896 \text{ cm g}^{-3}$ is the nominal density.

	cWGAN			cGAN-DF		
	$\frac{\ \rho_{\text{proj}} - \rho_{\text{GT}}\ _2}{\rho_0}$	$\frac{\ \rho_{\text{GT}} - \rho_{\text{GAN}}\ _2}{\rho_0}$	$\frac{\ \rho_{\text{GAN}} - \rho_{\text{proj}}\ _2}{\rho_0}$	$\frac{\ \rho_{\text{proj}} - \rho_{\text{GT}}\ _2}{\rho_0}$	$\frac{\ \rho_{\text{GT}} - \rho_{\text{GAN}}\ _2}{\rho_0}$	$\frac{\ \rho_{\text{GAN}} - \rho_{\text{proj}}\ _2}{\rho_0}$
Minimization with L2 distance						
J1	0.85	1.47	1.10	0.83	0.87	0.41
J2	1.26	2.46	1.19	2.07	1.76	0.54
J3	1.44	2.43	1.10	1.24	1.11	0.79
J4	1.94	3.00	1.22	0.70	0.69	0.65
J5	1.60	2.65	1.97	1.94	1.75	0.59
J6	2.20	3.08	1.65	1.41	1.18	0.60
J7	5.09	4.15	2.18	8.20	8.94	2.15
J8	5.03	5.00	1.08	9.51	9.69	1.99
J9	6.03	5.93	4.24	8.04	8.08	1.60
J10	5.83	5.30	1.59	8.00	8.61	1.87
Minimization with Wasserstein distance						
J1	1.08	0.95	1.10	0.42	0.48	0.49
J2	1.26	1.61	1.19	1.75	1.20	1.23
J3	1.23	2.45	1.10	1.09	0.76	1.00
J4	1.29	1.75	1.22	1.04	0.51	1.06
J5	1.09	1.43	1.97	1.38	1.22	0.51
J6	1.14	1.73	1.65	1.67	0.98	1.12
J7	3.51	2.09	2.18	3.50	3.20	1.56
J8	4.24	1.02	1.08	3.74	3.03	1.92
J9	4.32	3.50	4.24	4.37	4.10	1.36
J10	4.43	1.36	1.59	3.53	3.21	1.27




Fig. 7. Discriminator function values. Histogram shows the probability distribution in the training dataset. Delta functions for test cases J1 through J15 are shown in different colors.

profiles in the zoomed inset. Finally, on quantifying the effect of uncertainty in the shock and edge location due to scatter, we see the overall RMSE increases by a modest amount of about 14% assuming the stated uncertainties in the shock and edge positions.




Fig. 8. Effect on the reconstruction of density of uncertainty in the location of shock and edge attributable to various sources, in the same format as in Fig. 6. A shock location uncertainty of 300 microns and an edge location uncertainty of 100 microns lead to significant changes in the reconstructed density profiles. The density reconstruction error increases by a modest amount of about 14%.

5.2. cGAN-DF results

All models, including \tilde{R} , are optimized using Adam [32] with a learning rate of 2×10^{-4} and a β_1 parameter 0.5. The cGAN-DF training is terminated when shock locations appear sharp, at 5000 epochs, to prevent overfitting due to the L1 loss term. All results utilize, $\lambda_D = 1$, $\lambda_{L1} = 100$, $\lambda_{\tilde{F}} = 10$, and $\lambda_M = 1$. The loss functions are shown in Fig. 9.

Figure 10 compares the generated density profiles with real snapshots from J1, J8, and J11. Simulation J1 is produced using the MG EOS and EOS parameters within the 5D hypercube used to generate the training data. The corresponding density predicted by the cGAN-DF has low RMSE. Here a predicted density is an average of a 1000 generated densities for given set of features. Parameters associated with J8 lie outside the hypercube but the density is still captured accurately. The cGAN-DF is less robust to model mismatch in J11 as it fails to be as accurate as J1 or J8 when the MG EOS is replaced by a different model.

The introduction of the additional loss terms in Eq. (11) may steer the generator away from learning a true distribution match with P , but as noted in [35], the L1 term in particular is necessary in order to produce convincing results. During training, the cWGAN encounters multiple non unique feature locations as $F(\rho)$ are quantized into pixel width intervals. In this context, an L1 loss would learn to fit a centroid of plausible ρ s associated with the conditional input, resulting in blurry generated samples. However, we observe that a larger λ_{L1} weight results in sharper shocks in the samples generated by the cGAN-DF, in addition to improved texture matching. This is a result of unique $F(\rho)$ values produced by the subpixel feature extraction method. This difference in the nature of the training data and the additional constraints introduced by L_M and $L_{\tilde{F}}$ in Eq. (11) limits the variation in the samples generated by the cGAN-DF for a

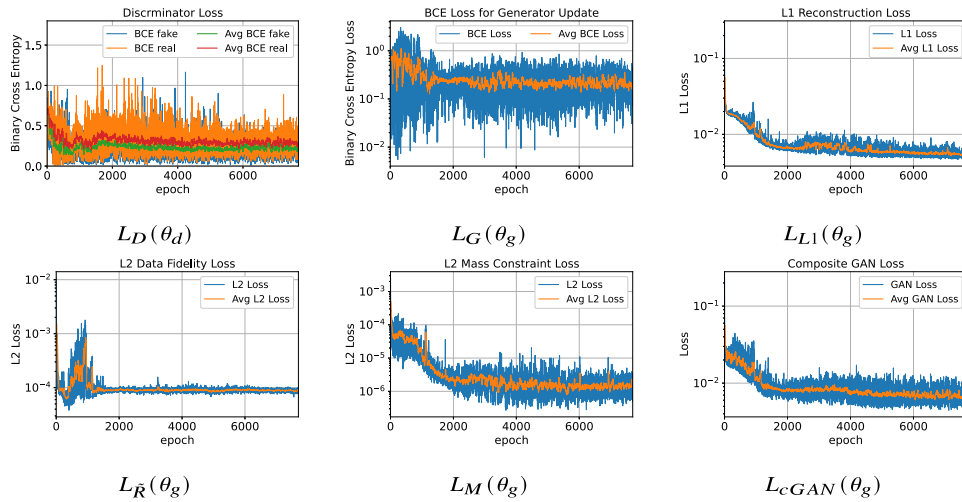


Fig. 9. cGAN-DF loss functions, plotted alongside running averages of 20 epochs for clarity.

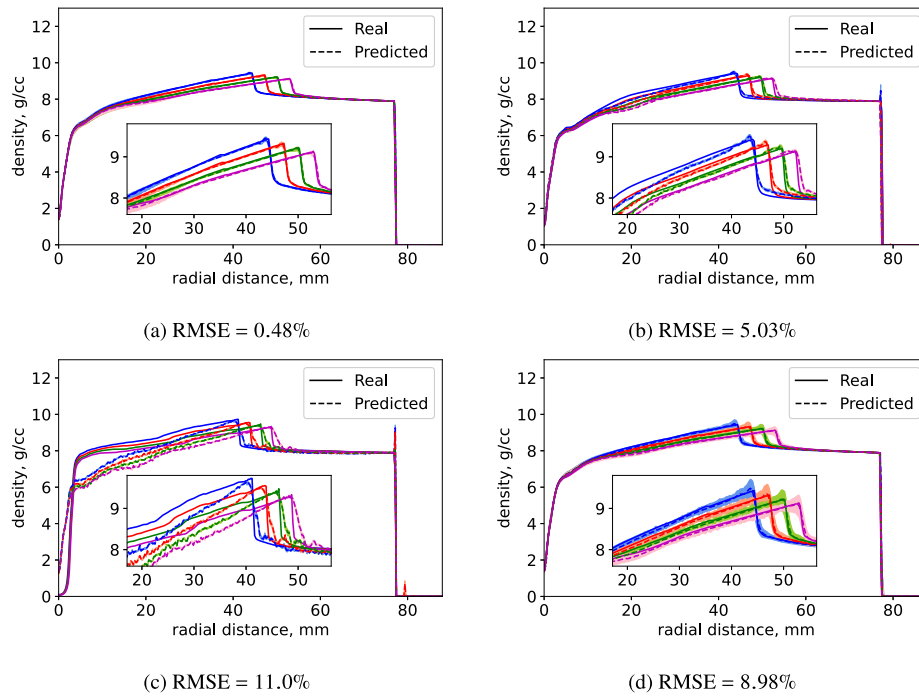


Fig. 10. Real (solid lines) and predicted (dashed lines) density profiles for four successive time steps: $63 \mu s$ (blue), $63.5 \mu s$ (red), $64 \mu s$ (green), and $64.5 \mu s$ (magenta). Inference on (a) J1, (b) J8, and (c) J11. (d) Inference on J1 using network with 30% dropout during testing. For all four cases, an ensemble of 1000 generated densities are plotted and are most visible in panels (a) and (d). The predicted result is an average of the 1000 generated samples.




Fig. 11. Left: errors in density profiles for test simulations J1-J10. Right: illustration of the L_2 distances between the ground truth (GT), GAN density reconstruction (GAN), and the projection of the GAN density profile (proj) onto the manifold (gray curve).

given set of features. As in [35], we also observe that dropout during testing can increase the entropy in generated samples (see Fig. 10(d)).

5.3. Projection of reconstructed densities onto hydrodynamic manifold

We have utilized the robust hydrodynamic features, i.e., the shocks and edges in conjunction with conditional GAN network architectures to learn the density sequences from a series of dynamic radiographs. While the density reconstructions appear to be satisfactory, there is no guarantee that these solutions will lie on a hydrodynamic manifold as given by the learned dynamic model.

Consequently, in this section we perform a projection of the density fields obtained from the learned model onto a hydrodynamic manifold by using parameter estimation. In this example we assume that we have a parametric form for the physics model, i.e., the MGR. Situations in which this is not the case may be addressed by learning a parametric model from the data or a neural network for example. [36–38]

Let Σ denote the space of admissible parameters for our test problem. Since we fix our initial conditions, each point $\sigma \in \Sigma$ is a vector of MG parameters $\sigma = (T_0, c_s, s_1, \Gamma_0, c_V)$. Given a numerical solution corresponding to the parameters σ and an initial measurement time t_0 in some discrete set T_{data} , we can then extract a density sequence $\rho_n^{(\sigma)}(r) = \rho(r, t_0 + [n-1]\Delta t)$ from this solution that is contained in the hydrodynamic manifold \mathcal{M} . For an unknown density sequence $\rho_{\text{target } n}$ we can estimate its parameters σ_{target} by solving the optimization problem

$$\sigma_{\text{target}} = \arg \min_{\sigma \in \Sigma_{\text{data}}} \|\rho_{\text{target}} - \rho_{\sigma}\|_2. \quad (3)$$

Here the combined L_2 -norm over the $N_t = 4$ time snapshots is defined as:

$$\|\delta\rho\|_2 := \sqrt{\frac{1}{N_t V} \sum_{n=1}^{N_t} \sum_{i=1}^{N_x} (\delta\rho(r_i))^2 \Delta V_i}, \quad (4)$$

with N_x being the total number of cells in radial direction, ΔV_i being the volume of an i -th cell, and V the total volume.

This method of parameter estimation gives accurate results for an ideal case when the density profiles (a) lie on a restricted hydrodynamic manifold; (b) contain no noise; and (c) are known to be within the range of parameteric hypercube of the simulations database. In such case, the accuracy of the parameter estimation is only limited by the error in interpolating between the neighboring entries in the database grid. We may also use parameter estimation for the predictions of the density sequences obtained from our cWGAN and cGAN-DF networks.

The results of the manifold projection by parameter estimation are presented in Table 3 for the test simulations J1-J10. Recall that only the first four simulations J1-J4 are within the bounds of the database ($\pm 10\%$ from the nominal values, see Table 2).

The three columns in the table list the L_2 density mismatch between three density profiles: ρ_{GT} corresponds to the “ground truth”, the true profile of the test; ρ_{GAN} is the GAN-reconstructed density profile, and ρ_{proj} represents the projection of the latter onto the dynamic manifold. As can be seen from the table and Fig. 11, in most cases the error of GAN reconstruction is comparable to the error of manifold reprojection.

6. Discussion and conclusion

In this study, we consider a parametric study of a spherically symmetric shell implosion problem with subsequent expansion, focusing on the shock expansion phase, in an attempt to reconstruct complete density profiles using only the discontinuous hydrodynamic features: shock and edge in conjunction with learned dynamics so as to enable dynamic density reconstructions. Our study is motivated by the fact that the features are normally much more easily identifiable in physical radiographs than removing scatter and other artifacts from the radiographic images, with the latter requiring solutions to the complex descattering, beam physics, detector modeling, and inversion problems to recover the density fields.

We have demonstrated that a conditional generative adversarial network (cGAN) can be successfully applied to recover the density from a sequence of four positions of shock and edge, produced from corresponding time sequence of radiographs. We used two different neural network architectures: conditional Wasserstein GAN (cWGAN, Section A.), and a cGAN with data fidelity (cGAN-DF, Section B.) and show that they produce results of comparable accuracy around 1% in RMSE, for models which lie on the restricted hydrodynamic manifold. For models which are beyond the boundaries of the hydrodynamic manifold, the networks do not perform as well (see Table 3), recovering density profiles with about 5%–8% accuracy. To put these numbers in perspective, Fig. 12 shows that achieving 3% RMSE requires removing the scatter down to $\lesssim 20\%$ level, which is challenging in many applications. Further, hands-on experience with radiography at DAHRT suggests that the RMS density error of standard descattering and reconstruction is around 0.8 g/cc for simple test objects with comparable scatter-to-direct ratio as in Fig. 2 (top left) to the objects used here; in comparison, our reconstruction errors are about 0.24 g/cc (see Table 3).

To better understand the properties of our underlying hydrodynamic simulations and also enable an understanding of the origin of the errors in the density reconstructions we examine the properties our hydrodynamic features, focusing on the velocity and acceleration of the shocks. Figure 13 depicts the point cloud of shock velocity and acceleration for all simulations in the database. Each point in the Figure represents a single simulation, with the velocity and acceleration of the shock fitted over a fixed time interval between 63.0 and 64.5 μs . It also marks the location of the test simulations J1-J10. Because they were performed with the same MG equation of state, they should also lie on the hydrodynamic manifold. One can clearly see from this visualization how J1-J6 are inside the point cloud, while J7-J10 (open circles) are not too far from it, in the sense that the data can be extrapolated to / resampled in that region to properly capture the shock behavior. This figure also shows the location of J11 on the acceleration-velocity map (red triangle), which was performed with an entirely different equation of state. It is clear that parameter estimation on this point can hardly produce meaningful results, because the point is too far away from the database. This figure illustrates the difficulties encountered when extrapolation from the manifold is required. Application of machine learning techniques cannot circumvent this problem, because extrapolation will still be necessary. However, to address this problem additional simulations using the same model that better enable the training data to include the testing data may be performed. Alternatively, alternative models may be utilized that

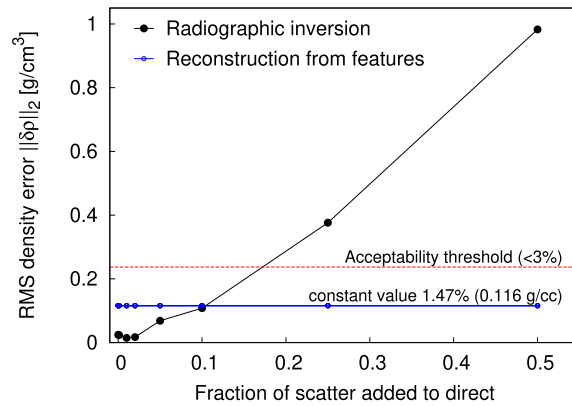


Fig. 12. The relative RMS error $\|\delta\rho\|_2/\rho_0$ of the reconstructed density as a function of scattered radiation added to the direct transmission, obtained with two different methods: by radiographic inversion (black circles), and from the features in radiograph using the cWGAN (blue circles). Effect of the added scatter on the location of the features is negligible, therefore the scatter does not affect the quality of the density reconstruction from features.

are closer to the desired location in the shock velocity acceleration phase space. This unique ability to examine the location of the observable features in the radiographs occurs to perform density reconstructions offers great potential in reducing density errors in reconstructions not available in the traditional dynamic radiographic reconstruction algorithms.




Fig. 13. Map of the fitted velocity and acceleration of the shock for all simulations in the Mie-Grüneisen database, test simulations J1-J10 and a test simulation J11 (red triangle) with a different equation of state (using Sesame tables: see Table 2). The fitting was performed in the same fixed time interval between 63 and 64.5 μs for all simulations. This plot illustrates why it may not be possible to meaningfully extrapolate the hydrodynamic path from the simulation database.

Finally, we remark that our investigation of the nature of degeneracies of the density fields has revealed that while sub-pixel information appears to remove some degeneracies, given by the results indicated by the cGAN investigations and introduction of the L1 loss, uncertainties undoubtedly will be present due to the finite resolution of the imaging system. Consequently, the use of the GAN architecture is a necessary component of the learning of features to density fields.

Appendix A. Conditional Wasserstein GAN (cWGAN)

As a first approach, we train a cWGAN conditioned on shock and edge locations. The discriminator loss function is

$$\mathbb{E}_{\rho \sim P, z \sim \mathcal{N}} [D_{\theta_d}(G_{\theta_g}(z, F(\rho)), F(\rho))] - \mathbb{E}_{\rho \sim P} [D_{\theta_d}(\rho, F(\rho))] + \lambda \mathbb{E}_{\rho \sim P_u} \left[(\|\nabla_{\rho} D_{\theta_d}(\rho)\|_2 - 1)^2 \right]. \quad (5)$$

Here, $\rho \sim P$ corresponds to real samples (that are numerically generated using the hydrodynamic solver and which are used for training the cWGAN) drawn uniformly at random from the training set, D_{θ_d} refers to the discriminator network with parameters θ_d , G_{θ_g} refers to the generator network with parameters θ_g , $z \sim \mathcal{N}(0, 1)$ is i.i.d. Gaussian noise, λ is a coefficient that corresponds to the penalty, and the final term in the sum is the gradient penalty that encourages the discriminator to be 1-Lipschitz; see [39] for further details.

In our experiments, the generator and discriminator were multi-layer perceptrons (MLP) with two hidden layers, each consisting of 512 neurons; a ReLU activation was used and a dropout of 10% of the nodes was implemented to discourage overfitting.

Appendix B. Conditional GAN with data fidelity constraints (cGAN-DF)

Unlike the cWGAN case, in the traditional GAN setting further modifications of the generator loss function were required to obtain meaningful results and these are described below. In this second approach, we build on a traditional cGAN by incorporating modules that enforce data fidelity constraints. Figure 14 presents the cGAN-DF architecture.




Fig. 14. Discriminator and generator architectures for cGAN-DF.

The discriminator loss function is

$$- \mathbb{E}_{\rho \sim P} [\log(D_{\theta_d}(\rho, F(\rho)))] - \mathbb{E}_{\rho \sim P, z \sim \mathcal{N}} [\log(1 - D_{\theta_d}(G_{\theta_g}(z, F(\rho)), F(\rho)))] \quad (6)$$

where $\rho \sim P$ denotes a density drawn uniformly at random from the training set, $z \in \mathbb{R}^{128}$ is i.i.d. Gaussian noise, i.e., $z \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{128})$. $F(\rho) \in \mathbb{R}^{2 \times 4}$ are instances of shock and edge locations on which the cGAN-DF is conditioned, and $\rho \in \mathbb{R}^{360 \times 4}$ are instances of 1D density profiles produced via simulation. Likewise, $\tilde{\rho} = G_{\theta_g}(z, F(\rho))$ are generated (also called fake) densities.

In a traditional cGAN, the loss function for the generator is

$$L_G(\theta_g) = \mathbb{E}_{\rho \sim P, z \sim \mathcal{N}} [\log(1 - D_{\theta_d}(G_{\theta_g}(z, F(\rho)), F(\rho)))] . \quad (7)$$

We augment this loss function with an $L1$ reconstruction loss term as in [35]

$$L_{L1}(\theta_g) = \mathbb{E}_{\rho \sim P, z \sim \mathcal{N}} [\|\rho - G_{\theta_g}(z, F(\rho))\|_1] \quad (8)$$

Prior to training the cGAN-DF, density simulations from the training set are processed using the method described in Appendix C. to produce subpixel shock and edge locations. We denote this mapping from density to features $F(\rho)$. Ideally, generated densities show backwards consistency with the features, i.e., $F(G_{\theta_g}(z, F(\rho))) = F(\rho)$. To enforce this constraint, we introduce an $L2$ penalty on differences in $F(\rho)$ and $F(G_{\theta_g}(z, F(\rho)))$. However, the edge and shock extraction procedure, F , is not differentiable and cannot facilitate back propagation during GAN training. Instead, a convolutional model, denoted \tilde{F} is pretrained on $(\rho, F(\rho))$ pairs to approximate this map. \tilde{F} is made up of five 1D convolutional layers of widths from 16 up to 256 with kernels of length three and Leaky Rectified Logic Unit (LReLU) with threshold parameter 0.3. The final layer is a Dense operation with a TanH activation and reshaping. The weights of \tilde{F} are held fixed while training of cGAN-DF. The corresponding loss term is given by

$$L_{\tilde{F}}(\theta_g) = \mathbb{E}_{\rho \sim P, z \sim \mathcal{N}} [\|F(\rho) - \tilde{F}(G_{\theta_g}(z, F(\rho)))\|_2] \quad (9)$$

Let $M(\tilde{\rho}) \in \mathbb{R}^4$ be the spherical integral of the radial density profile for the four distinct timeframes. We introduce an $L2$ penalty between the mass of the generated result and known mass of the steel capsule, M_ρ .

$$L_M(\theta_g) = \mathbb{E}_{\rho \sim P, z \sim \mathcal{N}} [\|M_\rho - M(G_{\theta_g}(z, F(\rho)))\|_2] \quad (10)$$

The resulting composite loss function for the generator update is

$$\mathcal{L}_{cGAN}(\theta_g) = \lambda_D L_G(\theta_g) + \lambda_{L1} L_{L1}(\theta_g) + \lambda_{\tilde{F}} L_{\tilde{F}}(\theta_g) + \lambda_M L_M(\theta_g) \quad (11)$$

Appendix C. Subpixel shock extraction

When recovering the shock position, it is convenient to use radial density $\mu(r) := 4\pi\rho r^2$ (with the units of g cm^{-1}). This quantity has an advantage of producing the mass directly when integrated over radius. At a first step, we locate the shock and the grid cell with the maximum magnitude of the radial density gradient, $|\nabla_r \mu|$. For the outgoing shock, this gradient is negative at the location of the shock and in a few neighboring cells. In the remaining parts of the domain, the density does not drop very fast and because of the $O(r^2)$ term in the definition of μ , the gradient of μ is positive almost everywhere. A small connected segment $[r_{j-}, r_{j+}]$ around the shock consisting of cells with the negative density gradient separates the domain into upstream and downstream parts. We fit the radial density in the downstream and upstream N -cell neighborhood of the shock with linear regression, obtaining radial density approximations $\mu_-(r)$ and $\mu_+(r)$ to the left and to the right of the shock, respectively. An accurate shock location can then be computed by assuming a discontinuous transition from $\mu_-(r)$ to $\mu_+(r)$ at some $r = r_s$ and solving the mass conservation equation for r_s :

$$\int_{r_{j-}}^{r_s} \mu_-(r) dr + \int_{r_s}^{r_{j+}} \mu_+(r) dr = \sum_{j=j-}^{j+} \mu_j \Delta r. \quad (12)$$

This again becomes a simple quadratic equation on r_s . Clearly, the root must be picked such that $r_s \in [r_{j-}, r_{j+}]$. Because of the imposed mass conservation, our technique is robust with respect to low-amplitude oscillating numerical artifacts near the shock.

Figure 15 shows evolution of the extracted subpixel shock positions for different resolutions of hydrodynamic model. This evolution demonstrates stable and convergent behavior of our subpixel feature extraction algorithm.




Fig. 15. Extracted position of the shock for the test problem with nominal parameter values, for different successively increasing resolutions of the hydrodynamic numerical model. The orange bar on the plot indicates the size of numerical cell for the coarsest resolution, demonstrating subpixel accuracy of shock and edge extraction.

Funding. National Science Foundation (CCF-1763896); National Nuclear Security Administration (89233218CNA000001).

Acknowledgments. The authors thank Luke Pfister and Jonah M. Miller for their helpful feedback. This work used resources provided by the LANL Institutional Computing Program. LANL is operated by Triad National Security, LLC, for the National Nuclear Security Administration of the U.S. DOE (Contract No. 89233218CNA000001). This work was also partially supported by NSF grant number CCF-1763896. This work is authorized for unlimited release under LA-UR-21-31230.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

1. J. Radon, "über die Bestimmung von Funktionen durch ihre Integralwerte längs gewisser Mannigfaltigkeiten," *Berichte über die Verhandlungen der Sächsischen Akademie der Wissenschaften* **69**, 262–277 (1917). 00000.
2. A. M. Cormack, "Representation of a function by its line integrals, with some radiological applications," *J. Appl. Phys.* **34**(9), 2722–2727 (1963).
3. L. A. Feldkamp, L. C. Davis, and J. W. Kress, "Practical cone beam algorithm," *J. Opt. Soc. Am. A* **1**(6), 612 (1984).
4. R. N. Bracewell and R. N. Bracewell, *The Fourier transform and its applications*, vol. 31999 (McGraw-Hill New York, 1986).
5. I. Elbakri and J. Fessler, "Statistical image reconstruction for polyenergetic X-ray computed tomography," *IEEE Trans. Med. Imaging* **21**(2), 89–99 (2002).
6. S. Ravishanker, J. C. Ye, and J. A. Fessler, "Image reconstruction: From sparsity to data-adaptive methods and machine learning," *Proc. IEEE* **108**(1), 86–109 (2020).
7. C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Atom-Photon Interactions: Basic Processes and Applications* (VCH PUBLN, 1998).
8. J. P. Stonestrom and A. Macovski, "Scatter considerations in fan beam computerized tomographic systems," *IEEE Trans. Nucl. Sci.* **23**(5), 1453–1458 (1976).
9. M. Sun and J. M. Star-Lack, "Improved scatter correction using adaptive scatter kernel superposition," *Phys. Med. Biol.* **55**(22), 6695–6720 (2010).
10. N. Bhatia, D. Tisseur, and J. M. Létang, "Convolution-based scatter correction using kernels combining measurements and monte carlo simulations," *J. X-Ray Sci. Technol.* **25**(4), 613–628 (2017).
11. D. Tisseur, N. Bhatia, N. Estre, L. Berge, D. Eck, and E. Payan, "Evaluation of a scattering correction method for high energy tomography," *EPJ Web Conf.* **170**, 06006 (2018).
12. J. Maier, S. Sawall, M. Knaup, and M. Kachelrieß, "Deep scatter estimation (DSE): Accurate real-time scatter estimation for X-ray CT using a deep convolutional neural network," *J. Nondestruct. Eval.* **37**(3), 57 (2018).

13. M. T. McCann, M. L. Klasky, J. L. Schei, and S. Ravishankar, "Local models for scatter estimation and descattering in polyenergetic X-ray tomography," *Opt. Express* **29**(18), 29423–29438 (2021).
14. E.-P. Rührschopf and K. Klingensbeck, "A general framework and review of scatter correction methods in X-ray cone-beam computerized tomography. part 1: Scatter compensation approaches," *Med. Phys.* **38**(7), 4296–4311 (2011).
15. E.-P. Rührschopf and K. Klingensbeck, "A general framework and review of scatter correction methods in cone beam CT. part 2: Scatter estimation approaches," *Med. Phys.* **38**(9), 5186–5199 (2011).
16. C. J. Werner, J. S. Bull, C. J. Solomon, F. Brown, G. McKinney, M. Rising, D. Dixon, R. Martz, H. Hughes, L. Cox, A. J. Zukaitis, J. Armstrong, R. A. Forster, and L. Casswell, "MCNP version 6.2 release notes," Tech. Rep. LA-UR-18-20808, Los Alamos National Laboratory (2018).
17. E. S. Hertel Jr. and G. I. Kerley, "Cth reference manual: The equation of state package," report SAND98-0947, Sandia National Laboratories (1998).
18. J. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Anal. Mach. Intell.* **PAMI-8**(6), 679–698 (1986).
19. D. Robinson and P. Milanfar, "Fast local and global projection-based methods for affine motion estimation," *J. Math. Imaging Vis.* **18**(1), 35–54 (2003).
20. J. Cao, D. Davis, S. Vander Wiel, and B. Yu, "Time-varying network tomography: Router link data," *J. Am. Stat. Assoc.* **95**(452), 1063–1075 (2000).
21. C. Jailin and S. Roux, "Dynamic tomographic reconstruction of deforming volumes," *Materials* **11**(8), 1395 (2018).
22. S. Sikdar, Q. Wei, and N. Cortes, "Dynamic ultrasound imaging applications to quantify musculoskeletal function," *Exercise Sport Sci. Rev.* **42**(3), 126–135 (2014).
23. J. Nemirovsky, A. Lifshitz, and I. Be'ery, "Tomographic reconstruction of incompressible flow," *Rev. Sci. Instrum.* **82**(5), 055115 (2011).
24. M. Krimerman, "Reconstruction of bubble trajectories and velocity estimation," Ph.D. thesis, University of British Columbia (2013).
25. J. Gregson, M. Krimerman, M. B. Hullin, and W. Heidrich, "Stochastic tomography and its applications in 3d imaging of mixing fluids," *ACM Trans. Graph.* **31**(4), 1–10 (2012).
26. I. Ihrke and M. Magnor, "Image-based tomographic reconstruction of flames," in *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*, (2004), pp. 365–373.
27. D. Steinberg, S. Cochran, and M. Guinan, "A constitutive model for metals applicable at high-strain rate," *J. Appl. Phys.* **51**(3), 1498–1504 (1980).
28. E. S. Hertel Jr. and G. I. Kerley, "Cth eos package: Introductory tutorial," report SAND98-0945, Sandia National Laboratories (1998).
29. J. Johnson, "The sesame database," Tech. rep., Los Alamos National Lab.(LANL), Los Alamos, NM (United States) (1994).
30. D. L. Preston, D. L. Tonks, and D. C. Wallace, "Model of plastic deformation for extreme loading conditions," *J. Appl. Phys.* **93**(1), 211–220 (2003).
31. I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," *Advances in neural information processing systems* **27** (2014).
32. D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," arXiv e-prints arXiv:1312.6114 (2013).
33. D. Rezende and S. Mohamed, "Variational inference with normalizing flows," in *International conference on machine learning*, (PMLR, 2015), pp. 1530–1538.
34. R. R. Coifman and S. Lafon, "Diffusion maps," *Appl. Computational Harmonic Anal.* **21**(1), 5–30 (2006).
35. P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, (2017), pp. 1125–1134.
36. B. Kashiwa, "The mggb equation-of-state for multifield applications: a numerical recipe for analytic expression of sesame eos data," Tech. rep., Los Alamos National Lab.(LANL), Los Alamos, NM (United States) (2010).
37. K. Zhu and E. A. Muller, "Generating a machine-learned equation of state for fluid properties," *J. Phys. Chem. B* **124**(39), 8628–8639 (2020).
38. F. Morawski and M. Bejger, "Neural network reconstruction of the dense matter equation of state derived from the parameters of neutron stars," *Astron. & Astrophys.* **642**, A78 (2020).
39. I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin, and A. Courville, "Improved training of wasserstein gans," arXiv preprint arXiv:1704.00028 (2017).