

Article

Physics-Constrained Neural Identification of Fragmentation Kernels: From Synthetic Recovery to Effective Daughter-Volume-Fraction Estimation in Droplet Breakup

Joseph El Maalouf ¹, Alain Ajami ² and Candy Abboud ^{3,*} 

¹ School of Science—Dubai Branch, Saint Joseph University of Beirut, Beirut 345008, Lebanon; joseph.maalouf2@usj.edu.lb

² Faculty of Engineering, Saint Joseph University of Beirut, Beirut 1107 2050, Lebanon; alain.ajami@usj.edu.lb

³ College of Engineering and Technology, American University of the Middle East, Egaila 15453, Kuwait

* Correspondence: candy.abboud@aum.edu.kw

Abstract

Fragmentation processes arise in many physical systems, including droplet breakup, aerosols, sprays, comminution, and granular media. A central difficulty in fragmentation modeling is the identification of the breakup law from observed particle-size distributions. In this work, we propose a physics-constrained neural framework for the inverse identification of fragmentation laws. Starting from a size-structured binary fragmentation equation, the kernel is decomposed into a breakup rate $B(s)$, depending on the parent size s , and a daughter-size distribution $\kappa(z)$, depending on the normalized daughter-size fraction z . The unknown functions B and κ are represented by neural networks designed to satisfy physical constraints, including non-negativity of the breakup rate, non-negativity and normalization of the daughter distribution, and first-moment conservation in the symmetric binary setting. The direct fragmentation equation is solved using a first-moment-conserving quadrature discretization, and the inverse problem is formulated as a regularized optimization problem constrained by the fragmentation dynamics. Synthetic experiments, generated independently of the inverse solver, show that the proposed method can recover B and κ from simulated particle-size distributions, with accurate recovery in the unimodal case and reasonable performance in the more challenging bimodal case. Robustness tests indicate stability with respect to moderate observational noise and highlight the importance of multiple observation times. For experimental droplet-breakup measurements without time-resolved particle-size distributions, the same constrained neural daughter-density representation is used to estimate effective marginal daughter-volume-fraction distributions from normalized daughter-to-parent volume fractions. The resulting distributions distinguish rim and node fragments and different breakup modes, while stratified cluster-bootstrap confidence bands quantify their sampling uncertainty. The experimental analysis is therefore interpreted as constrained density estimation from final fragment measurements rather than as full validation of the PDE-constrained inverse recovery.

Keywords: fragmentation equation; inverse problem; physics-constrained neural networks; breakup kernel; daughter-size distribution; droplet breakup



Academic Editor: Sandra Ferreira

Received: 4 June 2026

Revised: 13 July 2026

Accepted: 19 July 2026

Published: 20 July 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and

conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Fragmentation processes arise in a wide range of physical and engineering systems, including droplet breakup, sprays, aerosols, comminution, granular flows, polymer degra-

ation, and cell division. In such systems, particles, droplets, or aggregates break into smaller fragments, leading to the evolution of a size distribution over time. Mathematical models based on fragmentation and coagulation–fragmentation equations have been extensively studied in deterministic, probabilistic, and semigroup frameworks; see, for instance [1,2]. Physical fragmentation mechanisms and their role in fluid and multiphase systems have also been widely investigated [3].

A central quantity in these models is the fragmentation kernel. This kernel determines both the rate at which a parent particle breaks and the distribution of the resulting daughter fragments. In the forward problem, the fragmentation law is assumed to be known, and the goal is to predict the evolution of the particle-size distribution. However, in many practical situations, the kernel is not directly measurable. Instead, one observes particle-size distributions at different times or under different experimental conditions. This naturally leads to an inverse problem: recovering information about the fragmentation law from observed size-distribution data. Related inverse questions include the estimation of division rates and fragmentation kernels from size-distribution measurements [4], as well as the recovery of fragmentation kernels from short-time behavior [5].

From an application viewpoint, the observed particle-size distributions are not treated as prior knowledge of the fragmentation kernel, but as the measurable data used to identify it. This setting is relevant when direct measurements of breakup frequencies or daughter-size distributions are unavailable, while size-distribution data can be obtained from experiments, simulations, or imaging-based measurements. The recovered functions should therefore be interpreted as effective fragmentation laws associated with the available observation window and physical conditions.

The inverse identification of fragmentation kernels is challenging for several reasons. First, the kernel may depend on several variables, making the inverse problem high-dimensional. Second, different kernels may produce similar particle-size distributions, which creates identifiability and stability issues. Third, classical approaches often assume a prescribed parametric form of the kernel, reducing the inverse problem to the estimation of a small number of constants. While this can be effective in controlled settings, it may be too restrictive for complex physical fragmentation mechanisms.

In this work, we propose a physics-constrained neural framework for the inverse identification of fragmentation laws. Instead of attempting to recover an arbitrary two-variable kernel directly, we use a physically interpretable decomposition into a breakup rate $B(s)$, depending on the parent size s , and a daughter-size distribution $\kappa(z)$, depending on the normalized daughter-size fraction z . Here, the size variable is interpreted as an additive quantity, such as mass or volume. More precisely, the fragmentation kernel is written in the form

$$F(x, y) = B(x + y) \frac{1}{x + y} \kappa\left(\frac{x}{x + y}\right), \quad (1)$$

where $x + y$ is the size of the parent particle and $x/(x + y)$ is the daughter-size fraction. This decomposition separates the frequency of breakup events from the distribution of daughter fragments.

The unknown functions B and κ are represented by neural networks, but the parametrization is designed to respect the physical structure of the fragmentation equation. The breakup rate is constrained to be non-negative, while the daughter distribution is constrained to be non-negative and normalized. In this sense, the neural networks do not replace the physical model. Rather, they provide flexible representations of the unknown fragmentation law inside the governing equation. This philosophy is related to physics-informed and physics-constrained learning approaches, where neural representations are coupled with governing equations or physical laws [6]. Neural-network-based methods

have also been widely developed for inverse problems in physics and mechanics. Raissi and Karniadakis introduced hidden-physics models for learning nonlinear partial differential equations from data, illustrating how neural representations can be used to infer unknown dynamical laws [7]. More generally, physics-informed machine learning has provided a broad framework for combining data with governing equations in forward and inverse problems [8]. Data-driven approaches to inverse problems have also been reviewed from a broader mathematical and computational perspective by Arridge et al. [9]. In fluid mechanics, physics-informed neural networks have been used for flow reconstruction, parameter identification, and data assimilation [10], while in solid mechanics they have been applied to inversion and surrogate modeling of material and constitutive behavior [11]. Related developments include physics-informed deep learning for parameter and constitutive-law identification in subsurface flow [12] and neural-operator approaches for learning mappings between function spaces [13]. The present work follows this general direction, but focuses specifically on the inverse identification of fragmentation kernels under positivity, normalization, and mass-conservation constraints.

The experimental motivation of this work comes from aerodynamic droplet breakup, where a parent droplet fragments into child droplets through mechanisms such as bag, rim, and node breakup. Droplet-size distributions resulting from aerodynamic breakup have been studied experimentally and theoretically, notably in the work of Jackiw and Ashgriz [14]. Their supplementary measurements provide parent-drop properties, breakup regimes, and child-fragment sizes. Since the fragmentation equation conserves an additive size variable, the measured child-to-parent diameter ratios are converted into daughter-volume fractions before estimating effective daughter-volume-fraction distributions from the experimental data.

The contributions of this paper are as follows:

- We formulate a physics-constrained neural inverse framework for identifying fragmentation laws through the functions B and κ .
- We introduce neural parametrizations that enforce non-negativity of the breakup rate and non-negativity and normalization of the daughter-size distribution.
- We establish structural properties of the proposed neural framework, showing that the parametrization preserves admissibility and the mass-conservation structure of the fragmentation model, and that the associated regularized inverse problem is well defined under suitable assumptions.
- We validate the method on synthetic data generated independently of the inverse solver, including both unimodal and bimodal daughter distributions.
- We investigate the robustness of the method with respect to observational noise and the number of available observation times.
- We use the same positivity- and normalization-constrained neural daughter-density representation to estimate effective marginal daughter-volume-fraction distributions from experimental droplet-breakup measurements, and we quantify their sampling uncertainty using a stratified cluster-bootstrap analysis.

The experimental application is interpreted carefully. Since the available droplet-breakup dataset contains measured parent and child fragment sizes but does not provide full time-resolved particle-size distributions, the real-data part focuses on estimating an effective marginal daughter-volume-fraction distribution rather than solving the full PDE-constrained inverse problem or identifying a breakup-rate function. The synthetic experiments, where the true kernels are known and the observations are generated independently of the inverse solver, provide the validation of the full inverse recovery of both B and κ .

The remainder of the paper is organized as follows. Section 2 introduces the fragmentation equation and the kernel decomposition. Section 3 presents the physics-constrained

neural parametrization and the inverse problem formulation. Section 4 describes the numerical discretization of the direct problem. Section 5 gives the structural properties of the proposed neural inverse framework. Section 6 presents synthetic validation results. Section 7 studies robustness with respect to noise and observation richness. Section 8 estimates effective daughter-volume-fraction distributions from experimental droplet-breakup data using the constrained neural density representation. Finally, Section 9 discusses the results and Section 10 concludes the paper.

2. Fragmentation Model

Let $c(x, t)$ denote the density of particles of size $x > 0$ at time $t \geq 0$. Throughout the model, x denotes an additive size variable, such as mass or volume. We consider a size-structured binary fragmentation equation of the form

$$\frac{\partial c}{\partial t}(x, t) = -B(x)c(x, t) + 2 \int_x^\infty B(y)\kappa\left(\frac{x}{y}\right)\frac{1}{y}c(y, t) dy, \quad x > 0, t > 0. \quad (2)$$

The first term on the right-hand side represents the loss of particles of size x due to breakup. The integral term represents the gain of particles of size x produced by the fragmentation of larger particles $y \geq x$.

The function $B(y) \geq 0$ is the breakup rate of a parent particle of size y . The function $\kappa(z)$, defined for $z \in (0, 1)$, describes the probability density of the daughter-size fraction

$$z = \frac{x}{y}.$$

The factor $1/y$ appears from the change of variables between the daughter size x and the daughter fraction $z = x/y$. The factor 2 accounts for binary fragmentation: each breakup event produces two daughter fragments. With this convention, κ describes the size-fraction distribution of one daughter, while the factor 2 accounts for the contribution of both daughters to the gain term. The binary formulation is most directly suited to breakup events dominated by two principal daughters. It may also be used as an effective approximation of progressive fragmentation when a more complex process is represented as a sequence of binary splitting events; however, it is not a literal description of a single event producing several daughter fragments simultaneously.

Equivalently, if x and y denote the two daughter sizes, then the parent size is $x + y$. The two-variable fragmentation kernel can be written as

$$F(x, y) = B(x + y)\frac{1}{x + y}\kappa\left(\frac{x}{x + y}\right). \quad (3)$$

This representation separates the breakup frequency of a parent particle from the distribution of the daughter-size fraction. In this work, the inverse problem is formulated in terms of the two physically interpretable functions B and κ , rather than by attempting to identify a completely free two-variable function F .

We impose the admissibility conditions

$$B(x) \geq 0, \quad \kappa(z) \geq 0, \quad \int_0^1 \kappa(z) dz = 1. \quad (4)$$

In the symmetric binary case, we additionally assume

$$\kappa(z) = \kappa(1 - z), \quad z \in (0, 1). \quad (5)$$

Under the normalization and symmetry assumptions, the first moment of κ satisfies

$$\int_0^1 z\kappa(z) dz = \frac{1}{2}. \quad (6)$$

This condition expresses the fact that, on average, the two daughters preserve the size of the parent particle.

A key physical property of Equation (2) is conservation of total mass. Defining

$$M_1(t) = \int_0^\infty xc(x, t) dx, \quad (7)$$

one expects $M_1(t)$ to remain constant along the fragmentation dynamics, provided the daughter distribution satisfies (6). In contrast, the total number of particles,

$$M_0(t) = \int_0^\infty c(x, t) dx, \quad (8)$$

typically increases with time, since each breakup event replaces one parent particle by two daughters.

The inverse problem considered in this paper is the following. Given observations of the particle-size distribution at discrete times,

$$c_{\text{obs}}(x, t_1), \dots, c_{\text{obs}}(x, t_m),$$

we aim to identify an admissible fragmentation law represented by B and κ . In the synthetic setting, both B and κ are recovered and compared with known ground-truth functions. In the experimental droplet-breakup setting, full time-resolved size distributions are not available. Therefore, the real-data application focuses on the estimation of an effective marginal daughter-volume-fraction distribution from normalized daughter-to-parent volume fractions.

3. Physics-Constrained Neural Inverse Framework

The goal of the proposed method is to identify the unknown functions B and κ without prescribing a restrictive parametric form. Instead, we represent them by neural networks while enforcing the physical constraints in (4). The neural networks are therefore used as flexible parametrizations of admissible fragmentation laws inside the forward model, rather than as direct black-box predictors.

3.1. Neural Representation of the Breakup Rate

The breakup rate must be non-negative. To guarantee this property, we define

$$B_\theta(s) = \text{Softplus}(N_\theta(s)), \quad (9)$$

where N_θ is a neural network with parameter vector θ , and

$$\text{Softplus}(r) = \log(1 + e^r).$$

Therefore,

$$B_\theta(s) > 0 \quad \text{for all } s.$$

This construction avoids nonphysical negative breakup rates during the optimization process.

3.2. Neural Representation of the Daughter Distribution

The daughter-size distribution must be non-negative and normalized on $(0, 1)$. We therefore define

$$\kappa_{\eta}(z) = \frac{\exp(M_{\eta}(z))}{\int_0^1 \exp(M_{\eta}(u)) du}, \quad 0 < z < 1, \quad (10)$$

where M_{η} is a neural network with parameter vector η . By construction,

$$\kappa_{\eta}(z) > 0, \quad \int_0^1 \kappa_{\eta}(z) dz = 1.$$

Thus, the neural network output is transformed into an admissible probability density.

When symmetric binary fragmentation is assumed, we use the symmetrized density

$$\kappa_{\eta}^{\text{sym}}(z) = \frac{1}{2} [\kappa_{\eta}(z) + \kappa_{\eta}(1 - z)]. \quad (11)$$

This guarantees

$$\kappa_{\eta}^{\text{sym}}(z) = \kappa_{\eta}^{\text{sym}}(1 - z).$$

Moreover, since $\kappa_{\eta}^{\text{sym}}$ is normalized and symmetric, it satisfies

$$\int_0^1 z \kappa_{\eta}^{\text{sym}}(z) dz = \frac{1}{2}. \quad (12)$$

Since the fragmentation operator contains the factor 2, condition (12) is equivalently expressed as

$$2 \int_0^1 z \kappa_{\eta}^{\text{sym}}(z) dz = 1,$$

which is the first-moment conservation condition for binary fragmentation. Thus, normalization and symmetry enforce this condition structurally, and no additional moment penalty is required in the symmetric setting considered in the synthetic experiments. For a nonsymmetric binary daughter distribution, the same physical condition could instead be incorporated through a penalty such as

$$R_{\text{mom}}(\eta) = \left(2 \int_0^1 z \kappa_{\eta}(z) dz - 1 \right)^2.$$

For simplicity of notation, we continue to write κ_{η} in the sequel. When symmetric binary fragmentation is imposed in the synthetic experiments, κ_{η} denotes the symmetrized density in (11). In the experimental droplet-breakup application, where the observed process is multi-fragment and interpreted through an effective marginal daughter-volume-fraction distribution, the general normalized density (10) is used.

3.3. Forward Map

For a given pair $(B_{\theta}, \kappa_{\eta})$, let $c_{\theta, \eta}(x, t)$ denote the solution of the fragmentation equation

$$\frac{\partial c_{\theta, \eta}}{\partial t}(x, t) = -B_{\theta}(x)c_{\theta, \eta}(x, t) + 2 \int_x^{\infty} B_{\theta}(y)\kappa_{\eta}\left(\frac{x}{y}\right)\frac{1}{y}c_{\theta, \eta}(y, t) dy. \quad (13)$$

The map

$$(\theta, \eta) \longmapsto c_{\theta, \eta}(\cdot, t)$$

is the neural forward map. Evaluating this map requires solving the direct fragmentation equation for each candidate pair (B_θ, κ_η) . The inverse problem consists of choosing θ and η so that $c_{\theta, \eta}$ matches the observed distributions.

3.4. Inverse Objective

Assume that observations are available at times

$$0 < t_1 < \dots < t_m.$$

We define the data misfit

$$\mathcal{D}(\theta, \eta) = \sum_{n=1}^m \|c_{\theta, \eta}(\cdot, t_n) - c_{\text{obs}}(\cdot, t_n)\|_{L^2(\Omega_x)}^2, \quad (14)$$

where Ω_x is the computational size interval.

The use of observations at several times is particularly important for identifying the breakup rate $B(s)$. Indeed, $B(s)$ controls the rate at which particles of size s disappear and generate daughter fragments, and this rate information is encoded in the temporal evolution of the particle-size distribution. Therefore, full joint recovery of B and κ requires time-resolved or multi-stage size-distribution observations. When such observations are not available, as in the experimental droplet-breakup application considered in Section 8, the analysis is restricted to estimating an effective marginal daughter-volume-fraction distribution rather than a complete breakup-rate function.

Because inverse fragmentation problems are generally ill-conditioned, we add regularization terms. In this work, we use smoothness penalties of the form

$$R_B(\theta) = \int_{\Omega_x} \left| \frac{d}{ds} B_\theta(s) \right|^2 ds, \quad (15)$$

and

$$R_\kappa(\eta) = \int_0^1 \left| \frac{d}{dz} \kappa_\eta(z) \right|^2 dz. \quad (16)$$

The regularized inverse problem is then

$$\min_{\theta, \eta} J(\theta, \eta), \quad (17)$$

where

$$J(\theta, \eta) = \mathcal{D}(\theta, \eta) + \lambda_B R_B(\theta) + \lambda_\kappa R_\kappa(\eta). \quad (18)$$

The parameters $\lambda_B \geq 0$ and $\lambda_\kappa \geq 0$ control the strength of the regularization. Thus, positivity, normalization, symmetry, and the associated binary first-moment condition are enforced structurally, whereas smoothness is imposed softly through the derivative regularization terms.

The important point is that the optimization is constrained by the fragmentation dynamics. The neural networks do not directly map observations to kernels. Instead, they define admissible kernels, these kernels are inserted into the fragmentation equation, and the resulting solution is compared with the observed size distributions. If additional physical variables are available, the same construction can be extended by adding them as inputs to the neural parametrizations of B and κ , while preserving the same positivity and normalization constraints.

4. Numerical Discretization

In this section, we describe the deterministic numerical solver used for the direct fragmentation problem. This solver is embedded in the optimization procedure for the inverse problem.

Let

$$\Omega_x = [x_{\min}, x_{\max}]$$

be the computational size interval, and consider a uniform grid

$$x_i = x_{\min} + (i - 1)\Delta x, \quad i = 1, \dots, N.$$

We denote by $c_i(t)$ an approximation of $c(x_i, t)$. The gain term in (2) is approximated by a quadrature rule over the grid points $x_j \geq x_i$. For each parent class x_j , we first define the raw daughter weights

$$q_{ij} = \frac{1}{x_j} \kappa\left(\frac{x_i}{x_j}\right), \quad i \leq j. \quad (19)$$

The daughter weights used in the discrete gain term are then chosen as

$$p_{ij} = \alpha_j q_{ij}, \quad \alpha_j = \frac{x_j}{2 \sum_{\ell \leq j} w_\ell x_\ell q_{\ell j}}, \quad i \leq j, \quad (20)$$

with $p_{ij} = 0$ for $i > j$. Hence, for every parent class x_j , the weights satisfy the discrete first-moment balance

$$2 \sum_{i \leq j} w_i x_i p_{ij} = x_j. \quad (21)$$

This condition states that the total discrete volume carried by the two daughter fragments equals the volume of the parent class. It is in the same conservation-oriented spirit as moment-preserving sectional methods for population balance equations, such as the fixed-pivot technique of Kumar and Ramkrishna [15], while retaining the quadrature-based structure used in the present inverse solver.

With these weights, the gain term is approximated as

$$\int_{x_i}^{x_{\max}} B(y) \kappa\left(\frac{x_i}{y}\right) \frac{1}{y} c(y, t) dy \approx \sum_{j \geq i} w_j B(x_j) p_{ij} c_j(t). \quad (22)$$

This leads to the finite-dimensional system

$$\frac{dc_i}{dt} = -B(x_i) c_i(t) + 2 \sum_{j \geq i} w_j B(x_j) p_{ij} c_j(t), \quad i = 1, \dots, N. \quad (23)$$

The factor 2 is kept in the discrete equation to account for binary fragmentation.

The discrete first-moment conservation follows directly from (21). Indeed, defining

$$M_1^h(t) = \sum_{i=1}^N w_i x_i c_i(t),$$

multiplication of (23) by $w_i x_i$, followed by summation over i , gives

$$\frac{dM_1^h}{dt} = - \sum_{j=1}^N w_j x_j B(x_j) c_j(t) + 2 \sum_{j=1}^N w_j B(x_j) c_j(t) \sum_{i \leq j} w_i x_i p_{ij}.$$

Using (21), the two terms cancel, and therefore

$$\frac{dM_1^h}{dt} = 0.$$

Thus the discrete solver preserves the first moment, which corresponds to total mass or volume when x is an additive size variable.

For fixed B and κ , (23) is an ordinary differential equation system of the form

$$\frac{d\mathbf{c}}{dt} = \mathcal{A}_{B,\kappa}\mathbf{c}, \quad (24)$$

where

$$\mathbf{c}(t) = (c_1(t), \dots, c_N(t))^T.$$

The operator $\mathcal{A}_{B,\kappa}$ contains a diagonal loss term and an upper triangular gain contribution. In the numerical experiments, the resulting ODE system is integrated using an explicit fourth-order Runge–Kutta method.

The same discretization is used in the inverse problem. At each optimization step, the current neural approximations B_θ and κ_η are evaluated on the grid, the discrete operator $\mathcal{A}_{B_\theta, \kappa_\eta}$ is constructed, and the corresponding direct problem is solved. The computed solution is then compared with the observed size distributions in the objective function (18).

4.1. Implementation Details

The inverse computations reported in Sections 6 and 7 use the deterministic forward solver described above. The neural parametrizations B_θ and κ_η are implemented using one-hidden-layer feedforward networks with hyperbolic tangent activation functions. The breakup rate is transformed through the Softplus function, while the daughter distribution is transformed through the normalized exponential parametrization described in Section 3. The neural parameters are optimized by minimizing the regularized objective function (18). The exact numerical settings used in the synthetic and robustness tests are reported in Section 6.

4.2. Numerical Conservation Diagnostics

Before using the solver inside the inverse problem, we verify that it reproduces the expected qualitative properties of fragmentation dynamics. In particular, for symmetric daughter distributions satisfying (6), the total mass

$$M_1(t) = \int_{\Omega_x} xc(x, t) dx$$

should be conserved, while the total number of particles

$$M_0(t) = \int_{\Omega_x} c(x, t) dx$$

should increase.

In the discrete setting, these moments are approximated by

$$M_0^h(t) = \sum_{i=1}^N w_i c_i(t), \quad M_1^h(t) = \sum_{i=1}^N w_i x_i c_i(t). \quad (25)$$

The numerical tests reported below confirm that $M_1^h(t)$ remains conserved, while $M_0^h(t)$ increases with time. Figure 1 illustrates a representative synthetic direct simulation. The initial distribution is concentrated near larger sizes, and the solution progressively shifts

toward smaller sizes as fragmentation produces daughter particles. This behavior is consistent with the expected qualitative dynamics of the fragmentation equation.




Figure 1. Synthetic direct fragmentation simulation. The initial distribution is concentrated near large sizes, and the evolution shows the progressive production of smaller fragments.

Figure 2 reports the corresponding moment diagnostics. The normalized total particle number increases over time, whereas the normalized total mass remains equal to one throughout the simulation. This provides a basic validation of the direct solver before it is used in the inverse identification procedure.




Figure 2. Forward solver diagnostics. The total particle number increases, while the total mass remains conserved.

5. Structural Properties of the Neural Inverse Framework

This section establishes the structural properties of the proposed physics-constrained neural inverse framework. Some of the arguments rely on classical properties of fragmentation equations, such as mass conservation and stability of the forward map. They are recalled here and specialized to the neural parametrization in order to show that the learned kernels remain physically admissible and that the neural forward model used in the inverse problem is well defined.

The main point is not to claim new general theory for fragmentation equations. Rather, the goal is to show that the proposed neural construction preserves the physical structure of the model and provides a stable admissible class for inverse identification.

5.1. Admissibility of the Neural Kernels

Here and below, p_B and p_κ denote the dimensions of the trainable parameter vectors used in the neural parametrizations of B and κ , respectively.

Proposition 1 (Admissibility of the neural parametrization). *Let $\Omega_x \subset (0, \infty)$ be the size domain, let $\Theta \subset \mathbb{R}^{p_B}$ and $\mathcal{H} \subset \mathbb{R}^{p_\kappa}$ be parameter sets, and let*

$$N_\theta : \Omega_x \rightarrow \mathbb{R}, \quad M_\eta : [0, 1] \rightarrow \mathbb{R}, \quad \theta \in \Theta, \quad \eta \in \mathcal{H},$$

be continuous neural networks. Define B_θ and κ_η by

$$B_\theta(s) = \text{Softplus}(N_\theta(s)), \quad s \in \Omega_x, \quad (26)$$

and

$$\kappa_\eta(z) = \frac{\exp(M_\eta(z))}{\int_0^1 \exp(M_\eta(u)) du}, \quad 0 < z < 1. \quad (27)$$

Then

$$B_\theta(s) > 0, \quad \kappa_\eta(z) > 0, \quad \int_0^1 \kappa_\eta(z) dz = 1. \quad (28)$$

If the symmetrized form

$$\kappa_\eta^{\text{sym}}(z) = \frac{1}{2} [\kappa_\eta(z) + \kappa_\eta(1 - z)], \quad 0 < z < 1, \quad (29)$$

is used, then

$$\kappa_\eta^{\text{sym}}(z) = \kappa_\eta^{\text{sym}}(1 - z). \quad (30)$$

Moreover,

$$\int_0^1 z \kappa_\eta^{\text{sym}}(z) dz = \frac{1}{2}. \quad (31)$$

Proof. The Softplus function is strictly positive for every real input. Therefore

$$B_\theta(s) > 0$$

for all $s \in \Omega_x$. Similarly, since the exponential function is strictly positive, we have

$$\exp(M_\eta(z)) > 0$$

for all $z \in (0, 1)$. The denominator in (27) is also strictly positive, and hence

$$\kappa_\eta(z) > 0.$$

Moreover,

$$\int_0^1 \kappa_\eta(z) dz = \frac{\int_0^1 \exp(M_\eta(z)) dz}{\int_0^1 \exp(M_\eta(u)) du} = 1.$$

If κ_η^{sym} is defined by (29), then

$$\kappa_\eta^{\text{sym}}(1 - z) = \frac{1}{2} [\kappa_\eta(1 - z) + \kappa_\eta(z)] = \kappa_\eta^{\text{sym}}(z).$$

Finally, since $\kappa_{\eta}^{\text{sym}}$ is normalized and symmetric, we have, by the change of variables $u = 1 - z$,

$$\int_0^1 z \kappa_{\eta}^{\text{sym}}(z) dz = \int_0^1 (1 - z) \kappa_{\eta}^{\text{sym}}(z) dz.$$

Adding the two equal quantities gives

$$2 \int_0^1 z \kappa_{\eta}^{\text{sym}}(z) dz = \int_0^1 \kappa_{\eta}^{\text{sym}}(z) dz = 1.$$

Therefore,

$$\int_0^1 z \kappa_{\eta}^{\text{sym}}(z) dz = \frac{1}{2}.$$

This proves the result. $\square$

5.2. Mass Conservation of the Neural Forward Model

The following proposition is a direct adaptation of the classical mass-balance property of fragmentation equations. It is stated for the neural forward model to emphasize that the proposed parametrization preserves the physical conservation structure.

Proposition 2 (Mass conservation of the neural forward model). *Let $B_{\theta} : (0, \infty) \rightarrow [0, \infty)$ and $\kappa_{\eta} : (0, 1) \rightarrow [0, \infty)$ be neural kernels, with $\theta \in \Theta \subset \mathbb{R}^{p_B}$ and $\eta \in \mathcal{H} \subset \mathbb{R}^{p_{\kappa}}$. Let $c_{\theta, \eta}$ be a sufficiently regular non-negative solution on $(0, \infty) \times [0, T]$ of*

$$\frac{\partial c_{\theta, \eta}}{\partial t}(x, t) = -B_{\theta}(x)c_{\theta, \eta}(x, t) + 2 \int_x^{\infty} B_{\theta}(y)\kappa_{\eta}\left(\frac{x}{y}\right)\frac{1}{y}c_{\theta, \eta}(y, t) dy. \quad (32)$$

Assume that the neural daughter distribution satisfies

$$\int_0^1 z \kappa_{\eta}(z) dz = \frac{1}{2}. \quad (33)$$

Then the total mass

$$M_1(t) = \int_0^{\infty} x c_{\theta, \eta}(x, t) dx \quad (34)$$

is conserved in time.

Proof. Multiplying (32) by x and integrating over $x > 0$, we obtain

$$\frac{d}{dt} M_1(t) = - \int_0^{\infty} x B_{\theta}(x) c_{\theta, \eta}(x, t) dx + 2 \int_0^{\infty} x \int_x^{\infty} B_{\theta}(y) \kappa_{\eta}\left(\frac{x}{y}\right) \frac{1}{y} c_{\theta, \eta}(y, t) dy dx.$$

For the gain term, we exchange the order of integration:

$$2 \int_0^{\infty} B_{\theta}(y) c_{\theta, \eta}(y, t) \left[\int_0^y x \kappa_{\eta}\left(\frac{x}{y}\right) \frac{1}{y} dx \right] dy.$$

Using the change of variables

$$z = \frac{x}{y}, \quad x = yz, \quad dx = y dz,$$

the inner integral becomes

$$\int_0^y x \kappa_{\eta}\left(\frac{x}{y}\right) \frac{1}{y} dx = y \int_0^1 z \kappa_{\eta}(z) dz.$$

By (33), this gives

$$\int_0^y x \kappa_\eta \left(\frac{x}{y} \right) \frac{1}{y} dx = \frac{y}{2}.$$

Therefore, the gain contribution is

$$2 \int_0^\infty B_\theta(y) c_{\theta,\eta}(y, t) \frac{y}{2} dy = \int_0^\infty y B_\theta(y) c_{\theta,\eta}(y, t) dy.$$

This exactly cancels the loss contribution. Hence,

$$\frac{d}{dt} M_1(t) = 0.$$

□

5.3. Stability of the Neural Forward Map

The next result shows that the neural forward model is stable with respect to perturbations of the learned kernels. This property is useful for the inverse problem, because the optimization procedure changes the neural parameters and therefore changes B_θ and κ_η . The estimate is stated on the finite computational interval Ω_x , where the fragmentation operators are bounded under the assumptions below.

Theorem 1 (Stability of the neural forward map). *Let $\Omega_x \subset (0, \infty)$ be the computational size interval, and let*

$$(B_{\theta_1}, \kappa_{\eta_1}) \quad \text{and} \quad (B_{\theta_2}, \kappa_{\eta_2})$$

be two admissible neural fragmentation laws with

$$B_{\theta_\ell} \in L^\infty(\Omega_x), \quad \kappa_{\eta_\ell} \in L^\infty(0, 1), \quad \ell = 1, 2.$$

Let

$$c_{\theta_\ell, \eta_\ell} \in C([0, T]; L^2(\Omega_x)), \quad \ell = 1, 2,$$

be the corresponding solutions of the neural forward model with the same initial condition. Assume that, on the computational domains,

$$0 \leq B_{\theta_\ell}(x) \leq B_{\max}, \quad 0 \leq \kappa_{\eta_\ell}(z) \leq K_{\max}, \quad \ell = 1, 2.$$

Then, for every $T > 0$, there exists a constant $C_T > 0$ such that

$$\|c_{\theta_1, \eta_1}(t) - c_{\theta_2, \eta_2}(t)\|_{L^2(\Omega_x)} \leq C_T (\|B_{\theta_1} - B_{\theta_2}\|_\infty + \|\kappa_{\eta_1} - \kappa_{\eta_2}\|_\infty), \quad 0 \leq t \leq T. \quad (35)$$

Proof. We write the neural forward equation in abstract form as

$$\frac{dc}{dt} = \mathcal{A}_{\theta, \eta} c,$$

where $\mathcal{A}_{\theta, \eta}$ denotes the loss-gain fragmentation operator associated with B_θ and κ_η . Let

$$e(t) = c_{\theta_1, \eta_1}(t) - c_{\theta_2, \eta_2}(t).$$

Then

$$\frac{de}{dt} = \mathcal{A}_{\theta_1, \eta_1} c_{\theta_1, \eta_1} - \mathcal{A}_{\theta_2, \eta_2} c_{\theta_2, \eta_2}.$$

Adding and subtracting $\mathcal{A}_{\theta_1, \eta_1} c_{\theta_2, \eta_2}$, we get

$$\frac{de}{dt} = \mathcal{A}_{\theta_1, \eta_1} (c_{\theta_1, \eta_1} - c_{\theta_2, \eta_2}) + (\mathcal{A}_{\theta_1, \eta_1} - \mathcal{A}_{\theta_2, \eta_2}) c_{\theta_2, \eta_2}.$$

Using boundedness of the kernels on the finite computational interval, the first term is bounded by

$$L \|e(t)\|_{L^2(\Omega_x)},$$

for some constant $L > 0$. The second term is controlled by the difference between the two neural kernels:

$$\|(\mathcal{A}_{\theta_1, \eta_1} - \mathcal{A}_{\theta_2, \eta_2}) c_{\theta_2, \eta_2}\|_{L^2(\Omega_x)} \leq C(\|B_{\theta_1} - B_{\theta_2}\|_\infty + \|\kappa_{\eta_1} - \kappa_{\eta_2}\|_\infty),$$

where C depends on the computational domain, bounds on the solution, and the kernel bounds. Therefore,

$$\frac{d}{dt} \|e(t)\|_{L^2(\Omega_x)} \leq L \|e(t)\|_{L^2(\Omega_x)} + C(\|B_{\theta_1} - B_{\theta_2}\|_\infty + \|\kappa_{\eta_1} - \kappa_{\eta_2}\|_\infty).$$

Since both solutions have the same initial condition, $e(0) = 0$. Grönwall's inequality gives

$$\|e(t)\|_{L^2(\Omega_x)} \leq C_T(\|B_{\theta_1} - B_{\theta_2}\|_\infty + \|\kappa_{\eta_1} - \kappa_{\eta_2}\|_\infty),$$

for $0 \leq t \leq T$. This proves the result. $\square$

5.4. Existence of a Minimizer for the Regularized Neural Inverse Problem

The following result is a standard compactness argument, but it is useful because it shows that the proposed regularized neural inverse problem is well defined under natural boundedness assumptions on the parameter sets.

Theorem 2 (Existence of a regularized minimizer). *Assume that the neural parameters are restricted to compact sets,*

$$\theta \in \Theta \subset \mathbb{R}^{p_\theta}, \quad \eta \in \mathcal{H} \subset \mathbb{R}^{p_\eta},$$

where Θ and $\mathcal{H}$ are compact. Assume also that the neural forward map

$$(\theta, \eta) \mapsto c_{\theta, \eta}(\cdot, t_n) \in L^2(\Omega_x)$$

is continuous for every observation time t_n . Then the regularized inverse problem

$$\min_{\theta \in \Theta, \eta \in \mathcal{H}} J(\theta, \eta) \tag{36}$$

admits at least one minimizer.

Proof. The neural maps

$$\theta \mapsto B_\theta, \quad \eta \mapsto \kappa_\eta$$

are continuous with respect to the parameters. By assumption, the associated forward map is continuous at each observation time. Therefore, the data misfit

$$\mathcal{D}(\theta, \eta) = \sum_{n=1}^m \|c_{\theta, \eta}(\cdot, t_n) - c_{\text{obs}}(\cdot, t_n)\|^2$$

is continuous on $\Theta \times \mathcal{H}$. The regularization terms are also continuous under the same assumptions. Hence, the objective function

$$J(\theta, \eta) = \mathcal{D}(\theta, \eta) + \lambda_B R_B(\theta) + \lambda_\kappa R_\kappa(\eta)$$

is continuous on the compact set $\Theta \times \mathcal{H}$. By the Weierstrass theorem, J attains its minimum on $\Theta \times \mathcal{H}$. Therefore, the regularized neural inverse problem admits at least one minimizer. $\square$

5.5. Approximation Capacity of the Admissible Neural Class

The previous results show that the neural parametrization preserves the physical constraints. We now show that these constraints do not make the admissible neural class too restrictive.

Theorem 3 (Approximation of admissible fragmentation laws). *Let $\Omega_x \subset (0, \infty)$ be compact. Let $B^* \in C(\Omega_x)$ be strictly positive on Ω_x , and let $\kappa^* \in C([0, 1])$ be strictly positive on $[0, 1]$ and satisfy*

$$\int_0^1 \kappa^*(z) dz = 1.$$

Then, for every $\varepsilon > 0$, there exist finite-dimensional parameter vectors $\theta \in \mathbb{R}^{p_B}$ and $\eta \in \mathbb{R}^{p_\kappa}$, for suitable neural-network dimensions p_B and p_κ , defining neural networks $N_\theta : \Omega_x \rightarrow \mathbb{R}$ and $M_\eta : [0, 1] \rightarrow \mathbb{R}$ such that

$$\|B_\theta - B^*\|_\infty < \varepsilon, \quad \|\kappa_\eta - \kappa^*\|_\infty < \varepsilon,$$

where B_θ and κ_η are defined by (26) and (27).

Proof. Since B^* is strictly positive and continuous on the compact interval Ω_x , the function

$$g^*(s) = \log(e^{B^*(s)} - 1)$$

is continuous on Ω_x . By the universal approximation theorem [16,17], there exists a neural network N_θ such that

$$\|N_\theta - g^*\|_\infty$$

is arbitrarily small. Since the Softplus function is continuous, it follows that

$$B_\theta(s) = \text{Softplus}(N_\theta(s))$$

can approximate $B^*(s)$ uniformly on Ω_x .

Similarly, since κ^* is strictly positive and continuous on $[0, 1]$, the function

$$q^*(z) = \log \kappa^*(z)$$

is continuous. Again by the universal approximation theorem [16,17], there exists a neural network M_η such that M_η approximates q^* uniformly. The map

$$M \mapsto \frac{\exp(M(z))}{\int_0^1 \exp(M(u)) du}$$

is continuous with respect to uniform convergence on compact sets. Therefore, the normalized exponential parametrization can approximate κ^* uniformly. This proves the result. $\square$

5.6. Comment on Identifiability

The results above do not imply uniqueness of the recovered fragmentation law. Identifiability of fragmentation kernels from size-distribution observations is a delicate inverse problem and depends on the richness of the available data. In particular, the numerical results in Section 7 show that using multiple observation times substantially improves the reliability of the recovery of both B and κ . Thus, the theoretical results in this section should be understood as structural guarantees for the proposed neural inverse framework, while practical identifiability is investigated numerically. Accordingly, the recovered pair (B, κ) should be interpreted as an admissible effective fragmentation law associated with the available observations and the imposed physical constraints. In this sense, reliability is not based on a global uniqueness theorem, but on the combination of physical constraints, regularization, multiple observation times, and numerical sensitivity checks.

6. Synthetic Validation

In this section, we validate the proposed neural inverse framework using synthetic data. The purpose of the synthetic experiments is to test whether the method can recover known fragmentation laws. This step is essential because, for experimental data, the true fragmentation kernel is generally unknown.

Synthetic observations are generated as follows. We first choose ground-truth functions B_{true} and κ_{true} . We then solve the direct fragmentation equation using an independent conservative sectional solver on a fine grid. The resulting distributions

$$c_{\text{true}}(x, t_1), \dots, c_{\text{true}}(x, t_m)$$

are treated as observations. The inverse problem is then solved using the quadrature-based solver described in Section 4 on a coarser grid. This separation between the data-generation solver and the inverse solver is used to avoid an inverse crime. In some tests, small additive noise is introduced to emulate measurement uncertainty. The inverse method is then applied without using the knowledge of the true kernel, and the recovered functions are compared with the ground truth.

The independent data-generation solver is implemented as a conservative sectional method on a fine pivot grid. For a daughter size r , that falls between two neighboring pivots $X_\ell \leq r \leq X_{\ell+1}$, its contribution is distributed to the two neighboring size classes using the fixed-pivot weights

$$\omega_\ell(r) = \frac{X_{\ell+1} - r}{X_{\ell+1} - X_\ell}, \quad \omega_{\ell+1}(r) = \frac{r - X_\ell}{X_{\ell+1} - X_\ell}.$$

Hence,

$$\omega_\ell(r)X_\ell + \omega_{\ell+1}(r)X_{\ell+1} = r,$$

so that the first moment is preserved during the redistribution step. For each parent size X_j , daughter sizes zX_j are redistributed in this way and integrated over the ground-truth daughter distribution $\kappa_{\text{true}}(z)$, with the binary factor accounting for the two daughter fragments. The resulting fine-grid linear system is solved using matrix-exponential time propagation. The synthetic observations are then linearly interpolated onto the inverse grid before solving the inverse problem with the quadrature-based solver of Section 4. Thus, the observations and the inverse reconstruction are produced by different discrete forward operators.

Unless otherwise stated, the synthetic experiments are performed on a finite computational interval $\Omega_x = [x_{\min}, x_{\max}]$. The inverse problem is solved using the deterministic quadrature discretization of Section 4 and a fourth-order Runge–Kutta time integrator. The neural functions B_θ and κ_η are optimized by minimizing the regularized objective (18). When observation noise is added, it is scaled relative to the maximum value of the corresponding distribution and the noisy observations are projected onto the non-negative range.

The relative reconstruction errors for B and κ are measured using

$$E_B = \frac{\|B_{\text{true}} - \hat{B}\|_{L^2(\Omega_x)}}{\|B_{\text{true}}\|_{L^2(\Omega_x)}}, \quad (37)$$

and

$$E_\kappa = \frac{\|\kappa_{\text{true}} - \hat{\kappa}\|_{L^2(0,1)}}{\|\kappa_{\text{true}}\|_{L^2(0,1)}}. \quad (38)$$

We also evaluate the relative reconstruction error of the particle-size distributions at the observation times:

$$E_c(t_n) = \frac{\|c_{\text{true}}(\cdot, t_n) - c_{\text{rec}}(\cdot, t_n)\|_{L^2(\Omega_x)}}{\|c_{\text{true}}(\cdot, t_n)\|_{L^2(\Omega_x)}}. \quad (39)$$

6.1. Numerical Settings

The synthetic inverse experiments use the deterministic forward solver described in Section 4 for the inverse recovery. Synthetic observations are generated on a finer grid using the independent conservative sectional solver described above, while the inverse recovery is carried out on a coarser grid. The neural parameters are optimized by minimizing the regularized objective function (18). The main numerical settings are summarized in Table 1.

Table 1. Numerical settings used in the synthetic inverse experiments.

Experiment	Data Grid N	Inverse Grid N	Observation Times	Unknowns
Recover κ , known B	200	80	0.2, 0.5, 0.8, 1.2	κ
Recover B , known κ	200	80	0.2, 0.5, 0.8, 1.2	B
Joint recovery, unimodal κ	200	80	0.2, 0.5, 0.8, 1.2	B, κ
Joint recovery, bimodal κ	200	80	0.15, 0.35, 0.65, 1.0, 1.35	B, κ
Noise robustness tests	200	80	0.2, 0.5, 0.8, 1.2	B, κ
Observation-time sensitivity tests	200	80	subsets of 0.2, 0.5, 0.8, 1.2	B, κ

The neural parameters are optimized by minimizing the regularized objective function (18) using the L-BFGS-B algorithm [18]. The neural-network and optimization details used in the inverse experiments are reported in Table 2.

The regularization coefficients were selected from preliminary calibration runs by balancing the data misfit with the smoothness of the recovered functions, while avoiding visibly oscillatory kernels. The sensitivity checks focus on the numerical and observational choices most directly affecting the inverse recovery: grid size and time step are examined in Section 6.5, while noise level and number of observation times are examined in Sections 7.1 and 7.2, respectively. Sensitivity to initialization is mitigated by using multiple random restarts and selecting the lowest regularized objective value.

Table 2. Neural-network and optimization settings used in the inverse experiments.

Quantity	Setting
Architecture	One hidden layer, tanh activation
Hidden neurons	5 for B_θ ; 6 for κ_η except 10 in the bimodal test
B_θ constraint	Softplus $+10^{-6}$
κ_η constraint	Normalized exponential; symmetrized when needed
Parameter bounds	Box constraints $[-5, 5]$ on trainable neural parameters
Initialization	Normal initialization, std. 0.15–0.25
Optimizer	L-BFGS-B
Maximum iterations	up to 2200
Stopping tolerances	Projected gradient 10^{-7} , objective change 10^{-9}
Regularization	$\lambda_B = 10^{-3}$, $\lambda_\kappa = 10^{-6}$
Restarts	Multiple random initializations when needed
Selection criterion	Lowest regularized objective value
Computational cost	Approximately 10–40 s per inverse run
Sensitivity checks	Noise, observation-time, grid, and time-step tests

6.2. Recovery of the Daughter Distribution with Known Breakup Rate

We first consider the case where the breakup rate B is assumed known and only the daughter distribution κ is recovered. This test isolates the ability of the neural density representation (10) to identify the fragmentation ratio distribution.

The ground-truth breakup rate is chosen as

$$B_{\text{true}}(s) = 1.8s^{1.2}, \quad (40)$$

and the daughter distribution is chosen as a symmetric beta distribution,

$$\kappa_{\text{true}}(z) = \text{Beta}(z; 5, 5). \quad (41)$$

Synthetic observations are generated at several observation times, and the neural inverse problem is solved for κ_η .

Figure 3 shows the recovered daughter distribution together with the ground truth. The neural representation closely captures the shape of the distribution. The relative error is

$$E_\kappa = 0.0127.$$

Figure 4 compares the true and reconstructed particle-size distributions at the observation times. The mean reconstruction error is 0.0081.




Figure 3. Synthetic inverse test with known B : recovery of the daughter distribution $\kappa(z)$.




Figure 4. Synthetic inverse test with known B : comparison between true and reconstructed particle-size distributions.

6.3. Recovery of the Breakup Rate with Known Daughter Distribution

We next consider the complementary case where the daughter distribution κ is assumed known and the breakup rate B is recovered. This test isolates the ability of the neural parametrization (9) to recover a size-dependent breakup rate.

The same ground-truth functions (40) and (41) are used. The inverse problem is solved for B_θ , while κ is kept fixed at its true value.

Figure 5 shows the recovered breakup rate. The neural approximation follows the true increasing rate with relative error

$$E_B = 0.0324.$$

The corresponding reconstructed particle-size distributions are shown in Figure 6. The mean reconstruction error is 0.0089.




Figure 5. Synthetic inverse test with known κ : recovery of the breakup rate $B(s)$.




Figure 6. Synthetic inverse test with known κ : comparison between true and reconstructed particle-size distributions.

6.4. Joint Recovery of the Breakup Rate and Daughter Distribution

We now consider the fully coupled inverse problem in which both B and κ are unknown. This setting is more difficult because the breakup rate and the daughter distribution can partially compensate each other in the predicted size distributions.

The same ground-truth functions are used:

$$B_{\text{true}}(s) = 1.8s^{1.2}, \quad \kappa_{\text{true}}(z) = \text{Beta}(z; 5, 5).$$

Both B_θ and κ_η are trained simultaneously using several observation times.

Figures 7 and 8 show the joint recovery of B and κ , respectively. The relative errors are

$$E_B = 0.02554, \quad E_\kappa = 0.01109.$$

Although the joint problem is harder than the two previous tests, the recovered functions remain close to the ground truth. Figure 9 shows that the reconstructed size distributions agree well with the true distributions, with mean reconstruction error 0.00974.




Figure 7. Joint synthetic inverse test: recovery of the breakup rate $B(s)$.




Figure 8. Joint synthetic inverse test: recovery of the daughter distribution $\kappa(z)$.




Figure 9. Joint synthetic inverse test: comparison between true and reconstructed particle-size distributions.

6.5. Grid and Time-Step Sensitivity Check

To verify that the joint recovery results are not an artifact of a particular numerical discretization, we repeated the unimodal joint recovery test using different inverse grids and time steps, while keeping the synthetic observations generated on the independent fine grid. The results are reported in Table 3. The errors remain small and of the same order under grid refinement and time-step reduction, indicating that the reported recovery results are not tied to a particular discretization choice. The small non-monotonic variation is expected because each inverse recovery involves a nonconvex neural optimization problem.

Table 3. Grid and time-step sensitivity of the joint unimodal inverse recovery. Synthetic observations are generated on the independent fine grid, while the inverse problem is solved using different quadrature grids and time steps.

Inverse Grid N	Time Step Δt	E_B	E_κ	Mean Reconstruction Error
60	0.008	0.03262	0.01137	0.01210
80	0.006	0.02554	0.01109	0.00974
100	0.004	0.02972	0.00918	0.00972

6.6. Bimodal Daughter Distribution

To test the method beyond a unimodal daughter distribution, we consider a more challenging synthetic kernel. The breakup rate is chosen as

$$B_{\text{true}}(s) = 0.35 + 1.45s^{1.4}, \quad (42)$$

and the daughter distribution is chosen as the symmetric bimodal mixture

$$\kappa_{\text{true}}(z) = \frac{1}{2} \text{Beta}(z; 2.5, 8) + \frac{1}{2} \text{Beta}(z; 8, 2.5). \quad (43)$$

This distribution favors asymmetric splitting, producing one relatively small fragment and one relatively large fragment.

Figures 10 and 11 show the recovered B and κ . The method still captures the increasing breakup rate and the bimodal structure of the daughter distribution, although the joint bimodal case is more challenging than the previous synthetic tests. The relative errors are

$$E_B = 0.16372, \quad E_\kappa = 0.08123.$$

The reconstructed particle-size distributions are shown in Figure 12. The mean reconstruction error is 0.01583.




Figure 10. Bimodal synthetic test: recovery of the breakup rate $B(s)$.




Figure 11. Bimodal synthetic test: recovery of the bimodal daughter distribution $\kappa(z)$.




Figure 12. Bimodal synthetic test: comparison between true and reconstructed particle-size distributions.

6.7. Summary of Synthetic Validation

Table 4 summarizes the main synthetic validation results. The results show that the method can recover the daughter distribution when B is known, the breakup rate when κ

is known, and both functions jointly. The bimodal test further shows that the method is not restricted to simple unimodal daughter laws, although this case is more challenging than the unimodal setting.

Table 4. Summary of synthetic inverse validation errors.

Test	E_B	E_κ	Mean Reconstruction Error
Known B , recover κ	–	0.0127	0.0081
Known κ , recover B	0.0324	–	0.0089
Joint recovery, unimodal κ	0.02554	0.01109	0.00974
Joint recovery, bimodal κ	0.16372	0.08123	0.01583

7. Robustness and Sensitivity Analysis

The previous section shows that the proposed method can recover known fragmentation laws from synthetic observations. In this section, we further investigate the stability of the joint inverse recovery with respect to observational noise and the number of available observation times. These tests are performed in the joint recovery setting, where both B and κ are unknown, since this is the most challenging synthetic inverse configuration considered in this work. They are important because inverse fragmentation problems may be sensitive to measurement errors and incomplete observations.

7.1. Robustness with Respect to Observational Noise

We first test the influence of noise on the joint recovery of B and κ . Starting from the same synthetic setting used in the unimodal joint recovery test, we add noise to the observed distributions:

$$c_{\text{obs}}(x_i, t_n) = \max \left\{ c_{\text{true}}(x_i, t_n) + \sigma_{\text{noise}} \max_i c_{\text{true}}(x_i, t_n) \xi_{i,n}, 0 \right\}, \quad (44)$$

where $\xi_{i,n}$ are independent standard Gaussian perturbations and σ_{noise} controls the relative noise level. The projection onto the non-negative range is used to avoid nonphysical negative observations. The inverse problem is then solved for different noise levels.

Figure 13 shows the relative errors on B , on κ , and on the reconstructed particle-size distributions as the noise level increases. The results indicate that the proposed method remains stable for moderate noise levels. Even for a noise level of 3%, the mean reconstruction error remains close to 3%, and the recovered kernels remain within a reasonable error range.




Figure 13. Robustness of the joint neural inverse recovery of B and κ in the unimodal benchmark for relative noise levels 0%, 0.5%, 1%, and 3%. Independent Gaussian perturbations are scaled by the maximum value of the distribution at each observation time, as defined in (44), and negative perturbed values are projected to zero. The figure reports the relative errors E_B and E_κ , together with the mean reconstruction error; all remaining numerical settings are unchanged from the noise-free joint-recovery experiment.

The numerical values are summarized in Table 5. The overall behavior of the reconstruction error suggests that the regularized physics-constrained formulation is not overly sensitive to moderate perturbations of the observed distributions.

Table 5. Effect of observational noise on the joint recovery of B and κ .

Noise Level	E_B	E_κ	Mean Reconstruction Error
0.0%	0.02554	0.01109	0.00974
0.5%	0.01882	0.02632	0.01195
1.0%	0.03113	0.01317	0.01040
3.0%	0.11462	0.04044	0.02976

7.2. Influence of the Number of Observation Times

We next investigate the effect of the number of observed particle-size distributions. This is directly related to the identifiability of the inverse problem. If too few observations are available, different fragmentation laws may produce similar distributions, making the recovery unstable or non-unique.

We consider four cases: one observation time, two observation times, three observation times, and four observation times. In all cases, the same ground-truth fragmentation law is used, and the inverse method is applied to recover both B and κ .

Figure 14 shows that the recovery improves substantially when moving from a single observation time to multiple observation times. With only one observation time, the errors are

$$E_B = 0.38454, \quad E_\kappa = 0.43556.$$

The largest improvement occurs when moving from one observation time to two observation times. This behavior is expected, because two time levels provide information on the temporal change of the particle-size distribution, whereas a single observation is much less informative for identifying a rate-dependent kernel. In the present unimodal benchmark,

the two selected observation times already capture much of the relevant evolution. With four observation times, the errors decrease to

$$E_B = 0.02554, \quad E_\kappa = 0.01109.$$




Figure 14. Effect of the number and temporal placement of the observed particle-size distributions on the joint recovery of B and κ in the unimodal benchmark. The cases use one observation at $t = 0.8$, two observations at $t = 0.5, 1.2$, three observations at $t = 0.2, 0.8, 1.2$, and four observations at $t = 0.2, 0.5, 0.8, 1.2$. The figure reports the relative errors E_B and E_κ , together with the mean reconstruction error; all remaining numerical settings and the ground-truth fragmentation law are kept unchanged.

The corresponding numerical values are reported in Table 6. These results suggest that multiple observed distributions provide substantially more information about the fragmentation law than a single observation. The smaller and non-monotone changes observed when using three or four observation times indicate that the benefit of additional observations also depends on their temporal placement and on the nonconvex optimization landscape. Overall, this supports the use of time-resolved or multi-stage size-distribution data whenever available.

Table 6. Effect of the number of observation times on the joint inverse recovery.

Number of Times	Observation Times	E_B	E_κ	Mean Reconstruction Error
1	0.8	0.38454	0.43556	0.11113
2	0.5, 1.2	0.00856	0.01134	0.00633
3	0.2, 0.8, 1.2	0.04807	0.01514	0.01404
4	0.2, 0.5, 0.8, 1.2	0.02554	0.01109	0.00974

Overall, the robustness and sensitivity tests indicate that the proposed framework is stable under moderate observational noise and benefits from multiple observation times. These findings are consistent with the ill-posed nature of inverse fragmentation problems and motivate the use of sufficiently rich observation data.

8. Experimental Estimation of Effective Daughter-Volume-Fraction Distributions from Droplet-Breakup Data

The synthetic studies in Sections 6 and 7 validate the complete PDE-constrained inverse framework under controlled time-resolved observations, including the separate

and joint recovery of the breakup rate B and the daughter distribution κ . The experimental dataset considered in this section has a different observational structure: it contains parent-drop and final child-fragment measurements, but not time-resolved particle-size distributions $c_{\text{obs}}(x, t_i)$. Consequently, the analysis presented here does not solve the PDE-constrained inverse problem and does not identify the breakup-rate function $B(s)$.

Instead, the observed paired parent-droplet and daughter-fragment measurements are used to estimate an effective marginal daughter-volume-fraction distribution κ_{exp} by constrained maximum likelihood, using the same positive and normalized neural representation developed for the daughter distribution in the full inverse framework. The synthetic and experimental studies are therefore complementary parts of the paper: the former validates the complete inverse identification of B and κ , whereas the latter demonstrates the real-data applicability of the identifiable daughter-distribution component under the information available in the experimental dataset.

8.1. Experimental Data and Preprocessing

The experimental data are taken from the supplementary material of the droplet-breakup study of Jackiw and Ashgriz [14]. The dataset contains parent-drop properties, breakup regimes, and measured child-fragment sizes associated with different fragmentation mechanisms. In particular, the measurements include rim fragments and node fragments.

For each measured child fragment, we define the normalized child-volume fraction

$$\zeta_V = \left(\frac{d_{\text{child}}}{D_{\text{drop}}} \right)^3, \quad (45)$$

where d_{child} is the measured child-fragment diameter and D_{drop} is the corresponding parent-drop diameter. Since droplet volume is proportional to the cube of the diameter, ζ_V represents the fraction of the parent-drop volume carried by the measured child fragment. This preprocessing makes the experimental measurements dimensionless and expresses fragment size in terms of the daughter-to-parent volume fraction, which is compatible with the additive size variable used in the fragmentation model. Since the experimental density represents an effective marginal distribution for potentially multi-fragment events, no binary first-moment constraint is imposed on κ_{exp} .

8.2. Breakup Modes and Fragment Mechanisms

The experimental labels used in the dataset correspond to different aerodynamic breakup morphologies described in droplet-breakup studies [14,19]. In the Bag regime, the parent droplet deforms into a thin bag-like film bounded by a thicker rim; the rupture of this film produces small droplets. In the BS (bag–stamen) regime, the morphology is similar to bag breakup, but a central stamen or core remains during the breakup and may generate an additional larger fragment. In the MB (multi-bag) regime, several bag-like structures may form and rupture. These aerodynamic breakup regimes generally produce multiple child droplets; accordingly, the experimental distributions considered below are interpreted as effective marginal distributions of individual daughter-volume fractions rather than as literal binary daughter kernels.

In addition to these global breakup modes, the dataset distinguishes between local fragment mechanisms. Rim fragments are produced by the breakup of the liquid rim surrounding or retracting from the bag. Node fragments are associated with larger structures or nodes along the rim and typically produce larger child droplets. This distinction is important because different contributions, including bag rupture, rim fragmentation, and node formation, shape the final child-droplet size distribution [14].

In this terminology, Bag, BS (bag–stamen), and MB (multi-bag) are global breakup-regime labels, whereas rim and node identify local fragment-production mechanisms within the breakup process. In the present work, these labels are used to compare effective daughter-volume-fraction distributions across breakup regimes and fragment mechanisms.

Figure 15 summarizes the relation between the breakup-regime labels and the fragment-mechanism labels used in the experimental analysis.




Figure 15. Conceptual schematic of the breakup-regime and fragment-mechanism labels used in the experimental analysis.

After preprocessing, the dataset contains 944 normalized daughter-volume-fraction measurements. Among them, 701 correspond to rim fragments and 243 correspond to node fragments. The data also include the breakup-mode labels Bag, BS, and MB. Table 7 summarizes the main empirical statistics of the normalized daughter-volume fractions, both by local fragmentation mechanism and by global breakup mode.

Table 7. Summary statistics of the normalized daughter-volume fraction $\zeta_V = (d_{\text{child}}/D_{\text{drop}})^3$ in the experimental droplet-breakup dataset.

Group	Number of Measurements	Mean ζ_V	Median ζ_V
All fragments	944	0.0238	0.0102
Node fragments	243	0.0600	0.0416
Rim fragments	701	0.0113	0.0060
BS mode	523	0.0174	0.0075
Bag mode	354	0.0357	0.0151
MB mode	67	0.0120	0.0057

The clear difference between rim and node fragments already indicates that the daughter-volume-fraction distribution depends on the physical fragmentation mechanism. Rim fragments carry a smaller volume fraction on average, whereas node fragments carry a larger volume fraction.

8.3. Effective Daughter-Volume-Fraction Distribution

Since the experimental data provide parent and child sizes but not full time-resolved distributions $c(x, t)$, we focus on learning an effective daughter-volume-fraction distribution

$$\kappa_{\text{exp}}(\zeta_V), \quad (46)$$

where ζ_V is the normalized child-volume fraction defined in (45). This corresponds to the experimental analogue of the theoretical daughter distribution $\kappa(z)$, with the normalized daughter-size fraction specialized to the daughter-to-parent volume fraction.

We use the same normalized neural density parametrization introduced in (10):

$$\kappa_{\eta,\text{exp}}(\zeta_V) = \frac{\exp(M_\eta(\zeta_V))}{\int_0^1 \exp(M_\eta(u)) du}. \quad (47)$$

The parameters η are fitted to the empirical daughter-to-parent volume-fraction measurements. Given normalized measurements $\zeta_{V,1}, \dots, \zeta_{V,N_{\text{exp}}}$, the parameters are estimated by minimizing the negative log-likelihood

$$\mathcal{L}_{\text{exp}}(\eta) = -\frac{1}{N_{\text{exp}}} \sum_{i=1}^{N_{\text{exp}}} \log \kappa_{\eta,\text{exp}}(\zeta_{V,i}). \quad (48)$$

Unlike the PDE-constrained inverse problems considered in the synthetic sections, the objective (48) does not involve the forward fragmentation equation or the breakup-rate function B . It therefore defines a constrained neural density-estimation problem for the effective marginal daughter-volume-fraction distribution.

To quantify the sampling uncertainty of the learned experimental densities, we perform a stratified cluster bootstrap with 500 replicates. The experimental Test ID is used as the resampling unit so that all daughter fragments measured within the same breakup test remain together, and Test IDs are resampled with replacement separately within each of the Bag, BS, and MB breakup-mode strata. For every replicate, the effective densities are refitted using the same neural architecture and L-BFGS-B optimization settings as for the full dataset. The 2.5th and 97.5th percentiles of the bootstrap estimates at each value of ζ_V define 95% pointwise confidence bands.

Figure 16 shows the global effective daughter-volume-fraction distribution learned from all 944 child fragments, together with its 95% pointwise bootstrap confidence band. The distribution is concentrated near small daughter-volume fractions, indicating that most measured child fragments carry a small fraction of the parent-drop volume. The confidence band remains relatively narrow over most of the domain, supporting the stability of the estimated global distribution under cluster resampling of the experimental tests.

8.4. Dependence on Fragmentation Mechanism

We next separate the measurements by fragmentation mechanism. Figure 17 compares the effective distributions learned from rim and node fragments, together with their 95% pointwise bootstrap confidence bands.

The result is physically coherent. The rim-fragment distribution is concentrated at smaller values of ζ_V , with mean value approximately

$$\bar{\zeta}_{V,\text{rim}} = 0.0113.$$

By contrast, the node-fragment distribution is shifted toward larger values, with mean value approximately

$$\bar{\zeta}_{V,\text{node}} = 0.0600.$$

The separation between the two distributions remains clear when the bootstrap uncertainty is taken into account, supporting the ability of the neural density estimates to capture meaningful physical differences between fragmentation mechanisms.




Figure 16. Effective daughter-volume-fraction distribution $\kappa_{\text{exp}}(\zeta_V)$ learned from all 944 experimental child-fragment measurements. The histogram represents the empirical distribution, the solid curve is the full-sample neural estimate, and the shaded region denotes the 95% pointwise confidence band obtained from 500 stratified cluster-bootstrap replicates using Test ID as the resampling unit.




Figure 17. Effective daughter-volume-fraction distributions learned separately from 701 rim fragments and 243 node fragments. The solid curves represent the full-sample neural estimates, and the shaded regions denote the 95% pointwise confidence bands obtained from 500 stratified cluster-bootstrap replicates. Rim fragmentation produces smaller normalized fragments, whereas node fragmentation produces larger normalized fragments.

8.5. Dependence on Breakup Mode

The data also allow a comparison between different breakup regimes. Figure 18 shows the effective distributions obtained for the Bag, BS, and MB modes, together with their 95% pointwise bootstrap confidence bands.




Figure 18. Effective daughter-volume-fraction distributions learned separately from 354 Bag-mode, 523 BS-mode, and 67 MB-mode fragment measurements. The solid curves represent the full-sample neural estimates, and the shaded regions denote the 95% pointwise confidence bands obtained from 500 stratified cluster-bootstrap replicates. The distributions reflect differences in the normalized child-fragment sizes produced by the different breakup regimes.

The Bag mode has a larger mean normalized daughter-volume fraction than the BS and MB modes:

$$\bar{\zeta}_{V,\text{Bag}} = 0.0357, \quad \bar{\zeta}_{V,\text{BS}} = 0.0174, \quad \bar{\zeta}_{V,\text{MB}} = 0.0120.$$

The differences between the estimated distributions remain visible when bootstrap uncertainty is taken into account, suggesting that the learned effective daughter-volume-fraction distributions capture regime-dependent fragmentation patterns. The uncertainty is more pronounced near the sharp peak of the MB distribution, which should be interpreted cautiously because the MB subset contains substantially fewer observations than the Bag and BS subsets.

8.6. Interpretation and Limitations

The experimental results show that the proposed neural representation can extract smooth effective daughter-volume-fraction distributions from measured child-fragment data. The learned distributions are physically meaningful: rim fragments carry smaller daughter-volume fractions than node fragments, and different breakup modes lead to different normalized volume-fraction distributions. The bootstrap confidence bands support these qualitative differences, although the MB-mode estimate exhibits greater uncertainty near its sharp peak because this subset contains substantially fewer measurements.

However, the interpretation of the experimental application must be made carefully. The available dataset does not provide complete time-resolved size distributions of the form

$$c_{\text{obs}}(x, t_1), \dots, c_{\text{obs}}(x, t_m).$$

Therefore, the experimental data are not sufficient to identify a full breakup-rate function $B(s)$ through the fragmentation equation. Instead, the real-data analysis identifies an effective marginal daughter-volume-fraction distribution κ_{exp} from paired parent-droplet and daughter-fragment measurements.

The full recovery of both B and κ is validated in the synthetic experiments, where the true fragmentation law is known and time-resolved distributions are generated independently of the inverse solver. The experimental section then illustrates how the same

physics-constrained neural representation can be used to estimate meaningful daughter-volume-fraction distributions from real droplet-breakup measurements. Thus, the experimental application should be viewed as a real-data demonstration of the constrained neural density representation for final child-fragment measurements, rather than as a full kinetic validation of the progressive fragmentation model.

Finally, because real droplet breakup may generate several fragments in a single event, the learned density should be interpreted as an effective marginal daughter-volume-fraction distribution associated with the observed fragments. The binary fragmentation model is most directly applicable to breakup events dominated by two principal daughters, or as an effective description of a sequence of binary splitting events. It should not be interpreted as a literal representation of aerodynamic regimes in which several fragments are generated simultaneously. In particular, the Bag, BS, and MB regimes considered in the experimental dataset are not strictly binary, because their bag, rim, node, stamen, or multiple-bag structures can generate several child droplets; the binary model is therefore physically literal only for low-multiplicity events dominated by two principal daughters and otherwise serves as a sequential-cascade or reduced effective approximation. Extending the present framework to explicitly model multi-fragment breakup events is a natural direction for future work.

9. Discussion

The synthetic inverse results indicate that the proposed physics-constrained neural framework provides a flexible approach for inverse fragmentation problems, while the experimental results demonstrate the applicability of its constrained daughter-density representation to real fragment measurements. The key idea is to represent the unknown fragmentation law through the physically interpretable functions B and κ , while enforcing admissibility constraints directly in the neural parametrization. In this way, the neural networks provide flexibility without removing the physical structure imposed by the fragmentation equation.

From a practical viewpoint, the method is useful when direct measurement of the fragmentation kernel is difficult, but size-distribution data are available. In such situations, the framework identifies an effective fragmentation law that is consistent with the observed evolution and with the physical constraints of the model. Once identified, this effective law can be inserted into the forward fragmentation equation to reproduce, compare, or predict particle-size evolution under comparable conditions. The recovered functions should therefore be interpreted as effective kernels associated with the experimental or numerical conditions under consideration, rather than as universal material laws independent of the flow configuration.

This interpretation also defines the generalizability and predictive scope of the present approach. When the available data correspond to a single operating condition, the recovered size-dependent kernel is inherently tied to that condition. Therefore, it should not be interpreted as a general predictive law for different flow configurations, turbulence levels, material properties, or breakup regimes. In particular, a kernel identified from a specific particle-size or droplet-size distribution can be used for single-condition reconstruction and for forward prediction under the same, or sufficiently comparable, physical conditions, but it does not by itself provide predictive capability across different operating conditions. In the experimental droplet-breakup application, where only final parent-child fragment measurements are available, the real-data analysis should therefore be interpreted as a reconstruction of an effective daughter-volume-fraction distribution, not as a cross-condition predictive model. This distinction is important to avoid overstating the predictive capability of a kernel identified from a single dataset.

This limitation is not specific to the neural parametrization, but follows from the information content of the data: a single-condition dataset cannot identify how the breakage law varies with external physical parameters. Generalization would require particle-size or droplet-size distribution data collected across multiple operating conditions and the replacement of the present size-only laws by conditional multi-factor breakage kernels, for example

$$B(s; \varepsilon, \sigma, \mu, \rho, \text{We}, \text{Oh}, u, \dots), \quad \kappa(z \mid s, \varepsilon, \sigma, \mu, \rho, \text{We}, \text{Oh}, u, \dots),$$

where ε may denote the turbulence dissipation rate, σ the interfacial tension, μ the fluid viscosity, ρ the fluid density, We the Weber number, Oh the Ohnesorge number, and u a characteristic velocity or flow descriptor. These variables are listed as possible conditioning inputs, depending on which physical quantities are available in a multi-condition dataset. In such a multi-condition setting, the same physics-constrained neural construction could be used to regress fragmentation laws that vary with the operating regime, while preserving positivity, normalization, and conservation constraints. The present study should therefore be viewed as a single-condition inverse-identification framework and a foundation for conditional, multi-factor kernel learning.

The synthetic experiments demonstrate that the method can recover the daughter distribution κ when the breakup rate is known, recover the breakup rate B when the daughter distribution is known, and recover both functions jointly when both are unknown. The joint recovery problem is naturally more difficult because changes in the breakup rate and changes in the daughter distribution can have partially compensating effects on the resulting particle-size distributions. This reflects the ill-posed nature of inverse fragmentation problems. Accordingly, the present work does not claim global uniqueness of the recovered pair (B, κ) . Rather, the goal is to recover an admissible effective fragmentation law under the imposed physical constraints and the available observations. Nevertheless, the proposed regularized formulation recovers both functions with satisfactory accuracy in the unimodal synthetic tests and remains effective in the more challenging bimodal setting.

The bimodal synthetic test is particularly important. It shows that the method is not limited to smooth unimodal daughter distributions. The neural parametrization is able to capture a more complex daughter-size law involving two dominant daughter-size regimes. This is relevant for physical fragmentation processes where different mechanisms may contribute simultaneously to the final size distribution.

The robustness tests suggest that the method remains stable under moderate observational noise. This is important for applications, since experimental particle-size distributions are often affected by measurement noise, sampling uncertainty, and preprocessing errors. The observation-time sensitivity test also highlights a central feature of inverse fragmentation problems: a single observed distribution may not contain enough information to identify the kernel reliably. Multiple observation times significantly improve the reliability of the recovery of B and κ , supporting the importance of time-resolved or multi-stage observations for practical inverse identification.

The experimental droplet-breakup application illustrates how the constrained neural daughter-density representation can be used with real fragmentation measurements. Since the available dataset contains parent and child fragment sizes, but not full time-resolved particle-size distributions, the experimental analysis focuses on the effective marginal daughter-volume-fraction distribution κ_{exp} . The learned distributions are physically meaningful: rim fragments are concentrated at smaller daughter-volume fractions, while node fragments are shifted toward larger daughter-volume fractions. The corresponding 95% pointwise bootstrap confidence bands support these qualitative differences. The learned distributions also vary across breakup modes, indicating that the effective density estimates

capture regime-dependent fragmentation patterns, although the MB-mode estimate carries greater uncertainty because of its smaller sample size.

A limitation of the present experimental application is that it does not allow the full identification of the breakup rate $B(s)$. In the fragmentation equation, $B(s)$ represents a rate of breakup and therefore requires information about the temporal evolution of the size distribution, or repeated experiments over sufficiently varied parent sizes with measured breakup frequencies. The droplet-breakup dataset used here is well suited for estimating effective daughter-volume-fraction distributions, but not for identifying a complete size-dependent breakup rate. For this reason, the recovery of both B and κ is validated in the synthetic setting, while the real-data application is interpreted as the estimation of an effective marginal daughter-volume-fraction distribution.

Another limitation is that the mathematical model considered in this paper is a binary fragmentation equation, whereas aerodynamic droplet breakup can generate multiple child fragments through several mechanisms. The effective marginal daughter-volume-fraction distribution estimated from the experimental data should therefore be interpreted as a reduced representation of a more complex multi-fragment process. Extending the proposed neural inverse framework to multi-fragment fragmentation kernels is a natural direction for future work.

Despite these limitations, the proposed approach provides a useful bridge between inverse fragmentation equations and data-driven modeling. The neural networks provide flexible function representations, while the fragmentation equation imposes physical structure. Future work may include the recovery of breakup rates from time-resolved experimental distributions, the use of high-fidelity simulations such as molecular-dynamics or other physics-based breakup simulations to generate time-resolved fragmentation kinetics, the extension to multi-fragment breakup models, and the development of conditional multi-factor fragmentation laws using datasets collected across several operating conditions. Such extensions would make it possible to move from single-condition reconstruction toward more general and predictive data-driven fragmentation models.

10. Conclusions

We proposed a physics-constrained neural framework for the inverse identification of fragmentation laws from particle-size distribution data. Starting from a size-structured binary fragmentation equation, the fragmentation kernel was decomposed into a breakup rate $B(s)$ and a daughter-size distribution $\kappa(z)$. Both functions were represented by neural networks designed to satisfy physical constraints, including positivity of the breakup rate, positivity of the daughter distribution, and normalization of the daughter distribution.

The direct fragmentation equation was solved using a first-moment-conserving quadrature discretization, and the inverse problem was formulated as a regularized optimization problem constrained by the fragmentation dynamics. Structural properties of the proposed framework were established, showing that the neural parametrization preserves admissibility and the mass-conservation structure of the fragmentation model. The stability of the neural forward map, the existence of a regularized minimizer under compactness assumptions, and the approximation capacity of the admissible neural class were also discussed.

Synthetic experiments validated the method in several settings. The framework recovered the daughter distribution when the breakup rate was known, recovered the breakup rate when the daughter distribution was known, and recovered both functions jointly. The synthetic observations were generated independently of the inverse solver, which supports the reliability of the validation procedure. The method also remained effective for a more challenging bimodal daughter distribution, indicating that it can capture more complex

daughter-size laws. Robustness tests showed stability under moderate observational noise, while sensitivity tests indicated that multiple observation times improve kernel recovery.

The constrained neural daughter-density representation was then applied to experimental droplet-breakup measurements. From normalized daughter-to-parent volume fractions, effective marginal daughter-volume-fraction distributions were estimated globally, by fragmentation mechanism, and by breakup mode. The results were physically consistent: rim fragments were associated with smaller daughter-volume fractions, while node fragments were associated with larger daughter-volume fractions. The corresponding 95% pointwise bootstrap confidence bands supported these qualitative patterns, while indicating greater uncertainty for the smaller MB-mode subset. Since the experimental dataset does not provide full time-resolved particle-size distributions, the real-data application was interpreted as the estimation of an effective marginal daughter-volume-fraction distribution rather than the identification of a complete breakup-rate function. Accordingly, this experimental recovery should be understood as a single-condition reconstruction, not as a predictive model for different operating conditions.

The present work opens several perspectives. Future developments may include the recovery of breakup rates from time-resolved experimental distributions, extension to multi-fragment breakup models, learning conditional multi-factor fragmentation laws from datasets collected across several operating conditions, and applications to other fragmentation systems. Overall, the results suggest that physics-constrained neural representations provide a promising tool for inverse fragmentation problems, especially when combined with physically meaningful data and sufficiently rich observations.

Author Contributions: Conceptualization, J.E.M. and A.A.; methodology, J.E.M.; software, J.E.M. and C.A.; validation, A.A.; formal analysis, J.E.M. and A.A.; investigation, C.A. and A.A.; resources, J.E.M.; data curation, J.E.M.; writing-original draft preparation, C.A. and A.A.; writing-review and editing, J.E.M.; visualization, A.A. and C.A.; supervision, J.E.M.; project administration, J.E.M. and A.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The experimental droplet-breakup measurements used in this study are derived from the supplementary material associated with [14].

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Bertoin, J. *Random Fragmentation and Coagulation Processes*; Cambridge Studies in Advanced Mathematics; Cambridge University Press: Cambridge, UK, 2006; Volume 102. <https://doi.org/10.1017/CBO9780511617768>.
2. Banasiak, J.; Arlotti, L. *Perturbations of Positive Semigroups with Applications*; Springer Monographs in Mathematics; Springer: London, UK, 2006. <https://doi.org/10.1007/1-84628-153-9>.
3. Villiermaux, E. Fragmentation. *Annu. Rev. Fluid Mech.* **2007**, *39*, 419–446. <https://doi.org/10.1146/annurev.fluid.39.050905.110214>.
4. Doumic, M.; Escobedo, M.; Tournus, M. Estimating the Division Rate and Kernel in the Fragmentation Equation. *Ann. de l'Institut Henri Poincaré. C Anal. Non Linéaire* **2018**, *35*, 1847–1884. <https://doi.org/10.1016/j.anihpc.2018.03.004>.
5. Doumic, M.; Escobedo, M.; Tournus, M. An Inverse Problem: Recovering the Fragmentation Kernel from the Short-Time Behaviour of the Fragmentation Equation. *Ann. Henri Lebesgue* **2024**, *7*, 621–671. <https://doi.org/10.5802/ahl.207>.
6. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-Informed Neural Networks: A Deep Learning Framework for Solving Forward and Inverse Problems Involving Nonlinear Partial Differential Equations. *J. Comput. Phys.* **2019**, *378*, 686–707. <https://doi.org/10.1016/j.jcp.2018.10.045>.
7. Raissi, M.; Karniadakis, G.E. Hidden physics models: Machine learning of nonlinear partial differential equations. *J. Comput. Phys.* **2018**, *357*, 125–141. <https://doi.org/10.1016/j.jcp.2017.11.039>.
8. Karniadakis, G.E.; Kevrekidis, I.G.; Lu, L.; Perdikaris, P.; Wang, S.; Yang, L. Physics-informed machine learning. *Nat. Rev. Phys.* **2021**, *3*, 422–440. <https://doi.org/10.1038/s42254-021-00314-5>.

9. Arridge, S.; Maass, P.; Öktem, O.; Schönlieb, C.B. Solving inverse problems using data-driven models. *Acta Numer.* **2019**, *28*, 1–174. <https://doi.org/10.1017/S0962492919000059>.
10. Cai, S.; Mao, Z.; Wang, Z.; Yin, M.; Karniadakis, G.E. Physics-informed neural networks (PINNs) for fluid mechanics: A review. *Acta Mech. Sin.* **2021**, *37*, 1727–1738. <https://doi.org/10.1007/s10409-021-01148-1>.
11. Haghighat, E.; Raissi, M.; Moure, A.; Gomez, H.; Juanes, R. A physics-informed deep learning framework for inversion and surrogate modeling in solid mechanics. *Comput. Methods Appl. Mech. Eng.* **2021**, *379*, 113741. <https://doi.org/10.1016/j.cma.2021.113741>.
12. Tartakovsky, A.M.; Marrero, C.O.; Perdikaris, P.; Tartakovsky, G.D.; Barajas-Solano, D. Physics-informed deep neural networks for learning parameters and constitutive relationships in subsurface flow problems. *Water Resour. Res.* **2020**, *56*, e2019WR026731. <https://doi.org/10.1029/2019WR026731>.
13. Lu, L.; Jin, P.; Pang, G.; Zhang, Z.; Karniadakis, G.E. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nat. Mach. Intell.* **2021**, *3*, 218–229. <https://doi.org/10.1038/s42256-021-00302-5>.
14. Jackiw, I.M.; Ashgriz, N. Prediction of the Droplet Size Distribution in Aerodynamic Droplet Breakup. *J. Fluid Mech.* **2022**, *940*, A17. <https://doi.org/10.1017/jfm.2022.249>.
15. Kumar, S.; Ramkrishna, D. On the solution of population balance equations by discretization–I. A fixed pivot technique. *Chem. Eng. Sci.* **1996**, *51*, 1311–1332. [https://doi.org/10.1016/0009-2509\(96\)88489-2](https://doi.org/10.1016/0009-2509(96)88489-2).
16. Cybenko, G. Approximation by Superpositions of a Sigmoidal Function. *Math. Control. Signals Syst.* **1989**, *2*, 303–314. <https://doi.org/10.1007/BF02551274>.
17. Hornik, K.; Stinchcombe, M.; White, H. Multilayer Feedforward Networks Are Universal Approximators. *Neural Netw.* **1989**, *2*, 359–366. [https://doi.org/10.1016/0893-6080\(89\)90020-8](https://doi.org/10.1016/0893-6080(89)90020-8).
18. Byrd, R.H.; Lu, P.; Nocedal, J.; Zhu, C. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM J. Sci. Comput.* **1995**, *16*, 1190–1208. <https://doi.org/10.1137/0916069>.
19. Jackiw, I.M.; Ashgriz, N. On Aerodynamic Droplet Breakup. *J. Fluid Mech.* **2021**, *913*, A33. <https://doi.org/10.1017/jfm.2021.7>.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.