Mechanisms of spatiotemporal mode-locking

Logan G. Wright [1,6]*, Pavel Sidorenko[1], Hamed Pourbeyram[1], Zachary M. Ziegler[1], Andrei Isichenko[1], Boris A. Malomed[2,3], Curtis R. Menyuk [4], Demetrios N. Christodoulides [5] and Frank W. Wise[1]

Mode-locking is a process in which different modes of an optical resonator establish stable synchronization through non-linear interactions. This self-organization underlies light sources that enable many modern scientific applications, such as ultrafast and high-field optics and frequency combs. Despite this, mode-locking has almost exclusively referred to the self-organization of light in a single dimension—time. Here we present a theoretical approach—attractor dissection—to understand three-dimensional spatiotemporal mode-locking. The key idea is to find a specific, minimal reduced model for each distinct type of three-dimensional pulse, and thus identify the important intracavity effects responsible for its formation and stability. An intuition for the results follows from the minimum loss principle, the idea that a laser strives to find the configuration of intracavity light that minimizes loss (maximizes gain extraction). Through this approach, we identify and explain several distinct forms of spatiotemporal mode-locking. These phases of coherent laser light have no analogues in one dimension and are supported by measurements of the three-dimensional field, which reveals spatiotemporal mode-locked states that comprise more than 10^7 cavity modes. Our results should facilitate the discovery and understanding of new higher-dimensional forms of coherent light which, in turn, may enable new applications.

L ight fields created by passive mode-locking (ML) are unmatched for ultrashort duration, peak electric field strength, energy localization, low noise and frequency precision. These unique features make ML lasers essential for frontier science and technology in physics, biology, chemistry and materials science[1–4]. Today, practically viable implementations based on fibre or semiconductor gain media hold promise to bring these frontier-pushing technologies to mainstream use[3,5,6] as the state-of-the-art capabilities continue their remarkable growth[7].

ML in lasers was demonstrated shortly after the laser itself[8] and understood in the frequency domain as the synchronization of longitudinal modes. Longitudinal modes are patterns of the electromagnetic field structured along the length of the resonator, whose resonant lasing frequencies differ by a free spectral range, $\Delta = c/P$, where P is the optical path length of one round trip of the resonator. ML occurs when the laser oscillates simultaneously in multiple modes and the phase difference between them is stable. Of course, laser resonators may support many transverse as well as longitudinal modes. The optical path length for each transverse mode is different, so their free spectral range is different too (Supplementary Fig. 1b). It was realized early on that if several transverse modes have the right relationship between resonances, multitransverse-mode ML may occur. This was dubbed 'transverse ML' for the case that involved a single longitudinal mode, and 'total ML' when multiple modes of each kind were involved[8,9].

Why does ML occur? A now widespread intuition is that it represents the laser's solution to an optimization problem defined by the cavity. When lasing is initiated, field fluctuations present in all the modes are amplified simultaneously. All complex combinations of the laser's modes compete for the limited energy available from the gain medium. In passive ML, a saturable absorber (SA, an element with loss inversely related to the optical power) is added to the laser cavity. Then, the minimum-loss state becomes one in which

many longitudinal modes lase simultaneously, in phase, as this corresponds to a short pulse in time with a maximum peak power. As it experiences the least loss each round trip, this lowest-loss configuration is preferentially amplified and eventually dominates all the other possible mode configurations.

As the study of ML matured, Ginzburg–Landau equations[10–12] were adopted in recognition of the importance of phase modulations in addition to gain and loss. These models reveal the mode-locked pulse to be a (dissipative) soliton, most commonly when the anomalous group velocity dispersion of the resonator cancels the non-linear phase shift that arises from the Kerr non-linearity in the laser medium. Later, it was realized that ML also occurs in lasers with normal dispersion[12]. In this case, the pulses carry a linear chirp, and a spectral filter is used to periodically shorten the chirped pulse by attenuating edges of the spectrum (and, owing to the chirp, wings of the pulse in time).

Some works considered aspects of higher dimensions in ML[13–16], notably in the Kerr lens realization, in which weak coupling to higher-order modes is essential[17–21]. These works still view ML as a low-dimensional phenomenon. A notable exception is work on random lasers[22–24], but to date coherent self-organization in these systems has been limited. Finally, early studies of multidimensional ML[8,9] predated insights from non-linear dynamics and so represent only a minute fraction of the possibilities. We recently reported observations of spatiotemporal mode-locking (STML)[25] in the form of three-dimensional (3D) coherent pulses in lasers based on multimode fibres (MMFs). Although this work showed self-organized pulses in 3D lasers, our physical understanding was limited. The investigation of STML presents acute theoretical and experimental challenges. A direct simulation can be significantly complicated by the presence of disorder, spatiotemporal dispersion and multiscale non-linear interactions among a population of modes 10 to 10^4 times larger than those in traditional ML. Similar challenges occur

[1]School of Applied and Engineering Physics, Cornell University, Ithaca, NY, USA. [2]Department of Physical Electronics, School of Electrical Engineering, Faculty of Engineering, and the Center for Light-Matter Interaction, Tel Aviv University, Tel Aviv, Israel. [3]ITMO University, St Petersburg, Russia. [4]Department of Computer Science and Electrical Engineering, University of Maryland Baltimore County, Baltimore, MD, USA. [5]CREOL/College of Optics and Photonics, University of Central Florida, Orlando, FL, USA. [6]Present address: NTT Physics and Informatics Laboratory, NTT Research, Inc., East Palo Alto, CA, USA. *e-mail: lgw32@cornell.edu
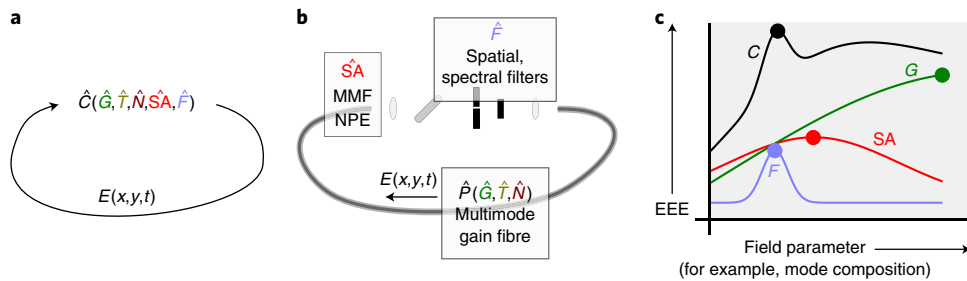
Fig. 1 | Conceptual outline of STML, attractor dissection and the spatiotemporal maximum-gain principle. a, A laser cavity can be described by a non-linear projection operator $\hat{C}$, which represents the entirety of the effects of propagation in the course of one round trip through the cavity on the intracavity electric field (equation (1)). **b**, How the different effects that compose $\hat{C}$ manifest in our experiments and our numerical model of a MMF laser. Here $\hat{SA}$, $\hat{F}$ and $\hat{P}$ are operators that describe the effects of the spatiotemporal SA, intracavity SF (F) and non-linear gain fibre ($\hat{P}$, which includes the inseparable effects of 3D gain, $\hat{G}$, dispersion, $\hat{T}$ and non-linear mode coupling, $\hat{N}$). In our experiment, the spatiotemporal SA is implemented by nonlinear polarization evolution (NPE) in the multimode fibre. **c**, An approximate intuition for the laser's route to steady state, an EEE surface C, which is the product of the EEE curves that arise from the component effects shown, with respect to an arbitrary field parameter (for example, the amplitude in a given 3D mode). EEE describes how much energy a field with given parameters extracts from the gain medium. A (spatiotemporally) mode-locked pulse $A(x,y,t)$ corresponds to an eigenpulse of $\hat{C}$: $A(x, y, t) = \hat{C}A(x, y, t)$, and is an attractive fixed point of the iterative map, that is, $A(x, y, t) = \lim_{n \to \infty} \hat{C}^n E(x, y, t)$. The eigenpulse typically maximizes EEE. Through attractor dissection, we identify the laser gain competition as an optimization approximately factored into subproblems. The example displayed in **c** illustrates a SF-dominated regime, in which the coordinates of the expected steady-state solution (the maximum of C, black circle) are closest to the attractor of F. In other regimes, G and SA play more important roles.

for applying traditional 1D measurement techniques to STML oscillators, as virtually all the standard experimental techniques implicitly assume a spatially single-mode (that is, 1D) field. Thus, our understanding of STML was restricted to observational reports with careful, but limited, adaptations of traditional theoretical and experimental methods.

We outline here a mechanistic description of STML (Fig. 1) based on a new theoretical approach that we call attractor dissection. We constructed simplified models for 3D mode-locked lasers to identify the key mechanisms for several novel 3D laser solitary waves. This approach was adapted from workhorse techniques of laser physics[10,26–28] and its results inform a similar intuition regarding the laser's steady state. That is, STML steady states can be intuited as the result of a multidimensional optimization of the laser field towards the eigenstate of the cavity that most effectively extracts energy from the gain medium. To address the experimental challenges, we utilized a flexible approach for 3D electric-field measurement, scanning off-axis digital holography, and performed an experimental modal decomposition.

Attractor dissection theory

We now outline the attractor dissection theory and apply it to understand the distinct regimes of the steady-state 3D pulse formation that occur as we vary one parameter of the cavity, the intracavity spatial filter (SF) size. Further details are given in Methods and Supplementary Sections 1–4. We consider the laser cavity as a non-linear projection operator that represents the transformation of the laser-field envelope, $E_i(x,y,t)$, through its propagation along the longitudinal coordinate z of the cavity through one round trip (Fig. 1a). The cavity operator $\hat{C}$ performs the mapping:

$$E_{i+1}(x, y, t) = \hat{C}E_i(x, y, t) \tag{1}$$

where the subscript of E is the round-trip number. $\hat{C}$ accounts for several different effects that are, in general, coupled inseparably. In this view, the laser physics is a composition of iterated non-linear projection operations. In each round trip, the non-linear dissipation of $\hat{C}$ selects from the field certain attributes, and the saturable laser gain provides a conditional (frequency- and energy-limited, and spatially localized) rescaling of the selected field.

We first considered a simplified multimode laser cavity containing seven transverse modes (listed in the legend at the top of Fig. 2), which consists of a normal-dispersion multimode gain fibre, an idealized spatiotemporal SA (a SA sensitive to the local field intensity in space–time, $|E(x, y, t)|^2$), a spectral filter and a SF (Fig. 1b). The fibre is a model of that used in the experiments, whose key feature is its relatively low modal dispersion. As the filtering, SA and fibre propagation occur in spatially separate regions, $\hat{C}$ can be factored into $\hat{C} = \hat{F}(x, y)\hat{F}(\omega)\hat{SA}(x, y, t)\hat{P}(x, y, t)$, where $\hat{F}(x, y)$ and $\hat{F}(\omega)$ are the spatial and spectral filter functions, $\hat{SA}(x, y, t)$ is the spatiotemporal SA transfer function and $\hat{P}(x, y, t)$ accounts for the effect of the pulse propagation through the 3D non-linear gain medium. $\hat{P}$ includes the inseparable effects of 3D gain, $\hat{G}$, linear effects, such as spatiotemporal dispersion, $\hat{T}$, and non-linear mode coupling $\hat{N}$. $\hat{P}$ is implemented by integrating a set of seven coupled non-linear partial differential equations.

Holding other parameters fixed, we simulated this cavity, starting from noise, with varying Gaussian SF sizes (positioned for simplicity on-axis with the fibre). Figure 2 shows the spatial mode composition of the steady-state pulses as a function of the SF width (solid dots). We refer to these stable 3D fields, $A(x,y,t)$, which are fixed points of equation (1), as eigenpulses of the cavity. This terminology emphasizes that they are eigenfunctions of the cavity operator: $\hat{C}A(x, y, t) = A(x, y, t)$ (but are not necessarily solitons, as occur in models that take the continuous z limit of equation (1)).

The key idea of attractor dissection is that, despite the many effects that occur on each round trip, a given eigenpulse may be principally selected by far fewer. This allows us to construct minimal models of the laser in which $\hat{C} = \hat{R}\hat{O}$, where $\hat{O}$ describes the effect(s) relevant to the specific eigenpulse's attractor, and $\hat{R}$ is a rescaling operator that acts as a simplified saturating gain, returning the field's energy to a fixed value once per round trip. $\hat{O}$ may be any effect or combination of effects in the cavity. By finding for a given eigenpulse the minimal $\hat{O}$ necessary to approximate it, we can identify the key effects responsible for its formation and stability. The fixed point attractor of the minimal model $\hat{C}$, which we call the field attractor of $\hat{O}$, provides insight into the role of $\hat{O}$ within the full laser dynamics (that is, what type of eigenpulses does $\hat{O}$ 'prefer'?). We can find the field attractors by computing the asymptotic solution to repeated iterations of $\hat{C}$ starting from a noise field, $E_0(x,y,t)$:
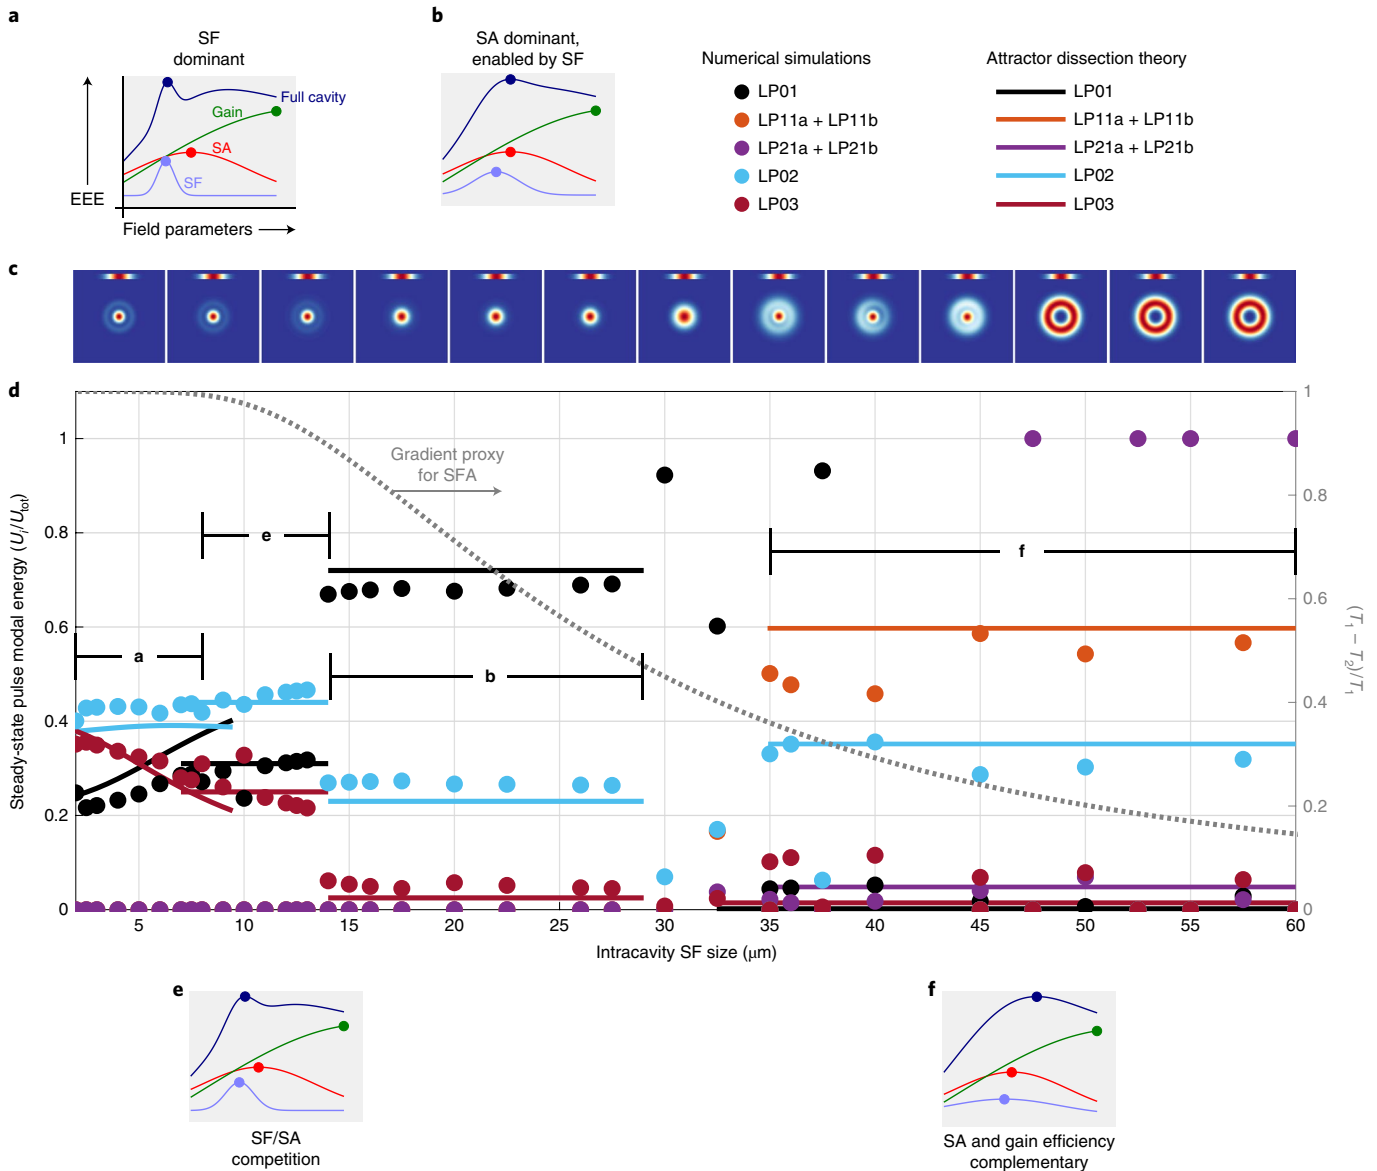
Fig. 2 | Identifying the mechanisms of 3D mode-locked pulses for varying intracavity SF size. a–f, The plot in **d** shows the numerically calculated mode composition (solid points show the energy in each transverse-mode family relative to the total pulse energy) of the steady-state pulse solution for a MMF-based cavity, for varying sizes of the SF. The theoretical modal compositions of the relevant attractor models are plotted alongside the numerical points (solid curves). For each regime, representative EEE surface 'cartoons' (**a,b,f**) are shown to provide intuition for the attractors. The right vertical axis (grey) shows the gradient proxy (see main text) for the SF's attractor (SFA), which estimates the relative importance of the SF in maximizing EEE. As the SF is increased, the steady-state operating regime changes from being dominated by the SF (**a**) to a SA-driven regime enabled by the SF (**b**), to competition between the SA and SF attractors (**e**) and finally to a regime characterized by a compromise between minimizing the SA loss while maximizing the gain efficiency (**f**). **c,** Representative beam profiles of steady-state pulses are shown above the plot in **d** for reference. Scale bars, Gaussian profile of the fundamental transverse mode (35 μm mode field diameter). An extended version of the figure and representative details of the steady-state regimes are presented in Supplementary Figs. 10–13.

$$A(x,y,t) = \lim_{n \to \infty} [\hat{R}\hat{O}]^n E_0(x,y,t) \qquad (2)$$

As an example, consider the field attractor of the SF. If we take $\hat{O} = \hat{F}$ in equation (2) and solve with the seven transverse modes of the cavity described above, the solution $A(x,y,t)$ has a straightforward interpretation. In the basis of the guided transverse modes of the fibre cavity, $\hat{F}$ takes the form of a matrix F with elements $F_{mn} = \langle \varphi_m \mid F(x,y) \mid \varphi_n \rangle$, where $F(x,y)$ is the transmission function of the filter and $\varphi_i = \varphi_i(x,y)$ are the transverse modes. The matrix F has eigenvectors $\boldsymbol{\nu}_i$, which are complex mode coefficients for fields that

pass through the filter with only a uniform loss, that is: $F\boldsymbol{\nu}_i = t_i \boldsymbol{\nu}_i$, where $T_i = |t_i|^2$ is the energy transmission through the filter. The relevant fixed point of equation (2), the attractor of $\hat{F}$, is the lowest-loss eigenvector of the SF in the guided mode basis (with a uniform distribution in time). The curves in Fig. 2 show the modal composition of this spatial filter attractor (SFA) alongside full numerical simulations from 1 to 9 μm SF sizes. We observe that, even though the dissected attractor model neglects most of the laser physics, it accurately predicts the composition of the eigenpulse in this regime. Thus, we identify the SF as the key mechanism for the eigenpulses in this regime, which can be referred to as SF-driven pulses.

We repeated this procedure for the spatiotemporal SA, where we used a minimal model $\hat{O} = \hat{S}\hat{A}\hat{T}$ (where $\hat{T}$ describes the linear spatiotemporal dispersion of the modes). We found that two distinct fixed point solutions to equation (2) exist, depending on whether the peak intensity of the pulse is below or above the saturation intensity of the SA. These dissected attractor models accurately predict the full numerical simulations for filter sizes between about 10 and 30 μm in Fig. 2 (thus, those eigenpulses can be mechanistically described as SA-driven pulses). Similarly, we found that a minimal model for many of the eigenpulses in the large-filter regime uses $\hat{O} = \hat{S}\hat{A}\hat{P}(\hat{T}, \hat{G})$, where $\hat{P}(\hat{T}, \hat{G})$ is the gain fibre-propagation operator, neglecting non-linear effects other than the gain, $\hat{G}$.

Maximum gain principle

To see how the dissected attractor theory mechanistically explains 3D laser physics, a helpful intuition is similar to the minimum-loss principle (this principle is established wisdom in laser science[29]). In a 3D laser, the expected winner of the gain competition is the field that, through a combination of low loss and high overlap with the gain medium, maximizes the energy extraction efficiency (EEE) from the gain medium (that is, maximizes the total energy gained by the field).

The idea of an expected winner of the gain competition is an effective intuition for the steady state in a 3D mode-locked laser: the steady state is the laser's attempt to solve the optimization problem of maximizing EEE while satisfying the periodic boundary condition and starting from low-amplitude noise. The path to a steady state can be visualized as a maximization on an abstract surface in the $2N$-dimensional space of the laser field's 3D mode coefficients, with N the number of 3D modes (Fig. 1c). We can imagine this EEE surface as being composed of distinct contributions from each effect in the cavity, that is, as a problem approximately factored into distinct subproblems. Each dissected subsurface has a maximum that corresponds to that effect's attractor, which is the steady-state field that minimizes (maximizes) the loss (gain) from that effect. The 1D EEE surface cartoons in Figs. 1c and 2 illustrate the intuition: subsurfaces cooperate or compete in producing the EEE maximum, whose coordinates specify the expected winner of the laser gain competition. In some cases, as with SF-driven pulses, the maximum of EEE is mostly determined by the maximum of a single subproblem. In most scenarios, however, multiple effects cooperate or compete to form the maxima of EEE.

A quantitative measure of the importance of an effect is the gradient of the EEE surface near its maximum. The solid grey curve in Fig. 2 (right y axis) shows a proxy measure of the gradient of the EEE surface around the SF's attractor: the normalized difference in transmission between the first and second filter eigenfunctions, $(T_1 - T_2)/T_1$. As the SF expands, the EEE gradient diminishes (from 5 to 10 μm filter widths, visualized in Fig. 2e). Thus, minimizing the loss through the SF becomes less important for maximizing EEE. Eventually, the SF and SA effects have nearby, similar-height contributions to the EEE surface, so eigenpulses form between the respective attractors. As the filter expands, the SA becomes more important to the EEE maximization, and the eigenpulses then resemble the SA's attractor (Fig. 2b). Finally, as the filter is made broader still, eigenpulses change from comprising the few radially symmetric modes (linear polarized LP modes LP01, LP02 and LP03) to a diverse set of higher-order modes. As implied by the EEE surfaces (Fig. 2f), this transition reflects the growing importance of the field's overlap with the 3D gain medium.

We call the eigenpulses in the large-filter regime SAGE (saturable absorber-gain efficiency) pulses, as in this regime, EEE maximization is approximately factored into (i) maximizing the gain efficiency via overlap with the gain medium (accomplished by occupying large-area, high-order modes) and (ii) minimizing the SA loss (accomplished by forming short, intense pulses). This is a multifaceted compromise, and leads to a range of novel 3D physics, which we summarize here (further details in Supplementary Sections 5–8). As the gain is slow (temporally non-local), EEE maximization can lead initially to numerous distinct 3D pulses that form independently from noise, which independently minimize the SA loss and collectively maximize the overlap with the gain medium. These pulses propagate with different group velocities (they contain different modes), so they eventually collide. Through these collisions, single-mode SAGE pulses form as the multimode pulses cross-modulate each other through the spatiotemporally local SA (Supplementary Video 1). Intuitively, single-mode pulses are promoted by the SA due to their minimum modal dispersive broadening. These single-mode SAGE pulses, which are visible in Fig. 2, for example, at 47.5 and 55 μm filter size in the LP21 mode group, are practically important. They reach a higher energy than is possible in single-mode oscillators, are free of spatiotemporal (modal) dispersion and, due to their unique mechanism, should also occur even in step-index fibres (Supplementary Section 8). They are, however, not the expected winner of the gain competition: in Fig. 2, the composition of these pulses deviates significantly from the expected SAGE attractor. Local–global competition of this kind is one reason that STML oscillators lack the same distinction between single- and multipulse regimes as that of conventional oscillators[30], and single-mode SAGE pulses highlight the limits of the approximate dissected attractor theory and EEE maximization intuition. For a given cavity configuration in the SAGE regime, other qualitatively different, multistable steady-states emerge as a result of different initial noise conditions. For example, we also observed multimode SAGE pulses that displayed a periodic breathing. In some cases, they are higher-order eigenpulses that satisfy $A(x, y, t) = \hat{C}^n A(x, y, t)$ for integer n, whereas others display breathing incommensurate with the cavity period due to an unconstrained transverse mode beating (Supplementary Section 7).

Building up reduced models of complex attractors

A motivation for attractor dissection is that it naturally extends to more advanced, but still reduced, models to describe experiments. The simplest models—the ones considered so far—are useful to identify the primary mechanisms for distinct 3D ML regimes in full numerical simulations with seven transverse modes. Simulations of fibre that supports 180 transverse modes (as in the experiments described below) is challenging, especially because of the complex role of disorder—an experimental feature we find essential for accurate modelling. Although these reduced models fail to describe the details of ultrafast pulse propagation, they permit efficient simulations and provide insight into experiments, which clarifies how effects like disorder affect the regimes identified in the few-mode, disorder-free models. To compare with experiments, we used a reduced model for $\hat{C} = \hat{R}\hat{T}_r\hat{F}\hat{S}\hat{A}\hat{G}$, where $\hat{G}$ is an approximation of the saturating gain operator and $\hat{T}_r$ is the linear transfer function of the cavity, extended to include random linear mode coupling (Methods and Supplementary Section 4 give details).

Experiments

Experimentally, we observed STML in lasers constructed with multimode graded-index Yb:doped gain fibre and some free-space sections, similar to those of Wright et al.[25]. The laser incorporates an adjustable SF, a spectral filter and a spatiotemporal SA, implemented using non-linear polarization evolution (NPE) in the multimode gain fibre. For a suitable choice of waveplate orientations in the cavity, NPE amounts to the action of an effective spatiotemporal SA, which imposes a loss inversely related to the spatiotemporal intensity of the circulating field. The electric field is measured in 3D using scanning off-axis digital holography (Methods) and is then decomposed into the transverse modes of the fibre (Fig. 3).
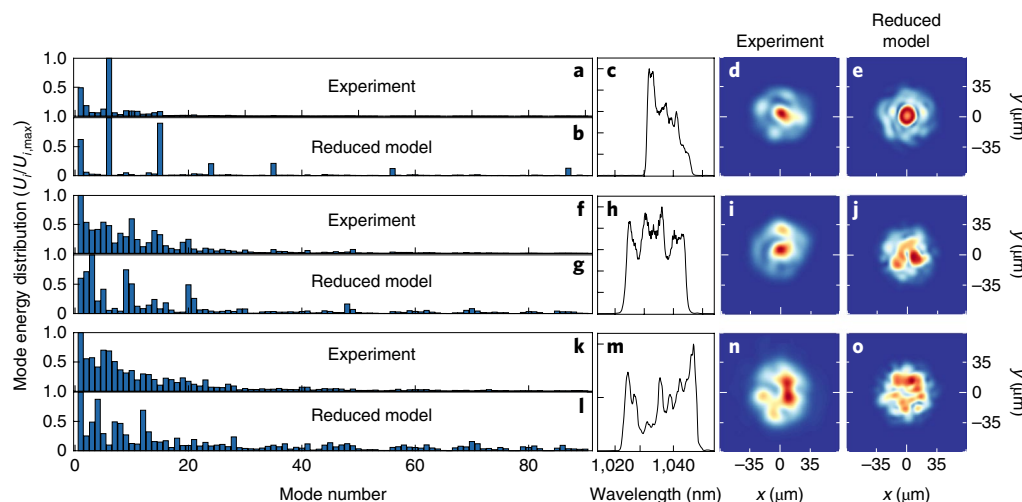
Fig. 3 | Experimental regimes of STML and results from a reduced laser model. We measured the 3D field of stable mode-locked pulses in a normal-dispersion multimode fibre laser for various intracavity SF dimensions. **a–e**, Results for a pulse with a narrow SF in the cavity (roughly a SF-driven pulse), showing the experimental mode decomposition of the 3D field into the transverse fibre modes (**a**), the same for a reduced model calculation (**b**), the experimental pulse spectrum (note the y axis is normalized to the mode with the maximum energy) (**c**) and the near-field beam profiles of the experimental (**d**) and reduced (**e**) model pulses. **f–o**, the same series of plots for mode-locked states with an intermediate-sized intracavity SF (**f–j**) and for a cavity without deliberate spatial filtering (roughly, SAGE pulses) (**k–o**). An extended version showing more typical experimental results is given in Extended Data Fig. 1, whereas Supplementary Figs. 16 and 17 show additional measurements and Supplementary Fig. 19 shows representative error estimations (see Methods). Here the reduced model neglects the Kerr non-linearity of the fibre and employs approximate operators to efficiently describe the saturating gain and linear fibre disorder. Despite these approximations, the reduced model reproduces coarse trends of the experiments, such as the reducing symmetry and broader modal distribution as the filter size is increased.

As we vary the SF experimentally, we observed, as in the seven-mode simulations in Fig. 2, various transitions in the STML steady-state pulses (Fig. 3 and Extended Data Fig. 1). These pulses have similar spectra, pulse duration, energy and other features as those in our simulations. SF-driven pulses (narrow SF) are narrowband, with triangular spectra and radially symmetric beams. With intermediate filter sizes, the pulses are similar to SA-driven pulses; they have beam profiles with near-radial symmetry and broad rectangular spectra, and concentrate energy in low-order modes. SAGE (broad or no SF) pulses are broadband and energetic, with rectangular spectra and multimode beam profiles. Qualitative dynamical features, such as multistability and multiperiodicity, were observed; much like the seven-mode simulations, these occur primarily in the large-filter regime.

Beyond these qualitative trends, the experimental results are more complex (Fig. 3, with more details in Supplementary Sections 4, 6 and 7). Reduced models are helpful for predicting how the effects of disorder, and the increased dimension of the optimization, affect the regimes of STML. Both in the experiment and reduced models, we see that with narrow SFs, pulses still comprise mostly radially symmetric modes. As the filter is widened, however, disorder and modal diversity incrementally result in more complex pulses. We routinely observed SAGE pulses that contained over 30 million locked 3D modes, about 10 times more than in any previous mode-locked oscillator[31,32]. This quantitative comparison fails to capture the qualitative strangeness of these pulses, however. Characterized by a disordered spatiotemporal structure, the experimental SAGE pulses nonetheless exhibit remarkable long-range order (evidenced by a high radio frequency contrast >80 dB (Supplementary Figs. 16 and 17). They are robust, with many being self-starting and stable over long durations and perturbations in the laboratory. Although principal modes of disordered fibres[33,34] suggest a first guess towards a detailed model of these pulses, reconciling the surprising complexity with the full multimode non-linear dynamics and periodic boundary conditions of the cavity will—as for many other new phenomena in 3D lasers—require advances in both theoretical and experimental methods.

Discussion

The present study focuses on a single narrow swath of parameter space—the variation of the SF in a normal dispersion, MMF-based STML oscillator. STML oscillators present many possibilities for exploring multidimensional non-linear dynamics and solitary waves, and fundamentally new forms of coherent light. For example, STML states that approximate multimode solitons[35,36] should form with anomalous dispersion, and other spatial and spatiotemporal attractors[13,37,38] should be observable by a judicious design of the cavity. Meanwhile, high-performance lasers should be possible in settings that range from few-mode step-index fibres (Supplementary Section 8) to free-space cavities with rod-like fibres, thin disks or other bulk gain media. Realizing STML with more systematic intracavity control (for example, intracavity spatial light modulators or deformable mirrors) and novel spatiotemporal SAs will be an important step. Advances in 3D measurement tools will be crucial for insightful experimental studies, and, hopefully, automated closed-loop control of STML lasers[39–42].

Although our intention here was to use the optimization-problem perspective to derive intuitive models for 3D laser physics, reconfigurable 3D lasers are promising as analogue computers and heuristic optimizers. Even at a modest scale of ~100 transverse modes, the cavity's natural computations are ~10–100 billion times faster (and require ~10^{13} times less energy) than full numerical simulations on a digital computer. 3D lasers have many more controllable degrees of freedom than 1D oscillators, and have been described approximately with spin-glass models[23,43] similar to coherent Ising machines[44]. Although the physics of competitive optimization (that is, Darwinism) is itself a universal model for a wide range of physics, from materials to biological evolution, extending mathematical spin-glass and synchronization analogies will be helpful for using 3D lasers as heuristic optimization and analogue simulators. For

experimentally realizing these and other applications, the control of coupling using active intracavity modulation[45] or by engineered modes in semiconductor lasers[46], multimode microresonators[47] or coupled microresonators[48] is promising.

Online content

References

1. Popmintchev, T. et al. Bright coherent ultrahigh harmonics in the keV X-ray regime from mid-infrared femtosecond lasers. *Science* **336**, 1287–1291 (2012).
2. England, R. J. et al. Dielectric laser accelerators. *Rev. Mod. Phys.* **86**, 1337–1389 (2014).
3. Xu, C. & Wise, F. W. Recent advances in fiber lasers for nonlinear microscopy. *Nat. Photon.* **7**, 875–882 (2013).
4. Cundiff, S. T. & Ye, J. Colloquium: femtosecond optical frequency combs. *Rev. Mod. Phys.* **75**, 325–342 (2003).
5. Orringer, D. A. et al. Rapid intraoperative histology of unprocessed surgical specimens via fibre-laser-based stimulated Raman scattering microscopy. *Nat. Biomed. Eng.* **1**, 0027 (2017).
6. Tilma, B. W. et al. Recent advances in ultrafast semiconductor disk lasers. *Light Sci. Appl.* **4**, e310–e310 (2015).
7. Fattahi, H. et al. Third-generation femtosecond technology. *Optica* **1**, 45–63 (2014).
8. Smith, P. W. Mode-locking of lasers. *Proc. IEEE* **58**, 1342–1357 (1970).
9. Auston, D. Transverse mode locking. *IEEE J. Quantum Electron.* **4**, 420–422 (1968).
10. Haus, H. A. Mode-locking of lasers. *IEEE J. Sel. Top. Quantum Electron.* **6**, 1173–1185 (2000).
11. Krausz, F. et al. Femtosecond solid-state lasers. *IEEE J. Quantum Electron.* **28**, 2097–2122 (1992).
12. Renninger, W., Chong, A. & Wise, F. Dissipative solitons in normal-dispersion fiber lasers. *Phys. Rev. A* **77**, 023814 (2008).
13. Akhmediev, N., Soto-Crespo, J. M. & Grelu, P. Spatiotemporal optical solitons in nonlinear dissipative media: from stationary light bullets to pulsating complexes. *Chaos* **17**, 037112 (2007).
14. Renninger, W. H. & Wise, F. W. Spatiotemporal soliton laser. *Optica* **1**, 101–104 (2014).
15. Kutz, J. N., Conti, C. & Trillo, S. Mode-locked X-wave lasers. *Opt. Express* **15**, 16022–16028 (2007).
16. Büttner, T. F. S. et al. Multicore, tapered optical fiber for nonlinear pulse reshaping and saturable absorption. *Opt. Lett.* **37**, 2469–2471 (2012).
17. Negus, D. K., Spinelli, L., Goldblatt, N. & Feugnet, G. Sub-100 femtosecond pulse generation by Kerr lens mode-locking in Ti:Al₂O₃. In *Proc. Topical Meeting on Advanced Solid State Lasers* (eds Dubé, G. & Chase, L.) SPL7 (OSA, 1991).
18. Herrmann, J. Theory of Kerr-lens mode locking: role of self-focusing and radially varying gain. *J. Opt. Soc. Am. B* **11**, 498–512 (1994).
19. Kalosha, V. P., Müller, M., Herrmann, J. & Gatz, S. Spatiotemporal model of femtosecond pulse generation in Kerr-lens mode-locked solid-state lasers. *J. Opt. Soc. Am. B* **15**, 535–550 (1998).
20. Jirauschek, C., Kärtner, F. X. & Morgner, U. Spatiotemporal Gaussian pulse dynamics in Kerr-lens mode-locked lasers. *J. Opt. Soc. Am. B* **20**, 1356–1368 (2003).
21. Dunlop, A. M., Firth, W. J. & Wright, E. M. Master equation for spatio-temporal beam propagation and Kerr lens mode-locking. *Opt. Commun.* **138**, 211–226 (1997).
22. Leonetti, M., Conti, C. & Lopez, C. The mode-locking transition of random lasers. *Nat. Photon.* **5**, 615–617 (2011).
23. Antenucci, F., Crisanti, A., Ibáñez-Berganza, M., Marruzzo, A. & Leuzzi, L. Statistical mechanics models for multimode lasers and random lasers. *Philos. Mag.* **96**, 704–731 (2016).
24. Nixon, M. et al. Real-time wavefront shaping through scattering media by all-optical feedback. *Nat. Photon.* **7**, 919–924 (2013).
25. Wright, L. G., Christodoulides, D. N. & Wise, F. W. Spatiotemporal mode-locking in multimode fiber lasers. *Science* **358**, 94–97 (2017).
26. Fox, A. G. & Tingye, Li Modes in a maser interferometer with curved and tilted mirrors. *Proc. IEEE* **51**, 80–89 (1963).
27. Kogelnik, H. & Li, T. Laser beams and resonators. *Appl. Opt.* **5**, 1550–1567 (1966).
28. Siegman, A. E. *Lasers* (University Science Books, 1986).
29. Haken, H *Light: Laser Light Dynamics* (North-Holland, 1985).
30. Qin, H., Xiao, X., Wang, P. & Yang, C. Observation of soliton molecules in a spatiotemporal mode-locked multimode fiber laser. *Opt. Lett.* **43**, 1982–1985 (2018).
31. Matos, L. et al. Direct frequency comb generation from an octave-spanning, prismless Ti:sapphire laser. *Opt. Lett.* **29**, 1683–1685 (2004).
32. Renninger, W. H., Chong, A. & Wise, F. W. Giant-chirp oscillators for short-pulse fiber amplifiers. *Opt. Lett.* **33**, 3025–3027 (2008).
33. Xiong, W. et al. Principal modes in multimode fibers: exploring the crossover from weak to strong mode coupling. *Opt. Express* **25**, 2709–2724 (2017).
34. Carpenter, J., Eggleton, B. J. & Schröder, J. Observation of Eisenbud–Wigner–Smith states as principal modes in multimode fibre. *Nat. Photon.* **9**, 751–757 (2015).
35. Renninger, W. H. & Wise, F. W. Optical solitons in graded-index multimode fibres. *Nat. Commun.* **4**, 1719 (2013).
36. Wright, L. G., Renninger, W. H., Christodoulides, D. N. & Wise, F. W. Spatiotemporal dynamics of multimode optical solitons. *Opt. Express* **23**, 3492–3506 (2015).
37. Guenard, R. et al. Kerr self-cleaning of pulsed beam in an ytterbium doped multimode fiber. *Opt. Express* **25**, 4783–4792 (2017).
38. Guenard, R. et al. Nonlinear beam self-cleaning in a coupled cavity composite laser based on multimode fiber. *Opt. Express* **25**, 22219–22227 (2017).
39. Florentin, R. et al. Shaping the light amplified in a multimode fiber. *Light Sci. Appl.* **6**, e16208 (2016).
40. Tzang, O., Caravaca-Aguirre, A. M. & Piestun, R. Wave-front shaping in nonlinear multimode fibers. *Nat. Photon.* **12**, 368–374 (2018).
41. Iegorov, R., Teamir, T., Makey, G. & Ilday, F. Ö. Direct control of mode-locking states of a fiber laser. *Optica* **3**, 1312–1315 (2016).
42. Woodward, R. I. & Kelleher, E. J. R. Towards 'smart lasers': self-optimisation of an ultrafast pulse source using a genetic algorithm. *Sci. Rep.* **6**, 37616 (2016).
43. Angelani, L., Conti, C., Ruocco, G. & Zamponi, F. Glassy behavior of light. *Phys. Rev. Lett.* **96**, 065702 (2006).
44. Wang, Z., Marandi, A., Wen, K., Byer, R. L. & Yamamoto, Y. Coherent Ising machine based on degenerate optical parametric oscillators. *Phys. Rev. A* **88**, 063853 (2013).
45. Weill, R., Fischer, B. & Gat, O. Light-mode condensation in actively-mode-locked lasers. *Phys. Rev. Lett.* **104**, 173901 (2010).
46. Gustave, F. et al. Observation of mode-locked spatial laser solitons. *Phys. Rev. Lett.* **118**, 044102 (2017).
47. Lucas, E. et al. Spatial multiplexing of soliton microcombs. *Nat. Photon.* **12**, 699–705 (2018).
48. Jang, J. K. et al. Synchronization of coupled optical microresonators. *Nat. Photon.* **12**, 688–693 (2018).

Methods

Theory. The main numerical and theoretical models are outlined in Supplementary Sections 2–4. For our few-mode simulations, we consider only 7 modes: the first 6 modes, plus the 15th one (LP03).

Simulations of the non-linear wave equation. For simulations presented in the main article, we considered a 50 cm long fibre, which was chosen to be a similar order of magnitude to that of the experiment, but short enough to permit thousands of round trips to be efficiently simulated for each data point in Fig. 2. We took the gain to be a spatially saturating function of the fluence (the integral of the intensity over the temporal window) at each point in the transverse plane. The spatiotemporal SA was represented by an idealized transfer function with a modulation depth of 1 and a saturation intensity of 50 GW cm^{-2}. Both numbers are of similar magnitude to those routinely used to describe normal-dispersion single-mode fibre lasers based on NPE. Besides the parameters noted, we assumed an output coupling ratio of 0.7, an additional lumped loss of 0.5 and a Gaussian spectral filter of 10 nm full-width-at-half-maximum. In addition, we rescaled the field after passing through the SF. This artificial step was done to keep the lasing threshold approximately constant as the filter size was varied, and thus allow the SF size to be adjusted without changing the gain saturation energy. As dissipative projection followed by the energy rescaling of the saturating gain is the essence of the laser physics, this approximation does not change the physics qualitatively.

Dissected attractor laser models for component attractors and reduced laser models. Supplementary Sections 3 and 4 describe the calculation of the attractors of each intracavity effect, and extensions of these approximate laser models to more complex attractors within slightly more realistic laser models. Like the main theoretical model, these ones describe the laser as the action of an iterative non-linear projection operation on the intracavity field. The key difference is that only one or a small number of specific effects are considered alongside idealized implementations of the laser cavity's physics, linear modal propagation and rescaling (from the saturating gain). In other words, all such models are based on an equation of the form:

$$A(x, y, t) = \lim_{n \to \infty} \hat{C}^n E_0(x, y, t) = \lim_{n \to \infty} [\hat{R}\hat{O}]^n E_0(x, y, t) \qquad (3)$$

where $\hat{O}$ represents an effect (or small number of effects) of interest, $E_0(x,y,t)$ is an initial noise field and $\hat{R}$ is the unconditional rescaling operation that represents an idealized saturating laser gain. $\hat{R}$ amounts to the following:

(1) Calculate the total energy of the field, $U = \iiint |E(x, y, t)|^2 \, dxdydt$.
(2) Multiply the field by $\sqrt{U_0/U}$, so that the energy of the field is restored to a constant U_0 regardless of the losses or gains elsewhere.

In most of the minimal models, we need the operator $\hat{T}$. $\hat{T}$ is the linear transmission operator that describes linear propagation (for example, the spatiotemporal dispersion) of the cavity's guided modes (defined by the presence of the fibre, the bounding mirrors and so on) through one round trip. The linear transmission $\hat{T}$ is accomplished by multiplying the components of the field in the frequency domain, $\tilde{A}$, by a transmission matrix T (whose elements are $T_{qp\omega}$), for each transverse mode q and each frequency ω:

$$\tilde{A}_{q\omega} \to \sum_{p=1}^{M} \tilde{A}_{p\omega} T_{qp\omega} \qquad (4)$$

For the calculation of the component attractors relevant to the few-mode simulations, $T_{qp\omega}$ is diagonal (as disorder is neglected), so $T_{qp\omega} = T_{q\omega}\delta_{qp}$ is derived as the solution to the linear fibre propagation over the fibre length from $z = 0$ to $z = L$:

$$\tilde{A}_{q\omega} \to \tilde{A}_{q\omega} \exp\left[i\delta\beta_0^{(q)}L - i\Omega\delta\beta_1^{(q)}L - \sum_{m=2}^{3} i^{m+1}\frac{\beta_m^{(q)}}{m!}\Omega^m L\right] \qquad (5)$$

where $\Omega = \omega - \omega_o$. In general, $\hat{T}$ may also account for linear coupling between the guided modes, such as is caused by disorder, bending or clamping a multimode fibre, by imperfections in the imaging from the fibre end facet to the end facet or possibly by coupling deliberately introduced, for example, by waveplates, 2D phase plates, dispersive delay lines or other components.

We typically obtained the solutions to equation (2) numerically, initializing the field as a noise input and performing a large number of iterations until a steady-state was attained.

As mentioned in the main text, such models can be naturally extended into a hierarchy of more accurate, but more complex, models of the laser. For example, to model the experiments with all the fibre's guided modes in Fig. 3, we applied this by neglecting details of non-linear pulse propagation. This reduced model corresponds to an equation of the form:

$$A(x, y, t) = \lim_{n \to \infty} [\hat{R}\hat{T}_r\hat{F}\hat{S}\hat{A}\hat{G}]^n E_0(x, y, t) \qquad (6)$$

where $\hat{T}_r$ is the linear transmission through the cavity as above, but includes disorder. $\hat{R}$ is obviously artificial and is not strictly required in reduced laser models; it is applied here practically to minimize the need to systematically vary the gain saturation energy when other parameters are modified. Here, we use a simplified model of the gain operator by employing an approximate solution to the field propagation (see Supplementary Section 4 for details). Disorder is implemented in $\hat{T}_r$ by multiplying the vector of modal coefficients at each frequency before and after applying the spatiotemporal dispersion by a matrix $M_D = I + rX$, where I is the M by M identity matrix, r is a parameter that characterizes the strength of disorder (between 0 and 1) and X is a random complex matrix (drawn from a uniform distribution from –1 to 1 for real and imaginary components). As experimental disorder manifests both in conservative (coupling between guided modes) and dissipative coupling (that is, disordered mode-dependent loss through disorder-induced coupling to radiating modes), X is not symmetric. We implicitly assumed that disorder only couples transverse modes, and take the disorder to be fixed across the timescale of the simulation. To account for the disordered coupling occurring mainly at the fibre clamps (and to approximate the distributed disorder within the fibre), we applied two different fixed matrices, $M_D^{(1)}$ and $M_D^{(2)}$, before and after applying the spatiotemporal dispersion each round trip ($\hat{T}$ as in equation (5)). For the simulations shown in Fig. 3, we additionally took the output field to be the rejected light from the SA, as in the experiment.

For the reduced model simulations in Fig. 3, for the narrow SF simulation we used a 6 μm full-width-at-half-maximum Gaussian SF, offset from the centre by 1 μm, a 40 μm gain dopant width parameter (see Supplementary Section 4 for the definition), $r = 0.05$ and a SA saturation intensity of 40 GW cm^{-2} and gain saturation fluence of 3 mJ cm^{-2}. For the intermediate filter simulations, these parameters were respectively 50 μm, 2.5 μm, 35 μm, 0.15, 30 GW cm^{-2} and 3 mJ cm^{-2}. Finally, for the broad filter simulation, the parameters were a 1,000 μm SF full-width-at-half-maximum, offset by 270 μm, with a gain dopant width parameter of 35 μm, $r = 0.2$, SA saturation intensity of 25 GW cm^{-2} and gain saturation fluence of 10 mJ cm^{-2}.

STML oscillator experiments. The STML oscillator is based on the same partially graded multimode gain fibre used by Wright et al. (ref. [25]). The laser design is shown in Fig. 1b; more detailed schematics are given in Supplementary Fig. 18, and a brief tour of the oscillator and measurements is shown in Supplementary Video 1. A 10 nm spectral filter was accomplished through a birefringent filter, a quartz plate between the output polarizing beam splitter (which also served as the rejection port for non-linear polarization rotation) and the isolator. An adjustable SF was realized with an adjustable iris on a three-axis translation stage. A 60 W, 976 nm pump diode was coupled into the outer cladding of the fibre and combined and/or separated with the laser signal using dichroic mirrors. The results presented here are based on a cavity design in which the SF was placed in the near-field image of the output fibre end facet, which was subsequently imaged back onto the input fibre end facet with a 1:1 total magnification. The SF was usually slightly defocused from the near-field image of the output fibre end facet, to achieve a slightly smoother SF with our hard aperture. However, details like this, or even the imaging system itself, are mentioned for completeness only—we obtained qualitatively similar results in a variety of other designs, which included ones with different spectral filter sizes, SF implementations, without imaging optics in the free-space section and with SFs in the far-field plane or intermediate locations.

We either angle cleaved the fibre at each end or employed angled end caps to prevent back reflections that disrupt the ML. Although end caps increase the laser's peak power handling and dramatically reduce the effects of clamp-induced linear mode coupling and disordered mode-dependent loss, we found evidence for beam distortions due to the relatively small aperture of our end caps (400 μm). To avoid the error this would introduce in mode decomposition, the results presented here use simply angle-cleaved fibres. We expect that the use of larger end caps in the future will alleviate several practical challenges with the current design.

The laser was aligned to maximize the continuous-wave output power and the SF was aligned to the centre of the beam by closing it incrementally and maximizing the output power by adjusting the transverse position. The SF was then opened and the pump power increased to a level at which ML was expected (in our experiments, this was >20 W). To mode-lock the laser, we employed the same techniques used to mode-lock single-mode fibre lasers based on non-linear polarization SAs, but also adjusted the SF. Generally, it is easy to achieve ML, but in some realizations finding a stable mode-locked state (that is, one that remains mode-locked for many hours, and through environmental perturbations over the course of a day) can be challenging. Empirically, we found that the use of a moderate or narrow SF was helpful; this may result from forcing the laser to mostly occupy low-modal-dispersion, low-order modes, and also by reducing the number of competing solutions. A high-quality, self-starting mode-locked state, stable enough to make measurements over several days, can be typically obtained by finding any mode-locked state, then slightly adjusting the pump power and SF size and/or position to ensure the single-pulse operation and/or stability of the pulse train.

For each mode-locked state, to verify single-pulsing operation without a continuous-wave background, we made the following measurements. First, the laser's autocorrelation was measured using a long-range intensity autocorrelator, and the pulse train from the laser was measured with a fast photodiode and

oscilloscope. Together, these measurements ensured that only a single pulse existed in the cavity. Note that, due to the spatiotemporal structure of the pulse, the intensity autocorrelation is not an accurate representation of the pulse but is, nonetheless, useful as a rough measure of the pulses' existence and temporal extent. Second, after spatially uniform sampling by glass windows, we coupled the entire field into a long highly multimode step-index fibre and then measured the spectrum using an optical spectrum analyser. This was done to average the spectrum across the entire field in space, to ensure that no continuous-wave background was present (given the numerous lasing modes, it is possible for the continuous-wave background to exist in modes uninvolved with the mode-locked state, and hence measurements that do not average across the entire field may yield misleading results). We then measured the radio frequency spectrum using a radio frequency spectrum analyser with a 10 Hz resolution bandwidth and a 1 Hz video bandwidth. Finally, 3D measurements and mode decomposition were measured as described in the next section. Before and after the measurements, we blocked the cavity and unblocked it to allow the laser to start from noise to determine whether the state is self-starting. We found that stable states were almost always self-starting (those reported here are), and are robust against small perturbations of the fibre, the cavity alignment and so on. Some states (especially in the SAGE regime) were bi- or multistable ones, and hence each time the laser was started it reached one of several different steady states. Supplementary Figs. 16 and 17 show the complete set of measurements for two representative STML states (both states were additionally verified to be self-starting).

In experiments, the rejected light from the NPE SA was used as the output coupled light (Supplementary Fig. 18). This means that the output field differed slightly from that which remained circulating in the cavity, unlike the case for the simulations. To account for this, we performed some experiments with a second additional output, released after the polarizing beam-splitter that realized the loss for the NPE SA. With this measure, no significant differences were observed, and therefore, for practical purposes, we maintained the simpler design of using the rejected light as the output coupled light.

3D pulse-field measurement and mode decomposition experiments. Our approach to measuring the 3D electric field (3D pulse) emitted by the laser was inspired by work on STRIPED FISH[49,50] and TERMITES[51], as well as by a large number of works on algorithm-based mode decomposition[52–54]. Many other works are published on both 3D field measurements and mode decompositions, each with their own advantages and disadvantages. The primary objective of our approach was to build a simple device that would have an easily adjustable space–time bandwidth product to allow the measurement of the range of possibly complex chirped pulses emitted directly by the laser. This was accomplished by the use of delay-scanned off-axis digital holography[55,56] (for a schematic of the device, see Supplementary Fig. 18) in which a spatially filtered version of the main pulse was used as a spatiotemporal reference. We measured the electric field $E(x,y,\tau)$ (where τ refers to the local time in the pulse reference frame) and performed a Fourier transform to obtain $\tilde{E}(x,y,\omega)$. The phase relationship between the frequency components was obtained by measuring the frequency-resolved optical gating trace of the spatially filtered reference field. To perform the mode decomposition, we then took the overlap of the field at each frequency with the calculated spatial modes, $\varphi_n(x,y)$, for the nominal refractive index profile of the multimode gain fibre, that is $c_n(\omega) = \iint \varphi_n^*(x,y)\tilde{E}(x,y,\omega)\mathrm{d}x\mathrm{d}y$. The plotted values in Fig. 3, $\frac{U_i}{U_{tot}}$, were then:

$$\frac{U_i}{U_{tot}} = \frac{\int |c_i(\omega)|^2 \, \mathrm{d}\omega}{\sum_{n=1}^{M} \int |c_n(\omega)|^2 \, \mathrm{d}\omega} \tag{7}$$

The mode decomposition has a notably high uncertainty following, first, from the inexact knowledge of the fibre index profile and, second, from the inability to locate the centre of the fibre in the image with a high precision. For a fibre with 90 modes, the representational capacity of the guided modes is high, and so the decomposition error associated with the latter effect can be considerable, even if one takes a rigorous approach to image calibration; note that a high-quality mode decomposition can be obtained even when these calibration parameters are wrong. (In contrast, it is much easier to unambiguously obtain the correct calibration for a few-mode fibre, with its very limited representational capacity.) Our calibration approach involved utilizing four independent, distinct calibration methods of the scale and centre position of the fibre and then averaging across these. To translate the uncertainty in the calibration of the scale and fibre centre position into the mode decomposition error, we compared the mode decomposition obtained with the nominal optimum calibration versus one in which the scale and fibre centre position were adjusted by the error of these calibration parameters (inferred by the deviation across the independent calibration methods). In the Supplementary Information, we analyse this error for two specific STML states for which all the measurements are presented. We found a typical error of about 6% (8% and 4% for the two examples considered). However, we note that this error estimate

is an average across the modes, and specific modes have a larger uncertainty (Supplementary Fig. 19); furthermore it does not account for the error associated with the real fibre's refractive index profile differing from our best estimate (which introduces additional error due to the difference between the decomposition modes and the true modes of the fibre). Altogether, these issues motivate our interpretation of the mode decomposition as a rough estimate: it yields a good measure of the overall distribution of energy (in low- versus high-order modes, say), but not a precise measurement of the exact modal distribution.

Note added in proof: Prior to publication of this work, we noticed a preprint that provides evidence for pulses resembling single-mode SAGE pulses[57].

Data availability
All data in the manuscript and supplementary are available from the corresponding author on reasonable request.

Code availability
The principal components of the codes used in the manuscript have been made publicly available with extensive documentation at https://github.com/WiseLabAEP/GMMNLSE-Solver-FINAL Additional codes mostly build on this core code, and may be requested from the corresponding author.

References
49. Guang, Z., Rhodes, M., Davis, M. & Trebino, R. Complete characterization of a spatiotemporally complex pulse by an improved single-frame pulse-measurement technique. *J. Opt. Soc. Am. B* **31**, 2736–2743 (2014).
50. Guang, Z., Rhodes, M. & Trebino, R. Measuring spatiotemporal ultrafast field structures of pulses from multimode optical fibers. *Appl. Opt.* **56**, 3319–3324 (2017).
51. Pariente, G., Gallet, V., Borot, A., Gobert, O. & Quéré, F. Space–time characterization of ultra-intense femtosecond laser beams. *Nat. Photon.* **10**, 547–553 (2016).
52. Shapira, O., Abouraddy, A. F., Joannopoulos, J. D. & Fink, Y. Complete modal decomposition for optical waveguides. *Phys. Rev. Lett.* **94**, 143902 (2005).
53. Lü, H., Zhou, P., Wang, X. & Jiang, Z. Fast and accurate modal decomposition of multimode fiber based on stochastic parallel gradient descent algorithm. *Appl. Opt.* **52**, 2905–2908 (2013).
54. Paurisse, M., Lévêque, L., Hanna, M., Druon, F. & Georges, P. Complete measurement of fiber modal content by wavefront analysis. *Opt. Express* **20**, 4074–4084 (2012).
55. Verrier, N. & Atlan, M. Off-axis digital hologram reconstruction: some practical considerations. *Appl. Opt.* **50**, H136–H146 (2011).
56. Kim, M. K. Principles and techniques of digital holographic microscopy. *J. Photon. Energy* **1**, 018005 (2010).
57. Ding, Y. et al. Spatiotemporal mode-locking in lasers with large modal dispersion. Preprint at http://arXiv.org/abs/1912.00161 (2019).

Author contributions
L.G.W. performed experiments and simulations and developed the theoretical models. P.S. constructed, with help from A.I. and L.G.W., the 3D pulse measurement and mode decomposition device and software. H.P. performed additional experiments. Z.M.Z. contributed to the numerical codes. B.A.M., C.R.M. and D.N.C. contributed to the formulation and analysis of the theoretical models. L.G.W. and F.W.W. wrote the paper, which was edited by all the authors. F.W.W. supervised the project.

Competing interests
L.G.W. and F.W.W. hold a US patent, number US20190207361A1 for STML.

Additional information
Extended data is available for this paper at https://doi.org/10.1038/s41567-020-0784-1.

Supplementary information is available for this paper at https://doi.org/10.1038/s41567-020-0784-1.

Correspondence and requests for materials should be addressed to L.G.W.

Peer review information *Nature Physics* thanks Philippe Grelu, Fatih Ilday and Heping Zheng for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.
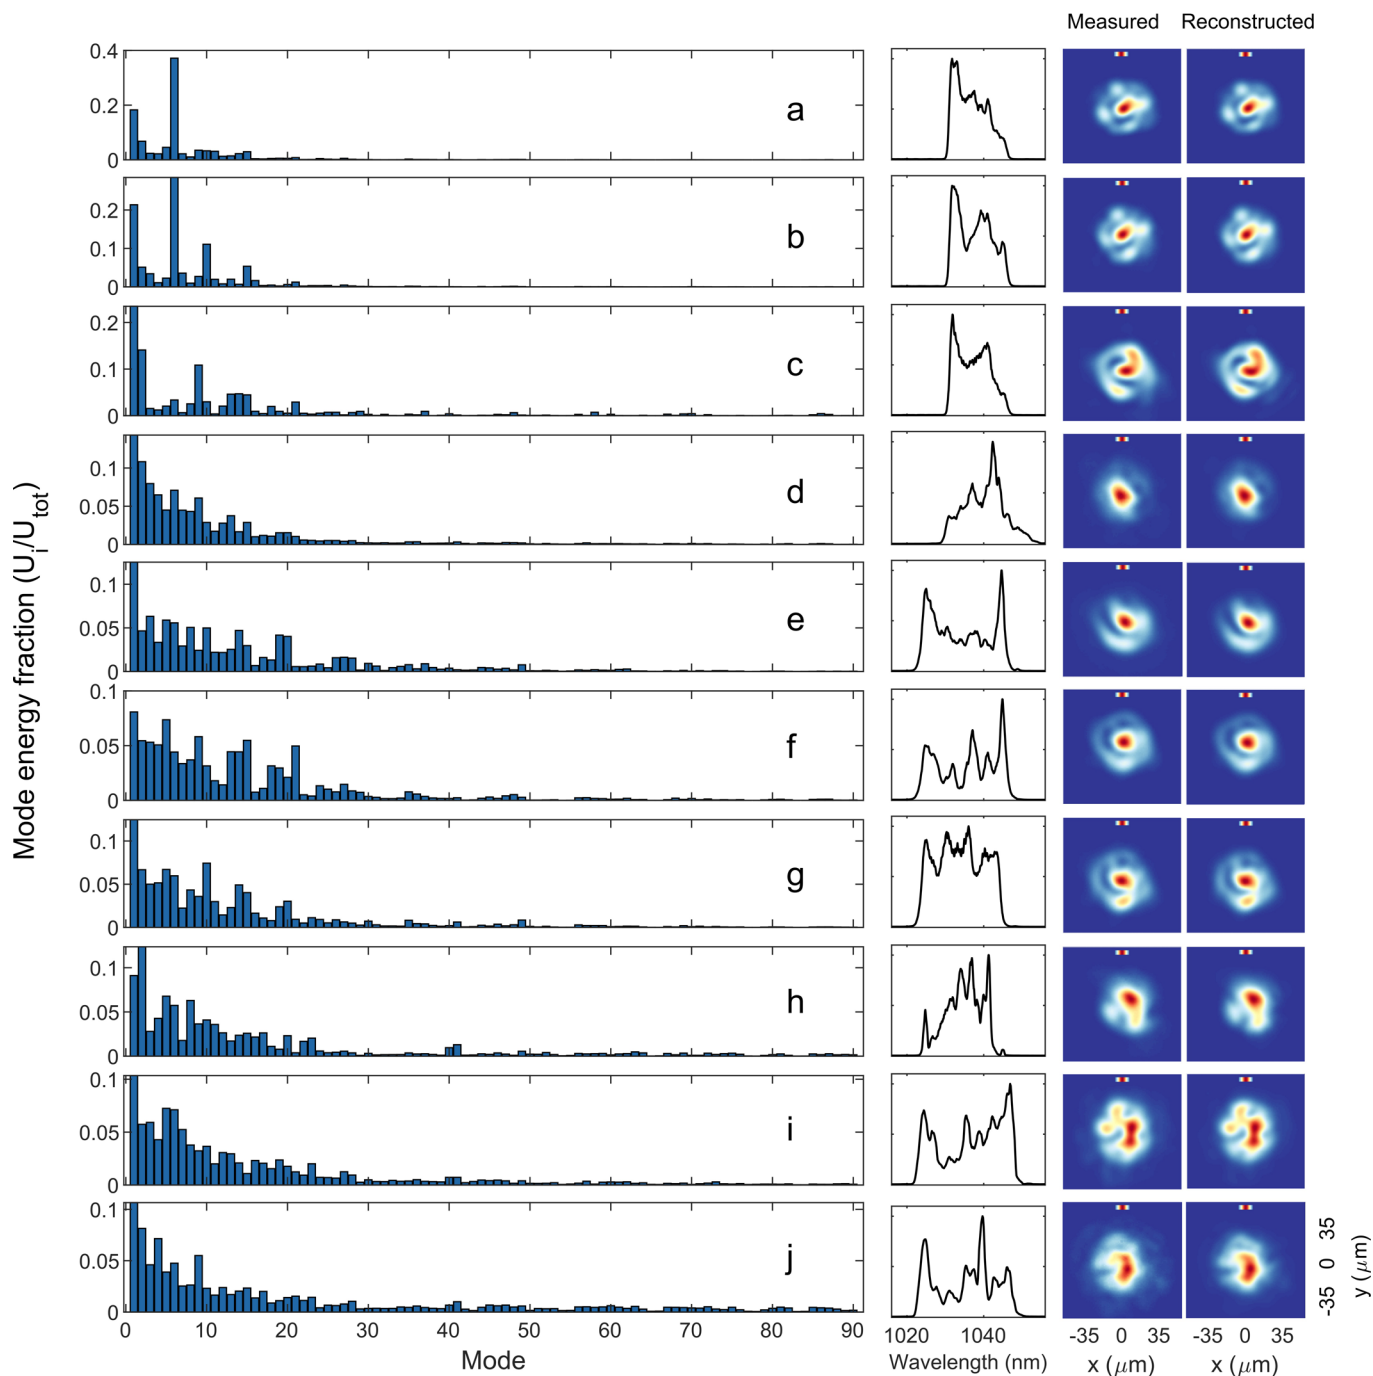
Extended Data Fig. 1 | Examples of experimental STML states ranging from the very narrow filter regime (a–c) through narrow-intermediate filter sizes (d–h) and large/no spatial filter (i-j). The panels (left to right) show the modal energy distribution, the whole-field spectrum, and the measured and reconstructed beam profiles. The reconstructed beam profiles obtained from the 3D field modal decomposition are shown to help qualify the mode decomposition. Overall, we see similar trends as in the few-mode simulations. In the narrow-filter regime, STML states concentrate energy into radially-symmetric modes and a small number of low-order modes. The spectra in this regime also resemble the narrow, asymmetric spectra observed in few-mode simulations. For intermediate filter sizes, we see STML states with intermediate features, still tending to concentrate energy into radially-symmetric modes, but with a broader distribution of modes and with broader and less asymmetric, more rectangular spectra. Finally, for the largest spatial filter (here, meaning no spatial filtering besides that occurring in the optical isolator and fibre coupling) we see the experimental manifestation of the SAGE regime: broad symmetric spectra with a broad distribution of modes that is most heavily weighted to low-order modes. The absence of radial symmetric features in the beam profiles also evidences the transition from radially-symmetric low-order modes to the degenerate mode families of less-symmetric higher-order modes that occurs in the SAGE regime. Considering that 90 transverse modes are present in the oscillator, and that disordered linear mode coupling is important especially for low-symmetry modes, the agreement with the few-mode trends is good. The experimental results also agree well with the predictions of the reduced models in Supplementary Material, Section 4.