A 2020s Vision of CMB Lensing

Marius Millea

BERKELEY CENTER *for*
COSMOLOGICAL PHYSICS

@marius311

@cosmicmar

With: SPT Collaboration, Ethan Anderes, Ben Wandelt, Uros Seljak

Cosmology from Home 2020

Main points

- CMB lensing now plays a role in nearly all aspects of CMB observations
- Techniques to analyze lensing are at a (necessary) watershed moment, requiring more sophistication
- I want to teach you how these work so you can do it yourself or you can understand the CMB data products that you may be interacting with in the future

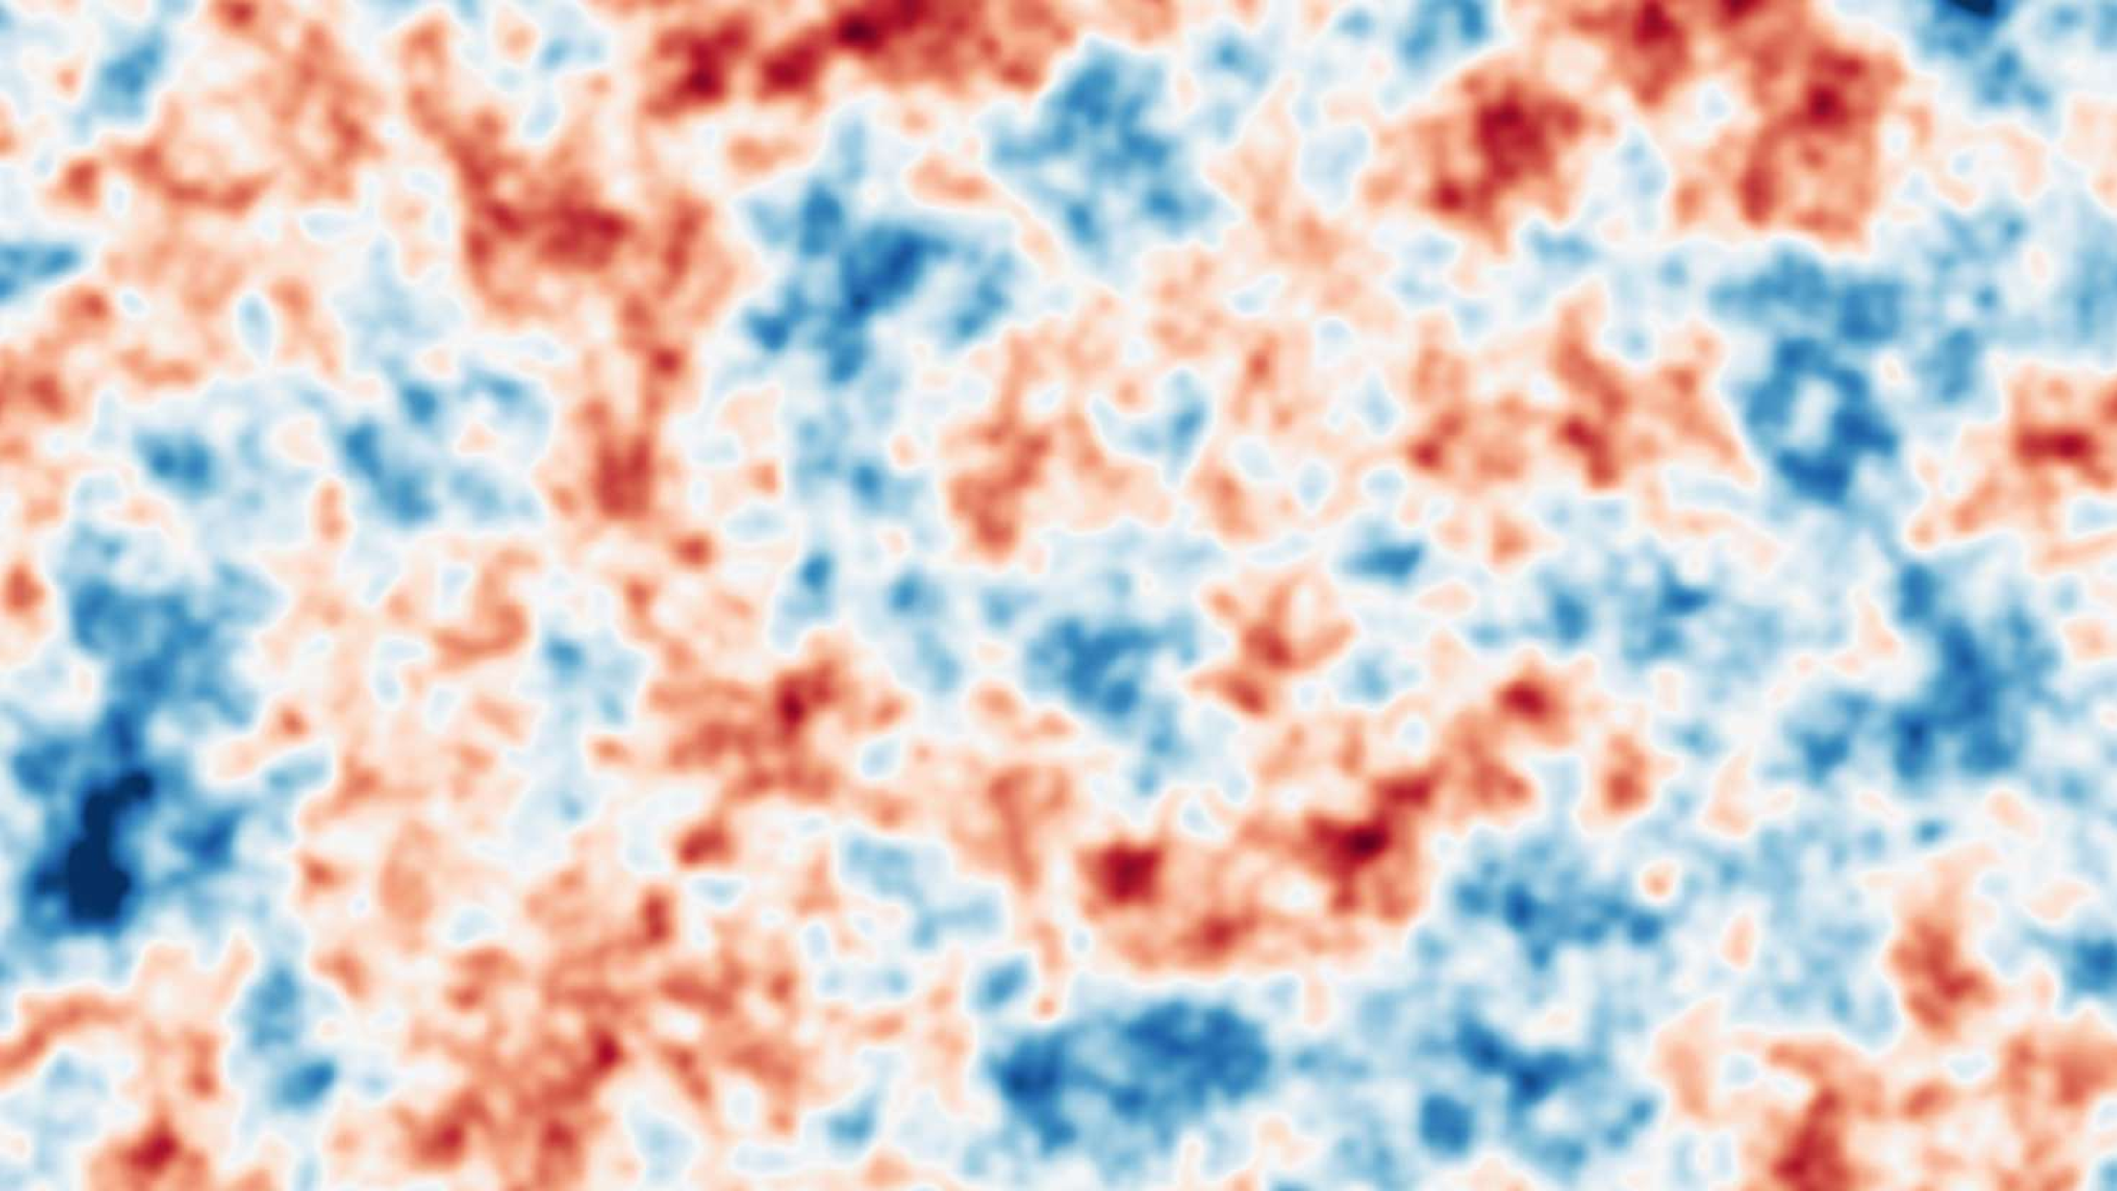

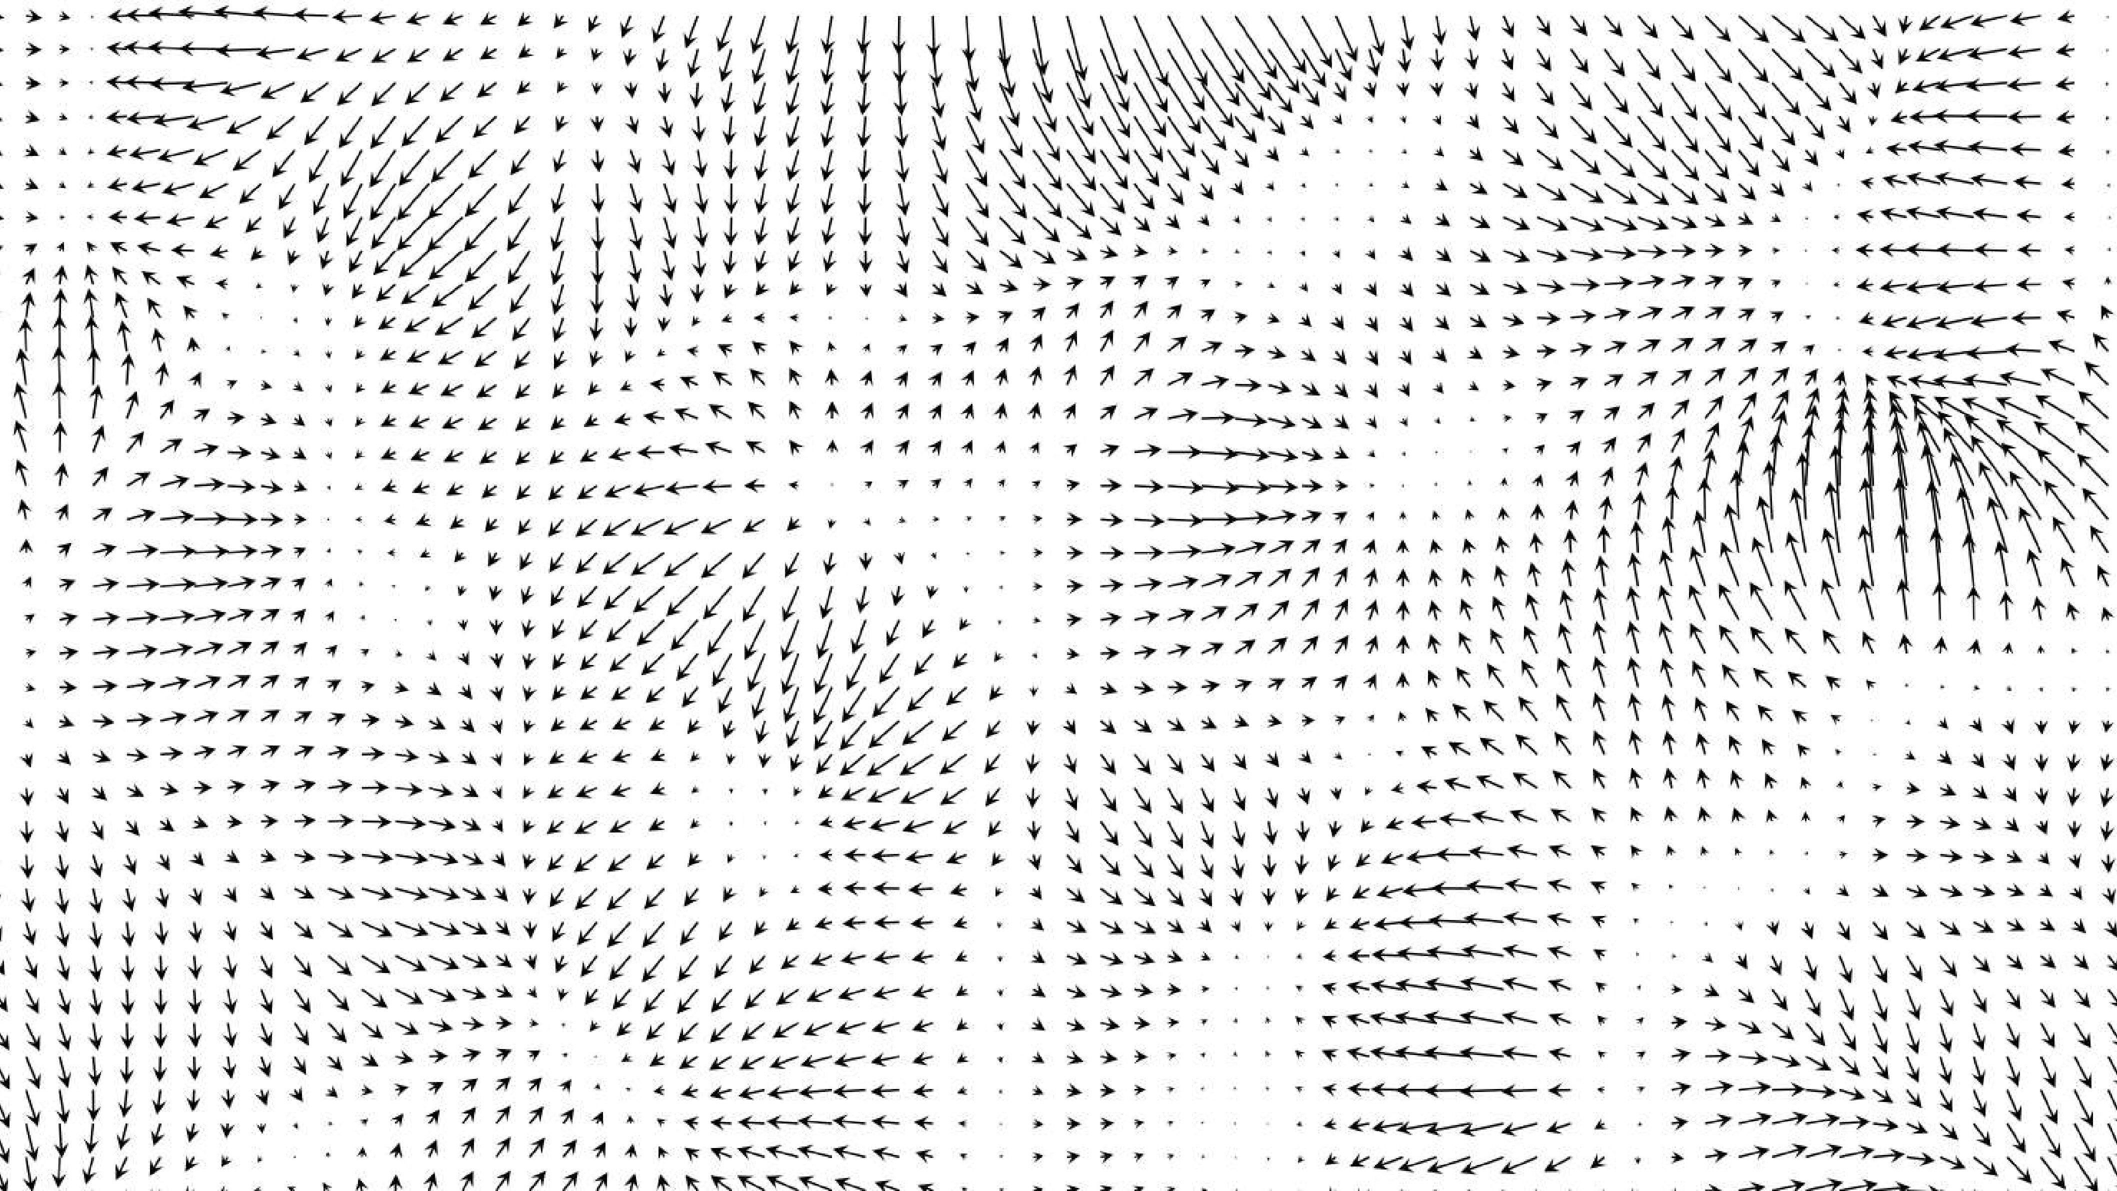

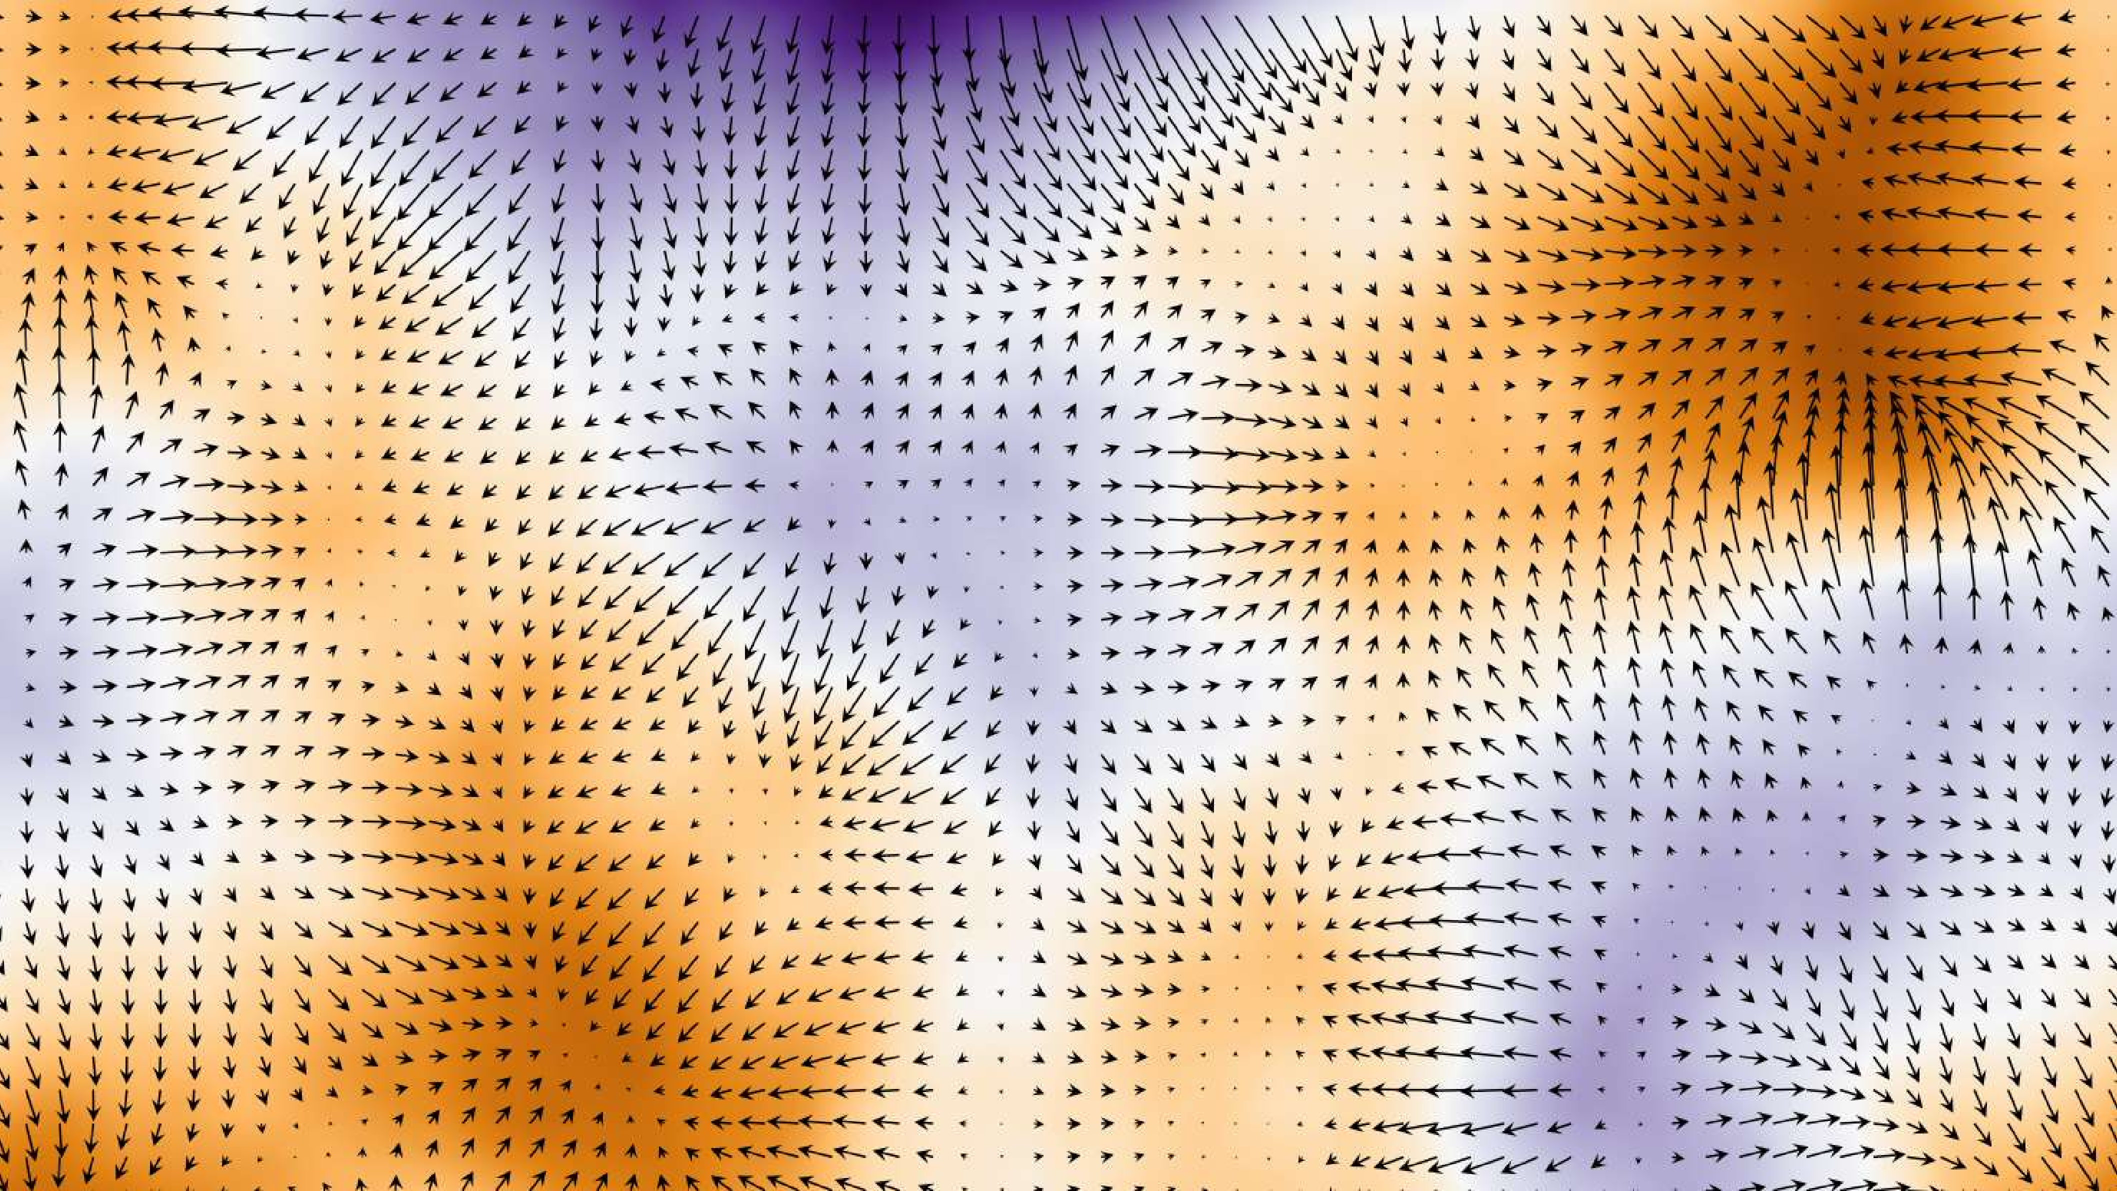

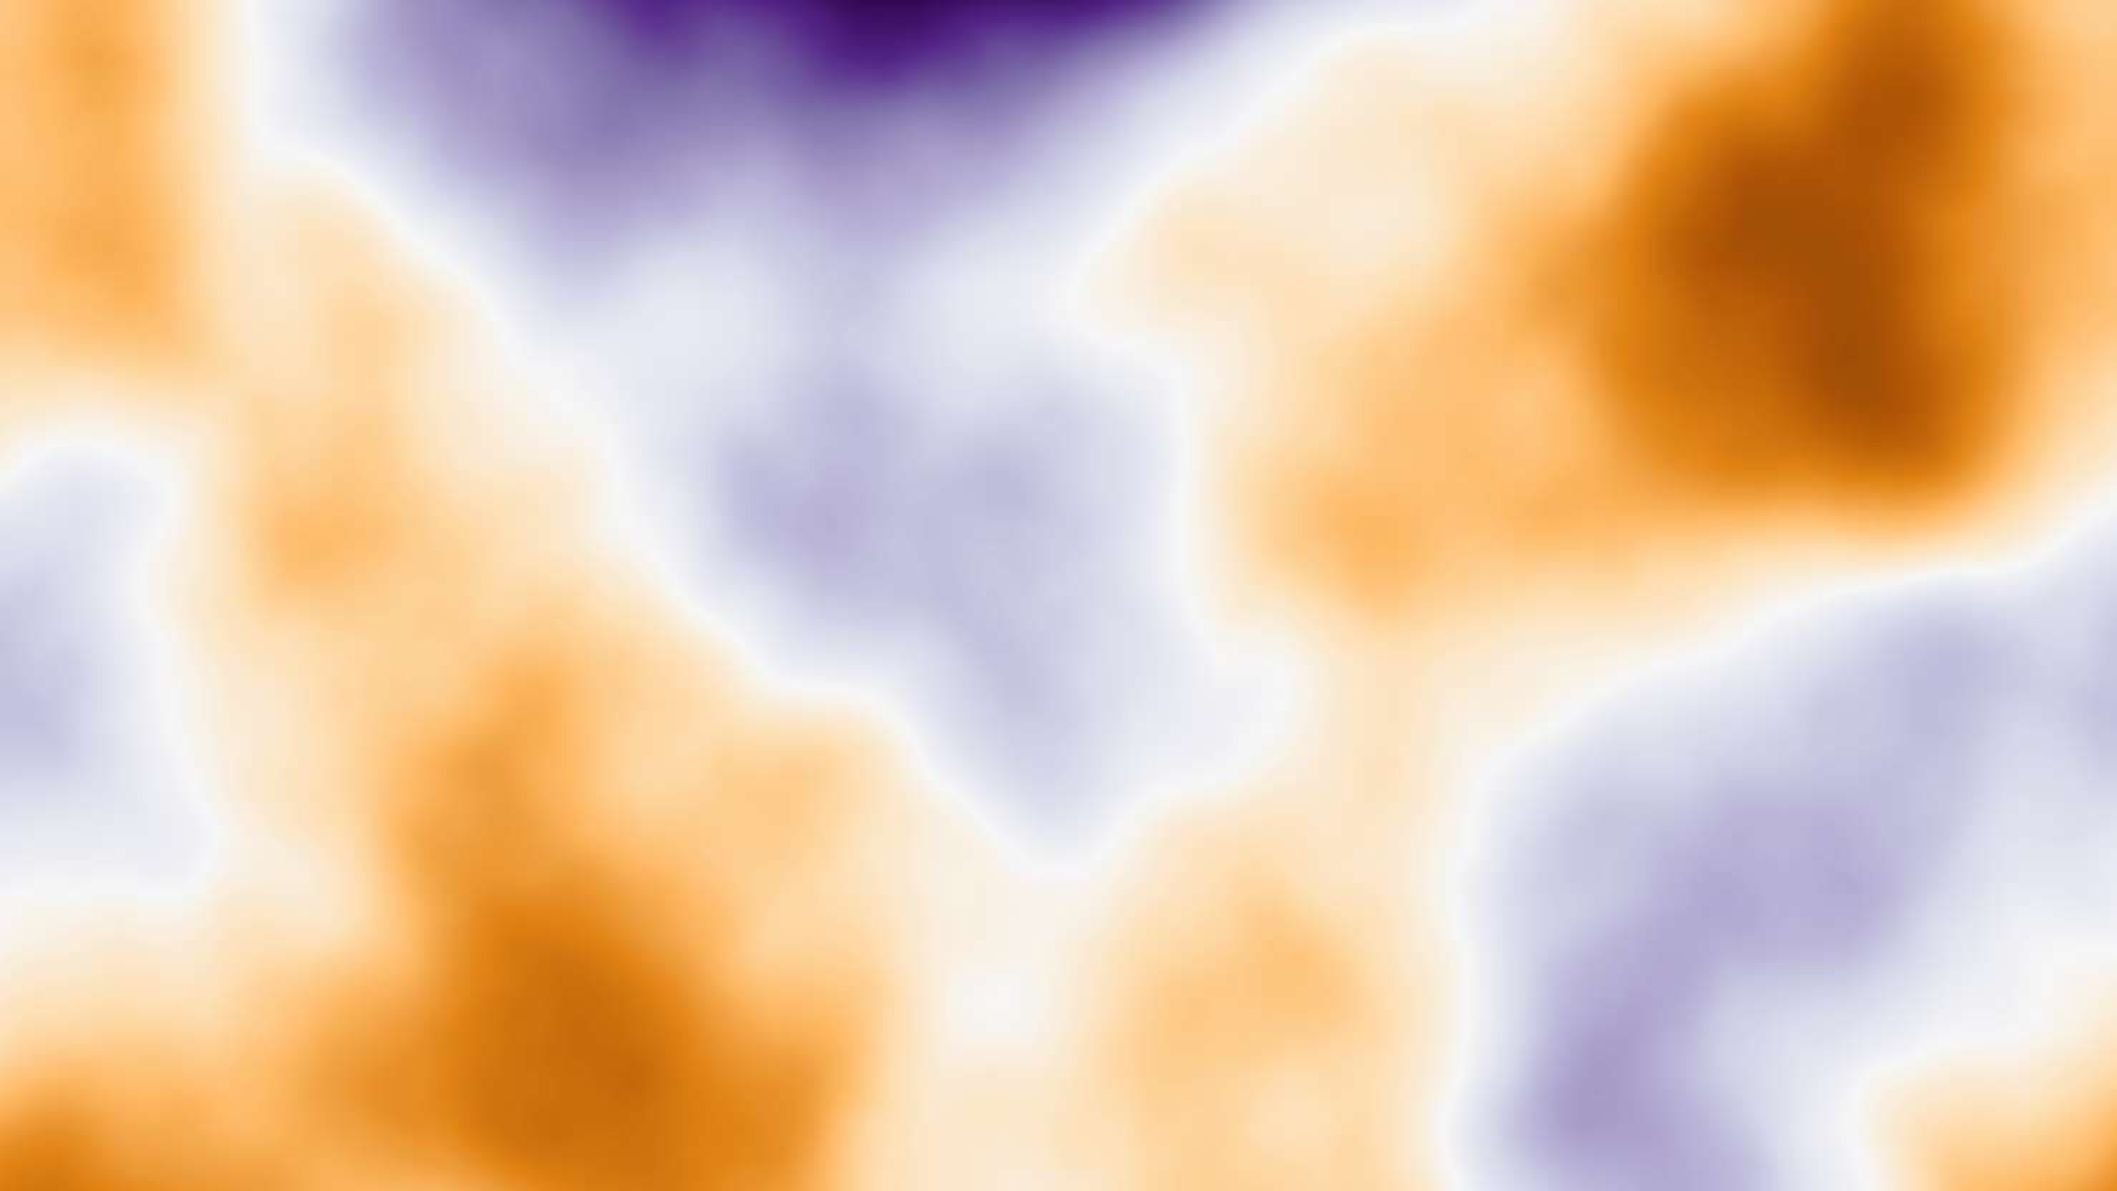

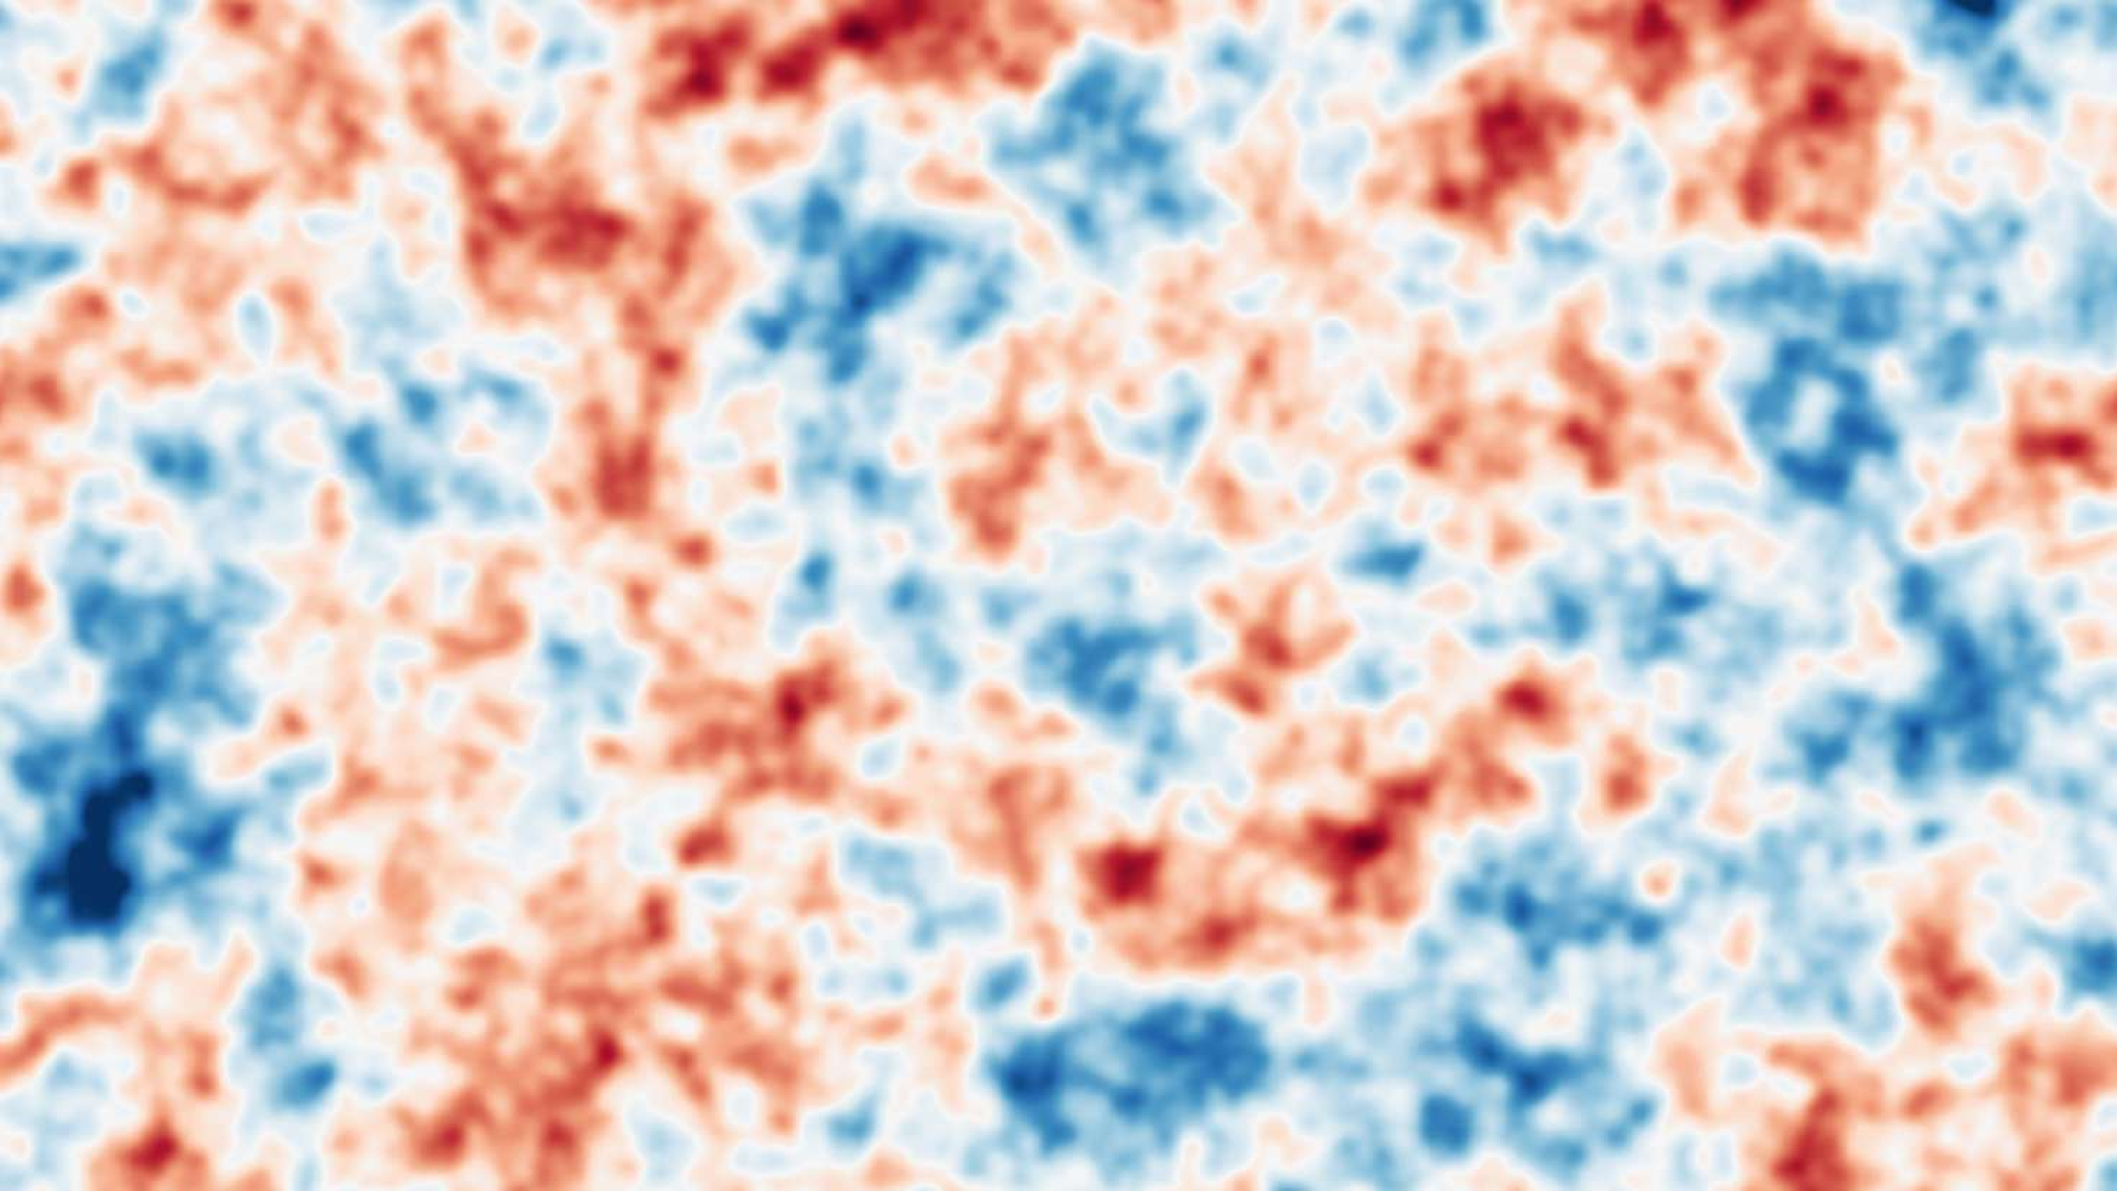

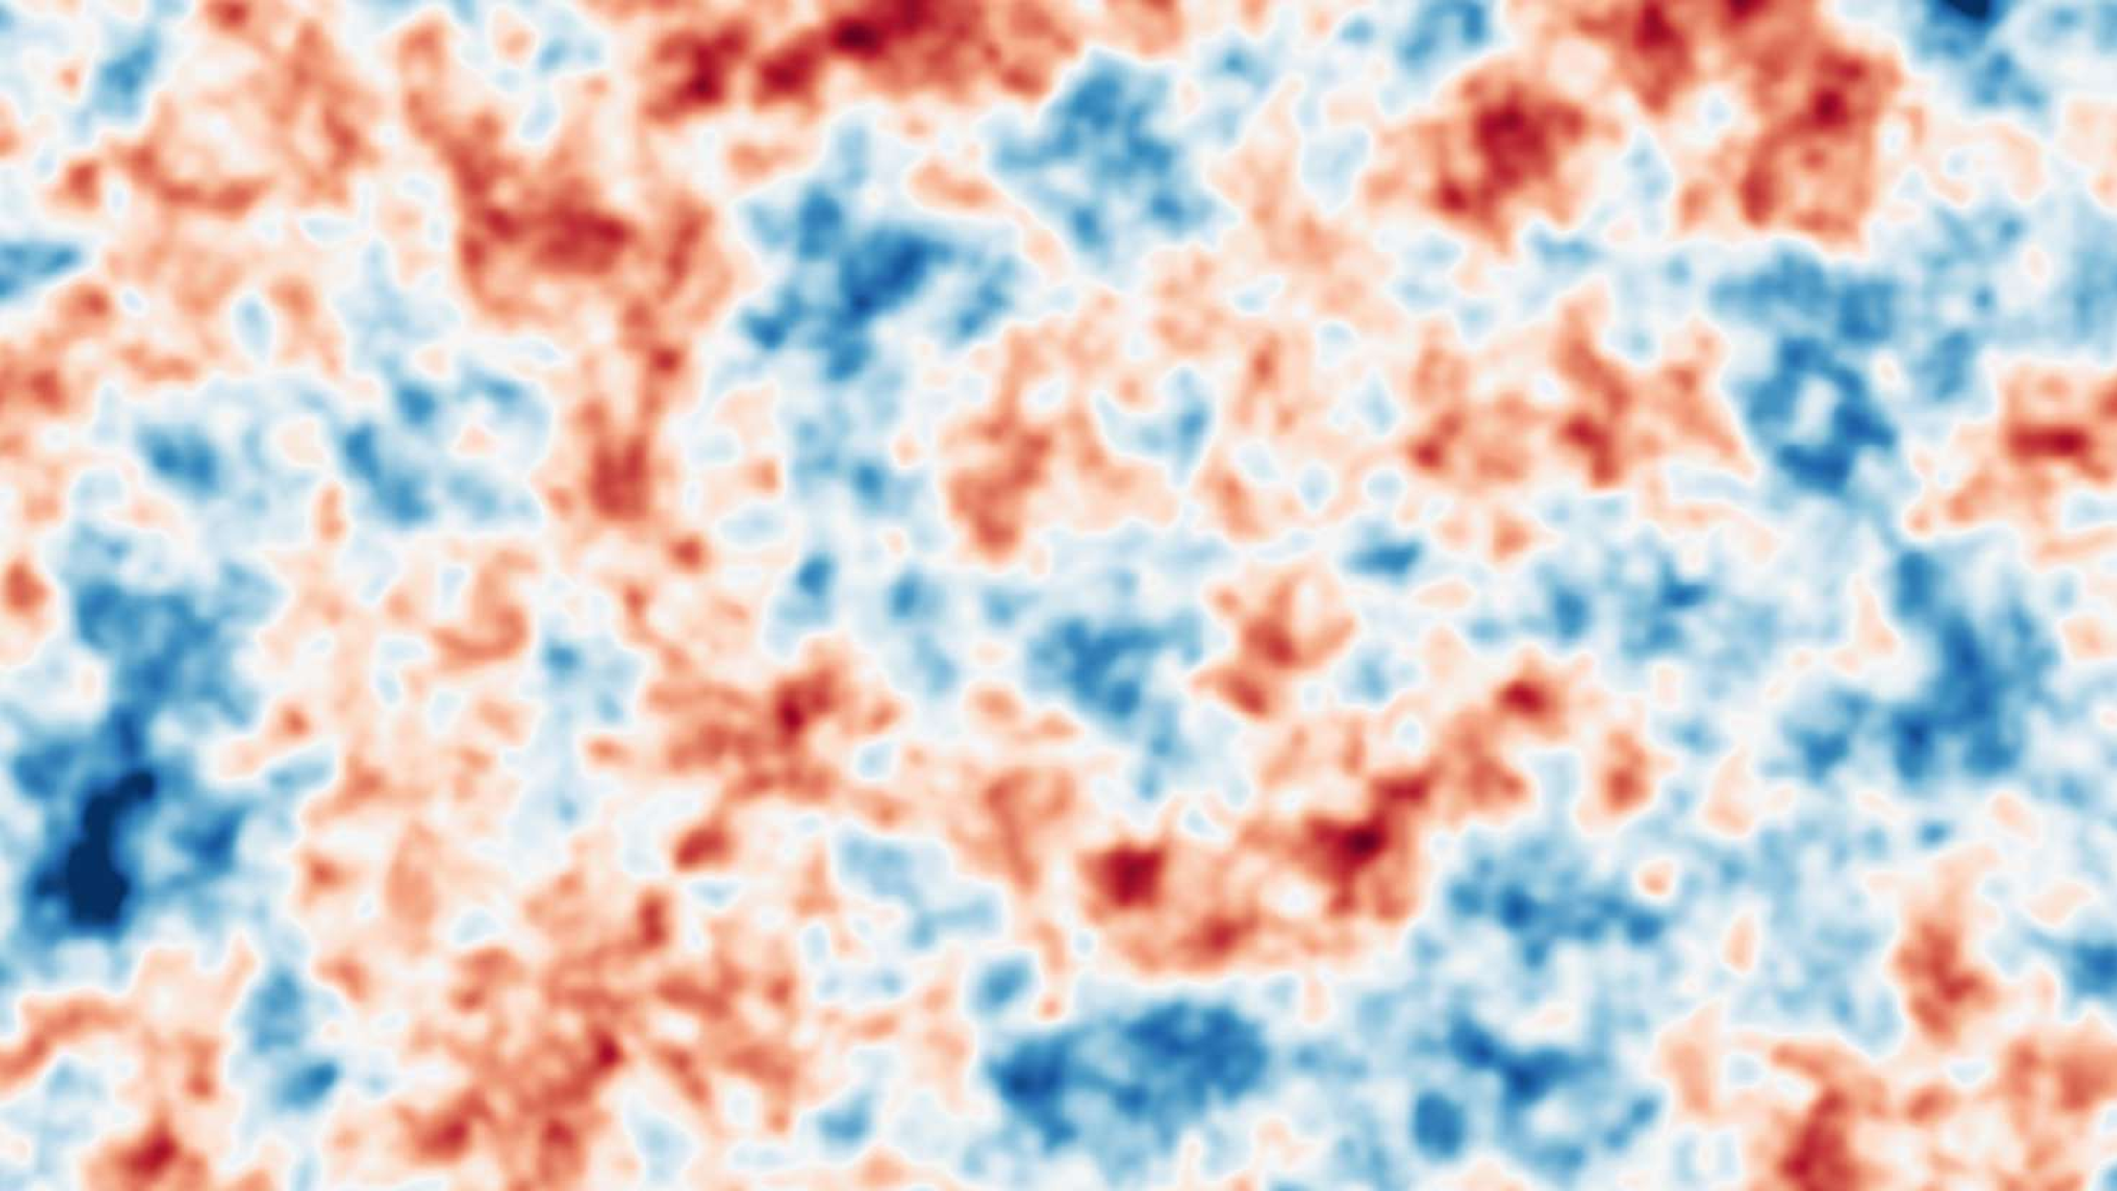

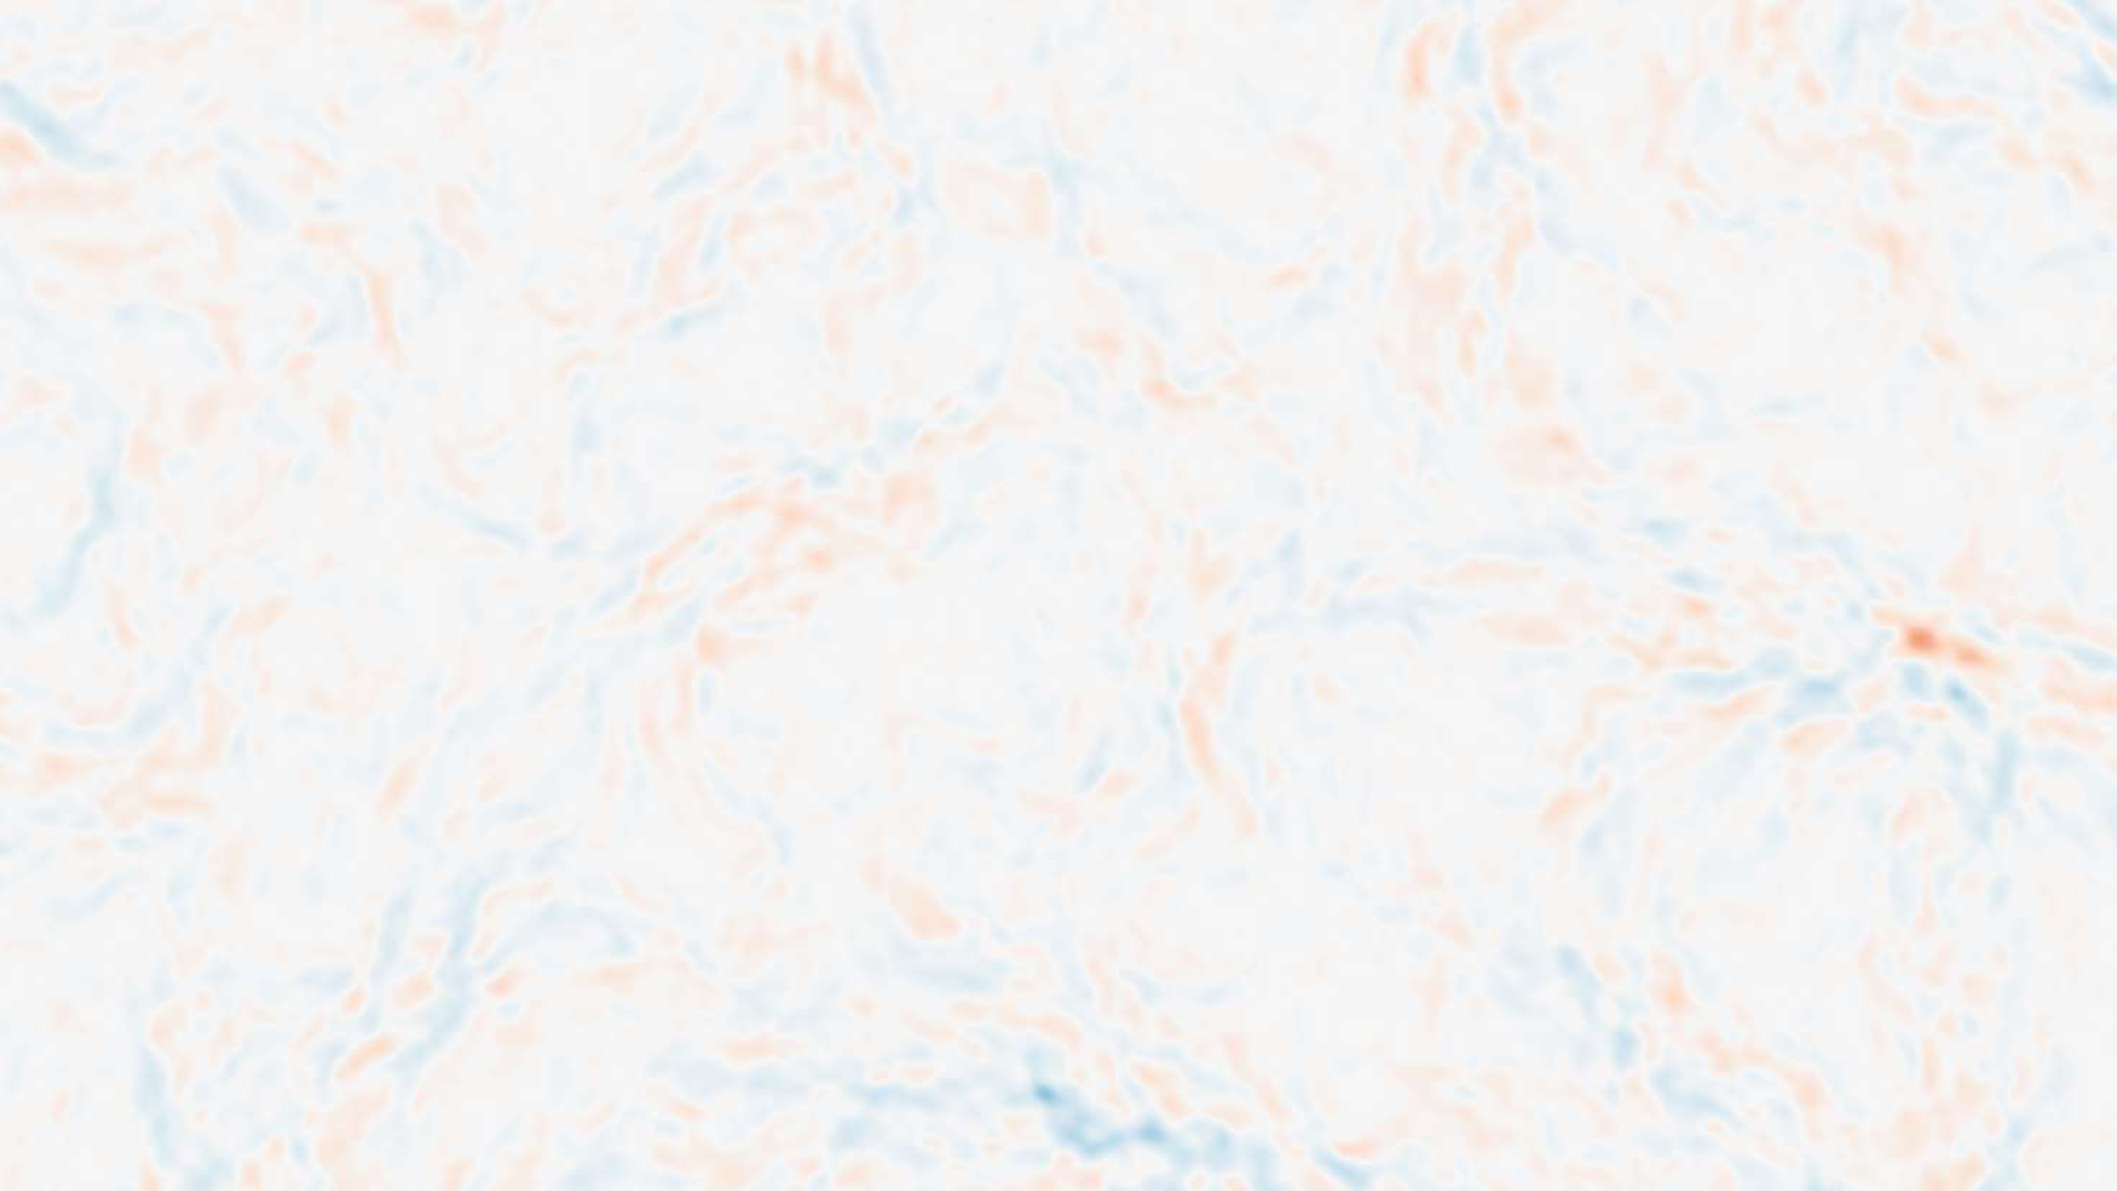

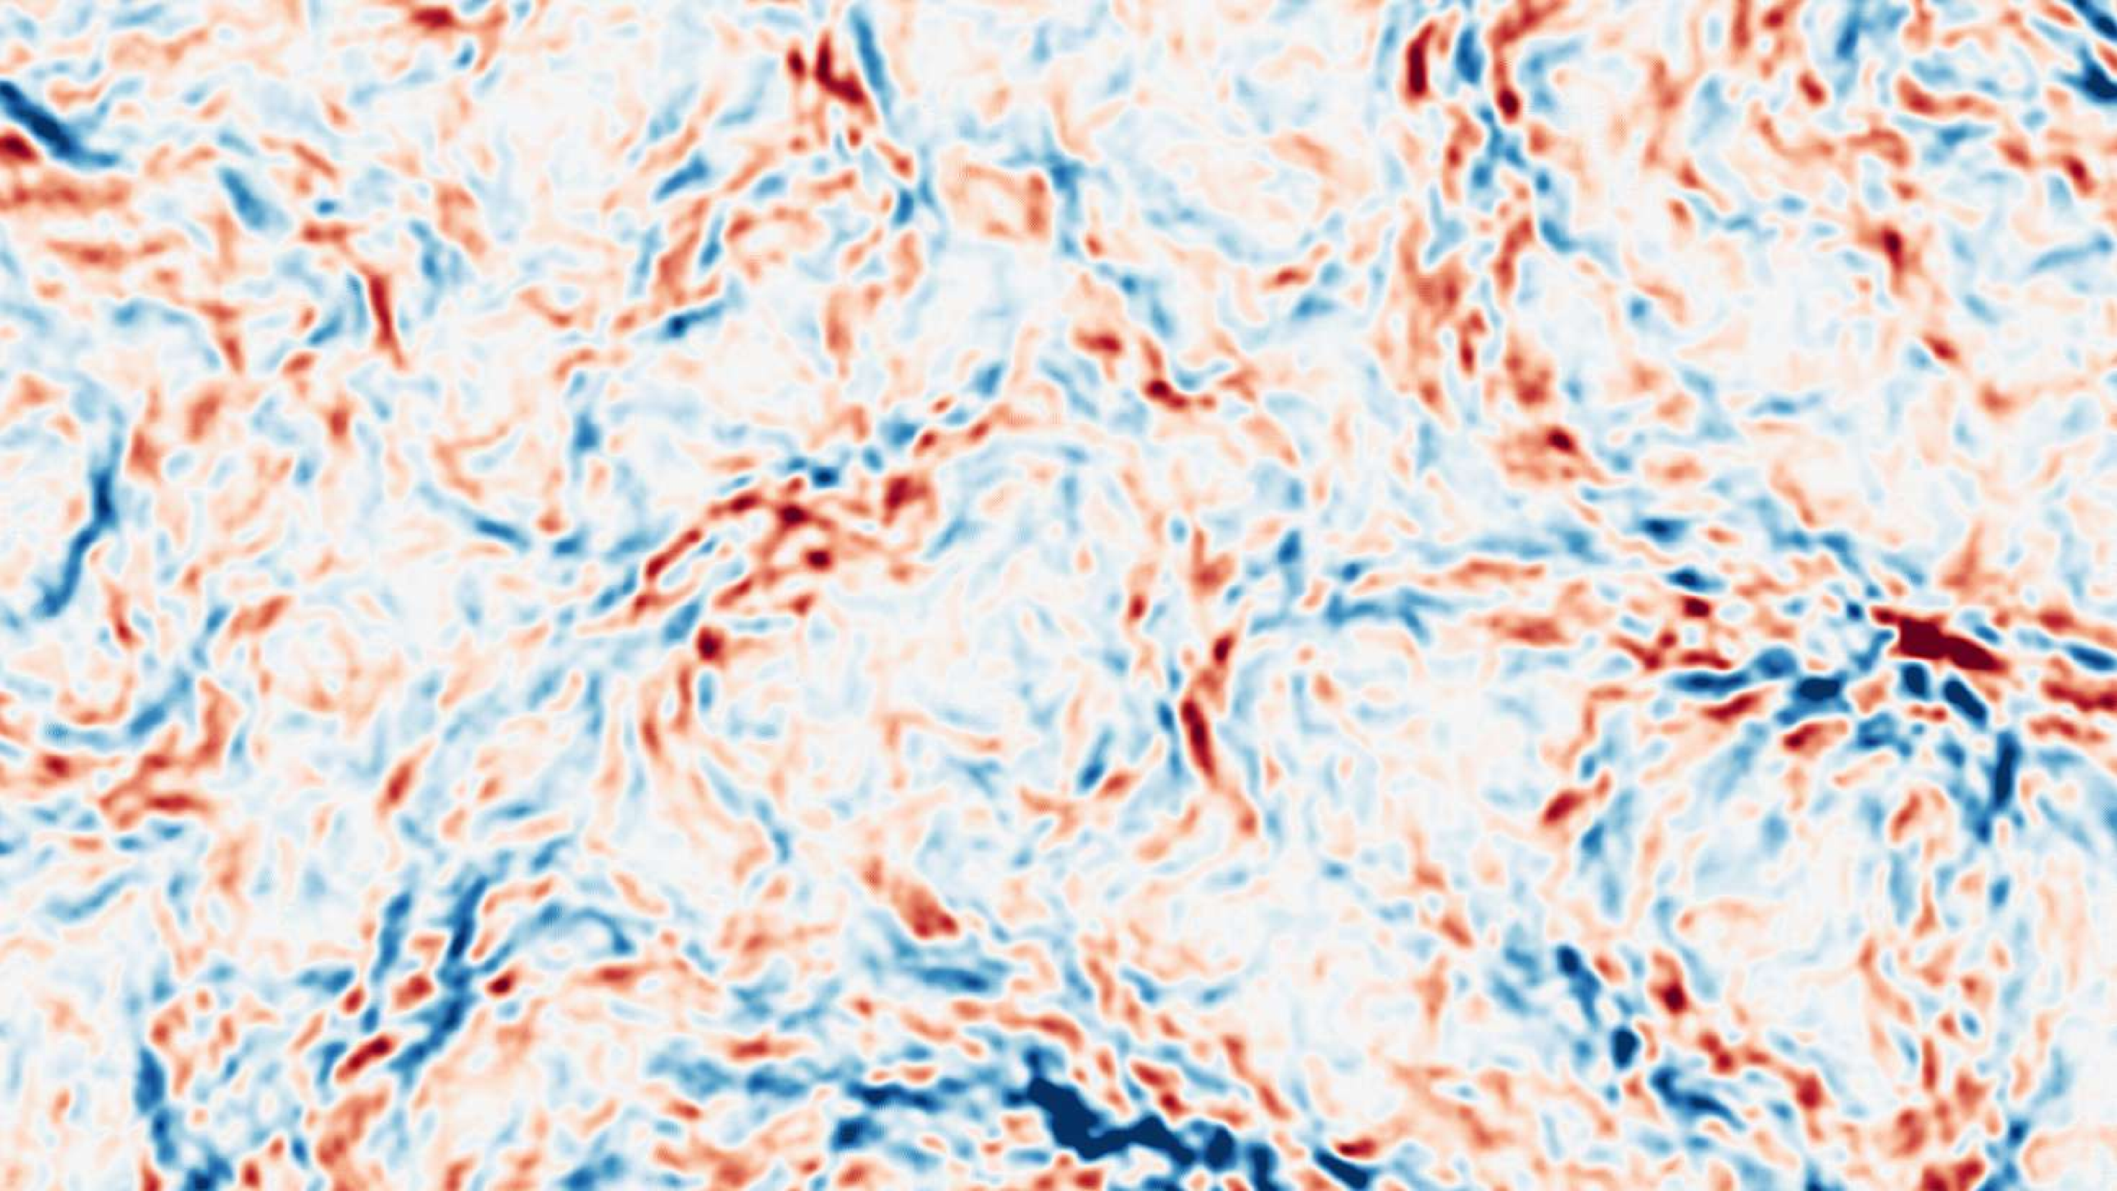

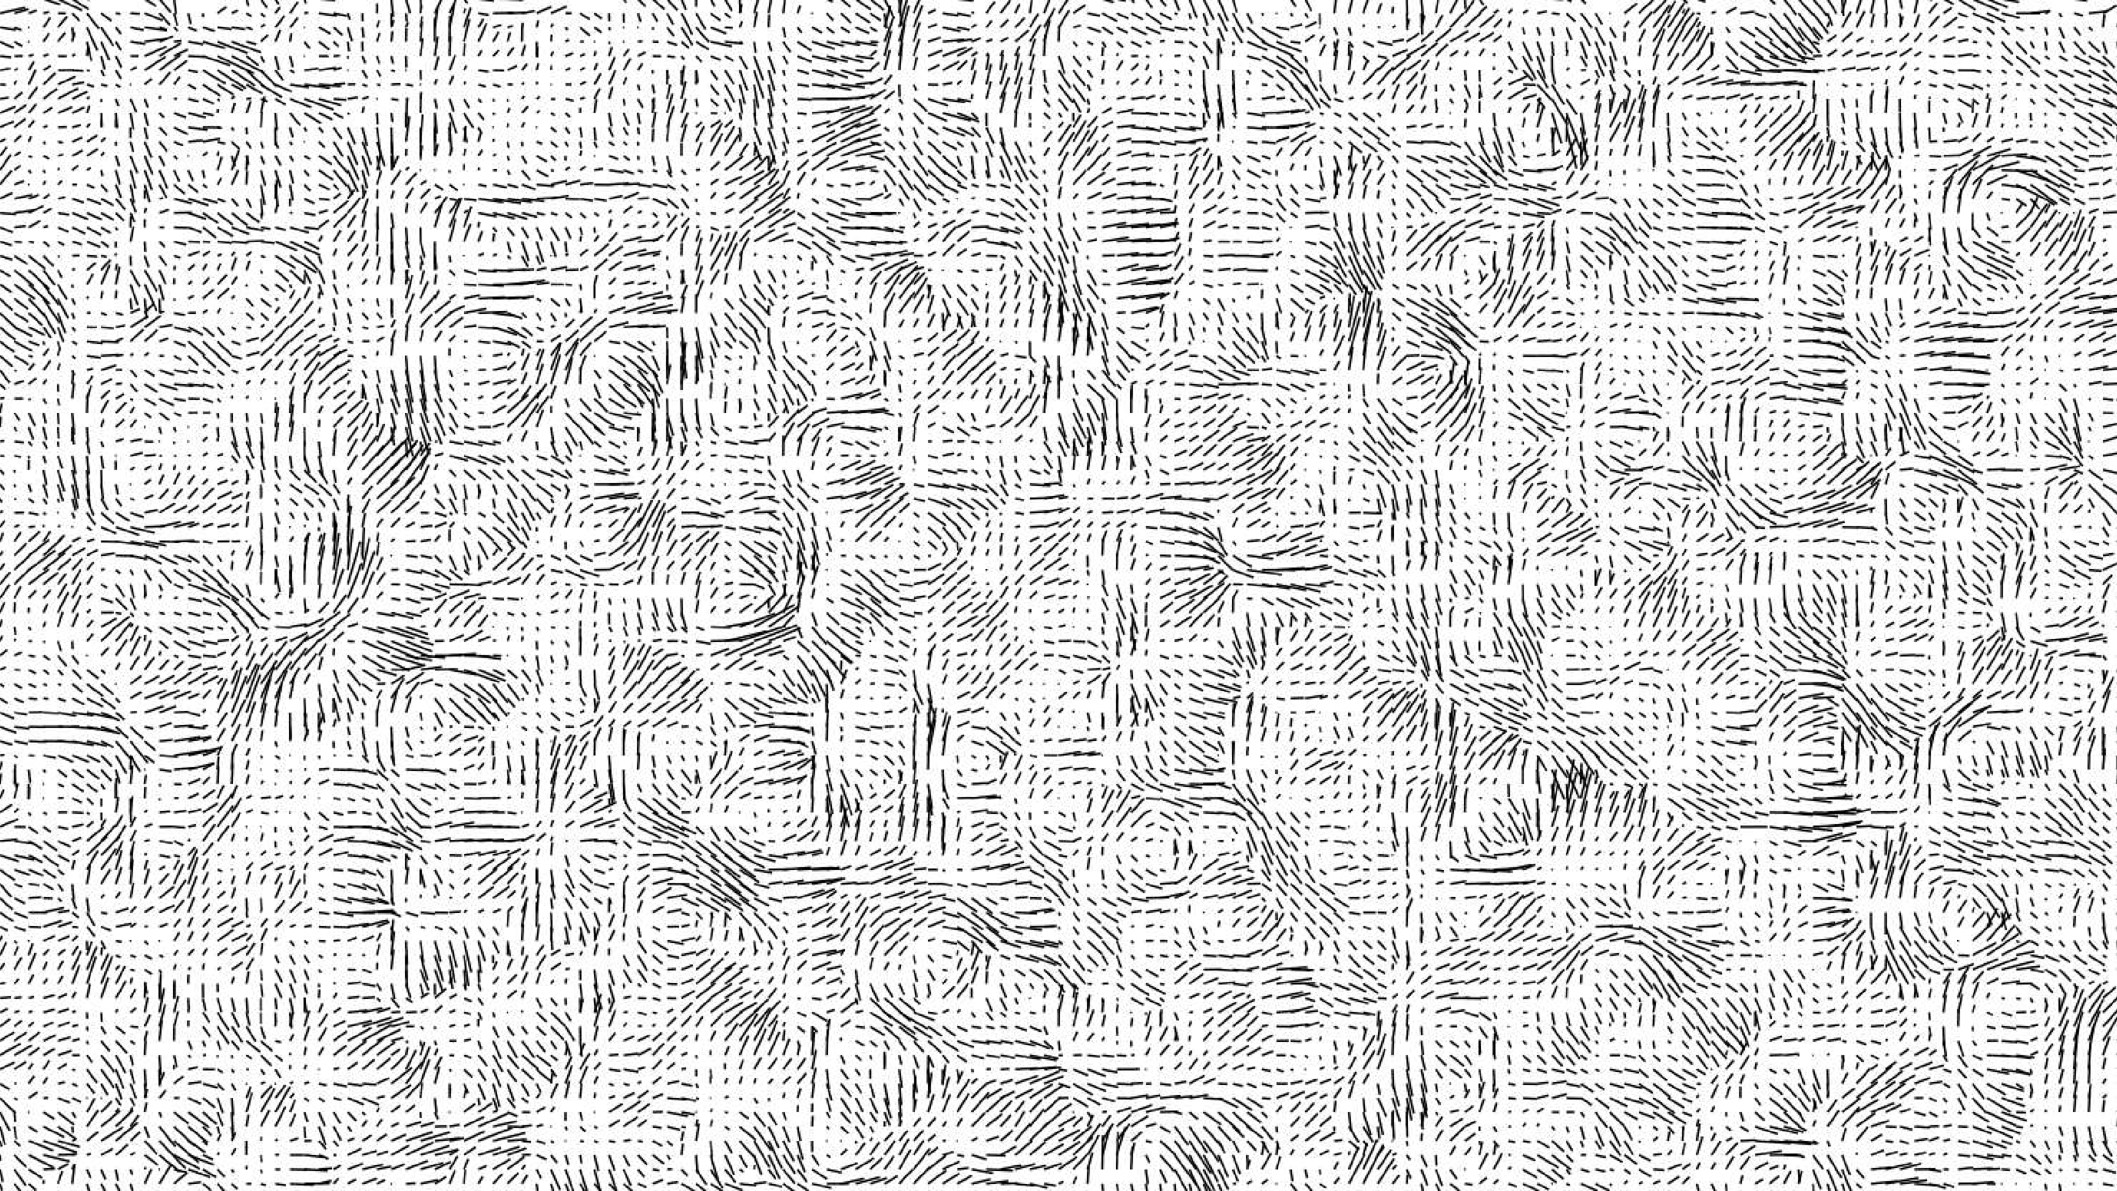

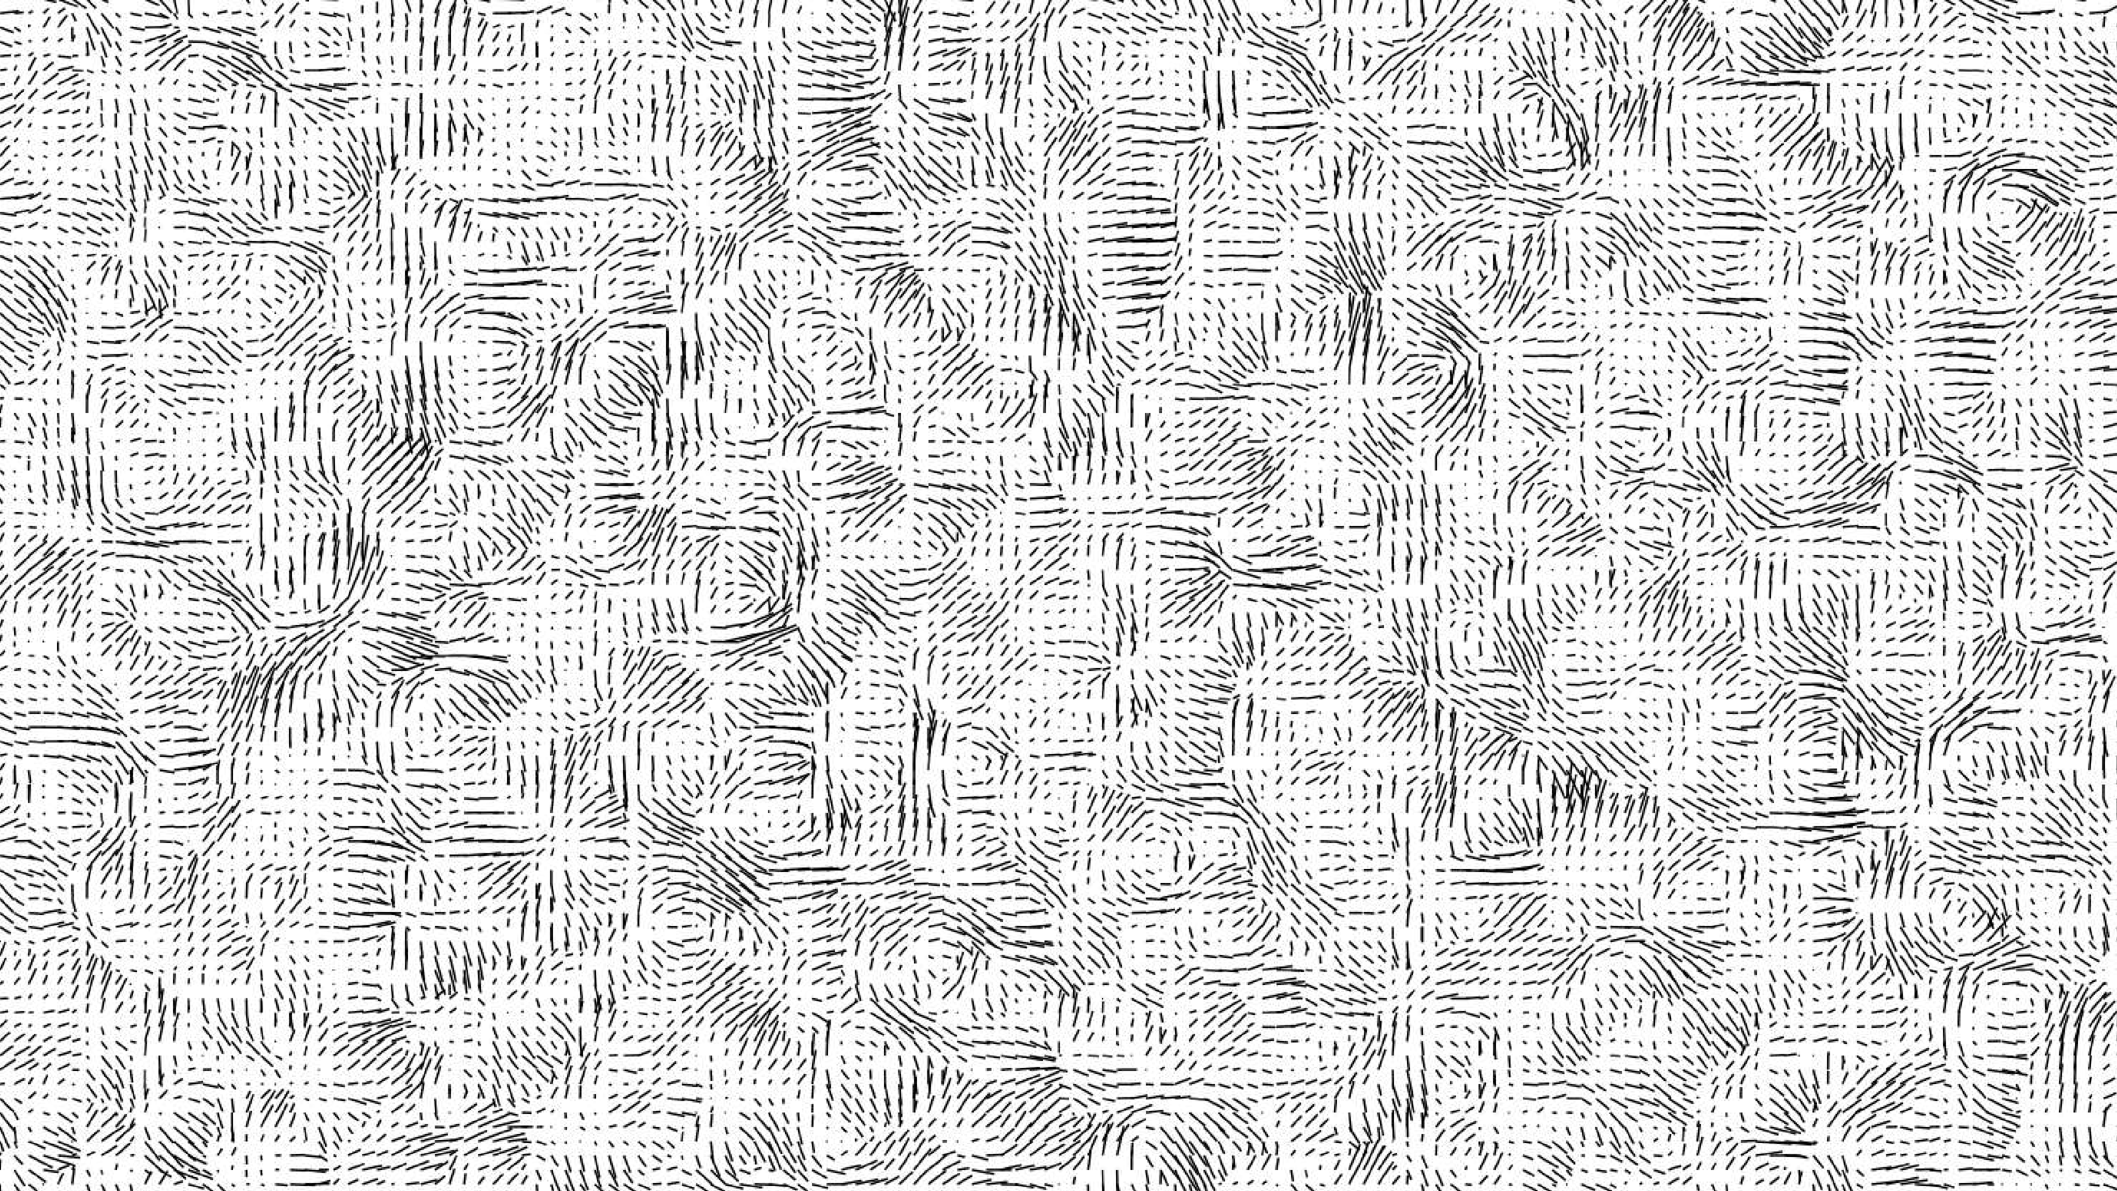

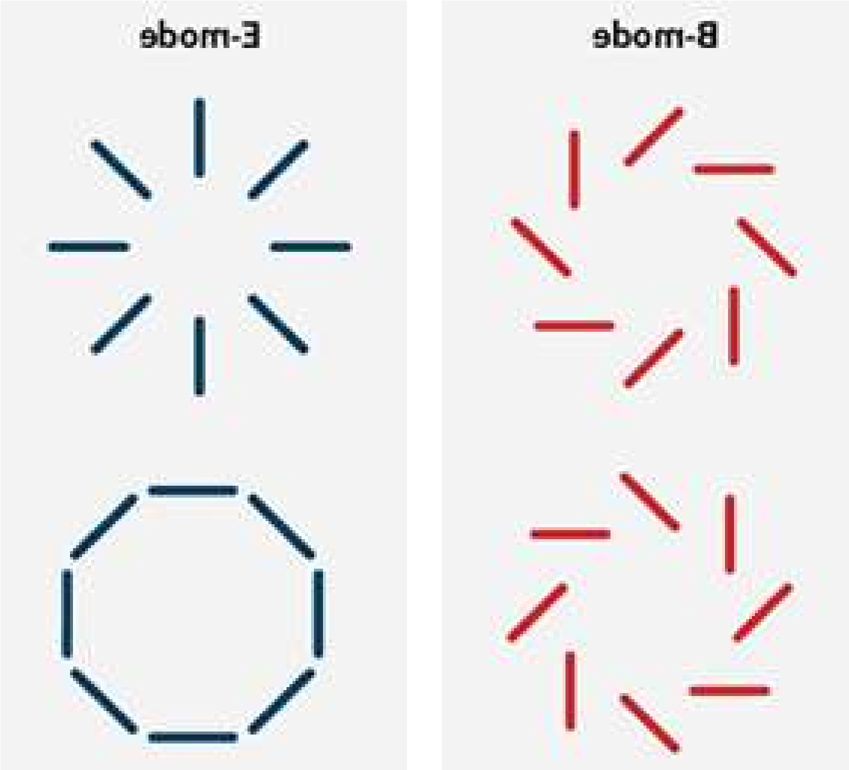

Sheared (lensed) E mode

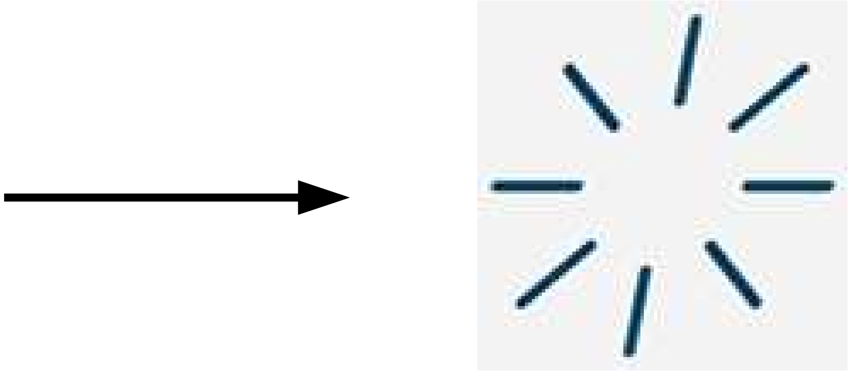

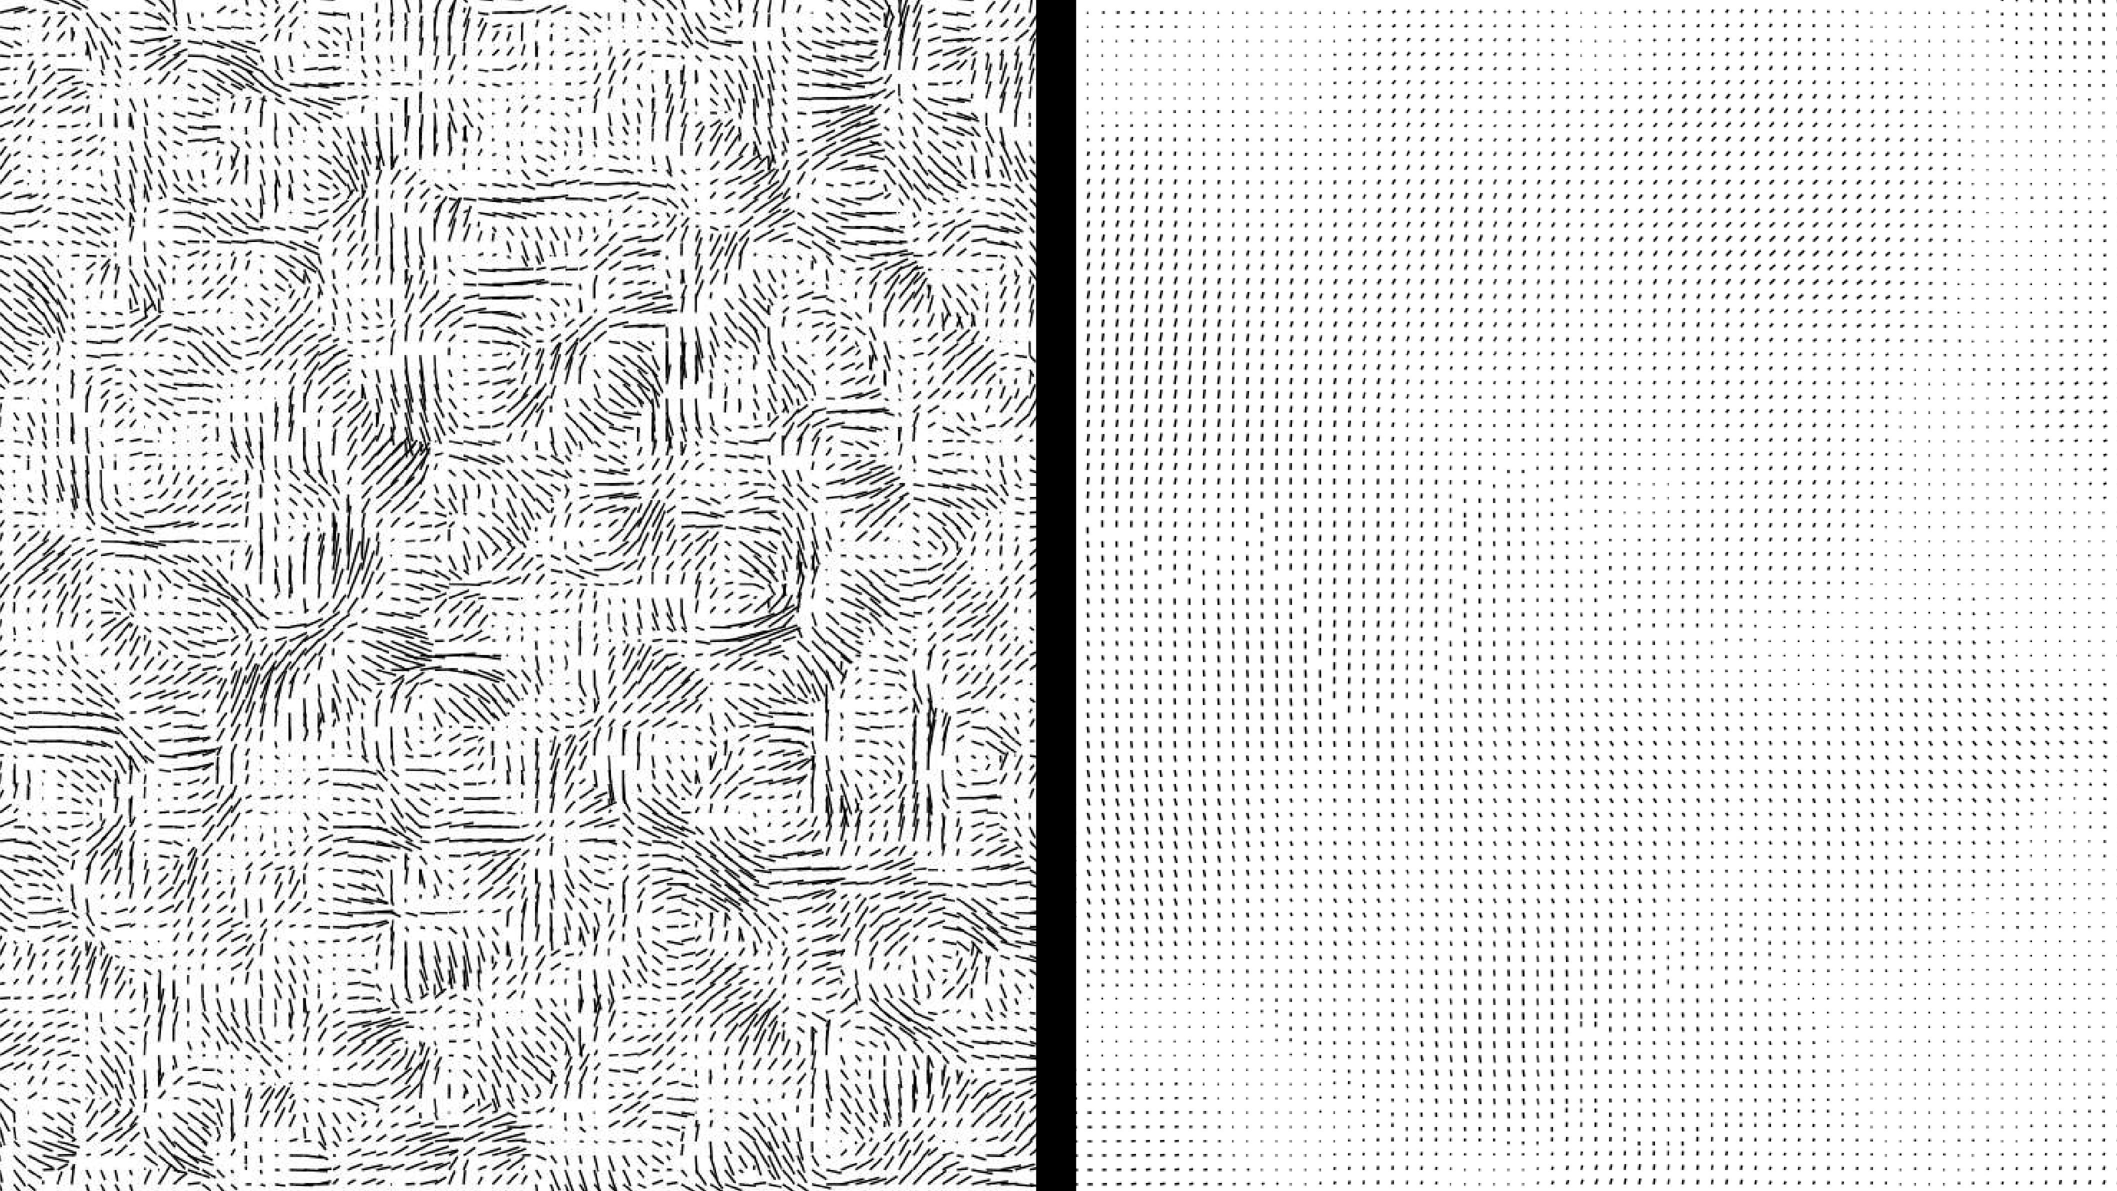

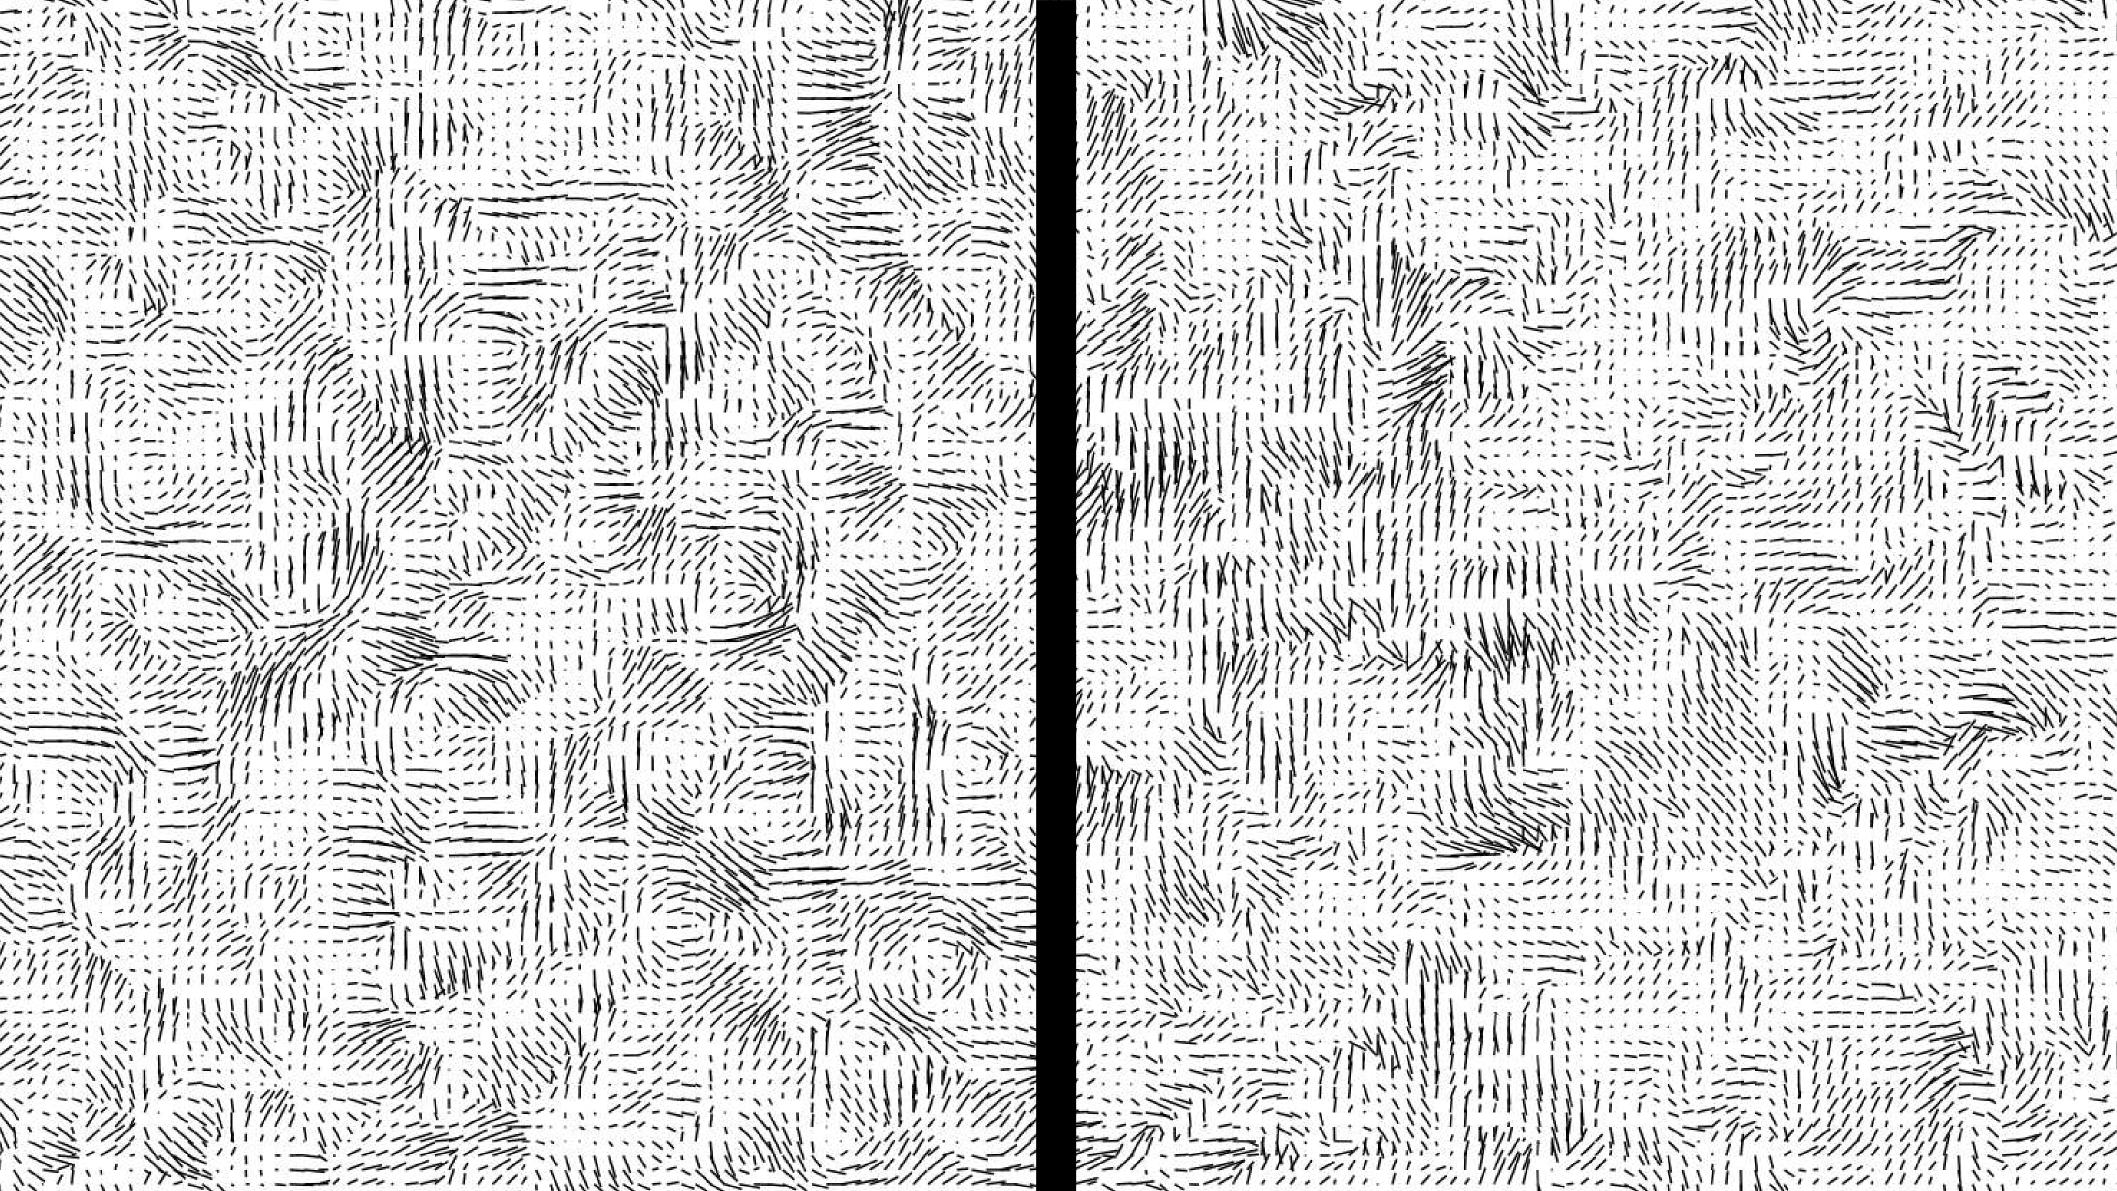

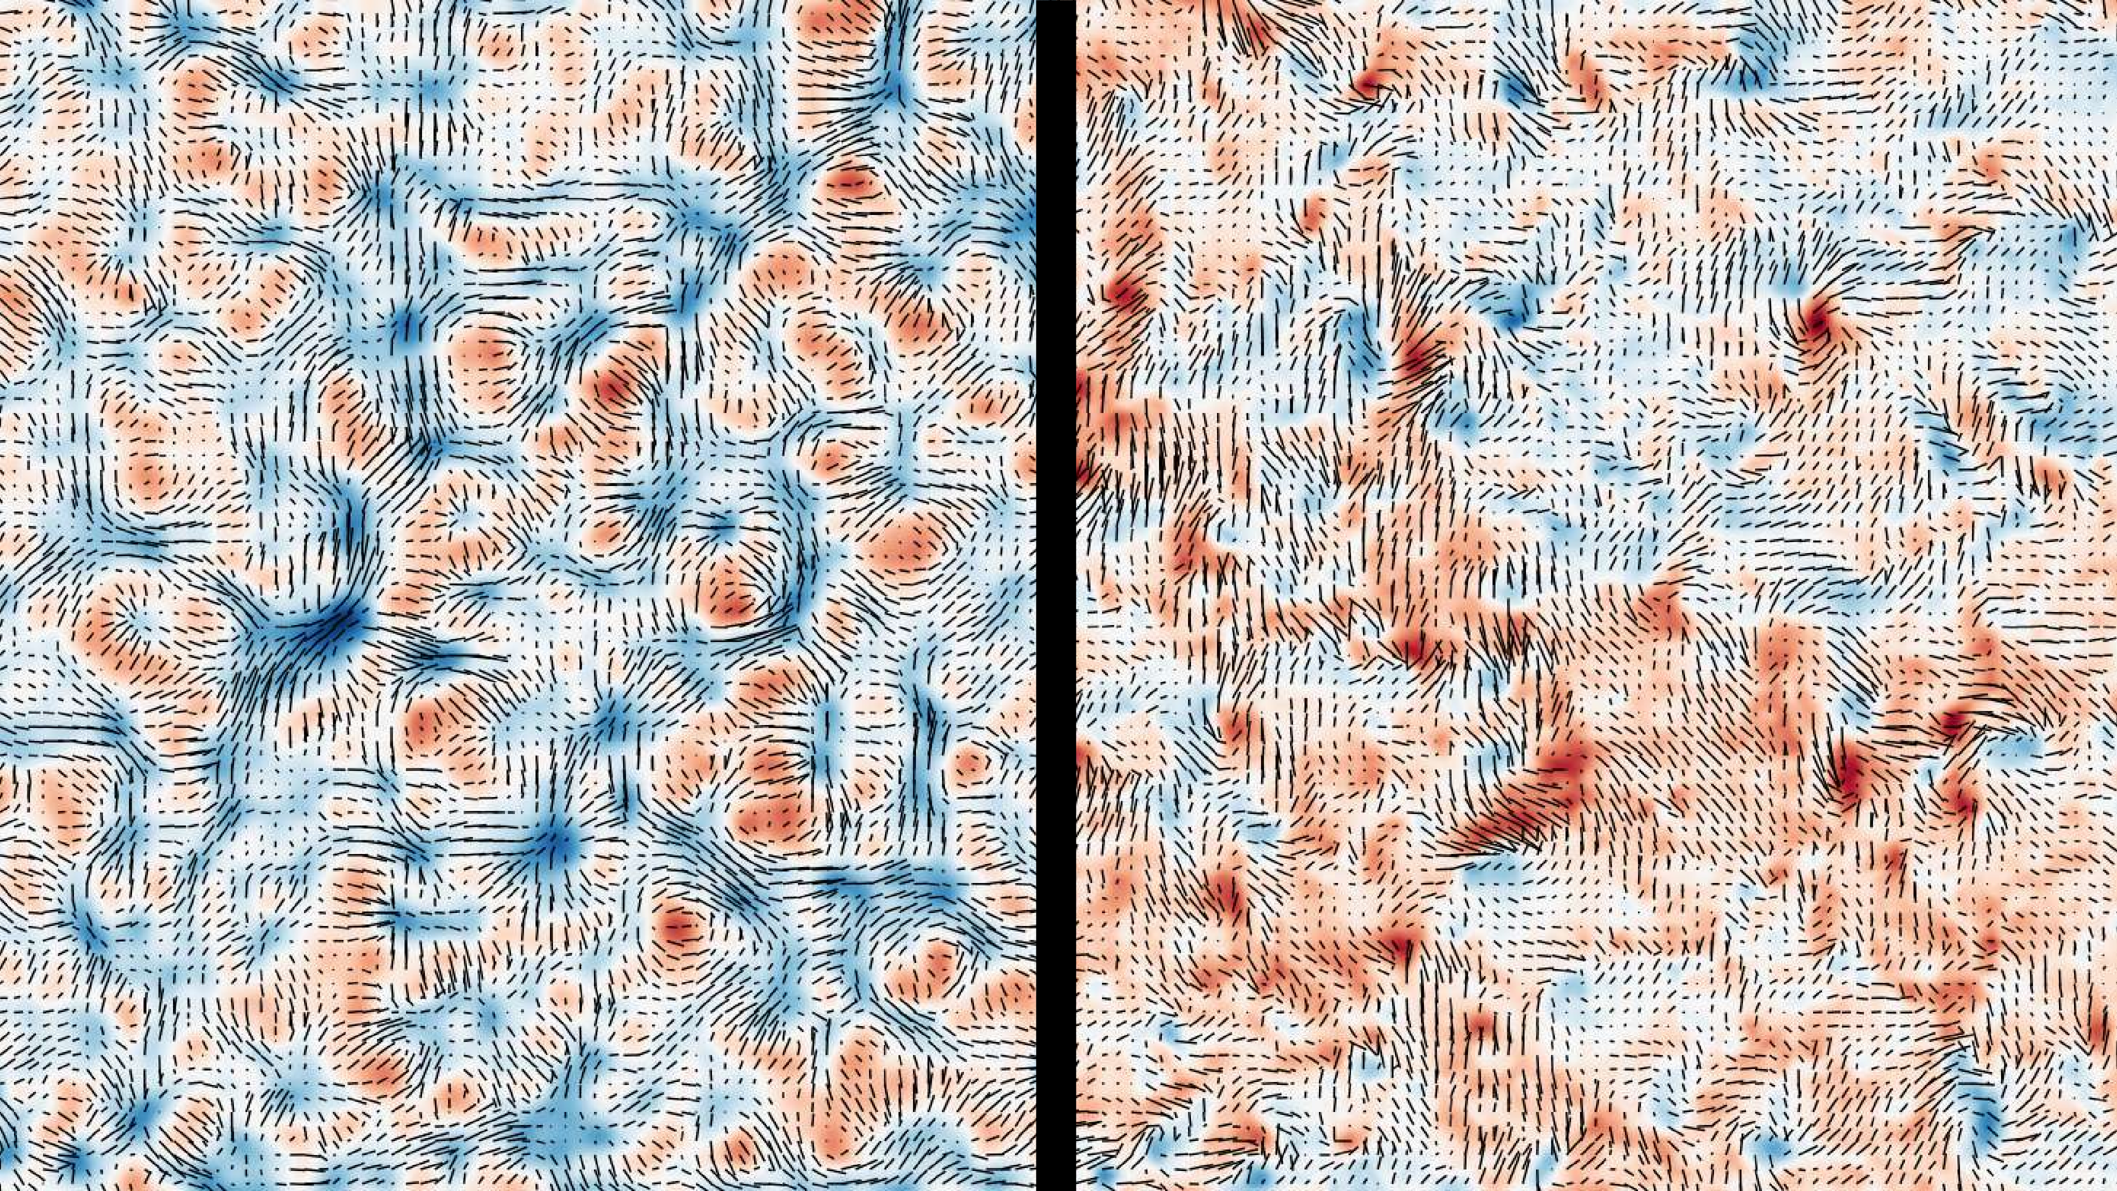

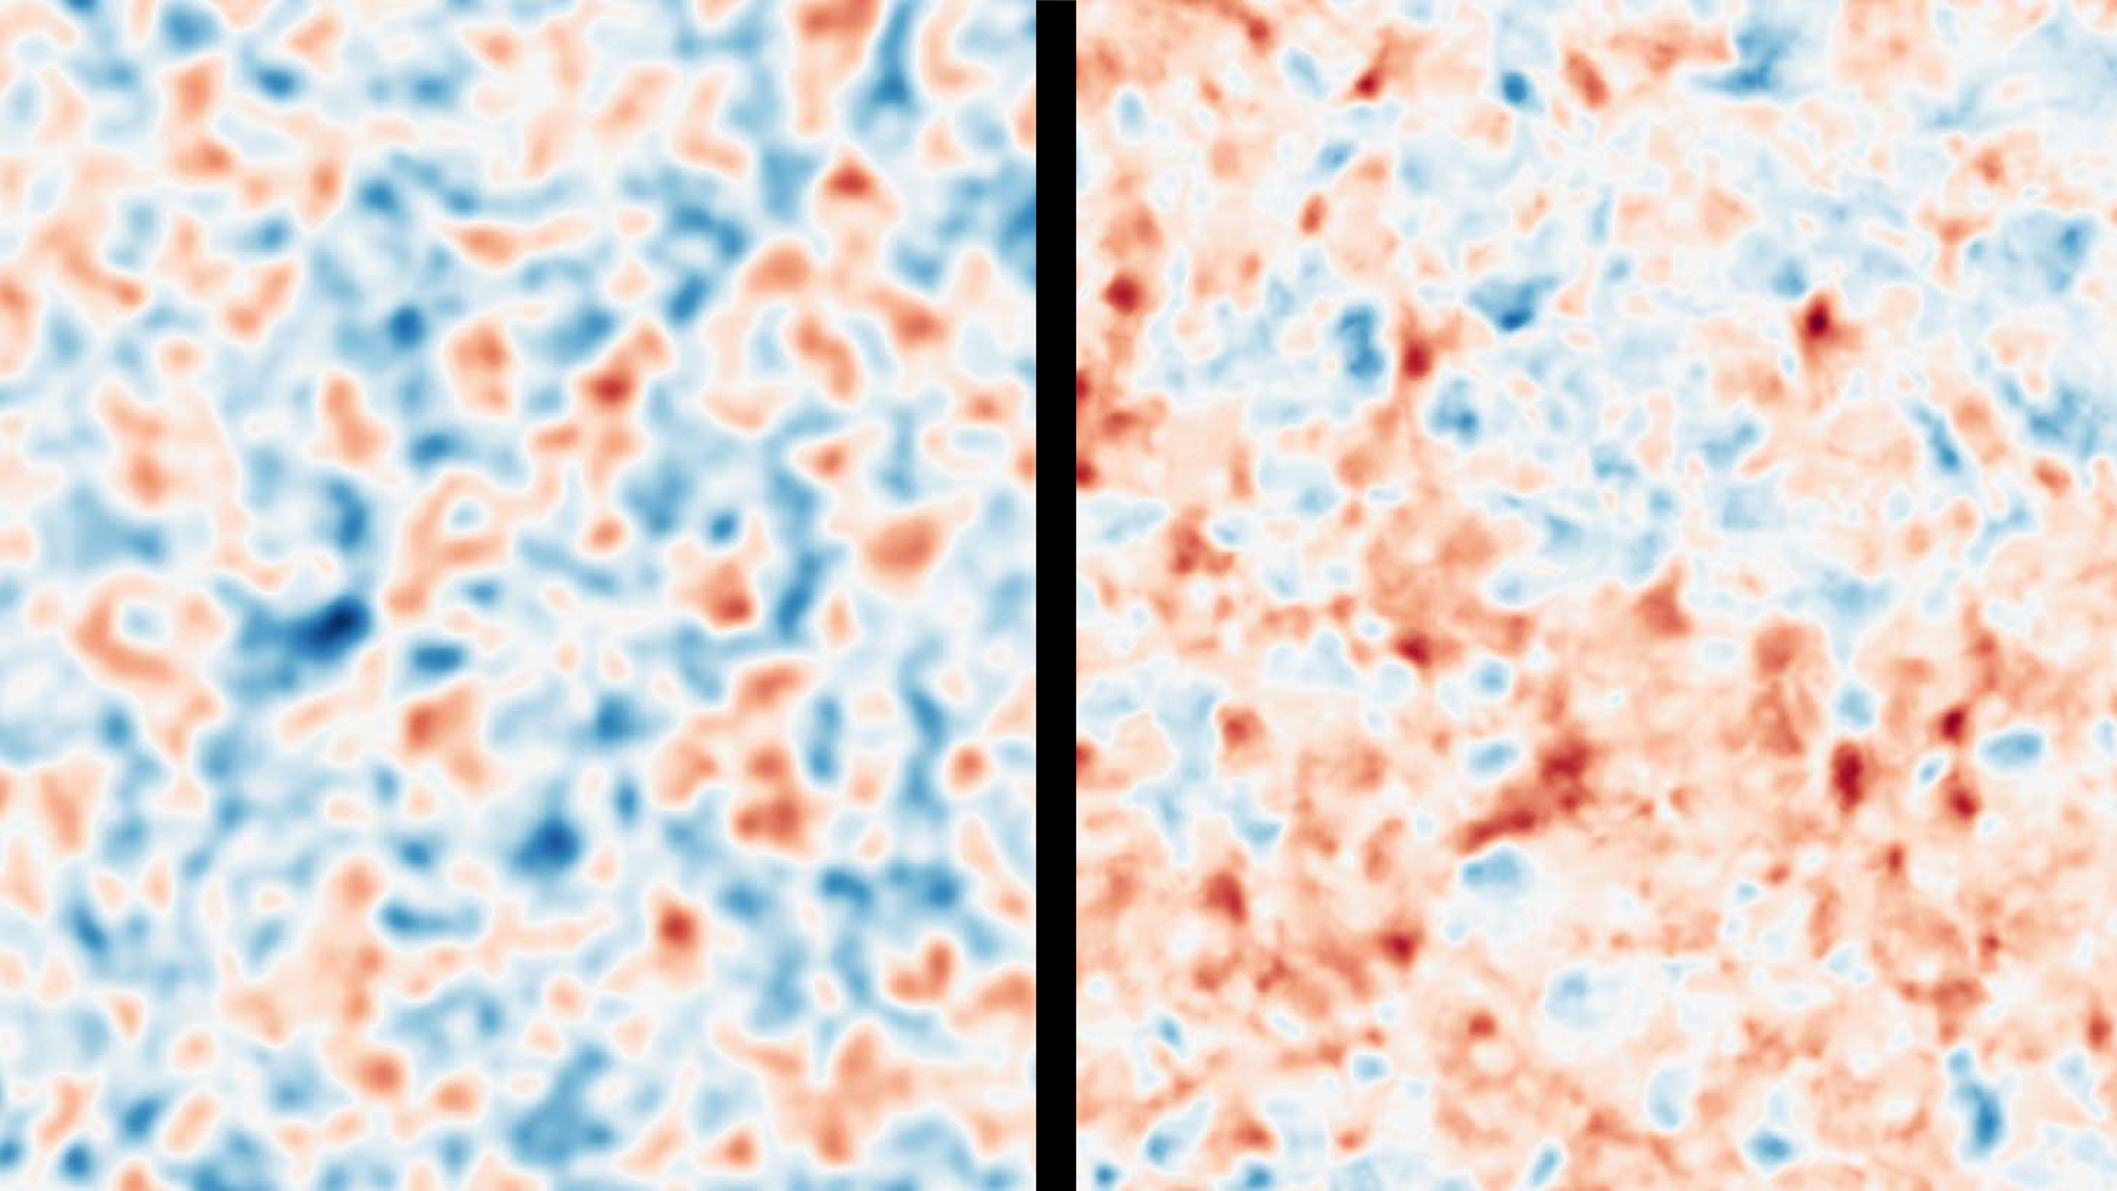

$T, E,$ and B $\qquad\qquad$ ϕ

Choi et al. 2019

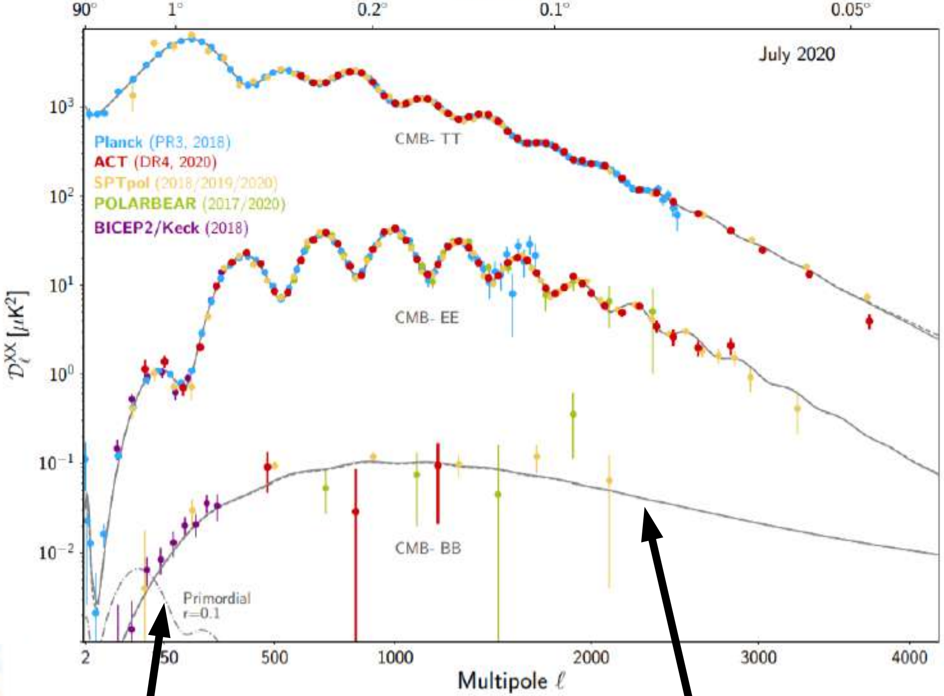

Wu et al. 2019

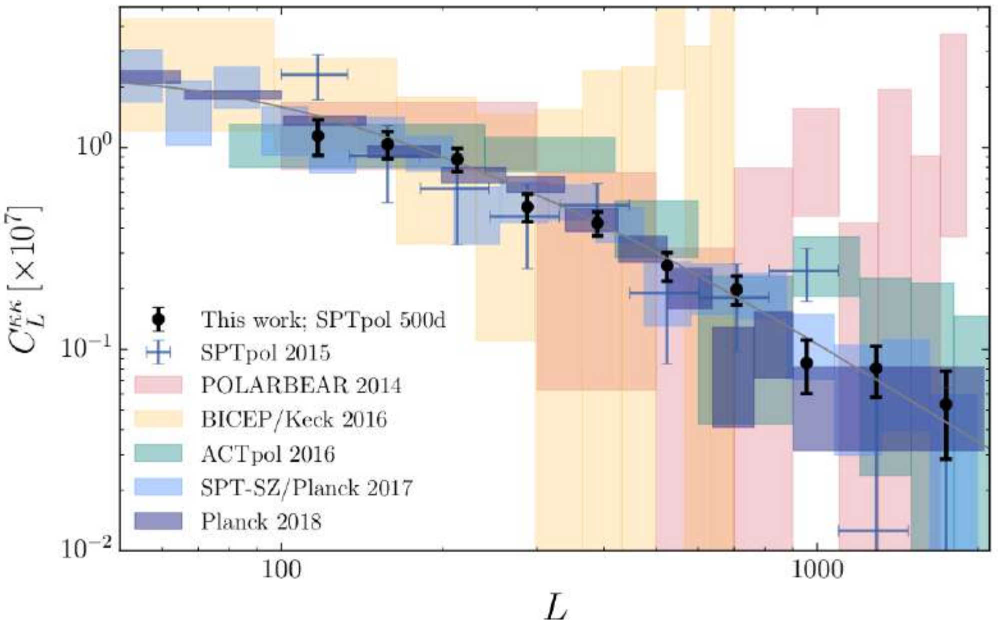

Lensing B modes

Primordial B modes

Isotropic Gaussian random fields:

$$\langle f(\vec{\ell}) f(\vec{\ell}+\vec{L})\rangle = \mathbb{C}_f(\ell)\delta(\vec{L})$$

$f \equiv (T, E, B)$
CMB "fields"

$\mathbb{C}_f(\ell)$ CMB 2D power-spectrum

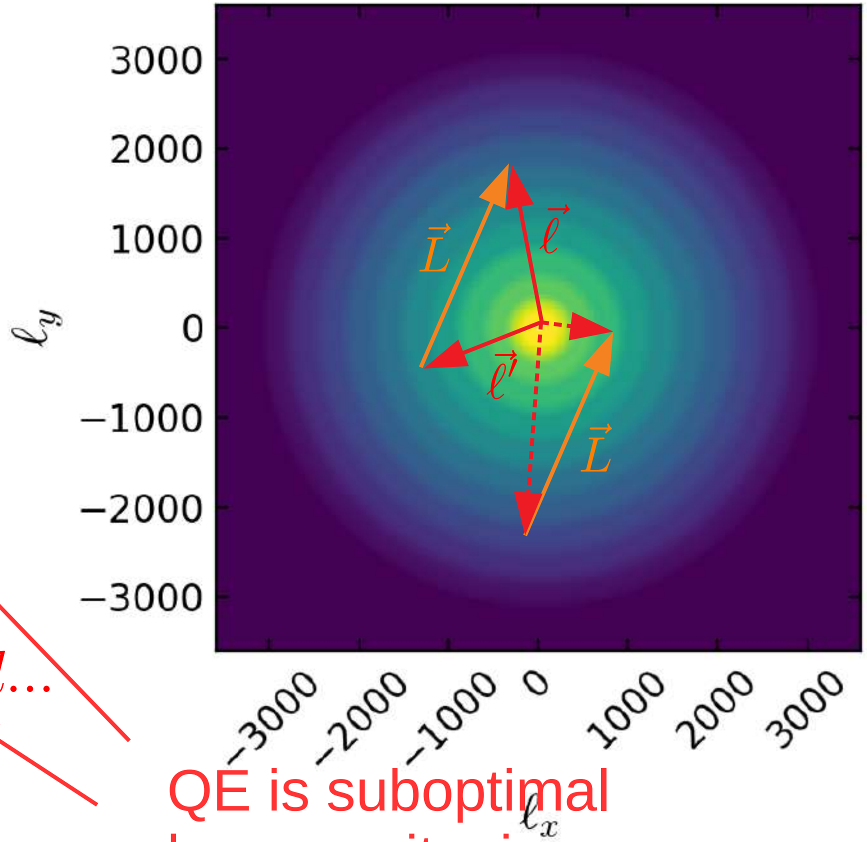

Lensed fields:

$$\langle \tilde{f}(\vec{\ell}) \tilde{f}(\vec{\ell}+\vec{L})\rangle \sim \phi(\vec{L}) \quad + \phi^2 + \ldots$$

Quadratic estimate (QE):

$$\hat{\phi}_{\rm QE}(\vec{L}) \sim \sum_{\vec{\ell}} d(\vec{\ell}) d(\vec{\ell}+\vec{L}) \quad \sum ddd\ldots$$

The data

$$\langle C_\ell^{\hat{\phi}_{\rm QE}}\rangle = A_\ell\, C_\ell^{\phi} + N_\ell^0 + N_\ell^1 + N_\ell^{3/2} + \ldots$$

Noise biases and normalization

QE is suboptimal because it misses this information

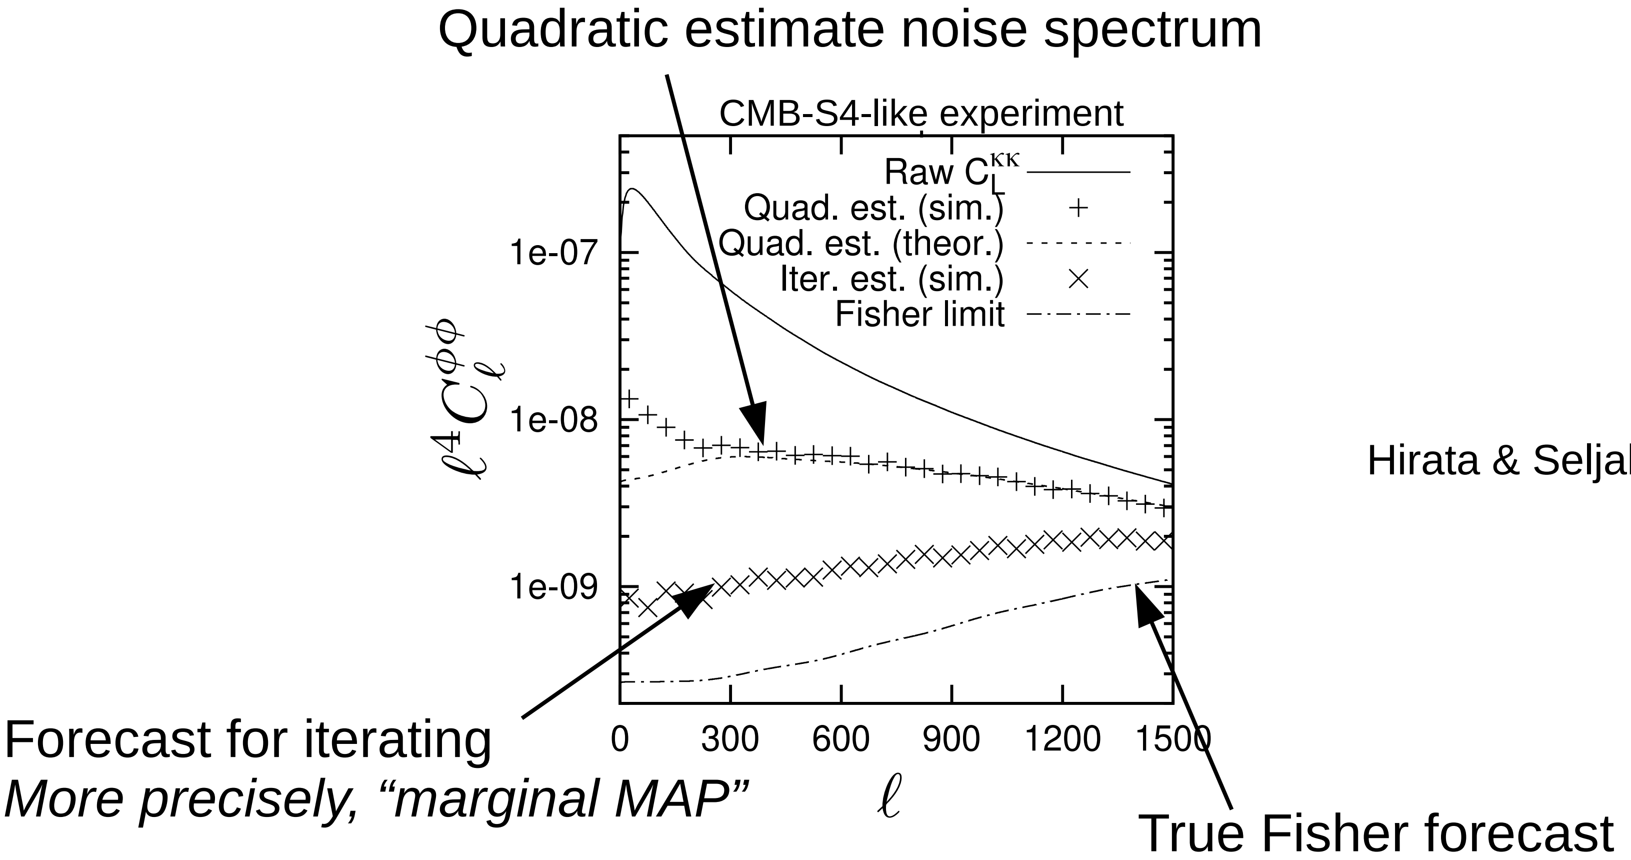

Hirata & Seljak (2003)

- These are forecasts or highly simplified analyses
- **Almost 20 years later, we are finally asking: How do we do this to real data?**

Towards optimality…

- **DeepCMB** (Caldeira et al. 2018)
  - Achieves noise levels comparable to the iterative-forecast
  - Challenges in extracting cosmological parameters
- **Gradient inversion** (Horowitz et al. 2018, Hadzhiyska et al. 2018)
  - Simple, but only optimal in the asymptotic limit of small scales
- **Optimal filtering** (Mirmelstein et al. 2019)
  - A way to more optimally filter a QE ϕ map before taking its power spectrum
  - May be useful mainly in the short term
- **Bayesian methods**
  - Guaranteed to be optimal, but computationally hard

Bayesian Lensing

Notation: $x^2/\mathbb{C} \equiv x^\dagger \mathbb{C}^{-1} x$

Cosmological parameters or theory spectra

Data model:

$$d = \mathbb{L}(\phi) f + n$$

Priors:

$$f \sim \mathrm{Gaussian}(0, \mathbb{C}_f(\theta)) \qquad n \sim \mathrm{Gaussian}(0, \mathbb{C}_n)$$

$$\phi \sim \mathrm{Gaussian}(0, \mathbb{C}_\phi(\theta)) \qquad \theta \sim \mathrm{Uniform}$$

“Joint” posterior (MM,Anderes,Wandelt 2018, 2020):

$$\mathcal{P}\big(f, \phi, \theta \,|\, d\big) = \frac{\exp\left\{-\dfrac{(d - \mathbb{L}(\phi) f)^2}{2\,\mathbb{C}_n}\right\}}{\det \mathbb{C}_n^{1/2}} \frac{\exp\left\{-\dfrac{f^2}{2\,\mathbb{C}_f(\theta)}\right\}}{\det \mathbb{C}_f(\theta)^{1/2}} \frac{\exp\left\{-\dfrac{\phi^2}{2\,\mathbb{C}_\phi(\theta)}\right\}}{\det \mathbb{C}_\phi(\theta)^{1/2}}$$

“Marginal” posterior (Hirata&Seljak 2003; Carron&Lewis 2018):

$$\mathcal{P}\big(\phi, \theta \,|\, d\big) = \frac{\exp\left\{-\dfrac{d^2}{2\,\mathbb{C}_d(\phi, \theta)}\right\}}{\det\, \mathbb{C}_d(\phi, \theta)^{1/2}} \frac{\exp\left\{-\dfrac{\phi^2}{2\,\mathbb{C}_\phi(\theta)}\right\}}{\det\, \mathbb{C}_\phi(\theta)^{1/2}}$$

$$\text{where } \; \mathbb{C}_d(\phi, \theta) \equiv \mathbb{L}(\phi) \mathbb{C}_f(\theta) \mathbb{L}(\phi)^\dagger + \mathbb{C}_n$$

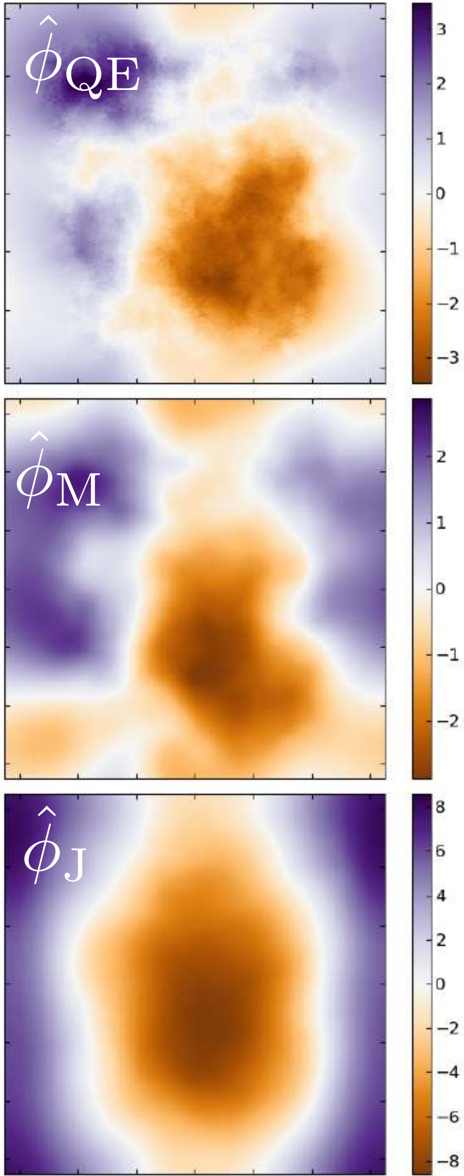

$$\hat{\phi}_{\rm M} \equiv \underset{\phi}{\rm argmax}\, \mathcal{P}(\phi, \theta_{\rm fid} \,|\, d)$$ *marginal MAP*

(1st step of this is ~QE)

$$\hat{f}_{\rm J}, \hat{\phi}_{\rm J} \equiv \underset{f,\phi}{\rm argmax}\, \mathcal{P}(f, \phi, \theta_{\rm fid} \,|\, d)$$ *joint MAP*

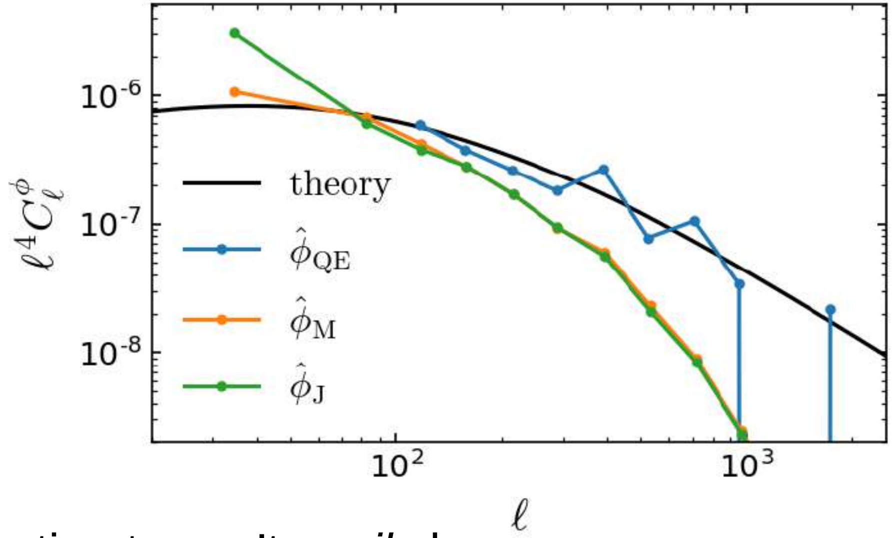

- MAP estimate can't *easily* be normalized to theory like the QE
- Even if done via MC, is cosmology-dependent, no QE *tricks* available
- Not the right path for power-spectrum or parameter estimation

“Joint” posterior (MM,Anderes,Wandelt 2018, 2020):

$$\mathcal{P}\big(\tilde{f},\phi,\theta\,|\,d\big) = \frac{\exp\left\{-\frac{(d - \quad \tilde{f} \quad)^2}{2\,\mathbb{C}_n}\right\}}{\det \mathbb{C}_n^{1/2}} \frac{\exp\left\{-\frac{(\mathbb{L}(\phi)^{-1}\tilde{f})^2}{2\,\mathbb{C}_f(\theta)}\right\}}{\det \mathbb{C}_f(\theta)^{1/2}} \frac{\exp\left\{-\frac{\phi^2}{2\,\mathbb{C}_\phi(\theta)}\right\}}{\det \mathbb{C}_\phi(\theta)^{1/2}} \frac{1}{\det\ \mathbb{L}(\phi)}$$

Instead of maximizing, *marginalize:*

$$\mathcal{P}(\theta\,|\,d) = \int \mathrm{d}f\,\mathrm{d}\phi\;\mathcal{P}(f,\phi,\theta\,|\,d)$$

This is guaranteed to be “optimal,” ie represent all the information that we can extract.

This ~million dimensional marginalization done with Hamiltonian Monte Carlo.

Traditional lensing:

$$\tilde{f}(x) = f(x + \nabla\phi) \approx f(x) + \nabla f(x)\nabla\phi(x) + ...$$

LenseFlow:

$$f_t(x) = f(x + t\nabla\phi) \qquad \frac{df_t(x)}{dt} = \underbrace{\nabla\phi(x) \cdot [\mathbb{1} + t\nabla\nabla\phi(x)]^{-1} \cdot \nabla}_{\mathbb{V}_t} f_t(x)$$

$$f_{t=1} = [\mathbb{1} + \Delta t\, \mathbb{V}_{t_n}] \; ... \; [\mathbb{1} + \Delta t \mathbb{V}_{t_1}] \; f_{t=0}$$

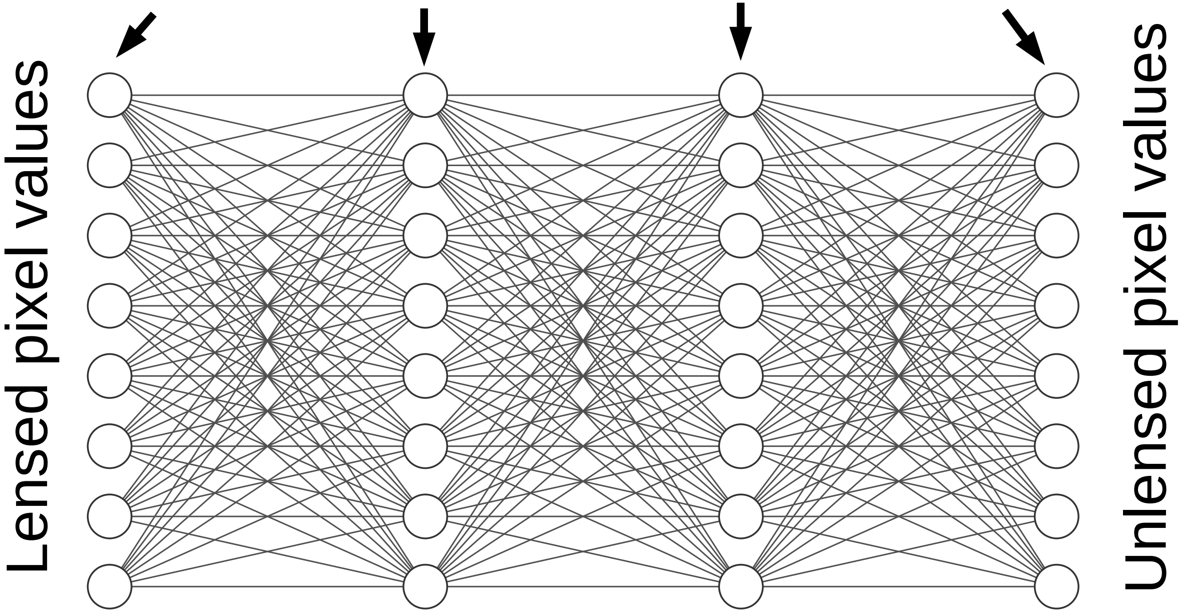

See also:
neural ODE
in machine
learning

Upcoming South Pole Telescope Analysis

With Cail Daley, Jody Ti-Lin Chou, SPT collaboration

Deepest 100deg^2 polarization measurements to-date at the angular scales most relevant for lensing.

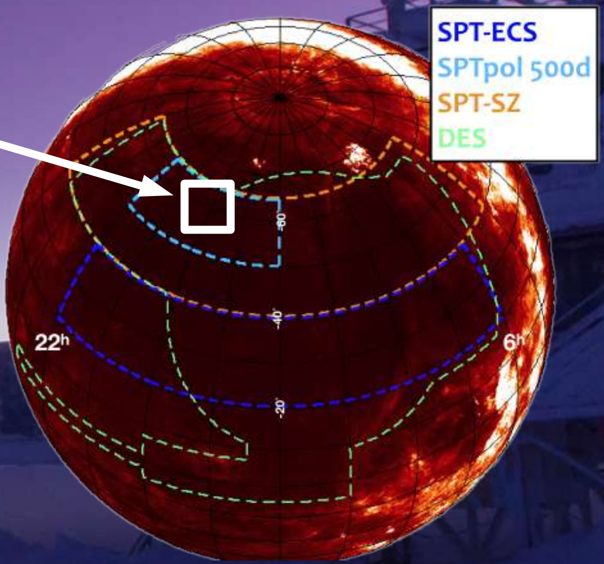

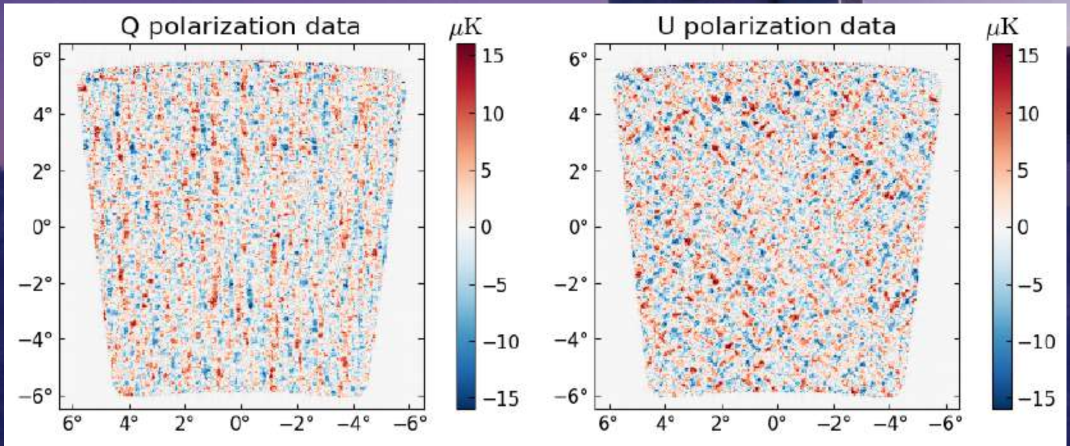

Data model for this analysis

Pixel mask

Pol. flattening

Fourier mask

Lensing

Unlensed CMB

Noise

$$d = \mathbb{M}_{\mathrm{f}} \, \mathbb{M}_{\mathrm{p}} \, \mathbb{R}_{\mathrm{obs}} \times \left[P_{\mathrm{cal}} \, \mathbb{R}(\psi_{\mathrm{pol}}) \, \mathbb{T} \, \mathbb{B} \, \mathbb{L}(\phi) \, f + \epsilon_{\mathrm{Q}} t_{\mathrm{Q}} + \epsilon_{\mathrm{U}} t_{\mathrm{U}} \right] + n$$

Beam & TF

TP leakage templates
and coefficients

Pol. calibration

Global pol. rotation

Trace of various quantities throughout the samples:

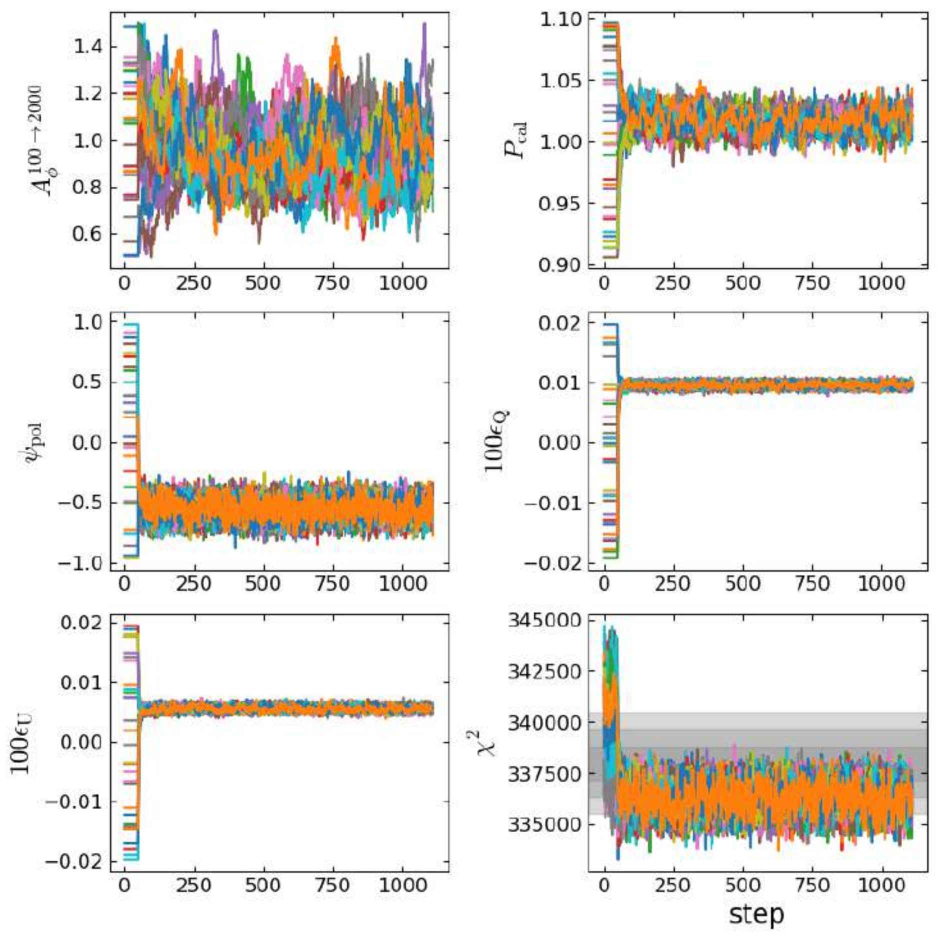

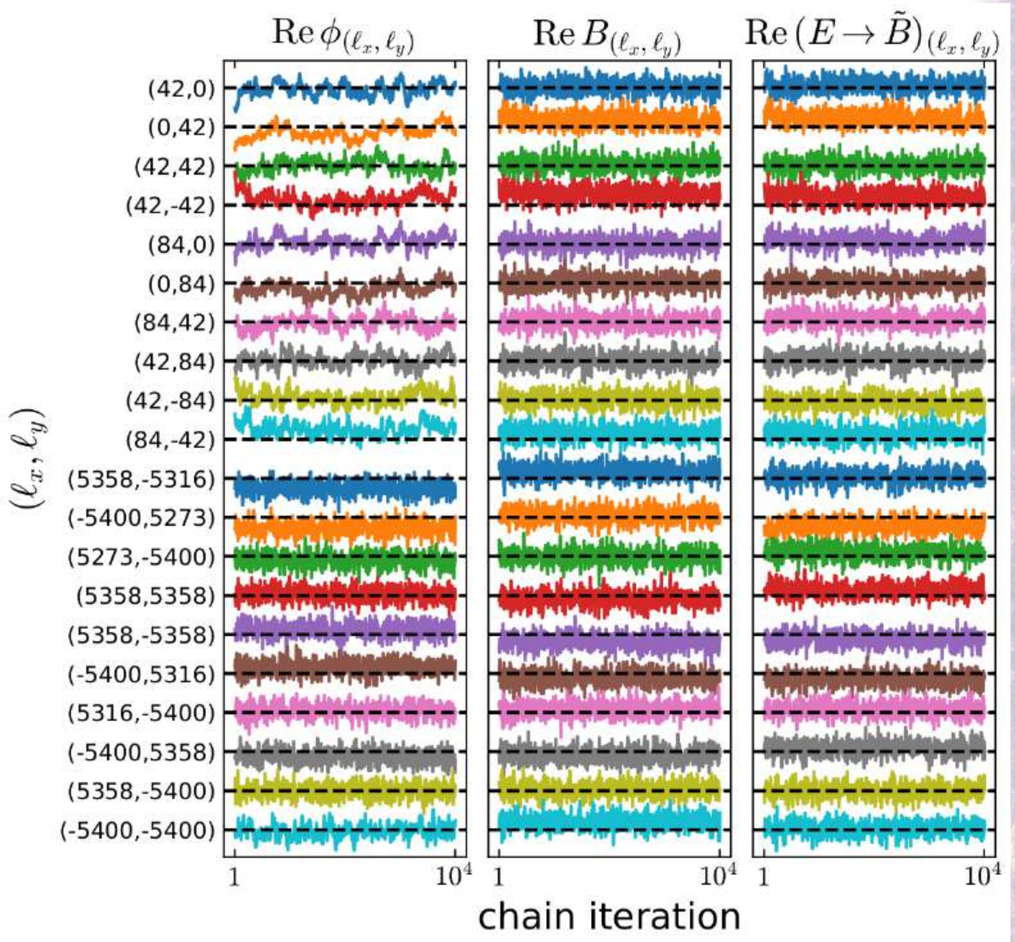

$$+\,3\times260^2\approx200{,}000$$

(every pixel in Q,U,ϕ maps)

Systematics

Preliminary

TABLE 3
SYSTEMATIC UNCERTAINTIES

Type	$\Delta A_{\rm MV}$	$\Delta A_{\rm POL}$	$\Delta A_{\rm T}$
$\Delta A_{\rm beam}$	0.008	0.010	0.005
$\Delta A_{\rm cal}$	0.023	0.039	0.008
$\Delta A_{\rm T\to P}$	$\ll 0.001$	$\ll 0.001$	N/A
$\Delta A_{\rm pol.rot.}$	< 0.001	< 0.001	N/A
$\Delta A_{\rm fg}$	0.004	N/A	0.008
$\Delta A_{\rm tot}$	0.025	0.040	0.012

Previous SPT analysis
(Wu et al. 2019)

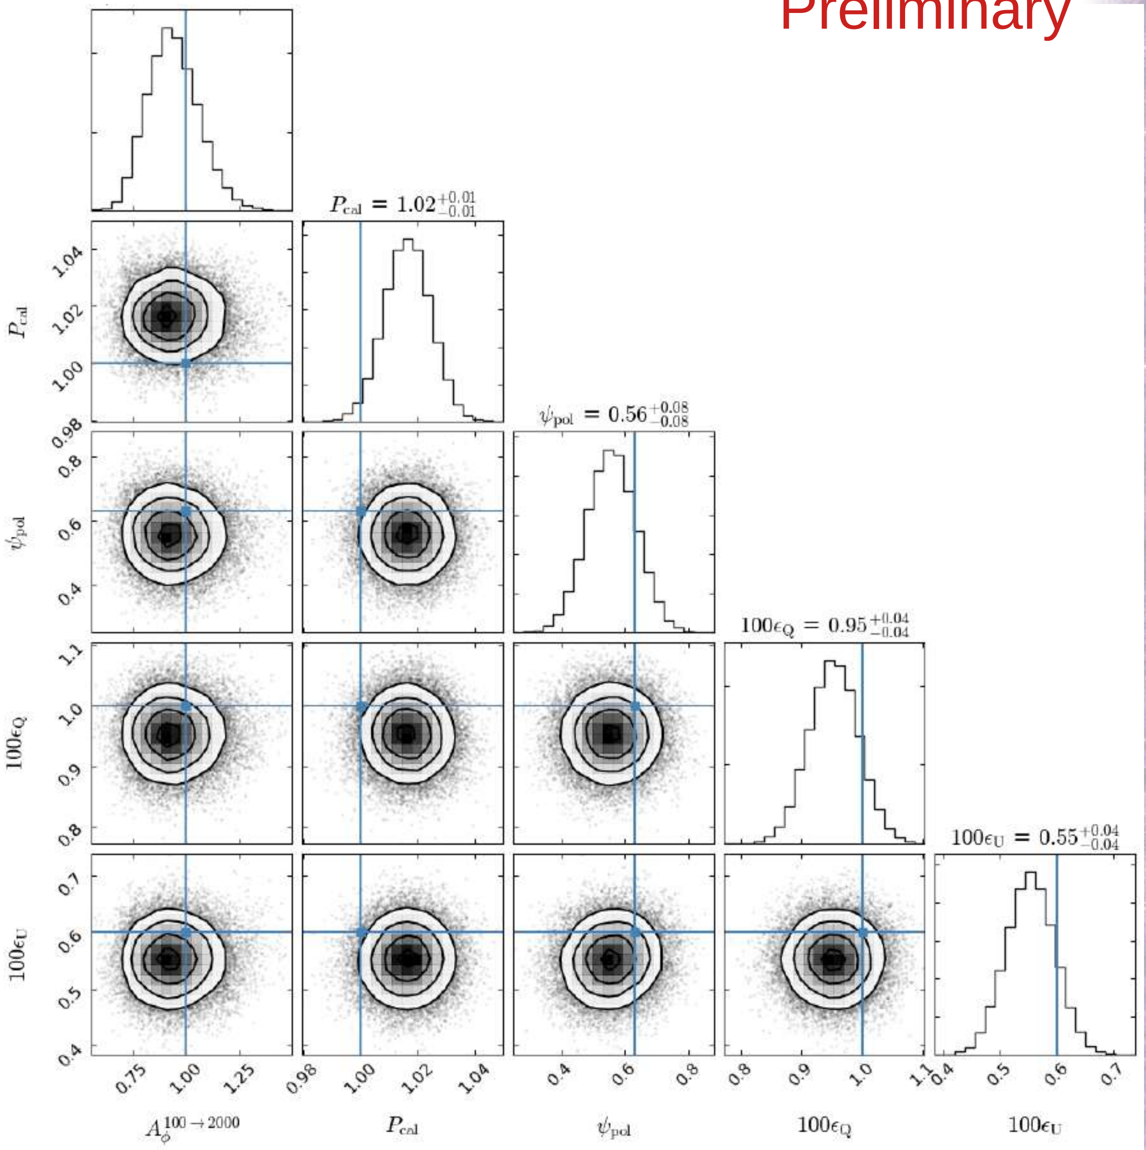

- Systematics are now a simple and self-consistent part of the analysis instead of a secondary ad-hoc thing
- As a bonus, we've reduced the impact of Pcal from ½σ to effectively zero.

Comparison to QE

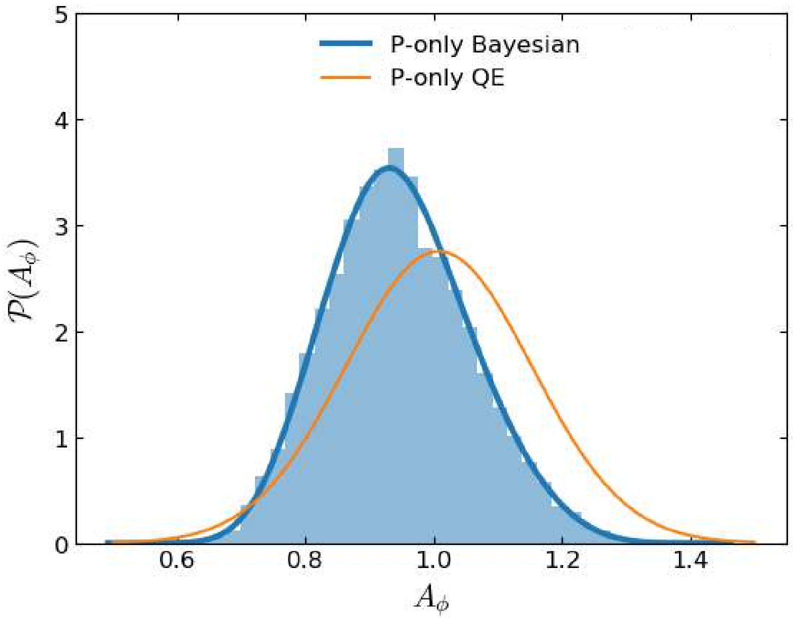


```
Bayesian:   0.9459(75) ± 0.1123(50)
QE:         1.01       ± 0.134
```

- 23% tighter error bars, in line with expectations from forecasts
- First time cosmological parameter extracted from optimal lensing reconstruction

Bayesian lensing paves the way for solving future problems in CMB analysis

Foreground contamination

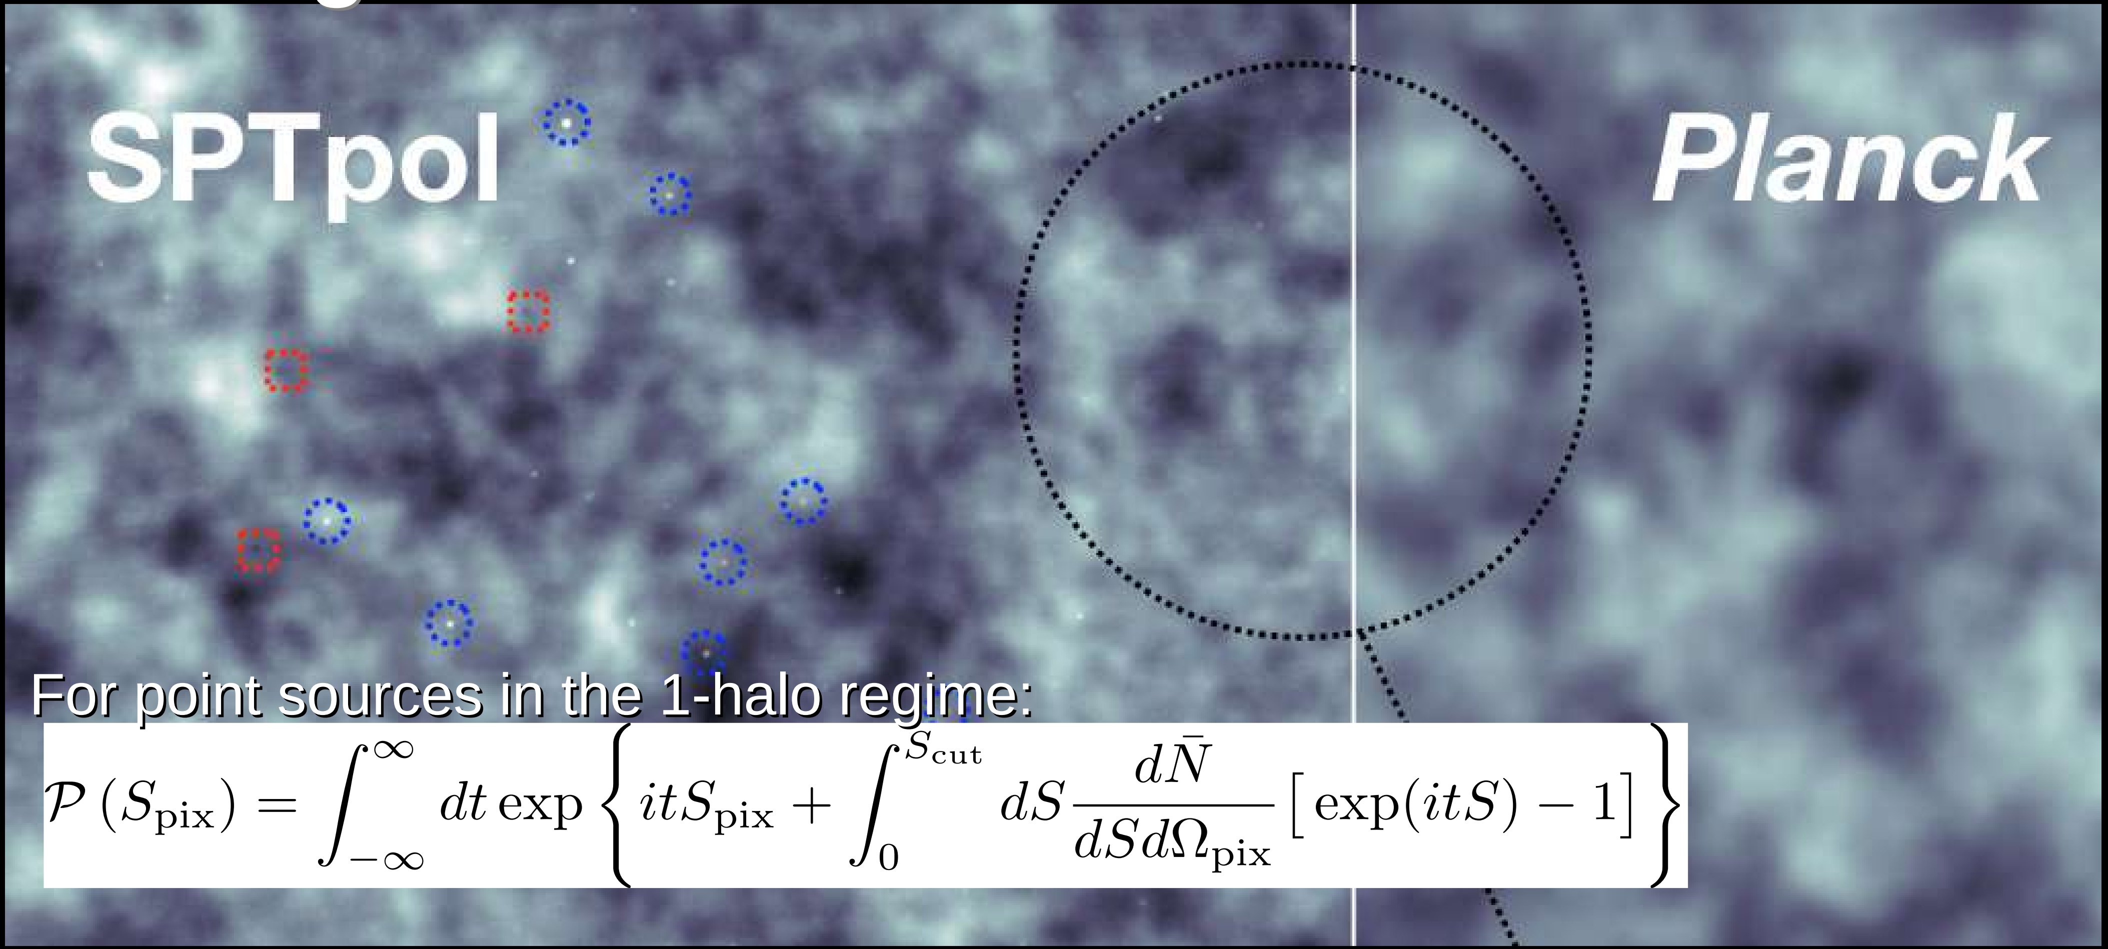

For point sources in the 1-halo regime:

$$\mathcal{P}\left(S_{\rm pix}\right) = \int_{-\infty}^{\infty} dt \exp\left\{ itS_{\rm pix} + \int_0^{S_{\rm cut}} dS \frac{d\bar{N}}{dS d\Omega_{\rm pix}} \left[\exp(itS) - 1\right] \right\}$$

work with Emmanuel Schaan

Machine learning models

Aylor et al. 2019

Generative neural network

$\mathrm{Gaussian}(0, \mathbb{I}_{64})$

Prior distribution

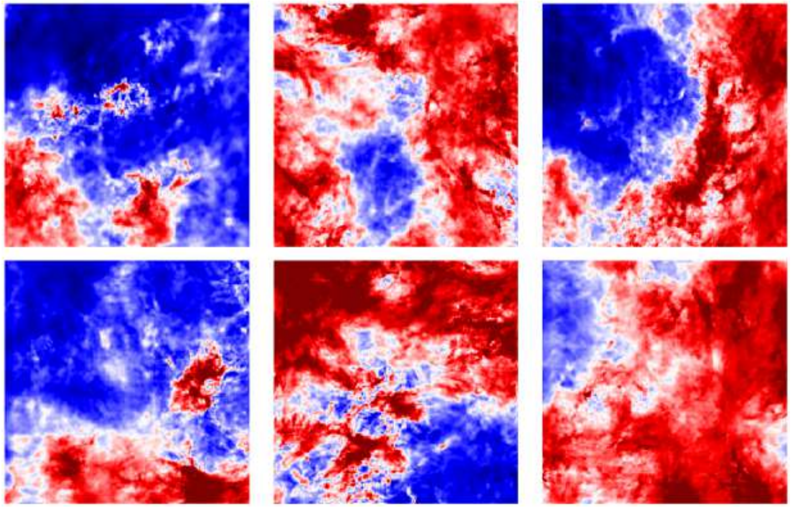

work with Ben Thorne, Lloyd Knox

Joint bandpower estimation MM, Seljak, in prep

$$\mathcal{P}(C_\ell^{\phi\phi}, C_\ell^{EE}, \,|\, d) = \int \mathrm{d}f \mathrm{d}\phi \, \mathcal{P}(f, \phi, C_\ell^{\phi\phi}, C_\ell^{EE} \,|\, d)$$

We have developed **approximations** to this integral. Using our ability to get the *exact* answer via sampling, we can **validate** these approximations given realistic data.

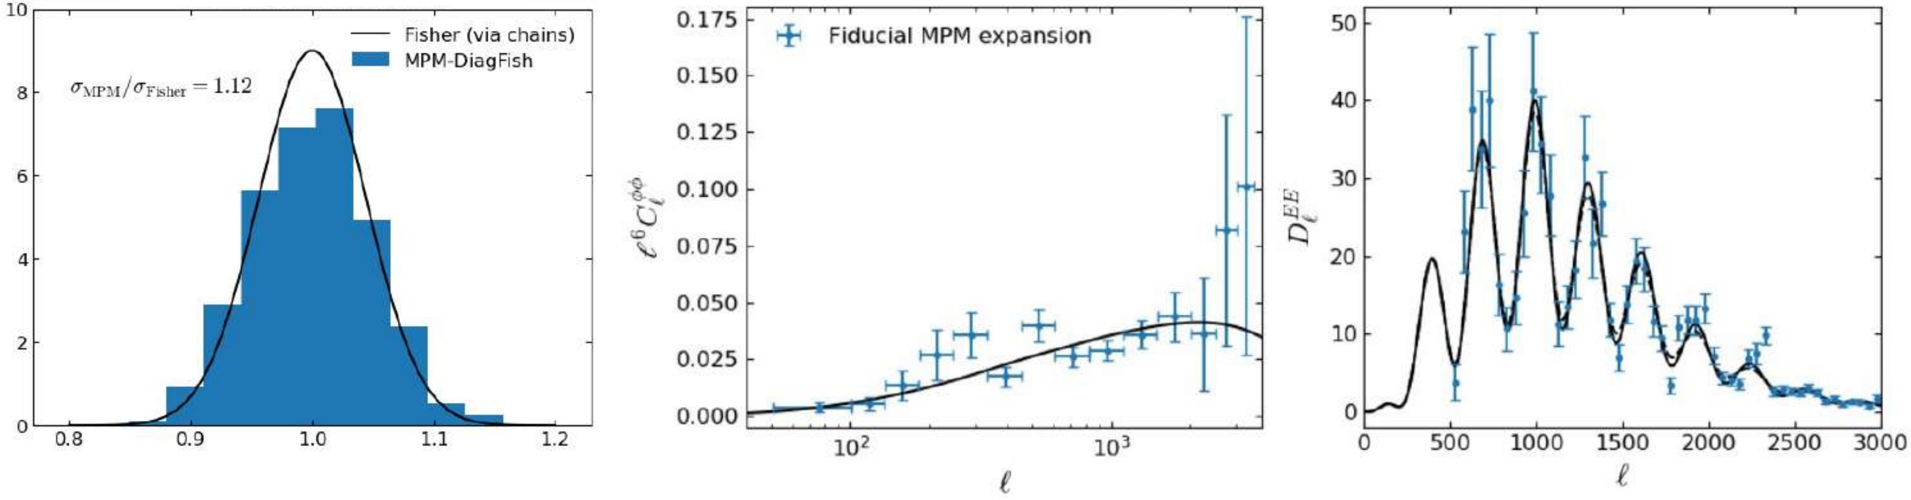

Conclusions

- Through the 2020s, all lensing analyses will eventually go beyond the QE
- The Bayesian solution is a promising way forward
- For exploring the Bayesian posterior, check out CMBLensing.jl

cosmicmar.com/CMBLensing.jl

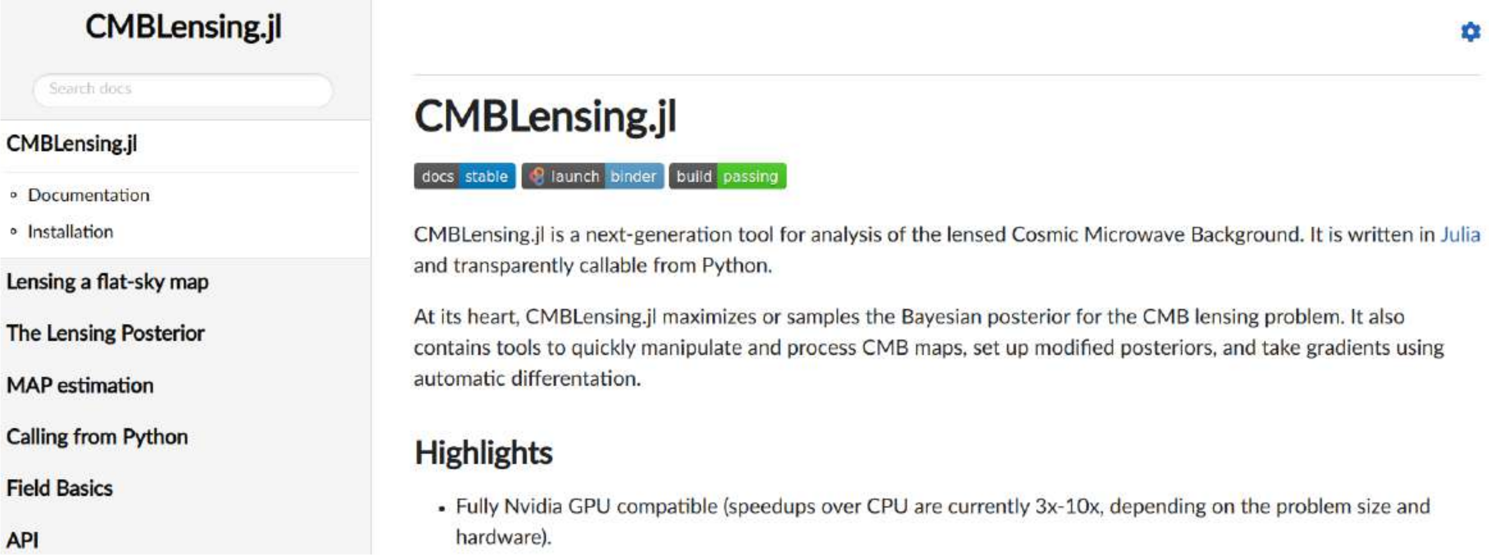

CMBLensing.jl

Search docs

CMBLensing.jl

- Documentation
- Installation

Lensing a flat-sky map

The Lensing Posterior

MAP estimation

Calling from Python

Field Basics

API

CMBLensing.jl

docs stable | launch binder | build passing

CMBLensing.jl is a next-generation tool for analysis of the lensed Cosmic Microwave Background. It is written in Julia and transparently callable from Python.

At its heart, CMBLensing.jl maximizes or samples the Bayesian posterior for the CMB lensing problem. It also contains tools to quickly manipulate and process CMB maps, set up modified posteriors, and take gradients using automatic differentation.

Highlights

- Fully Nvidia GPU compatible (speedups over CPU are currently 3x-10x, depending on the problem size and hardware).