High-convenience Interferometric Measurement of the Spectrum of a Picosecond Industrial-grade Nd:YAG Laser Using a PID Linear Drive

Kristian Cvecek[*1,2] and Michael Schmidt[1,2]

[1]*Institute of Photonic Technologies, Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany*
[2]*Erlangen Graduate School in Advanced Optical Technologies (SAOT), Friedrich-Alexander-Universität Erlangen-Nürnberg, Germany*
[*]*Corresponding author's e-mail: kristian.cvecek@lpt.uni-erlangen.de*

We present a high-resolution spectroscopic characterization of the frequency-doubled output (532 nm) of an industrial-grade picosecond Nd:YAG laser. A Michelson interferometer with a PID-controlled linear stage was used to provide the path delay change during taking the interferogram, despite its non-uniform motion which in general prohibits direct Fourier-based spectral reconstruction. We introduce a correction scheme combining inverse FFT artefact suppression and a matched-filter approach to recover the true spectral content with sub-nanometer accuracy. The most probable emission peak was determined at 531.844 nm with a FWHM of 0.0437 nm, in good agreement with literature values for Nd:YAG fluorescence peaks. Our method enables precise spectral measurements, without resorting to phase-stabilized scanning methods or calibrated spectrometers.
DOI: 10.2961/ilmn.2025.03.2005
Keywords: interferometric spectroscopy, frequency recovery from non-uniform sampling, ultra-short pulsed lasers, Michelson interferometer, PID-controlled stage

1. Introduction

Ultra-short pulsed (USP) lasers exhibit many highly interesting properties that make them highly sought for in material processing. Of course, materials processing is just one branch of possible applications of these lasers. Another highly interesting field is, among others, sensing and microscopy, where USP lasers can be used e.g. for 2-photon microscopy, Raman microscopy and similar spectroscopy techniques that are utilizing Stimulated Raman Scattering (SRS) [1] or Coherent Anti-stokes Raman Scattering (CARS) [2]. Certain variants of USP-laser-based SRS or CARS methods require often additional specialized equipment such as narrow-band notch filters to block the wavelength of the pump pulse in order to be able to measure the Raman response of the material close to the pump wavelength. Since the blocking bandwidth of the filter is typically on the order of 10 cm^{-1} (0.28 nm at 532 nm) or below, the filter must be designed and manufactured specifically to the laser's wavelength and width.

Unfortunately, laser laboratories equipped with USP lasers for materials processing do not know these values (peak wavelength and wavelength/frequency bandwidth) as typically the laser manufacturers do not provide the spectral data either at all or at least not in the required accuracy. This is problematic especially when new interdisciplinary projects are to be initiated and preliminary results are necessary to demonstrate the proof of principle. In such cases, a fast and inexpensive experimental setup becomes necessary. In this work we describe high-precision measurements of the spectrum of a frequency doubled Nd:YAG laser running at 1064 nm (Coherent HyperRapid NXT) providing pulses with a duration typically between 9 and 15 ps (FWHM).

To obtain the spectrum, we use a Michelson interferometer with an unstabilized, proportional-integral-derivatively (PID) controlled linear translation stage and a camera. In contrast to earlier approaches (e.g. [3,4]) relying on spectrally resolved detection or phase-stabilized optical delay scans to reconstruct the full spectral phase and amplitude. Moreover, our system is operated without any reference interferometer or wavelength-calibrated spectrometer. In our case, the compromise is that any sensitivity towards the phase is discarded and we obtain only the spectral envelope. In spite of such non-optimal measurement conditions, we nevertheless attempt to recover the relevant spectral information with sub-nanometer precision.

Our idea is to utilize several aspects of the interferometer's trace as a kind "across-the-envelope approximation" to obtain this precision, i.e. based on plausibility and probability arguments. To this end, we follow a multi-layered approach. First, we transform the interferogram data by a Fast-Fourier-Transform (FFT) to obtain the spectrum. Since it is known that a PID-controlled linear stage does not move linearly, the non-uniform sampling of the data will generate several artefact-peaks in the spectrum that have no correspondence whatsoever to the wavelength of the laser, as can be easily checked by even a "low-resolution" optical spectrum analyzer. We identify the unphysical artifact regions and transform only these back (not the signal peak) into the delay-domain by an inverse-FFT, to obtain a correction for the original interferometer data.

After applying the correction we process the resulting signal with the so-called matched-filter approach to obtain the strength of individual frequency-components without regard to whether they are in- or out-of-phase in the corrected data. Although the matched-filter approach is commonly used in Fourier-domain OCT [5,6], in which spectrally resolved signals are transformed into the time-delay or spatial delay domain (time/spatial domain), we apply here the

matched-filter principle within the time/spatial domain itself: for each synthetic wavelength, we calculate its "coherence degree" w.r.t. the real signal, see Section 3. Importantly, any unexpected phase shifts that might occur during the experiment for even an ideally monochromatic wave will not affect the identification of that wavelength's total contribution to the overall spectrum.

This provides several sharp peaks in the spectral region of interest. Although these peaks are also caused by the non-uniform stage motion and data noise, they lie too close to the sought wavelength to be categorically ruled out by the optical spectrometer. Therefore, we evaluate their "fitness" to represent most of the interferogram's data by examining their match to a fitted spectral envelope generated from the whole amplitude envelope of the interferogram data, i.e. the envelope which corresponds to the interference contrast. The best match, in terms of several quantitative measures, is then assumed to be the sought spectral distribution, while the others are assumed to be only artefacts.

The manuscript is structured as follows. In section 2 we discuss the experimental setup, the way how the interferogram data were obtained and point out the underlying issues. In Section 3 we describe our method to evaluate the data. Section 4 provides and discusses the results.

2. Experimental setup and data acquisition

2.1 The setup

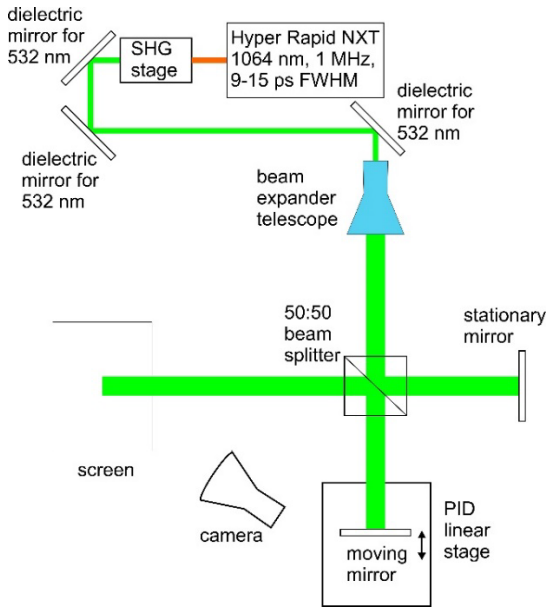

Fig. 1 Schematic setup of the experiment.

The schematic of the experimental setup is shown in Fig. 1. The laser source was a Coherent Hyper Rapid NXT, an industrial-grade picosecond fiber laser. It operates at a wavelength of 1064 nm and delivers pulses, according to the manufacturer with a typical duration between 9 and 15 ps. The beam was fed into a self-assembled nonlinear conversion stage to generate second harmonic radiation at 532 nm. To determine the pulse duration of the SHG pulses we conducted a sum-frequency-generation (SFG) autocorrelation using the fundamental and the SHG pulses. To this end, the SHG beam and the fundamental beam were collinearly focussed into a nonlinear crystal to generate 355 nm. The path length of the SHG beam was delayed by a variable delay line and the resulting power was measured as a function of the delay. The resulting autocorrelation trace had a full-width at half-maximum (FWHM) of 11.5 ps, resulting in a FWHM pulse duration of 8.9 ps for the fundamental, assuming a hyperbolic secant pulse shape. Assuming no chirp and no group velocity mismatch the corresponding FWHM pulse duration of the SHG pulses has to be consequently 6.1 ps [7].

While the laser is capable of delivering up to 100 W average power, we chose the input power of the fundamental to be such that the laser beam at 532 nm was around 50 mW, at a pulse repetition rate of 1 MHz. These low pulse energies (50 nJ) were chosen in order to be able to neglect as far as possible any nonlinear spectral broadening or shifting effects that migh be caused during the pulse propagation by nonlinear effects such as self-phase modulation etc.

In order to filter out the fundamental wavelength to the best of our possibilities, several dielectric mirrors with high-reflectivity coatings for 532 nm were implemented to reflect the beam after the SHG stage and guide it into the Michelson interferometer. Before entering the Michelson interferometer, the diameter of the collimated laser beam was increased to 15 mm using a telescope.

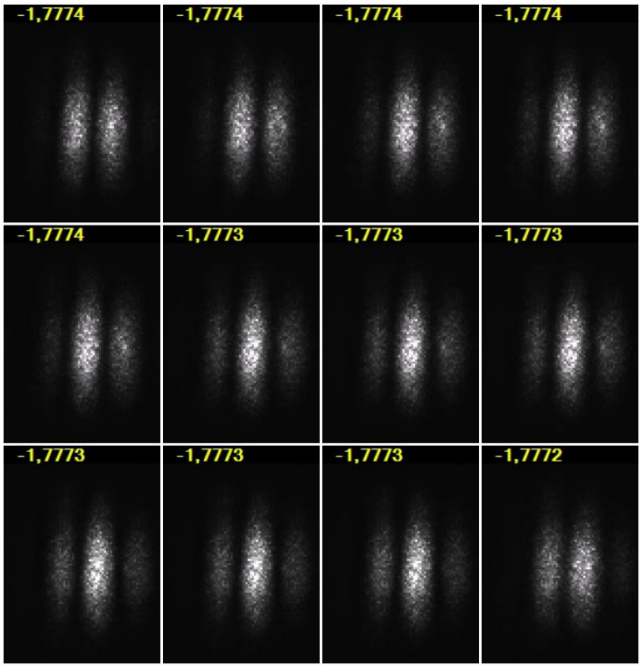

Fig. 2 A sequence of successive camera frames showing the interference pattern at the screen and the corresponding position readout by the stage's control software.

Inside the interferometer the mirror responsible for generating the time delay was placed onto a PID-controlled linear stage. Both beams were aligned in such a way that at the screen, which was monitored by a camera running at 30 fps, a linear fringe pattern with well-resolvable fringes became apparent in the regions of high interference contrast, see Fig. 2. This was done to obtain an additional spatial signal relating to the actually induced time delay by the moving mirror. Furthermore, care was taken to adjust the sensitivity of the camera to not exceed its dynamic range in the vicinity of high interference contrast in order to prevent the observed interferometer's signal from saturation in the camera.

The PID controlled linear drive stage was programmed to move in a rectilinearly uniform motion at a speed of 0.5 µm/s. To obtain the data on the stage's current position and at the same time record the current interferogram at the screen we displayed the motion control program and the camera stream side-by-side on the computer screen and used screen capture software to record these data synchronously. While this approach had the benefit of a well-synchronised stage position's display w.r.t. the interferogram's fringe pattern on the screen, the disadvantage was that the screenshot of the stage position needed to be run through OCR for each frame to obtain machine-readable data for the stage position. A short sequence of the stage position data (as image) and the corresponding camera frame is shown in Fig. 2.

2.2 Data in the ideal case

In principle, it should be possible to calculate the spectrum just from the amplitude evolution of a single pixel, ideally in the middle of the fringe pattern in Fig. 2, across all frames by means of a Fourier-transformation. Of course, the signal should be corrected, i.e. the background amplitude should be subtracted from each interferogram to obtain a signal with a zero average value prior to the Fourier-transform in order to prevent the DC-part of the spectrogram to be larger than the signal peak we want to measure. A good starting point for the background correction is to take an interferogram from a camera frame that is as far away as possible from the frames with high interference contrast, i.e. ideally zero interference contrast, and subtract this image from all other camera frames. The resulting amplitude will provide a very high visibility of the fringes in regions of high interference contrast, which is maximum when the path length difference between both arms of the interferometer is zero.

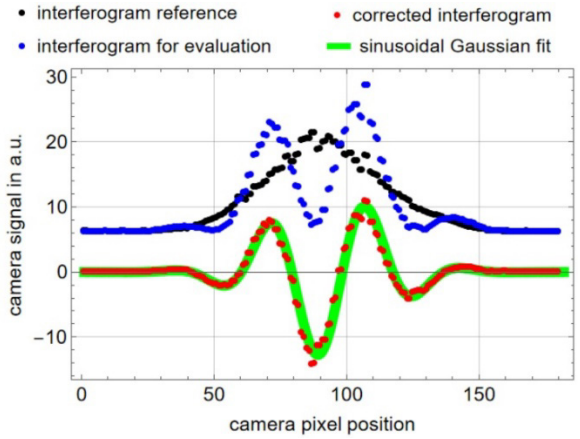

Fig. 3 Exemplary data for calculating interferogram parameters from the spatial fringe pattern. The to-be-evaluated interferogram (blue) is background-corrected by subtracting the reference run of the interferogram at a mirror delay that gives close to zero interference contrast (black). The result, an interferogram run with zero mean (red), is fitted by a sinusoidally modulated Gaussian (solid green line).

As the camera noise always compounds the amplitude readout, we decided to reduce the noise by averaging. To this end, we adjusted the two partial beams inside the interferometer in such a way that the resulting interferogram displayed straight-line fringes. This allowed us to sum up the signal contributions along the fringes and "collapse" the actual two-dimensional interferogram on the camera sensor to just a one-dimensional array with a width of about 180 pxl for each frame, see Fig. 3.

This provides us effectively with a denoised interferogram, because the read-out-noise of the pixels is randomly distributed around a mean value – adding all pixel values up along the fringe direction makes the pixel noise decrease. A second benefit is that this dimensional collapse of the interferogram reduces also the computational complexity later on. In fact, in the case of no camera and no phase noise, no vibrations and a perfectly uniform motion of the moving mirror, even a single pixel would be sufficient (corrected for the background) to compute the spectrum of the laser via a Fourier transformation. However, as soon as camera noise (even if denoised), non-equidistant motion of the moving mirror or mechanical vibrations between the mirrors etc. are present during the experiment, just a simple Fourier transform of the raw signal's amplitude will not yield a sufficiently precise spectrum (see Fig. 7 for illustration).

The only option for improving the precision lies in the fact that also other sources of information on the interferogram, such as the stage's position read-out can be used to refine or correct the acquired amplitude data. To do so, it becomes necessary to be able to compare the induced path delay of the stage (e.g. in units of the carrier wavelength) to the accumulated phase of the amplitude's signal. Hence, we must first separate the interferometer's signal, i.e. the instantaneous amplitude, into its envelope and its phase.

2.3 Phase and envelope recovery

One direct approach to separate an oscillatory signal into its envelope and its phase is to use the Hilbert transformation, see e.g. [8]. The Hilbert transform provides the so-called analytic signal to an oscillatory data, i.e. it infers from an experimentally available, oscillatory signal also the corresponding (not experimentally available) imaginary part, so that the instantaneous amplitude and phase of the signal can be calculated for each frame.

However, to perform the Hilbert transform requires having already an initial estimate of the envelope's local amplitude as well as its phase. We obtain these estimates by low-pass-filtering the raw, background-corrected data for the phase and by applying a localized trend approximation with subsequent low-pass-filtering of the magnitude, i.e. the absolute value, of the background-corrected data. These filtering steps are necessary because the Hilbert transform expects the signal to be approximately narrowband and locally sinusoidal, whereas our experimental signal deviates significantly from these assumptions, especially in regions with low interference contrast. Similar to the simple Fourier transform discussed in Section 2.2, the Hilbert transform can be applied to just only one pixel of the one-dimensional data. We chose here the pixel in the 72nd column, as it is far above the background noise even at rather low interference contrasts.

The above process of applying the Hilbert transformation is illustrated in Fig. 4. Obviously, the estimated envelope does not differ strongly from the envelope calculated by the Hilbert transform, on the one hand. On the other hand, the phase noise is much smaller in the now recovered phase, making the phase run to be much more plausible to match

with the expected average translational speed of the linear stage (0.5 µm/s) and the camera's frame rate of 30 fps. Indeed, these hardware parameters predict 15.96 video frames to complete one full oscillation of the electromagnetic (EM) field for a wavelength of 532 nm. In contrast, in Fig. 4, it is possible to recognize 12 "slow" oscillations (and several very fast ones which can be neglected) of the phase within 200 successive camera frames, which gives around 16.67 for one full EM field cycle.

However, while the Hilbert transform provides an amplitude envelope and corresponding phase for every frame, its application involves certain necessary pre-processing steps. Besides this, as is evident from the remaining noise in the recovered phase in Fig. 4, the Hilbert transform should be considered more as a method of last resort when dealing with signals that are strongly affected by noise.

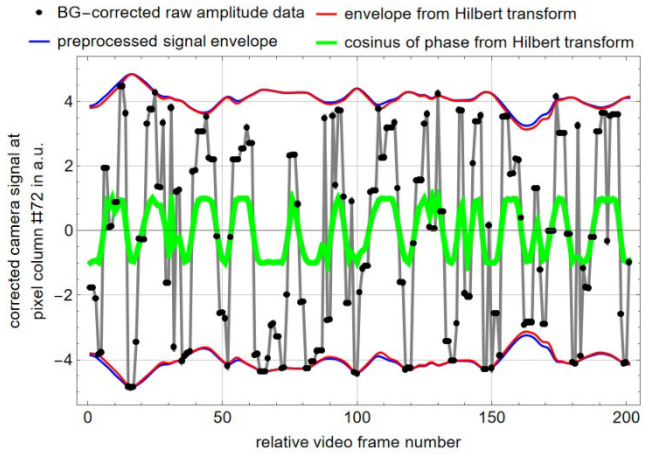

Fig. 4 Illustration of the phase and envelope determination by the Hilbert transform. The original data (black dots) are used to estimate the local amplitude the envelope (blue). After conducting the Hilbert transform one obtains the envelope (red) and the phase, of which the cosine is displayed (green), for a better comparison to the raw data (black dots/gray lines).

For a more reliable phase and amplitude recovery it would help to consider the whole run of the one-dimensional signal at the camera, see Fig. 3, instead of just a single pixel, to help distinguish whether the phase increases or decreases during the stage's motion. This is because if all that is available is just the signal from a single column or pixel, then it is especially the maxima and minima of the oscillatory signal that give rise to an ambiguity of the mirror's direction of motion – the mirror moving forward or backward from the maximum reduces in both cases the detected signal. The mirror's direction of motion can't be inferred from such data only. Nevertheless, the ambiguity of the mirror's direction of motion can be resolved when the whole interferogram's spatial structure is taken into account.

To do so, we fitted a sinusoidally modulated Gaussian (green solid line in Fig. 3) on the background corrected signal (red dots in Fig. 3) for each video frame. As the sinusoidally modulated Gaussian is an analytic function, it is easy to obtain the phase shift together with the direction of motion between two subsequent frames as long as the moving mirror did not "jump" further than one wavelength or 2π. In addition, this analytic function provides also maximum value of the spatial envelope, even if the phase shift is such that the spatial interferogram does not reach the maximum value. Unfortunately, due to noise issues and reduced interference contrast at larger delays, it was not possible to reliably fit the sinusoidal Gaussian. For all frames where this fit was not possible we used the data obtained from the Hilbert transform.

2.4 Phase unwrap

The phase data obtained in Section 2.3 are wrapped between 0 and 2π as both phase recovery algorithms (Hilbert transform and the fit function) are only able to determine the modulus of the remainder of the actual phase divided by 2π. To be able to infer any additional information from e.g. the stage's position, it is necessary to obtain the accumulated phase, i.e. the phase run where also e.g. the mirror's direction of motion is taken into account. Therefore, a phase unwrapping procedure must be performed. As there exist several possible methods for phase unwrapping we chose the phase unwrapping via max-flow/min-cut (PUMA), described in [9], and modified to our one-dimensional problem. The result, plotted against the path delay of the stage, is shown in Fig. 5.

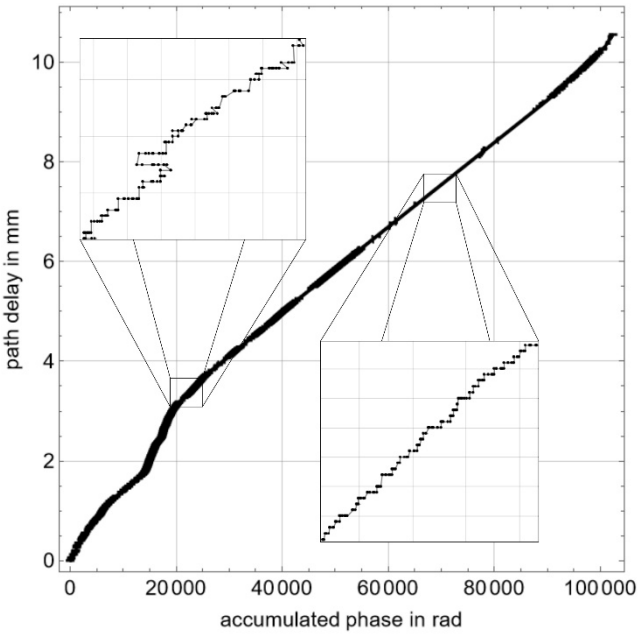

Fig. 5 The path delay provided by the stage (derived from the stage's position read-out) plotted against the accumulated (unwrapped) phase.

This way the linearity of the stage's motion w.r.t. the accumulated phase or vice versa can be estimated. As can be seen, the result is not a linear line as should be expected, but follows a somewhat random path instead. When zooming closer in (insets in Fig. 5), the phase seems to move even back and forth between consecutive frames. Even in the most linear section of the curve (right inset) the stage seems to move rather erratically to a new position while the phase increases. Nevertheless, this observation is explainable by a continuous (but not uniform) stage motion while the stage's position read-out is rounded in steps of 100 nm by the stage's control software.

A closer analysis of the linear portion of the signal reveals that the stage's position indeed jitters throughout its

motion around an "average" position. However, as the stage is controlled by the PID technique, this behavior is indeed expected, see Fig. 6. As apparent, the stage is for certain video frame intervals ahead of the accumulated phase and then falls again behind the path delay calculated from the accumulated phase for the assumed wavelength of 532 nm. Note that setting the assumed wavelength to a different value would lead to an average, non-zero slope of the deviation. In addition, Fig. 6 evidences significant noise, which might be caused by the phase jitter, the non-uniform position or its read-out (or both) of the stage. At this point, it is not possible to distinguish between the individual noise contributions to the deviation jitter.

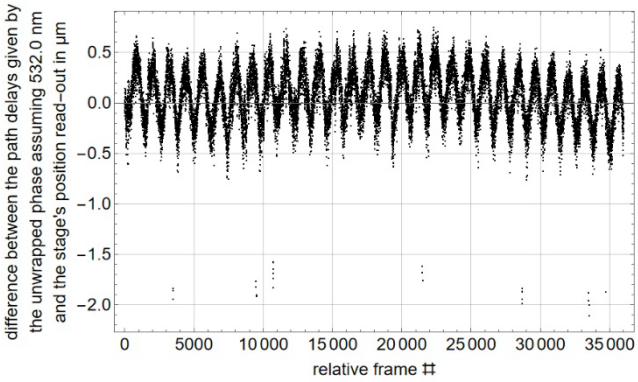

Fig. 6 The deviation of the path delay given by the stage's position read-out to a virtual path delay position estimated from the accumulated phase for a wavelength of 532.0 nm.

3. Determination of the spectrum

As has been shown in Section 2, the obtained interferometer signals, respective the accumulated phase but also the stage's position exhibit strong jitter and phase jumps and it is not possible to distinguish between them without additional hardware. The spectrum obtained by applying the FFT to these data exhibits clearly artefact-peaks besides the sought wavelength, which can be distinguished from the sought wavelengt by a simple optical spectrometer. This is shown in Fig. 7.

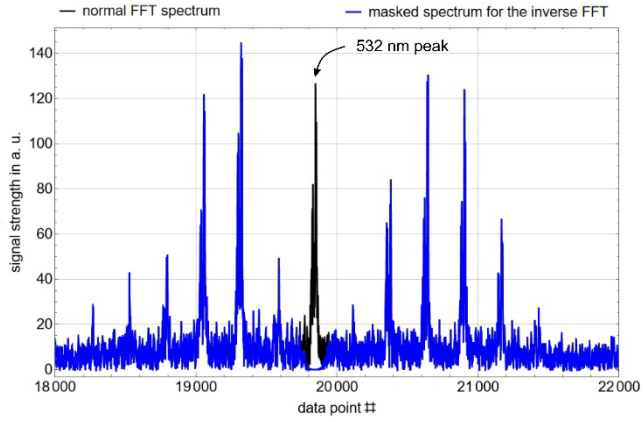

Fig. 7 The spectral distribution of the raw interferometer signal calculated by the FFT (black). Spectrum with only the artefact peaks for calculating the correction signal by inverse-FFT (blue). Note, that the blue curve has outside the 532 nm peak runs identically as the black curve, and thus the black curve is hidden behind the blue one.

Suppressing the 532 nm peak data in the complex Fourier-domain on both sides, (Fig. 7 shows only an excerpt of the FFT's absolute amplitude) by multiplying the FFT data (black) with a suitable mask with, allows to remove the artefact-peak-contributions from the raw interferometer signal after performing an inverse FFT:

$$\mathbf{s}(\nu) = \mathrm{FFT}\{\boldsymbol{a}(t)\}\ , \qquad (1)$$

$\boldsymbol{a}(t)$ is the oscillatory, raw, background-corrected interferogram data, with t the time delay induced by the shift of the mirror, i.e. signifying the time domain. $\mathbf{s}(\nu)$ is the (complex) Fourier transformed data of $\boldsymbol{a}(t)$ in the spectral domain ν. $\boldsymbol{M}_{532\mathrm{nm}}$ is the masking vector to suppress the valid peak at 532 nm to calculate by the inverse Fourier transform (FFT^{-1}) the waveform in the time domain that describes the artefact peaks $\boldsymbol{a}_{ARTEFACT}(t)$:

$$\boldsymbol{a}_{ARTEFACT}(t) = \mathrm{FFT}^{-1}\{\boldsymbol{M}_{532\mathrm{nm}} \cdot \mathbf{s}(\nu)\}\ . \qquad (2)$$

The real part, signified by $\Re\{\dots\}$, of the complex vector $\boldsymbol{a}_{ARTEFACT}(t)$ is subtracted from the original data to obtain the artefact-suppressed time domain signal $\boldsymbol{a}_{CLEAN}(t)$:

$$\boldsymbol{a}_{CLEAN}(t) = \boldsymbol{a}(t) - \Re\{\boldsymbol{a}_{ARTEFACT}(t)\}\ . \qquad (3)$$

The Hilbert transform $\mathrm{HT}\{\dots\}$ is then used on the corrected time domain signal $\boldsymbol{a}_{CLEAN}(t)$ to extract its phase run in the time domain $\boldsymbol{a}^{Phase}_{CLEAN}(t)$. (The envelope $\boldsymbol{a}^{Envelope}_{CLEAN}(t)$ provided by the Hilbert transform is not used in further calculations.):

$$\left\{\boldsymbol{a}^{Envelope}_{CLEAN}(t), \boldsymbol{a}^{Phase}_{CLEAN}(t)\right\} = \mathrm{HT}\{\boldsymbol{a}_{CLEAN}(t)\}\ . \qquad (4)$$

With these preparations the matched-filter approach can be implemented. We first need a perfectly equidistantly spaced phase in the time domain, $\boldsymbol{a}_{Virtual\lambda_i}$, that is generated from the number of all frames $\#VidFrames$ of the whole interferogram, the traveled distance d of the stage, assuming the (virtual) wavelength λ_i:

$$\boldsymbol{a}_{Virtual\lambda_i} = \left\{\frac{2\pi}{\lambda_i} \cdot k \cdot \frac{2\,d}{\#VidFrames} \,\middle|\, k \in \mathbb{N}_0\right\}\ . \qquad (5)$$

The "coherence degree" $\boldsymbol{CD}$ can be calculated for each (virtual) wavelength λ_i either using the signal without artefact-peaks using eq. (6) or from the phase run of the original raw $\boldsymbol{a}^{Phase}_{RAW}(t)$ (extraction procedure described in Sections 2.3 and 2.4) using eq. (7):

$$\boldsymbol{CD}_{CLEAN}(\lambda_i) = \left[\left(\sin\left(\boldsymbol{a}^{Phase}_{CLEAN}(t)\right).\sin\left(\boldsymbol{a}_{Virtual\lambda_i}\right)\right)^2 + \left(\sin\left(\boldsymbol{a}^{Phase}_{CLEAN}(t)\right).\cos\left(\boldsymbol{a}_{Virtual\lambda_i}\right)\right)^2\right]^{1/2}\ , \qquad (6)$$

$$\boldsymbol{CD}_{RAW}(\lambda_i) = \left[\left(\sin\left(\boldsymbol{a}^{Phase}_{RAW}(t)\right).\sin\left(\boldsymbol{a}_{Virtual\lambda_i}\right)\right)^2 + \left(\sin\left(\boldsymbol{a}^{Phase}_{RAW}(t)\right).\cos\left(\boldsymbol{a}_{Virtual\lambda_i}\right)\right)^2\right]^{1/2}\ . \qquad (7)$$

Plotting the coherence degree against each tested virtual wavelength gives for both cases the "spectrum" shown in Fig. 8. For comparison, the Fourier-transformed signal $\boldsymbol{a}_{CLEAN}(t)$ is shown additionally. All curves are scaled to be the same height. As can be seen, the matched-filter ansatz provides in general a higher signal level across the whole region of interest than the Fourier-transform of the cleaned signal. This is because the Fourier-transform is sensitive to phase shifts and signals at the same wavelength might have a phase shift of π in the time domain signal and suppress or at least diminish the corresponding spectral component. The matched filer is insensitive to any phase shifts and simply accumulates all signal parts that are coherent to the chosen virtual wavelength.

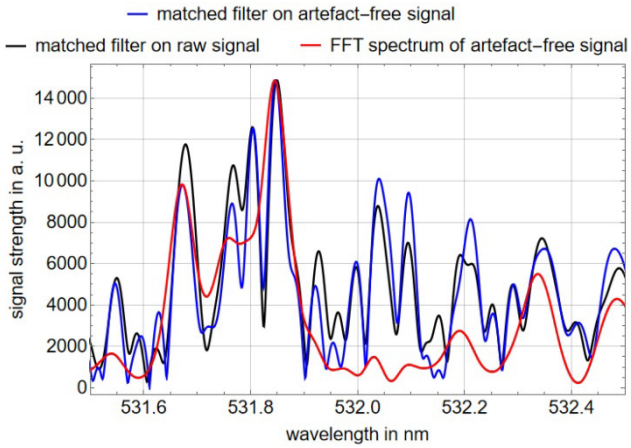

Fig. 8 The matched-filter approach applied to the raw signal (black) and the artefact-free signal (blue). Spectrum of the Fourier-transformed artefact-free signal (red) as comparison.

Choosing the most probable peak:

Fig. 8 shows several peaks that might be generated by the laser. However, due to the known non-uniform motion of the stage, the phase and camera noise, there remain still small fluctuations in the cleaned interferometer signal.

Importantly, no spectral peak can be narrower than the limit set by the total width of the interferometric envelope in the delay domain. This defines a fundamental lower bound for the achievable spectral resolution, regardless of whether the peak corresponds to the true laser wavelength or a spurious artifact.

This means that any frequency component that contributes coherently over the full delay range, either due to a uniform sampling of the laser wavelength, or due to a mechanical vibration that constantly modulates the signal, can appear as a spectrally narrow peak. In contrast, incoherent disturbances such as phase jitter, random noise, or irregular stage motion can only broaden spectral features; they cannot narrow them down beyond the envelope-imposed limit.

Consequently, observing a spectrally narrow peak close to the minimal theoretical width is a strong indicator of a coherent contribution to the interferogram, either due to the actual laser wavelength or a physically plausible, stable modulation. Broader peaks, by comparison, are more likely to result from incoherent noise or temporally localized distortions.

In our implementation, the matched-filter output is computed using a quadrature-based approach that combines the in-phase and quadrature components of the interferogram with respect to each virtual wavelength. This yields a phase-invariant coherence measure for each tested wavelength. While this method ensures that contributions with arbitrary local phase shifts are still detected, it does not change the fundamental resolution limit: the minimal peak width is still governed by the total delay span over which coherent contributions exist. In this sense, the quadrature matched filter acts as a robust detector of spectral content while remaining strictly limited by the effective temporal coherence of the interferometric signal.

Because of this we have now a very good tool at hand to find plausible peaks in Fig. 8 that might indeed correspond to the sought wavelength. All that is needed is just a comparison of all the peaks to the Fourier-transformed of the signal's envelope, shown in Fig. 9, being defined in Section 2.3.

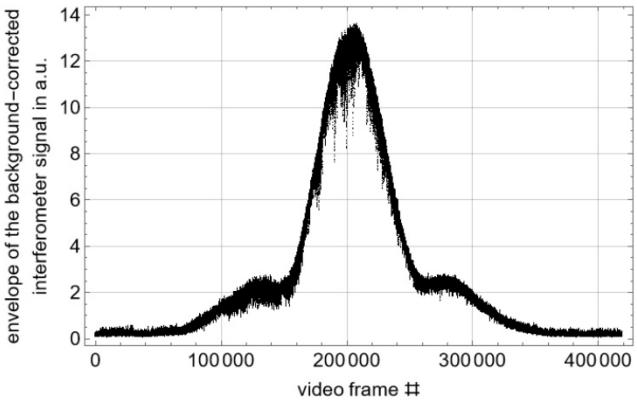

Fig. 9 Recovered envelope for all video frames.

We calculate the spectral envelope from the interferometric amplitude shown in Fig. 9. Since the envelope itself is not oscillatory, it cannot directly be compared to the spectrally resolved peaks shown in Fig. 8. To bridge this, we multiply the envelope with a sinusoidal carrier wave at the nominal center wavelength (532 nm), producing a quasi-oscillatory signal suitable for Fourier transformation.

The resulting spectral envelope is computed in the wavenumber domain, allowing direct comparison to the matched-filter response spectrum. Since the envelope spectrum is a smooth function, we use an interpolated version to enable efficient parametric fitting. The fitting function is defined as a shifted and scaled version of this interpolated envelope spectrum and depends only on the amplitude and the spectral shift. The optimal fit is obtained by minimizing the squared error between the matched-filter spectrum and the shifted envelope spectrum in the spectral region of interest. Performing the fit in the wavenumber domain ensures that the spectral shape (e.g., width) remains invariant under shifts, which would not be the case in wavelength space due to its nonlinear scaling. The results for the three best peaks is shown together with the two-dimensional cost function for the relevant spectrum in Fig. 10.

The frequency shift of zero in the graph in Fig. 10 (a) corresponds to 532 nm, which is 18796.9925 cm^{-1} ≈ 18797 cm^{-1} in Fig. 10 (b). As can be seen, the peaks with the lowest error are quite close to 18797 cm^{-1}. Lu et al. ([10]) report an Nd:YAG (ceramics) wavelength between 1064.16 and 1064.31 nm (9395.76 - 9397.08 cm^{-1}) depending on the Nd^{3+}-ion concentration with fluorescence widths of below 0.86 nm (7.6 cm^{-1}). In addition, their comparison to single-

crystal Nd:YAG the same emission wavelength and width for a Nd:YAG single crystal with a Nd^{3+}-ion concentration of 1%. The bandwidth reported in [11] is with up to 11 cm^{-1} is somewhat larger, but does, importantly, not show any significant influence of nanocrystalline, ceramic or monocrystalline structures of Nd:YAG. In addition, a random Coherent HyperRapid NXT spectrum provided by Coherent to the authors gives a peak wavelength of 1064.092 nm with a width of 0,354 nm. Consequently, it is reasonable to assume that our Coherent Hyperrapid NXT laser unit has similar properties.

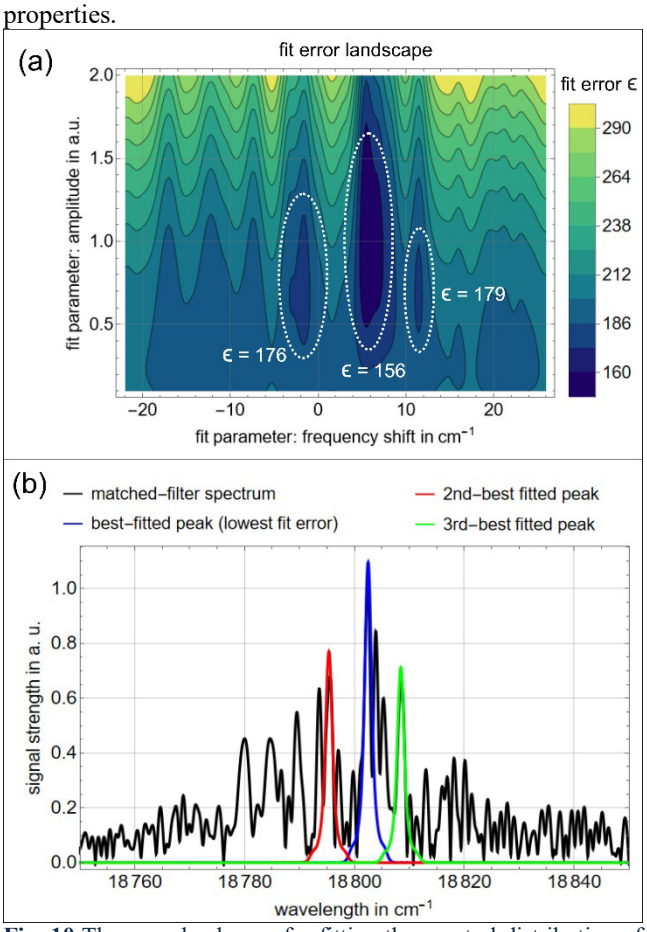

Fig. 10 The error landscape for fitting the spectral distribution of the interferogram's envelope on the matched-filter spectrum (a). The positions with the lowest achieved fit error are marked by dashed white ellipses. The corresponding fitted peaks are shown in (b).

Nevertheless, to better assess how well each candidate envelope peak represents the laser's true spectrum, we consider first the fitting error. In addition, we evaluate also the shape similarity and the spectral energy fraction, which provide a more complete picture of each peak's relevance:

The fitting error is directly obtained from the fitting procedure described above. The peak's central wavelength is derived from its wavenumber, which is also determined in the fitting step. The shape similarity measure (SSM) between the candidate peak and the measured spectrum is based on the normalized cross-correlation. It quantifies how well the shapes of the two curves agree, regardless of their absolute amplitudes. A high value indicates a strong match in shape, even if the peak amplitude differs from that of the signal. This measure is particularly sensitive to the peak's position and width, but robust to global amplitude scaling or background variations.

Furthermore, we compute the spectral energy fraction (SEF), which quantifies how much of the total spectral energy lies within the support of the candidate peak. A high value here indicates that the peak accounts for a substantial portion of the overall signal, regardless of how precisely its shape matches the spectrum. Together, these two measures complement the fitting error: while the error function minimizes pointwise deviation, it does not reveal whether the candidate truly reflects the structure of the spectrum or only coincides with it locally or coincidentally. The results are summarized in Table 1.

Table 1 Properties of the fitted peaks (see text for details). The best values are marked in green.

Peak #	λ_{PEAK} in nm	Fit error	SSM	SEF
1	531.844	**156.1**	**0.54**	**0.24**
2	532.046	176.7	0.25	0.10
3	531.676	179.7	0.29	0.12

As can be seen, the peak at 531.84 nm, has, besides having the smallest fit error, also the best correspondence to the shape and carries most of the spectral energy. Its wavelength 531.844 nm would correspond to a wavelength of 1063.688 nm in the fundamental. Though this is 0.312 nm below 1064.0 nm it would nevertheless fit into an envelope of a Nd:YAG fluorescence line at 1064.16 nm with a spectral FWHM width of 0.86 nm [10] or, even larger, up to 11 cm^{-1} as in [11]. Hence, it is logical to assume this peak to represent the wavelength of the laser. In Fig. 11 we show this peak in the wavelength-space separately.

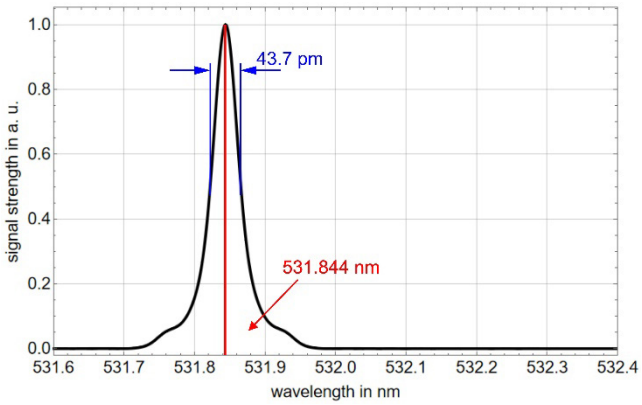

Fig. 11 The most probable peak shown as a function of the wavelength.

The full-width-at-half-maximum (FWHM) bandwidth of the peak is calculated to be 0.0437 nm (1.54 cm^{-1} or 46.26 GHz). Assuming the pulse shape is a hyperbolic secant, this would allow for 6.7 ps FWHM pulse duration, which is very close to the pulse duration of 6.1 ps (that would have a spectral bandwidth of 50.8 GHz/0.048 nm at 532 nm in the case of being time-bandwidth-limited) estimated from the SFG-autocorrelation measurement given at the beginning of Section 2.

Since non-uniform stage motion inevitably produces artifact peaks near the true signal, our algorithm relies on a probability analysis that incorporates a-priori knowledge of the laser's basic spectral characteristics. Consequently, it is not universally applicable to arbitrary light sources. It performs robustly whenever the spectrum contains a single dominant peak or several well-separated peaks, whose approximate positions and envelope are known in advance. Conversely, it becomes unreliable for densely wavelength-multiplexed sources (e.g., high-power diode arrays or telecom signals with closely spaced lines), where multiple unresolved modes would produce overlapping artifact and real peaks that cannot be unambiguously discriminated by our matched-filter and envelope-based criteria.

4. Conclusion

The interferometric measurement allowed us to determine the spectral characteristics of a frequency-doubled, industrial-grade picosecond Nd:YAG laser. The most probable emission peak was found at 531.844 nm, corresponding to 1063.688 nm at the fundamental wavelength, with a full-width at half-maximum (FWHM) bandwidth of 0.0437 nm. These values align well with previously reported emission characteristics of Nd:YAG crystals and fall within the expected fluorescence envelope.

The underlying data were acquired using a Michelson interferometer with a PID-controlled linear stage, whose non-uniform motion introduced significant sampling irregularities. To address this, a multi-step correction process was applied: spectral artefacts were first removed via selective inverse FFT filtering, followed by a matched-filter algorithm applied to the corrected time-domain signal. This approach allowed us to identify coherent spectral components even in the presence of phase noise and sampling distortion.

A key advantage of the proposed method lies in its high spectral resolution without the need for phase-stabilized delay lines or calibrated spectrometers. By leveraging the full interferometric envelope and employing a phase-insensitive quadrature-based coherence analysis, we were able to isolate the actual laser emission from noise and artefacts with high confidence and sub-nanometer accuracy.

Acknowledgements

The authors gratefully acknowledge funding of the Erlangen Graduate School in Advanced Optical Technologies (SAOT) by the Bavarian State Ministry for Science and Art.

Disclosures

In the experimental process of this work, the authors used the tools "ChatGPT" to assist in programming and editing tasks. After using this tool, the authors reviewed and edited the content as necessary and they take full responsibility for the content of the publication.

References

[1] G. Batignani, D. Bossini, N. Di Palo, C. Ferrante, E. Pontecorvo, G. Cerullo, A. Kimen, and T. Scopigno: Nat. Photonics, 9, (2015) 506.
[2] M. Kerstan, I. Makos, S. Nolte, A. Tünnermann, and R. Ackermann: Appl. Phys. Lett., 110, (2017) 021116.
[3] Y. Meng, C. Zhang, C. Marceau, A.Y. Naumov, P.B. Corkum, and D.M. Villeneuve: J. Mod. Opt., 63, (2016) 1661.
[4] M. Wen and J. Houlihan: Opt. Commun., 540, (2023) 129491.
[5] H.C. Hendargo, M. Zhao, N. Shepherd, and J.A. Izatt: Opt. Express, 17, (2009) 5039.
[6] L. Song, G. Shi, H. Liu, H. Lin, F. Zhang, and D. Sun: Sensors, 22, (2022) 9021.
[7] S. Luan, M.H.R. Hutchinson, R.A. Smith, and F. Zhou: Meas. Sci. Technol., 4, (1993) 1426.
[8] B. Boashash: Proc. IEEE, 80, (1992) 520.
[9] J.M. Bioucas-Dias and G. Valadão: IEEE Trans. Image Process., 16, (2007) 698.
[10] J. Lu, M. Prabhu, J. Song, C. Li, J. Xu, K. Ueda, A.A. Kaminskii, H. Yagi, and T. Yanagitani: Appl. Phys. B, 71, (2000) 469.
[11] M. Pokhrel, N. Ray, G.A. Kumar, and D.K. Sardar: Opt. Mater. Express, 2, (2012) 235.

(Received: June 30, 2025, Accepted: September 28, 2025)