Noise and dynamical sensitivity in critical planar percolation

Based on C. Garban, G. Pete & O. Schramm:
The Fourier spectrum of critical percolation, *Acta Math.* 2010.

Bernoulli(p) bond and site percolation

Given an (infinite) graph $G = (V, E)$ and $p \in [0, 1]$. Each site (or bond) is chosen open with probability p, closed with $1 - p$, independently of each other. Consider the open connected clusters. $\theta(p) := \mathbf{P}_p[0 \longleftrightarrow \infty]$.

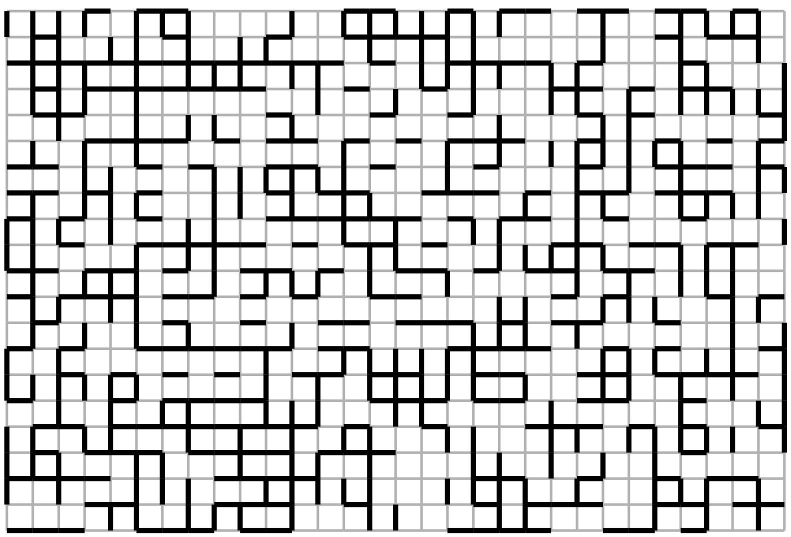

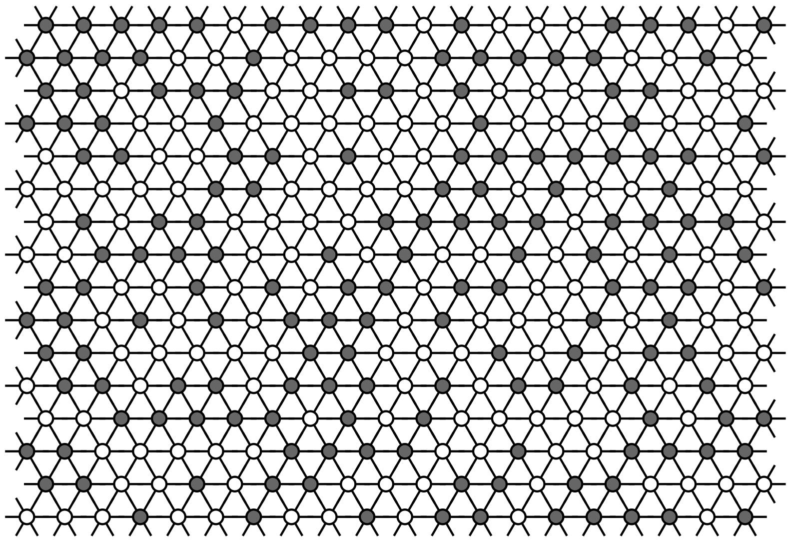

Theorem (Harris 1960 and Kesten 1980).

$p_c(\mathbb{Z}^2, \text{bond}) = p_c(\Delta, \text{site}) = 1/2$, and $\theta(1/2) = 0$.

For $p > 1/2$, there is a.s. one infinite cluster.

Crossing probabilities and criticality

$p \approx 0.8$ $\quad$ $p \approx 0.55$ $\quad$ $p = 0.5$ $\quad$ $p \approx 0.45$

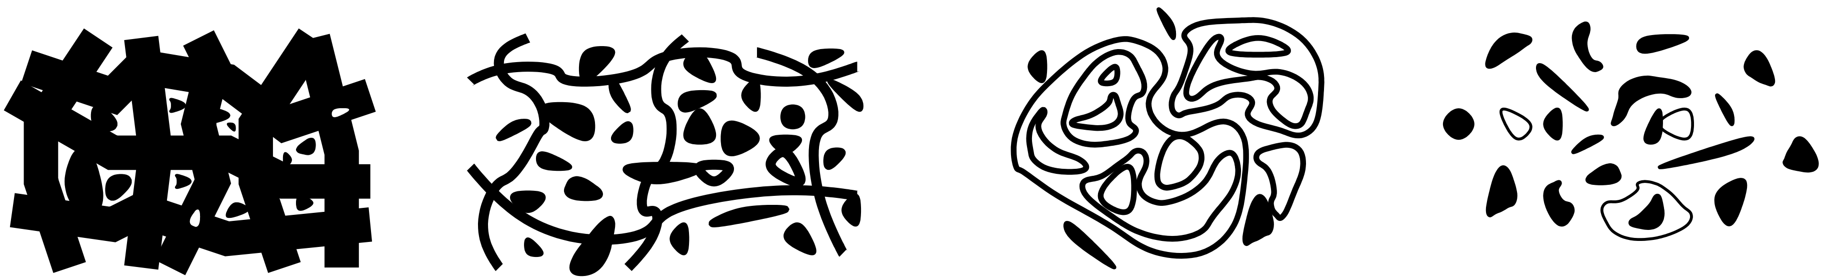

Theorem (Russo 1978 and Seymour-Welsh 1978). In critical percolation on almost any planar lattice, for $L, n > 0$,

$$0 < a_L < \mathbf{P}[\,\text{left-right crossing in } n \times Ln\,] < b_L < 1.$$

Same holds for annulus-crossings.

By repeating this on all scales, and gluing the pieces by FKG:

$$(r/R)^{\alpha} < \mathbf{P}[\,\partial B_r \longleftrightarrow \partial B_R\,] < (r/R)^{\beta}.$$

Moreover, for the (polychromatic) ℓ-arm probabilities

$$\alpha_\ell(r,R) := \mathbf{P}[\,\partial B_r \stackrel{\ell}{\longleftrightarrow} \partial B_R\,],$$

again have quasi-multiplicativity: $\alpha_\ell(r,R) \asymp \alpha_\ell(r,\rho)\,\alpha_\ell(\rho,R)$, and thus $c_\ell\,(r/R)^{C_\ell} < \alpha_\ell(r,R) < C_\ell\,(r/R)^{c_\ell}$.

But these are non-monotone events, so cannot use just FKG to prove this q-multiplicativity. Need Separation Lemma: conditioned on having ℓ arms from ∂B_r to ∂B_R, the collection of interfaces both in $B_R \setminus B_{R/2}$ and $B_{2r} \setminus B_r$ are well-separated. Then we can glue.

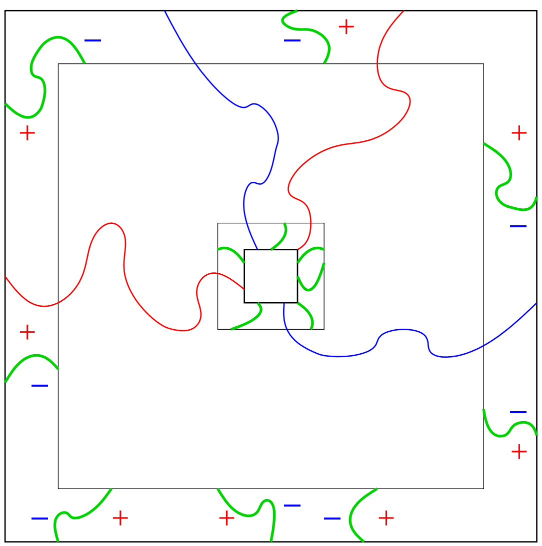

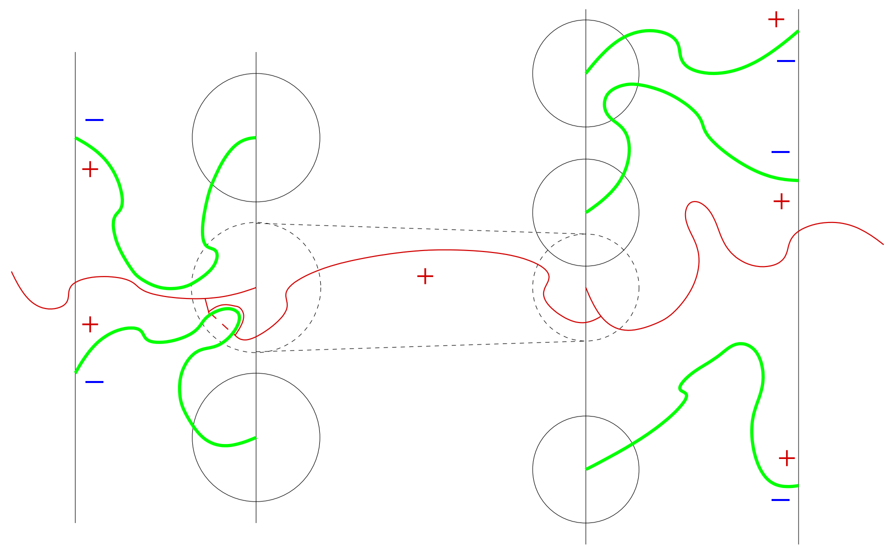

Bernoulli$(1/2)$ bond and site percolation

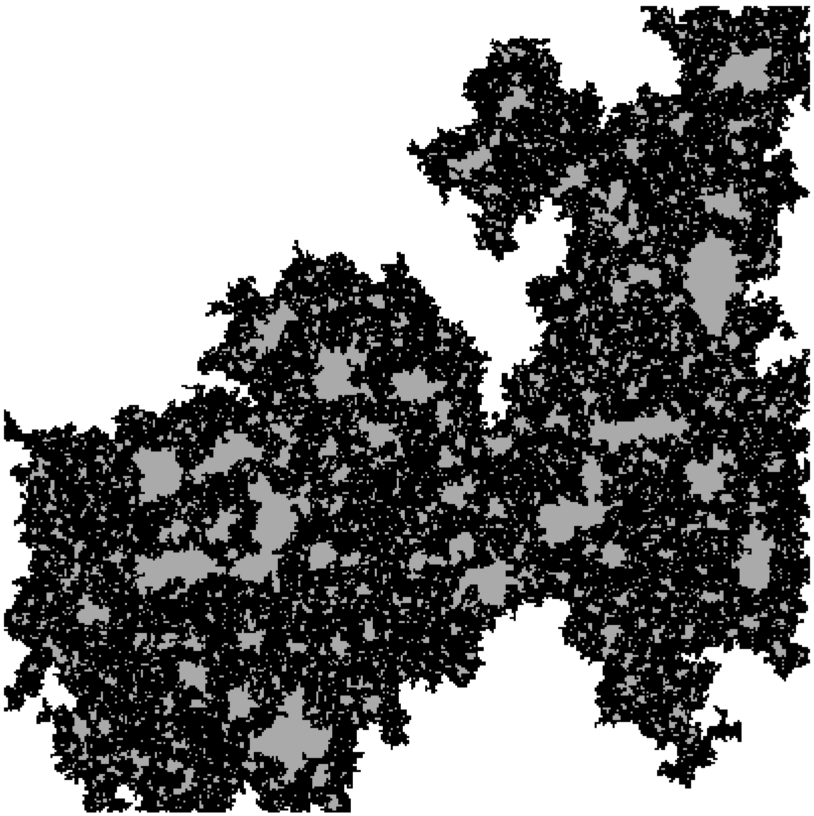

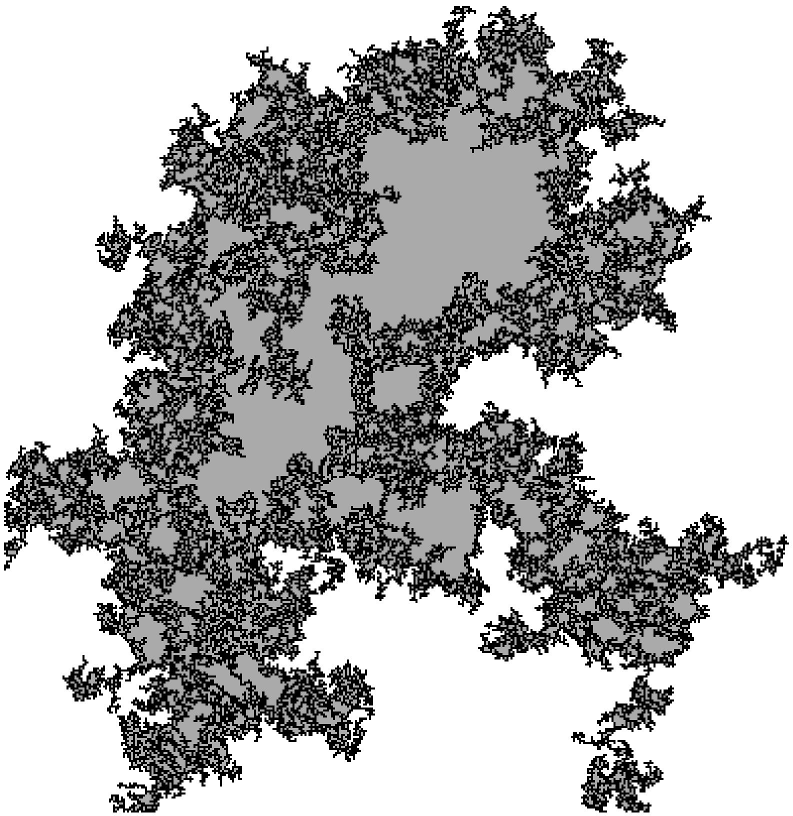

Conformal invariance on Δ

Theorem (Smirnov 2001). For $p = 1/2$ site percolation on Δ_η, and $\mathcal{Q} \subset \mathbb{C}$ a piecewise smooth quad (simply connected domain with four boundary points $\{a, b, c, d\}$),

$$\lim_{\eta \to 0} \mathbf{P}\Big[ab \longleftrightarrow cd \text{ inside } \mathcal{Q}, \text{ in percolation on } \Delta_\eta \Big]$$

exists, is strictly between 0 and 1, and conformally invariant.

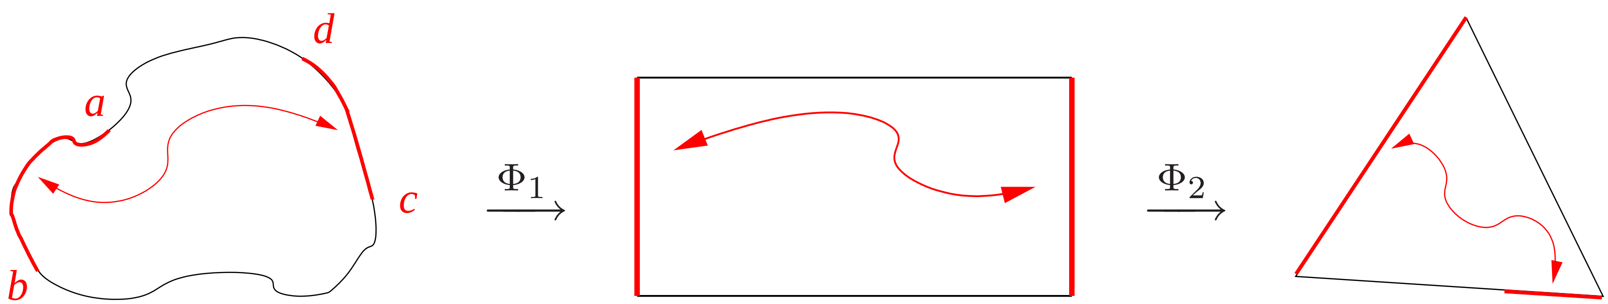

Calls for a continuum scaling limit, encoding macroscopic connectivity, cluster boundaries, etc. Aizenman '95, Schramm '00, Camia-Newman '06, Sheffield '09, Schramm-Smirnov '10. In physics, correlation functions.

SLE_6 exponents

Given the conformal invariance, the exploration path converges to the Stochastic Loewner Evolution with $\kappa = 6$ (Schramm 2000).

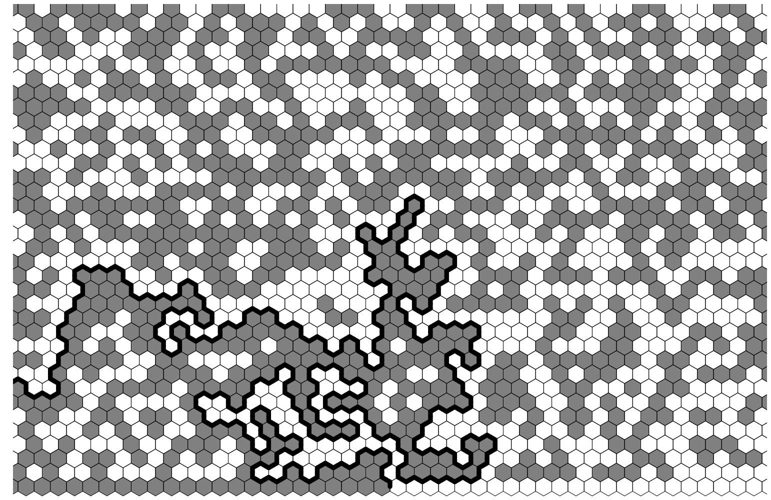

Using the SLE_6 curve, several critical exponents can be computed (Lawler-Schramm-Werner, Smirnov-Werner 2001, plus Kesten 1987), e.g.:

$$\alpha_4(r, R) := \mathbf{P}\Big[$$

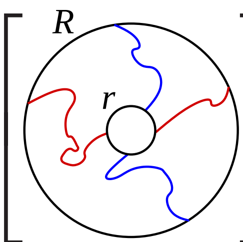

$$\Big] = (r/R)^{5/4+o(1)},$$

$\alpha_1(r, R) = (r/R)^{5/48+o(1)}$, and $\theta(p_c + \epsilon) := \mathbf{P}_{p_c+\epsilon}[0 \longleftrightarrow \infty] = \epsilon^{5/36+o(1)}$. Here $\beta = 5/36 = \frac{5/48}{2-5/4} = \frac{\xi_1}{2-\xi_4}$. (Will explain this numerology tomorrow.)

Percolation and noise

Take an ω critical percolation configuration. Let ω^ϵ be a new configuration, where each site (or bond) is resampled with probability ϵ, independently. (The ϵ-noised version of ω.)

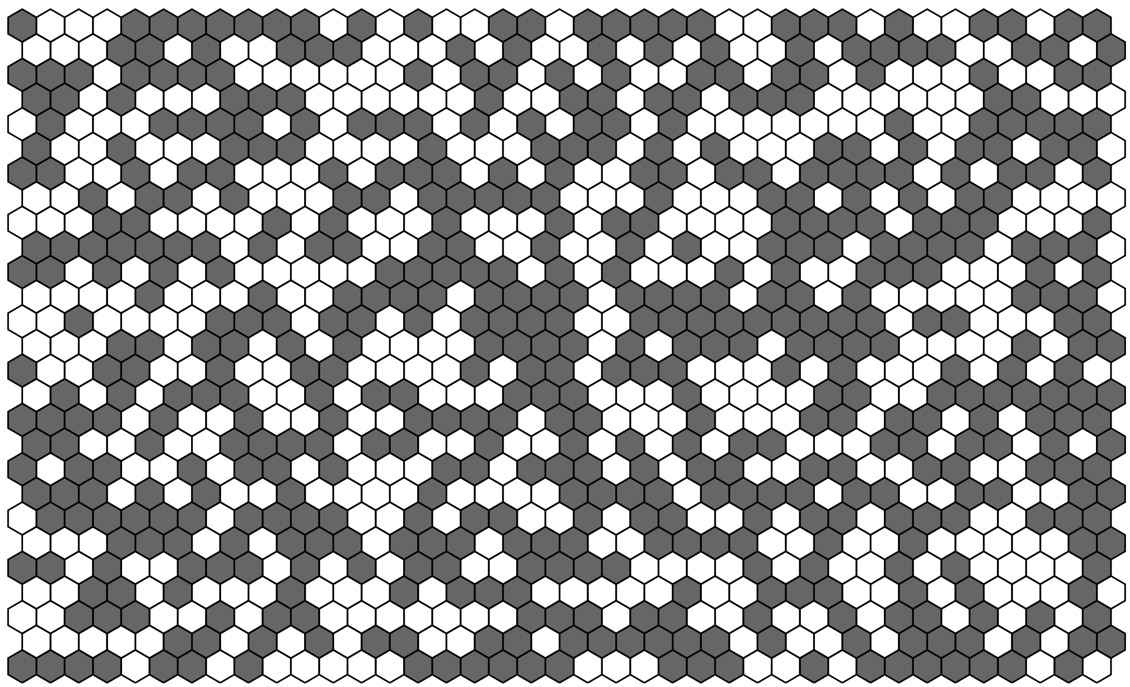

For how large an ϵ can we still predict from ω whether there is a left-right crossing in ω^ϵ?

Percolation and noise

Take an ω critical percolation configuration. Let ω^ϵ be a new configuration, where each site (or bond) is resampled with probability ϵ, independently. (The ϵ-noised version of ω.)

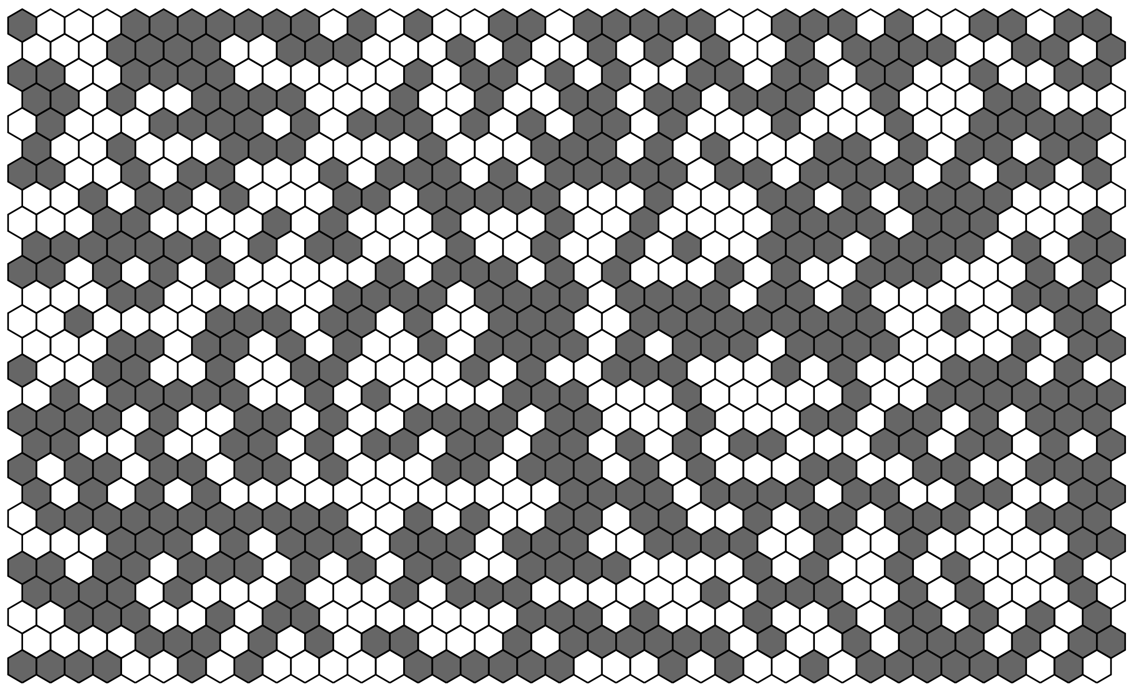

For how large an ϵ can we still predict from ω whether there is a left-right crossing in ω^ϵ?

Naive idea: how many pivotals are there?

A site (or bond) is pivotal in ω, if flipping it changes the existence of a left-right crossing.
$\mathbf{E}|\mathrm{Piv}_n| \asymp n^2\,\alpha_4(n) \quad (= n^{3/4+o(1)})$.

Furthermore, $\mathbf{E}[\,|\mathrm{Piv}_n|^2\,] \leqslant C\,(\mathbf{E}|\mathrm{Piv}_n|)^2$.
So, $\mathbf{P}\big[\,|\mathrm{Piv}_n| > \lambda\mathbf{E}|\mathrm{Piv}_n|\,\big] < C/\lambda^2$, any λ.

And not only $\exists\epsilon$ $\mathbf{P}\big[\,|\mathrm{Piv}_n| > \epsilon\mathbf{E}|\mathrm{Piv}_n|\,\big] > \epsilon$,
but $\mathbf{P}\big[\,0 < |\mathrm{Piv}_n| < \epsilon\mathbf{E}|\mathrm{Piv}_n|\,\big] \asymp \epsilon^{11/9+o(1)}$,
as $\epsilon \to 0$ (exponent only for Δ).

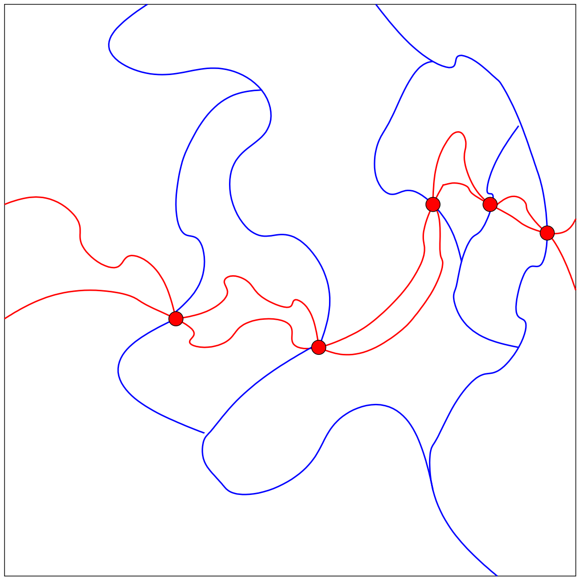

Cannot have many pivotals $\Longrightarrow$ If $\epsilon_n\mathbf{E}[\,|\mathrm{Piv}_n|\,] \to 0$, then we don't hit any pivotals (even in expectation) $\Longrightarrow$ Asymptotically full correlation.

Cannot have few pivotals (if there's any) $\Longrightarrow$ If $\epsilon_n\mathbf{E}[\,|\mathrm{Piv}_n|\,] \to \infty$, hit many pivotals (at least in expectation). But $\not\Longrightarrow$ asymptotic independence!

Noise sensitivity of percolation

All results use Fourier analysis of Boolean functions:

Theorem (Benjamini, Kalai & Schramm 1998). If $\epsilon > 0$ is fixed, and f_n is the indicator function for a left-right percolation crossing in an $n \times n$ square, then as $n \to \infty$

$$\mathbf{E}\big[\, f_n(\omega)\, f_n(\omega^\epsilon)\,\big] - \mathbf{E}\big[\, f_n(\omega)\,\big]^2 \to 0.$$

This holds for all $\epsilon = \epsilon_n > c/\log n$.

Theorem (Schramm & Steif 2005). Same if $\epsilon_n > n^{-a}$ for some positive $a > 0$. If triangular lattice, may take any $a < 1/8$.

Theorem (Garban, P & Schramm 2008). Same holds if and only if $\epsilon_n\, \mathbf{E}\big[\, |\mathrm{pivotals}|\,\big] \to \infty$. For triangular lattice, this threshold is $\epsilon_n = n^{-3/4+o(1)}$.

Dynamical percolation

Each variable is resampled according to an independent Poisson(1) clock. This is a Markov process $\{\omega(t) : t \in [0, \infty)\}$, in which $\omega(t+s)$ is an ϵ-noised version of $\omega(t)$, with $\epsilon = 1 - \exp(-s)$.

An exceptional time is such a (random) t, at which an almost sure property of the static process fails for $\omega(t)$.

Main example: (Non-)existence of an infinite cluster in percolation.

Toy example: Brownian motion on the circle does sometimes hit a given point, as opposed to its static version: a uniform random point.

In this toy example, the set of exceptional times is a random Cantor set of Lebesgue measure zero (because of Fubini) and Hausdorff-dimension $1/2$.

If the static event is not extremely unlikely, and it is very sensitive to noise, then we may have some chance to see an exceptional time.

Dynamical percolation results

Theorem (Häggström, Peres & Steif 1997).

- No exceptional times when $p \neq p_c$.
- No exceptional times when $p = p_c$ for bond percolation on $\mathbb{Z}^d$, $d \geqslant 19$.

The latter is essentially due to Hara-Slade '90 on the off-critical exponent $\beta = 1$:

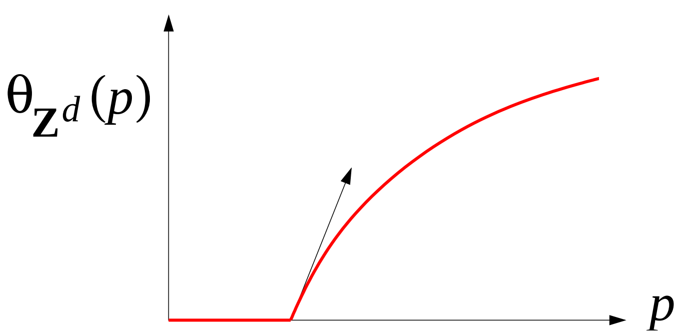

$\theta(p_c + \epsilon) < C\epsilon$, hence, even switching asymmetrically, $\mathbf{E}\big[$number of ϵ-subintervals of $[0,1]$ with $0 \longleftrightarrow \infty\big] \leqslant C$. But this exceptional set is closed without isolated points, so this number should blow up, if non-zero.

Theorem (Schramm & Steif 2005).

- There are exceptional times (a.s.) for critical site percolation on the triangular lattice.
- They have Hausdorff dimension in $[1/6, 31/36]$.

Theorem (Garban, P & Schramm 2008).

- There are exceptional times also on $\mathbb{Z}^2$.
- On the triangular grid they have Hausdorff dimension $31/36$.
- On the triangular grid, there are exceptional times with an infinite white **and** an infinite black cluster simultaneously. $(1/9 \leqslant \dim \leqslant 2/3)$

What is the Fourier spectrum and why is it useful?

$f_n : \{\pm 1\}^{V_n} \longrightarrow \{\pm 1\}$ indicator of left-right crossing, $V = V_n$ vertices.

$(N_\epsilon f)(\omega) := \mathbf{E}[\, f(\omega^\epsilon) \mid \omega\,]$ is the noise operator, acting on the space $L^2(\Omega, \mu)$, where $\Omega = \{\pm 1\}^V$, μ uniform measure, inner product $\mathbf{E}[\, fg\,]$.

Correlation: $\mathbf{E}[\, f(\omega^\epsilon) f(\omega)\,] - \mathbf{E}[\, f(\omega)\,]\mathbf{E}[\, f(\omega^\epsilon)\,] = \mathbf{E}[\, f(\omega) N_\epsilon f(\omega)\,] - \mathbf{E}[\, f(\omega)\,]^2$. So, we would like to diagonalize the noise operator N_ϵ.

Let χ_i be the function $\chi_i(\omega) = \omega(i)$, $\omega \in \Omega$.

For $S \subset V$, let $\chi_S := \prod_{i \in S} \chi_i$, the parity inside S. Then

$$N_\epsilon \chi_i = (1-\epsilon)\,\chi_i\,; \qquad N_\epsilon \chi_S = (1-\epsilon)^{|S|}\,\chi_S.$$

Moreover, the family $\{\chi_S,\ S \subseteq V\}$ is an orthonormal basis of $L^2(\Omega, \mu)$.

Any function $f \in L^2(\Omega, \mu)$ in this basis (Fourier-Walsh series):

$$\hat{f}(S) := \mathbf{E}[\, f \chi_S \,]\,; \qquad f = \sum_{S \subseteq V} \hat{f}(S)\, \chi_S\,.$$

The correlation:

$$\mathbf{E}[\, f N_\epsilon f\,] - \mathbf{E}[\, f\,]^2 = \sum_S \sum_{S'} \hat{f}(S)\, \hat{f}(S')\, \mathbf{E}\big[\, \chi_S\, N_\epsilon \chi_{S'} \,\big] - \mathbf{E}[\, f \chi_\emptyset \,]^2$$

$$= \sum_{\emptyset \neq S \subseteq V} \hat{f}(S)^2\, (1-\epsilon)^{|S|} = \sum_{k=1}^{|V_n|} (1-\epsilon)^k \sum_{|S|=k} \hat{f}(S)^2.$$

By Parseval, $\sum_S \hat{f}(S)^2 = \mathbf{E}[\, f^2\,] = 1$. So can define probability measure $\mathbf{P}\big[\, \mathscr{S}_f = S \,\big] := \hat{f}(S)^2 / \mathbf{E}[\, f^2\,]$, the spectral sample $\mathscr{S}_f \subseteq V$.

If, for some functions f_n and numbers k_n, we have $\mathbf{P}\big[0 < |\mathscr{S}_n| < t k_n\big] \to 0$ as $t \to 0$, uniformly in n, then $(1-\epsilon)^k \sim \exp(-\epsilon k)$ implies that for $\epsilon_n \gg 1/k_n$ we have asymptotic independence. Maybe with $k_n = \mathbf{E}|\mathscr{S}_n|$?

Pivotals versus spectral sample

$\nabla_i f(\omega) := f(\sigma_i(\omega)) - f(\omega) \in \{-2, 0, +2\}$ gradient.

$\nabla_i f(\omega) = \sum_S \hat{f}(S)[\chi_S(\sigma_i(\omega)) - \chi_S(\omega)]$, hence $\widehat{\nabla_i f}(S) = -2\hat{f}(S)\mathbf{1}_{i\in S}$.

$\mathbf{P}[\, i \in \mathrm{Piv}_f \,] = \frac{1}{4}\|\nabla_i f\|_2^2 = \frac{1}{4}\sum_S \widehat{\nabla_i f}(S)^2 = \sum_{S\ni i} \hat{f}(S)^2 = \mathbf{P}[\, i \in \mathscr{S}_f \,]$.

Thus, $\mathbf{E}|\mathscr{S}_f| = \mathbf{E}|\mathrm{Piv}_f|$. So, the pivotal upper bound for noise sensitivity is sharp if there is tightness around $\mathbf{E}|\mathscr{S}|$.

Will see $\mathbf{P}[\, i, j \in \mathrm{Piv}_f \,] = \mathbf{P}[\, i, j \in \mathscr{S}_f \,]$, hence $\mathbf{E}|\mathscr{S}_f|^2 = \mathbf{E}|\mathrm{Piv}_f|^2$.

Not for more points and higher moments! Both random subsets measure the "influence" or "relevance" of bits, but in different ways.

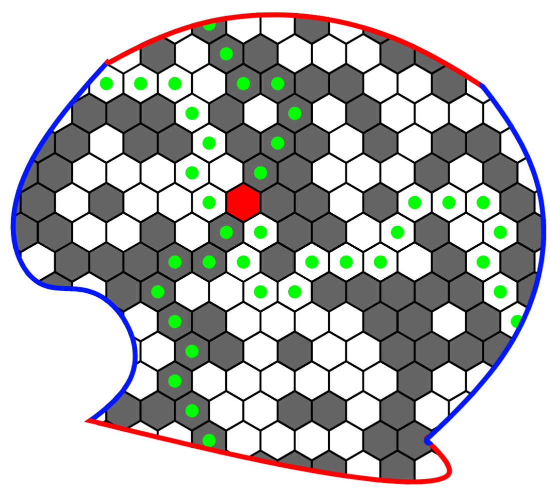

For percolation, $\mathbf{E}[\, |\mathrm{Piv}_n|^2 \,] \leqslant C\,(\mathbf{E}|\mathrm{Piv}_n|)^2$, hence $\exists c > 0$ s.t. $\mathbf{P}\big[\, |\mathscr{S}_n| > c\,\mathbf{E}|\mathscr{S}_n| \,\big] > c$. That's why one hopes for tightness around mean.

Three very simple examples

$\text{Dictator}_n(x_1, \ldots, x_n) := x_1$.
Here $\text{Cov}\big[\, \text{Dic}_n(x), \text{Dic}_n(x^\epsilon) \,\big] = 1 - \epsilon$, so noise-stable.
And $\mathbf{P}\big[\, \mathscr{S}_n = \{x_1\} \,\big] = 1$.

$\text{Majority}_n(x_1, \ldots, x_n) := \operatorname{sgn}(x_1 + \cdots + x_n) \approx \frac{1}{\sqrt{n}}(x_1 + \cdots + x_n)$.
Here $\text{Cov}\big[\, \text{Maj}_n(x), \text{Maj}_n(x^\epsilon) \,\big] = 1 - O(\epsilon)$, so noise-stable.
And $\mathbf{P}\big[\, \mathscr{S}_n = \{x_i\} \,\big] \asymp 1/n$, most of the weight is on singletons.
On the other hand, $\mathbf{E}|\mathscr{S}_n| = \mathbf{E}|\text{Piv}_n| \asymp \frac{1}{\sqrt{n}}\, n \asymp \sqrt{n}$.

$\text{Parity}_n(x_1, \ldots, x_n) := x_1 \cdots x_n$
Here $\text{Cov}\big[\, \text{Par}_n(x), \text{Par}_n(x^\epsilon) \,\big] = (1 - \epsilon)^n$, the most sensitive to noise.
And $\mathbf{P}\big[\, \mathscr{S}_n = \{x_1, \ldots, x_n\} \,\big] = 1$.

Benjamini, Kalai & Schramm 1998

Theorem. A sequence f_n of monotone Boolean functions is noise sensitive, i.e., for any fixed $\epsilon > 0$,

$$\mathbf{E}\big[\, f_n(\omega)\, f_n(\omega^\epsilon)\,\big] - \mathbf{E}\big[\, f_n(\omega)\,\big]^2 \to 0$$

as $n \to \infty$, iff it is asymptotically uncorrelated with all weighted majorities $\mathrm{Maj}_w(x_1, \ldots, x_n) = \mathrm{sign} \sum_{i=1}^n x_i w_i$. Also, not very slow decorrelation with all subset-majorities is enough for sensitivity.

Theorem. The left-right percolation crossing in an $n \times n$ square is noise sensitive, even with $\epsilon = \epsilon_n > c/\log n$.

Schramm & Steif 2005

Theorem. If $f : \Omega \longrightarrow \mathbb{R}$ can be computed with a randomized algorithm with revealment δ (each bit is read only with probability $\leqslant \delta$), then

$$\sum_{S:|S|=k} \hat{f}(S)^2 \leqslant \delta\, k\, \|f\|_2^2\,.$$

For left-right crossing in $n \times n$ box on the hexagonal lattice, exploration interface with random starting point gives revealment $n^{-1/4+o(1)}$ (it has length $n^{7/4+o(1)}$, given by 2-arm exponent), while $\sum_{k \leqslant m} k \asymp m^2$, thus:

Theorem. Left-right crossing on the triangular lattice is noise sensitive under $\epsilon_n > n^{-a}$, with any $a < 1/8$. Even on square lattice, can take some positive $a > 0$.

The revealment is at least $n^{-1/2+o(1)}$ for *any* algorithm computing the crossing, hence this method can give only $n^{-1/4+o(1)}$-sensitivity, far from the conjectured $\epsilon_n = n^{-3/4+o(1)}$.

The GPS approach, 2008

Although goal is to understand size, Gil Kalai suggested trying to understand entire distribution of $\mathscr{S}_f$. A strange random set of bits.

Effective sampling? If f is an effectively computable Boolean function, then there is an effective quantum algorithm for $\mathscr{S}_f$ [Bernstein-Vazirani 1993].

For $\mathscr{S}_{\mathcal{Q},n}$ (left-right crossing in a conformal rectangle $\mathcal{Q}$, mesh $1/n$), [Smirnov '01] $+$ [Tsirelson '04] $+$ [Schramm-Smirnov '11] implies that it has a conformally invariant scaling limit.

How to prove tightness for the size of strange random fractal-like sets?

Basic properties of the spectral sample

For $A \subseteq V$: $\mathbf{E}\big[\, \chi_S \mid \mathscr{F}_A \,\big] = \begin{cases} \chi_S & S \subseteq A\,, \\ 0 & \text{otherwise}\,. \end{cases}$

Therefore, $\mathbf{E}\big[\, f \mid \mathscr{F}_A \,\big] = \sum_{S \subseteq A} \widehat{f}(S)\, \chi_S$, a nice projection.

Also, for $T \subseteq A$: $\mathbf{E}\big[\, f\, \chi_T \mid \mathscr{F}_{A^c} \,\big] = \sum_{S \subseteq A^c} \widehat{f}(T \cup S)\, \chi_S$, hence

$$\mathbf{E}\Big[\mathbf{E}\big[\, f\, \chi_T \mid \mathscr{F}_{A^c} \,\big]^2\Big] = \sum_{S \subseteq A^c} \widehat{f}(T \cup S)^2 = \mathbf{P}\big[\, \mathscr{S} \cap A = T \,\big]\,.$$

This is the Random Restriction Lemma of Linial-Mansour-Nisan ‘93. E.g.,

$$\begin{aligned} \mathbf{P}\big[\, i,j \in \mathscr{S}_f \,\big] &= \mathbf{E}\Big[\mathbf{E}\big[\, f \chi_{\{i,j\}} \mid \mathscr{F}_{\{i,j\}^c} \,\big]^2\Big] \\ &= \frac{1}{4}\mathbf{P}\big[\, \omega|_{\{i,j\}^c} \text{ is such that } i,j \text{ each may be pivotal} \,\big] \\ &= \mathbf{P}\big[\, i,j \in \mathrm{Piv}_f \,\big]\,. \end{aligned}$$

How does $\left[\mathscr{S}_n \cap B \,\middle|\, \mathscr{S}_n \cap B \neq \emptyset\right]$ look like?

B as set has to be pivotal.

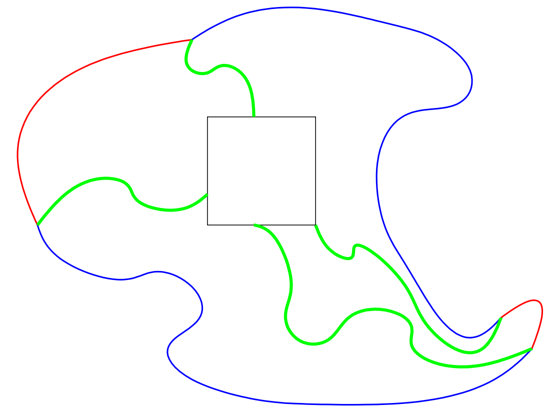

Strong Separation Lemma. For $\mathrm{d}(B, \partial Q) > \mathrm{diam}(B)$, conditioned on the 4 interfaces to reach ∂B, with *arbitrary starting points*, with a uniformly positive conditional probability the interfaces are well-separated around ∂B. Very bad separation is very unlikely. [Simple proof by Damron-Sapozhnikov '09, following Kesten '87. Also explained in Appendix to GPS '11.]

Corollary 1. $\mathbf{P}\Big[\mathscr{S}_n \cap B_r \neq \emptyset\Big] \asymp \alpha_4(r,n)\,.$

Corollary 2. $\mathbf{E}\Big[|\mathscr{S}_n \cap B_r| \,\Big|\, \mathscr{S}_n \cap B_r \neq \emptyset\Big] \asymp r^2\,\alpha_4(1,r) \asymp \mathbf{E}|\mathscr{S}_r|\,.$

Self-similarity for left-right crossing of $n \times n$ square

$$\mathbf{E}|\mathscr{S}_n| = \mathbf{E}|\mathrm{Piv}_n| \asymp n^2\,\alpha_4(1,n) \overset{\Delta}{\asymp} n^{3/4+o(1)}\,,$$

$$\mathbf{E}|\mathscr{S}_n(r)| := \mathbf{E}\Big[\,\#\big\{r\text{-boxes } \mathscr{S}_n \cap B_r \neq \emptyset\big\}\,\Big] \asymp \frac{n^2}{r^2}\,\alpha_4(r,n) \asymp \mathbf{E}|\mathscr{S}_{n/r}|\,,$$

$$\mathbf{E}\Big[\,|\mathscr{S}_n \cap B_r| \;\Big|\; \mathscr{S}_n \cap B_r \neq \emptyset\Big] \asymp r^2\,\alpha_4(1,r) \asymp \mathbf{E}|\mathscr{S}_r|\,.$$

Of course, $r^2\,\alpha_4(1,r) \cdot \frac{n^2}{r^2}\,\alpha_4(r,n) \asymp n^2\,\alpha_4(1,n)$, by quasi-multiplicativity.

Self-similarity for left-right crossing of $n \times n$ square

$$\mathbf{E}|\mathscr{S}_n| = \mathbf{E}|\mathrm{Piv}_n| \asymp n^2\,\alpha_4(1,n) \overset{\Delta}{\asymp} n^{3/4+o(1)}\,,$$

$$\mathbf{E}|\mathscr{S}_n(r)| := \mathbf{E}\Big[\#\big\{r\text{-boxes } \mathscr{S}_n \cap B_r \neq \emptyset\big\}\Big] \asymp \frac{n^2}{r^2}\,\alpha_4(r,n) \asymp \mathbf{E}|\mathscr{S}_{n/r}|\,,$$

$$\mathbf{E}\Big[\,|\mathscr{S}_n \cap B_r| \;\Big|\; \mathscr{S}_n \cap B_r \neq \emptyset\Big] \asymp r^2\,\alpha_4(1,r) \asymp \mathbf{E}|\mathscr{S}_r|\,.$$

Of course, $r^2\,\alpha_4(1,r)\cdot\frac{n^2}{r^2}\,\alpha_4(r,n) \asymp n^2\,\alpha_4(1,n)$, by quasi-multiplicativity.

Similar to the zero-set of simple random walk: $\mathbf{E}|\mathcal{Z}_n| \asymp n\,n^{-1/2} = n^{1/2}$,

$$\mathbf{E}|\mathcal{Z}_n(r)| := \mathbf{E}\Big[\#\big\{r\text{-intervals } \mathcal{Z}_n \cap I_r \neq \emptyset\big\}\Big] \asymp \frac{n}{r}\,(n/r)^{-1/2} \asymp \mathbf{E}|\mathcal{Z}_{n/r}|\,,$$

$$\mathbf{E}\Big[\,|\mathcal{Z}_n \cap I_r| \;\Big|\; \mathcal{Z}_n \cap I_r \neq \emptyset\Big] \asymp r\,r^{-1/2} \asymp \mathbf{E}|\mathcal{Z}_r|\,.$$

These results are related to the existence of scaling limits.

What concentration can we expect?

$\mathscr{S}_n$ is very different from uniform set of similar density:
i.i.d. $\mathbf{P}[\, x \in \mathscr{U}_n \,] = n^{-5/4}$. Hence $\mathbf{E}|\mathscr{U}_n| = n^{3/4}$.

For large r ($\gg n^{5/8}$), this $\mathscr{U}_n$ intersects every r-box;
for small r, if it intersects one, there is just one point there.

Concentration of size: roughly within $\sqrt{\mathbf{E}|\mathscr{U}_n|} = n^{3/8}$.

A bit more similar: for $i = 1, \ldots, (n/r)^2$, i.i.d. $\mathbf{P}\big[\, X_i = r^{3/4} \,\big] = (n/r)^{-5/4}$,
$X_i = 0$ otherwise. Then $S_{n,r} := \sum_i X_i$. Hence $\mathbf{E}|S_{n,r}| = n^{3/4}$.

For $r = n^\gamma$, size $|S_{n,r}|$ is concentrated within $n^{3/8(1+\gamma)}$, still $o(\mathbf{E}|S_{n,r}|)$.

For self-similar sets, we expect only tightness around the mean:
$\mathbf{P}\big[\, 0 < |\mathscr{S}_n| < \lambda\, \mathbf{E}|\mathscr{S}_n| \,\big] \to 0$ as $\lambda \to 0$, uniformly in n.

Proving tightness with a lot of independence

Assume we have the following ingredients, true for the zeroes:

(1) $\mathbf{P}\Big[\, |\mathcal{Z}_n \cap I_r| > c\,\mathbf{E}|\mathcal{Z}_r| \;\Big|\; \mathcal{Z}_n \cap I_r \neq \emptyset,\ \mathscr{F}_{[n]\setminus I_r} \Big] \geqslant c > 0.$

(2) $\mathbf{P}\big[\, |\mathcal{Z}_n(r)| = k \,\big] \leqslant g(k)\,\mathbf{P}\big[\, |\mathcal{Z}_n(r)| = 1 \,\big]$, with sub-exponential $g(k)$:

when the r-intervals intersected are scattered, have to pay k times to get to and leave them, and this cost is not balanced by combinatorial entropy.

Proving tightness with a lot of independence

Assume we have the following ingredients, true for the zeroes:

(1) $\mathbf{P}\Big[\, |\mathcal{Z}_n \cap I_r| > c\,\mathbf{E}|\mathcal{Z}_r| \,\Big|\, \mathcal{Z}_n \cap I_r \neq \emptyset,\ \mathscr{F}_{[n]\setminus I_r} \Big] \geqslant c > 0.$

(2) $\mathbf{P}\big[\, |\mathcal{Z}_n(r)| = k \,\big] \leqslant g(k)\, \mathbf{P}\big[\, |\mathcal{Z}_n(r)| = 1 \,\big]$, with sub-exponential $g(k)$:

when the r-intervals intersected are scattered, have to pay k times to get to and leave them, and this cost is not balanced by combinatorial entropy.

$$\mathbf{P}\big[\, 0 < |\mathcal{Z}_n| < c\,\mathbf{E}|\mathcal{Z}_r| \,\big] = \sum_{k \geqslant 1} \mathbf{P}\Big[\, 0 < |\mathcal{Z}_n| < c\,\mathbf{E}|\mathcal{Z}_r|\,,\ |\mathcal{Z}_n(r)| = k \,\Big]$$

$$\text{by (1):} \quad \leqslant \sum_{k \geqslant 1} (1-c)^k\, \mathbf{P}\big[\, |\mathcal{Z}_n(r)| = k \,\big]$$

$$\text{by (2):} \quad \leqslant O(1)\, \mathbf{P}\big[\, |\mathcal{Z}_n(r)| = 1 \,\big] \asymp (n/r)^{1-3/2},$$

which, using $\lambda = \frac{c\,\mathbf{E}|\mathcal{Z}_r|}{\mathbf{E}|\mathcal{Z}_n|} \asymp (r/n)^{1/2}$, reads as $\mathbf{P}\big[\, 0 < |\mathcal{Z}_n| < \lambda\,\mathbf{E}|\mathcal{Z}_n| \,\big] \asymp \lambda$.

But we know much less independence for $\mathscr{S}_n$

$$(1') \quad \mathbf{P}\Big[\, |\mathscr{S}_n \cap B_r/3| > c\,\mathbf{E}|\mathscr{S}_r| \;\Big|\; \mathscr{S}_n \cap B_r \neq \emptyset = \mathscr{S}_n \cap W \,\Big] \geqslant c > 0,$$

for any W that is not too close to B_r.

Why only this negative conditioning? Inclusion formula:

$$\mathbf{P}[\,\mathscr{S}_f \subset U\,] = \sum_{S \subset U} \hat{f}(S)^2 = \mathbf{E}\Big[\Big(\sum_{S \subset U} \hat{f}(S)\,\chi_S\Big)^2\Big] = \mathbf{E}\Big[\mathbf{E}[\,f \mid \mathscr{F}_U\,]^2\Big].$$

From this, for disjoint subsets A and B,

$$\begin{aligned}\mathbf{P}[\,\mathscr{S}_f \cap B \neq \emptyset = \mathscr{S}_f \cap A\,] &= \mathbf{P}[\,\mathscr{S}_f \subseteq A^c\,] - \mathbf{P}[\,\mathscr{S}_f \subseteq (A \cup B)^c\,] \\ &= \mathbf{E}\Big[\mathbf{E}[\,f \mid \mathscr{F}_{A^c}\,]^2 - \mathbf{E}[\,f \mid \mathscr{F}_{(A\cup B)^c}\,]^2\Big] \\ &= \mathbf{E}\Big[\big(\mathbf{E}[\,f \mid \mathscr{F}_{A^c}\,] - \mathbf{E}[\,f \mid \mathscr{F}_{(A\cup B)^c}\,]\big)^2\Big].\end{aligned}$$

So, what are we going to do?

With quite a lot of work for both items,

(1') $\mathbf{P}\Big[\, |\mathscr{S}_n \cap B_r/3| > c\,\mathbf{E}|\mathscr{S}_r| \;\Big|\; \mathscr{S}_n \cap B_r \neq \emptyset = \mathscr{S}_n \cap W \,\Big] \geqslant c > 0.$

(2) $\mathbf{P}[\, |\mathscr{S}_n(r)| = k \,] \leqslant g(k)\,\mathbf{P}[\, |\mathscr{S}_n(r)| = 1 \,]$, with sub-exponential $g(k)$.

We could repeat (1') for many r-boxes only if "not enough points in one box" meant "we found nothing in that box".

So, take an independent random dilute sample: $\mathbf{P}[\, x \in \mathcal{R} \,] = 1/\mathbf{E}|\mathscr{S}_r|$ i.i.d.
Then, $|\mathscr{S}_n \cap B_r/3|$ is small $\implies \mathcal{R} \cap \mathscr{S}_n \cap B_r/3 = \emptyset$ is likely,
and $|\mathscr{S}_n \cap B_r/3|$ is large $\implies \mathcal{R} \cap \mathscr{S}_n \cap B_r/3 \neq \emptyset$ is likely.

So, what are we going to do?

With quite a lot of work for both items,

(1') $\quad \mathbf{P}\Big[\, |\mathscr{S}_n \cap B_r/3| > c\,\mathbf{E}|\mathscr{S}_r| \,\Big|\, \mathscr{S}_n \cap B_r \neq \emptyset = \mathscr{S}_n \cap W \,\Big] \geqslant c > 0.$

(2) $\quad \mathbf{P}\big[\, |\mathscr{S}_n(r)| = k \,\big] \leqslant g(k)\, \mathbf{P}\big[\, |\mathscr{S}_n(r)| = 1 \,\big]$, with sub-exponential $g(k)$.

We could repeat (1') for many r-boxes only if "not enough points in one box" meant "we found nothing in that box".

So, take an independent random dilute sample: $\mathbf{P}[\, x \in \mathcal{R}\,] = 1/\mathbf{E}|\mathscr{S}_r|$ i.i.d.
Then, $|\mathscr{S}_n \cap B_r/3|$ is small $\implies \mathcal{R} \cap \mathscr{S}_n \cap B_r/3 = \emptyset$ is likely,
and $|\mathscr{S}_n \cap B_r/3|$ is large $\implies \mathcal{R} \cap \mathscr{S}_n \cap B_r/3 \neq \emptyset$ is likely.

But $\mathbf{P}\Big[\, \mathscr{S}_n \neq \emptyset = \mathcal{R} \cap \mathscr{S}_n \,\Big|\, |\mathscr{S}_n(r)| = k \,\Big]$ is still problematic conditioning.

A strange large deviations lemma solves the issue.

The strange large deviation lemma

Suppose $X_i, Y_i \in \{0,1\}$, $i = 1, \ldots, n$, and that $\forall J \subset [n]$ and $\forall i \in [n] \setminus J$

$$\mathbf{P}\big[\, Y_i = 1 \,\big|\, \forall_{j\in J} Y_j = 0 \,\big] \geqslant c\,\mathbf{P}\big[\, X_i = 1 \,\big|\, \forall_{j\in J} Y_j = 0 \,\big].$$

Then

$$\mathbf{P}\Big[\forall_i Y_i = 0\Big] \leqslant c^{-1}\,\mathbf{E}\Big[\exp\Big(-(c/e)\sum_i X_i\Big)\Big].$$

We use this with $X_j := 1_{\{\mathscr{S}\cap B_j \neq \emptyset\}}$ and $Y_j := 1_{\{\mathscr{S}\cap B_j \cap \mathcal{R} \neq \emptyset\}}$.

Proof: Instead of sequential scan, average everything together. Choose $J \subset [n]$ randomly, Bernoulli$(1-p)$. Get $\mathbf{E}\big[\, Y\, p^Y \,\big] \geqslant c\,\mathbf{E}\big[\, X\, p^{Y+1} \,\big]$.

So, $\mathbf{E}[\, Z\,] \geqslant 0$, where $Z := (Y - c\,p\,X)\,p^Y$. Choose $p := e^{-1}$. Maximize Z over Y, and get the bound $Z \leqslant \exp(-1 - c\,X/e)$. Altogether, $c\,e^{-1}\,\mathbf{P}\big[\, Y = 0 < X \,\big] \leqslant \mathbf{E}\big[\, 1_{X>0}\exp(-1 - c\,X/e)\,\big]$, and done.

Final result for the spectral sample

If $r \in [1, n]$, then $\big\{|\mathscr{S}_n| < \mathbf{E}|\mathscr{S}_r|\big\}$ is basically equivalent to being contained inside some $r \times r$ sub-square:

$$\mathbf{P}\big[\,0 < |\mathscr{S}_n| < \mathbf{E}|\mathscr{S}_r|\,\big] \asymp \alpha_4(r,n)^2 \left(\frac{n}{r}\right)^2.$$

In particular, on the triangular lattice Δ,

$$\mathbf{P}\big[\,0 < |\mathscr{S}_n| < \lambda\,\mathbf{E}|\mathscr{S}_n|\,\big] \asymp \lambda^{2/3}.$$

The *scaling limit* of $\mathscr{S}_n$ is a conformally invariant Cantor-set with Hausdorff-dimension $3/4$.

GPS (2010-12) proves that the scaling limit of dynamical percolation exists as a Markov process; for mesh $1/n$ the time-scale is $tn^{-3/4+o(1)}$. The above implies that this process is ergodic, with correlations decaying as $t^{-2/3}$.

Question 1: Can one build similar proofs for other Boolean functions?

Question 2: Self-similarity of Piv_n and $\mathscr{S}_n$ is a lot of restriction on these random sets. And it's not only because of conformal invariance: Gil Kalai noticed that the spectral sample of recursive 3-wise majority is the leaves of a GW-tree! This phenomenon might be somehow general:

Influence-Entropy conjecture [Friedgut-Kalai 1996]: For some universal constant C, for any Boolean function f,

$$\mathrm{SpecEnt}(f) := \sum_{S\subset[n]} \widehat{f}(S)^2 \log\frac{1}{\widehat{f}(S)^2} \leqslant C \times$$

$$\times\, \mathrm{Influence}(f) := \mathbf{E}|\mathscr{S}_f| = \mathbf{E}|\mathrm{Piv}_f| = \sum_{S\subset[n]} \widehat{f}(S)^2 |S|\,.$$

I.e., there is no $\log$ factor in the entropy as it would be in uniform.

I think I can do it for Piv_n, but not enough independence is known in $\mathscr{S}_n$.