Research Article | Vol. 31, No. 22 / 23 Oct 2023 / Optics Express 36915
Optics EXPRESS

Evolutionary optimization and long-term stabilization of a white-light seeded two-stage OPCPA seed laser

TIMO EICHNER,[1,2,*] THOMAS HÜLSENBUSCH,[1,2] GUIDO PALMER,[1] AND ANDREAS R. MAIER[1]

[1]Deutsches Elektronen-Synchrotron DESY, Notkestr. 85, 22607 Hamburg, Germany
[2]Department of Physics, University of Hamburg, Luruper Chaussee 149, 22761 Hamburg, Germany
[*]timo.eichner@desy.de

Abstract: Ultrafast laser systems, such as optical parametric chirped pulse amplifiers (OPCPA), are complex tools. Optimizing laser performance for a given application is often plagued by intricate couplings between different output parameters, making simultaneous control of multiple pulse properties difficult. Here, we experimentally demonstrate an autonomous tuning procedure of a white-light seeded two-stage OPCPA using an evolutionary strategy to reliably reach an optimized working point. We use the data collected during the tuning procedure to calibrate a performance model of the laser system, which we then apply to stabilize the intricately coupled laser output energy and spectrum simultaneously. Our approach ensures reliable day-to-day operation at optimized working points without manual tuning. We demonstrate shot-to-shot energy stability of <0.18 % rms, in combination with <25 pm rms wavelength stability and <0.2 % rms bandwidth stability during multi-day operation.

1. Introduction

Operating a laser system at an optimal working point and ensuring the reproducibility of that working point is crucial for arguably any laser application. However, in real-world operation, perfect day-to-day reproducibility is difficult to achieve: a changing laboratory environment may have subtle influences on the laser and change its operation conditions. Manually recovering the system performance is often highly dependent on the experience and skill of the lasers operator.

This is especially true for complex ultrafast laser systems. A prime example are optical parametric chirped pulse amplifiers (OPCPA) [1,2]. The flexibility in laser pulse properties makes them an excellent source for a wide range of applications. However, the properties of the amplified pulses are typically coupled and adjusting only one input parameter, e.g., the energy of a pump pulse, can simultaneously affect many parameters, e.g., the energy, bandwidth, and spectral shape of the amplified pulse. This complexity is caused by an intricate and often non-intuitive interplay of different nonlinear and dispersive effects. The complexity quickly increases when extending the system, e.g., by combining multiple amplification stages to achieve higher overall gain. While this interplay provides a large number of tunable parameters that enable flexibility of the output parameters, it also makes it difficult to ensure an optimally tuned laser. Environmental influences, such as a warm-up of the laser hardware, can further reduce the reproducibility of performance. The success of manual tuning to recover the performance often highly depends on the operator's experience and knowledge of the specific details of the laser system.

To be independent from such factors, an automated tuning and control procedure is desirable. In the past, machine learning techniques have been used in several areas of photonics [3]. In particular, the application to the self-tuning, optimization and mode-locking of ultrafast fiber lasers

#504694 https://doi.org/10.1364/OE.504694
Journal © 2023 Received 30 Aug 2023; revised 9 Oct 2023; accepted 9 Oct 2023; published 17 Oct 2023

Optics EXPRESS

[4,5] and the application of deep reinforcement learning to optimize nonlinear processes such as white light generation [6] demonstrate the potential benefits of machine learning techniques in the generation and amplification of ultrafast laser pulses. However, to our knowledge, such methods have not been applied to more complex multi-stage laser systems.

Here, we apply an evolutionary strategy to optimize the operating point of a two-stage ultrafast OPCPA system that has been designed to seed a Ti:sapphire-based high-intensity laser system. We demonstrate automated optimization of the laser for a variety of performance goals. Our approach can be used to both (i) find a working point for a specific application and to (ii) improve the day-to-day reproducibility of the laser performance. We use the data collected during the optimization process to calibrate a linear performance model of the laser system, which we then use to stabilise the system using full-state feedback control. While relatively simple, our approach allows for precise long-term control of the laser performance and is transferable to other laser technologies.

2. Experimental setup

Our experiments use a two-stage white-light seeded OPCPA system – named MALCOLM – that has previously been described in [7]. It delivers up to 50 μJ pulses with a spectrum centered at around 800 nm, and at a repetition rate of 1 kHz. The spectrum of the amplified pulses supports a minimum pulse duration of 25 fs. The pulses are amplified in lithium triborate (LBO) crystals, that are colinearly phase-matched to ensure a high spatio-temporal beam quality.

As illustrated in Fig. 1, the OPCPA is driven by a commercial Yb-laser (Pharos, LightConversion), that delivers 1 mJ pulses at a center wavelength of 1030 nm. The pulses are compressible to approximately 170 fs, but we drive the parametric amplifier with a stretched output of 500 fs pulse duration. These pulses are frequency doubled in a 1 mm long barium borate (BBO) crystal to provide a 515 nm pump beam for the OPCPA stages. A few-μJ fraction of the fundamental pulses are split off before the second harmonic generation (SHG), recompressed to the 170 fs Fourier-limit in a transmission grating compressor, and used to drive the white light generation (WLG) in YAG to provide the broadband seed that is amplified in the OPCPA stages. The beam coming from the Pharos laser system is actively stabilised into the OPCPA setup by a piezo-based beam stabilisation system.

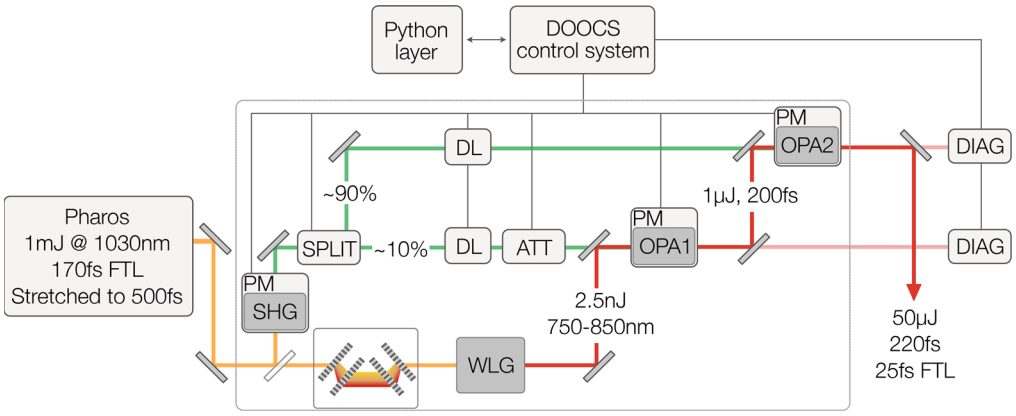

Fig. 1. Layout of the optical and diagnostic setup of the OPCPA system. SHG: Second harmonic generation, WLG: White light generation, OPA: Optical parametric amplification stage, DL: Delay line, SPLIT: Variable beam splitter, ATT: Variable attenuator, PM: Phase matching angle, DIAG: Diagnostic setup

Optics EXPRESS

The phase matching angles of the SHG and the two OPCPA stages can be adjusted using motorized rotation stages. In addition, the relative delay between the pump and seed pulses in both OPCPA stages can be adjusted with motorized delay lines in the respective pump arms, and the pump energy of each OPCPA stage can be adjusted with a variable beam splitting unit and an attenuator consisting of a motorized half-wave plate and a thin-film polarizer (TFP). All of these actuators are piezo-crawler stages. In total, the laser system thus has 7 degrees of freedom that can be used for optimization.

Software limits for each stage ensure safe operation of the laser. These limits are chosen to correspond to mechanical constraints in the case of phase matching angles and delay lines, and for the pump attenuators to keep the pump intensity below 60 GW/cm^2 to prevent laser-induced damage of optics and ensure long-term reliable operation.

The spectrum and energy of the signal pulses are monitored by online diagnostics that use the spectrally flat 2 % leakage of dedicated beam samplers after each OPCPA stage.

All diagnostics and stages are implemented in the DOOCS accelerator control system available at DESY [8] and can be controlled using Python. All diagnostics operate at a repetition rate of 10 Hz.

3. Self-optimization

A number of numerical methods have been successfully used in the past to optimize the performance of fiber lasers where the beam propagates in a well-controlled environment shielded from outside influences. For laser systems such as optical parametric chirped pulse amplifiers with substantial free-space propagation, the laser is more susceptible to changes in ambient conditions, leading to increased noise and potential inaccuracies in the measurement of the laser properties. In addition, small daily changes in the environmental conditions result in the inability to create a sufficiently accurate and complete numerical model of the laser system that could be used for numerical optimization of laser performance.

3.1. Overview of suitable optimisation strategies

Many black-box optimization algorithms, that do not rely on any knowledge about the system, have been developed and are well suited for optimization tasks with changing or unknown parameters of the system. For example, Bayesian optimization has been successfully applied to laser wakefield acceleration [9,10], which poses a similar optimization problem to that of multi-stage laser systems. However, such surrogate-based methods have the disadvantage of limited scalability to high-dimensional problems – i.e., problems with a large number of input variables. In addition, fitting the surrogate model to the gathered data is computationally expensive, making Bayesian optimization more appropriate for applications where the time between sampling is long.

Evolutionary methods, such as genetic algorithms or evolutionary strategies, do not rely on a surrogate model and can therefore outperform surrogate-based methods in applications where sampling is fast [11] – such as high repetition rate laser systems. These methods optimize a function mimicking biological evolution by comparing the performance of individuals (inputs to the function) within a population, and continuously adjusting the population according to a set of rules so that the overall performance of the population moves toward an optimum of the function. Many such algorithms scale well to large numbers of input variables [11]. Within evolutionary algorithms, evolutionary strategies are particularly well suited for optimization in continuous parameter spaces [12], such as the positions of motorized stages.

A versatile and commonly used evolutionary strategy is the Covariance Matrix Adaptation Evolutionary Strategy (CMA-ES) [13], in which the population consists of a distribution of randomly sampled points. The mean and covariance matrix of this distribution are continuously adapted, such that the distribution moves towards an optimum – similar to gradient descent

Optics EXPRESS

methods – with the spread of the distribution continuously decreasing. CMA-ES works particularly well with relatively few sampled points, ill-conditioned problems, and non-separable problems [14], meaning that the input variables are strongly interdependent, as is the case in OPCPA systems. While there are optimization methods with similar characteristics and possibly even better performance in certain aspects, CMA-ES also has a small number of hyperparameters, i.e., parameters that control the algorithm itself, that may need to be tuned to make it optimise it for a specific application. The resulting lack of need for extensive parameter tuning of the algorithm makes the implementation of the algorithm relatively easy. For our application, we therefore used a Python-based open source CMA-ES implementation [15].

The optimization process in CMA-ES starts from a point X_0 within the search space, that can either be randomly chosen or user-defined. A normally distributed population of N points – called x_i – with a standard deviation of σ_0 around the mean X_0 is then sampled to determine the local variation of the fitness $f(x_i)$ of the function to be optimized. In the case of a laser system, the fitness can be determined by some performance characteristic, such as, e.g., the output energy, and the sampled point x_i can be a set of input parameters. In simple implementations, the sampled points are then ranked according to their respective fitness values and the covariance matrix of the μ best points is used to determine the mean X_1 and the width σ_1 of the distribution of points of the next iteration. Newly sampled points within this updated distribution are then added to the population to form the next generation of the evolutionary process. Over several such iterations, the mean of the population converges to a global optimum, with the width of the distribution decreasing once the optimum lies within the sampled distribution [14].

Due to the inherently statistical treatment of sampled points, CMA-ES tends to perform better than many other evolutionary methods in the optimization of noisy functions [16]. However, there is a trade-off between speed and noise robustness, as an increased robustness is achieved by a larger number of sampled points per generation.

In our specific case we used the previously mentioned 7 actuators of our OPCPA system – the phase matching angles of the SHG and both OPA stages, as well as the pump energy and pump delay of both OPA stages – as input variables to the optimization algorithm and used a population size of 9. Due to speed limitations of the actuators, the stage positions could only be updated at a rate of 1 Hz. In this configuration, the potential speed of the evolutionary strategies algorithm could therefore not be fully exploited, but we still used this approach to ensure scalability in future systems with upgraded hardware.

3.2. Self-tuning from a random initial state

A useful and desirable goal when seeding a Ti:Sapphire laser system is to maximize the bandwidth and energy of the seed laser, while centering the spectrum at a wavelength that pre-compensates the red-shift in the subsequent Ti:Sapphire amplifiers. We can use these desired parameters to define a fitness function, describing laser output in a single value. For this particular objective we chose

$$f = \Delta\lambda \cdot E - C(\lambda_{\mathrm{COG}} - \lambda_0)^2, \tag{1}$$

where $\Delta\lambda$ is the FWHM bandwidth of the output spectrum, E is the output energy, λ_{COG} is the center of gravity wavelength of the output spectrum and λ_0 is the desired center wavelength. C is an experimentally determined weighting constant. The product of the spectral bandwidth and the output energy is roughly proportional to the peak power of the fully compressed pulses of the OPCPA, which we want to maximize.

Figure 2 shows the evolution of the fitness function, the pulse energy, the bandwidth and the central wavelength over the duration of an optimization run, starting from a random configuration within the safety limits of the motorised stages. Initially, there is no amplification in the OPA stages, and only after around 40 generations of randomly sampling the parameter space, a configuration with a low output energy (within the dynamic range of the diagnostics) is found

Optics EXPRESS

(compare also Fig. 3). The optimizer then starts to iteratively improve on this configuration and to move the laser towards a state with high output energy. After around 200 iterations, the optimizer converges to an optimal value, with further search leading to only very minor improvements The evolution of the spectrum is also visible in Fig. 3.

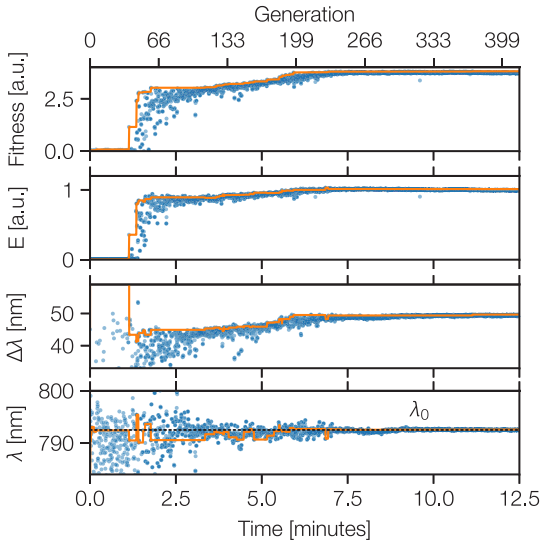

Fig. 2. Evolution of pulse properties during optimization from a random initial point. The top panel shows the fitness function calculated using Eq. (1). Below are the normalized output energy, the FWHM bandwidth $\Delta\lambda$ and the center of mass wavelength of the output spectrum. The target wavelength of 792 nm is indicated by the dashed line. The orange lines show the respective values best configuration that has been found so far.

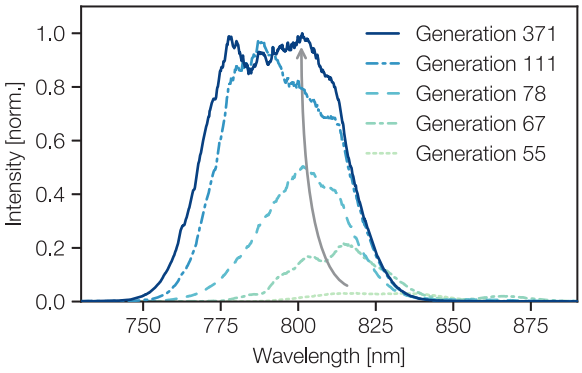

Fig. 3. Snapshots of the output spectrum at different stages during optimization. The arrow indicates the direction of evolution of the spectrum.

This is not only significantly faster than manual tuning, but also typically results in a 5-10% increase in output energy and bandwidth compared to manual tuning. While the comparison is somewhat subjective, it underscores the potential of automated tuning procedures.

The optimizer not only adjusts the parameters of the final OPCPA stage, whose output is considered in the fitness function, but also actively adjusts the properties of the seed and pump pulses, e.g. through slight adjustments of the phase matching angle of the SHG that generates the pump pulse. Due to its narrow phase matching bandwidth, this adjustment has a significant

Optics EXPRESS

influence on the spectral and temporal shape of the pump pulse, which then affects the spectral properties of the fully amplified pulse. Adjustments to the seed pulse that comes from the first OPA stage include its energy, but also the bandwidth and, most interestingly, the center wavelength of the seed, which is shifted toward shorter wavelengths by about 10 nm to pre-compensate a red shift due to temporal walk-off in the final OPA stage.

In our specific laser system, the spatial overlap of the pump and seed beams cannot be adjusted in an automated way, but an iterative tuning procedure with alternating manual alignment of the mirrors and automated tuning of delays, pump energies and phase matching angles can nevertheless lead to an overall acceleration of the alignment of the laser system, where changes in the spatial alignment can also influence the spectral properties, e.g. due to changes in the optical path length of the pump and seed pulses.

3.3. Improving day-to-day reproducibility

While the previously described approach works well to maximize the performance at any given time, changes in the lab conditions can affect performance. In our specific case, performance drifts on a minute time scale were observed despite high temperature stability of the laboratory environment of ±0.1 K. To ensure the highest possible repeatability of performance from day to day and to minimize the remaining performance drifts, it may be better to use a different fitness function. An initial optimization can indeed be done with the fitness function described above, but with slightly adjusted limits of the parameter space, in order to leave some room for compensation of drifts in performance. This is especially necessary in the case of the pump energies, since these are typically limited by the conversion efficiency of the SHG and only decrease with changes in the system. Once an optimal working point is reached, the output spectrum can be stored as a reference for future optimization runs.

By defining a fitness function that characterises the deviation of the output spectrum from the un-normalized reference, one can reach the same spectral shape and amplitude on a daily basis. The fitness function we minimized to achieve this can be written as

$$f = \text{mean}\left[(I(\lambda) - I_{\text{ref}}(\lambda))^2\right]. \tag{2}$$

Figure 4 shows the performance of the laser system over a five day period with daily fine-tuning of the operating point of the laser system. During the fine-tuning, which typically happended in the morning, Eq. (2) was minimized starting from the current operating point and with the initial spread of the population limited to 5% of the extent of the safety limits. This restricts the optimization to a small search space and allows faster convergence, which is typically achieved in around 2 minutes. To ensure that the optimizer has indeed converged, we let the it run for a further minute beyond this typical time and stop the optimisation procedure only after 3 minutes.

From Fig. 4 we can also see that while the operation point is successfully recovered at each optimization, there are some drifts in laser performance between the tuning times. To compensate for these drifts, the laser system needs to be adapted more frequently, which would be possible with a feedback loop that stabilizes the relevant laser parameters.

Optics EXPRESS

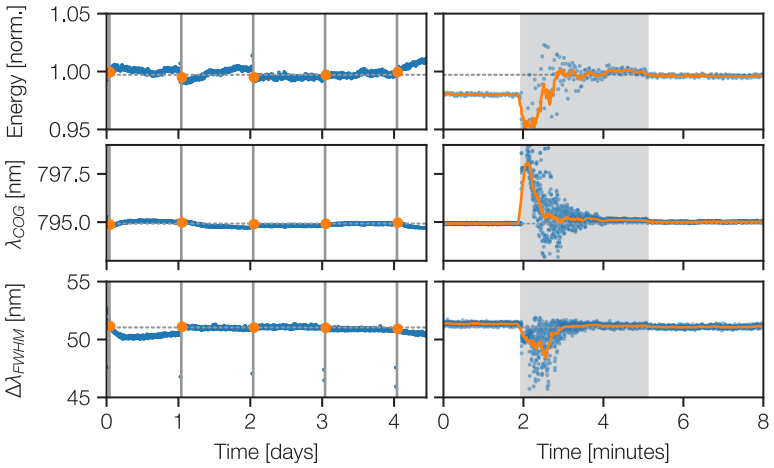

Fig. 4. Reproducibility of daily optimization over a 4.5-day period. The vertical lines show the periods of the daily optimization. The dots in the left column show the respective optimisation outcomes. The right column shows a zoom into an example of one of the 3-minute-long optimisation periods, with the solid line showing a rolling mean of the respective pulse properties.

4. Full state feedback control at an optimized working point

Stabilising a system with intricately coupled parameters faces a similar challenge as optimising such a system: The output properties are strongly coupled, and multiple output properties can be influenced by a single actuator. As a result, independent feedback loops based on conventional single-input-single-output (SISO) feedback loops that aim to stabilize separate laser properties can interfere with each other, potentially leading to resonances and unwanted fluctuations in some of the laser parameters.

One approach to avoid such interference is to use a centralised controller that stabilizes the overall state of the laser system, that is defined by the set of monitored output parameters. Stabilization of multiple-input multiple-output (MIMO) systems is a common problem throughout engineering, and a widely used solution is full-state feedback (FSF) control. This is based on a linear state-space representation of the system for which an analytical solution exists, that brings the system to a desired state.

4.1. Overview of state space control

State space control is based on a simple linear model of the system to be controlled. Since the 10 Hz operating frequency of our diagnostic system naturally introduces discrete time-steps, we will use the discrete-time formulation to describe the temporal evolution of the laser system.

For our laser system, we can simplify the general description of discrete-time state-space-models [17] to

$$x_{k+1} = x_k + Bu_k \tag{3}$$

where x_k is the state vector of the system at time k, and $x_k = x_{k,\mathrm{sp}} - x_{k,\mathrm{absolute}}$ is the deviation of the actually measured state $x_{k,\mathrm{absolute}}$ from the setpoint $x_{k,\mathrm{sp}}$. $B \in \mathbb{R}^{n\times m}$ is the input matrix, which describes the influence of a given input on the state of the system. u_k is the relative change of the actuator positions. x_k and u_k have dimensions n and m respectively.

Optics EXPRESS

In our specific case, the absolute state-vector

$$x_{k,\text{absolute}} = \begin{bmatrix} \text{E}_{\text{OPA1}} \\ \text{E}_{\text{OPA2}} \\ \lambda_{\text{OPA2}} \\ \Delta\lambda_{\text{OPA2}} \end{bmatrix} \tag{4}$$

is given by the output energy of both OPA stages, as well as the center wavelength and bandwidth of the second OPA, while the input vector

$$u_k = \begin{bmatrix} \tau_{\text{OPA1}} \\ \tau_{\text{OPA2}} \\ \text{E}_{\text{P,OPA1}} \\ \text{E}_{\text{P,OPA2}} \end{bmatrix} \tag{5}$$

consists of the pump delays τ and the pump energies E_P of the two OPA stages, respectively.

For such a system to be fully controllable by the inputs u_k, it must satisfy the controllability condition which in our case can be written as [17]

$$\text{rank}(B) = n. \tag{6}$$

This means that for the system to be movable to any arbitrary state x_{k+1} for which the state space model is valid, the different axes of the system must be linearly independent. For a system in which each input influences only a single output parameter, this is condition is always fulfilled, since each input-output-relation is independent from one another. However, for systems in which individual inputs affect multiple outputs – such as in OPCPA systems as we described earlier – this is not as trivial, and controllability must be explicitly verified.

If the system is controllable, it can also be stabilized. This can be achieved by making the inputs to the system dependent on the current state, thus creating a feedback loop

$$x_{k+1} = (\mathbb{I}^n + BK)x_k, \tag{7}$$

where K is the feedback matrix. To stabilize the system, i.e., to ensure that $x_{k+1} \simeq 0$ is fulfilled over long durations, we have to solve this equation and calculate the required feedback gains K.

4.2. *Linear performance model of* Malcolm

A first step in solving Eq. (7) is to determine B, which describes how the actuators act on the laser performance. During the daily optimization (see section 3.2) we typically explore around 1000 different actuator positions around the optimized working point. We have recorded these actuator positions together with the corresponding state of the system, i.e., the values measured by the online diagnostics of the laser, resulting in a dataset that relates input parameters to laser performance. For small deviations from the optimum, i.e., for data points close to the optimum working point, the relationship can be approximated to be linear and the input matrix B can be calculated by multivariate linear regression.

We performed the linear regression on the FWHM bandwidth, pulse energy, and center wavelength of the two OPA stages, using the 650 data points with the best fitness value according to Eq. (2). This restricts the data set to points closest to the optimum working point, which ensures a linear relationship between the parameters by omitting points heavily influenced by the

Optics EXPRESS

nonlinear behavior of the laser system. We further modify the measured values by subtracting their mean and normalize them by their standard deviation to scale the parameters to a similar magnitude and allow meaningful statistical analysis.

Since the most relevant parameters to stabilize are the output parameters of OPA2, namely its output energy, center wavelength, and spectral bandwidth, we left the spectral properties of OPA1 unstabilized to not over-constrain the system. As a result, we have four stabilized output parameters: The center wavelength, the spectral bandwidth of OPA2, and pulse energy of OPA2 and OPA1. We found that these parameters are mainly influenced by the respective pump energy and pump pulse delay in the two OPA stages. Therefore, we limited our linear regression to these sets of four input and output parameters.

The laser performance is well described by the linear model, as illustrated in Fig. 5. The blue line shows the measured laser parameters and, for comparison, the laser performance as predicted from the stage positions for the same shots (orange). As expected, the bulk of the performance variations is determined by the four input parameters. In particular, the output energy of OPA stages 1 and 2, and center wavelength of the second OPA stage, are well represented by the model with coefficients of determination of $R^2 = 0.8$, $R^2 = 0.73$, and $R^2 = 0.63$, respectively. In the case of the 10%-bandwidth, the coefficient is $R^2 = 0.48$, which is reasonable considering the complex couplings in the OPCPA that affect the bandwidth.

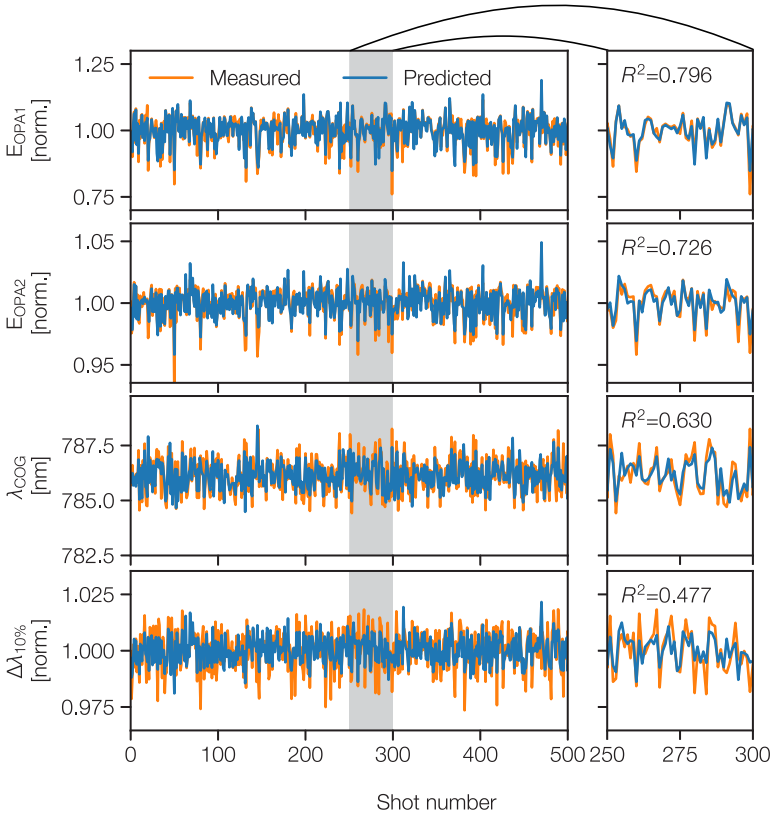

Fig. 5. Blue solid: measured performance parameters – pulse energy, center of mass wavelength λ_{COG}, and 10%-bandwidth $\Delta\lambda$ of the laser. Orange solid: laser performance as predicted by the linearized model using the stage positions at the respective laser shot. The right column panels show a zoom to the shaded area in the plots on the left.

Optics EXPRESS

To quantify the relevance of individual input parameters on the laser performance, we calculate the p-value, i.e., the significance level of the observed correlations,

$$p = \begin{array}{c} \\ E_{OPA1} \\ E_{OPA2} \\ CWL_{OPA2} \\ BW_{OPA2} \end{array} \begin{array}{c} \begin{array}{cccc} \tau_{OPA1} & \tau_{OPA2} & E_{P,OPA1} & E_{P,OPA2} \end{array} \\ \begin{pmatrix} 0.00 & 0.84 & 0.08 & 0.09 \\ 0.00 & 0.00 & 0.15 & 0.03 \\ 0.00 & 0.00 & 0.67 & 0.33 \\ 0.00 & 0.00 & 0.00 & 0.00 \end{pmatrix} \end{array} \tag{8}$$

Lower p-values indicate that the observed correlation is more likely to be caused by actual changes in the inputs rather than by random causes or other non-linear behaviour of the laser that is not described by the linear model. For our application, we have defined a cut-off of $p = 0.1$, above which we consider a correlation as spurious. The corresponding elements in Eqs. (8) and (9) are highlighted in red.

As an example, we can see from the low p-values in the bottom row of matrix (8), that the bandwidth of OPA2 is determined by all of the four input parameters (pump delays τ_{OPA1} and τ_{OPA2}, as well as pump energies $E_{P,OPA1}$, $E_{P,OPA2}$) – again highlighting the complex couplings in OPCPA that influence this particular parameter –, while the center wavelength of OPA2 is determined only by the pump delays of the two OPA stages.

To ensure the controllability of the system and to eliminate spurious correlations from influencing the feedback loop, we simplify matrix B by eliminating parameters with $p>0.1$:

$$B = \begin{array}{c} \\ E_{OPA1} \\ E_{OPA2} \\ CWL_{OPA2} \\ BW_{OPA2} \end{array} \begin{array}{c} \begin{array}{cccc} \tau_{OPA1} & \tau_{OPA2} & E_{P,OPA1} & E_{P,OPA2} \end{array} \\ \begin{pmatrix} -0.88 & -0.01 \to 0 & 0.04 & 0.03 \\ -0.75 & -0.49 & 0.01 \to 0 & 0.08 \\ 0.37 & 0.66 & 0.04 \to 0 & 0.03 \to 0 \\ -0.32 & 0.62 & 0.10 & 0.11 \end{pmatrix} \end{array} \tag{9}$$

Despite this simplification, the overall coefficient of determination of our model remains at 0.66, indicating that the high predictability of the model is maintained. The resulting matrix now has a rank of 4 and thus satisfies the controllability condition (Eq. (6)). Note that the values in the input matrix B (Eq. (9)) are in normalized units.

4.3. Calculating stabilization gains

While Eq. (7) can be solved unambiguously now that we know the system response to a given input, the direct solution of the equation generally does not provide a feedback matrix K that has any tolerance to model inaccuracy or measurement noise – both of which are present in real-world implementations.

An alternative method for the calculation of K that results in a more cautious response to changes of the system is pole placement [17]. In this method, the eigenvalues s of the closed-loop system $(1 - BK)$ are chosen to be $|s|<1$, such that Eq. (7) converges towards zero over several time-steps. Adjusting the eigenvalues in the range $-1<s<1$ further allows to tune the response of the feedback loop. An eigenvalue of $s = 0$ would correspond to an aggressive response of the system that aims to eliminate the entire deviation from the setpoint within a single time-step. For values of $s \neq 0$, the response is more cautious, with convergence to a stabilized state occurring over multiple time steps. This provides an opportunity to tune the response of the system and to include some tolerance to an imperfect representation of the laser dynamics by the model.

Optics EXPRESS

In our specific case, we empirically found $s = 0.8$ to be a reasonable compromise between a fast response to external perturbations of the system and tolerance to noise in the measured laser properties, as well as towards drifts of the model over time.

4.4. Results of full-state-feedback-stabilization

Based on this approach, we implemented a full-state-feedback loop, to demonstrate the successful simultaneous stabilization of the energy, bandwidth, and center wavelength of the laser system over 4 hours, as shown in Fig. 6. Due to the limited response time of the actuators, the stabilization was implemented as a simple drift compensation with the feedback loop operating at a rate of 0.5 Hz. In combination with the chosen eigenvalue of the feedback matrix of $s = 0.8$, this means that after a step-like disturbance to the system, the desired setpoint to within 10% of the initial deviation is recovered after 20 seconds.

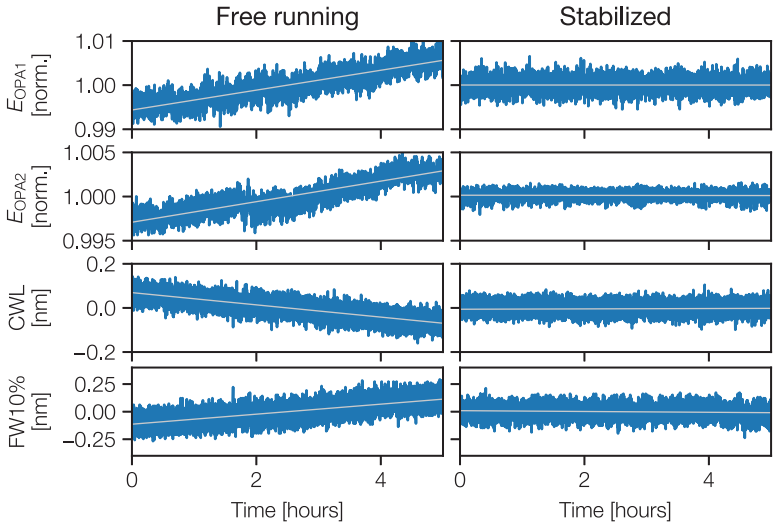

Fig. 6. Free-running (left) and stabilized (right) laser performance using full-state feedback control. The grey lines show a linear fit of the data to guide the eye and illustrate the effectivness of the stabilization.

In particular the stabilization of the bandwidth demonstrates the ability of model-based control to influence certain output parameters without the presence of an individual actuator that directly influence the property. Instead, a combination of all actuators is used, with major components being the saturation level of both amplifier stages and the timing of the pump pulses in the two OPA stages. This is also reflected by the matrices Eq. (8) and Eq. (9).

5. Long term stability

To verify the long term stability of MALCOLM, the laser was operated over period of 50 hours with the stabilization running. The resulting trends in output energy, center wavelength, and bandwidth are shown in Fig. 7. For the measurements, we were running the stabilization using the lasers built-in diagnostics, and measured the laser output performance with an independent set of diagnostics. For the energy measurement, all shots were recorded at the full repetition rate of 1 kHz, while the spectra were recorded at a rate of 1 Hz (triggered) integrating over one single shot.

The pulse energy stability was 0.18 % rms over the entire 50 hour measurement, while the wavelength stability was 25 pm rms. The bandwidth at 10 % intensity level was 0.19 % rms. The

Optics EXPRESS

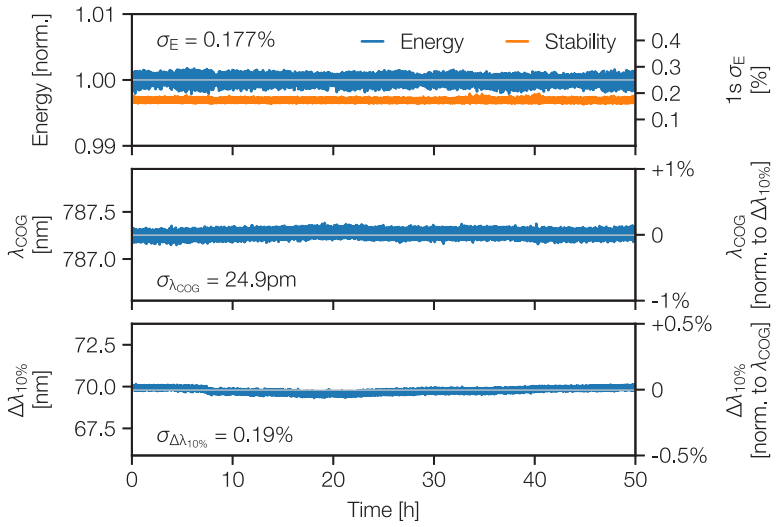

Fig. 7. Long term stability of the MALCOLM OPCPA seed laser, characterized with an out-of-loop diagnostics. To guide the eye, the grey lines show the respective mean values.

short-term energy jitter – calculated over 1 second, i.e., 1000 shots – remained below 0.19 % throughout the entire measurement, with an average of 0.17 %.

The slight drifts in the laser spectrum are actually caused by a subtle pointing drift within the laser, which affected the coupling into the spectrometer. We therefore expect the actual stability of the spectrum to be better than what is indicated in the plot.

The measured energy stability is largely limited by the resolution of our photodiode-based energy heads, and in fact the rms variation of the integrated counts of the single-shot spectra is only 0.13% over the entire 50-hour measurement. The integration over many spectrometer pixels results in higher resolution in counts and lower relative intrinsic noise than with the energy heads. Nevertheless, the value obtained from the spectrometer counts should be taken as an indication only, since the spectra are only recorded at 1 Hz. However, short-term measurements of the spectrum at the full 1 kHz repetition rate, show no difference in stability compared to measurements at 1 Hz, except for additional small drifts in the long-term measurement due to the aforementioned slow changes in coupling into the spectrometer.

6. Conclusion and outlook

In summary, we have demonstrated an automated tuning procedure of a two-stage OPCPA system that allows the laser system to operate at an optimum working point. The working point can be freely chosen. This not only speeds up the tuning procedure, but is also largely independent of the laser operator's experience with the system, and routinely allows higher output performance to be achieved than is possible with manual tuning. The tuning procedure can also be used to ensure a high degree of reproducibility of the laser's performance on a day-to-day basis. We have also shown that data collected during the optimization process can be used to calibrate a linear performance model of the laser system, which can be used to stabilize the laser output at the optimized working point using a full-state-feedback controller. Using this approach, we could demonstrate continuous operation of our laser system over 50 hours with excellent stability of pulse energy and spectral characteristics.

Our experiments show that even relatively simple methods that have been known for decades are effective for precise and reliable control of laser systems. This allows the full potential of a

Optics EXPRESS

laser system's performance to be exploited, despite highly complex couplings between various input and output parameters. As a result, highly complex laser systems such as optical parametric amplifiers can be used as reliable tools for science. We believe that this potential is not limited to OPCPA, but will also play a crucial role in other large laser systems that often face similar challenges in reproducibility.

Funding. Deutsche Forschungsgemeinschaft (491245950).

Acknowledgments. We thank Frida Brogren (DESY) for valueable discussions.

Disclosures. The authors declare no conflicts of interest.

Data availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

References

1. A. Dubietis, G. Jonušauskas, and A. Piskarskas, "Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal," Opt. Commun. **88**(4-6), 437–440 (1992).
2. C. Manzoni and G. Cerullo, "Design criteria for ultrafast optical parametric amplifiers," J. Opt. **18**(10), 103501 (2016).
3. G. Genty, L. Salmela, J. M. Dudley, D. Brunner, A. Kokhanovskiy, S. Kobtsev, and S. K. Turitsyn, "Machine learning and applications in ultrafast photonics," Nat. Photonics **15**(2), 91–101 (2021).
4. R. I. Woodward and E. J. R. Kelleher, "Towards 'smart lasers': self-optimisation of an ultrafast pulse source using a genetic algorithm," Sci. Rep. **6**(1), 37616 (2016).
5. C. Sun, E. Kaiser, S. L. Brunton, and J. N. Kutz, "Deep reinforcement learning for optical systems: A case study of mode-locked lasers," Mach. Learn.: Sci. Technol. **1**(4), 045013 (2020).
6. C. M. Valensise, A. Giuseppi, G. Cerullo, and D. Polli, "Deep reinforcement learning control of white-light continuum generation," Optica **8**(2), 239 (2021).
7. T. Eichner, T. Hülsenbusch, J. Dirkwinkel, T. Lang, L. Winkelmann, G. Palmer, and A. R. Maier, "Spatio-spectral couplings in saturated collinear OPCPA," Opt. Express **30**(3), 3404 (2022).
8. "DOOCS control system," https://doocs.desy.de.
9. R. J. Shalloo, S. J. D. Dann, and J.-N. Gruse, *et al.*, "Automation and control of laser wakefield accelerators using bayesian optimization," Nat. Commun. **11**(1), 6355 (2020).
10. S. Jalas, M. Kirchen, P. Messner, P. Winkler, L. Hübner, J. Dirkwinkel, M. Schnepp, R. Lehe, and A. R. Maier, "Bayesian optimization of a laser-plasma accelerator," Phys. Rev. Lett. **126**(10), 104801 (2021).
11. N. Hansen, A. Auger, R. Ros, S. Finck, and P. Pošík, "Comparing results of 31 algorithms from the black-box optimization benchmarking," in *Proceedings of the 12th annual conference companion on Genetic and evolutionary computation*, (2010), pp. 1689–1696.
12. O. Kramer, "Genetic algorithms," in *Genetic algorithm essentials*, (Springer, 2017), pp. 11–19.
13. N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," Evol. Comput. **9**(2), 159–195 (2001).
14. N. Hansen, "The cma evolution strategy: A comparing review," in *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea, eds. (Springer Berlin Heidelberg, 2006).
15. J. Blank and K. Deb, "pymoo: Multi-objective optimization in python," IEEE Access **8**, 89497–89509 (2020).
16. D. V. Arnold and H.-G. Beyer, *Noisy optimization with evolution strategies*, vol. 8 (Springer Science & Business Media, 2002).
17. B. Friedland, *Control system design: an introduction to state-space methods* (Courier Corporation, 2012).