Spectrochimica Acta Part B: Atomic Spectroscopy 246 (2026) 107635

Contents lists available at ScienceDirect

Spectrochimica Acta Part B: Atomic Spectroscopy

journal homepage: www.elsevier.com/locate/sab

Simulation-based calibration-free laser-induced breakdown spectroscopy using gradient optimization

Christoph H. Egerland [a,b,*], Kristin Rammelkamp [a], Elise Clavé [a,c], Ana Lomashvili [a,d], Peder Bagge Hansen [a], Fabian Seel [a], Susanne Schröder [a], Heinz-Wilhelm Hübers [a,b]

[a] Institute of Space Research, German Aerospace Center, DLR, Rutherfordstr. 2, 12489, Berlin, Germany
[b] Department of Physics, Humboldt-Universität zu Berlin, Newtonstraße 15, 12489, Berlin, Germany
[c] Universite Claude Bernard Lyon 1, ENS de Lyon, CNRS, UJM, LGL-TPE, UMR 5276, Villeurbanne cedex, F-69622, France
[d] BIFOLD and TU Berlin, Berlin, 12489, Germany

ARTICLE INFO

Dataset link: https://github.com/DLR-WR/libskit, https://github.com/DLR-WR/libs_gd_paper

Keywords:
LIBS
Optimization
Machine learning

ABSTRACT

We implement a local gradient-based scheme to fit Laser-Induced Breakdown Spectroscopy (LIBS) spectra to the emission of a uniform, isothermal, stationary plasma in local thermal equilibrium. The gradients of the loss function are obtained by automatic differentiation. We demonstrate the self-consistency and robustness of our fitting scheme through systematic analysis of synthetic spectra and find that, for a wide range of plasma parameters, no additional minimum of the loss function was detected within the investigated parameter domain, enabling reliable recovery of the true parameters.

1. Introduction

Laser-induced breakdown spectroscopy (LIBS) is a method from the field of atomic emission spectroscopy used for qualitative and quantitative elemental analysis [1]. A pulsed laser is focused onto the target material where material is ablated and transformed into a transient plasma. Spectral analysis of the plasma emission gives information about the elemental composition of the plasma that comprises typically both elements of the sample but also some contribution of the ambient atmosphere. Quantitative information is typically derived through calibration-based approaches by measuring matrix-matched reference samples with the same or very similar hardware in representative ambient conditions and correlating spectral features to known concentrations [2–4]. LIBS is increasingly used for terrestrial and extraterrestrial in-situ analysis of geomaterials with the most famous examples being ChemCam and SuperCam on the two NASA Mars rovers Curiosity and Perseverance, respectively [5–9]. Due to the enormous compositional and mineralogical diversity of geologic materials, it is practically impossible to obtain fully matrix-matched reference materials for all sample types, which limits the applicability of conventional calibration-based quantification methods. Calibration-free LIBS (CF-LIBS) offers an alternative to calibration-based techniques where no matrix-matched or other laboratory data is required [10]. Here, physics-based plasma models are used as opposed to the, often purely, data-driven approaches in calibration LIBS.

Several plasma models have been investigated in the literature. The simplest case is that of a uniform, isothermal and stationary plasma [11, 12]. To accurately describe complex plasma phenomena, more sophisticated plasma models are needed. Two-layer models are more appropriate to capture effects of inhomogeneity within the plasma as it interacts with the atmosphere [13–15]. Due to instrumental restrictions in field missions, short-time resolution of the plasma evolution cannot always be achieved, which is where the need to model time-integrated spectra arises.

Increasing complexity due to the incorporation of spatial inhomogeneity, time integration and plasma chemistry into plasma models raises the issue of fitting the spectra of these complex plasma models, since they have more parameters. Existing fitting schemes suffer from the curse of dimensionality and are rather slow global models or they require expert knowledge on the target composition and emission lines such as in the iterative approach of [16]. This calls for a fitting scheme that is both fast, scales well with the number of plasma parameters and does not require expert knowledge.

In this work we demonstrate the utility of gradient-based local optimization for CF-LIBS spectral fitting. We implement a uniform, isothermal, stationary plasma model in local thermal equilibrium (LTE) and use automatic differentiation to obtain mathematically exact gradients for the L-BFGS optimization algorithm. We validate our approach

* Corresponding author at: Institute of Space Research, German Aerospace Center, DLR, Rutherfordstr. 2, 12489, Berlin, Germany.
E-mail address: christoph.egerland@dlr.de (C.H. Egerland).

https://doi.org/10.1016/j.sab.2026.107635
Received 23 June 2026; Received in revised form 28 August 2026; Accepted 28 August 2026
Available online 5 September 2026

through comprehensive testing on synthetic spectra, establishing a robust foundation for future application to experimental data.

1.1. Plasma model

We follow the notation of [17] for the plasma model formulation, which is consistent with other CF-LIBS implementations [11,16,18].

We implement a plasma that is in local thermal equilibrium (LTE), obeying the following assumptions: (1) isothermal, (2) homogeneous, (3) stationary. Note that only the plasma emission is modeled, not the ablation process.

With these assumptions, a plasma of K elements can be described by $3+(K-1)$ parameters: the temperature T, the electron number density n_e, the plasma length L, and the mass concentrations $c_k (k = 0, \dots K-1)$ where one of the concentrations is redundant because they have to sum to unity. Equivalently we could give the elemental number densities of all elements and then calculate the electron density using the neutrality condition. To calculate the plasma's emitted irradiance we first need to calculate the plasma species' densities. Hereafter we denote with n the number densities (m^{-3}), with index $k = 1, \dots K$ the element and with index $z = 0, \dots, Z_k$ the charge of the ion (where $z = 0$ is the neutral atom and Z_k is the maximum ionization of element k).

1.1.1. Elemental number densities

Ionic population fraction $\beta_{n,k}$. We first calculate the fraction of element k that is in a specified ionic state. This is given by $\beta_{k,z} = n_{k,z}/n_k$ which can be found by solving the following system of Saha equations

$$\frac{\beta_{k,z+1} n_e}{\beta_{k,z}} = \frac{2}{\Lambda_e(T)^3} \frac{U_{k,z+1}(T)}{U_{k,z}(T)} \exp\left(-\frac{\chi_{k,z}}{k_B T}\right), \tag{1}$$

$$\sum_{z=0}^{Z_k} \beta_z = 1. \tag{2}$$

Here $\chi_{k,z}$ is the ionization potential and Λ_e is the thermal deBroglie wavelength of the electron given by

$$\Lambda_e(T) = \sqrt{\frac{2\pi\hbar^2}{m_e k_B T}}, \tag{3}$$

with the mass of the electron m_e, the Boltzmann constant k_B, the reduced Planck constant $\hbar$ and the partition function $U_{k,z}(T)$ given by

$$U_{k,z}(T) = \sum_{E_i < \chi_{k,z}} g_i \exp\left(-\frac{E_i}{k_B T}\right). \tag{4}$$

Eq. (2) is a normalization condition since all fractions must sum to one. The system of equations can be solved numerically, but there is an explicit solution, as shown in Appendix A

$$\beta_{k,z} = \frac{\prod_{i=0}^{z-1} \Omega_{k,i}}{\sum_{j=0}^{Z_k} \prod_{i=0}^{j-1} \Omega_{k,i}}, \tag{5}$$

where

$$\Omega_{k,i} = \frac{2}{n_e \Lambda_e^3} \frac{U_{k,i+1}}{U_{k,i}} \exp\left(-\frac{\chi_{k,i}}{k_B T}\right) \tag{6}$$

Elemental number density n_k. The elemental number densities n_k can be calculated from the given mass concentrations c_k through the following equations. Since the plasma is neutral it holds that

$$n_e = \sum_{k=0}^{K-1} \sum_{z=0}^{Z_k} z \beta_{k,z} n_k \tag{7}$$

The total number density of the whole plasma n is given by

$$n = \sum_k n_k \tag{8}$$

where we may also express the elemental number density n_k via the total number of particles N_k, the volume V and the molar mass M_k like so:

$$n_k = \frac{N_k}{V} = \rho N_A \frac{c_k}{M_k} \tag{9}$$

(with ρ being the density of the whole plasma and N_A Avogadros constant). After inserting Eq. (9) into Eq. (8) and rearranging for ρN_A we find

$$n_k = n \frac{\frac{c_k}{M_k}}{\sum_k \frac{c_k}{M_k}} \tag{10}$$

Inserting this expression into Eq. (7) and rearranging for the total number density yields

$$n = \frac{n_e \sum_k \frac{c_k}{M_k}}{\sum_k \left(\frac{c_k}{M_k} \sum_{z=0}^{Z_k} z \beta_{k,z}\right)} \tag{11}$$

Finally, we find an expression for the elemental number density that only depends on known quantities by inserting this back into Eq. (10)

$$n_k = n_e \frac{c_k}{M_k} \left(\underbrace{\sum_{k=1}^{K} \left(\frac{c_k}{M_k} \sum_{z=1}^{Z_k} z \beta_{k,z} \right)}_{:= F_{\mathrm{sum}}} \right)^{-1} \tag{12}$$

Plasma species number density $n_{k,z,i}$. The number density of an element k in the ionic state z and the electronic state i is obtained by the product

$$n_{k,z,i} = n_k \cdot \beta_{k,z} \cdot m_{k,z,i} \tag{13}$$

where we obtain n_k from Eq. (12), $\beta_{k,z}$ from Eq. (5) and $m_{k,z,i}$ is given by the Boltzmann distribution

$$m_{k,z,i} = \frac{g_{k,z,i}}{U_{k,z}} \exp\left(-\frac{E_{k,z,i}}{k_B T}\right) \tag{14}$$

1.1.2. Emission coefficient

We only consider bound-bound transitions and neglect the effects from continuum radiation. For this case the emission coefficient is given by:

$$\varepsilon_{bb}(\lambda) = \sum_{k,z,j,i<j} \frac{q_e}{4\pi} (E_j - E_i) A_{ji} n_{k,z,j} P_{k,z,i,j}(\lambda), \tag{15}$$

where A_{ji} is the transition probability of the transition $j \to i$, $P_{k,z,i,j}(\lambda)$ is the line profile.

1.1.3. Emission line broadening

The line profile function $P_{k,z,i,j}(\lambda)$ defines the shape of an emission line. It is a probability density function and satisfies

$$\int_0^\infty P_{k,z,i,j}(\lambda) d\lambda = 1 \tag{16}$$

We consider Doppler and Stark broadening as the most dominant effects, so the broadening is given by the convolution of a Gauss function G and a Lorentz function Lor, yielding a Voigt profile

$$P_{k,z,i,j}(\lambda) = V(\lambda; \sigma, \gamma) = \int_{-\infty}^{\infty} G(\lambda'; \sigma) \mathrm{Lor}(\lambda - \lambda'; \gamma) d\lambda' \tag{17}$$

The broadening parameters σ and γ are specific to every emission line and may also depend on the wavelength λ itself.

1.1.4. Irradiance via radiative transfer equation

To calculate the irradiance from the plasma we have to solve the one-dimensional radiative transfer equation (RTE)

$$\frac{dI(\lambda, x)}{dx} = \varepsilon(\lambda, x) - \alpha(\lambda, x) I(\lambda, x) \tag{18}$$

where $\varepsilon(\lambda, x)$ is the emission and $\alpha(\lambda, x)$ is the absorption coefficient of the plasma. We consider a homogeneous plasma, allowing us to drop the spatial dependence in the coefficients, enabling us to solve the RTE explicitly:

$$I(\lambda) = B(\lambda)\left(1 - \exp\left(-\frac{\varepsilon(\lambda)}{B(\lambda)}L\right)\right) \tag{19}$$

where we have used the fact that in a thermal equilibrium the absorption and emission are related via the blackbody radiation

$$B(\lambda) = \frac{\varepsilon(\lambda)}{\alpha(\lambda)} \tag{20}$$

$$B(\lambda, T) = \frac{2hc^2}{\lambda^5}\left(\exp\left(\frac{hc}{k_B T \lambda}\right) - 1\right)^{-1} \tag{21}$$

1.1.5. Instrumental broadening

The recorded spectrum is further broadened on the optical path that the photons travel on between plasma and detector due to the optical elements. This is quantified in an instrumental broadening function that is similar to the line broadening function in Eq. (17). After we convolute the computed irradiance with the instrumental broadening function we obtain the irradiance as it is detected in the spectrometer:

$$S(\lambda) = \int d\lambda P_{\text{instr}}(\lambda - \lambda')I(\lambda') \tag{22}$$

1.1.6. Parameters

We extracted all needed spectral line parameters as well as the ionization energies and levels from the NIST Atomic Spectra Database [19]. The Stark widths and shifts parameters were extracted from the Stark-B database [20]. The two data sources were combined by matching lines for ions that are present in both database. This significantly reduces the number of emission lines as there is a lack of documented Stark parameters for a lot of emission lines. Contrary to [14,17] we do not impute missing parameters by defining some default Stark width/shift, but rather do not consider these lines at all.

1.2. Spectral fitting

1.2.1. General problem setting

Generally in calibration-free LIBS, a "synthetic" spectrum, i.e. a spectrum stemming from a simulation, is computed and then, in a second step, this simulated spectrum is compared to the experimental spectrum. Mathematically, we can express this as a minimization problem [11]:

$$\underset{\theta}{\text{minimize}} \quad f(\theta) = f(I_{exp}, I_{syn}(\theta)) \tag{23}$$

where f is the cost (or loss) function which measures the accuracy of the fit between the synthetic and experimental spectrum, I_{syn} is the plasma model and θ is a vector of plasma parameters. In our case of a uniform and stationary LTE plasma we have

$$\theta = [T, n_e, L, c_0, c_1, \ldots, c_{K-1}], \tag{24}$$

where T is the plasma temperature, n_e is the electron density, L is the plasma length and c_i $(i \in 0, \ldots, K-1)$ are the elemental mass concentrations. The overall fitting procedure is illustrated in Fig. 1.

In any case, the optimization scheme that is used to perform this minimization should converge on the same set of true parameters regardless of initialization. The scheme should also ideally possess these properties: (1) It should be reasonably fast, enabling the analysis of large experimental datasets in a feasible time frame (2) Ideally, it should not require further knowledge on the target composition

Convergence. In the general case presented above we cannot know a priori whether there is only one global minimum. This depends on the plasma model, the loss function and the experimental spectrum itself. Other groups have argued that in these cases, local optimization methods might get caught in local minima [21], therefore favoring global solvers. We address this concern empirically with two complementary diagnostics: a multi-start convergence study (Section 2.2) and a direct sampling of the loss landscape (Sections 1.2.5 and 2.3), both of which indicate a single basin for the LTE model on the studied parameter range.

No knowledge. The iterative method presented in [16] presupposes knowledge of which elements are in the target, requiring a preprocessing step. As with many other CF-LIBS schemes the iterative method also relies on the selection of optically thin lines for the determination of the plasma temperature. In an ideal scenario, a fitting scheme utilizes all available lines, optically thin and self-absorbed, in the spectrum to extract the maximum amount of information, thus requiring no knowledge a priori.

Reasonably fast. Rapid inference is required for data intensive tasks, where a larger number of spectra should be analyzed. Other groups have grappled this issue with parallelization [21].

1.2.2. Loss functions

Different loss functions have been considered in the literature. The so-called L2-Loss (or sum of squared errors) measures the "distance" that the guess spectrum is off from the true spectrum and is used in standard curve-fitting methods like Levenberg–Marquardt [14].

$$f_{\text{L2}} = \sum (y_{\text{true}} - y_{\text{predicted}})^2. \tag{25}$$

It is commonly applied in spectral fitting, e.g. [15,16,22,23]. A variant of the cosine similarity has been employed by [11,21]. A custom loss function based on a correlation function employing Pearsons correlation coefficient has been used in [17].

The advantage of using L2-Loss is that one can straightforwardly apply (non-)linear least squares methods for fitting. The disadvantage is that, numerically speaking, bigger peaks have a bigger impact on the loss function and therefore the fitting quality of minor elements may suffer. This may be alleviated by using appropriate weighting [14,18]. In the case of other loss functions, one has to resort to other fitting methods like grid search [17] or Monte-Carlo methods [11,21].

In this study we primarily use the L2 loss function f_{L2} defined above, combined with an analytic normalization strategy that computes the optimal scale factor $c_{\text{opt}} = \langle I_{\text{sim}}, I_{\text{sim}} \rangle / \langle I_{\text{sim}}, I_{\text{exp}} \rangle$ before evaluating the loss. This avoids the need for an explicit normalization step and ensures that all spectral features contribute to the fit proportionally. We investigate the effect of different loss functions in Section 2.5 and the effect of different normalization strategies in Section 2.6.

1.2.3. Local vs. global optimization

Fitting schemes may be divided into local and global optimization schemes. Global schemes search over the whole parameter range, whereas local optimization searches for a minimum around the region of a provided initial guess. Global optimization schemes usually suffer from the curse of dimensionality, i.e. their computational complexity scales exponentially with the number of plasma parameters. Local optimization schemes may become stuck in local, not global minima. In our case this would mean that the scheme finds a minimum that does not correspond to the true physical parameters. We discuss this point in Sections 2.2 and 2.3; for the LTE model studied here both diagnostics indicate that this concern does not materialize.

In the context of LIBS, global fitting schemes have typically been used so far by different groups [11,17,21], with the exception of the iterative scheme presented in [16], which does not use gradients.

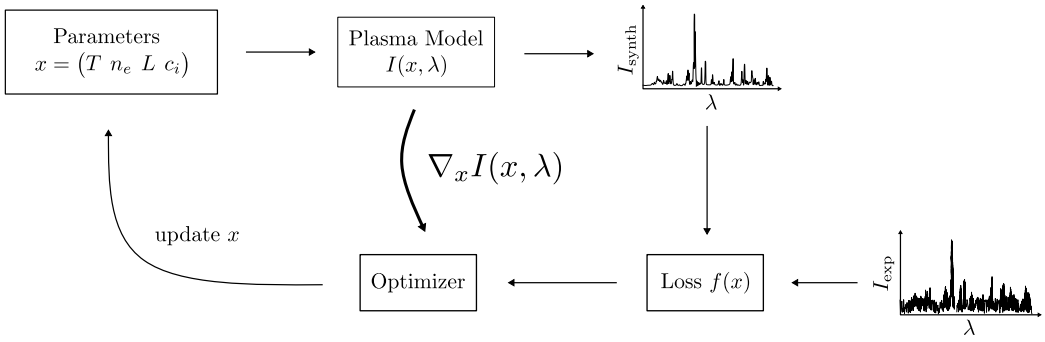

Fig. 1. General scheme for direct CF-LIBS. For a chosen plasma model, we insert the plasma parameters to obtain a "synthetic" LIBS spectrum. Then, the experimental spectrum has to be compared with the synthetic spectrum. An update mechanism updates the parameters to improve the fit between simulation and experiment.

1.2.4. Internal parametrization

The physical plasma parameters span multiple orders of magnitude ($T \sim 10^3$–10^5 K, $n_e \sim 10^{22}$–10^{24} m^{-3}, $L \sim 10^{-5}$–10^{-2} m) and for the element mass fractions we have $\sum_i c_i = 1$, $c_i \geq 0$. Optimizing directly in physical units is difficult for the following reasons: the large parameter ranges lead to ill-conditioned optimization landscapes, where the loss function changes very rapidly in some directions (e.g., n_e) and very slowly in others (e.g., T), resulting in elongated, banana-shaped basins of attraction. Furthermore the constraint that the mass fractions must sum to one and be non-negative introduces additional complexity, as it requires either projection methods or penalty terms to enforce these constraints during optimization. Both effects produce stiff and soft directions in parameter space and make L-BFGS line searches inefficient.

We therefore reparametrize the problem in internal coordinates that have more similar scales The physical parameters T, n_e and L are scaled as $\log_{10} T$, $\log_{10} n_e$ and $\log_{10} L$, respectively, which compresses the wide parameter ranges into a more manageable scale. For the element mass fractions we introduce logits via the softmax function, which automatically enforces $\sum_i c_i = 1$ and $c_i > 0$ without explicit constraints. Note that while the softmax strictly enforces $c_i > 0$, elements absent from the sample converge to numerically negligible concentrations in practice, effectively reaching zero (as demonstrated in Section 2.9). Since the softmax is invariant under a uniform shift of all logits (adding the same constant to every u_i leaves the concentrations unchanged), we fix $u_0 = 0$ and optimize over the $K - 1$ free logits $u_1, \ldots, u_{K-1}$:

$$c_i = \frac{\exp(u_i)}{\sum_{j=0}^{K-1} \exp(u_j)}, \quad u_0 \equiv 0 \tag{26}$$

1.2.5. Loss landscape sampling and uniqueness diagnostic

A local optimizer converges to a stationary point of f, i.e. a point where $\nabla f = 0$. A spurious local minimum in the domain Ω would therefore appear as a second point in Ω at which the gradient norm collapses to (near) zero, distinct from the true optimum $\theta^\star$. This motivates a purely empirical diagnostic: sample the parameter space Ω densely, evaluate $\|\nabla f\|$ at each sample, and check whether any sample comes close to being stationary. If every sampled point has a gradient norm that is orders of magnitude larger than $\|\nabla f(\theta^\star)\|$, no additional basin of attraction was detected on that domain. We emphasize that sampling can neither prove uniqueness of the global minimum nor rule out spurious basins that are narrower than the sampling resolution; our diagnostic provides empirical evidence that no such additional basin was encountered for the specific synthetic case studied here. We show this in Section 2.3.

1.2.6. Gradient-based vs. derivative-free algorithms

Local optimization schemes can be made more efficient by providing gradients of the loss function in addition to its value. In this study we use the Limited-memory Broyden–Fletcher–Goldfarb–Shanno (L-BFGS) algorithm as implemented in the NLOpt library [24], a quasi-Newton method that approximates the inverse Hessian matrix using gradient information from previous iterations. L-BFGS updates the parameter vector θ_i in internal coordinates (Section 1.2.4) by computing a search direction from the current gradient $\nabla f(\theta_i)$ and the approximated curvature information, followed by a line search to determine the step size [25].

The gradients $\nabla f(\theta_i)$ are computed using automatic differentiation (see Section 1.2.7). We terminate the optimization when the relative change in parameter values falls below a prespecified threshold $\delta = 10^{-4}$. Independently of this stopping rule, convergence to a stationary point can be verified by inspecting the gradient norm $\|\nabla f(\theta^\star)\|$.

1.2.7. Automatic differentiation

In Section 1.2.6 we have shown that gradient descent uses the gradient of the loss function to update the parameters. Generally speaking, there are multiple ways of computing gradients, the most common being a finite-difference approximation

$$\partial f / \partial x \approx [f(x + \Delta x) - f(x - \Delta x)]/2\Delta x \tag{27}$$

where Δx is a small step size for x. This method is, however, susceptible to round-off errors and we need to evaluate the function f twice.

In this work we calculate gradients using reverse-mode automatic differentiation (AD) via the Enzyme AD framework [26]. Enzyme operates as a compiler plugin for LLVM (Low Level Virtual Machine, a widely used open-source compiler infrastructure) and transforms the compiled intermediate representation of a function into its reverse-mode derivative. Given a scalar loss function $\mathcal{L}(x)$ with $x \in \mathbb{R}^N$, a single reverse pass produces the full gradient $\nabla_x \mathcal{L}$ at a cost proportional to that of one forward pass, without requiring a separate differentiation pass for each of the N parameters. This is in contrast to forward-mode AD, which requires N forward passes (one per parameter), and to finite differences, which require $2N$ function evaluations. Because Enzyme differentiates the optimized machine code directly, the generated derivative code benefits from the same compiler optimizations as the original function.

1.3. Test case

Our method is fully generic with respect to the composition of samples. To test the fitting scheme, we use a synthetic spectrum of a mixture of four elements (Ca, Mg, Mn, Na), which are representative of elements commonly analyzed with LIBS in terrestrial and planetary applications. Parameters are given in Table 1.

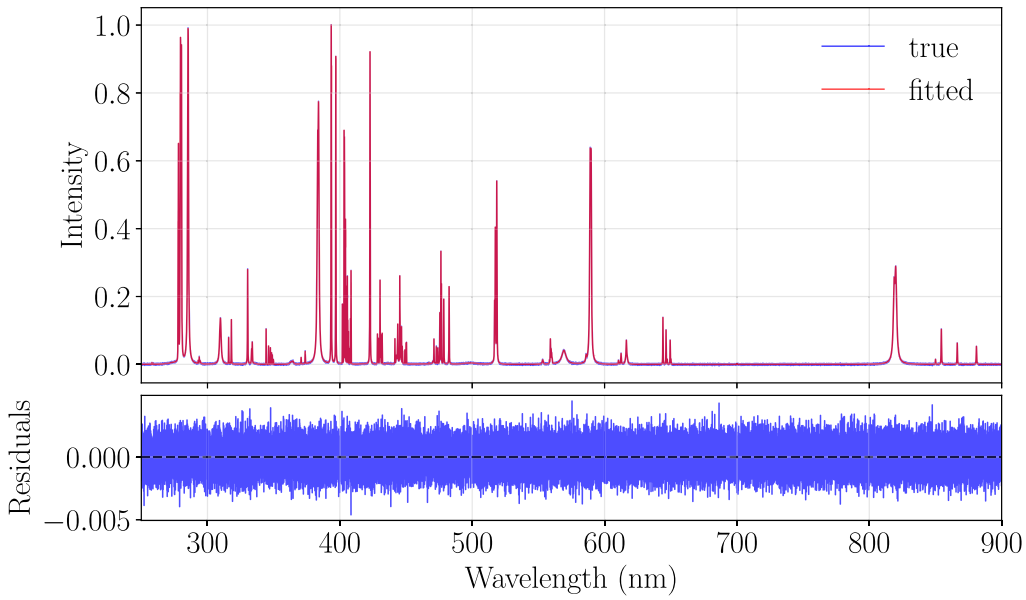

Fig. 2. Spectral fitting of a synthetic spectrum. The true parameters and the parameters for the initial guess are given in Table 1. The top panel shows the true and fitted spectrum over the whole wavelength range (normalized by their maxima). The bottom panel shows the residuals between true and fitted spectrum.

Table 1
Parameters for the fit of the synthetic spectrum. True parameters indicates the target values, initial values are where the optimization started, and fitted values are the final values after the optimization finished. We chose a relative tolerance of $\delta = 10^{-4}$ as the stopping criterion.

Parameter	True value	Initial value	Fitted value
T	8300 K	10 000 K	8299.7 K
n_e	2.3×10^{23} m^{-3}	1×10^{23} m^{-3}	2.3001×10^{23} m^{-3}
L	3 mm	1 mm	2.9981 mm
c_{Ca}	5%	25%	5.0026%
c_{Mg}	20%	25%	20.0095%
c_{Mn}	25%	25%	24.9996%
c_{Na}	50%	25%	49.9883%

2. Results

2.1. Fitting of simulated spectra

To evaluate the robustness of our method we have checked its internal consistency by fitting a simulated spectrum of four elements (Ca, Mg, Mn, Na). We used the plasma parameters and composition given in Table 1. We added Gaussian noise with a standard deviation of 2% of the maximum peak value to mimic experimental noise. Gaussian noise is an appropriate model for noise in CCD-based spectrometers. We then chose an initial guess that is reasonable for LIBS plasmas and agnostic with respect to the element concentrations, see Table 1. The results are shown in Fig. 3. The optimization converges within 39 function evaluations and reproduces the target spectrum, which is best illustrated by the residuals in Fig. 2 that are on the order of the added noise.

The parameter values after convergence are given in Table 1 and show that all parameters are recovered to within 0.1% of their true values. The L2 loss function combined with analytic normalization yields essentially exact parameter recovery on synthetic data at the noise level used here.

2.2. Convergence for different initial guesses

To gauge the soundness of our method we test how different initial guesses impact the fitting result. We used the same target spectrum as in Section 2.1. To ensure an even distribution of initial guesses over the parameter domain, we used latin hypercube sampling. We sampled 256 points in the following ranges:

$$T \in [5000\,\mathrm{K}, 50\,000\,\mathrm{K}] \tag{28}$$

$$n_e \in [10^{22}\,\mathrm{m}^{-3}, 10^{24}\,\mathrm{m}^{-3}] \tag{29}$$

$$L \in [10^{-5}\,\mathrm{m}, 10^{-2}\,\mathrm{m}] \tag{30}$$

$$c_i \in [0\%, 100\%] \quad i = \{\mathrm{Ca}, \mathrm{Mg}, \mathrm{Mn}, \mathrm{Na}\} \tag{31}$$

These ranges extend well beyond values that are typically expected, resulting in a thorough test of the robustness of the fitting scheme.

All 256 runs converge to the same global minimum with a median of 57 function evaluations (see Fig. 4). The lowest final loss across the successful runs is $f^\star = 6.5126 \times 10^{-2}$ and the largest is 6.5129×10^{-2}, so the spread of final loss values is below 5×10^{-5} relative to $f^\star$, which is well below the noise floor of the synthetic spectrum. Every plasma parameter is recovered to better than 0.12% of its true value, with median relative errors below 7×10^{-4} for T, n_e, L and the four element fractions.

A complementary view of the loss landscape, showing why almost every initial guess in this wide parameter box leads to the same global minimum, is given in Section 2.3.

2.3. Empirical loss-landscape survey

To complement the multi-start convergence study of Section 2.2 with a direct view of the loss landscape, we apply the diagnostic introduced in Section 1.2.5 to the LTE L2 loss. We drew $N = 200,000$ Latin hypercube samples on the same parameter box used in Section 2.2, drawn uniformly and with the sum-to-one constraint applied to the element fractions. All samples were evaluated in the internal coordinates of Section 1.2.4, and $(f, \|\nabla f\|)$ were computed via Enzyme reverse-mode autodifferentiation. This way we investigate the loss as the optimizer sees it, i.e. in the coordinates that are actually used for the L-BFGS updates. For $f^\star$ we used the loss at the known true parameters of the test case, which lies below the minimum loss of any sampled point.

Fig. 5 shows the resulting $(f - f^\star, \|\nabla f\|)$ in a density plot of the 200,000 samples, with the gradient norm at the true optimum, $\|\nabla f(\theta^\star)\| \approx 4.6 \times 10^{-2}$, overlaid as a dashed horizontal line. The

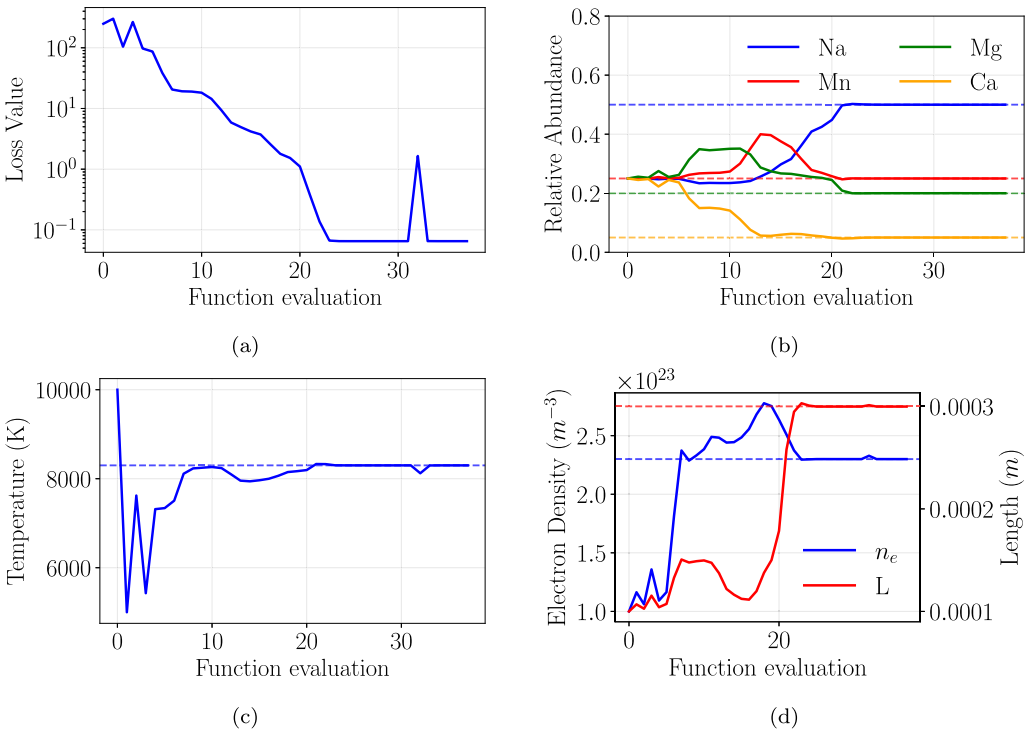

Fig. 3. Parameter evolution during fitting of a synthetic spectrum. For each function evaluation (a) shows the loss value, (b) shows the evolution of the mass concentrations for each element, (c) shows the evolution of the plasma temperature, (d) shows the evolution of the electron density and plasma length.

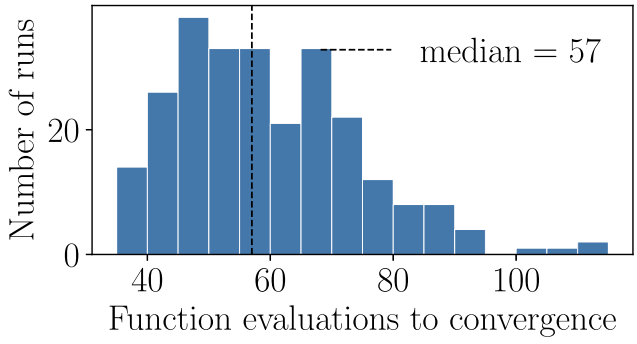

Fig. 4. Distribution of the number of function evaluations to convergence over the 256 multi-start runs of Section 2.2. The dashed line marks the median (57 evaluations). All 256 runs reach the same global minimum.

color scale indicates the density of samples in each hexagonal bin, with lighter colors corresponding to higher densities. The smallest gradient norm observed across all samples is 1.3, more than an order of magnitude above the optimum value, and the bulk of the samples sits several orders of magnitude higher still. No sample comes close to satisfying the stationarity condition $\nabla f = 0$: the sampled cloud is separated from the dashed reference line by a clear gap. Empirically, therefore, no second basin of attraction was detected anywhere in this wide, physically-plausible parameter box. Combined with the multi-start evidence in Section 2.2, this supports the conclusion that the global minimum is unique for this test case on the studied domain. We emphasize that this is empirical evidence only: the finite sample can neither prove global uniqueness nor rule out spurious basins narrower than the sampling resolution.

2.4. Comparison to other optimization schemes

To highlight the advantages of the gradient-based L-BFGS approach we compare it against six other optimization algorithms from the NLopt library [24], using the same target spectrum and initial guess as in Table 1. The algorithms span three categories (for details on the algorithms see [24] and references in the NLopt documentation):

Local algorithms. L-BFGS and Nelder–Mead are both local optimization algorithms that require an initial guess. They differ in that L-BFGS uses gradients to inform its search direction, while Nelder–Mead is a derivative-free method.

Global algorithms. DIRECT, DIRECT-L, and StoGO are global optimization algorithms that do not require an initial guess and instead search the entire parameter space.

Global algorithms with local sub-optimizations. Controlled Random Search (CRS) and Multi-Level Single-Linkage (MLSL) are global optimization algorithms that incorporate local optimization steps to refine candidate solutions. These methods are similar to our multiple initial guess approach, but they automate the selection of starting points and the local optimization process.

We restrict the number of function evaluations to 50.000 for all algorithms. The results are shown in Fig. 6.

Both local algorithms converge to the same final parameter values, given in Table 1. With reverse-mode AD via Enzyme, each L-BFGS step requires only a single function evaluation (one reverse pass for the full gradient), the same as each Nelder–Mead step. L-BFGS uses 28 function evaluations compared to Nelder–Mead's 367, a factor of roughly 13.

Among the global methods, only CRS reaches a loss comparable to the local algorithms, requiring on the order of 3×10^3 evaluations. CRS performs local optimization from randomly chosen starting points, which explains its success: the loss landscape around the optimum is smooth and well-suited for local search, so any starting point that falls within the basin of attraction converges quickly. MLSL, although it also incorporates local sub-optimizations, does not converge to the

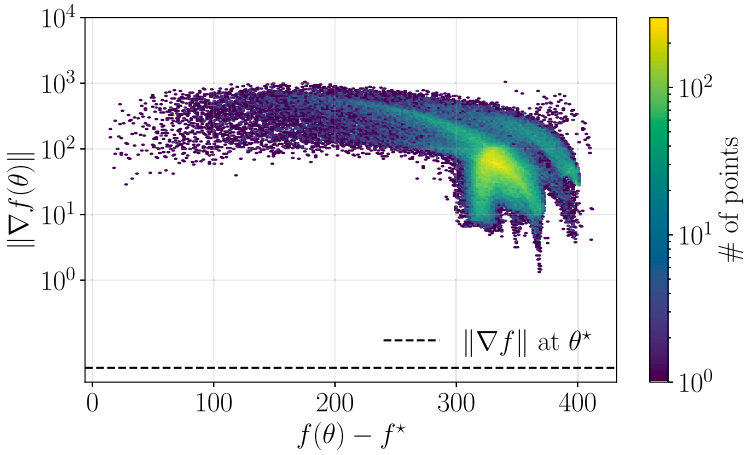

Fig. 5. Empirical loss-landscape survey for the LTE L2 loss. Hexbin density of $N = 200,000$ Latin hypercube samples on the wide physically-plausible parameter box of Section 2.2, evaluated in the internal coordinates of Section 1.2.4. The dashed horizontal line marks the gradient norm at the true optimum $\theta^\star$. All samples have gradient norms more than an order of magnitude above this reference and none approach the stationarity condition $\nabla f = 0$, providing empirical evidence that no additional basin of attraction was detected on this domain.

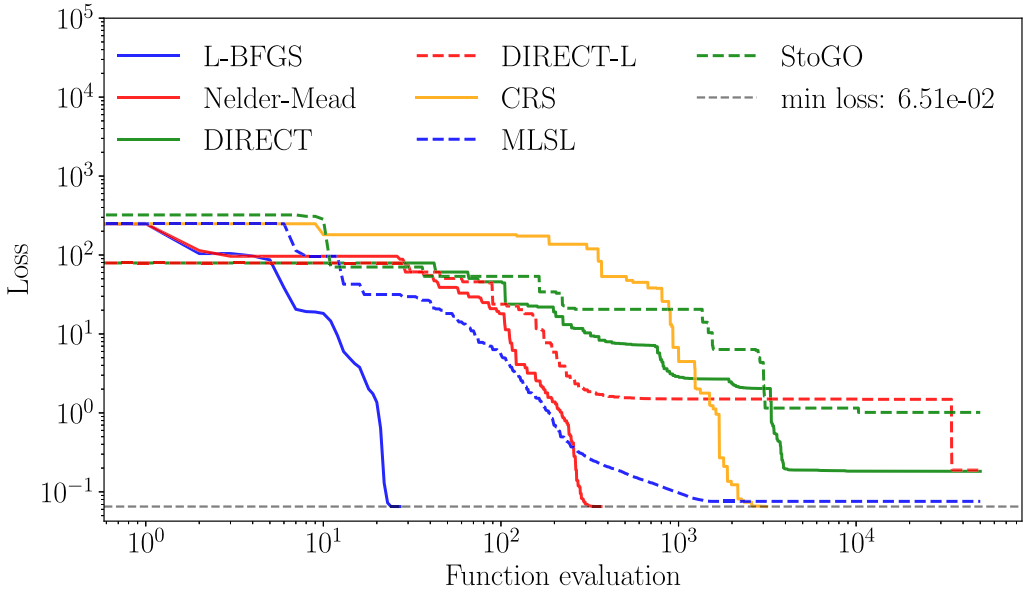

Fig. 6. Algorithm comparison. Loss as a function of function evaluation for seven optimization algorithms, using the same target spectrum and initial guess as in Table 1. To enhance visibility, we plot the best loss achieved up to each evaluation. The two local algorithms (L-BFGS, Nelder–Mead) converge rapidly, with L-BFGS reaching the minimum loss an order of magnitude faster than Nelder–Mead. Among the global methods, only CRS—which performs local optimization from random starting points—reaches a competitive loss within 50.000 evaluations. The purely partition-based schemes (DIRECT, DIRECT-L, StoGO) fail to converge within 50.000 evaluations.

global minimum within 50,000 evaluations, reaching a final loss approximately 16% above the global best. The purely partition-based schemes — DIRECT, DIRECT-L, and StoGO — do not converge within 50,000 evaluations and reach final loss values orders of magnitude higher than the local methods. This suggests that the loss landscape cannot be efficiently resolved by deterministic space-partitioning alone, consistent with the high evaluation counts reported for global solvers in previous work [21].

This highlights the importance of local and gradient-based optimization for efficient convergence.

2.5. Comparison of loss functions

Different loss functions induce different optimization landscapes and may therefore affect convergence speed. To quantify this, we compare three loss functions using the same synthetic spectrum and the same initial guess as in Table 1: the standard L2 loss $f_{\mathrm{L2}} = \sum_i (y_i - \hat{y}_i)^2$, the L1 loss $f_{\mathrm{L1}} = \sum_i |y_i - \hat{y}_i|$, and the log-domain L2 loss $f_{\log} = \sum_i [\ln(1+y_i) - \ln(1+\hat{y}_i)]^2$. All three loss functions are used with the optimized scale normalization described in Section 2.6.

The results are shown in Fig. 7. All three loss functions recover the plasma parameters to within 0.2% of their true values for all seven parameters (T, n_e, L and the four elemental concentrations). This indicates that the choice of loss function does not affect the accuracy of the final result for this test case. The differences lie in convergence speed: the log-domain L2 loss converges fastest (36 function evaluations), followed closely by the L2 loss (42 function evaluations) and L1 loss (43 function evaluations). We note that the computational cost per evaluation may differ slightly between loss functions (e.g. the log transform requires additional operations), but these differences are small compared to the dominant cost of the forward model evaluation.

The log-domain variant compresses the range of the spectral residuals, i.e. it gives more weight to weaker emission lines, which may

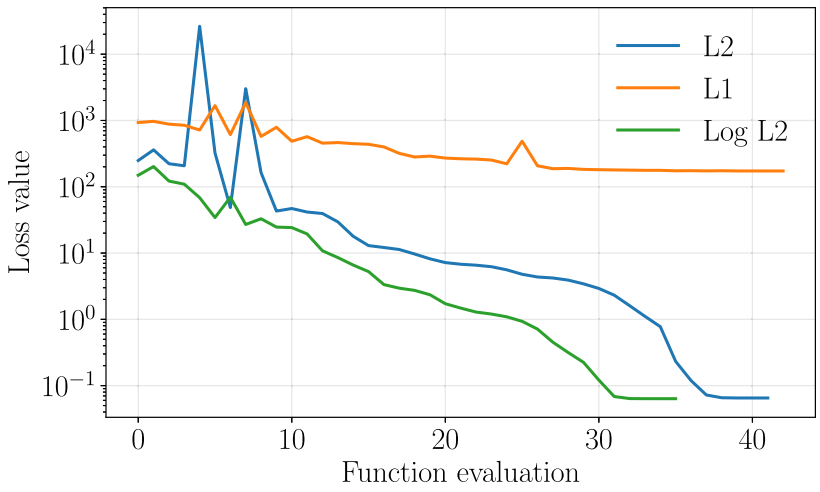

Fig. 7. Loss convergence for three loss functions when fitting the synthetic spectrum from Section 2.1 with L-BFGS. All runs use the same initial guess, convergence criterion ($\delta = 10^{-4}$), and optimized scale normalization. All three loss functions recover the plasma parameters to within 0.1% of their true values; differences are solely in convergence speed.

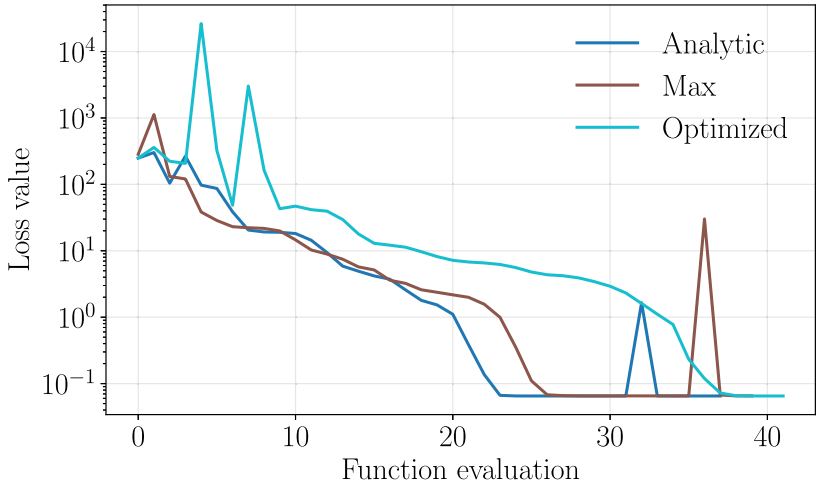

Fig. 8. Loss convergence for the three normalization strategies when fitting the synthetic spectrum from Section 2.1 with L2 loss. All three strategies recover the plasma parameters to within 0.1% of their true values; differences are solely in convergence speed. The final number of function evaluations for convergence are: 38 for analytic scale factor normalization, 39 for max normalization and 42 for optimized scale factor normalization.

explain its faster convergence in this case. We note that the differences between the loss functions are not large and that, as we show in the next section (Section 2.6), the choice of normalization strategy also has a significant impact on convergence speed. Therefore, we keep the L2 loss as our default choice for the remainder of this work, but we note that the log-domain variant may be advantageous due to its accounting of weaker spectral features.

2.6. Normalization strategy

In many LIBS applications, the recorded spectra are not radiometrically calibrated, while the physical plasma simulation yields the plasma emission in units of spectral radiance. By using normalization schemes during the fitting process, only a relative spectral response of the experimental data is required.

We implement three normalization strategies and compare their effect on convergence:

1. **Max normalization:** The simulated and experimental spectra are each independently divided by their respective maxima, $I_{\text{sim}}/\max(I_{\text{sim}})$ and $I_{\text{exp}}/\max(I_{\text{exp}})$. This is simple but normalizes each spectrum independently, which may distort relative peak ratios when the highest peak in the simulation does not coincide with the highest peak in the experiment. It also places a disproportionate emphasis on the wavelength bin with the highest peak, which may not be ideal if it is not well modeled.
2. **Optimized scale factor normalization:** In general we can treat the scale factor c as an additional free parameter and optimize it jointly with the plasma parameters by the gradient descent algorithm. The spectra are normalized as $y_{\text{exp}} = I_{\text{exp}}/\max(I_{\text{exp}})$ and $y_{\text{sim}} = I_{\text{sim}}/(c \cdot \max(I_{\text{exp}}))$, and the loss is evaluated on these normalized spectra.
3. **Analytic scale factor normalization:** In case of the L2 loss, the optimal scale factor can be computed analytically. The loss after normalization takes the form $\|I_{\text{sim}}/c - I_{\text{exp}}/\max(I_{\text{exp}})\|^2$, and minimizing with respect to c gives:

$$c_{\text{opt}} = \frac{\langle I_{\text{sim}}, I_{\text{sim}} \rangle}{\langle I_{\text{sim}}, I_{\text{exp}} \rangle}. \tag{32}$$

Then the experimental spectrum is normalized by its maximum, $y_{\text{exp}} = I_{\text{exp}}/\max(I_{\text{exp}})$, and the simulated spectrum is normalized by $y_{\text{sim}} = I_{\text{sim}}/(c_{\text{opt}} \cdot \max(I_{\text{exp}}))$. This is fully differentiable, making it compatible with gradient-based optimization.

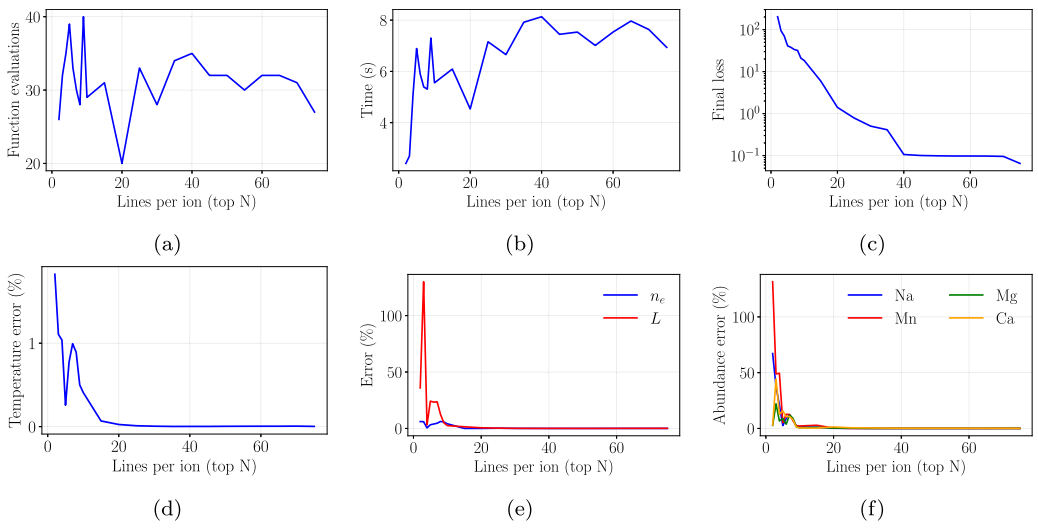

Fig. 9. Effect of the number of model lines on the fitting of a synthetic full spectrum. (a) Number of function evaluations to convergence, (b) wall-clock time, (c) final loss value, (d) temperature recovery error, (e) electron density and plasma length recovery error, (f) elemental abundance recovery error. All quantities are plotted as a function of the number of emission lines per ion as ranked by their intensity.

Fig. 8 compares the three strategies on the synthetic spectrum using the L2 loss. All three normalizations recover the plasma parameters to within 0.1% of their true values; the differences are solely in convergence speed. The analytic scale factor normalization converges fastest (38 function evaluations), since the scale factor is always set to its optimal value at each step, leaving the optimizer to focus entirely on the physical parameters. The max normalization (39 function evaluations) and the optimized scale factor normalization (42 function evaluations) are slightly slower; the max normalization independently normalizes both spectra and can therefore distort the relative peak heights, whereas the optimized scale factor normalization treats the scale factor as an additional optimization variable that the optimizer must converge alongside the physical parameters. For the loss function comparison in Section 2.5, we use the optimized scale factor normalization, as it is the most general strategy and does not assume that the analytic projection is exact. Throughout the remainder of this work the analytic scale factor normalization is used as the default, as it provides the fastest convergence on complete spectral models.

2.7. Scaling with the number of spectral lines

In practice, the number of emission lines included in the fitting model is limited by the availability of spectral parameters in the databases. To investigate how the number of modeled lines affects the optimization behavior, we generated a synthetic test spectrum using all available lines and then fitted it with models that include only the top-N strongest lines per ion, ranked by their computed peak emission coefficient from the full LTE model at the true plasma conditions (see Eq. (15)). This mimics the realistic scenario where the measured spectrum contains contributions from all emission lines, but the fitting model only considers a subset of them. We used the same synthetic composition and plasma parameters as in Section 2.1 and swept the number of lines per ion from 1 to 80. Due to the fact that not all ions have the same number of lines available, the total number of lines considered in the model varies between 13 (top 2 lines per ion) and 243 (top 80 lines per ion).

The results are shown in Fig. 9. All wall-clock times reported in this and subsequent sections were measured on a single core of an Intel i9-14900 K processor and are intended to illustrate relative scaling trends; absolute timings are implementation- and hardware-dependent. As the number of model lines increases, the number of function evaluations required for convergence first decreases as more and more lines help to constrain the parameter space (Fig. 9(a)). After a sufficient number of lines are included, the number of evaluations stabilizes, as does the wall-clock time per fit (Fig. 9(b)).

The final loss value decreases monotonically with the number of lines (Fig. 9(c)), reflecting the improved capacity of the model to reproduce the full experimental spectrum. This decrease is most pronounced when going from very few lines (< 10) to a moderate number ($\sim$40), after which the loss plateaus (notice the logarithmic scale in Fig. 9(c)), indicating diminishing returns from including additional lines.

The recovery accuracy of the plasma parameters shows a similar trend. The temperature error (Fig. 9(d)) drops below 1% once the model includes more than approximately 10 lines per ion, and below 0.2% for 20 or more lines per ion. The electron density and plasma length (Fig. 9(e)) follow a comparable improvement, with both parameters recovered to within a few percent for models with 20 or more lines per ion. The elemental abundance recovery (Fig. 9(f)) reveals that relatively few lines are needed for an accurate recovery. Already with 10 lines per ion, the elemental abundances are recovered to within 2% of their true values.

After around 40 lines per ion, the parameter recovery errors plateau, indicating that the model has sufficient complexity to capture the relevant spectral features for accurate parameter inference.

2.8. Effect of internal wavelength resolution

In order to accurately capture the shape of emission line profiles during spectral fitting, the model spectrum must be evaluated on a grid that is fine enough to resolve the narrowest features. A natural scale for this is provided by the full width at half maximum (FWHM) of the narrowest Voigt profile in the model, which depends on the plasma parameters and the spectral lines considered. Rather than requiring to manually specify a wavelength grid resolution, we automatically generate a fine wavelength grid using an adjustable parameter, the number of points per FWHM of the narrowest Voigt line profile, denoted N_{FWHM}. Given N_{FWHM} and the minimum Voigt FWHM $\Delta\lambda_{\mathrm{min}}$ across all lines in the model, the internal fine grid spacing is

$$\delta\lambda = \frac{\Delta\lambda_{\mathrm{min}}}{N_{\mathrm{FWHM}}}. \tag{33}$$

The model spectrum is computed on this fine grid and then interpolated back onto the coarser grid matching the experimental spectrum. This ensures that line profiles are always adequately sampled regardless of the output grid resolution.

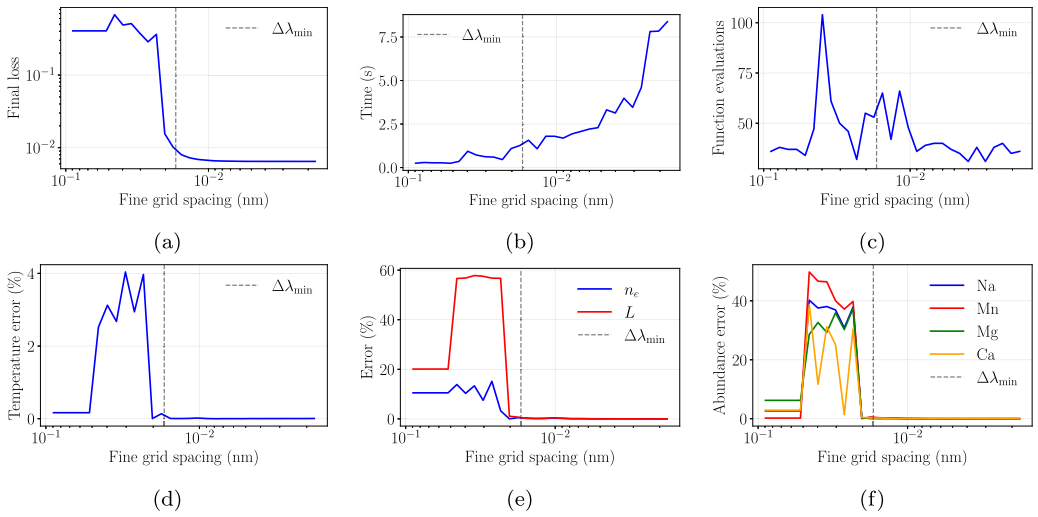

Fig. 10. Effect of the internal fine grid density on fitting quality. (a) Final loss value, (b) wall-clock time, (c) number of function evaluations to convergence, (d) temperature recovery error, (e) electron density and plasma length recovery error, (f) elemental abundance recovery error. All quantities are plotted as a function of the fine grid spacing $\delta\lambda$. The dashed vertical line marks the minimum Voigt FWHM $\Delta\lambda_{\min} = 0.017\,\text{nm}$, which provides a natural scale for resolving line profiles. Accurate parameter recovery is achieved once $\delta\lambda$ falls below this FWHM, i.e. $N_{\text{FWHM}} \geq 1$.

To investigate the effect of N_{FWHM} on fitting quality, we generated a synthetic ground-truth spectrum at high internal resolution ($N_{\text{FWHM}} = 10$) using the same plasma parameters as in Section 2.1 and an output grid with $\Delta\lambda = 0.1\,\text{nm}$. We then performed the fitting at 20 different values of N_{FWHM} logarithmically spaced between 0.1 and 5, corresponding to fine grid spacings from approximately $0.09\,\text{nm}$ down to $0.002\,\text{nm}$.

The results are summarized in Fig. 10. For very coarse internal grids $\delta\lambda > 0.02\,\text{nm}$, the line profiles are undersampled, leading to elevated loss values and erroneous parameter recovery (Figs. 10(a) and 10(d) to 10(f)). The optimizer still converges, but to incorrect parameters, with temperature errors exceeding 15% and electron density errors above 40% in the worst cases.

As the wavelength resolution approaches the FWHM of the narrowest line ($\Delta\lambda_{\min} = 0.017\,\text{nm}$), all parameter recovery errors decrease sharply and parameters are recovered to within a few percent of their true values. Further decreasing of the internal grid spacing does not significantly improve the recovery accuracy but does increase computational cost. The wall-clock time (Fig. 10(b)) increases with finer wavelength resolutions, while the number of function evaluations to convergence (Fig. 10(c)) remains relatively constant. The wall-clock time (Fig. 10(b)) increases approximately linearly with the number of grid points, from a few seconds at $0.2\,\text{nm}$ to over $1000\,\text{s}$ at $0.0005\,\text{nm}$. The number of function evaluations (Fig. 10(c)), however, remains roughly constant across all resolutions, confirming that the gradient-based optimizer converges in a comparable number of steps regardless of grid size.

This demonstrates that an internal fine grid wavelength resolution on the order of the narrowest line FWHM is sufficient for accurate parameter recovery. For all optimizations shown in the rest of this work, we use a default value of $N_{\text{FWHM}} = 5$, which provides a safety margin while keeping computational cost manageable.

2.9. Robustness to surplus elements in the model

In practice, the elemental composition of a sample is often unknown a priori, and one must include a superset of candidate elements in the fitting model. This raises the question of whether including elements that are not actually present in the sample degrades the quality of the fit or the recovery of the true parameters.

To test this, we repeated the synthetic fit from Section 2.1, but extended the guess model with four elements that regularly occur in geologic samples (Si, Al, Fe, K). The initial guess assigned equal concentration fractions (1/8) to all eight elements, and all other settings (initial plasma parameters, loss function, normalization, convergence criterion) were identical to the baseline experiment.

The optimizer converges in 60 evaluations and recovers the plasma parameters to within 0.01% (temperature), 0.01% (electron density) and 0.1% (plasma length) of their true values. All four surplus elements are driven to numerically negligible concentrations (below 10^{-10}), while the true elemental fractions (Na, Mn, Mg, Ca) are recovered to within 0.1% of their true values. The parameter evolution is shown in Fig. 11: the surplus elements are rapidly eliminated within the first ~ 20 evaluations, after which the optimization proceeds as in the baseline case.

For the tested element and plasma parameters, the gradient-based optimizer is robust to the inclusion of surplus elements in this test case: the additional degrees of freedom do not appear to degrade convergence or parameter recovery. We note that while we observe successful convergence with equal-fraction initialization, the presence of additional local minima cannot be ruled out in general and would require a more systematic exploration of the loss landscape. Critically, the surplus elements do not absorb any significant concentration, confirming that the spectral fitting is selective: elements whose emission lines are not present in the data are suppressed by the optimizer.

3. Discussion

3.1. Uniqueness and recommended workflow

A local optimization scheme does not by itself guarantee that the recovered minimum is the global one. We therefore tested uniqueness empirically in two complementary ways: (i) the loss-landscape survey of Section 2.3, in which $N = 200,000$ samples spanning a wide physically-plausible parameter box all show gradient norms more than an order of magnitude above the value at the true optimum, so that no additional near-stationary point is detected, and (ii) the multi-start convergence study of Section 2.2, where all of the 256 runs consistently recover the same parameters. Together, both diagnostics provide empirical evidence that only one global minimum is present for our chosen loss function on the studied parameter domain. It is therefore sufficient to run a single optimization from a reasonable initial guess to recover the true parameters. Additionally, stationarity at the returned solution can be verified by inspecting $\|\nabla f\|$. We note that the controlled experiments in Sections 2.4 to 2.8 use a single, reasonable

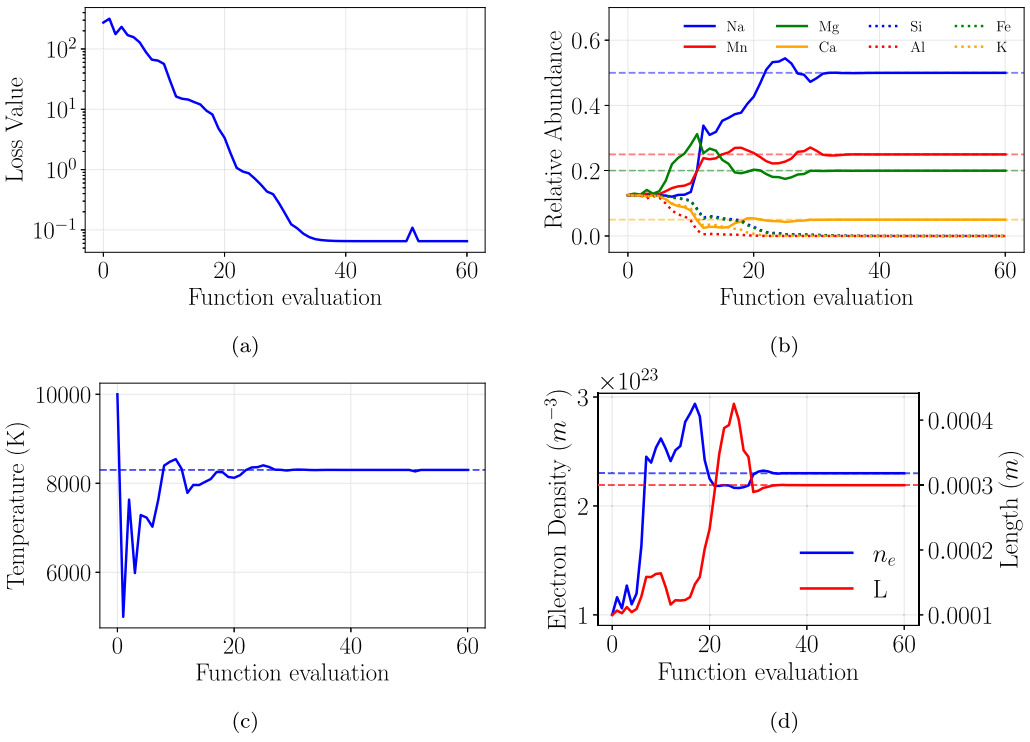

Fig. 11. Parameter evolution during fitting of the synthetic spectrum with an extended model containing four surplus elements (Si, Al, Fe, K) not present in the true sample. (a) Loss convergence, (b) element concentration evolution (solid lines: true elements; dotted lines: surplus elements), (c) temperature evolution, (d) electron density and plasma length evolution. All surplus elements are driven to zero; all true parameters are recovered accurately.

initial guess to isolate the effect of a single variable at a time. The internal logarithmic and softmax parametrization of Section 1.2.4 are essential for this behavior: in earlier experiments ran on the physical parameters directly, we observed that the optimizer would get stuck near the parameter space boundaries or along the slow direction of a banana-shaped valley.

Let us make the following observations about the used plasma model and the fitting scheme: (1) LIBS spectra contain highly redundant information on plasma properties in the sense that a single element may show multiple peaks and therefore it is not possible to change the height or width of a single peak without impacting the others. This is particularly true for major elements; minor and trace elements may exhibit fewer lines and thus less redundancy. (2) The plasma parameters are interdependent. Increasing the concentration of one element also decreases the concentration of all other elements, changing the plasma temperature impacts all peaks widths and heights. This puts further constraints on the possible solutions in the parameters space. (3) Even in the case of models that are out of scope for this publication, such as the two zone model [13–15], one can impose constraints on the parameters through additional equations, further restricting the space of feasible solutions.

We also note an inherent identifiability limitation: in spectral ranges without optically thick lines, the plasma length L acts primarily as a scale factor and becomes degenerate with the normalization constant, so that L cannot be independently identified from normalized spectra. In practice, at least one optically thick line — as is typical for the broad wavelength ranges used here — suffices to make L identifiable [16]; for narrow ranges or spectra with exclusively optically thin lines, L should be determined by an independent measurement such as absolute intensity calibration or a complementary measurement like plasma imaging.

3.2. Comparison to other fitting methods

Fitting methods used in the literature can be broadly categorized into global and local optimization schemes. Amongst the local schemes, they may be further divided into gradient-based and derivative-free methods. We provide an overview of common CF-LIBS fitting methods in Table 2 and discuss the differences between these schemes and ours below.

Our method is, to the best of our knowledge, the only gradient-based local optimization scheme for CF-LIBS. The only other local optimization scheme employed so far is the iterative method presented in [16] and differs in the following ways: (1) The iterative method presupposes knowledge of which elements are in the target, as well as which lines are optically thin. This is needed to calculate the temperature from line ratios. (2) The iterative method uses a single peak from each element to infer the element concentration iteratively, whereas our method uses the whole spectrum to infer all parameters simultaneously. (3) The iterative method is derivative-free, so it does not use gradient information to update the parameters.

All global optimization schemes have in common that they search the whole parameter space for the global minimum and use the whole spectrum to infer all parameters simultaneously. They mostly differ in the way they search the parameter space and which loss functions they use. Our method, being a local optimization scheme, scales much better with the number of parameters than these global schemes. This is because global schemes need to search the whole parameter space, which scales exponentially with the number of parameters (this is also called the curse of dimensionality). The benefit of this is that they supposedly always find the global minimum. In practice however, it is not possible to prove that this is the case and, as we have shown in Section 2.2, in a local scheme, one may use multiple initial guesses to ensure convergence to the true parameters, resulting in the same effect.

3.3. Convergence

The empirical loss-landscape survey of Section 2.3 and the multi-start convergence study of Section 2.2 jointly indicate that, on the physically-plausible parameter range studied here, every investigated initial guess leads to the same global minimum of the LTE loss. A

Table 2
An overview of common CF-LIBS fitting schemes. GD = Gradient Descent.

Global	Local (derivative-free)	Local (gradient-based)
"Monte Carlo LIBS" [11,21]	Iterative approach [16]	LIBS+GD (this work)
Simulated Annealing [14,18]		
"NLOpt Global methods" [15]		

single optimization run from a reasonable initial guess is therefore likely sufficient within the studied parameter domain. The internal parametrization (Section 1.2.4) was important in order to achieve this: without the logarithmic and softmax scaling, the same loss would exhibit a banana-shape, hurting local convergence regardless of the optimizer used.

3.4. Gradient descent methods

We have chosen the L-BFGS algorithm as the gradient descent method for our fitting scheme, as it performed best in our tests. Since we use the NLOpt library [24], it is possible to easily swap out L-BFGS for other optimizers provided by the library (as shown in Section 3.2).

3.5. Use with experimental data

As with any other CF-LIBS method, the quality of the fit for experimental data is not only dependent on the fitting method, but, more crucially, on the availability of accurate spectral parameters. In this work we chose to only consider lines for which Stark broadening parameters are reported in the Stark-B database. However, not only does this significantly reduce the number of lines considered, but the parameters reported there stem from theoretical calculations and are provided without uncertainties. Inaccuracies in the Stark parameters is one possible explanation for deviating peak widths and shifts observed in experimental spectra.

For the other spectral parameters, we used the NIST ASD database [19], where we considered all lines that have matching data in the Stark-B database. NIST reports uncertainties for the transition probabilities, which may be up to 50% for the accuracy category D and are a further possible explanation for residuals observed in experimental spectra.

Whether or not it is better to exclude lines with missing or uncertain parameters — as opposed to including them and even imputing missing data with sensible default values — is an open question and out of scope for this publication.

4. Summary and conclusion

This work presents a local gradient-based optimization scheme for calibration-free LIBS spectral fitting. We tested our method on synthetic spectra of a mixture of four elements (Ca, Mg, Mn, Na) and demonstrated accurate recovery of the plasma parameters. A multi-start convergence study and a direct sampling of the loss landscape both indicate empirically that, on the physically-plausible parameter range we studied, every investigated initial guess leads to the same global minimum of the LTE loss; a single optimization run from a reasonable initial guess therefore likely recovers the true plasma parameters within this parameter domain. The method scales favorably with the number of spectral lines included in the model, with the computational cost growing roughly linearly and the number of function evaluations remaining nearly constant (see Section 2.7). The internal fine grid density, controlled by the number of samples per FWHM of the narrowest Voigt line, needs to be at least one sample per FWHM for accurate parameter recovery, while the default of five samples per FWHM provides a practical safety margin (see Section 2.8).

This method represents a promising alternative to the global optimization schemes used in the literature, as it scales better with the number of plasma parameters and is easily parallelizable. This is especially important for field LIBS analysis, where (1) the plasma models needed to accurately describe the laser-induced plasma are more complex than the homogeneous LTE model and (2) a large number of spectra may need to be analyzed in a reasonable time frame.

This method is agnostic to the plasma model chosen as long as it is differentiable. To model in-situ plasma, one may want to consider a plasma with two (or more) zones, which this method is open to. In that case, additional physical constraints are needed to ensure that a unique global minimum is found as it was shown that already a two-zone model can have multiple minima [14].

CRediT authorship contribution statement

Christoph H. Egerland: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Conceptualization. **Kristin Rammelkamp:** Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization. **Elise Clavé:** Writing – review & editing, Validation. **Ana Lomashvili:** Writing – review & editing, Validation. **Peder Bagge Hansen:** Writing – review & editing, Software, Conceptualization. **Fabian Seel:** Writing – review & editing. **Susanne Schröder:** Writing – review & editing. **Heinz-Wilhelm Hübers:** Supervision, Funding acquisition.

Declaration of competing interest

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

Acknowledgments

The authors would like to thank Peter Jung for fruitful discussions on gradient methods. C.H.E. is thankful for advice on software publication from Carina Haupt.

Appendix A. Proof for formula ionic population fraction

We give a short proof of Eq. (5) and note it shares the same form as the one found in the supplement of [17].

Let us drop the subindex k and rewrite the problem as follows:

$$\beta_{z+1} = \Omega_z \beta_z \tag{A.1}$$

$$\sum_{z=0}^{Z_k} \beta_z = 1 \tag{A.2}$$

where Ω_z is the right hand side of the Saha equation Eq. (6) divided by n_e. By writing out the first few steps explicitly one can deduce a general formula for β_z:

$$\beta_1 = \Omega_0 \beta_0 \tag{A.3}$$

$$\beta_2 = \Omega_1 \beta_1 = \Omega_1 \Omega_0 \beta_0 \tag{A.4}$$

$$\beta_3 = \Omega_2 \beta_2 = \Omega_2 \Omega_1 \Omega_0 \beta_0 \tag{A.5}$$

$$\beta_z = \beta_0 \prod_{i=0}^{z-1} \Omega_i \tag{A.6}$$

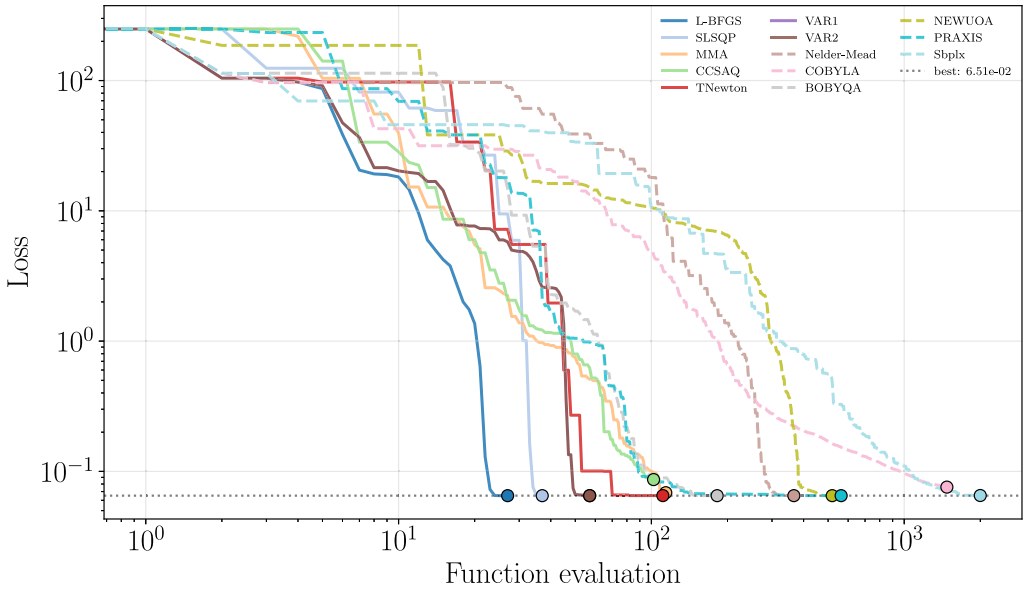

Fig. B.12. Convergence of local optimization algorithms on the spectral fitting problem. Solid lines denote gradient-based methods; dashed lines denote derivative-free methods. Filled circles mark the final function evaluation of each algorithm. The dotted horizontal line indicates the overall best loss.

Inserting this into the normalization condition

$$\sum_{z=0}^{Z_k} \beta_z = \beta_0 \sum_{z=0}^{Z_k} \prod_{i=0}^{z-1} \Omega_i = 1 \tag{A.7}$$

yields an expression for β_0

$$\beta_0 = \frac{1}{\sum_{z=0}^{Z_k} \prod_{i=0}^{z-1} \Omega_i} \tag{A.8}$$

which, inserted back into Eq. (A.6) results in our expression for β_z:

$$\beta_z = \frac{\prod_{i=0}^{z-1} \Omega_i}{\sum_{j=0}^{Z_k} \prod_{i=0}^{j-1} \Omega_i} \tag{A.9}$$

Appendix B. Local optimization algorithm comparison

To supplement the algorithm comparison in Fig. 6, we benchmark a broader set of 13 local optimization algorithms from the NLopt library [24] on the same spectral fitting problem (6 free parameters: T, n_e, L, and three independent compositional logits u_1, u_2, u_3). Gradient-based methods compute the full gradient via Enzyme reverse-mode AD in a single pass per optimization step; derivative-free methods likewise require only one function evaluation per step. All algorithms start from the same initial point.

Fig. B.12 shows the best-so-far loss as a function of total function evaluations, with filled circles marking each algorithm's termination point. Table B.3 lists the total function evaluations, wall-clock time and best loss achieved by each method.

Most algorithms converge to losses within a factor of two of the global best (6.51×10^{-2}). Among gradient-based methods, L-BFGS, SLSQP, VAR1 and VAR2 reach the global minimum with the fewest evaluations (28–58) and the shortest wall times (6.9 s–12.8 s), making them the most efficient choices. TNewton, MMA and CCSAQ converge in 103–115 evaluations but to slightly elevated losses (6.52–8.65×10^{-2}). Among derivative-free methods, BOBYQA achieves the global minimum in 183 evaluations. NEWUOA and PRAXIS converge to the global minimum in 520 and 564 evaluations respectively; Nelder–Mead reaches it in 367 evaluations. COBYLA terminates at a loss of 7.6×10^{-2}, slightly above the global best, while Sbplx terminates at its evaluation limit of 2000 at a loss of 6.53×10^{-2}.

Table B.3
Summary of local optimization algorithms. Function evaluations count the number of optimization steps (for gradient-based methods, each step internally computes the full gradient via reverse-mode AD in a single pass).

Algorithm	Type	Evals	Wall time (s)	Best loss
L-BFGS	gradient	28	6.9	6.51×10^{-2}
SLSQP	gradient	38	6.8	6.51×10^{-2}
MMA	gradient	115	25.8	6.86×10^{-2}
CCSAQ	gradient	103	24.1	8.65×10^{-2}
TNewton	gradient	112	24.9	6.52×10^{-2}
VAR1	gradient	58	12.4	6.51×10^{-2}
VAR2	gradient	58	12.8	6.51×10^{-2}
Nelder–Mead	deriv-free	367	16.4	6.51×10^{-2}
COBYLA	deriv-free	1476	69.2	7.58×10^{-2}
BOBYQA	deriv-free	183	9.1	6.51×10^{-2}
NEWUOA	deriv-free	520	24.8	6.51×10^{-2}
PRAXIS	deriv-free	564	27.6	6.51×10^{-2}
Sbplx	deriv-free	2000	92.7	6.53×10^{-2}

Data availability

The full simulation and fitting code has been published as open-source software [27] and is available at https://github.com/DLR-WR/libskit. The specific code and data used to generate the results in this publication are available at https://github.com/DLR-WR/libs_gd_paper.

References

[1] David A. Cremers, Leon J. Radziemski, Handbook of Laser-Induced Breakdown Spectroscopy, first ed., Wiley, 2013, http://dx.doi.org/10.1002/9781118567371.
[2] David W. Hahn, Nicoló Omenetto, Laser-induced breakdown spectroscopy (LIBS), Part I: Review of basic diagnostics and plasma—Particle interactions: Still-challenging issues within the analytical plasma community, Appl. Spectrosc. 64 (12) (2010) 335A–336A, http://dx.doi.org/10.1366/000370210793561691.
[3] David W. Hahn, Nicoló Omenetto, Laser-Induced Breakdown Spectroscopy (LIBS), Part II: Review of instrumental and methodological approaches to material analysis and applications to different fields, Appl. Spectrosc. 66 (4) (2012) 347–419, http://dx.doi.org/10.1366/11-06574.
[4] Ryan B. Anderson, Olivier Forni, Agnes Cousin, Roger C. Wiens, Samuel M. Clegg, Jens Frydenvang, Travis S.J. Gabriel, Ann Ollila, Susanne Schröder, Olivier Beyssac, Erin Gibbons, David S. Vogt, Elise Clavé, Jose-Antonio Manrique, Carey Legett, Paolo Pilleri, Raymond T. Newell, Joseph Sarrao, Sylvestre

Maurice, Gorka Arana, Karim Benzerara, Pernelle Bernardi, Sylvain Bernard, Bruno Bousquet, Adrian J. Brown, César Alvarez-Llamas, Baptiste Chide, Edward Cloutis, Jade Comellas, Stephanie Connell, Erwin Dehouck, Dorothea M. Delapp, Ari Essunfeld, Cecile Fabre, Thierry Fouchet, Cristina Garcia-Florentino, Laura García-Gómez, Patrick Gasda, Olivier Gasnault, Elisabeth M. Hausrath, Nina L. Lanza, Javier Laserna, Jeremie Lasue, Guillermo Lopez, Juan Manuel Madariaga, Lucia Mandon, Nicolas Mangold, Pierre-Yves Meslin, Anthony E. Nelson, Horton Newsom, Adriana L. Reyes-Newell, Scott Robinson, Fernando Rull, Shiv Sharma, Justin I. Simon, Pablo Sobron, Imanol Torre Fernandez, Arya Udry, Dawn Venhaus, Scott M. McLennan, Richard V. Morris, Bethany Ehlmann, Post-landing major element quantification using SuperCam laser induced breakdown spectroscopy, Spectrochim. Acta Part B: At. Spectrosc. 188 (2022) 106347, http://dx.doi.org/10.1016/j.sab.2021.106347.

[5] Roger C. Wiens, Sylvestre Maurice, Bruce Barraclough, Muriel Saccoccio, Walter C. Barkley, James F. Bell, Steve Bender, John Bernardin, Diana Blaney, Jennifer Blank, Marc Bouyé, Nathan Bridges, Nathan Bultman, Phillippe Caïs, Robert C. Clanton, Benton Clark, Samuel Clegg, Agnes Cousin, David Cremers, Alain Cros, Lauren DeFlores, Dorothea Delapp, Robert Dingler, Claude D'Uston, M. Darby Dyar, Tom Elliott, Don Enemark, Cecile Fabre, Mike Flores, Olivier Forni, Olivier Gasnault, Thomas Hale, Charles Hays, Ken Herkenhoff, Ed Kan, Laurel Kirkland, Driss Kouach, David Landis, Yves Langevin, Nina Lanza, Frank LaRocca, Jeremie Lasue, Joseph Latino, Daniel Limonadi, Chris Lindensmith, Cynthia Little, Nicolas Mangold, Gerard Manhes, Patrick Mauchien, Christopher McKay, Ed Miller, Joe Mooney, Richard V. Morris, Leland Morrison, Tony Nelson, Horton Newsom, Ann Ollila, Melanie Ott, Laurent Pares, René Perez, Franck Poitrasson, Cheryl Provost, Joseph W. Reiter, Tom Roberts, Frank Romero, Violaine Sautter, Steven Salazar, John J. Simmonds, Ralph Stiglich, Steven Storms, Nicolas Striebig, Jean-Jacques Thocaven, Tanner Trujillo, Mike Ulibarri, David Vaniman, Noah Warner, Rob Waterbury, Robert Whitaker, James Witt, Belinda Wong-Swanson, The ChemCam Instrument Suite on the Mars Science Laboratory (MSL) Rover: Body Unit and Combined System Tests, Space Sci. Rev. 170 (1) (2012) 167–227, http://dx.doi.org/10.1007/s11214-012-9902-4.

[6] S. Maurice, R.C. Wiens, M. Saccoccio, B. Barraclough, O. Gasnault, O. Forni, N. Mangold, D. Baratoux, S. Bender, G. Berger, J. Bernardin, M. Berthé, N. Bridges, D. Blaney, M. Bouyé, P. Caïs, B. Clark, S. Clegg, A. Cousin, D. Cremers, A. Cros, L. DeFlores, C. Derycke, B. Dingler, G. Dromart, B. Dubois, M. Dupieux, E. Durand, L. d'Uston, C. Fabre, B. Faure, A. Gaboriaud, T. Gharsa, K. Herkenhoff, E. Kan, L. Kirkland, D. Kouach, J.-L. Lacour, Y. Langevin, J. Lasue, S. Le Mouélic, M. Lescure, E. Lewin, D. Limonadi, G. Manhès, P. Mauchien, C. McKay, P.-Y. Meslin, Y. Michel, E. Miller, H.E. Newsom, G. Orttner, A. Paillet, L. Parès, Y. Parot, R. Pérez, P. Pinet, F. Poitrasson, B. Quertier, B. Sallé, C. Sotin, V. Sautter, H. Séran, J.J. Simmonds, J.-B. Sirven, R. Stiglich, N. Striebig, J.-J. Thocaven, M.J. Toplis, D. Vaniman, The ChemCam Instrument Suite on the Mars Science Laboratory (MSL) Rover: Science Objectives and Mast Unit Description, Space Sci. Rev. 170 (1) (2012) 95–166, http://dx.doi.org/10.1007/s11214-012-9912-2.

[7] Roger C. Wiens, Sylvestre Maurice, Scott H. Robinson, Anthony E. Nelson, Philippe Cais, Pernelle Bernardi, Raymond T. Newell, Sam Clegg, Shiv K. Sharma, Steven Storms, Jonathan Deming, Darrel Beckman, Ann M. Ollila, Olivier Gasnault, Ryan B. Anderson, Yves André, S. Michael Angel, Gorka Arana, Elizabeth Auden, Pierre Beck, Joseph Becker, Karim Benzerara, Sylvain Bernard, Olivier Beyssac, Louis Borges, Bruno Bousquet, Kerry Boyd, Michael Caffrey, Jeffrey Carlson, Kepa Castro, Jorden Celis, Baptiste Chide, Kevin Clark, Edward Cloutis, Elizabeth C. Cordoba, Agnes Cousin, Magdalena Dale, Lauren Deflores, Dorothea Delapp, Muriel Deleuze, Matthew Dirmyer, Christophe Donny, Gilles Dromart, M. George Duran, Miles Egan, Joan Ervin, Cecile Fabre, Amaury Fau, Woodward Fischer, Olivier Forni, Thierry Fouchet, Reuben Fresquez, Jens Frydenvang, Denine Gasway, Ivair Gontijo, John Grotzinger, Xavier Jacob, Sophie Jacquinod, Jeffrey R. Johnson, Roberta A. Klisiewicz, James Lake, Nina Lanza, Javier Laserna, Jeremie Lasue, Stéphane Le Mouélic, Carey Legett, Richard Leveille, Eric Lewin, Guillermo Lopez-Reyes, Ralph Lorenz, Eric Lorigny, Steven P. Love, Briana Lucero, Juan Manuel Madariaga, Morten Madsen, Soren Madsen, Nicolas Mangold, Jose Antonio Manrique, J.P. Martinez, Jesus Martinez-Frias, Kevin P. McCabe, Timothy H. McConnochie, Justin M. McGlown, Scott M. McLennan, Noureddine Melikechi, Pierre-Yves Meslin, John M. Michel, David Mimoun, Anupam Misra, Gilles Montagnac, Franck Montmessin, Valerie Mousset, Naomi Murdoch, Horton Newsom, Logan A. Ott, Zachary R. Ousnamer, Laurent Pares, Yann Parot, Rafal Pawluczyk, C. Glen Peterson, Paolo Pilleri, Patrick Pinet, Gabriel Pont, Francois Poulet, Cheryl Provost, Benjamin Quertier, Heather Quinn, William Rapin, Jean-Michel Reess, Amy H. Regan, Adriana L. Reyes-Newell, Philip J. Romano, Clement Royer, Fernando Rull, Benigno Sandoval, Joseph H. Sarrao, Violaine Sautter, Marcel J. Schoppers, Susanne Schröder, Daniel Seitz, Terra Shepherd, Pablo Sobron, Bruno Dubois, Vishnu Sridhar, Michael J. Toplis, Imanol Torre-Fdez, Ian A. Trettel, Mark Underwood, Andres Valdez, Jacob Valdez, Dawn Venhaus, Peter Willis, The SuperCam Instrument Suite on the NASA Mars 2020 Rover: Body Unit and Combined System Tests, Space Sci. Rev. 217 (1) (2020) 4, http://dx.doi.org/10.1007/s11214-020-00777-5.

[8] S. Maurice, R.C. Wiens, P. Bernardi, P. Caïs, S. Robinson, T. Nelson, O. Gasnault, J.-M. Reess, M. Deleuze, F. Rull, J.-A. Manrique, S. Abbaki, R.B. Anderson, Y. André, S.M. Angel, G. Arana, T. Battault, P. Beck, K. Benzerara, S. Bernard, J.-P. Berthias, O. Beyssac, M. Bonafous, B. Bousquet, M. Boutillier, A. Cadu, K. Castro, F. Chapron, B. Chide, K. Clark, E. Clavé, S. Clegg, E. Cloutis, C. Collin, E.C. Cordoba, A. Cousin, J.-C. Dameury, W. D'Anna, Y. Daydou, A. Debus, L. Deflores, E. Dehouck, D. Delapp, G. De Los Santos, C. Donny, A. Doressoundiram, G. Dromart, B. Dubois, A. Dufour, M. Dupieux, M. Egan, J. Ervin, C. Fabre, A. Fau, W. Fischer, O. Forni, T. Fouchet, J. Frydenvang, S. Gauffre, M. Gauthier, V. Gharakanian, O. Gilard, I. Gontijo, R. Gonzalez, D. Granena, J. Grotzinger, R. Hassen-Khodja, M. Heim, Y. Hello, G. Hervet, O. Humeau, X. Jacob, S. Jacquinod, J.R. Johnson, D. Kouach, G. Lacombe, N. Lanza, L. Lapauw, J. Laserna, J. Lasue, L. Le Deit, S. Le Mouélic, E. Le Comte, Q.-M. Lee, C. Legett, R. Leveille, E. Lewin, C. Leyrat, G. Lopez-Reyes, R. Lorenz, B. Lucero, J.M. Madariaga, S. Madsen, M. Madsen, N. Mangold, F. Manni, J.-F. Mariscal, J. Martinez-Frias, K. Mathieu, R. Mathon, K.P. McCabe, T. McConnochie, S.M. McLennan, J. Mekki, N. Melikechi, P.-Y. Meslin, Y. Micheau, Y. Michel, J.M. Michel, D. Mimoun, A. Misra, G. Montagnac, C. Montaron, F. Montmessin, J. Moros, V. Mousset, Y. Morizet, N. Murdoch, R.T. Newell, H. Newsom, N. Nguyen Tuong, A.M. Ollila, G. Orttner, L. Oudda, L. Pares, J. Parisot, Y. Parot, R. Pérez, D. Pheav, L. Picot, P. Pilleri, C. Pilorget, P. Pinet, G. Pont, F. Poulet, C. Quantin-Nataf, B. Quertier, D. Rambaud, W. Rapin, P. Romano, L. Roucayrol, C. Royer, M. Ruellan, B.F. Sandoval, V. Sautter, M.J. Schoppers, S. Schröder, H.-C. Seran, S.K. Sharma, P. Sobron, M. Sodki, A. Sournac, V. Sridhar, D. Standarovsky, S. Storms, N. Striebig, M. Tatat, M. Toplis, I. Torre-Fdez, N. Toulemont, C. Velasco, M. Veneranda, D. Venhaus, C. Virmontois, M. Viso, P. Willis, K.W. Wong, The SuperCam Instrument Suite on the Mars 2020 Rover: Science Objectives and Mast-Unit Description, Space Sci. Rev. 217 (3) (2021) 47, http://dx.doi.org/10.1007/s11214-021-00807-w.

[9] Roger C. Wiens, Agnes Cousin, Samuel M. Clegg, Olivier Gasnault, Zhaopeng Chen, Sylvestre Maurice, Rong Shu, Geochemistry of Mars with Laser-Induced Breakdown Spectroscopy (LIBS): ChemCam, SuperCam, and MarSCoDe, Minerals 15 (8) (2025) 882, http://dx.doi.org/10.3390/min15080882.

[10] Igor Gornushkin, Calibration-Free Quantitative Analysis, in: Gábor Galbács (Ed.), Laser-Induced Breakdown Spectroscopy in Biological, Forensic and Materials Sciences, Springer International Publishing, Cham, 2022, pp. 67–99, http://dx.doi.org/10.1007/978-3-031-14502-5_3.

[11] A. Demidov, S. Eschlböck-Fuchs, A.Ya. Kazakov, I.B. Gornushkin, P.J. Kolmhofer, J.D. Pedarnig, N. Huber, J. Heitz, T. Schmid, R. Rössler, U. Panne, Monte Carlo standardless approach for laser induced breakdown spectroscopy based on massive parallel graphic processing unit computing, Spectrochim. Acta Part B: At. Spectrosc. 125 (2016) 97–102, http://dx.doi.org/10.1016/j.sab.2016.09.016.

[12] Jörg Hermann, Christoph Gerhard, Miloš Burger, Valentin Craciun, Frédéric Pelascini, Progress in calibration-free laser-induced breakdown spectroscopy, Spectrochim. Acta Part B: At. Spectrosc. 200 (2023) 106595, http://dx.doi.org/10.1016/j.sab.2022.106595.

[13] Jörg Hermann, Antonella Lorusso, Alessio Perrone, Francesco Strafella, Christophe Dutouquet, Béatrice Torralba, Simulation of emission spectra from nonuniform reactive laser-induced plasmas, Phys. Rev. E 92 (5) (2015) 053103, http://dx.doi.org/10.1103/PhysRevE.92.053103.

[14] P.B. Hansen, S. Schröder, S. Kubitza, K. Rammelkamp, D.S. Vogt, H.-W. Hübers, Modeling of time-resolved LIBS spectra obtained in Martian atmospheric conditions with a stationary plasma approach, Spectrochim. Acta Part B: At. Spectrosc. 178 (2021) 106115, http://dx.doi.org/10.1016/j.sab.2021.106115.

[15] S.M. Zaitsev, I.N. Krylov, A.M. Popov, T.A. Labutin, Two-Zone Model of Laser-Induced Plasma, J. Appl. Spectrosc. 90 (6) (2024) 1183–1189, http://dx.doi.org/10.1007/s10812-024-01651-w.

[16] C. Gerhard, J. Hermann, L. Mercadier, L. Loewenthal, E. Axente, C.R. Luculescu, T. Sarnet, M. Sentis, W. Vi öl, Quantitative analyses of glass via laser-induced breakdown spectroscopy in argon, Spectrochim. Acta Part B: At. Spectrosc. 101 (2014) 32–45, http://dx.doi.org/10.1016/j.sab.2014.07.014.

[17] Sergey M. Zaytsev, Andrey M. Popov, Timur A. Labutin, Stationary model of laser-induced plasma: Critical evaluation and applications, Spectrochim. Acta Part B: At. Spectrosc. 158 (2019) 105632, http://dx.doi.org/10.1016/j.sab.2019.06.002.

[18] Peder Bagge Hansen, Modeling of LIBS Spectra Obtained in Martian Atmospheric Conditions (Ph.D. thesis), Humboldt-Universität zu Berlin, 2022, http://dx.doi.org/10.18452/25303.

[19] A. Kramida, Yu. Ralchenko, J. Reader, NIST ASD Team, 2024, NIST Atomic Spectra Database (ver. 5.12), [Online]. Available: https://physics.nist.gov/asd [2016, January 31]. National Institute of Standards and Technology, Gaithersburg, MD.

[20] S. Sahal-Bréchot, M.S. Dimitrijević, N. Moreau, N. Ben Nessib, The STARK-B database VAMDC node: A repository for spectral line broadening and shifts due to collisions with charged particles, Phys. Scr. 90 (5) (2015) 054008, http://dx.doi.org/10.1088/0031-8949/90/5/054008.

[21] Igor B. Gornushkin, Tobias Völker, Intrinsic performance of Monte Carlo Calibration-Free Algorithm for laser-induced breakdown spectroscopy, Sensors 22 (19) (2022) 7149, http://dx.doi.org/10.3390/s22197149.

[22] Aya Taleb, Chao Shen, David Mory, Katarzyna Cieślik, Sven Merk, Muhammad R. Aziz, Anna Paola Caricato, Christoph Gerhard, Frédéric Pelascini, Jörg Hermann, Echelle spectrometer calibration by means of laser plasma, Spectrochim. Acta Part B: At. Spectrosc. 178 (2021) 106144, http://dx.doi.org/10.1016/j.sab.2021.106144.

[23] Jörg Hermann, David Grojo, Emanuel Axente, Christoph Gerhard, Miloš Burger, Valentin Craciun, Ideal radiation source for plasma spectroscopy generated by laser ablation, Phys. Rev. E 96 (5) (2017) 053210, http://dx.doi.org/10.1103/PhysRevE.96.053210.

[24] Steven G. Johnson, The NLopt nonlinear-optimization package, 2007, https://github.com/stevengj/nlopt.

[25] Dong C. Liu, Jorge Nocedal, On the limited memory BFGS method for large scale optimization, Math. Program. 45 (1989) 503–528, http://dx.doi.org/10.1007/bf01589116.

[26] William S. Moses, Valentin Churavy, Ludger Paehler, Jan Hückelheim, Sri Hari Krishna Narayanan, Michel Schanen, Johannes Doerfert, Reverse-mode automatic differentiation and optimization of GPU kernels via Enzyme, in: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '21, ACM, 2021, http://dx.doi.org/10.1145/3458817.3476165.

[27] Christoph H. Egerland, LIBSkit: A software tool for the simulation and analysis of laser-induced breakdown spectroscopy (LIBS) spectra, 2026, http://dx.doi.org/10.5281/zenodo.21788186.