Check for updates

PERSPECTIVES

Filamentary plasma eruptions and their control on the route to fusion energy

Christopher Ham, Andrew Kirk, Stanislas Pamela and Howard Wilson

Abstract | The tokamak is the most advanced approach to fusion and is approaching operation under power-plant conditions, promising sustainable, low-emission, baseload power to the grid. As the heating power of a tokamak is increased above a threshold, the plasma suddenly bifurcates to a state of high confinement, creating a region of plasma with a large pressure gradient at its edge. This bifurcation results in a repetitive sequence of explosive filamentary plasma eruptions called edge-localized modes (ELMs). ELMs on next-step tokamaks, such as ITER, will likely cause excessive erosion to plasma-facing components and must be controlled. We present what is understood about how ELMs form, their filamentary nature and the mechanisms that transport heat and particles to the first wall of the tokamak. We also discuss methods to control ELMs, including magnetic perturbations.

Fusion is a promising source of sustainable, low-emission, baseload power to meet the energy needs of future generations. The fuel required, deriving ultimately from deuterium, contained in seawater, and lithium, used to breed tritium, is plentiful and available. The radioactive waste generated would be relatively small if the process used low-activation materials that have already been identified. Further optimization of the materials used in the construction of the fusion power plant could reduce waste even more, completely eliminating long-lived isotopes. Fusion power plants would produce no greenhouse gases and, thus, have the potential to contribute a significant fraction of the future energy mix without a detrimental impact on climate change. The realization of fusion energy is one of the grand challenges facing humanity, with enormous societal benefits on offer, and the fusion-research community is approaching a solution to the remaining challenges.

The most advanced fusion reactor design is the tokamak, which is used by projects such as ITER. The tokamak is not only a key facility on the pathway to fusion energy but it also offers a rich variety of plasma physics, exhibiting a range of fundamental phenomena, including magnetic reconnection, turbulence and features of complex systems, such as bifurcations and self-organization. In this Perspective, we focus on one particularly interesting and important aspect of tokamak plasmas — repetitive, violent, filamentary eruptions called edge-localized modes (ELMs)[1].

A tokamak uses a magnetic field to confine the energy and charged particles of a plasma in a toroidal chamber — a configuration known as a magnetic trap. The magnetic field has two orthogonal components: one that is created by a system of current-carrying, often D-shaped, coils around the plasma and one that is created by a current that is induced in the plasma. The resulting magnetic field lines spiral around a set of nested toroidal flux surfaces, providing an effective plasma-confinement system that can be heated to conditions approaching those required for fusion energy (FIG. 1). Experiments were carried out at the Joint European Torus (JET) tokamak in 1997 using deuterium and tritium, producing 16 MW of fusion energy while requiring 25 MW of plasma heating[2].

At low heating powers, the edge of the plasma is characterized by fine-scale plasma turbulence; this leads to a loss of heat and particles, and, hence, limits the core plasma pressure that is achievable with a given heating power[3]. This is the so-called low-confinement mode, or L-mode, of operation. If the plasma heating power is increased above a threshold, the plasma spontaneously self-organizes on a timescale of milliseconds into an improved, high-confinement (H-mode) regime (see BOX 1, REF.[4] and references therein). In the H-mode regime, the edge-plasma turbulence is dramatically suppressed. The mechanism is believed to be associated with a strong plasma flow shear at the plasma edge, which breaks up the turbulent eddies, reducing the particle and energy transport to generate what is called an edge transport barrier[4]. This barrier, typically a few centimetres wide in a plasma that is 1–2 m across, leads to a steep pressure gradient at the plasma edge. This region acts as an insulating envelope around the plasma, raising the pressure of the entire core. The increase in core pressure leads to an increase in overall confinement of a factor of approximately 2; thus, even though the edge transport barrier comprises only the outer few per cent of the plasma, it is vital for fusion performance and key for ITER to achieve its fusion performance targets.

However, despite the benefits of edge transport barriers, if their pressure gradient exceeds a certain limit, it is sufficient to drive explosive plasma instabilities, known as ELMs[1,5,6]. ELMs are sequences of fast, repetitive filamentary plasma eruptions. They are a concern for next-step tokamaks, such as ITER, because they are predicted to cause excessive erosion of the surrounding material surfaces[7]; they are also interesting from a fundamental plasma physics point of view, as discussed here.

In this Perspective, we describe the tokamak and the self-organized H-mode with its edge transport barrier that improves confinement but also brings about ELMs. We discuss the theory, modelling and experimental evidence for ELMs. We describe aspects such as their filamentary, explosive behaviour and the mechanisms by which energy is lost from the confined plasma during an ELM. We also show some of the strategies for controlling ELMs that will help achieve viable fusion power plants. This article is intended to give physicists from all disciplines an introduction and appreciation of the work and progress on ELMs.

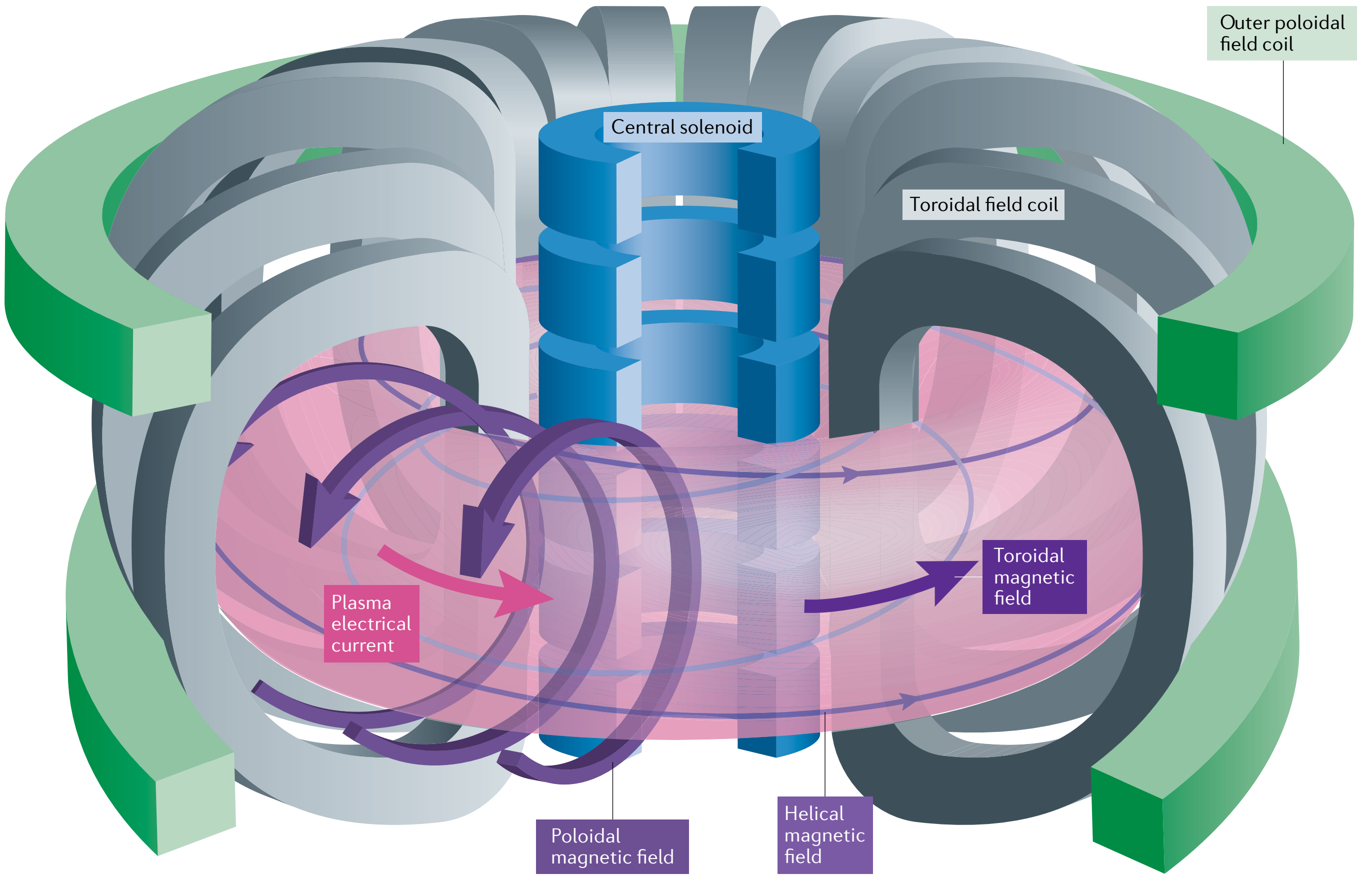

Fig. 1 | **Currents and fields in a tokamak.** A typical tokamak configuration uses current-carrying coils to generate part of the required magnetic field. A current is induced in the plasma to generate further magnetic fields, which result in helical magnetic field lines. To leading order, the charged particles of the plasma follow these field lines, creating an effective plasma-confinement system. Image courtesy of EUROfusion.

ELMs

Physics of filament formation. With features reminiscent of solar eruptions, ELMs are explosive, filamentary events, aligned along the magnetic field[8]. ELMs eject large amounts of energy and particles from the confined region — typically 5–15% of the total energy stored in the plasma — in a short amount of time (100–300 μs), resulting in large heat fluxes to the vessel structures surrounding the plasma. How these heat fluxes scale to future, larger tokamaks, including ITER, is a cause for concern[7]. Thus, as well as the strong motivation driven by inherent scientific curiosity, there is a high-priority need to understand ELMs, to quantify their impact in ITER and future tokamaks, and to develop control techniques; indeed, there is concern that ELMs will need to be completely avoided in tokamaks beyond ITER (such as demonstration power plants). However, there are some positive aspects to ELMs, such as their role in controlling plasma density and impurities, and the removal of the helium ash from the fusion reaction.

Over the past decade, much-improved fast-camera technology has made it possible to directly observe these filaments and their characteristics, building on the first observations that were made on the Mega Ampere Spherical Tokamak (MAST)[9]. A detailed picture of the physical properties of ELMs has been provided by a variety of diagnostic instruments employed on a number of the world's tokamaks[10]. For example, 2D electron cyclotron emission imaging has been used on the Korea Superconducting Tokamak Advanced Research (KSTAR) device to visualize filament dynamics[11].

Images from MAST at two different stages of the eruption show that, during an ELM event, narrow plasma filaments push out from the edge transport barrier region of the confined plasma into the scrape-off layer (SOL) (FIG. 2). The SOL is a region of open magnetic field lines that enclose the hot core and that guide the plasma to the armoured target plates of the divertor at the top and bottom of the tokamak, where the plasma deposits much of its exhausted energy. The filaments subsequently separate from the edge of the plasma and travel out radially towards the vacuum vessel wall, carrying with them particles and energy. The filaments exist for the time over which particles are being released into the SOL. In the early stages of their evolution, they rotate toroidally with the edge plasma, before decelerating toroidally as they accelerate radially outwards towards the vacuum vessel wall. As the filaments move outwards, they twist to remain aligned with the local magnetic field lines.

The basic ELM mechanism can be understood from the ideal magnetohydrodynamics (MHD) model[12], which approximates the plasma as a fluid with effectively infinite electrical conductivity. There are two basic ideal MHD instabilities associated with ELMs: the ballooning mode[13] and the kink (or peeling) mode (see BOX 2, REF.[14] and references therein). The ballooning mode is driven by the steep pressure gradient of the edge plasma. The kink or peeling mode results from the strong so-called bootstrap current density[15]. The bootstrap current is generated because the plasma is confined in a toroidal magnetic field. The magnetic field is weaker on the outboard side of the machine than on the inboard side; thus, particles with little velocity component along the magnetic field lines become magnetically trapped on the outboard side. The density and temperature gradients in the plasma mean that there are more particles towards the core and this imbalance generates a current parallel to the magnetic field. Collisions generate friction between the passing, or untrapped, particles and the trapped particles, and it is this friction that produces that bootstrap current[16]. As a result, both ballooning and peeling modes are destabilized and they couple to drive peeling–ballooning modes. Large ELMs are believed to be triggered by these peeling–ballooning modes[17–19].

The following simple picture represents an ELM cycle. At the beginning of the ELM cycle, just after the previous ELM, both the pressure gradient and current density are low and the plasma is stable. The plasma is still being heated and, so, the pressure gradient and edge current amplitudes steadily increase as the edge transport barrier re-establishes itself, typically broadening at the same time. This proceeds until the stability limit of the peeling–ballooning mode is reached, at which point the ELM is triggered, causing a crash in the edge pressure and current, and the cycle to start again.

Although this simple 'cartoon' picture of the ELM cycle is helpful, the plasma dynamics between ELMs is more complicated. It is often found experimentally that, following an ELM, the edge gradients first recover only in a narrow region in the immediate vicinity of the plasma edge, and this region steadily broadens as the full edge gradient region builds[20–22]. The pressure gradient is initially held below the peeling–ballooning instability threshold by transport driven by residual turbulence in the edge transport barrier region. However, as the transport barrier width broadens, the pressure gradient required for the peeling–ballooning instability reduces until an ELM is triggered[23,24]. A further issue is that of second stability. For sufficient current density, the pressure gradient to drive a ballooning mode is substantially increased, allowing a steeper edge pressure gradient and influencing the evolution of the edge transport barrier width between ELMs, for instance, as observed on JET[25,26]. Furthermore, theory and initial experimental evidence from the DIII-D tokamak show that an improved H-mode can be found, called a Super H-mode. This mode requires the plasma to be strongly shaped and produces a larger pedestal height and width[27].

The simple model that we describe here applies to the most violent type of ELM, so-called type I ELMs[28], but there exist other types, with smaller eruptions and higher frequencies. Their physics is different to type I ELMs and they are much less well understood, but their smaller size would be desirable for ITER and future devices[29]. Because these smaller ELM types do not erupt so explosively from the plasma, and there is no universally accepted model, we consider them to be beyond the scope of this Perspective and only address the most common type I ELMs.

Analytical theory of the early nonlinear evolution of the ballooning mode provides an explanation of why the ELM is such an explosive event[30]. For a ballooning mode to grow, it must bend the magnetic field lines that thread through the plasma; doing so requires a lot of energy. Instability arises when the free energy associated with the pressure gradient exceeds this field-line-bending energy. The nonlinear theory predicts that, as the ballooning mode grows, it modifies the magnetic field structure in such a way that the stabilizing effect of field-line bending is reduced, thus enhancing the net drive and accelerating the ballooning mode ever harder, even at fixed pressure gradient. This positive feedback mechanism drives explosive growth, consistent with the violent eruptions that are observed experimentally. The stabilizing field-line bending is minimized through the formation of filamentary structures, which are elongated in the direction along magnetic field lines, but increasingly narrow in the perpendicular direction (in the magnetic flux surface) as they erupt. Thus, the theoretical picture is one of narrow filaments of plasma erupting violently from the edge transport barrier on a typical timescale of the order of 50–100 μs.

This theory for the early nonlinear phase of the ELM has been further developed to a fully nonlinear ideal MHD model. The perfect plasma conductivity of this model means that the magnetic-field topology

Box 1 | **Equilibrium pressure and current profiles, and tokamak regions**

The plasma in a tokamak can be divided into two regions, defined by the topology of the magnetic field in them (see the left panel of the figure). The magnetic geometry is produced by placing a current-carrying coil underneath the divertor plates. The core, or confined, region has magnetic field lines that lie within a series of nested toroidal flux surfaces. These nested flux surfaces can sustain a pressure gradient. Outside the confined region, there is a region of open field lines that connect to the exhaust or divertor plates. This is called the scrape-off layer. The confined region and the scrape-off layer are separated by a flux surface known as the separatrix, at which the change in topology from closed to open field lines also introduces a magnetic X-point.

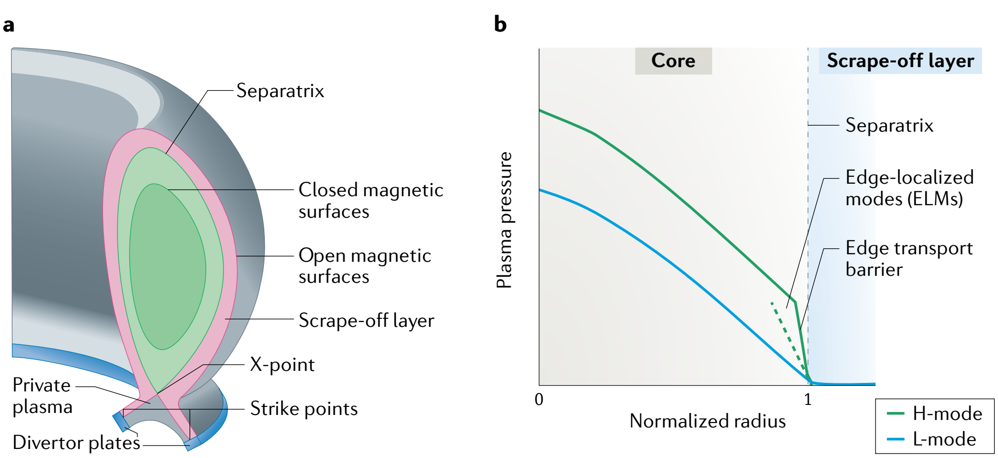

The equilibrium configuration of a tokamak plasma can be described in terms of the shape of flux surfaces and the pressure and current density as a function of the flux. The pressure profile is determined by the transport processes in the plasma. The transport across the core region is 'stiff', in the sense that it is difficult to change the pressure gradient in this region (see the right panel of the figure, which shows pressure profiles of modes of operation as a function of the minor radius across flux surfaces from the centre of the plasma to the edge). In the low-confinement mode (L-mode), the pressure decreases smoothly towards zero at the edge of the confined region. In the high-confinement mode (H-mode), there is a suppression of edge turbulence, leading to an edge transport barrier and strong edge pressure gradient region. The edge pressure gradient region raises the core pressure, thus improving the overall confinement.

Figure part **a** courtesy of EUROfusion.

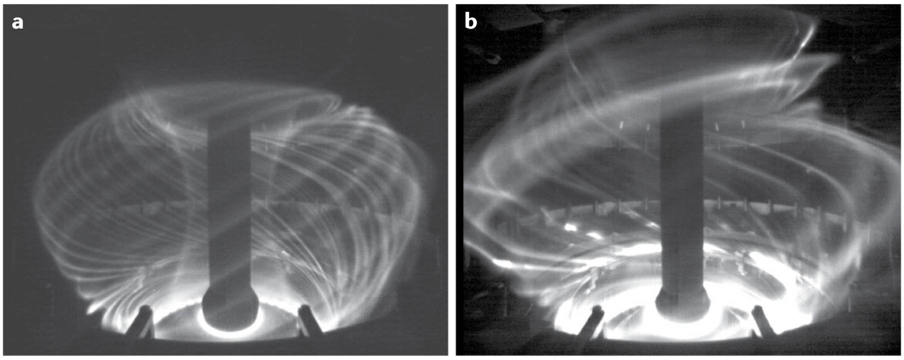

Fig. 2 | **Visible images of an edge-localized mode captured on the Mega Ampere Spherical Tokamak. a** | The start of an edge-localized mode. **b** | The eruption of the filamentary structures. Both images were taken using a 5-μs exposure time. Reprinted with permission from REF.[87], Taylor & Francis.

must be conserved — the field lines cannot break and reconnect. The displacement of the resulting filaments is, thus, predicted to saturate on a relatively fast timescale. Non-ideal processes, possibly involving reconnection, thus become important in determining how and where the heat and particles in the filament get to the first wall and divertor of the tokamak vessel[31,32].

A wide range of diagnostics are used to understand the state of the plasma[33] and combining the results from different instruments reveals further properties[34]. For example, the filaments are composed of hot, dense plasma, with density and temperature comparable to those of the edge transport barrier plasma that the ELM eruption originated from. In addition, the filaments carry a significant amount of current from the edge plasma, which is consistent with the strong magnetic signature of ELMs, as measured by magnetic sensors around the device[35].

The ELM energy-loss mechanism. We have, thus far, provided a theoretical picture of ELMs in terms of an ideal MHD model. A consequence of this model is that there can be no net transport: the plasma and magnetic field remain tied together within the filament as it erupts. However, in reality, there is transport as a result of the ELM, and quantifying it requires going beyond the ideal MHD description. Direct evidence that non-ideal effects become important includes the observation that plasma filaments eventually disconnect from the confined region and travel all the way out to the first wall; this cannot happen if the plasma in the filament remains ‘frozen’ to the magnetic field. Understanding particle and energy losses is important for predicting the potentially damaging consequences for future fusion devices (including ITER) and developing control and avoidance strategies.

Although it is clear that energy is transferred from the filament to the open field lines in the SOL, the precise physical mechanism by which this happens remains unclear. Nevertheless, various ideas have been proposed, some of which require magnetic reconnection[14]. One possibility is that the hot filaments could simply break off from the confined plasma (either through reconnection or by drifting across magnetic field lines) and decay in the SOL. Measurements of thermal energy stored in the filaments (obtained, for instance, by using Thomson scattering on MAST) show that, at any one time, the filaments collectively only carry approximately one-fifth of the total thermal energy lost in an individual ELM event[34,36]. Studies also suggest that most of the energy is lost in the direction parallel to the magnetic field[37]. These observations suggest that transport mechanisms are at play, in addition to the observed breaking off of filaments. The filaments must either directly act as a conduit for the hot plasma to travel from the confined region to the SOL or there is a degradation of the transport barrier leading to an enhanced flux during the ELM. In other words, mechanisms in addition to the direct transport from the filaments are required, as we now discuss.

In the ‘leaky-hosepipe’ model, the filament is assumed to remain connected to the confined plasma on the inboard (high-field) side. However, because of the ballooning nature of the instability, the filament pushes out into the SOL on the outboard side. It, therefore, provides a conduit to rapidly transport heat from the confined plasma into the SOL, where it leaks into the exhaust region. This diffusion is enhanced by the increase in the difference in pressure between the plasma inside and outside the filament as it erupts[34,38]. A similar model to the leaky hosepipe is the ‘squirting hosepipe’, in which the magnetic field lines within the filament reconnect, to join with those of the SOL, to create a continuous path for hot plasma to siphon from the confined plasma directly onto the divertor target plate. In a mechanism that we call the ‘ergodized-edge model’, the current-carrying filaments perturb the magnetic-field structure as they erupt from the edge transport barrier; this perturbation generates chains of magnetic islands by forced magnetic reconnection, which then overlap to produce an ergodized magnetic field in the edge region. This ergodic magnetic field fills a volume between two nested toroidal surfaces. Magnetic field lines wander through this volume so that there are no longer nested magnetic surfaces. This ergodized field has reduced confinement and, thus, plasma is lost to the SOL[39–41].

Among the experimental evidence for different mechanisms, there are indications that magnetic reconnection does occur during ELMs. Measurements of microwave bursts were obtained on MAST[42], suggesting that electrons are accelerated at the beginning of the ELM, presumably by electric fields that could be created by magnetic reconnection of the filaments. On the Axially Symmetric Divertor Experiment (ASDEX) Upgrade (AUG) tokamak, accelerated fast ions were also measured, together with soft X-ray and electron cyclotron emission bursts at the beginning of the ELMs[43,44], consistent with a reconnection event resulting in particle acceleration.

As the capability of high-performance computers advances, high-fidelity, nonlinear MHD simulations are becoming feasible, with the potential to reveal much more of the physics of ELM filaments. These simulations are extremely challenging, requiring multiple timescales and length scales to be resolved — from the microsecond eruption timescale to the tens of milliseconds or longer timescale for the evolution of the edge transport barrier. An even greater challenge is to resolve the extremely narrow layers that form in the vicinity of resonant surfaces (flux surfaces where the magnetic field lines close on themselves after a number of toroidal and poloidal turns) in low-resistivity plasmas, a result of which is that MHD simulations at realistic plasma resistivity are extremely challenging (and the resistivity influences the physics).

In the past decade, much progress has been made, particularly in terms of pushing resistivity closer to the experimental conditions[45,46]. Results from a number of codes have added weight to the international consensus that type I ELMs are indeed linked to peeling–ballooning modes and that the filamentation of the edge plasma plays a major role in the nonlinear dynamics of the instability[47,48]. FIGURE 3 shows nonlinear MHD simulations of an ELM filament using the JOREK MHD code[49] compared with camera images from an experiment. JOREK is a nonlinear extended MHD code that can simulate realistic tokamak geometries, including the main plasma, SOL and divertor region.

The array of experimental measurements mentioned above has been used to qualitatively validate these nonlinear simulations. Hypotheses that have been validated include the poloidal rotation of the filaments during ELM precursors, typically at rotational velocities of 1–10 km s^{-1}. Simulations have demonstrated that filaments rotate toroidally because of the momentum injected into the plasma by energetic neutral beam injection heating systems and poloidally owing to diamagnetic and neoclassical effects[48]. Another feature that is well reproduced by MHD modelling is the current carried by filaments, which is typically of the level of the bootstrap current at the plasma edge before the ELM onset[50]. In addition, the number of filaments that erupt in a given ELM (equivalent to the toroidal mode number n) has been shown to be described well by simulations[51,52].

There are some key measurements that can also provide more quantitative comparisons between the experiments and simulations and that, most importantly, address the fundamental characteristics of ELMs: their size (that is, how much energy they expel from the plasma), their velocity and their impact on wall surfaces; these are essential to understand for future reactors. On the JET tokamak, JOREK simulations reproduce quite well a range of experiments with ELM sizes between 20 and 250 kJ, and peak wall heat fluxes from 25 to 350 MW m^{-2} (REFS[46,53]). An important contribution of these simulations is that they add weight to the theoretical interpretation, discussed above, that most of the energy is lost along ergodic magnetic field lines that connect the plasma with the divertor plates (FIG. 4).

The encouraging agreement between theory, simulations and experiments[54] has shed light on the transport mechanisms of ELMs, but the picture is not yet complete and there remain substantial challenges for simulations. For example, the ergodic magnetic field that is predicted to be formed at the plasma edge, which provides enhanced transport during ELMs, can only occur via magnetic reconnection[55]. In an MHD model, this process is slow. Thus, even though simulations can reproduce the energy-loss mechanism and the filamentation relatively well, they typically exhibit a slightly lower rate of energy loss (approximately a factor of 2), which results in a longer duration of the ELM activity than that observed experimentally[46]. It is not yet clear which physics effect (or effects) that could accelerate the predicted ELM timescales is missing from current MHD simulation models, but we highlight two candidates.

The first is the nonlinear stability of the ELMs. This effect has been predicted theoretically and reproduced in simulations[45,53,56,57]. It has become clear that, as the plasma pressure progressively approaches the critical stability limit of the MHD peeling–ballooning modes, ELMs emerge out of underlying fluctuations at the edge of the plasma. This coupling between filaments and fluctuations of different wavelengths could be a fundamental effect underlying ELM onset and its violent nature.

The second candidate is related to kinetic effects that cannot be captured by MHD models alone, such as a role for electron inertia. The fluid approach of MHD relies on the assumption that the plasma is sufficiently collisional to maintain a Maxwellian velocity distribution to leading order, but this is not true for very hot plasmas at the edge of large tokamaks, with several consequences. The transport of energy along magnetic field lines in fluid models is described by a local collisional diffusivity, but kinetic effects, such as non-local transport due to fast electrons at the tail of their distribution, would lead to very different energy-transport mechanisms to the divertor[58]. This modification of transport has another consequence in turn, namely that, if hot electrons escape the plasma much faster than the ions, the result is that significant return currents are required to maintain the quasi-neutrality of the plasma; these additional currents could, perhaps, contribute to the degradation of the edge confinement.

In summary, progress in computational capabilities is enabling ever-higher-resolution MHD simulations, which are needed to capture nonlinear effects in highly conducting plasmas; as simulations advance, the agreement with experimental observations continues to improve. Nevertheless, fluid simulations have their limitations, and hybrid kinetic–MHD models are also emerging. These could finally deliver a truly predictive capability for filamentary eruptions — both for tokamaks and in astrophysics — and provide the basis from which to develop avoidance and control strategies.

ELM-control strategies

The heat and particles deposited by type I ELMs on the divertor and first-wall components of future devices, such as

Box 2 | **Ballooning modes and kink modes**

Two magnetohydrodynamics instabilities are key to edge-localized modes: the ballooning mode and the kink (or peeling) mode, both of which can be understood from the ideal magnetohydrodynamics model. The kink mode is a current-driven instability. It causes a straight cylindrical plasma column to become helical or kinked (figure panel **a**). The mode is not localized to the low-field side of the flux surfaces but also perturbs the inner, high-field side. The perturbation grows because, in the concave regions of the plasma column, magnetic pressure is increased, and in convex regions, the magnetic pressure is reduced, and this imbalance causes the deformation to grow further. The ballooning mode is driven by the pressure gradient in the plasma, combined with the curved magnetic field in the torus. On the inboard (high magnetic field) side of the tokamak, the curvature of the magnetic field lines stabilizes the pressure-driven mode, but not on the low-field side. If the free energy in the pressure gradient exceeds the energy required to bend magnetic field lines, the mode grows in this unstable region of the tokamak and produces the structure characteristic of the ballooning mode (figure panel **b**). The instability is aligned with the magnetic field and localized to the outboard side of the tokamak.

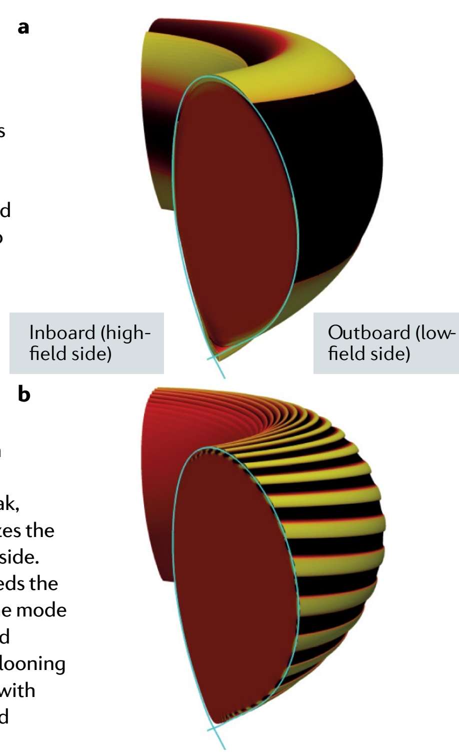

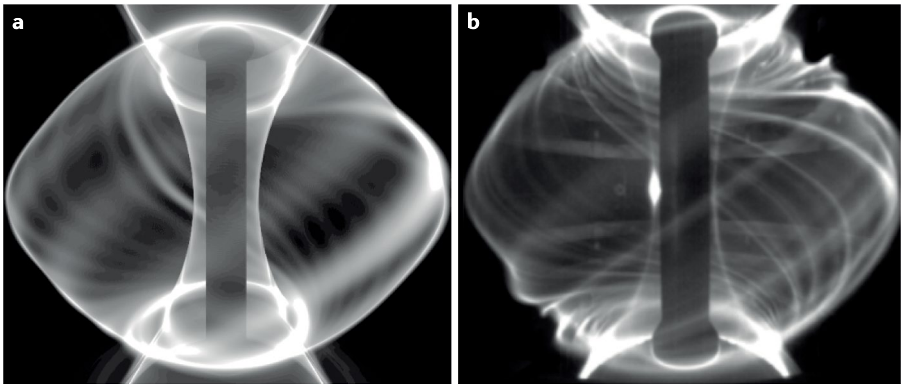

Fig. 3 | **Comparison of an edge-localized mode in simulation and experimental observation.** **a** | Light emission that is calculated using the simulation of an edge-localized mode and knowledge of emission from neutral particles — in other words, a synthetic diagnostic (for more details, see REF.[88]). **b** | A fast-visible-camera image of an edge-localized mode in an experiment at the Mega Ampere Spherical Tokamak.

ITER, will cause unacceptable damage. Therefore, strategies to avoid or control ELMs have become important. The most straightforward way to avoid ELMs is, of course, to run in L-mode. However, unless an alternative pathway to improve the confinement time in an L-mode plasma can be found, doing so would result in much lower core pressure than H-mode operation in a device of a given size and, hence, much lower fusion power. The question, then, is how an understanding of the physics of type I ELMs be used to mitigate or control them.

An appropriate place to start to answer this question is with the evolution of the operational point (that is, the edge current density and pressure gradient) towards the stability boundary for peeling–ballooning modes[59]; that is, the value at which an ELM is expected to be triggered. If the boundary is reached sooner — either because of an increase in the speed at which the edge transport barrier evolves or a decrease in the threshold for instability — then the ELMs would be more frequent (assuming that the ELM crashes to the same pressure gradient and current density). There is a robust experimental relationship between the ELM frequency (f_{ELM}) and the type I ELM size (ΔW_{ELM}): $f_{ELM}\Delta W_{ELM} = 0.2–0.4\,P$, where P is the input power to the plasma[60]. Therefore, increasing the ELM frequency can lead to less energy being released by each ELM; this is called ELM mitigation. The question then moves to how to increase the ELM frequency. Several techniques have been successful, such as pellet injection, in which pellets of deuterium ice are fired into the plasma one after another, in quick succession. Each one causes a rapid rise in the local pressure gradient[61–63], driving the edge pressure locally across the instability threshold to trigger an ELM; the ELM frequency is, thus, paced at the frequency of pellet injection. This strategy has been very successfully demonstrated on the DIII-D tokamak, at which a 12-fold increase in the ELM frequency has been documented with corresponding reduction in the ELM energy deposited at the divertor[62]. Furthermore, the core plasma was relatively unchanged by the pellets. ELM-pacing pellets are expected to be used on ITER, but for high-performance plasmas, several tens of thousands of pellets will be required per shot, which may present technical difficulties[64].

An alternative means of increasing the ELM frequency, which requires active control, is to apply fast vertical motions of the plasma (kicks) using the magnetic coils that vertically position the plasma. These kicks result in an increase of the edge current density, which then triggers ELMs at the frequency of the plasma vertical motions. This technique was first demonstrated on the Tokamak à configuration variable (TCV) tokamak[65]; it was later demonstrated on ASDEX Upgrade[66]. It has also been demonstrated on JET[67], the National Spherical Torus Experiment (NSTX)[68] and the KSTAR device[69]. The technique has become routine on JET but it may not be appropriate for use on the high-performance shots on ITER because of the demands on the vertical control system[64].

Both techniques have successfully reduced the size of ELMs as their frequency increases[70]. It is important to note that, although the size of type I ELMs reduces as the ELM frequency increases, the area on the divertor that the ELM interacts with is also reduced[37,71]. Therefore, the peak power density on the area of the divertor that the ELM interacts with remains constant. In other words, although increasing the ELM frequency reduces the damage they cause, the damage is only truly mitigated by leaving the type I ELM regime by either suppressing ELMs or moving to a small-ELM regime[72].

An additional technique to control ELMs is the application of non-axisymmetric edge magnetic-field perturbations. This technique can be used to produce smaller, more frequent ELMs or to suppress them completely. ELM suppression was first achieved on the DIII-D tokamak[73] and then on other tokamaks, including ASDEX Upgrade[74] (FIG. 5), KSTAR[75] and the Experimental Advanced Superconducting Tokamak (EAST)[76]. It is believed that resonant magnetic perturbations (RMPs) are able to suppress type I ELMs at ITER-relevant ‘low collisionality’ (which corresponds to a low ratio of density to the

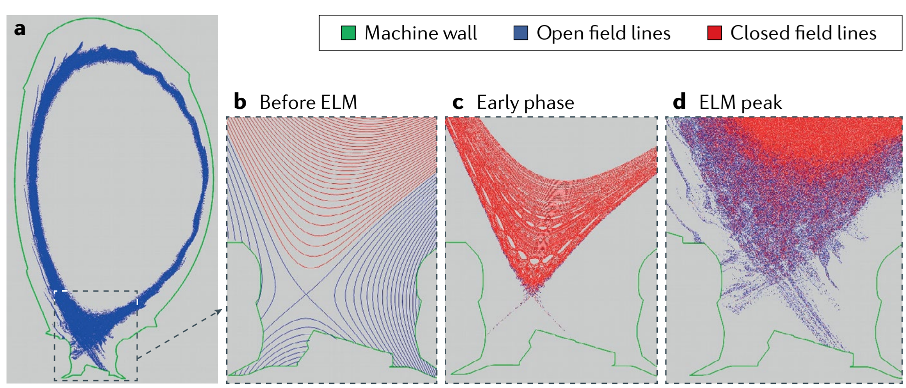

Fig. 4 | **Field lines traced in a JOREK simulation of a Joint-European-Torus-like plasma during an edge-localized mode.** **a** | The plasma cross section and region of interest. **b–d** | The evolution of the field lines during the edge-localized mode (ELM). Before the ELM, the confined region has closed (that is, infinitely long) field lines, whereas the field lines in the scrape-off layer are open and relatively short (panel **b**). In the early phase of the ELM, the field lines are perturbed and magnetic islands are produced (panel **c**). At the ELM peak, these magnetic islands have become large enough to mix open and closed field lines in the edge region of the plasma (panel **d**).

square of temperature) because the pressure gradient in the edge is held below the peeling–ballooning stability limit. The goal of the fusion community is to eliminate ELMs altogether while holding the plasma in the highest performance H-mode. However, the price paid for suppression of the ELMs is a drop of 10–20% in the overall stored energy of the plasma (FIG. 5).

It is still uncertain how this suppressed state is established and, in particular, how transport arises to stop the edge transport barrier evolving back to the parameter regime in which the peeling–ballooning mode is unstable. Part of the mechanism is that the edge pressure gradient decreases — due mainly to the density drop in the edge region, the so-called 'pump-out effect' — while the edge temperature does not drop and might even increase. Experimentally, pump-out is observed once the magnetic perturbations are large enough and well enough aligned with the pitch of the magnetic field lines[77]. Although this enhanced density transport may be caused by ergodic fields in the plasma edge, it may also be due to the lack of axisymmetry induced by the magnetic fields on small-scale instabilities that drive turbulence. Another part of the suppression mechanism that has been proposed is that a magnetic island forms close to the top of the edge transport barrier, which then prevents the edge transport barrier width broadening, thus keeping the peeling–ballooning instability threshold high[78]. However, the production of such a well-localized island places stringent constraints on the flow profiles of the plasma that would produce very strict access conditions. Recent work has shown that such flow profiles are not a necessary constraint for ELM suppression[74,79]. These studies suggest that 3D plasma distortions may lead to turbulent transport effects playing a role. For example, recent work has shown experimentally and theoretically that ballooning instabilities can be localized toroidally when magnetic perturbations are applied for ELM control, and such localization may lead to enhanced transport[80].

Our understanding of what is required for ELM suppression has advanced significantly in recent years, but a challenge is to reliably extrapolate these results to the conditions required for ELM suppression on ITER. Developing such a predictive capability for an RMP ELM-control strategy is an active area of research.

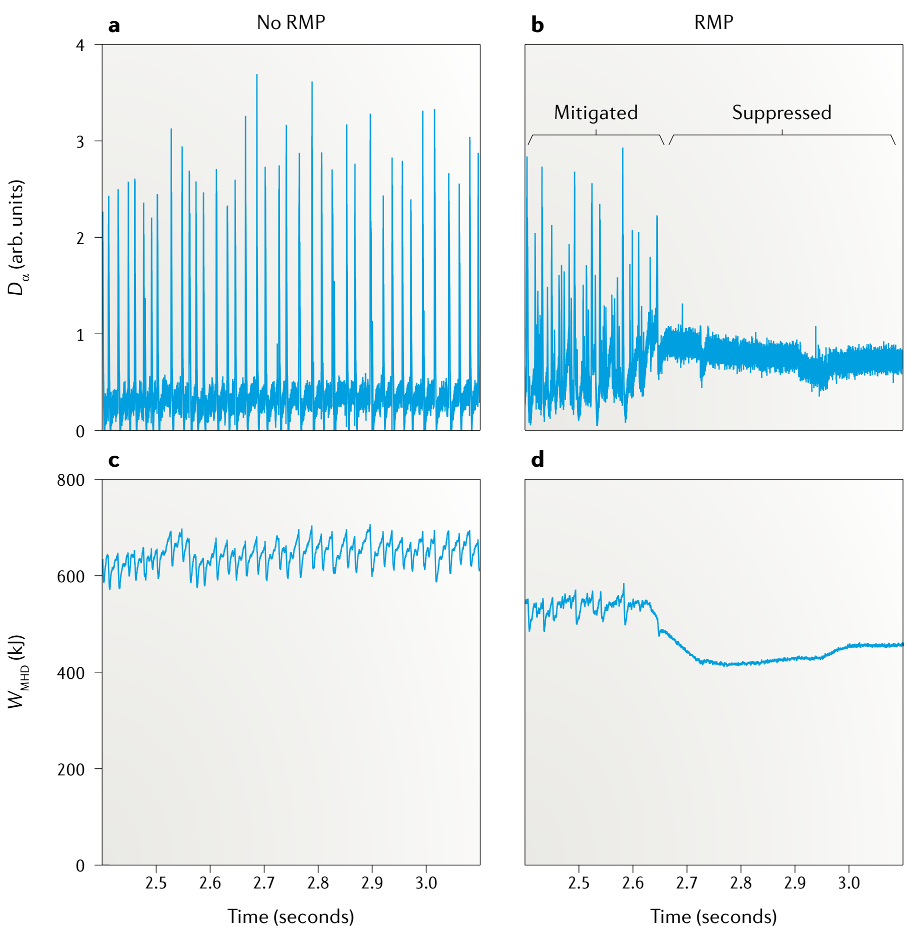

Fig. 5 | **Suppression of edge-localized modes at the Axially Symmetric Divertor Experiment Upgrade tokamak using resonant magnetic perturbations. a,b** | Time traces of D_α light emission from the divertor, which is a characteristic signature of an edge-localized mode event, without a resonant magnetic perturbation (RMP) (panel **a**) and with an RMP applied (panel **b**). **c,d** | The stored energy in the plasma (W_{MHD}) as calculated from a magnetohydrodynamics equilibrium code, without (panel **c**) and with (panel **d**) an RMP. Data extracted from REF.[74].

Progress and open questions

Progress on ELMs and their control. The understanding of ELMs and their mitigation/suppression in tokamak plasmas has advanced significantly over the past two decades, enabled largely through targeted, collaborative international research programmes. That understanding has revealed a rich variety of plasma physics that challenges theory, experiment and advanced simulation to provide a quantitative predictive capability. This challenge pushes the boundaries of plasma physics, driven both by inherent scientific curiosity and a need for a viable solution for future fusion reactors.

Looking back approximately 20 years at the JET DT experiments of 1997 (REF.[2]), in those experiments, an H-mode with ELM activity was used with no real concern about ELMs in ITER and future machines. By 2002, concern had grown about how the power loadings would scale to ITER and future tokamaks, and, hence, the damage that ELMs would cause to the divertor and the consequent shortening of the divertor lifetime[81]. This concern stimulated efforts to look at ELM-control strategies. ELM suppression using RMPs was first demonstrated by the DIII-D tokamak in 2004 (REF.[73]).

The work on understanding and modelling ELMs was also well under way. Fast visible cameras on MAST confirmed the filamentary nature of the ELMs in 2004 (REF.[9]), which had been predicted by the nonlinear ballooning theory a little earlier[30]. Several computational approaches were investigated, such as NIMROD[82], BOUT++[3,83] and JOREK[49]. In particular, the JOREK code was designed and built specifically to carry out nonlinear MHD modelling of the ELMs, and the first paper using this code was published in 2007 (REF.[49]).

The work over the past 10 years has been to refine and explore these initial important steps. ELM suppression has now been demonstrated on many machines. More diagnostics have been installed on tokamaks around the world and these have demonstrated the filamentary nature of the instability and have also made it possible to quantify how heat and particles get to the divertor plates. Particle acceleration, which is indicative of magnetic reconnection, has been observed. Modelling has become increasingly sophisticated, with more physics now routinely included, and higher-fidelity simulations have recovered experimental results, providing more confidence in extrapolation to ITER. ELM-free regimes, which may be necessary for future machines, have been identified.

Outlook. ELM suppression will almost certainly be needed in future fusion power plants, but it must be achieved with the smallest possible loss of fusion performance. RMPs suppress ELMs by increasing transport, preventing the plasma from reaching the peeling–ballooning stability limit, but understanding and extrapolating the effects on confinement remains an open challenge. It is likely that modelling these effects will require capturing the challenging physics associated with reconnection and a stochastic magnetic field, which will probably need physics models that go beyond the relatively simple fluid models of today's simulations. Such models will make it possible to design and test ELM-control schemes in silico and, thus, enable confidence in the design and operation of reactor regimes.

Although they have not been the subject of this Perspective, which focuses primarily on ELM filament physics and control, there are, in fact, naturally existing ELM-free plasma operational regimes in current devices that, if extrapolable, could provide the high level of confinement required for ITER to achieve its fusion-power objectives. Amongst these are the so-called I-mode (first documented on Alcator C-Mod[84]) and the QH-mode (as demonstrated on DIII-D[85]). There are also regimes with much smaller ELMs, such as 'grassy ELMs', seen on JT-60U[86] for example, which are also being investigated and may extrapolate favourably for certain regimes. If these operational regimes can be used, it will be possible to avoid installing new hardware into the machine, but it must be certain that they exist within a robust operating space and that they do not degrade the confinement too much. These regimes are also the subject of much research to assess whether their beneficial properties extrapolate to ITER and demonstration fusion power plants.

The plasma in a fusion reactor has many degrees of freedom and there are many actuators. Therefore, it is possible that new techniques or operating regimes will be found that optimize fusion performance without ELMs. A synergy of theoretical understanding, computer modelling and experiments will help achieve this.

Christopher Ham[1] ✉, *Andrew Kirk*[1], *Stanislas Pamela*[1] *and Howard Wilson*[1,2]

[1]*UKAEA – CCFE, Culham Science Centre, Abingdon, UK.*

[2]*York Plasma Institute, Department of Physics, University of York, Heslington, UK.*

✉*e-mail: christopher.ham@ukaea.uk*

https://doi.org/10.1038/s42254-019-0144-1

Published online 14 February 2020

1. Leonard, A. W. Edge-localized-modes in tokamaks. *Phys. Plasmas* **21**, 090501 (2014).
2. JET Team (prepared by Watkins, M. L.). Physics of high performance JET plasmas in DT. *Nucl. Fusion* **39**, 1227–1244 (1999).
3. Dudson, B. D. et al. Experiments and simulation of edge turbulence and filaments in MAST. *Plasma Phys. Control. Fusion* **50**, 124012 (2008).
4. Wagner, F. A quarter-century of H-mode studies. *Plasma Phys. Control. Fusion* **49**, B1–B33 (2007).
5. Connor, J. W. Edge-localized modes — physics and theory. *Plasma Phys. Control. Fusion* **40**, 531–542 (1998).
6. Suttrop, W. The physics of large and small edge localized modes. *Plasma Phys. Control. Fusion* **42**, A1–A14 (2000).
7. Loarte, A. et al. Characteristics of type I ELM energy and particle losses in existing devices and their extrapolation to ITER. *Plasma Phys. Control. Fusion* **45**, 1549–1569 (2003).
8. Cowley, S. C., Wilson, H., Hurricane, O. & Fong, B. Explosive instabilities: from solar flares to edge localized modes in tokamaks. *Plasma Phys. Control. Fusion* **45**, A31–A38 (2003).
9. Kirk, A. et al. Spatial and temporal structure of edge-localized modes. *Phys. Rev. Lett.* **92**, 245002 (2004).
10. Kirk, A. et al. Physics of ELM power fluxes to plasma facing components and implications for ITER. *J. Nucl. Mater.* **390–391**, 727–732 (2009).
11. Yun, G. S. et al. Two-dimensional visualization of growth and burst of the edge-localized filaments in KSTAR H-mode plasmas. *Phys. Rev. Lett.* **107**, 045004 (2011).
12. Freidberg, J. P. *Ideal MHD* (Cambridge Univ. Press, 2014) ISBN 978-1-107-00625-6.
13. Connor, J. W., Hastie, R. J. & Taylor, J. B. High mode number stability of an axisymmetric toroidal plasma. *Proc. R. Soc. Lond. A* **365**, 1–17 (1979).
14. Wilson, H. R., Cowley, S. C., Kirk, A. & Snyder, P. B. Magneto-hydrodynamic stability of the H-mode transport barrier as a model for edge localized modes: an overview. *Plasma Phys. Control. Fusion* **48**, A71–A84 (2006).
15. Bickerton, R. J., Connor, J. W. & Taylor, J. B. Diffusion driven plasma currents and bootstrap tokamak. *Nat. Phys. Sci.* **229**, 110–112 (1971).
16. Peeters, A. G. The bootstrap current and its consequences. *Plasma Phys. Control. Fusion* **42**, B231–B242 (2000).
17. Hegna, C. C., Connor, J. W., Hastie, R. J. & Wilson, H. R. Toroidal coupling of ideal magnetohydrodynamic instabilities in tokamak plasmas. *Phys. Plasmas* **3**, 584–592 (1996).
18. Connor, J. W., Hastie, R. J. & Wilson, H. R. Magnetohydrodynamic stability of tokamak edge plasmas. *Phys. Plasmas* **5**, 2687–2700 (1998).
19. Snyder, P. B. et al. Edge localized modes and the pedestal: a model based on coupled peeling–ballooning modes. *Phys. Plasmas* **9**, 2037–2043 (2002).
20. Burckhart, A. et al. Inter-ELM behaviour of the electron density and temperature pedestal in ASDEX upgrade. *Plasma Phys. Control. Fusion* **52**, 105010 (2010).
21. Dickinson, D. et al. Kinetic instabilities that limit β in the edge of a tokamak plasma: a picture of an H-mode pedestal. *Phys. Rev. Lett.* **108**, 135002 (2012).
22. Hatch, D. R. et al. Gyrokinetic study of ASDEX upgrade inter-ELM pedestal profile evolution. *Nucl. Fusion* **55**, 063028 (2015).
23. Snyder, P. B. et al. A first-principles predictive model of the pedestal height and width: development, testing and ITER optimization with the EPED model. *Nucl. Fusion* **51**, 103016 (2011).
24. Saarelma, S. et al. Integrated modelling of H-mode pedestal and confinement in JET-ILW. *Plasma Phys. Control. Fusion* **60**, 014042 (2018).
25. Maggi, C. F. et al. Pedestal confinement and stability in JET-ILW ELMy H-modes. *Nucl. Fusion* **55**, 113031 (2015).
26. Bowman, C. et al. Pedestal evolution physics in low triangularity JET tokamak discharges with ITER-like wall. *Nucl. Fusion* **58**, 016021 (2018).
27. Snyder, P. B. et al. Super H-mode: theoretical prediction and initial observations of a new high performance regime for tokamak operation. *Nucl. Fusion* **55**, 083026 (2015).
28. Zohm, H. Edge localized modes (ELMs). *Plasma Phys. Control. Fusion* **38**, 105–128 (1996).
29. Oyama, N. et al. Energy loss for grassy ELMs and effects of plasma rotation on the ELM characteristics in JT-60U. *Nucl. Fusion* **45**, 871–881 (2005).
30. Wilson, H. R. & Cowley, S. C. Theory for explosive ideal magnetohydrodynamic instabilities in plasmas. *Phys. Rev. Lett.* **92**, 175006 (2004).
31. Ham, C. J., Cowley, S. C., Brochard, G. & Wilson, H. R. Nonlinear stability and saturation of ballooning modes in tokamaks. *Phys. Rev. Lett.* **116**, 235001 (2016).
32. Ham, C. J., Cowley, S. C., Brochard, G. & Wilson, H. R. Nonlinear ballooning modes in tokamaks: stability and saturation. *Plasma Phys. Control. Fusion.* **60**, 075017 (2018).
33. Hutchinson, I. H. *Principles of Plasma Diagnostics* (Cambridge Univ. Press, 2002)
34. Kirk, A. et al. Evolution of the pedestal on MAST and the implications for ELM power loadings. *Plasma Phys. Control. Fusion* **49**, 1259–1275 (2007).
35. Vianello, N. et al. Direct observation of current in Type-I edge-localized-Mode filaments on the ASDEX upgrade tokamak. *Phys. Rev. Lett.* **106**, 125002 (2011).
36. Beurskens, M. N. A. et al. Pedestal and scrape-off layer dynamics in ELMy H-mode plasmas in JET. *Nucl. Fusion* **49**, 125006 (2009).
37. Eich, T. et al. ELM divertor peak energy fluence scaling to ITER with data from JET, MAST and ASDEX upgrade. *Nucl. Mater. Energy* **12**, 84–90 (2017).
38. Becoulet, M. et al. Edge localized mode physics and operational aspects in tokamaks. *Plasma Phys. Control. Fusion* **45**, A93–A113 (2003).
39. Alladio, F., Mancuso, A. & Micozzi, P. Rotating twisted filaments buoyancy: comparison between the convective region of the sun and the edge of a tokamak plasma. *Plasma Phys. Control. Fusion* **50**, 124019 (2008).
40. Evans, T. E. et al. A conceptual model of the magnetic topology and nonlinear dynamics of ELMs. *J. Nucl. Mater.* **390**, 789–792 (2009).
41. Rack, M. et al. Thermoelectric currents and their role during ELM formation in JET. *Nucl. Fusion* **52**, 074012 (2012).
42. Freethy, S. J. et al. Electron kinetics inferred from observations of microwave bursts during edge localized modes in the mega-amp spherical tokamak. *Phys. Rev. Lett.* **114**, 125004 (2015).
43. Galdon-Quiroga, J. et al. Beam-ion acceleration during edge localized modes in the ASDEX upgrade tokamak. *Phys. Rev. Lett.* **121**, 025002 (2018).
44. Galdon-Quiroga, J. et al. Velocity space resolved absolute measurement of fast ion losses induced by a tearing mode in the ASDEX Upgrade tokamak. *Nucl. Fusion* **58**, 036005 (2018).
45. Pamela, S. J. P. et al. Nonlinear MHD simulations of edge-localized-modes in JET. *Plasma Phys. Control. Fusion* **53**, 054014 (2011).
46. Pamela, S. J. P. et al. Recent progress in the quantitative validation of JOREK simulations of ELMs in JET. *Nucl. Fusion* **57**, 076006 (2017).
47. Snyder, P. B. et al. Pedestal stability comparison and ITER pedestal prediction. *Nucl. Fusion* **49**, 085035 (2009).
48. Bécoulet, M. et al. Non-linear MHD modelling of edge localized modes dynamics in KSTAR. *Nucl. Fusion* **57**, 116059 (2017).
49. Huysmans, G. T. A. & Czarny, O. MHD stability in X-point geometry: simulation of ELMs. *Nucl. Fusion* **47**, 659–666 (2007).
50. Ebrahimi, F. Nonlinear reconnecting edge localized modes in current-carrying plasmas. *Phys. Plasmas* **24**, 056119 (2017).

51. Kirk, A. et al. Evolution of filament structures during edge-localized modes in the MAST tokamak. *Phys. Rev. Lett.* **96**, 185001 (2006).
52. Mink, A. F. et al. Nonlinear coupling induced toroidal structure of edge localized modes. *Nucl. Fusion* **58**, 026011 (2018).
53. Pamela, S. J. P. et al. Non-linear MHD simulations of ELMs in JET and quantitative comparisons to experiments. *Plasma Phys. Control. Fusion* **58**, 014026 (2016).
54. Kirk, A. et al. Recent progress in understanding the processes underlying the triggering of and energy loss associated with type I ELMs. *Nucl. Fusion* **54**, 114012 (2014).
55. Huysmans, G. T. A., Pamela, S. J. P., van der Plas, E. & Ramet, P. Non-linear MHD simulations of edge localized modes (ELMs). *Plasma Phys. Control. Fusion* **51**, 124012 (2009).
56. Henneberg, S. A., Cowley, S. C. & Wilson, H. R. Explosive ballooning mode instability in tokamaks: modelling the ELM cycle. *41st EPS Conference, Berlin, Germany P1.066* (2014).
57. Henneberg, S. A., Cowley, S. C. & Wilson, H. R. Interacting filamentary eruptions in magnetised plasmas. *Plasma Phys. Control. Fusion* **57**, 125010 (2015).
58. Brodrick, J. P. et al. Testing nonlocal models of electron thermal conduction for magnetic and inertial confinement fusion applications. *Phys. Plasmas* **24**, 092309 (2017).
59. Snyder, P. B. et al. ELMs and constraints on the H-mode pedestal: peeling–ballooning stability calculation and comparison with experiment. *Nucl. Fusion* **44**, 320–328 (2004).
60. Herrmann, A. Overview on stationary and transient divertor heat loads. *Plasma Phys. Control. Fusion* **44**, 883–903 (2002).
61. Lang, P. et al. ELM pace making and mitigation by pellet injection in ASDEX Upgrade. *Nucl. Fusion* **44**, 665–677 (2004).
62. Baylor, L. R. et al. Reduction of edge-localized mode intensity using high-repetition-rate pellet injection in tokamak H-mode plasmas. *Phys. Rev. Lett.* **110**, 245001 (2013).
63. Futatani, S. et al. Non-linear MHD modelling of ELM triggering by pellet injection in DIII-D and implications for ITER. *Nucl. Fusion* **54**, 073008 (2014).
64. Lang, P. T. et al. ELM control strategies and tools: status and potential for ITER. *Nucl. Fusion* **53**, 043004 (2013).
65. Degeling, A. W. et al. Magnetic triggering of ELMs in TCV. *Plasma Phys. Control. Fusion* **45**, 1637 (2003).
66. Lang, P. T. et al. Frequency control of type-I ELMs by magnetic triggering in ASDEX Upgrade. *Plasma Phys. Control. Fusion* **46**, L31–L39 (2004).
67. de La Luna, E. et al. Understanding the physics of ELM pacing via vertical kicks in JET in view of ITER. *Nucl. Fusion* **56**, 026001 (2015).
68. Gerhardt, S. P. et al. First observation of ELM pacing with vertical jogs in a spherical torus. *Nucl. Fusion* **50**, 064015 (2010).
69. Jayhyun Kim et al. ELM control experiments in the KSTAR device. *Nucl. Fusion* **52**, 114011 (2012).
70. Garzotti, L. et al. Investigating pellet ELM triggering physics using the new small size pellet launcher at JET. *37th EPS Conference, Dublin, Ireland P2.131* (2010).
71. Jakubowski, M. et al. Overview of the results on divertor heat loads in RMP controlled H-mode plasmas on DIII-D. *Nucl. Fusion* **49**, 095013 (2009).
72. Viezzer, E. Access and sustainment of naturally ELM-free and small-ELM regimes. *Nucl. Fusion* **58**, 115002 (2018).
73. Evans, T. E. et al. Suppression of large edge-localized modes in high-confinement DIII-D plasmas with a stochastic magnetic boundary. *Phys. Rev. Lett.* **92**, 235003 (2004).
74. Suttrop, W. et al. Experimental conditions to suppress edge localised modes by magnetic perturbations in the ASDEX Upgrade tokamak. *Nucl. Fusion* **58**, 096031 (2018).
75. Jeon, Y. M. et al. Suppression of edge localized modes in high-confinement KSTAR plasmas by nonaxisymmetric magnetic perturbations. *Phys. Rev. Lett.* **109**, 035004 (2012).
76. Sun, Y. et al. Nonlinear transition from mitigation to suppression of the edge localized mode with resonant magnetic perturbations in the EAST tokamak. *Phys. Rev. Lett.* **117**, 115001 (2016).
77. Liu, Y. et al. Modelling of plasma response to resonant magnetic perturbation fields in MAST and ITER. *Nucl. Fusion* **51**, 083002 (2011).
78. Wade, M. R. et al. Advances in the physics understanding of ELM suppression using resonant magnetic perturbations in DIII-D. *Nucl. Fusion* **55**, 023002 (2015).
79. Paz-Soldan, C. et al. The effect of plasma shape and neutral beam mix on the rotation threshold for RMP-ELM suppression. *Nucl. Fusion* **59**, 056012 (2019).
80. Willensdorfer, M. et al. Field-line localized destabilization of ballooning modes in three-dimensional tokamaks. *Phys. Rev. Lett.* **119**, 085002 (2017).
81. Federici, G., Loarte, A. & Strohmayer, G. Assessment of erosion of the ITER divertor targets during type I ELMs. *Plasma Phys. Control. Fusion.* **45**, 1523 (2003).
82. Pankin, A. Y. et al. Modelling of ELM dynamics for DIII-D and ITER. *Plasma Phys. Control. Fusion* **49**, S63–S75 (2007).
83. Dudson, B. D. et al. Simulation of edge localized modes using BOUT++. *Plasma Phys. Control. Fusion* **53**, 054005 (2011).
84. Whyte, D. G. et al. I-mode: an H-mode energy confinement regime with L-mode particle transport in Alcator C-Mod. *Nucl. Fusion* **50**, 105005 (2010).
85. Burrell, K. H. et al. Discovery of stationary operation of quiescent H-mode plasmas with net-zero neutral beam injection torque and high energy confinement on DIII-D. *Phys. Plasmas* **23**, 056103 (2016).
86. Kamada, K. et al. Pedestal characteristics and extended high-βp ELMy H-mode regime in JT-60U. *Plasma Phys. Control. Fusion* **44**, A279–A286 (2002).
87. Kirk, A. Nuclear fusion: bringing a star down to Earth. *Contemp. Phys.* **57**, 1–18 (2016).
88. Pamela, S. et al. Resistive MHD simulation of edge-localized modes for double-null discharges in the MAST device. *Plasma Phys. Control. Fusion* **55**, 095001 (2013).

Acknowledgements
The authors gratefully acknowledge the comments on this manuscript given by Samuli Saarelma, William Morris and Jack Connor. This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014–2018 and 2019–2020 under grant agreement no. 633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. This work has been part-funded by the RCUK Energy Programme (grant number EP/P012450/1). To obtain further information on the data and models underlying this paper, please contact PublicationsManager@ukaea.uk. Simulations published in this review have benefited from the support of EUROfusion on the Marconi HPC cluster (CINECA, Italy), as well as the support of PRACE on the MareNostrum HPC cluster (BSC, Spain).

Author contributions
All authors contributed to all aspects of this Perspective.

Competing interests
The authors declare no competing interests.

Reviewer information
Nature Reviews Physics thanks F. Jenko and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Publisher's note
Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.