

A quasi-real-time analysis of CPP fs-CARS spectra using Convolutional Neural Networks for temperature measurement

Célia Bonnefoy*, Benoît Barviau, Alexandre Poux, Frédéric Grisch

INSA Rouen Normandie, Univ Rouen Normandie, CNRS, Normandie Univ, CORIA UMR 6614, F-76000 Rouen, France

*Corresponding author: celia.bonnefoy@coria.fr

Keywords: Neural Networks, Chirped Probe Pulse Femtosecond Coherent anti-Stokes Raman scattering, Thermometry.

ABSTRACT

High-repetition-rate nonlinear optical diagnostics for combustion require high-speed data processing methods capable of handling large experimental datasets. A deep learning approach based on a one-dimensional convolutional neural network (1D-CNN) is proposed for the analysis of chirped-probe-pulse femtosecond coherent anti-Stokes Raman scattering (CPP fs-CARS) spectra. The objective is to directly infer temperature and relevant physical parameters from single-shot spectra without relying on iterative inversion methods. The model is trained and benchmarked in a supervised framework using a theoretical database of two million spectra generated from a detailed CARS physical model, including molecular rovibrational physics and laser pulse characteristics. A spectral truncation of the CARS theoretical spectra is performed to reproduce realistic experimental spectral windows, and preprocessing is applied to homogenize simulated and experimental data. After training, the CNN is assessed on unseen theoretical spectra and on experimental measurements recorded in air and in a laminar CH_4 /air flame at atmospheric pressure. Temperature retrieval demonstrates high accuracy in simulations and robustness under partial spectral information, with experimental temperatures remaining consistent with expected values. Reconstructed spectra from predicted parameters agree well with measurements, confirming physical consistency. Once trained, the network processes thousands of spectra within seconds, providing several orders of magnitude speed-up compared to genetic algorithms. This enables near real-time analysis of CPP fs-CARS data and opens perspectives for diagnostics in turbulent reactive flows.

1. Introduction

In the context of the climate emergency, optimizing energy production systems and combustion processes is a major challenge for limiting pollutant emissions while improving energy efficiency. The emergence of new combustion chamber designs, combined with the growing use of low-carbon alternative fuels, has increased the complexity of combustion phenomena. Such complexity is driven by the strong interactions between chemical kinetics, species transport, and flame

structure, requiring a detailed understanding of the underlying physicochemical mechanisms, especially at the short time scales characteristic of turbulent reactive flows. In this context, the development of non-intrusive optical diagnostics at high repetition rates is essential to obtain reliable, time-resolved quantitative measurements (Ehn et al., 2017; Bhattacharya, 2024). Among the available techniques, Coherent Anti-Stokes Raman Scattering (CARS) is a third-order nonlinear optical diagnostic capable of providing accurate measurements of temperature and concentration of major species in complex environments, such as high-temperature reactive media (Eckbreth, 2022). Temperature, as a fundamental scalar quantity, plays a central role in controlling combustion processes, directly influencing reaction kinetics, pollutant formation, and the overall efficiency of systems. Thus, the use of advanced measurement diagnostics such as CARS opens up significant opportunities for the analysis and optimization of next-generation combustion systems.

Initially developed in the nanosecond pulse regime, CARS was validated and subsequently widely used for combustion (Roy et al., 2010) and complex media (El-Diasty, 2011) studies. However, the emergence of high-energy, high-repetition-rate femtosecond laser sources has opened up new possibilities for the analysis of transient phenomena (Lucht et al., 2007). In this regime, several configurations have been developed to access temporal information. Time-resolved approaches (Lang et al., 1999; Gord et al., 2008) allow for the reconstruction of molecular coherence dynamics, but remain limited in turbulent or non-stationary conditions due to long acquisition times (Lang & Motzkus, 2002). Conversely, configurations such as fs/ps CARS (Prince et al., 2006; Miller et al., 2010) or CPP fs-CARS (Lang & Motzkus, 2002; Roy et al., 2009) allow for instantaneous measurements, which are better suited for studying turbulent flows. The CPP fs-CARS (Chirped Probe Pulse) diagnostic relies on the use of a temporally stretched and frequency-chirped probe pulse to probe the molecular coherence generated by two ~ 100 femtosecond laser pulses. This approach enables single-shot acquisition of high-repetition-rate spectra, paving the way for single-shot temperature measurements in complex environments (D. R. Richardson et al., 2010; Song et al., 2022). In combustion applications, these spectra are typically obtained by probing the rovibrational transitions of N_2 , which is used as a thermometric tracer (Legros et al., 2021).

However, quantitative exploitation of CPP fs-CARS spectra relies on the comparison between experimental measurements and theoretical spectra computed from detailed physical models that incorporate both molecular spectroscopic properties and the characteristics of the incident laser pulses (D. Richardson et al., n.d.). Conventional approaches are based on iterative optimization techniques, such as genetic algorithms, to adjust a large number of parameters (D. R. Richardson et al., 2011). Although robust, these methods are computationally expensive, which becomes a major limitation for high-repetition-rate diagnostics producing large datasets (D. R. Richardson et al., 2010, 2017). In this context, machine learning methods emerge as a promising alternative to accelerate experimental data processing. In particular, neural networks provide the flexibility required to model complex nonlinear relationships and are well-suited for the analysis of large datasets (Yann Lecun, 2023; Liu et al., 2024; Gariso et al., 2025). They are therefore particularly

effective for the interpretation of signals arising from nonlinear physical processes. Although their use has already been explored for CARS spectra in the nanosecond regime (van der Steen & Black, 1997; Fantoni et al., 2000; Junjuri & Bocklitz, 2025), their application to CPP fs-CARS spectra has, to the best of our knowledge, not yet been investigated.

The objective of this study is to develop a deep learning-based post-processing strategy enabling quasi-real-time temperature measurements from high-repetition-rate CPP fs-CARS spectra. This approach aims to significantly reduce computational cost while maintaining a level of accuracy compatible with the requirements of diagnostics in reactive flow environments.

2. Physical principle of CPP fs-CARS and experimental setup

The CARS diagnostic is a pump–probe optical technique based on a third-order nonlinear interaction between laser pulses and molecules of the probed medium. Two excitation pulses, referred to as pump and Stokes, interact simultaneously with the medium. When the frequency difference between both laser beams matches a rovibrational transition of the target molecule, a macroscopic molecular coherence is induced. This coherence is then probed by a third laser pulse, referred to as the probe. The interaction between this pulse and the coherent molecular polarization leads to the generation of a higher-frequency signal, known as the anti-Stokes signal. This coherent signal carries information about the thermodynamic state of the medium, in particular temperature, through the distribution of molecular energy levels (Begley et al., 1974; Hancock et al., 1992).

In the CPP fs-CARS configuration, the temporal dynamics of the coherence, occurring on the femtosecond timescale, are mapped into the frequency domain. This approach relies on the use of a picosecond probe pulse whose spectral components are temporally stretched. The frequency spreading over time implies a continuous shift of the instantaneous frequency throughout the duration of the pulse, thereby encoding its temporal evolution into the anti-Stokes spectrum (D. R. Richardson et al., 2011). To minimize non-resonant contributions, the probe pulse is temporally delayed with respect to the pump and Stokes pulses, reducing temporal overlap (D. Richardson et al., n.d.).

The experimental implementation is based on a femtosecond laser source delivering ~ 100 fs pulses at 800 nm at a 1 kHz repetition rate (Legros et al., 2021). The Stokes pulse is centered at 800 nm, while an optical parametric amplifier generates 675 nm pulses used for both the pump and probe beams. The probe pulse is temporally stretched through propagation in a dispersive element, yielding a duration of ~ 3 ps. The pump and Stokes pulses remain close to 100 fs and are tuned to excite the rovibrational transitions of N_2 . The three beams are synchronized with adjustable delay lines and spatially arranged in a folded BOXCARS geometry to satisfy phase-matching conditions (Shirley et al., 1980). The generated CARS signal is therefore spatially separated from the excitation beams, and recorded on a spectrograph coupled to a 512-pixel EMCCD camera, and recorded for post-processing. Each spectrum provides information on the rovibrational population distribution

and, therefore, on the local temperature of the medium.

3. Modeling and generation of CARS spectra

Experimental CARS spectra do not provide direct access to the temperature of the probed medium. Estimating this scalar quantity, therefore, requires solving an inverse problem based on the comparison between experimental spectra and theoretical spectra. In this context, a physical model of the CARS signal has been implemented to accurately reproduce spectra by accounting for both molecular spectroscopic properties and the characteristics of the incident laser pulses. This model is based on the CARS spectral intensity given in Eq. 1 (D. R. Richardson et al., 2011).

$$I_{\text{CARS}}(\omega) = \left| \mathcal{F} \left\{ E_{pr}(t) \left[\alpha E_p(t) E_s(t) + \beta \left(\int_{-\infty}^t E_p(t') E_s(t') dt' \right) \right. \right. \right. \\ \left. \left. \left. \times \sum_i \Delta N_i \left(\frac{d\sigma}{d\Omega} \right)_i \cos(\omega_i t - \phi) e^{-\Gamma_i t} \right] \right\} \right|^2 \quad (1)$$

This formulation highlights two main contributions. The non-resonant contribution depends only on the electric fields of the pump $E_p(t)$, Stokes $E_s(t)$, and probe $E_{pr}(t)$ pulses. The resonant contribution is associated with rovibrational transitions of nitrogen and carries the thermodynamic information of the medium.

The resonant contribution, described in Eq. 1, further depends on molecular spectroscopic properties. For each transition i , ΔN_i denotes the population difference between rovibrational levels, $\left(\frac{d\sigma}{d\Omega} \right)_i$ the Raman cross section, ω_i the transition frequency, and Γ_i the linewidth broadening coefficient. The ratio β/α weights the resonant and non-resonant contributions, while ϕ represents their relative phase shift.

A detailed description of the temporal electric fields is a key element of the model. While pulse spectra are experimentally measured, these pulses are not Fourier-transform limited, so knowledge of the spectral amplitude alone is insufficient to reconstruct their temporal evolution. The spectral phase must therefore be introduced to retrieve temporal profiles via inverse Fourier transform. Within the CPP fs-CARS framework, the spectral phase is modeled using a fourth-order Taylor expansion around the angular frequency ω_0 (Eq. 2). The higher-order terms are required due to the spectral amplitude profiles, which exhibit asymmetric tails and non-Gaussian shapes. The coefficients β_2 , β_3 , and β_4 describe dispersion effects for each of the three laser pulses, while the delay parameters t_0 control the relative temporal positions of the pump/Stokes and pump/probe pulses. These parameters directly influence the simulated CARS signal. In practice, experimental drifts such as thermal variations, laser instabilities, or optical heating can induce fluctuations in these quantities, which must be accounted for in the model.

$$E(\omega) = A(\omega) \exp[-it_0(\omega - \omega_0)] \exp\left[-i\left(\frac{\beta_2}{2}(\omega - \omega_0)^2 + \frac{\beta_3}{6}(\omega - \omega_0)^3 + \frac{\beta_4}{24}(\omega - \omega_0)^4\right)\right] \quad (2)$$

Overall, CARS spectrum simulation relies on a set of 14 parameters, including spectroscopic quantities, gas temperature, temporal and spectral characteristics of the laser pulses. Several of these parameters are not directly accessible experimentally and must be adjusted numerically. Moreover, small variations in operating conditions can lead to significant changes in spectral shape, requiring accurate parameter characterization to ensure reliable simulations.

Temperature estimation is performed via a post-processing step based on least-squares minimization of the discrepancy between experimental and simulated spectra. This inverse problem is typically solved using multiparameter optimization schemes such as genetic algorithms (Storn & Price, 1997). Although these methods are robust and widely validated (Legros et al., 2021), their computational cost remains prohibitive (D. R. Richardson et al., 2010, 2017), particularly when applied to large datasets comprising several thousand single-shot spectra. In such cases, the analysis requires from several days up to a week of CPU time, which significantly limits their use for time-constrained or high-throughput applications. This limitation is partly alleviated by a two-step strategy, in which a first calibration is performed on representative spectra recorded under well-characterized conditions. The resulting laser parameters are then kept fixed for the remainder of the dataset, and only a reduced set of parameters is retrieved, under the assumption of phase stability of the laser pulses during acquisition (D. R. Richardson et al., 2011, 2017).

These limitations motivate the development of alternative strategies aimed at significantly accelerating temperature retrieval while preserving the accuracy of the physical model. In this context, machine learning approaches, and in particular neural networks, emerge as a promising solution, as discussed in the following section.

4. Deep learning methodology for CPP fs-CARS spectral analysis

The method developed in this study is based on the use of a CNN, which is well known for its ability to extract local patterns in data (Lopez Pinaya et al., 2020; Bhatt et al., 2021). The objective is to establish a direct relationship between CARS spectra obtained from CPP fs-CARS and the physical parameters of the probed medium, in particular temperature. The goal is to bypass conventional iterative inversion methods by providing a fast, direct estimation of the quantities of interest from spectral information, while avoiding the need to fix parameters or introduce simplifying assumptions.

The model is trained in a supervised learning framework using a dataset of theoretical spectra for which the associated physical parameters are known. These data are propagated through a network composed of successive layers applying nonlinear transformations, enabling progressive extraction of relevant spectral features. The network parameters (weights and biases) are optimized

by minimizing a loss function that quantifies the discrepancy between predicted and reference values, using gradients backpropagation (Mijwel et al., 2023).

The overall processing strategy is structured around three main steps: the design of a neural network architecture adapted to spectral data, the generation of a large database of theoretical CARS spectra reproducing realistic experimental conditions, and the training of the model, followed by its validation on both simulated and experimental datasets.

4.1. Network architecture

The neural network architecture is a one-dimensional convolutional neural network (1D-CNN) (Ige & Sibiya, 2024) tailored for spectral signal processing. The model input consists of CARS spectra represented as fixed-size 1D vectors corresponding to the number of spectral points.




Figure 1. Convolutional neural network architecture for physical parameter estimation. Illustration of the processing from the input signal to the prediction of the parameters θ representing the temperature (θ_1) and the associated physical parameters (θ_2 to θ_{16}).

As illustrated in Fig. 1, the network is built as a succession of convolutional blocks designed to perform hierarchical feature extraction from the spectral data. Each block progressively transforms the raw input signal into increasingly discriminative representations with respect to the target physical parameters.

Each block begins with a 1D convolution layer, whose role is to determine local patterns within the spectrum. The convolutional kernels act as adaptive filters capable of identifying localized spectral features. The convolution output is then normalized using instance normalization, which stabilizes the distribution of activations for each sample by reducing scale variations across spectra (Ulyanov et al., 2016). A SiLU (Sigmoid Linear Unit) activation function is subsequently applied to introduce the non-linearity required to model the complex relationship between spectral signatures and physical parameters. Compared to more classical activation functions, SiLU provides smooth differentiability and improves gradient propagation in deep architectures, thereby enhancing training convergence (Elfwing et al., 2018). Each block is completed by a max-pooling

operation, which reduces the spectral resolution by selecting the maximum value within fixed, non-overlapping spectral bins, while preserving the dominant features (Zafar et al., 2022). This hierarchical reduction is essential to connect fine local structures to global spectral trends. Within this architecture, the number of filters increases with depth, enabling a progression from simple feature extraction (local intensity variations) to more abstract representations combining multiple spectral structures. This gradual increase in representational capacity is essential for capturing the complexity of CARS signals.

The feature maps produced by the convolutional blocks are then flattened and passed through fully connected layers. These dense layers perform nonlinear combinations of extracted features to build a global representation of the spectrum. The final layer performs regression and simultaneously outputs the estimated temperature as well as the associated physical model parameters.

Overall, the architecture contains approximately 12 million trainable parameters, providing high representational capacity while remaining suitable for supervised training on large theoretical datasets.

4.2. Training dataset

The development of a high-performing neural network requires a sufficiently large dataset that is representative of the variability of experimental conditions. In CPP fs-CARS, experimental spectra are not used for training due to their limited availability. Moreover, the physical parameters associated with measured spectra cannot be determined with sufficient accuracy, as their estimation relies on inversion procedures that remain inherently uncertain. In this context, a theoretical database of two million CARS spectra has been generated using a dedicated numerical solver. Each of the 14 physical parameters is uniformly sampled over ranges consistent with experimental conditions, as summarized in Table 1.

Each spectrum is represented as a one-dimensional vector describing the intensity over the spectral range $[16793, 18501]\text{cm}^{-1}$. The number of spectral points was set to 2048 to ensure an optimal balance between spectral resolution, memory, and computation time. This database was generated using a computer with a parallel architecture comprising 56 processors, with a computation time of approximately ~ 2 weeks.

To ensure the stability of the learning process and the network's numerical performance, the CARS spectra are preprocessed before training (Pyle, 1999). Since the raw intensities span several orders of magnitude, a multi-step normalization procedure is implemented to reduce scale effects and ensure a balanced contribution from the full spectral content (Anggoro, 2019; Laska & Yolanda, 2024). First, the spectra are normalized using min-max scaling to bring them into a bounded interval (Eq. 3) (Han et al., 2012). A logarithmic transformation is then applied to reduce the dynamic range between high and low-intensity regions, thereby preventing high-intensity features

Name	Description	Range	Name	Description	Range
T	Temperature	[275 ; 2500] K	β_2^s	2nd-order spectral phase of the Stokes pulse	[1200 ; 2800] fs ²
$\frac{\beta}{\alpha}$	Amplitude ratio between the resonant and non-resonant contributions	[1 ; 20] $\times 10^{-13}$	β_3^s	3-rd-order spectral phase of the Stokes pulse	[-50000 ; 50000] fs ³
ϕ	Phase shift between the resonant and non-resonant contributions	[0 ; 2π]	β_4^s	4th-order spectral phase of the Stokes pulse	$[-10^6 ; 10^6]$ fs ⁴
β_2^p	2nd-order spectral phase of the Pump pulse	[1000 ; 2001] fs ²	t_0^{pr}	Time Delay between pump and probe pulses	[-2 ; 3] ps
β_3^p	3-rd-order spectral phase of the Pump pulse	[-50000 ; 50000] fs ³	β_2^{pr}	2-nd-order spectral phase of the Probe pulse	[72900 ; 75900] fs ²
β_4^p	4th-order spectral phase of the Pump pulse	$[-10^6 ; 10^6]$ fs ⁴	β_3^{pr}	3-rd-order spectral phase of the Probe pulse	[-75000 ; 75000] fs ³
t_0^s	Time Delay between pump and Stokes pulses	[-15 ; 15] fs	β_4^{pr}	4th-order spectral phase of the Probe pulse	$[-10^6 ; 10^6]$ fs ⁴

Table 1. Realistic simulation parameter ranges, representative of the experimental conditions, were used to generate a diverse database enabling the neural network to effectively perform its training.

from dominating and potentially masking weaker spectral information (Eq. 4) (Attar et al., 2018; Yang et al., 2022). Finally, z-score normalization is performed to center the data and reduce its variance across the entire training dataset (Eq. 5) (Jayalakshmi & Santhakumaran, 2011).

$$X_1 = \frac{X - \min(X)}{\max(X) - \min(X)} \quad (3)$$

$$X_2 = \log(\alpha + X_1(1 - \alpha)) \quad (4)$$

$$X_3 = \frac{X_2 - \mu(X_2)}{\sigma(X_2)} \quad (5)$$

α is set to 10^{-2} to obtain a balanced contribution in terms of spectral amplitude for each part of the CARS spectrum, X denotes a given CARS spectrum, X_3 the resulting processed spectrum, and μ and σ correspond to the mean and standard deviation computed over the full CARS library.

Particular attention is paid to the compatibility between simulated and experimental data to ensure consistent learning by the CNN. The goal is to ensure a common representation of the spectra, despite differences in resolution and acquisition conditions between simulation and experiment. In practice, each theoretical spectrum is represented as a vector of 2048 points corresponding to the preprocessed CARS intensity. In contrast, experimental measurements are acquired using a 512-pixel emCCD camera, with this resolution dictated by the characteristics of the spectrograph and

the detection system. To standardize the network inputs, a common spectral grid is defined over the interval $[16793\text{--}18501]\text{ cm}^{-1}$. The experimental spectra are first interpolated onto this grid, then padded with zeros to reach 2048 points. Conversely, the theoretical spectra are restricted to the same spectral window and also padded with zeros, thus ensuring an identical input dimension for all samples. It is important to note that during the training phase, the truncation applied to the theoretical spectra is not fixed but variable. Each spectrum is thus associated with a different spectral window, rather than a common truncation. This design encourages the network to learn representations that are robust to spectral position within the input space, rather than relying on absolute alignment. Finally, this strategy ensures a consistent and homogeneous input representation for the CNN while preserving the physically accessible spectral content. Examples of this truncation procedure are shown in Fig. 2, where the preprocessing leads to an apparent pseudo-logarithmic intensity distribution.




Figure 2. Example of truncated, normalized theoretical spectra. The dashed line corresponds to the intensity-weighted spectral mass center.

Finally, the output parameters are also normalized to avoid numerical imbalances during training, due to their different orders of magnitude. Thus, excepted for the phase shift parameter, all variables are normalized using a min-max transformation to bring them into the interval $[0,1]$ (Han et al., 2012). This normalization prevents a variable with a large amplitude from dominating the learning process. The phase shift parameter, meanwhile, is represented in trigonometric form, via its sinusoidal and cosinusoidal components, to ensure its continuity in parameter space. This representation naturally introduces two variables, and significantly improves the stability of the learning process. Furthermore, an additional parameter, denoted c_i and corresponding to the normalized, intensity-weighted spectral centroid, is introduced to enrich the physical information provided to the network (Favre et al., 2025). This parameter, shown as vertical dashed lines in Fig. 2, is calculated before spectral normalization and incorporated as an additional target variable.

Ultimately, the preprocessed spectra serve as input to the CNN, whose output consists of an extended set of 16 normalized physical parameters, including those reported in Table 1, with the phase contributions separated into two terms and supplemented by additional c_i coefficients.

4.3. Neural network learning process

The CNN is trained in a supervised learning framework on a database of theoretical CARS spectra paired with their corresponding reference physical parameters. The objective is to optimize the network weights to minimize the discrepancy between the predicted and reference values of the physical parameters, thereby enabling the learning of a direct approximation of the relationship between CARS spectra and the associated physical quantities, without resorting to iterative optimization.

To this end, the database is divided into two distinct sets: 80% of the spectra (1,600,000) are used for model training, while 20% (400,000) are reserved for validation to assess the model's generalization ability.

Network optimization is performed using the Adam optimization algorithm, chosen for its robustness and efficiency in high-dimensional problems (Kingma & Ba, 2014). Training is performed over 125 epochs, which represents a trade-off between computational cost and model convergence. The cost function used corresponds to a weighted mean squared error applied to all output parameters, with a higher weight assigned to temperature, which represents the primary physical quantity of interest.

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \sum_{j=0}^{16} (\omega_j (\theta_{i,j} - t_{i,j}))^2 \quad (6)$$

In Eq. 6, $\theta_{i,j}$ and $t_{i,j}$ are the predicted and target values for the j^{th} parameter of the i^{th} spectrum, N is the total number of spectra, and the weight ω_j is set to 3 for temperature and 1 for all other parameters.

The learning process is gradually stabilized by adjusting the learning rate and batch size (Goodfellow et al., 2016; Masters & Luschi, 2018). The learning rate controls the magnitude of the network weight updates during backpropagation. A high value promotes exploration of the whole domain parameter range, while a lower value allows for gradual convergence towards a stable minimum. This one is initially set to 10^{-4} and then reduced by a factor of 10 at epochs 70 and 100 to improve convergence accuracy. Complementarily, the batch size, corresponding to the number of samples processed together to compute a single gradient update, directly impacts the trade-off between numerical stability and generalization ability. A larger batch size leads to a more stable gradient estimate, while a smaller batch size introduces stochastic noise that improves generalization. In this study, the batch size is manually reduced from 32 to 16 and then to 8 at epochs 40 and 115. This heuristic choice is made empirically and aims at progressively stabilizing the training process and refining the weight updates. These step-wise reductions are shown in Fig. 3 as a function of epoch number on the secondary y-axis.

These combined strategies enable effective convergence of the model while preserving the risk of

convergence toward local minima. The evolution of training and validation losses, depicted with red and blue curves, respectively (Fig. 3), shows a simultaneous and steady decrease, indicating the network's strong generalization ability and the absence of overfitting. This behavior confirms that the network retains its ability to effectively process unknown spectra.




Figure 3. Evolution of the error as a function of the number of epochs. The red curve shows the evolution of the error on the spectra selected for the training phase, while the blue curve shows the evolution of the error on the synthetic spectra reserved for the testing phase. The errors are quantified on the left axis. On the right axis, the evolution of the learning rate is shown in black, and the batch size in green. Their evolution is represented by dotted lines in the same colors.

The model was implemented in PyTorch 2.5.1 and trained on the Austral supercomputer (CRI-ANN, Normandy, France), equipped with 8 NVIDIA A100 80 GB GPUs, with a training time of approximately 8 hours until satisfactory convergence was achieved. Once trained, the model is capable of simultaneously predicting all physical parameters, particularly temperature, from a single CARS spectrum.

The strategy relies on direct propagation through the network, without any iterative optimization procedure, resulting in a computational time that is negligible compared to conventional methods. This feature makes the approach particularly well-suited for the fast, or even real-time, processing of large amounts of experimental data, as shown in the following section.

5. Results and Validation

The CNN's performance is assessed using three datasets to assess its generalization ability under conditions of increasing complexity. The first dataset consists of theoretical CARS spectra from the test database that were not used during training, thereby ensuring evaluation on completely unseen data. The corresponding error estimate is associated with the last blue dot in Fig. 3. The second set consists of theoretical spectra generated at temperatures of 300 K, 1000 K, and 2000 K, with all other parameters randomly sampled within their respective ranges to validate the model's ability to reproduce known thermodynamic conditions. Finally, a set of 1,000 experimental spectra

is used to quantify performance under real measurement conditions.

The network outputs parameters are then denormalized to convert the predicted values back to their original physical scale by applying the inverse of the normalization transformations used during the training phase. The spectra reconstructed from the predicted parameters are compared to the reference spectra following the preprocessing and normalization procedure described in Section 4.2. The estimated temperatures are analyzed to quantify the model's overall accuracy.

5.1. Results on unseen synthetic test set

The first validation step consists in testing the model on unseen theoretical N_2 CARS spectra from the test dataset, ensuring a fully unbiased assessment on data not used during training. In this configuration, the same truncation procedure as used during training is applied, allowing the network to learn spectral patterns independently of their absolute position. However, the retained spectral window varies from one sample to another, leading to an inhomogeneous amount of information across spectra. The performance obtained for all parameters is reported in Table 2.

Parameter	T	ϕ	β/α	β_2^p	β_3^p	β_4^p	t_0^{ss}	β_2^s
$R^2(\%)$	99.67	61.75	96.54	63.13	93.12	89.95	66.77	64.55

Parameter	β_3^s	β_4^s	t_0^{pr}	β_2^{probe}	β_3^{pr}	β_4^{pr}	c_i
$R^2(\%)$	68.28	17.10	99.77	58.55	89.36	68.88	98.15

Table 2. R^2 values for each of the 15 parameters predicted by the neural network, obtained from the test dataset after training on truncated synthetic spectra. The predicted phase parameters have been grouped together here.

Despite the limitations induced by truncation, the network achieves excellent overall performance, with coefficients of determination exceeding 99% for the dominant parameters, in particular temperature and the probe delay $t_{0,pr}$. This high accuracy in temperature is partly promoted by the higher weight assigned to this parameter in the loss function (weight of 3 in Eq. 6), thereby strengthening the learning of this key variable. A decrease in performance is nevertheless observed for certain parameters. This degradation is mainly attributed to the reduced amount of spectral information available after truncation. Furthermore, parameters with a limited effect on the CARS spectral shape, such as β_4^s , are more difficult to estimate, whereas those governing the overall structure of the signal remain predicted with high accuracy. Overall, these results demonstrate the ability of the network to robustly extract relevant physical information, even when the spectral data is partially incomplete.

5.2. Parametric robustness study

To further assess robustness under controlled thermodynamic conditions, three datasets of 1,000 theoretical spectra are generated for fixed temperatures of 300 K, 1000 K, and 2000 K, while all other parameters are randomly sampled within their respective ranges.

In this configuration, and in order to faithfully reproduce the experimental constraints, the truncation applied to the simulated spectra is adjusted so that the retained spectral content closely matches that observed experimentally. To validate the model's ability to recover precise thermodynamic conditions, this truncation is made less restrictive than that used during training. It ensures that the synthetic spectra preserve all relevant physical information without loss of essential spectral features. This modified truncation strategy therefore enables the neural network to be tested under realistic conditions, closely approximating the experimental environment.

The temperature distributions predicted by the network are presented in Fig. 4 as probability density functions. The model reproduces the target temperatures with high accuracy, with mean values of 299.0 K, 999.6 K, and 2002.3 K. The relative errors remain below 1%, and the standard deviations are 17 K, 14 K, and 13 K, respectively.




Figure 4. Temperature Pdf issued from CNN analysis of 1,000 theoretical spectra, with meaningful truncation with respect to experimental spectra, calculated at 300K, 1,000K and 2,000K

These results show that the network robustly learns the relationship between spectral structure and temperature and confirm that the information contained in the CARS spectra is sufficient to effectively constrain the temperature estimate. A slight degradation in performance is observed at low temperatures, where spectral modulations are less pronounced. Conversely, at high temperatures, the increase in the number of active transitions enriches the signal and improves the accuracy of the estimate. This behavior is consistent with the physics underlying the generation of the CARS spectrum (Goss et al., 1980; Martinsson et al., 1993).

The quality of the spectrum reconstruction is illustrated in Fig. 5, which compares the target spectra with those reconstructed from the predicted parameters. Good agreement is formed across the full temperature range studied. The reconstruction quality is quantified using the mean squared

error, as commonly done in post-processing approaches based on genetic algorithms. It is computed on the preprocessed spectra and restricted to the truncated spectral window, and remains low (on the order of 10^{-3}). Residuals and direct superposition further confirm the accuracy of the reconstruction.




Figure 5. Comparison of theoretical CARS spectra (black) and simulated spectra (red) for three target temperatures: 300 K, 1000 K, and 2000 K. The bottom part of each subplot shows the normalized residual ϵ between the experimental data and the simulation.

In general, the neural network reconstructs CARS spectra with high accuracy in regions of high spectral intensity, where the information content is strongest at low wavenumbers. The reconstruction quality gradually decreases at higher wavenumbers, where weaker spectral features become more difficult to recover. This effect is partially mitigated by the preprocessing stage, in particular, the logarithmic transformation, which helps the rebalance of the contribution of different intensity levels. However, discrepancies persist for the weakest spectral modulations, which are more sensitive to higher-order phase parameters and to the reduced information content associated with spectral truncation. This behavior is consistent with the quantitative results reported in Table 2, where lower R^2 values are observed for parameters associated with weaker spectral contributions. Therefore, the results are in good agreement with Table 2. In addition, the use of a truncation strategy more representative of experimental spectra provides sufficient spectral content to reduce ambiguities previously observed under random truncation, resulting in higher R^2 values and allowing a great reconstruction accuracy.

5.3. Performance on experimental spectra

After validation using truncated theoretical CARS spectra, the neural network was applied on experimental data recorded under various conditions. Measurements were performed in ambient air (approximately 300 K) as well as in a stoichiometric CH_4 /air premixed laminar flame at

atmospheric pressure. For each condition, 1,000 single-shot CARS spectra were recorded at a frequency of 1 kHz. To ensure consistency with the training phase, each experimental spectrum is first convolved with the instrument response function of the detection system, then subjected to a smoothing step and the same preprocessing as that applied to the synthetic data (section 4.2). This procedure ensures optimal compatibility between the experimental data and the training dataset used to train the network.




Figure 6. a) Temperature probability density function obtained with the CNN from an experimental spectrum recorded in a laminar CH_4/air flame. b) Comparison between the experimental CARS spectrum (black) and the corresponding simulated spectrum (red). The bottom part of each subplot shows the normalized residual ϵ between the experiment and the simulation.

At room temperature, the network predicts a mean temperature of $\sim 326 \text{ K}$ with a standard deviation of 31 K . The corresponding temperature probability density function is shown in Fig. 6, illustrating the distribution of single-shot predictions obtained from experimental spectra recorded at room temperature. Compared to the results obtained on theoretical datasets (Section 5.1), where relative errors remained below 1%, this performance degradation reflects the increased complexity of experimental conditions.

Figure 6 also shows an example of an experimental CARS spectrum together with its reconstruction based on the parameters predicted by the CNN. Despite being trained exclusively on noise-free synthetic spectra, the model successfully reproduces the main spectral features of the measured signal, demonstrating its ability to generalize to real experimental data. The normalized reconstruction error is reported and remains moderate, although residual discrepancies are observed, particularly in the fine spectral structures.

These deviations can be attributed to several experimental sources of uncertainty. First, Poisson noise associated with the detection process imposes a fundamental statistical limit, leading to temperature fluctuations of a few kelvins for signal levels in the range of 30,000–40,000 photoelectrons (Pealat et al., 1985). In addition, shot-to-shot variations in the femtosecond pulse spectral profile contribute to further uncertainty, mainly due to nonlinear propagation effects such as self-phase

modulation, which induces spectral broadening (Gu et al., 2019). By contrast, beam pointing instabilities are considered negligible under the present experimental conditions due to the stability of the probed volume.

Finally, the remaining discrepancies in the reconstructed spectra may also arise from small inaccuracies in the estimation of secondary parameters, such as phase parameters. Although each of these parameters has a limited individual impact on the spectral shape, their combined effect influences the reproduction of fine spectral features, which can amplify residual differences in the reconstruction.

Despite these limitations, the overall performance remains robust. The least-squares reconstruction error computed on the experimental spectra remains low, confirming that the CNN retains strong predictive capability even under noisy and complex experimental conditions. These results demonstrate the robustness of the proposed approach and its applicability to real single-shot CARS measurements.




Figure 7. a) Temperature probability density function obtained with the CNN from an experimental spectrum recorded in a laminar CH_4/air flame. b) Comparison between the experimental CARS spectrum (black) and the corresponding simulated spectrum (red). The bottom part of each subplot shows the normalized residual ϵ between the experiment and the simulation.

Measurements were then taken in a premixed laminar CH_4/air flame, within the combustion products located a few millimeters above the fresh gas cone, in a region characterized by homogeneous temperature and composition fields. At high temperature, the spectra exhibit more pronounced modulations as well as higher relative intensities in the high-wavenumber region ($> 17,050 \text{ cm}^{-1}$), as shown in Fig. 7. This behavior is explained by an increase in the ratio of resonant to non-resonant contributions, leading to a more pronounced signal structure. This spectral richness improves the network's ability to relate the shape of the spectrum to physical parameters. Under these conditions, the model predicts an average temperature of 2192 K with a standard deviation of 66 K. This improvement in performance at high temperature is consistent with the trend observed on theoretical spectra of different temperatures (Section 5.2). More intense modulations

provide more discriminating information, facilitating parameter determination.

Figure 7 also shows a comparison between experimental spectra and spectra reconstructed from the parameters predicted by the CNN. Overall, good agreement is observed across the entire spectral window, confirming the model's ability to reproduce the dominant spectral structures despite the presence of measurement noise and experimental fluctuations. The residual deviations ϵ are mainly localized in low-intensity regions, consistent with the limitations identified previously in the reconstruction of weak spectral features.

For comparison, the performance of conventional post-processing methods based on genetic algorithms is typically reported to yield similar levels of accuracy under comparable experimental conditions, with standard deviations on the order of approximately 2.5% in laminar flames (Thomas et al., 2017). In these approaches, the parameter estimation is formulated as a global optimization problem, where a population of candidate solutions is iteratively evolved through selection, crossover, and mutation operations. Each candidate requires the determination of a physical CARS model, making the procedure computationally intensive and requiring independent optimization for each measured spectrum.

5.4. Computation time and benefits of the CNN approach

A major advantage of the proposed CNN-based framework lies in the drastic reduction of computational cost associated with post-processing. In contrast to genetic algorithms, which rely on solving a high-dimensional nonlinear optimization problem for each spectrum independently, the CNN performs a direct and simultaneous inference of all physical parameters from a single pass.

In conventional genetic algorithm approaches, the complexity of the CARS model, combined with the high dimensionality of the parameter space, leads to computation times of several hours per spectrum (D. R. Richardson et al., 2010, 2017). Although this limitation can be partially mitigated by a two-step strategy relying on the assumption of parameter stability over the experiment, it remains significant (D. R. Richardson et al., 2010, 2017). Despite this simplification, the method remains computationally demanding and severely limits its applicability to large datasets and time-resolved measurements. By comparison, once trained, the CNN provides near-instantaneous predictions of all parameters in a single forward pass, without requiring iterative optimization or assumptions of parameter constancy.

Method	Training	1 Spectrum	1,000 Spectra
CNN	~ 8h (GPU)	~ 0.06s (CPU)	~ 2.6s
Iterative	None	~ 19h (CPU)	Days to week

Table 3. Comparison of typical computation times: Conventional Genetic Algorithm vs. CNN approach.

Table 3 summarizes the computation times associated with both approaches, highlighting the

trade-off between the one-time cost of training the neural network and the prohibitive cost of iterative inversion methods when applied at scale. In particular, the processing of 1,000 single-shot experimental spectra by CNN requires ~ 2.6 seconds, corresponding to a gain of several orders of magnitude compared to iterative optimization strategies. This drastic reduction in computational cost not only enables rapid estimation of mean thermodynamic quantities but also opens the way to time-resolved analysis of reactive flows, including highly dynamic or turbulent conditions.

6. Conclusions

This study presents a deep learning-based strategy for the analysis of CPP fs-CARS spectra, aimed at the simultaneous determination of temperature and a set of physical parameters governing signal formation. A one-dimensional convolutional neural network is trained on a large synthetic database spanning a wide range of thermodynamic conditions and experimental configurations, enabling the learning of a high-dimensional mapping between spectral variations and underlying physical parameters. The assessment of the model is performed through both parameter estimation accuracy and spectral reconstruction quality.

Results obtained from the theoretical spectra demonstrate that the network can accurately retrieve the temperature and reproduce the overall structure of the CARS signal with relative errors typically below 1%. The training framework includes a variable spectral truncation strategy that provides incomplete and spatially varying spectral information to emulate experimental constraints. This forces the network to learn robust representations from partial spectral windows whose position and extent differ across samples.

Validation on experimental data confirms the ability of the model to perform data processing in real measurement conditions. In both room-temperature and high-temperature CH_4 /air flame configurations, the CNN provides physically consistent information of temperature and achieves good reconstruction of measured spectra. This is allowed by the higher weighting assigned to temperature in the loss function, which prioritizes its accuracy during training. Remaining discrepancies are mainly localized in low-intensity spectral regions, which are more sensitive to higher-order phase contributions and experimental noise.

Overall, the results show that CNN does not merely perform statistical interpolation, but instead learns physically meaningful correlations between spectral features and thermodynamic parameters. This combination of accuracy, robustness, and computational efficiency makes the proposed approach a promising tool for advanced CARS-based diagnostics. In particular, the drastic reduction in computational cost is a key advantage: once trained, the network can process thousands of spectra within seconds, compared to several hours or even days required by traditional iterative optimization methods. This significant acceleration enables high-throughput CARS data analysis and opens the way toward near-real-time applications in complex and highly dynamic reactive

environments.

Nevertheless, several areas for development are being considered. First, the training dataset could be enriched by incorporating more realistic effects, including noise and other experimental variations, in order to better reproduce measurement conditions. Second, the incorporation of explicit physical constraints into the network architecture could further improve robustness, particularly under challenging or extreme operating conditions. Finally, performance could be enhanced through fine-tuning on annotated experimental data, which would allow the model to better adapt to real measurement conditions and further improve both accuracy and stability.

Acknowledgements

The authors acknowledge the support of the CRIANN supercomputing center (Normandie, France) for providing computational resources. We also thank S. Legros for the experimental database and A. Abad for valuable discussions and technical assistance.

References

- Anggoro, D. A. (2019, November). Improving Accuracy by applying Z-Score Normalization in Linear Regression and Polynomial Regression Model for Real Estate Data. *International Journal of Emerging Trends in Engineering Research*, 7(11), 549–555. doi: <https://doi.org/10.30534/ijeter/2019/247112019>
- Attar, T., Bhattacharjee, T., Kumar, R., & Dey, N. (2018, November). Parametric Optimization of Logarithmic Transformation using GWO for Enhancement and Denoising of MRI Images. , 1–7. doi: 10.1109/ICRAIE.2018.8710393
- Begley, R. F., Harvey, A. B., & Byer, R. L. (1974, October). Coherent anti-Stokes Raman spectroscopy. *Applied Physics Letters*, 25(7), 387–390. doi: <https://doi.org/10.1063/1.1655519>
- Bhatt, D., Patel, C., Talsania, H., Patel, J., Vaghela, R., Pandya, S., ... Ghayvat, H. (2021, October). CNN Variants for Computer Vision: History, Architecture, Application, Challenges and Future Scope. *Electronics*, 10(20), 2470. doi: <https://doi.org/10.3390/electronics10202470>
- Bhattacharya, S. (2024). ADVANCED OPTICAL DIAGNOSTIC TECHNOLOGY APPLICATION PROGRESS IN COMBUSTION TESTING OF ENERGETIC MATERIALS. *Eurasia Journal of Science and Technology*, 6(2). doi: 10.61784/ejstv6n209
- Eckbreth, A. C. (2022). *Laser Diagnostics for Combustion Temperature and Species*. CRC Press. doi: <https://doi.org/10.1201/9781003077251>

- Ehn, A., Zhu, J., Li, X., & Kiefer, J. (2017, February). Advanced Laser-Based Techniques for Gas-Phase Diagnostics in Combustion and Aerospace Engineering. *Applied Spectroscopy*, 71(3), 341–366. doi: 10.1177/0003702817690161
- El-Diasty, F. (2011, January). Coherent anti-Stokes Raman scattering: Spectroscopy and microscopy. *Vibrational Spectroscopy*, 55(1), 1–37. doi: <https://doi.org/10.1016/j.vibspec.2010.09.008>
- Elfwing, S., Uchibe, E., & Doya, K. (2018, November). Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Networks*, 107, 3–11. doi: <https://doi.org/10.1016/j.neunet.2017.12.012>
- Fantoni, R., Cola, F., Dominicis, D., Giorgi, M., D'Apice, M., Giammartini, S., & der Steen, H. J. L. V. (2000, June). On-line nitrogen CARS thermometry on a 130 kW burner by using a neural network approach. *Journal of Raman Spectroscopy*, 32(6), 697–701. doi: <https://doi.org/10.1063/5.0007821>
- Favre, A., Abad, A., Poux, A., Gosse, L., Berjaoui, A., Morel, V., & Bultel, A. (2025, February). Towards real-time calibration-free LIBS supported by machine learning. *Spectrochimica Acta Part B: Atomic Spectroscopy*, 224, 107082. doi: <https://doi.org/10.1016/j.sab.2024.107082>
- Gariso, R., Coutinho, J. P., Rato, T. J., & Reis, M. S. (2025, April). A comparative analysis of deep learning and chemometric approaches for spectral data modeling. *Analytica Chimica Acta*, 1347, 343766. doi: <https://doi.org/10.1016/j.aca.2025.343766>
- Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep Learning*. MIT Press. (<http://www.deeplearningbook.org>)
- Gord, J. R., Meyer, T. R., & Roy, S. (2008, July). Applications of Ultrafast Lasers for Optical Measurements in Combusting Flows. *Annual Review of Analytical Chemistry*, 1(1), 663–687. doi: <https://doi.org/10.1146/annurev.anchem.1.031207.112957>
- Goss, L. P., Fleming, J. W., & Harvey, A. B. (1980, August). Pure rotational coherent anti-Stokes Raman scattering of simple gases. *Optics Letters*, 5(8), 345. doi: <https://doi.org/10.1364/OL.5.000345>
- Gu, M., Satija, A., & Lucht, R. P. (2019, November). Effects of self-phase modulation (SPM) on femtosecond coherent anti-Stokes Raman scattering spectroscopy. *Optics Express*, 27(23), 33954. doi: <https://doi.org/10.1364/OE.27.033954>
- Han, J., Kamber, M., & Pei, J. (2012). *Data Mining Concepts and Techniques. 3rd Edition*. Morgan Kaufmann Publishers.

- Hancock, R. D., Boyack, K. W., & Hedman, P. O. (1992). *Coherent Anti-Stokes Raman Spectroscopy (CARS) in Pulverized Coal Flames*. Springer US. doi: https://doi.org/10.1007/978-1-4899-3671-4_15
- Ige, A. O., & Sibiya, M. (2024). State-of-the-Art in 1D Convolutional Neural Networks: A Survey. *IEEE Access*, 12, 144082–144105. doi: 10.1109/ACCESS.2024.3433513
- Jayalakshmi, T., & Santhakumaran, A. (2011). Statistical Normalization and Back Propagation for Classification. *International Journal of Computer Theory and Engineering*, 89–93. doi: 10.7763/IJCTE.2011.V3.288
- Junjuri, R., & Bocklitz, T. (2025, June). Review of Coherent Anti-Stokes Raman Scattering Non-resonant Background Removal and Phase Retrieval Approaches: From Experimental Methods to Deep Learning Algorithms. *Advanced Photonics Research*, 6(9). doi: <https://doi.org/10.1002/adpr.202500035>
- Kingma, & Ba. (2014). Adam: A Method for Stochastic Optimization. doi: <https://doi.org/10.48550/arXiv.1412.6980>
- Lang, T., Kompa, K.-L., & Motzkus, M. (1999, August). Femtosecond CARS on H₂. *Chemical Physics Letters*, 310(1–2), 65–72. doi: [https://doi.org/10.1016/S0009-2614\(99\)00787-3](https://doi.org/10.1016/S0009-2614(99)00787-3)
- Lang, T., & Motzkus, M. (2002, February). Single-shot femtosecond coherent anti-Stokes Raman-scattering thermometry. *Journal of the Optical Society of America B*, 19(2), 340. doi: <https://doi.org/10.1364/JOSAB.19.000340>
- Laska, R. R., & Yolanda, A. M. (2024). A Comparative Study of Z-Score and Min-Max Normalization for Rainfall Classification in Pekanbaru. *Journal of Data Science*, 2024(1). doi: 10.61453/jods.v2024no04
- Legros, S., Brunet, C., Domingo-Alvarez, P., Malbois, P., Salaun, E., Godard, G., ... Grisch, F. (2021, February). Combustion for aircraft propulsion: Progress in advanced laser-based diagnostics on high-pressure kerosene/air flames produced with low-NO_x fuel injection systems. *Combustion and Flame*, 224, 273–294. doi: <https://doi.org/10.1016/j.combustflame.2020.12.036>
- Liu, X., An, H., Cai, W., & Shao, X. (2024, March). Deep learning in spectral analysis: Modeling and imaging. *TrAC Trends in Analytical Chemistry*, 172, 117612. doi: <https://doi.org/10.1016/j.trac.2024.117612>
- Lopez Pinaya, W. H., Vieira, S., Garcia-Dias, R., & Mechelli, A. (2020). Convolutional neural networks. , 173–191. doi: <https://doi.org/10.1016/B978-0-12-815739-8.00010-9>

- Lucht, R. P., Kinnius, P. J., Roy, S., & Gord, J. R. (2007). Theory of femtosecond coherent anti-Stokes Raman scattering spectroscopy of gas-phase transitions. *The Journal of Chemical Physics*, 127(4). doi: <https://doi.org/10.1063/1.2751184>
- Martinsson, L., Bengtsson, P.-E., Aldén, M., Kröll, S., & Bonamy, J. (1993, August). A test of different rotational Raman linewidth models: Accuracy of rotational coherent anti-Stokes Raman scattering thermometry in nitrogen from 295 to 1850 K. *The Journal of Chemical Physics*, 99(4), 2466–2477. doi: <https://doi.org/10.1063/1.466197>
- Masters, D., & Luschi, C. (2018). Revisiting Small Batch Training for Deep Neural Networks. doi: <https://doi.org/10.48550/arXiv.1804.07612>
- Mijwel, M. M., Esen, A., & Shamil, A. (2023, August). Overview of Neural Networks. *Babylonian Journal of Machine Learning*, 2023, 42–45. doi: <https://doi.org/10.58496/BJML/2023/008>
- Miller, J. D., Slipchenko, M. N., Meyer, T. R., Stauffer, H. U., & Gord, J. R. (2010, July). Hybrid femtosecond/picosecond coherent anti-Stokes Raman scattering for high-speed gas-phase thermometry. *Optics Letters*, 35(14), 2430. doi: <https://doi.org/10.1364/OL.35.002430>
- Pealat, M., Bouchardy, P., Lefebvre, M., & Taran, J.-P. (1985, April). Precision of multiplex CARS temperature measurements. *Applied Optics*, 24(7), 1012. doi: <https://doi.org/10.1364/AO.24.001012>
- Prince, B. D., Chakraborty, A., Prince, B. M., & Stauffer, H. U. (2006, July). Development of simultaneous frequency- and time-resolved coherent anti-Stokes Raman scattering for ultrafast detection of molecular Raman spectra. *The Journal of Chemical Physics*, 125(4). doi: <https://doi.org/10.1063/1.2219439>
- Pyle, D. (1999). *Data Preparation for Data Mining*. California: Morgan Kaufmann Publishers.
- Richardson, D., Lucht, R., Roy, S., & Gord, J. (n.d.). Single-Pulse Femtosecond Coherent Anti-Stokes Raman Scattering (CARS) Spectroscopy for Temperature Measurements at Data Rates of 1 kHz, year = 2009, month = jan, booktitle = 47th AIAA Aerospace Sciences Meeting including The New Horizons Forum and Aerospace Exposition, doi = <https://doi.org/10.2514/6.2009-527>, publisher = American Institute of Aeronautics and Astronautics,.
- Richardson, D. R., Lucht, R. P., Kulatilaka, W. D., Roy, S., & Gord, J. R. (2010). Theory of Chirped-Probe Pulse Single-Shot Femtosecond Coherent Anti-Stokes Raman Scattering Thermometry in Flames at 1000 Hz. , LTuB5. doi: <https://doi.org/10.1364/LACSEA.2010.LTuB5>
- Richardson, D. R., Lucht, R. P., Kulatilaka, W. D., Roy, S., & Gord, J. R. (2011, April). Theoretical modeling of single-laser-shot, chirped-probe-pulse femtosecond coherent anti-Stokes Raman scattering thermometry. *Applied Physics B*, 104(3), 699–714. doi: 10.1007/s00340-011-4489-0

- Richardson, D. R., Stauffer, H. U., Roy, S., & Gord, J. R. (2017, February). Comparison of chirped-probe-pulse and hybrid femtosecond/picosecond coherent anti-Stokes Raman scattering for combustion thermometry. *Applied Optics*, 56(11), E37. doi: 10.1364/AO.56.000E37
- Roy, S., Gord, J. R., & Patnaik, A. K. (2010, April). Recent advances in coherent anti-Stokes Raman scattering spectroscopy: Fundamental developments and applications in reacting flows. *Progress in Energy and Combustion Science*, 36(2), 280–306. doi: 10.1016/j.pecs.2009.11.001
- Roy, S., Kulatilaka, W. D., Richardson, D. R., Lucht, R. P., & Gord, J. R. (2009, December). Gas-phase single-shot thermometry at 1 kHz using fs-CARS spectroscopy. *Optics Letters*, 34(24), 3857. doi: <https://doi.org/10.1364/OL.34.003857>
- Shirley, J. A., Hall, R. J., & Eckbreth, A. C. (1980, September). Folded BOXCARS for rotational Raman studies. *Optics Letters*, 5(9), 380. doi: <https://doi.org/10.1364/OL.5.000380>
- Song, Y., Wu, H., Zhu, G., Zeng, Y., Yu, G., & Yang, Y. (2022, August). A femtosecond time-resolved coherent anti-Stokes Raman spectroscopy thermometry for steady-state high-temperature flame. *Combustion and Flame*, 242, 112166. doi: <https://doi.org/10.1016/j.combustflame.2022.112166>
- Storn, R., & Price, K. (1997, December). Differential Evolution – A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *Journal of Global Optimization*, 11(4), 341–359. doi: <https://doi.org/10.1023/A:1008202821328>
- Thomas, L. M., Satija, A., & Lucht, R. P. (2017, October). Technique developments and performance analysis of chirped-probe-pulse femtosecond coherent anti-Stokes Raman scattering combustion thermometry. *Applied Optics*, 56(31), 8797. doi: <https://doi.org/10.1364/AO.56.008797>
- Ulyanov, D., Vedaldi, A., & Lempitsky, V. (2016). Instance Normalization: The Missing Ingredient for Fast Stylization. doi: <https://doi.org/10.48550/arXiv.1607.08022>
- van der Steen, H. J. L., & Black, J. D. (1997, December). Temperature analysis of Coherent Anti-Stokes Raman spectra using a neural network approach. *Neural Computing & Applications*, 5(4), 248–257. doi: <https://doi.org/10.1007/BF01424230>
- Yang, Z., Yang, Y., Li, C., Zhou, Y., Zhang, X., Yu, Y., & Liu, D. (2022, June). Tasseled Crop Rows Detection Based on Micro-Region of Interest and Logarithmic Transformation. *Frontiers in Plant Science*, 13. doi: <https://doi.org/10.3389/fpls.2022.916474>
- Yann Lecun, G. H., Yoshua Bengio. (2023, May). Deep learning. *Nature*, 521(7553), 436–444. doi: <https://doi.org/10.1038/nature14539>

Zafar, A., Aamir, M., Mohd Nawi, N., Arshad, A., Riaz, S., Alruban, A., . . . Almotairi, S. (2022, August). A Comparison of Pooling Methods for Convolutional Neural Networks. *Applied Sciences*, 12(17), 8643. doi: <https://doi.org/10.3390/app12178643>