

Sharp phase transition for Cox percolation^{*†}

Christian Hirsch[‡] Benedikt Jahnel[§] Stephen Muirhead[¶]

Abstract

We prove the sharpness of the percolation phase transition for a class of Cox percolation models, i.e., models of continuum percolation in a random environment. The key requirements are that the environment has a finite range of dependence, satisfies a local boundedness condition and can be constructed from a discrete iid random field, however the FKG inequality need not hold. The proof combines the OSSS inequality with a coarse-graining construction that allows us to compare different notions of influence.

Keywords: Cox point process; OSSS inequality; continuum percolation; Poisson point process; random environment; Delaunay tessellation.

MSC2020 subject classifications: 60F10; 60K35.

Submitted to ECP on March 11, 2022, final version accepted on September 13, 2022.

Supersedes arXiv:2203.01251.

1 Introduction

The field of continuum percolation deals with the existence and properties of giant connected components of a geometric graph on a stochastic system of points scattered at random in Euclidean space [13]. Since its early days, continuum percolation has attracted attention from researchers in wireless communication [7]. This appeal is based on the prospect of using the asymptotic theory to predict the behavior of large systems of devices that interact in a peer-to-peer fashion.

^{*}The authors are indebted to R. Lachièze-Rey and E. Cali for illuminating discussions and remarks on early versions of the manuscript.

[†]CH would like to acknowledge the financial support of the CogniGron research center and the Ubbo Emmius Funds (Univ. of Groningen). BJ was supported by the German Research Foundation under Germany's Excellence Strategy MATH+: The Berlin Mathematics Research Center, EXC-2046/1 project ID: 390685689, and the Leibniz Association within the Leibniz Junior Research Group on Probabilistic Methods for Dynamic Communication Networks as part of the Leibniz Competition. SM was supported by the Australian Research Council (ARC) Discovery Early Career Researcher Award DE200101467.

[‡]Department of Mathematics, Aarhus University, Ny Munkegade, 118, 8000, Aarhus C, Denmark; Bernoulli Institute for Mathematics, Computer Science and Artificial Intelligence, University of Groningen (Univ Groningen), Nijenborgh 9, NL-9747 AG Groningen, Netherlands; CogniGron (Groningen Cognitive Systems and Materials Center), University of Groningen (Univ Groningen), Nijenborgh 4, NL-9747 AG Groningen, Netherlands. E-mail: hirsch@math.au.dk

[§]Technische Universität Braunschweig, Universitätsplatz 2, 38106 Braunschweig, Germany, and Weierstrass Institute for Applied Analysis and Stochastics, Mohrenstraße 39, 10117 Berlin, Germany. E-mail: benedikt.jahnel@tu-braunschweig.de

[¶]School of Mathematics and Statistics, University of Melbourne, Melbourne, Australia. E-mail: smui@unimelb.edu.au

So far, the majority of results assume that the devices are scattered entirely at random in the infinite Euclidean plane, in the sense that they form a homogeneous Poisson point process. However, this assumption is in stark contrast with the topology of modern wireless networks, where devices are located predominantly on the streets of large cities, see Figure 1. This discrepancy has motivated research in the direction of *Cox percolation*, which can be thought of as a model of continuum percolation where the node distribution is governed by a random environment [9]. Another option is to allow the connection radii to be governed by the random environment [1].

Motivated by the need to extend the basic findings from percolation theory to random environments, [9, Theorems 2.4, 2.6] developed general conditions on the random environment ensuring a non-trivial phase transition. However, any progress beyond these findings was limited by the fact that the vast majority of environments relevant for applications do not satisfy the FKG inequality, which is a basic building block in percolation theory.

Recently, it has been discovered [5] that the OSSS-inequality (O’Donnell, Saks, Schramm & Servedio) is a powerful tool for analyzing percolation systems with complex spatial correlations [10, 16]. Since the OSSS inequality does not rely on the FKG inequality, it is particularly attractive for analyzing Cox percolation.

Our main result shows how to apply the OSSS inequality to establish the sharpness of the phase transition in a Cox percolation model built on a random environment subject to a set of general conditions. The main features that we require from the environment are a ‘factor of iid’ representation with finite range of dependence, and uniform local boundedness of the node intensity. We always work in the annealed model where the percolation probabilities average both over the environment and the particle placement. As a prototypical illustration of the general methodology, we apply our results to properly modified models. Here, before building the Delaunay triangulation, we first superimpose the Poisson process of vertices with a sparse grid and we impose an upper bound on the spatial device intensity.

As announced, the key tool to establish the main result is the OSSS inequality. However, in the context of Cox percolation, the standard approaches from the literature do not apply immediately since there are two sources of randomness: (i) the random configuration of nodes, and (ii) the random environment. Applying the OSSS inequality leads to a variance bound involving influences with respect to both sources of randomness. In contrast, when applying Russo’s formula for the derivative of the percolation probability, only influences with respect to random node locations appear. In order to convert one type of influence to another, we will rely on a coarse-graining strategy. This approach is broadly similar to arguments provided by [4, 18] for different discrete and continuum percolation models. The general idea is that, when comparing pivotality with respect to different parameters, it can be substantially simpler to work with coarse-grained pivotality rather than with point-wise pivotality. In [18], the pivotality with respect to the percolation parameter is compared to the pivotality with respect to a parameter controlling the range of the model. Since [18] deals with unbounded range models, the individual steps are delicate and the arguments need to be very quantitative. In

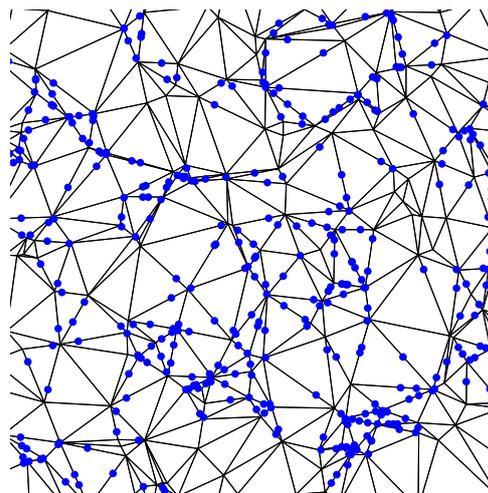


Figure 1: Devices (blue) scattered at random on the edges of a Poisson–Delaunay triangulation.

contrast, since we work only on a single scale, our proof is much more concise and we do not need to invoke the FKG inequality. We also note that [3] considers percolation in non-FKG but finite-range models. However, while one of the major innovations of [3] is a Russo-type inequality used in the Gaussian setting, there is no need to do any comparison of pivotals.

The rest of the manuscript is organized as follows. In Section 2, we define a general framework for ‘factor of iid’ representations of Cox processes, and provide Delaunay-based examples that are covered by this framework. Next, we state the main results, namely Proposition 3.2 and Theorem 3.3 on the non-triviality and the sharpness of the phase transition. Finally, these results are proven in Sections 4 and 5 respectively.

2 Factor of iid representations of Cox point processes

Recall that a *Cox point process* is a Poisson point process whose intensity measure (‘environment’) is a random Borel measure on \mathbb{R}^d . In this section, we define a general class of Cox point processes in which both, the environment and the Poisson points, may be ‘locally’ constructed from (independent) iid processes in the background. This representation will be crucial in the proof of our main results, where we will also assume stronger properties such as finite-range dependence and uniform boundedness.

The framework is broad enough to cover environments that are absolutely continuous with respect to the Lebesgue measure, e.g., supported on a random closed set of full dimension (see Remark 2.2), and also singular environments, e.g., supported on lower-dimensional structures such as hyperplanes or line-segments. In Section 2.2 below, we give examples of the both types motivated by wireless communications.

2.1 Definition

First we introduce a class \mathfrak{F} of environments which possess a ‘factor of iid’ representation. We fix a large scale $M \geq 1$ and a fine scale b such that b^{-1} is an integer exceeding $2\sqrt{db}M$; the scale M will later encode the dependency range of the environment, whereas the environment will be constructed on the scale Mb . Then, a random Borel measure \mathcal{E} on \mathbb{R}^d is in class \mathfrak{F} if it is of the form

$$\mathcal{E}(\cdot) := \sum_{x \in b\mathbb{Z}^d} \mathcal{E}_x(\cdot; Y), \quad \mathcal{E}_x(\cdot; Y) := \mathcal{U}_x(Y) \mathcal{Q}_x(\cdot; Y),$$

where

- (i) $Y := \{Y_z\}_{z \in \mathbb{Z}^d}$ is an iid family of random elements taking values in some measurable space,
- (ii) $\mathcal{U}_x := \mathcal{U}_x(Y)$ is a random variable with values in $[0, \infty)$,
- (iii) $\mathcal{Q}_x(\cdot) := \mathcal{Q}_x(\cdot; Y)$ is a random probability measure in the cube $Q(x; b, M) := Mx + [0, Mb]^d$, and
- (iv) $\mathcal{Q}_x(\cdot; Y)$ and $\mathcal{U}_x(Y)$ are translation covariant, i.e., $\mathcal{Q}_{x+z_0}(A + Mz_0; \{Y_{z+z_0}\}_{z \in \mathbb{Z}^d}) = \mathcal{Q}_x(A; \{Y_z\}_{z \in \mathbb{Z}^d})$ for any $z_0 \in \mathbb{Z}^d$, $x \in b\mathbb{Z}^d$ and measurable event $A \subset Q(x; b, M)$, and similarly for \mathcal{U}_x .

The intuition behind this construction is that the randomness of the environment enters via the iid field Y . Then, \mathcal{Q}_x and \mathcal{U}_x encode respectively the distribution of the locations and the mean number of particles in the cube $Q(x; b, M)$.

For an environment $\mathcal{E} \in \mathfrak{F}$ and intensity parameter $\lambda > 0$, we next construct a Cox point process X^s with intensity $\lambda\mathcal{E}$. To this end we fix a bimeasurable bijection $\Gamma: [0, Mb]^d \rightarrow [0, 1]$, and note that the push-forward $\Gamma_*(\mathcal{Q}_x(\cdot; Y))$ is a random probability

measure on $[0, 1]$ whose cumulative distribution function will henceforth be denoted by $\Phi_x(t) := \Phi_x(t; Y)$. Next, we let

$$V := \{V_x\}_{x \in b\mathbb{Z}^d} := \{(V_{x,i}, U_{x,i})_{i \geq 1}\}_{x \in b\mathbb{Z}^d} \tag{2.1}$$

be a family of iid homogeneous Poisson point processes on $[0, 1] \times [0, \infty)$ with intensity λ . Then, we define the point process

$$X^s := \Psi(V, Y) := \bigcup_{x \in b\mathbb{Z}^d} \Psi_x(V_x; Y), \tag{2.2}$$

where the configuration of X^s inside $Q(x; b, M)$ is given as

$$\Psi_x(V_x; Y) := \{Mx + \Gamma^{-1}(\Phi_x^{-1}(V_{x,i}; Y)) : U_{x,i} \leq \mathcal{U}_x(Y)\}. \tag{2.3}$$

That is, $X^s \cap Q(x; b, M)$ consists of all shifted and transformed points $\Gamma^{-1}(\Phi_x^{-1}(V_{x,i}; Y))$ with mark $U_{x,i}$ at most $\mathcal{U}_x(Y)$. Here, $\Phi_x^{-1}(v) := \inf_{y \geq 0} \{\Phi_x(y) \geq v\}$ denotes the inverse distribution function.

One observes that, conditionally on Y , $X^s \cap Q(x; b, M)$ is a Poisson point process with intensity $\lambda \mathcal{U}_x$, and hence X^s is a Cox process with intensity $\lambda \mathcal{E}$. Moreover, conditionally on Y , the processes $X^s \cap Q(x; b, M)$ are independent for different $x \in b\mathbb{Z}^d$.

Remark 2.1 (Stationarity). By construction, the Cox process X^s defined above is invariant under translations of the lattice $M\mathbb{Z}^d$. In certain examples (see Section 2.2) it may also be invariant under translations of \mathbb{R}^d .

Remark 2.2 (Full-dimensional environments). Encoding the Cox point process via the bimeasurable bijection Γ is a bit cumbersome from a technical perspective but it allows us to cover both non-singular and singular environments simultaneously. In the former case, it may be more natural to rely on simpler alternative constructions. For instance, for environments that are uniformly distributed on a random closed subset of full dimension (as in the example in Section 2.2.3 below), one could define $\mathcal{E}_x(Y)$ to be the intersection of this random set with $Q(x; b, M)$, and then let $X^s \cap Q(x; b, M)$ be given by $X_x \cap \mathcal{E}_x(Y)$, where X_x is a homogeneous Poisson point process in $Q(x; b, M)$ with intensity $\lambda > 0$.

2.2 Examples

We next present examples motivated by wireless communications networks; in these examples X^s encodes the location of devices placed on a street system embedded in \mathbb{R}^2 .

2.2.1 Delaunay network

In our first example, devices are placed uniformly on the edges of a Delaunay triangulation formed from an underlying independent Poisson point process. More precisely, let \mathcal{P} be a homogeneous Poisson point process on \mathbb{R}^2 with intensity $\lambda_{\text{Del}} > 0$, and let $\text{Del}(\mathcal{P})$ denote the Delaunay triangulation with vertices given by \mathcal{P} . Then, for an intensity parameter $\lambda > 0$, we consider a Cox point process X^s with intensity $\lambda \mathcal{E} := \lambda \text{Leb}(\text{Del}(\mathcal{P}))$, where $\text{Leb}(\text{Del}(\mathcal{P}))$ denotes the 1D-Lebesgue measure on $\text{Del}(\mathcal{P})$.

It is easy to see that $\mathcal{E} \in \mathfrak{F}$, and so one can construct X^s using the general framework introduced above. More precisely, fix $M \geq 1$ and $b \leq 1/(2\sqrt{2}M)$ such that b^{-1} is an integer, and for $z \in \mathbb{Z}^2$ define $Y_z := [0, Mb]^2 \cap (\mathcal{P} - Mbz)$, so that $Y := \{Y_z\}_{z \in \mathbb{Z}^2}$ is an iid family of Poisson point processes and $\mathcal{P} = \bigcup_{z \in \mathbb{Z}^2} (Y_z + Mbz)$. Then, $\mathcal{Q}_x(\cdot; Y)$ and $\mathcal{U}_x(Y)$ are respectively defined as the uniform distribution on $\text{Del}(\mathcal{P}) \cap Q(x; b, M)$ and the total length of $\text{Del}(\mathcal{P}) \cap Q(x; b, M)$. Further, X^s may be defined as in (2.2).

While this is a natural model for a wireless network, it has two major drawbacks:

- (a) The network has infinite range of dependence.
- (b) Since we view the streets as having width 0, the environment is singular, and there is no deterministic upper bound for the total intensity of X^s in a finite sampling window.

This motivates us to introduce the following variants of the Delaunay network.

2.2.2 Delaunay network superimposed with sparse grid

To obtain a network with finite-range dependence, we superimpose the underlying Poisson point process with a sparse grid. More precisely, defining a large parameter $L \geq 1$, and setting $M = M'L$ for a positive integer M' , we let \mathcal{P} and Y be as before, but replace $\text{Del}(\mathcal{P})$ in the definition of \mathcal{Q} and \mathcal{U} with $\text{Del}_L(\mathcal{P}) := \text{Del}(\mathcal{P} \cup LZ^2)$. With this change, we then construct X^s in an identical manner.

Since every triangle in the Delaunay tessellation $\text{Del}_L(\mathcal{P})$ has diameter at most cL for a suitable $c \geq 1$, by choosing M' sufficiently large we ensure that the construction of the environment is 1-dependent (on the scale $M = M'L$), i.e., setting $I^+(z) := z + \{-1, 0, 1\}^2$ we have for all $z \in \mathbb{Z}^2$ and $x \in I_b(z)$ that $\mathcal{Q}_x(\{Y_{z'}\}_{z' \in \mathbb{Z}^2}) = \mathcal{Q}_x(\{Y'_{z'}\}_{z' \in \mathbb{Z}^2})$ if $Y_{z''} = Y'_{z''}$ for every $z'' \in I^+(z)$.

2.2.3 Delaunay network of edges with positive width

To obtain a model with bounded intensity, one option is to consider the streets as having a non-zero ‘thickness’ $w_0 > 0$. Precisely, let \mathcal{P} , Y and L be as before, and define the random closed set $\text{Del}_L^{w_0}(\mathcal{P}) := \{x \in \mathbb{R}^2 : \text{dist}(x, \text{Del}_L(\mathcal{P})) \leq w_0\}$. Now, replace $\text{Del}(\mathcal{P})$ in the definition of \mathcal{Q} and \mathcal{U} with $\text{Del}_L^{w_0}(\mathcal{P})$. Then, the intensity of X^s is bounded by λ .

2.2.4 Delaunay network of edges with capped density

A second option to obtain a model with bounded intensity, but which retains the singular street structure, is to enforce a cap on the total intensity in any cube $Q(x; b, M)$. More precisely, let \mathcal{P} , Y and L be as before, and let $\rho > 0$ be a parameter. Then, define \mathcal{Q}_x to be the uniform distribution on $\text{Del}_L(\mathcal{P}) \cap Q(x; b, M)$, and define $\mathcal{U}_x := \rho \wedge |\text{Del}_L(\mathcal{P}) \cap Q(x; b, M)|$, i.e., we first measure the edge length of $\text{Del}_L(\mathcal{P})$ in $Q(x; b, M)$ and then manually cap the resulting length at $\rho > 0$. By construction, the total intensity of X^s in the cube $Q(x; b, M)$ is bounded by $\lambda\rho$.

Remark 2.3. The Delaunay network in Section 2.2.1 is stationary with respect to translations in \mathbb{R}^2 and hence also the associated measures $\mathcal{Q}_x(\cdot; Y)$ and $\mathcal{U}_x(Y)$ are translation covariant. However, after introducing the sparse grid LZ^2 , the model is only stationary with respect to shifts in LZ^2 . One option to enforce the \mathbb{R}^2 -stationarity could be to replace LZ^2 with $L(\mathcal{V} + \mathbb{Z}^2)$, where \mathcal{V} is uniformly distributed on $[0, 1]^2$. It could also be interesting to work with a finite-range model that is intrinsically \mathbb{R}^2 -stationary.

Let us also mention that the Delaunay networks mentioned in this section can be also defined in higher dimensions, giving rise to networks of hyperplanes.

3 Main result

In this section, we state our main result (Theorem 3.3) on the sharpness of the phase transition for Cox percolation models, i.e., the continuum percolation model built from a Cox point process X^s .

First, let us be precise about the definition of the Cox percolation model. There are two equivalent ways to proceed: we can view the model as the subset of \mathbb{R}^d formed by

the union of balls of radius $1/2$ centred at the points X^s ; equivalently, we can consider the random geometric graph whose vertex set is X^s and whose edge-set contains all pairs of points in X^s at mutual distance less than 1. For concreteness we will work with the former definition. We say that the model *percolates* if it has an infinite connected component.

3.1 Conditions on the environment

In order to state our result, we introduce a set of conditions on the environment.

We assume that X^s is defined as in Section 2 for an intensity parameter $\lambda > 0$ and an environment $\mathcal{E} = \mathcal{E}(Y) \in \mathfrak{F}$. This implies in particular that \mathcal{E} (and correspondingly X^s) is invariant under translations by $M\mathbb{Z}^d$. Moreover, we assume the dependence of \mathcal{E} on Y is finite-range, and that \mathcal{E} has uniformly bounded local intensity. Precisely, define $I_b(z) := b\mathbb{Z}^d \cap (z + [0, 1)^d)$ and recall that $I^+(z) = z + \{-1, 0, 1\}^d$, we assume

- (i) **1-dependence**, i.e., for all $z \in \mathbb{Z}^d$ and $x \in I_b(z)$ we have that $\mathcal{E}_x(\{Y_{z'}\}_{z' \in \mathbb{Z}^d}) = \mathcal{E}_x(\{Y_{z'}\}_{z' \in \mathbb{Z}^d})$ if $Y_{z''} = Y_{z'}$ for every $z'' \in I^+(z)$, and
- (ii) **uniformly bounded local intensity**, i.e., there exists a $\rho > 0$ such that $\mathcal{U}_x \leq \rho$ for every $x \in b\mathbb{Z}^d$.

We note that the models considered in [6, 4, 18] allow for a finite-range decomposition. Hence, instead of imposing 1-dependence, it was enough there to require a sufficiently fast decay of dependence. However, the lack of the FKG inequality in our setting prevents us from being able to compare pivotality at different scales, which is used in [6, 4, 18] to handle such models.

Next, we impose some natural connectivity conditions on the environment \mathcal{E} . Observe that, if the $1/2$ -neighborhood of the support of \mathcal{E} does not percolate, then, there is also no chance of percolation in the Cox model. To make this precise, we say that a site $x \in b\mathbb{Z}^d$ is *non-empty* if $\mathcal{U}_x(Y) > 0$. Analogously, we say that $z \in \mathbb{Z}^d$ is non-empty if x is non-empty for some $x \in I_b(z)$; this defines a finitely dependent site percolation model on \mathbb{Z}^d . We recall from [12, Theorem 0.0] that there exists $q_0(d) \in (0, 1)$ such that any 6-dependent site percolation model on \mathbb{Z}^d with marginal probability at least $q_0(d)$ percolates, where the 6-dependence is chosen for convenience of the proofs. Hence we assume in the following

- (iii) **coverage**, i.e., it holds that $\mathbb{P}(o \in \mathbb{Z}^d \text{ is non-empty}) > q_0(d)$.

Finally, we also need a more delicate connectivity condition that is described by the notion of *essential connectedness* [2]. To make this precise, two sites $x, x' \in b\mathbb{Z}^d$ are *adjacent* if they are at d_∞ -distance b . For $\eta > 0$, we say that $x \in b\mathbb{Z}^d$ is η -*supported* if $\mathcal{U}_x \geq \eta$. Introduce the enlarged cubes $I_b^+(z) := b\mathbb{Z}^d \cap (z + [-1, 2)^d)$ and $I_b^{++}(z) := b\mathbb{Z}^d \cap (z + [-2, 3)^d)$. Then we assume

- (iv) **essential connectedness**, i.e., that there exists $\eta > 0$ such that with probability 1, for all $z \in \mathbb{Z}^d$, any non-empty $x, x' \in I_b^+(z)$ are connected by a chain of adjacent η -supported sites in $I_b^{++}(z)$.

Remark 3.1 (Examples). All the above conditions are satisfied for the Delaunay-based examples in Section 2.2.3 and 2.2.4 by choosing M of the form $M = M'L$, for M' sufficiently large. On the other hand, Conditions (i), (ii) and (iv) fail for the example in Section 2.2.1, and Condition (ii) fails for the example in Section 2.2.2.

3.2 Statement of the main result

Write \mathbb{P}_λ for the law of X^s with parameter λ . For $A, B \subset \mathbb{R}^d$ we write $\{A \rightsquigarrow B \text{ in } X^s\}$ for the event that A and B are connected in the corresponding Cox percolation model,

i.e., there exists a path between A and B in the set $\bigcup_{x \in X^s} \{x + B(1/2)\}$, where $B(r)$ is the Euclidean ball of radius r . Then, defining $\Lambda_M = [-M, M]^d$, we let $\lambda_c := \inf\{\lambda > 0: \theta(\lambda) > 0\}$ be the critical intensity for percolation, where for $\partial\Lambda_n := \Lambda_{n-M} \setminus \Lambda_{n-2M}$ we let

$$\theta(\lambda) := \mathbb{P}_\lambda(\Lambda_{3M} \leftrightarrow \infty \text{ in } X^s) := \lim_{n \rightarrow \infty} \mathbb{P}_\lambda(\Lambda_{3M} \leftrightarrow \partial\Lambda_n \text{ in } X^s)$$

denote the probability of the percolation event.

Note that in classical continuum percolation, $\theta(\lambda)$ is defined as $\mathbb{P}_\lambda(0 \leftrightarrow \infty \text{ in } X^s)$. In the M -discretized setting the addition of a point at the origin is no longer natural, and hence we rely on a definition of $\theta(\lambda)$ that is better adapted to the M -discretized model.

As a preliminary step, we verify that Cox percolation exhibits a non-trivial phase transition:

Proposition 3.2 (Non-triviality). *It holds that $0 < \lambda_c < \infty$.*

Let us mention that Proposition 3.2 can be established for more general models under significantly weaker assumptions on the environment, see [9].

The main result of the paper states that this phase transition is sharp, i.e., there is exponential decay of connectivity in the subcritical phase, and the percolation probability grows at least linearly in the supercritical phase.

Theorem 3.3 (Sharpness). *The phase transition for Cox percolation is sharp:*

- (i) $\limsup_{n \uparrow \infty} n^{-1} \log \mathbb{P}_\lambda(\Lambda_{3M} \leftrightarrow \partial\Lambda_n \text{ in } X^s) < 0$ holds for every $\lambda < \lambda_c$, and
- (ii) $\liminf_{\lambda \downarrow \lambda_c} \theta(\lambda)/(\lambda - \lambda_c) > 0$.

This generalizes a known result for the standard continuum percolation model [14, 15, 19], i.e., the homogeneous case in which X^s is a Poisson point process.

3.3 Possible extensions

3.3.1 The Delaunay network

As mentioned, the Delaunay network model in Section 2.2.1 does not satisfy the finite-range and uniform boundedness conditions, and so our result does not apply. It would be interesting to weaken these conditions so as to cover this model, but it would require new ideas. Indeed, while in [4, 18] a sprinkling trick is used to deal with models with fast decay of dependence, the lack of the FKG inequality does not allow to adapt these arguments in the present setting. Moreover, it is not clear how to define a finite-range approximation such that the essential connectedness remains valid.

3.3.2 Random connectivity radii

A natural generalization of the Cox percolation model would be to equip each point in X^s with a ball whose radius is drawn independently from a certain radius distribution, rather than a ball of radius $1/2$. If the radius distribution is bounded, the proof of Theorem 3.3 works unchanged, but it does not if the distribution is unbounded. In the homogeneous case where X^s is a Poisson point process, an extension to unbounded radii was achieved in [6], and it would be interesting to generalize this to Cox percolation.

3.3.3 Varying connectivity radius

For applications it may also be important to vary the radius of the connectivity instead of the intensity parameter (i.e., we fix λ but connect all points in X^s within distance μ , for varying $\mu > 0$). Unlike in the homogeneous case, these are not equivalent up to global rescaling. The proof from Theorem 3.3 breaks down in this case because of a lack of a

Russo formula expressing the derivative of percolation probabilities as integrals over pivotal intensities (see (5.3)). However it may be possible to define an alternate notion of pivotal intensity wrt increasing the connectivity radius, and compare this to other relevant notations of influence.

4 Non-triviality: Proof of Proposition 3.2

To establish the non-triviality of the percolation phase transition, we build on classical techniques to deal with finite-range dependent percolation processes, most notably the stochastic domination criterion [12, Theorem 0.0]. Nevertheless, to make the manuscript more self-contained, we include some details.

Proof of Proposition 3.2, $\lambda_c > 0$. Recall the family $\{V_x\}_{x \in b\mathbb{Z}^d}$ of iid Poisson point processes on $[0, 1] \times [0, \infty)$ with intensity λ from (2.1). Recall also the uniformly bounded intensity condition, which implies that we can and will restrict each process V_x to $[0, 1] \times [0, \rho]$ without change to the Cox point process X^s .

Call a site $z \in \mathbb{Z}^d$ *bad* if $V_x \neq \emptyset$ for some $x \in I_b(z)$, and observe that the probability that a site is bad tends to 0 as $\lambda \rightarrow 0$. Moreover, Cox percolation of X^s implies the percolation of the Bernoulli site percolation process of bad sites in \mathbb{Z}^d . Then the claim follows since Bernoulli site percolation has a subcritical regime [8, Theorem 1.10]. \square

Proof of Proposition 3.2, $\lambda_c < \infty$. Call a site $x \in b\mathbb{Z}^d$ *populated* if $X^s \cap Q(x; b, M) \neq \emptyset$. Moreover, say that $z \in \mathbb{Z}^d$ is *good* if (i) some $x \in I_b(z)$ is populated and (ii) any populated $x, x' \in I_b^+(z)$ are connected by a chain of adjacent populated sites in $I_b^{++}(z)$. Since $2\sqrt{db}M \leq 1$, the percolation of good sites on \mathbb{Z}^d implies Cox percolation of X^s . Moreover, the good sites are 6-dependent, and the coverage and essential-connectedness conditions imply that the marginal probability exceeds $q_0(d)$ as $\lambda \rightarrow \infty$. Hence, an application of [12, Theorem 0.0] shows that for sufficiently large $\lambda > 0$, the good sites percolate. \square

5 Sharpness: Proof of Theorem 3.3

To prove the sharpness of the phase transition, our general strategy is to proceed in the vein of [5, 6] and rely on Russo’s formula and the OSSS technique to derive a key differential inequality (see Proposition 5.1). The particular challenges of Cox percolation are that (i) the model depends on an underlying environment, and that (ii) the FKG inequality does not necessarily hold.

5.1 The differential inequality

To reflect that the model inherently depends on the scale M , we define

$$\theta_n(\lambda) := \mathbb{P}(\Lambda_{3M} \rightsquigarrow \partial\Lambda_{Mn} \text{ in } X^s),$$

setting $\theta_i := 1$ for $i \leq 4$ to avoid ambiguities. We will deduce the sharpness of the phase transition as a consequence of the following differential inequality, see [6, Lemma 1.7].

Proposition 5.1 (Differential inequality). *Let $\lambda' > \lambda_c$. Then, there exists $c_{\text{Diff}} > 0$ such that for every $n \geq 1$ and $\lambda' > \lambda > \lambda_c$,*

$$\frac{d}{d\lambda} \theta_n(\lambda) \geq c_{\text{Diff}} \frac{n}{\sum_{s \leq n} \theta_s(\lambda)} \theta_n(\lambda) (1 - \theta_n(\lambda)).$$

Proof of Theorem 3.3. Once Proposition 5.1 is established, we argue as in [5] to show that there exists $\lambda_1 \in [\lambda_c, \lambda']$ such that $\limsup_{n \uparrow \infty} n^{-1} \log(\theta_n(\lambda)) < 0$ for every $\lambda < \lambda_1$, and $\liminf_{\lambda \downarrow \lambda_1} \theta(\lambda)/(\lambda - \lambda_1) > 0$. Therefore, $\lambda_c = \lambda_1$. \square

It remains to deduce Proposition 5.1. The key idea is to apply the OSSS inequality from [17] to the indicator of the event $\{\Lambda_{3M} \rightsquigarrow \partial\Lambda_{Mn} \text{ in } X^s\}$. The fact that the Cox point process X^s has a ‘factor of iid’ representation is crucial in implementing this strategy.

Recall the family $\{V_x\}_{x \in b\mathbb{Z}^d}$ of iid Poisson point processes from (2.1), and define the collection $X = \{X_z\}_{z \in \mathbb{Z}^d}$ with $X_z := \{V_x\}_{x \in I_b(z)}$. Then, we observe that $\theta_n(\lambda)$ may be considered as the expectation of a function $f_n(X, Y)$ of the discrete iid field $Z := \{Z_z\}_{z \in \mathbb{Z}^d} := \{(X_z, Y_z)\}_{z \in \mathbb{Z}^d}$. Applying the OSSS inequality to an algorithm T determining f_n gives that

$$\theta_n(\lambda)(1 - \theta_n(\lambda)) = \text{Var}(f_n(Z)) \leq \frac{1}{2} \sum_{z \in \mathbb{Z}^d} \delta_z(T) \text{Inf}_z(f_n), \tag{5.1}$$

where

- (i) $\delta_z(T) := \mathbb{P}(T \text{ reveals } Z_z)$ is the probability that the algorithm T reveals the value of Z_z , and
- (ii) $\text{Inf}_z(f_n) := \mathbb{P}(f_n(Z) \neq f_n(Z'(z)))$ denotes the (*resampling*) *influence* of Z_z , where $Z'(z)$ is formed from Z by replacing Z_z with an independent copy.

By the Efron–Stein inequality, which in our case becomes a union bound for two events,

$$\text{Inf}_z(f_n) \leq \text{Inf}_z^X(f_n) + \text{Inf}_z^Y(f_n),$$

where $\text{Inf}_z^X(f_n) := \mathbb{P}(f_n(X, Y) \neq f_n(X'(z), Y))$ and $\text{Inf}_z^Y(f_n) := \mathbb{P}(f_n(X, Y) \neq f_n(X, Y'(z)))$ denote the analogous (*resampling*) *influences* of X_z and Y_z respectively. Hence,

$$\theta_n(\lambda)(1 - \theta_n(\lambda)) \leq \frac{1}{2} \sum_{z \in \mathbb{Z}^d} \delta_z(T) (\text{Inf}_z^X(f_n) + \text{Inf}_z^Y(f_n)). \tag{5.2}$$

Our task is to relate the right-hand side of (5.2) to the derivative of $\theta_n(\lambda)$. For this we rely on Lemmas 5.2–5.4 below, whose proofs will be given at the end of the section.

Recall that the uniformly bounded intensity condition implies that we may restrict each Poisson point process V_x in (2.1) to $[0, 1] \times [0, \rho]$ without change to the Cox point process X^s . Then, since the event $\{\Lambda_{3M} \rightsquigarrow \partial\Lambda_{Mn} \text{ in } X^s\}$ is increasing, the infinitesimal Russo–Margulis formula (see e.g., [11, Theorem 19.1]) gives

$$\frac{d}{d\lambda} \theta_n(\lambda) = \lambda \sum_{x \in b\mathbb{Z}^d} \int_{[0,1] \times [0,\rho]} \text{Piv}_x(r, u) d(r, u), \tag{5.3}$$

with

$$\text{Piv}_x(r, u) := \mathbb{P}(f_n(V_{x,r,u}, Y) \neq f_n(X, Y)),$$

where $V_{x,r,u}$ is the collection $\{V_{x'}\}_{x'}$ determined by $V_{x'} = V_x$ if $x' \neq x$ and $V_x = V_x \cup \{(r, u)\}$. In words, $\text{Piv}_x(r, u)$ is the probability that the event $\{\Lambda_{3M} \rightsquigarrow \partial\Lambda_{Mn} \text{ in } X^s\}$ does not occur, but does occur after adding (r, u) to V_x . Very loosely speaking, formula (5.3) can be thought of as a way to interchange \mathbb{E} and the derivative.

Our first lemma states that one can control the integrated infinitesimal pivotal probabilities in (5.3) in terms of the corresponding discrete influences Inf_z^X :

Lemma 5.2 (Pivotality and influence). *Let $\lambda' > 0$. Then, there exists $c_{\text{Piv}} = c_{\text{Piv}}(\lambda') > 0$ such that for every $n \geq 1$, $\lambda \leq \lambda'$ and $z \in \mathbb{Z}^d$,*

$$\text{Inf}_z^X \leq c_{\text{Piv}} \lambda \sum_{x \in I_b(z)} \int_{[0,1] \times [0,\rho]} \text{Piv}_x(r, u) d(r, u).$$

Our second lemma states that we can bound Inf_z^Y in terms of $\text{Inf}_{z'}^X$ for $z' \in I^{++}(z) := z + \{0, \pm 1, \pm 2\}^d$.

Lemma 5.3 (Poisson- and environment influences). *There exists $c_{\text{Inf}} > 1$ such that for every $n \geq 1$, $\lambda > 0$ and $z \in \mathbb{Z}^d$,*

$$\text{Inf}_z^Y \leq c_{\text{Inf}} \sum_{z' \in I^{++}(z)} \text{Inf}_{z'}^X.$$

Our final lemma bounds the revelation probabilities of a suitable randomized exploration algorithm:

Lemma 5.4 (Revelation probabilities). *For every $n \geq 16$ there exists a randomized exploration algorithm T determining f_n such that for every $\lambda > 0$ and $z \in \mathbb{Z}^d$,*

$$\delta_z(T) \leq \frac{8}{n} \sum_{s \leq n} \theta_s(\lambda).$$

Let us use Lemmas 5.2–5.4 to conclude the proof of Proposition 5.1:

Proof of Proposition 5.1. First, Lemmas 5.3 and 5.4 show that

$$\frac{n}{\sum_{s \leq n} \theta_s(\lambda)} \sum_{z \in \Lambda_n \cap \mathbb{Z}^d} \delta_z(T) (\text{Inf}_z^X + \text{Inf}_z^Y) \leq 8 \sum_{z \in \Lambda_n \cap \mathbb{Z}^d} (\text{Inf}_z^X + \text{Inf}_z^Y) \leq 8(5^d + 1)c_{\text{Inf}} \sum_{z \in \Lambda_{n+2} \cap \mathbb{Z}^d} \text{Inf}_z^X.$$

Hence, invoking Lemma 5.2 together with the Russo–Margulis formula (5.3) concludes the proof. \square

5.2 Proof of the auxiliary lemmas

First, we establish Lemma 5.3, i.e., the domination of Inf_z^Y by a multiple of Inf_z^X , for which we adapt a coarse-graining strategy from [4, 18].

Proof of Lemma 5.3. We present a detailed proof in the case where $3 \leq |z|_\infty \leq n - 3$, noting that the arguments in the remaining cases are very similar. Let $Y'(z) = \{Y'_w\}_{w \in \mathbb{Z}^d}$ be the z -resampling of $Y = \{Y_w\}_{w \in \mathbb{Z}^d}$, i.e., the component of $\{Y_w\}_{w \in \mathbb{Z}^d}$ with index z is replaced by an independent copy. Note that, by the assumption of 1-dependence, the resampling does not modify the environment $\mathcal{E}(Y)$ outside $I^+(z)$. We introduce a ‘coarse grained’ version of the event $\{f_n(X, Y) \neq f_n(X, Y'(z))\}$, which depends on X only through the configuration of V_x for $x \notin I_b^{++}(z)$. First, we let $E_{\text{coarse}, -}$ denote the event that $\Lambda_{3M} \not\rightsquigarrow \partial\Lambda_{Mn}$ holds if $V_x \cap ([0, 1] \times [0, \rho]) = \emptyset$ for all $x \in I_b^{++}(z)$. In particular, $E_{\text{coarse}, -}$ occurs under the event $\{f_n(X, Y) \neq f_n(X, Y'(z))\}$. Next, we let $E_{\text{coarse}, +}$ denote the event that if $V_x \cap ([0, 1] \times [0, \eta]) \neq \emptyset$ for every $x \in I_b^{++}(z)$, then $\Lambda_{3M} \rightsquigarrow \partial\Lambda_{Mn}$. Now, setting $Q' := (Mz + [-M, 2M]^d)$, we note that, under the event $\{f_n(X, Y) \neq f_n(X, Y'(z))\}$, there are particles $X_i, X_j \in X^s \setminus Q'$ such that (i) X_i connects to Λ_{3M} outside Q' , and (ii) X_j connects to $\partial\Lambda_{Mn}$ outside Q' . Hence, we conclude from the essential-connectedness condition that also the event $E_{\text{coarse}, +}$ occurs under the event $\{f_n(X, Y) \neq f_n(X, Y'(z))\}$.

Defining $E_{\text{coarse}} := E_{\text{coarse}, -} \cap E_{\text{coarse}, +}$, we have shown that $\text{Inf}_z^Y \leq \mathbb{P}(E_{\text{coarse}})$, and so it suffices to show that $\mathbb{P}(E_{\text{coarse}}) \leq c_{\text{Inf}} \sum_{z' \in I^{++}(z)} \text{Inf}_{z'}^X$. To that end, we let $E_{\text{fine}, -} := \{V_x \cap ([0, 1] \times [0, \rho]) = \emptyset \text{ for all } x \in I_b^{++}(z)\}$. Similarly, we let $E_{\text{fine}, +}$ denote the event that $V_x \cap ([0, 1] \times [0, \eta]) \neq \emptyset$ for every $x \in I_b^{++}(z)$. Finally, we write X^* for the Poisson point process obtained by resampling $\{V_x\}_{x \in I_b^{++}(z)}$. Then, by the independence property of

Poisson point processes,

$$\begin{aligned} & \mathbb{P}(f_n(X, Y) \neq f_n(X^*, Y)) \\ & \geq \mathbb{P}(\{V_x\}_{x \notin I_b^{++}(z)}, Y \in E_{\text{coarse}}, \{V_x\}_{x \in I_b^{++}(z)} \in E_{\text{fine}, -}, \{V_x^*\}_{x \in I_b^{++}(z)} \in E_{\text{fine}, +}) \\ & = \mathbb{P}(\{V_x\}_{x \notin I_b^{++}(z)}, Y \in E_{\text{coarse}}) \mathbb{P}(\{V_x\}_{x \in I_b^{++}(z)} \in E_{\text{fine}, -}) \mathbb{P}(\{V_x^*\}_{x \in I_b^{++}(z)} \in E_{\text{fine}, +}), \end{aligned}$$

where the second and third factor in this product are bounded away from 0. Finally, by the Efron–Stein inequality, we see that $\mathbb{P}(f_n(X, Y) \neq f_n(X^*, Y)) \leq \sum_{z' \in I^{++}(z)} \text{Inf}_{z'}^X$, thereby concluding the proof. \square

Next, we prove Lemma 5.2 through a short computation using tail estimates for Poisson random variables:

Proof of Lemma 5.2. We condition on Y and note that by the superposition theorem ([11, Theorem 3.3]) we may think of the collection $X_z = \{V_x\}_{x \in I_b(z)}$ as a homogeneous Poisson point process on $[0, 1] \times [0, \rho] \times I_b(z)$ with intensity λ . Hence, $X_z = \{P_i\}_{i \leq N}$ where N is a Poisson random variable with parameter $\lambda_* := \lambda \rho b^{-d}$ and the $\{P_i\}_{i \geq 1}$ are iid uniform on $[0, 1] \times [0, \rho] \times I_b(z)$. Moreover, writing $X' := \{X_{z'}\}_{z' \neq z}$, we let

$$K := \sup \{k \geq 0 : \Lambda_{3M} \not\rightsquigarrow \partial \Lambda_{Mn} \text{ in } \Psi(\{P_i\}_{i \leq k} \cup X')\}$$

denote the maximum number of Poisson points that can be added to $I_b(z)$ such that the percolation event does not happen. Note that K may also take the values $-\infty$ or ∞ . The introduction of the quantity K has the advantage that the pivotal probabilities can be concisely represented via $\mathbb{P}(K = N) = b^d \rho^{-1} \sum_{x \in I_b(z)} \int_{[0,1] \times [0,\rho]} \text{Piv}_x(r, u) d(r, u)$. Indeed, expressing the uniform distribution of P_i on $[0, 1] \times [0, \rho] \times I_b(z)$ in integral form,

$$\begin{aligned} \mathbb{P}(K = N) &= b^d \rho^{-1} \sum_{x \in I_b(z)} \int_{[0,1] \times [0,\rho]} \mathbb{P}(f_n(V_{x,r,u}, Y) \neq f_n(X, Y)) d(r, u) \\ &= b^d \rho^{-1} \sum_{x \in I_b(z)} \int_{[0,1] \times [0,\rho]} \text{Piv}_x(r, u) d(r, u). \end{aligned}$$

Similarly, we can bound the influences through $\text{Inf}_z^X \leq 2\mathbb{P}(N > K \geq 0)$. Using the tail probabilities of a Poisson random variable and that K and N are independent shows that for every $k \geq 0$,

$$\frac{\mathbb{P}(N > k)}{\lambda_* \mathbb{P}(N = k)} = \sum_{\ell \geq k+1} \frac{k!}{\ell!} \lambda_*^{\ell-k-1}.$$

Noting that the right-hand side is bounded by $\exp(\lambda_*)$ concludes the proof. \square

Remark 5.5. Another approach to proving Lemma 5.2 which may be less sensitive to the Poisson assumption would be to first discretize space and then to invoke the Efron–Stein inequality to aggregate the coordinates.

To finish, we prove Lemma 5.4, i.e., we describe the randomized algorithm leading to the asserted bound on the revelation probabilities. In essence, the proof can be adapted from previous results in the literature, e.g., [6, Lemma 3.3]. Nevertheless, to make the presentation self-contained, we provide a brief overview.

Proof of Lemma 5.4. For every $6 \leq m \leq n-3$ we construct an algorithm T^m determining $\{\Lambda_{3M} \rightsquigarrow \partial \Lambda_{Mn}\}$ as follows. During the algorithm, a site $z \in \mathbb{Z}^d$ is called *active* if it is revealed but the neighborhood $I^{++}(z)$ is not yet entirely revealed.

Sharp phase transition for Cox percolation

- (i) First, reveal Z_z for all z with $\|z\|_\infty - m \leq 3$. This determines the point configuration for all z with $\|z\|_\infty - m \leq 2$. Let S denote the union of all connected components intersecting $\partial\Lambda_{Mm}$.
- (ii) Pick an active z with $S \cap Q(x; b, M) \neq \emptyset$ for some $x \in Q(z, b)$, reveal $I^{++}(z)$, and grow the components from S with the particles from $X^s \cap Q(x'; b, M)$ for $x' \in I_b^+(z)$.
- (iii) Continue this exploration until there is no more active z with $S \cap (Mz + [0, M]^d) \neq \emptyset$, or a connection from Λ_{3M} to $\partial\Lambda_{Mn}$ is found.

Next, we note that the revealment probabilities $\delta_z(T^m)$ are bounded above by the percolation probabilities in the sense that

$$\mathbb{P}(T^m \text{ reveals } z) \leq \mathbb{P}((Mz + \Lambda_{3M}) \leftrightarrow \partial\Lambda_{Mm}) \leq \theta_{|m - \|z\|_\infty|}.$$

Thus, picking $m \in \{6, \dots, n - 3\}$ at random shows the asserted $\delta_z(T) \leq \frac{2}{n-8} \sum_{m \leq n} \theta_m \leq \frac{4}{n} \sum_{m \leq n} \theta_m$. \square

References

- [1] D. Ahlberg and J. Tykesson. Gilbert's disc model with geostatistical marking. *Adv. Appl. Probab.*, 50(4):1075–1094, 2018. MR3881110
- [2] D. J. Aldous. Which connected spatial networks on random points have linear route-lengths? *arXiv preprint arXiv:0911.5296*, 2009.
- [3] V. Dewan and S. Muirhead. Upper bounds on the one-arm exponent for dependent percolation models. *arXiv preprint arXiv:2102.12123*, 2021.
- [4] H. Duminil-Copin, S. Goswami, P.-F. Rodriguez, and F. Severo. Equality of critical parameters for percolation of Gaussian free field level-sets. *arXiv preprint arXiv:2002.07735; to appear in Duke Math. J.*, 2020.
- [5] H. Duminil-Copin, A. Raoufi, and V. Tassion. Sharp phase transition for the random-cluster and Potts models via decision trees. *Ann. of Math. (2)*, 189(1):75–99, 2019. MR3898174
- [6] H. Duminil-Copin, A. Raoufi, and V. Tassion. Subcritical phase of d -dimensional Poisson-Boolean percolation and its vacant set. *Ann. Henri Lebesgue*, 3:677–700, 2020. MR4149823
- [7] E. N. Gilbert. Random plane networks. *J. Soc. Indust. Appl. Math.*, 9:533–543, 1961. MR0132566
- [8] G. Grimmett. *Percolation*. Springer, Berlin, second edition, 1999. MR1707339
- [9] C. Hirsch, B. Jahnel, and E. Cali. Continuum percolation for Cox point processes. *Stochastic Process. Appl.*, 129(10):3941–3966, 2019. MR3997667
- [10] G. Last, G. Peccati, and D. Yogeshwaran. Phase transitions and noise sensitivity on the Poisson space via stopping sets and decision trees. *arXiv preprint arXiv:2101.07180*, 2021.
- [11] G. Last and M. D. Penrose. *Lectures on the Poisson Process*. Cambridge University Press, Cambridge, 2016. MR3791470
- [12] T. M. Liggett, R. H. Schonmann, and A. M. Stacey. Domination by product measures. *Ann. Probab.*, 25(1):71–95, 1997. MR1428500
- [13] R. Meester and R. Roy. *Continuum Percolation*. Cambridge University Press, Cambridge, 1996. MR1409145
- [14] R. Meester, R. Roy, and A. Sarkar. Nonuniversality and continuity of the critical covered volume fraction in continuum percolation. *J. Statist. Phys.*, 75(1-2):123–134, 1994. MR1273055
- [15] M. V. Men'shikov. Coincidence of critical points in percolation problems. *Dokl. Akad. Nauk SSSR*, 288(6):1308–1311, 1986. MR0852458
- [16] S. Muirhead and H. Vanneuville. The sharp phase transition for level set percolation of smooth planar Gaussian fields. *Ann. Inst. Henri Poincaré Probab. Stat.*, 56(2):1358–1390, 2020. MR4076787

Sharp phase transition for Cox percolation

- [17] R. O’Donnell, M. Saks, O. Schramm, and R. A. Servedio. Every decision tree has an influential variable. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’05)*, pages 31–39. IEEE, 2005.
- [18] F. Severo. Sharp phase transition for Gaussian percolation in all dimensions. *arXiv preprint arXiv:2105.05219*, 2021.
- [19] S. A. Zuev and A. Sidorenko. Continuous models of percolation theory. I. *Theor. Math. Phys.*, 62(1):51–58, 1985. MR0782099

Acknowledgments. We thank the anonymous referee for the thorough reading of the manuscript. The remarks and suggestions helped to improve the quality of the material substantially.

Electronic Journal of Probability

Electronic Communications in Probability

Advantages of publishing in EJP-ECP

- Very high standards
- Free for authors, free for readers
- Quick publication (no backlog)
- Secure publication (LOCKSS¹)
- Easy interface (EJMS²)

Economical model of EJP-ECP

- Non profit, sponsored by IMS³, BS⁴, ProjectEuclid⁵
- Purely electronic

Help keep the journal free and vigorous

- Donate to the IMS open access fund⁶ (click here to donate!)
- Submit your best articles to EJP-ECP
- Choose EJP-ECP over for-profit journals

¹LOCKSS: Lots of Copies Keep Stuff Safe <http://www.lockss.org/>

²EJMS: Electronic Journal Management System: <https://vtex.lt/services/ejms-peer-review/>

³IMS: Institute of Mathematical Statistics <http://www.imstat.org/>

⁴BS: Bernoulli Society <http://www.bernoulli-society.org/>

⁵Project Euclid: <https://projecteuclid.org/>

⁶IMS Open Access Fund: <https://imstat.org/shop/donation/>