

RESEARCH ARTICLE

MARSFT: Efficient fitting of CARS spectra using a library-based genetic algorithm

M. Greifenstein¹ | A. Dreizler¹

Department of Mechanical Engineering,
Reactive Flows and Diagnostics, Technical
University Darmstadt, Otto-Berndt-Str. 3,
Darmstadt, 64287, Germany

Correspondence

M. Greifenstein, Technical University
Darmstadt, Institute for Reactive Flows
and Diagnostics, Otto-Berndt-Str. 3,
DE-64287 Darmstadt.
Email: greifenstein@rsm.tu-darmstadt.de

Funding information

Deutsche Forschungsgemeinschaft,
Grant/Award Number: DR374/24-1

Abstract

A loss-less compressed library scheme is presented in this publication that allows for computationally efficient fitting of coherent anti-Stokes Raman spectra with no restriction to the number of degrees of freedom for the spectral fit. The compression is achieved by convolving the squared modulus and the real part of the complex susceptibility with a Gaussian kernel narrower than the experimental apparatus function. This effectively reduces library size while allowing to convolve to the final experimental linewidth during the fit. For the optimization procedure, a gradient-free mixed-integer genetic algorithm was implemented due to its ability to extract library spectra without interpolation. We demonstrate the ability of the code by comparing it to CARSFT in terms of dependency on starting solution, computational cost and accuracy using simulated spectra with varying noise contribution.

KEYWORDS

CARS, genetic algorithm, library, spectral fit

1 | INTRODUCTION

Coherent anti-Stokes Raman spectroscopy (CARS) is a widely used non-linear diagnostic tool to measure temperature and species concentrations. Due to its high spatial and temporal resolution, CARS is a very valuable technique especially in turbulent reactive flows.^[1] Typically, CARS signals are evaluated in the spectral domain for which several strategies can be found in the literature. Traditionally, codes such as CARSFT^[2] use a gradient-based non-linear least squares fit to match a physical model to experimental data in order to retrieve the quantities of interest. Spectra are generated during the fit, requiring time consuming calculations for each evaluation of the objective function. Additionally, gradient-based solvers may suffer from susceptibility to local minima; that is, there is a dependence on the provided initial solution. To circumvent these issues, a machine learning approach

using neural networks trained with theoretical spectra was developed by Van der Steen et al.^[3] that even allows on-line fitting at 8-Hz repetition rate. However, as the authors also point out, the neural network gives no indication of the plausibility of the result. Another very common approach to avoid the issue with local minima and to optimize the computational effort for spectral evaluation is the use of spectral libraries generated prior to the actual data evaluation. These codes further split into two categories: (a) fully packed and (b) sparse libraries. One example for the former category is actually the quickfitter function of CARSFT (or—to be more precise—FTCARS). Here, the real part and the square of the resonant susceptibility is tabulated at given temperature increments to quickly generate interpolated spectra during the fit process.^[2] Unfortunately, no further information about this approach could be found in the literature. Especially the question, which fitting parameters have to be set constant to generate the

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2021 The Authors. Journal of Raman Spectroscopy published by John Wiley & Sons Ltd

library is unknown to the authors. Judging from the output of the file generated by CARSFT (which we have access to) for FTCARS (which we do not have access to), at least the line width multiplier and the convolution with pump and probe is applied prior to generating the library and are thus constant for subsequent fitting of spectra. For a discussion of what this implies, see Section 2.2. A prominent example of sparse libraries is the approach from Cutler et al.^[4] to evaluate multiple resonant species using a dual-pump approach. Within the sparse library approach, the user is left with the temperature, resonant species mole fraction and a horizontal shift as fitting parameters,^[4] all other degrees of freedom of the physical model, which will be discussed later on in the paper, are already included in the library spectra and are thus inaccessible at runtime, i.e. during the fit process of spectra.

Summarizing the above, when it comes to evaluation based on fitting, so far one had to choose between being flexible and slow (as in direct fitting approaches) or being a little more restricted in terms of available fitting parameters and fast (as in library based approaches). In this publication, we present a method which tries to circumvent all the mentioned drawbacks by employing a global fitting scheme using a mixed-integer genetic algorithm (MI-GA) using compressed fully-packed libraries as described in Section 2.2. Within Section 3, we demonstrate that the presented algorithm is computationally efficient, retains all degrees of freedom and is to a great extent independent of the provided initial solution. Our code can be downloaded and used under MIT license either as supporting information (in the version used for this publication) or at our public github repository.^[5] The code presented herein is implemented in MATLAB (tested in versions R2017b and R2019b) and makes use of the optimization toolbox, however ports to other languages should be straightforward.

2 | METHOD DESCRIPTION

This section is split into two parts. Section 2.1 provides a concise description how CARS spectra are simulated, that is, the physical model required for fitting. For the purpose of this paper, we assume a typical experimental setup for single-shot broadband ro-vibrational CARS measurements with a narrowband pump ($\approx 1 \text{ cm}^{-1}$ bandwidth) and a broadband Stokes laser ($> 100 \text{ cm}^{-1}$ bandwidth) and acquisition of the signal using a spectrometer. However, these are not fundamental limitations of the code. The library-based genetic algorithm is described in Section 2.2.

2.1 | Simulation of CARS spectra

Physical quantities are derived from the spectral shape of the CARS signal, which is proportional to the squared

modulus of the third-order susceptibility $\chi^{(3)}$ given in Equation (1), where $\chi^{(3)}$ is summed over all transitions q . a_q and Γ_q denote the amplitude and line width of transition q and $\Delta\omega_q = \omega_q - \omega_{\text{pump}} + \omega_{\text{Stokes}}$ represents detuning from the Raman resonance. $\chi_{\text{NR}}^{(3)}$ represents the non-resonant contribution to the CARS signal, which is purely real and wavelength-independent.

$$|\chi^{(3)}|^2 = \left| \sum_q \frac{a_q}{\Delta\omega_q - i\Gamma_q/2} + \chi_{\text{NR}}^{(3)} \right|^2 \quad (1)$$

The amplitudes a_q of the resonant susceptibility are given in Equation (2), where N is the molecule number density of the resonant specie, c the speed of light, \hbar the reduced Planck's constant, ω_S the Stokes wavenumber, Δ_q the population difference between the initial and final state and finally the differential Raman cross-section $\left. \frac{\delta\sigma}{\delta\Omega} \right|_q$, which includes the Hermann-Wallis-factors for branch dependent vibration-rotation-interaction. All equations are adapted from Marrocco.^[6]

$$a_q = \frac{Nc^4}{\hbar\omega_S^4} \Delta_q \left. \frac{\delta\sigma}{\delta\Omega} \right|_q \quad (2)$$

In addition to the transitional line widths Γ_q , a measured CARS signal consists not only of the spectrum summed over all transitions q but also depends on the combined apparatus function, that is, the line width and shape of the pump and probe laser as well as the spectrometer resolution. Here, they are modeled as a combined Gaussian G_i , which is used as convolution kernel for $|\chi^{(3)}|^2$; see Equation (3). In principle, arbitrary line shapes can be included in the code.

Also, the spectral profile of the Stokes laser is typically not flat. In this approach, this is accounted for by assuming a non-resonant signal in, for example, methane or argon has been acquired that is used to normalize the CARS signal.

$$I_{\text{CARS}}/I_{\text{Stokes}} = I_{\text{CARS}}^* \propto G_i \otimes |\chi^{(3)}|^2 \quad (3)$$

For the transitional line widths of N_2 at high temperatures, the modified exponential gap model from Rahn et al.^[7] is implemented, which we believe is the same model that CARSFT uses. At the moment, only three models for the transitional line shapes are implemented in MARSFT. These are (a) isolated lines, (b) Voigt profile, and (c) rotational diffusion.^[8] Within the isolated lines and the rotational diffusion model, each transition is modeled as an individual Lorentzian without contributions from Doppler broadening or line mixing. The Voigt line shape includes Doppler broadening by convolving every Lorentzian transition with a Gaussian with the Doppler width individually. A line width multiplier to account for line broadening mechanisms is introduced in the following section. For more details on the computation of

$\chi^{(3)}$, please refer to the appendix provided as supporting information and the code.

2.2 | Library-based genetic algorithm fitting algorithm

2.2.1 | Degrees of freedom in fit

Within the code, seven degrees of freedom (DOF) are implemented for fitting. These split into two categories:

Physical: Temperature T , resonant species mole fraction X_R , line width multiplier n and instrumental function G_i

In order to approximate foreign gas broadening, a line width multiplier is introduced in the same fashion as in CARSFT, that is, by multiplying the Lorentzian line width Γ_q of every Raman transition with a J independent constant n .^[2] While this is far from ideal, it is a method to approximate collisional interaction as usually no information about the chemical composition within the measurement volume is available. At the moment, pressure is not included as a fit parameter, that is, the code is limited to constant pressure applications. However, the line widths in the implemented model scale linearly with pressure, that is, pressure and line width multiplier are actually mathematically equivalent. Equation (4) shows the dependence of the physical model to the fitting parameters.

$$I_{\text{CARS}}^* \propto G_i \otimes |\chi_R^{(3)}(T, X_R, n) + \chi_{NR}^{(3)}(X_R)|^2 \quad (4)$$

Note that due to the approximation of foreign gas broadening with the aforementioned line width multiplier, the total number density N does not appear in the equation because both the resonant and non-resonant contribution scale linearly with N (i.e., $|\chi^{(3)}|^2 \propto N^2$). Instead, the resonant species mole fraction basically sets the ratio between *resonant* and *non-resonant* contribution of the signal. This is modeled by multiplying the resonant contribution with X_R and computing the non-resonant contribution from the non-resonant susceptibility of the resonant species ($\chi_{NR,R}^{(3)}$) and a presumed non-resonant susceptibility of the mixture of *all other* species ($\chi_{NR,b}^{(3)}$). Precise knowledge of this parameter is required to obtain physically meaningful results for the mole fraction. With this, Equation (4) reads:

$$I_{\text{CARS}}^* \propto G_i \otimes |X_R \chi_R^{(3)}(T, n) + (1 - X_R) \chi_{NR,b}^{(3)} + X_R \chi_{NR,R}^{(3)}|^2 \quad (5)$$

Experimental: Intensity expansion (for very noisy data), wavenumber shift, wavenumber expansion

To account for noisy data, an intensity expansion allows to scale the model vertically (after normalizing the measured and simulated spectra to the area under the signal). Wavenumber shift and expansion are implemented

to account for imperfect spectral calibration of the spectrometer. Naturally, these two DOFs are non-orthogonal, introducing many local minima during optimization. In practical applications, it is advisable to fix the wavenumber expansion after calibration of the spectrometer, as this parameter is in fact determined by the dispersion of the spectrometer only. As a result of beam-steering, the wavenumber shift may vary due to different focal positions of the CARS signal within the entrance slit of the spectrometer.

2.2.2 | Library generation and compression

When creating a library, one has to weigh the loss of generality—that is, to how many problems is one instance of the library applicable—and precision against the complexity and size of the library. Additionally, accessing a library spectrum should obviously be significantly more efficient than simulating it directly.

Looking at Equations (1) and (3), the time consuming part of the calculation is before the convolution of the theoretical susceptibility due to the high wavenumber resolution required to adequately resolve the smallest Lorentzian within the spectral region of interest. For example, simulating the first two vibrational bands of a N_2 ro-vibrational CARS squared theoretical susceptibility at a temperature of 2000 K requires a spectral resolution of 0.0006 cm^{-1} to discretize the narrowest transition with 10 datapoints. This results in an array with a length of approximately 485,000 elements (or 1.8 MB RAM at single precision) for the entire spectrum, summed over all transitions. Thus, from a memory perspective, tabulating the squared modulus of the theoretical susceptibility in a temperature range between 300 and 2300 K, which corresponds to approximately 4 GB, seems feasible. Unfortunately, this requires the resonant species mole fraction, χ_{NR} and the line width multiplier to be constant, see Equation (4), which is an unacceptable restriction in most practical cases. If instead of the squared modulus the *complex* resonant susceptibility $\chi_R^{(3)}(T, n)$ is tabulated, the dependency of the library on mole fraction X_R (and thus χ_{NR}) vanishes while doubling the memory requirement as two arrays (for the real and imaginary component) have to be stored separately.

However, the implication of a constant line width multiplier prevails. In a turbulent combustion context, where the chemical composition changes on a single-shot basis, this may be problematic, so ideally, the library should not impose any a priori assumptions on this parameter. As the line width multiplier affects the spectrum at a transitional level, one either has to store every transition independently, drastically increasing memory requirements, or treat the line width multiplier as an additional dimension

for the table. Unfortunately, storing the complex susceptibility at 2000 temperature increments and 100 linewidth increments results in $\approx 10^{11}$ - 10^{12} datapoints, corresponding to $\mathcal{O}(1\text{ TB})$ of RAM at single precision.

To optimize memory performance without compromising precision, a loss-less library compression scheme was developed. The basic idea is that the spectral resolution of a measured spectrum is limited by the apparatus and typically orders of magnitude larger than the spectral resolution required to resolve the narrowest lines. To make use of this fact, consider the following properties of convolution: (a) convolution is *associative* and (b) convolution is *distributive*. Additionally, the convolution kernel for the line shape is normalized to unity area, that is,

$$G_i \otimes a = \int_{-\infty}^{\infty} G_i(\hat{\omega}) \underbrace{a(\omega - \hat{\omega})}_{=\text{const.}} d\hat{\omega} = a \int_{-\infty}^{\infty} \underbrace{G_i(\hat{\omega})}_{=1} d\hat{\omega} = a$$

for any real constant a .

Performing the squared modulus in Equation (4) explicitly and making use of the distributivity and associativity leads to the form in Equation (6):

$$G_1 \otimes G_2 \otimes |\chi_R^{(3)} + \chi_{NR}^{(3)}|^2 = G_1 \otimes \left(G_2 \otimes |\chi_R^{(3)}|^2 + 2 \left(G_2 \otimes \Re \chi_R^{(3)} \right) \chi_{NR}^{(3)} + \left(\chi_{NR}^{(3)} \right)^2 \right) \text{ with } G_i = G_1 \otimes G_2 \quad (6)$$

Based on this idea, a Gaussian convolution kernel G_2 with a width *narrower* than the total instrumental line width of G_i is used to convolve the *squared modulus* and the *real part* of the resonant susceptibility independently prior to the tabulation. This effectively reduces the memory requirement in the aforementioned example to 1–2 GB, depending on the chosen preconvolution kernel width which effectively depends on the available equipment. In a later step, when a spectrum is generated from a library, the convolution to the final instrumental resolution is performed. As this process is mathematically equivalent to a one-step convolution, no information is lost during compression.

This scheme allows to retain the flexibility of a non-library based approach, that is, all seven DOFs are still accessible for the optimizer, albeit with a slightly reduced precision for the parameters *temperature* and *line width multiplier*.

2.2.3 | Spectral fit

During the fit process, spectra are generated from the library, and the experimental DOFs are applied by scaling the spectra vertically and interpolating them on the experimental grid, accounting for wavenumber shift and

expansion. This process is implemented in a vectorized fashion, allowing very efficient generation of spectra. For the optimization, we chose a mixed-integer genetic algorithm (MI-GA)^[9] due to its low susceptibility to local minima and the ability to handle non-continuous parameters. Unlike for a gradient-based approach, there is no requirement on differentiability of the objective function. This allows to avoid a costly interpolation between spectra by treating temperature and line width multiplier by their library index only. Effectively, this reduces CPU time while increasing requirements on RAM size and bandwidth as the library has to be more densely populated. Additionally, this reduces the achievable precision in those two parameters to the library resolution, which should be sufficient for most cases as experimental noise limits the temperature precision to a few K anyway. The implemented algorithm adapts the population size for a given generation and the maximum number of stall generations—that is, the amount of generations that do not produce a change of the quality of the solution larger than a defined function tolerance value—based on the number of free parameters (25 individuals in each per free fitting parameter; 3 max. stall generations per free parameter) and uses a Gaussian mutation scheme to provide genetic diversity and

broaden the search space. For more details regarding the GA implementation, please refer to the code.

Different variance-weighting strategies are implemented for the objective function to avoid excessive noise-fitting. These are (a) standard deviation based for cases where the only variance of the signal is noise (e.g., in a laminar flame), (b) shot-noise, that is, the variance is proportional to the signal, and (c) a custom noise model that allows to account for, for example, measured Stokes noise. The implemented residual given in Equation (7) corresponds to a modified version of the reduced chi-squared residual R proposed in Cutler and Magnotti,^[4] where S denotes the square root¹ of the CARS signal I , σ^2 the variance used for weighting, N_p the number of pixels in the spectrum, and $\nu = N_p$ minus the number of fit parameters. In Cutler and Magnotti,^[4] the CARS signal is fitted to the squared modulus of the resonant susceptibility. For the investigations presented in the following section, we found fitting the square root of the CARS signal to be more robust. This may be different for real measurements

¹To be precise, due to noise \sqrt{I} could return complex results which obviously does not make sense. So $\text{sign}(\sqrt{|I|})$ is used instead to preserve the sign while taking the square root of the signal magnitude only.

with stronger variations in resonant species mole fraction. Care must be taken to use the variance in the appropriate units, i.e. as the square root of the CARS signal is fitted, the variance must be computed in the same units.

$$R = \frac{1}{\nu} \sum_k^{N_p} \frac{(S_k - |\chi_k^{(3)}|)^2}{\sigma_k^2} \quad (7)$$

2.2.4 | Limitations

Currently, only N_2 ro-vibrational CARS is implemented. Extending the code to other diatomic molecules or pure rotational CARS should be straightforward. In principle, the code can also be extended to multiple species, however this needs some adaption of the library structure. As the complex theoretical susceptibility is tabulated—and not its squared modulus—there is no dependency on the resonant species' mole fraction in the library. This has the obvious advantage, that RAM requirement scales approximately linearly with (depending on the number of model parameters that other species may need) and not as the power of the number of species such as in Cutler and Magnotti.^[4] Note that this is only true as long as interactions on a transitional level, such as collisional broadening, are approximated without dependence on other species' mole fractions.

3 | RESULTS

In this section, numerical experiments with simulated spectra are performed to determine the performance of MARSFT in terms of robustness, efficiency and accuracy. Throughout the entire section, MARSFT uses a library with $\Delta T = 1$ K temperature and $\Delta n = 0.04$ line width increments corresponding to a size of approximately 2 GB. To put the results into a context, the same spectra have been evaluated using CARSFT. We chose this way because we believe that CARSFT and similar codes are widely used and experienced readers might be able to better interpret the results when they see a known code working on the same data set. It is explicitly not intended to depreciate the value of CARSFT in any way nor is it intended to imply that CARSFT represents all codes that use gradient-based optimization.

3.1 | Dependency on initial solution

In order to test the dependency on the provided initial solution, three cases are considered, which vary in the number of free parameters for the fit and the initial solution given to the respective code:

Case 1a: All parameters free, perfect starting solution

Case 1b: All parameters free, 100 random starting solutions

Case 1c: CARSFT only: Same as b, but run fit 6 times, using the previous solution as the starting solution for the next run

Case 1d: Fixed wavenumber shift and expansion to perfect solution. Other parameters as in b

Case 1e: CARSFT only: Same as d, but run it 6 times as in c

For all cases, synthetic spectra without additional noise at 100 K intervals in the range 300–2200 K were evaluated. All other model parameters were kept constant at the following values: $X_{N_2} = 0.8$, $n = 3$, $G_i = 1$, wavenumber shift = 0, wavenumber expansion = 0. The boundaries for the random starting solutions were $T \in [280, 2400]$, $X_{N_2} \in [0.3, 1]$, $n \in [2, 4]$, $G_i \in [1, 2]$, wavenumber shift $\in [-5, 5]$, wavenumber expansion $\in [-3, 3]$. Random starting solutions are uniformly distributed within these bounds.

Figure 1 shows the results for Cases 1a–e. As expected, Case 1a, which intends to show that the codes do not walk off the perfect solution, shows no significant deviation from the expected result with both codes. Interestingly, due to the intrinsic randomness of the GA, Case 1a—even though the perfect solution was used as a starting solution and as such is a guaranteed part of the initial population—the temperature error is non-zero. For some realizations, it is even outside the 2σ range of Case 1b. Fixing the wavenumber shift and expansion, the uncertainty is limited only by the precision of the generated library. While MARSFT stays mostly independent from the provided initial solution and number of free parameters, CARSFT has a strong dependency, especially when wavenumber expansion and shift are free (Case 1b). The results improve when the code is allowed to run multiple times with iteratively improved starting solutions (Case 1c), but in order to achieve a useable robustness, the wavenumber shift and expansion have to be fixed (Case 1d). But even then, the errorbar spans a range of 300 K. After 6 successive runs (Case 1e), this reduces to a few K. It is important to note that the conclusion *a GA based solver is always better than a gradient-based solver* cannot be drawn by this simplified comparison. It is meant to demonstrate the extent to which MARSFT is independent of the provided initial solution compared to a well-established code within the combustion community. Modern implementations of gradient-based solvers may have similar properties as the proposed MI-GA.

3.2 | Accuracy and precision

To assess the accuracy and precision of MARSFT in a controlled environment, the simulated temperature profile of




FIGURE 1 Dependency of the temperature error $\Delta T = T_{fit} - T_{actual}$ on initial solution for Cases 1a–e. Note the different y-axis limits. Errorbars span a 2σ interval, dots represent single realizations. Top row: MARSFT, bottom row: CARSFT [Colour figure can be viewed at wileyonlinelibrary.com]

a laminar premixed atmospheric methane/air flame at $\Phi = 1$ (CANtera,^[10] GRI-MECH 3.0 reaction mechanism^[11]) was superimposed with a noise model with contributions from dark noise, shot noise and spectral noise from the broadband dye laser.^[4] In most situations, spectral noise of the Stokes laser is the dominant source term; see Figure 2. For this reason, the relative standard deviation of a *measured* non-resonant CARS signal was used to add this noise contribution. The simulated signal intensity was scaled in a way that it matches the approximate signal levels that we typically achieve in our experiments which is in the order of 2000 counts on a state-of-the-art CCD camera at the first vibrational band head at ≈ 2200 K. For lower temperatures, the signals were scaled by $(2200/T)^2$, as the signal intensity is proportional to N^2 . The simulated temperature profile is discretized at 50 locations, each consisting of 100 noisy spectra. To account for varying non-resonant background, the real buffer gas susceptibility was calculated from the chemical composition from the laminar flame simulation². The varying buffer gas susceptibility leads, even though the physical N_2 mole fraction does not change much across the flame front, necessarily to a varying result of the fitted mole fraction. The value for the buffer gas susceptibility was kept fixed for both codes at the default value of 8.5. All spectra were simulated with the following model parameters: $n = 3$, $G_i = 1$, wavenumbershift = 0, wavenumberexpansion = 0. Temperature and mole fraction were taken from the laminar flame simulation.

As this is a simulation of a more practical situation, the following assumptions have been made: (a) The instrumental line width and wavenumber expansion are known. This is justified because they can be assumed constant at least on a day-by-day basis and thus can be determined with good accuracy. Additionally, beam-steering effects are assumed to be negligible such that the wavenumber shift

is fixed as well; (b) because the flame is laminar, chemical composition and as such the line width multiplier and N_2 mole fraction does not change. This parameter is determined from fitting the mean spectrum first and then kept constant for fitting the single shots; (c) The temperature of the mean fit is used as a starting solution for the evaluation of the single shots.

Note that as MARSFT normalizes the experiment and simulation to the area under the curve rather than the peak (what CARSFT appears to be doing), CARSFT needed the intensity expansion to be a free parameter to improve precision. For MARSFT, no influence could be observed so it was fixed to unity for all cases in this section. With the reduced set of fitting parameters and a proper selection of the initial solution approximately centered in the bounds, CARSFT converged after one run to a reasonable solution.


The following cases have been simulated which vary in the definition of the variance:

- Case 2a:** No noise weighting, that is, $\sigma_k^2 = 1$
- Case 2b:** Shot-noise weighting, that is, $\sigma_k^2 = S_k$
- Case 2c:** Standard-deviation based noise weighting (MARSFT only), that is, $\sigma_k^2 = \sigma_{S_k}^2$, that is, the squared shot-to-shot standard deviation of the square root of the CARS signal
- Case 2d:** Modeled noise weighting (using the noise model of Figure 2, MARSFT only)

The results are summarized in Figure 3. Generally, the accuracy lies within ≈ 10 K across the entire temperature range and across all cases except the CARSFT evaluation for Case 2b, that is, with shot-noise variance weighting which seems to cause some trouble at higher temperatures, where the fit overestimates the temperature by up to 100 K (not shown in figure) and Case 2d, where MARSFT overestimates the temperature by up to 15 K around 1200 K. Interestingly, MARSFT and CARSFT show an opposite trend for Case 2a. While CARSFT tends to overestimate temperatures ≥ 1200 K, MARSFT under-

²Non-resonant contributions: $\text{CO}_2 = 12.59$, $\text{H}_2\text{O} = 18.5$, $\text{N}_2 = 8.5$, $\text{O}_2 = 8.5$, $\text{CO} = 8.5$,^[2] $\text{CH}_4 = 26$,^[12] all units 10^{-18} cm³/(erg amagat).

FIGURE 2 Left: noise model used for the simulated experiment. Right: Example of simulated CARS signal intensity in counts. Black line represents average of 100 shots, gray lines are 10 exemplary single shots to give an impression of the noise [Colour figure can be viewed at wileyonlinelibrary.com]



estimates the same spectra by approximately the same amount. As already pointed out by Cutler and Magnotti,^[4] adequate variance weighting leads to increased accuracy, as shown in the high temperature range for Case 2c–d.

As expected, precision of the evaluation is reduced for higher temperatures as SNR decreases. Without noise-weighting, the performance of both codes is similar. For Case 2b, where shot-noise weighting was applied, CARSFT shows a better performance than MARSFT, even though the accuracy diminished tremendously. In other words: CARSFT fits a wrong temperature value with great repeatability. The best precision was achieved using the standard-deviation based noise weighting and the application of the noise model (Cases 2c–d). It is not surprising that applying the very same noise model during the fit that was used to add the noise in the first place yields the best results. However, for actual experiments, the noise characteristic of the detection system can easily be obtained by measurements with an integrating sphere for dark noise and shot-noise and measurements of the spectral noise of the broadband dye laser.

3.3 | Computational efficiency

To assess the computational efficiency, the total runtime for every case and the average runtime per spectrum was measured; see Figure 4. All evaluations with MARSFT were performed twice, once on a single-core and once using the Parallel Computing Toolbox on 12 cores in parallel (Intel Xeon E5-2630 @ 2.6 GHz). As expected, even on a single-core, MARSFT performs faster than the reference code. Only Case 1a and Case 1b show a slightly worse performance, where all fitting parameters are free. However, since the accuracy of CARSFT was only acceptable and robust for Cases 1d and 1e, one should rather compare these cases where a speed up of approximately ≈ 5 times is achieved for single-core and ≈ 45 times for multi-core evaluation. The computational efficiency for Cases 2a–2d allows close to real-time fitting at 10 Hz. However, for actual real-time fitting in a turbulent flame, at least the mole fraction and the line width multiplier have to be fitted on a single-shot additionally to the temperature. In principle, one could buffer a set of spectra and evaluate them in a multi-core environment for on-line evaluation.

Due to the many differences between implementations, a mere comparison of total runtime may indicate a good estimate of the net processing duration which is a valuable information for practical applications. However, with this global comparison, it is not possible to identify the merits of individual aspects of this approach. As a guidance, we report the approximate computation time per spectrum in our implementation for different strategies: (a) direct simulation without library: $\approx 1\text{--}3$ s per spectrum. There is a temperature dependency in calculation time due to the fact that the code only includes transitions with amplitudes $\geq 0.1\%$ of the maximum amplitude. Obviously, including fewer transitions reduces computational cost but increases inaccuracies in the spectrum. As this temperature dependency vanishes for library-based spectrum generation, the simulation of individual spectra is optimized for precision rather than computational cost. (b) Interpolation-based library access: ≈ 5 ms per spectrum. This would be necessary if the optimization procedure uses a solver that requires differentiability and/or continuous input variables. (c) Direct library access: ≤ 0.1 ms. Direct library access requires a solver that can handle mixed-integer optimization.

The resulting computational cost per fitted spectrum additionally depends on the number of function calls. A genetic algorithm tends to require more function calls than a gradient-based solver, however, as this is highly dependent on implementation and application, no systematic investigation was performed. In addition, there are numerous influences on computational efficiency arising from the chosen programming language (compilation, architecture optimization, vectorization and many more) and obviously achievable memory bandwidth, which becomes more important for library based approaches.

4 | CONCLUSION

Within this publication we described our novel approach to evaluate coherent anti-Stokes Raman spectra based on fully-packed compressed libraries to achieve a high computational efficiency with an accuracy and precision comparable or better than traditional fitting at the cost of discretizing temperature and linewidth multiplier. To achieve a compact library size, the squared modulus and the real




FIGURE 3 Top: Accuracy, represented by error in the temperature evaluation $\Delta T = T_{fit} - T_{actual}$ for Cases 2a–d. Bottom: Precision represented by the temperature standard deviation of 100 single-shot evaluations $\sigma(T)$. Dashed lines: CARSFT, solid lines: MARSFT [Colour figure can be viewed at wileyonlinelibrary.com]




FIGURE 4 Computational cost for the evaluation of all spectra in this publication. Top: MARSFT; left bars: evaluation on a single core, right bars: parallel evaluation on 12 cores. Bottom: CARSFT, only single-core evaluation available. Numbers on top of the bars represent average runtime per spectrum [Colour figure can be viewed at wileyonlinelibrary.com]

part of the resonant susceptibility are convolved with an arbitrary Gaussian compression kernel that has a narrower width than the resulting instrumental linewidth of the experiment prior to tabulation. This retains all degrees of freedom for the genetic algorithm fit at significantly lower computational cost. Additionally, the proposed method is practically independent from the provided starting solution. Currently, the implementation is done in MATLAB because of its wide use in academics and the fact that a decent GA implementation supporting mixed-integer optimization is available. However, porting to different languages should be straightforward using the provided code base. Although the actual implementation is currently restricted to ro-vibrational N_2 CARS, the proposed method of loss-less library compression may be generalized to apply to any situation where a convolution step separates the time consuming from the time efficient computations.

ACKNOWLEDGEMENTS

We would like to thank the colleagues from Sandia National Labs for creating CARSFT without which the presented code in this publication would not exist. The name

MARSFT (MATLAB-based coherent Anti-Stokes-Raman Fitting Tool) was chosen to honor the extraordinary amount of effort necessary to create CARSFT. Special thanks go to Leon Kleffmann for the fruitful discussions on optimizing code performance.

We highly appreciate funding from *Deutsche Forschungsgemeinschaft* through project DR 374/24-1. Open Access funding enabled and organized by ProjektDEAL.

ORCID

M. Greifenstein  <https://orcid.org/0000-0001-5624-3165>

A. Dreizler  <https://orcid.org/0000-0001-5803-7947>

REFERENCES

- [1] R. J. Hall, A. C. Eckbreth, J. F. Ready, R. K. Erf, 5, Academic Press New York **1984**.
- [2] R. E. Palmer, Sandia National Labs., Livermore, CA (USA) **1989**.
- [3] H. J. L. Van der Steen, J. D. Black, *Neural. Comput. Applic.* **1997**, 5(4), 248.
- [4] A. D. Cutler, G. Magnotti, *J. Raman. Spectrosc.* **2011**, 42(11), 1949.
- [5] M. Greifenstein, A. Dreizler, <https://github.com/tudarsm/marsft>
- [6] M. Marrocco, *J. Raman. Spectrosc.* **2009**, 40(7), 741.

- [7] L. A. Rahn, R. E. Palmer, M. L. Koszykowski, D. A. Greenhalgh, *Chem. Phys. Lett.* **1987**, 133(6), 513.
- [8] R. J. Hall, D. A. Greenhalgh, *Opt. Commun.* **1982**, 40(6), 417.
- [9] K. Deep, K. P. Singh, M. L. Kansal, C. Mohan, *Appl. Math. Comput.* **2009**, 212(2), 505.
- [10] D. G. Goodwin, H. K. Moffat, R. L. Speth, <http://www.cantera.org>, Version 2.1.0, **2015**.
- [11] G. P. Smith, http://www.me.berkeley.edu/gri_mech/, **2011**.
- [12] J. W. Hahn, E. S. Lee, *J. Opt. Soc. Am. B* **1995**, 12(6), 1021.

How to cite this article: Greifenstein M, Dreizler A. MARSFT: Efficient fitting of CARS spectra using a library based genetic algorithm. *J Raman Spectrosc.* 2021;52:655–663. <https://doi.org/10.1002/jrs.6046>

SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.