

Bayesian Models of Graphs, Arrays and Other Exchangeable Random Structures

PETER ORBANZ AND DANIEL M. ROY

Abstract. The natural habitat of most Bayesian methods is data represented by exchangeable sequences of observations, for which de Finetti’s theorem provides the theoretical foundation. Dirichlet process clustering, Gaussian process regression, and many other parametric and nonparametric Bayesian models fall within the remit of this framework; many problems arising in modern data analysis do not. This expository paper provides an introduction to Bayesian models of graphs, matrices, and other data that can be modeled by random structures. We describe results in probability theory that generalize de Finetti’s theorem to such data and discuss the relevance of these results to nonparametric Bayesian modeling. With the basic ideas in place, we survey example models available in the literature; applications of such models include collaborative filtering, link prediction, and graph and network analysis. We also highlight connections to recent developments in graph theory and probability, and sketch the more general mathematical foundation of Bayesian methods for other types of data beyond sequences and arrays.

1. Introduction. For data represented by exchangeable sequences, Bayesian nonparametrics has developed into a flexible and powerful toolbox of models and algorithms. Its modeling primitives—Dirichlet processes, Gaussian processes, etc.—are widely applied and well-understood, and can be used as components in hierarchical models [59] or dependent models [48] to address a wide variety of data analysis problems. One of the main challenges for Bayesian statistics and machine learning is arguably to extend this toolbox to the analysis of data sets with more interesting structure, such as graph, network, and relational data.

In this article, we consider structured data—sequences, graphs, trees, matrices, etc.—and ask:

What is the appropriate class of statistical models for a given type of structured data?

Representation theorems for exchangeable random structures lead us to an answer, and they do so in a very precise way: They characterize the class of possible Bayesian models for the given type of data, show how these models are parameterized, and even provide basic convergence guarantees. The probability literature provides such results for dozens of exchangeable random structures, including sequences, graphs, partitions, arrays, trees, etc. The purpose of this article is to explain how to interpret these results and how to translate them into a statistical modeling approach.

Overview

A statistical model is usually defined as a family \mathcal{P} of distributions on a sample space \mathbf{X} —the family of all Gaussians on $\mathbf{X} = \mathbb{R}$, say. The distributions P_θ in \mathcal{P} are indexed by a parameter θ (e.g., the mean and variance of the Gaussian). In a typical statistical inference problem, we would observe data X_1, \dots, X_n , and then try to determine which distribution in \mathcal{P} best explains the data. There are different ways of approaching this problem:

- One can seek an estimator (such as a maximum likelihood estimator) for the parameter.
- One can take a Bayesian approach, modeling θ with a random variable Θ and computing the conditional distribution of Θ given observed data (the posterior distribution).

However, the distribution of a sample X_1, \dots, X_n is a *joint* distribution on \mathbf{X}^n . Estimating the joint distribution is hopeless without further assumptions, since the entire sample constitutes just a single observation from this distribution.

Both approaches mentioned above resolve this problem by relying on an independence assumption—an i.i.d. assumption in the frequentist case or a conditional i.i.d. assumption (the observations are i.i.d. conditionally given the parameter value) in the Bayesian case. The independence property factorizes the joint distribution, which reduces the modeling problem to distributions on \mathbf{X} . In the Bayesian case, the assumption is justified if and only if the data can be represented as an exchangeable sequence—informally, if the order in which observations are recorded does not matter. This characterization is given by de Finetti’s theorem:

The joint distribution of an exchangeable sequence of random values in \mathbf{X} is characterized by a distribution of a random probability measure on \mathbf{X} .

If we assume a specific model \mathcal{P} , the random distribution mentioned in the theorem is an element P_Θ of \mathcal{P} , determined at random by the random parameter Θ .

Although this is clearly an important theoretical result, we would hardly regard de Finetti’s theorem as a modeling tool: Arguably, it revisits a modeling approach we would have regarded as natural anyway, and justifies it in hindsight. The situation is very different for structured data, where it is often far from obvious what a suitable class of models would be. Suppose, for example, that the sample is a large graph. Should we assume its edges are i.i.d. vari-

ables? In graphs, that implies a single one-parameter family of distributions, the Erdős-Rényi model, and is hence hardly a general modeling approach. If we assume the graph to be exchangeable—where we put off the precise definition for now—we obtain a much larger class of distributions. They are again characterized by an exchangeability theorem, but in this case, the implications for statistical models are rather more surprising than above:

The distribution of an exchangeable simple graph is characterized by a distribution on the space of functions from $[0, 1]^2$ to $[0, 1]$.

Hence, any specific function $w : [0, 1]^2 \rightarrow [0, 1]$ defines a distribution P_w on graphs (we will see in Section 3 how we can sample a graph from P_w).

For modeling purposes, this means that *any* statistical model of exchangeable graphs is a family of such distributions P_w . Density estimation in exchangeable graphs can therefore be formulated as a regression problem: It is equivalent to recovering the function w from data. Once again, we can choose a frequentist approach (define an estimator for w) or a Bayesian approach (define a prior distribution on a random function W); we can obtain nonparametric models by choosing infinite-dimensional subspaces of functions, or parametric models by keeping the dimension finite.

Since a graph can be regarded as a special type of matrix (the adjacency matrix), we can ask more generally for models of exchangeable matrices, and obtain a similar result:

The distribution of an exchangeable two-dimensional array is characterized by a distribution on the space of functions from $[0, 1]^3$ to $[0, 1]$.

There is a wide variety of random structures for which exchangeability can be defined; Table 1 lists some important examples. Borrowing language from [5], we collectively refer to such random objects as **exchangeable random structures**. This article explains representation theorems for exchangeable random structures and their implications for Bayesian statistics and machine learning. The overarching theme is that the key implications of de Finetti’s theorem can be generalized to many types of data, and that these results are directly applicable to the derivation and interpretation of statistical models.

Contents

TABLE 1
Exchangeable random structures

Random structure	Theorem of	Ergodic distributions \mathbf{p}_θ	Statistical application
Exchangeable sequences	de Finetti [22, 23] Hewitt and Savage [29]	product distributions	most Bayesian models [e.g. 57]
Processes with exchangeable increments	Bühlmann [17]	Lévy processes	
Exchangeable partitions	Kingman [39]	“paint-box” distributions	clustering
Exchangeable arrays	Aldous [2] Hoover [34] Kallenberg [35]	sampling schemes Eq. (6.4), Eq. (6.10)	graph-, matrix- and array-valued data (e.g., [31]); see Section 4
Block-exchangeable sequences	Diaconis and Freedman [19]	Markov chains	e.g. infinite HMMs [9, 24]

Section 2: reviews exchangeable random structures, their representation theorems, and the role of such theorems in Bayesian statistics.

Section 3: introduces the generalization of de Finetti’s theorem to models of graph- and matrix-valued data, the Aldous-Hoover theorem, and explains how Bayesian models of such data can be constructed.

Section 4: surveys models of graph- and relational data available in the machine learning and statistics literature. Using the Aldous-Hoover representation, models can be classified and some close connections emerge between models which seem, at first glance, only loosely related.

Section 5: describes recent development in the mathematical theory of graph limits. The results of this theory refine the Aldous-Hoover representation of graphs and provide a precise understanding of how graphs converge and how random graph models are parametrized.

Section 6: explains the general Aldous-Hoover representation for higher-order arrays.

Section 7: discusses sparse random structures and networks, why these models contradict exchangeability, and open questions arising from this contradiction.

Section 8: provides references for further reading.

2. Bayesian Models of Exchangeable Structures.

The fundamental Bayesian modeling paradigm based on exchangeable sequences can be extended to a very general approach, where data is represented by a random structure. Exchangeability properties are then used to deduce valid statistical models and useful parametrizations. This section sketches out the ideas underlying this approach, before we focus on graphs, matrices, and arrays in Section 3.

2.1. Basic example: Exchangeable sequences. The simplest example of an exchangeable random structure is an exchangeable sequence. We use the customary shorthand notation $(x_i) := (x_1, x_2, \dots)$ for a sequence, and similarly (x_{ij}) for a matrix, etc. Suppose (X_i) is an infinite sequence of random variables in a sample space \mathbf{X} . We call (X_i) **exchangeable** if its joint distribution satisfies

$$\begin{aligned} \mathbb{P}(X_1 \in A_1, X_2 \in A_2, \dots) \\ = \mathbb{P}(X_{\pi(1)} \in A_1, X_{\pi(2)} \in A_2, \dots) \end{aligned} \quad (2.1)$$

for every permutation π of $\mathbb{N} := \{1, 2, \dots\}$ and every collection of sets A_1, A_2, \dots . Expressing distributional equalities this way is cumbersome, and we can write (2.1) more

concisely as

$$(X_1, X_2, \dots) \stackrel{d}{=} (X_{\pi(1)}, X_{\pi(2)}, \dots), \quad (2.2)$$

or even $(X_i) \stackrel{d}{=} (X_{\pi(i)})$, where the notation $Y \stackrel{d}{=} Z$ means that the random variables Y and Z have the same distribution. Informally, exchangeability means that the probability of observing a particular sequence does not depend on the order of the elements in the sequence.

If the elements of a sequence are exchangeable, de Finetti's representation theorem implies they are conditionally i.i.d. The conditional independence structure is represented by a **random probability measure**, a random variable with values in the set $\mathbf{M}(\mathbf{X})$ of probability distributions on \mathbf{X} .

THEOREM 2.1 (de Finetti). *Let (X_1, X_2, \dots) be an infinite sequence of random variables with values in a sample space \mathbf{X} . The sequence X_1, X_2, \dots is exchangeable if and only if there is a random probability measure Θ on \mathbf{X} such that the X_i are conditionally i.i.d. given Θ and*

$$\mathbb{P}(X_1 \in A_1, X_2 \in A_2, \dots) = \int_{\mathbf{M}(\mathbf{X})} \prod_{i=1}^{\infty} \theta(A_i) \nu(d\theta) \quad (2.3)$$

where ν is the distribution of Θ . \square

The integral on the right-hand side of (2.3) can be interpreted as a two-stage sampling procedure:

1. Sample $\Theta \sim \nu$, i.e., draw a probability distribution at random from the distribution ν .
2. Given Θ , sample observations X_i conditionally i.i.d. as

$$X_1, X_2, \dots | \Theta \sim_{\text{iid}} \Theta. \quad (2.4)$$

The theorem says that *any* exchangeable sequence can be sampled by such a two-stage procedure; the distribution of the sequence is determined by the choice of ν . The random measure Θ is called the **directing random measure** and is said to **direct** the exchangeable sequence X . Its distribution ν is called the **mixing measure** or **de Finetti measure**.

Statistical inference is only possible if the distribution of the data, or at least some of its properties, can be recovered from observations. For i.i.d. random variables, this is ensured by the law of large numbers. The proof of de Finetti's theorem also implies a law of large numbers for exchangeable sequences:

THEOREM 2.2. *If the sequence (X_i) is exchangeable, the empirical distributions*

$$\hat{S}_n(\cdot) := \frac{1}{n} \sum_{i=1}^n \delta_{X_i}(\cdot) \quad (2.5)$$

converge to Θ , in the sense that

$$\hat{S}_n(A) \rightarrow \Theta(A) \quad \text{as } n \rightarrow \infty \quad (2.6)$$

holds with probability 1 for every set A . \square

The two theorems have fundamental implications for Bayesian modeling. If we assume the data can be represented by (some finite prefix) of an exchangeable sequence, this implies without any further assumptions:

- Conditioned on a random probability measure Θ representing an unknown distribution θ , every sample X_1, X_2, \dots is i.i.d. with distribution Θ .
- Every exchangeable sequence model is characterized by a unique distribution ν on $\mathbf{M}(\mathbf{X})$.
- A statistical model can be taken to be some subset of $\mathbf{M}(\mathbf{X})$ rather than $\mathbf{M}(\mathbf{X}^\infty)$, which we would have to consider for a general random sequence.
- Statistical inference is possible in principle: With probability one, the empirical distributions \hat{S}_n converge to the distribution Θ generating the data, according to (2.6).

A modeling application might look like this: We consider data generated by a specific data source or measurement process, and assume that data generated by this source can be represented as an exchangeable sequence. The definition of exchangeability for an infinite sequence does not mean we have to observe an infinite number of data points to invoke de Finetti's theorem; rather, it expresses the assumption that samples of *any finite* size generated by the source would be exchangeable. Hence, exchangeability is an assumption on the data source, rather than the data.

According to de Finetti's theorem, the data can then be explained by the two-stage sampling procedure above, for *some* distribution ν on $\mathbf{M}(\mathbf{X})$. A Bayesian model is specified by choosing a specific distribution ν , the prior distribution. In this abstract formulation of the prior as a measure on $\mathbf{M}(\mathbf{X})$, the prior also determines the observation model, as the smallest set $\mathcal{P} \subset \mathbf{M}(\mathbf{X})$ on which ν concentrates all its mass—since Θ then takes values in \mathcal{P} , and the sequence (X_i) is generated by a distribution in \mathcal{P} with probability 1. If $\mathbf{X} = \mathbb{R}$, for example, we could choose ν to concentrate on the set of all Gaussian distributions on \mathbb{R} , and would obtain a Bayesian model with a Gaussian likelihood and prior ν .

Given observations X_1, \dots, X_n , we then compute the posterior distribution, by conditioning ν on the observations. Theorem 2.2 implies that, if the empirical measure converges asymptotically to a specific measure $\theta \in \mathbf{M}(\mathbf{X})$, the posterior converges to a point mass at θ . This result has to be interpreted very cautiously, however: It only holds for a sequence (X_i) which was actually generated from the measure ν we use as a prior. In other words, suppose someone generates (X_i) from a distribution ν_1 on $\mathbf{M}(\mathbf{X})$ by the two-stage sampling procedure above, without disclosing ν_1 to us. In the sampling procedure, the variable $\Theta \sim \nu_1$ assumes as its value a specific distribution θ_1 , from which the data is then generated independently. We model the observed sequence by choosing a prior ν_2 . The posterior under ν_2 still converges to a point mass, but there is *no guarantee* that it is a point mass at θ_1 , and (2.6) only holds if $\nu_2 = \nu_1$.

Thus, there are several important questions that ex-

changeability does not answer:

- The de Finetti theorem says that there is *some* prior which adequately represents the data, but provides no guidance regarding the choice of ν : Any probability measure ν on $\mathbf{M}(\mathbf{X})$ is the prior for some exchangeable sequence.
- Theorem 2.2 *only* guarantees convergence for sequences of random variables generated from the prior ν .
- Theorem 2.2 is a first-order result: It provides no information on how quickly the sequence converges. Results on convergence rates can only be obtained for more specific models; the set of all exchangeable distributions is too large and too complicated to obtain non-trivial statements.

Answers to these questions typically require further modeling assumptions.

2.2. The general form of exchangeability results. Many problems in machine learning and modern statistics involve data which is more naturally represented by a random structure that is not a sequence: often a graph, matrix, array, tree, partition, etc. is a better fit to the structure of the data. If it is possible to define a suitable notion of exchangeability, the main features of de Finetti’s theorem typically generalize. Although results differ in their details, there is a general pattern, which we sketch in this section before considering specific types of exchangeable structures.

The setup is as follows: The product space \mathbf{X}^∞ of infinite sequences is substituted by a suitable space \mathbf{X}_∞ of more general, infinite structures. An infinite random structure X_∞ is a random variable with values in \mathbf{X}_∞ . Each element of \mathbf{X}_∞ can be thought of as a representation of an infinitely large data set or “asymptotic” sample. An actual, finite sample of size n is modeled as a substructure X_n of X_∞ , such as a the length- n prefix of an infinite sequence or a n -vertex subgraph of an infinite graph.

The first step in identifying a notion of exchangeability is to specify what it means to permute components of a structure $x_\infty \in \mathbf{X}_\infty$. If x_∞ is an infinite matrix, for example, a very useful notion of exchangeability arises when one considers all permutations that exchange the ordering of rows/columns, rather than the ordering of individual entries. Exchangeability of a *random* structure X_∞ then means that the distribution of X_∞ is invariant under the specified family of permutations.

Once a specific exchangeable random structure X_∞ is defined, the next step is to invoke a representation theorem that generalizes de Finetti’s theorem to X_∞ . Probability theory provides such theorems for a range of random structures; see Table 1 for examples. A representation theorem can be interpreted as determining (1) a natural parameter space \mathbf{T} for exchangeable models on \mathbf{X}_∞ , and (2) a special family of distributions on \mathbf{X}_∞ , which are called the **ergodic distributions** or **ergodic measures**. Each element $\theta \in \mathbf{T}$ determines an ergodic distribution, and we

denote this distribution as \mathbf{p}_θ . The set of ergodic distributions is

$$\{\mathbf{p}_\theta : \theta \in \mathbf{T}\} \subset \mathbf{M}(\mathbf{X}_\infty). \quad (2.7)$$

The distribution of any exchangeable random structure X_∞ can then be represented as a mixture of these ergodic distributions,

$$\mathbb{P}(X_\infty \in \cdot) = \int_{\mathbf{T}} \mathbf{p}_\theta(\cdot) \nu(d\theta). \quad (2.8)$$

In the specific case of exchangeable sequences, (2.8) is precisely the integral representation (2.3) in de Finetti’s theorem, and the ergodic measures are the distributions of i.i.d. sequences, that is,

$$\mathbf{T} := \mathbf{M}(\mathbf{X}) \quad \text{and} \quad \mathbf{p}_\theta(\cdot) = \theta^\infty(\cdot). \quad (2.9)$$

For more general random structures, the ergodic measures are not usually product distributions, but they retain some key properties:

- They are particularly simple distributions on \mathbf{X}_∞ , and form a “small” subset of all exchangeable distributions.
- They have a conditional independence property, in the sense that a random structure X_∞ sampled from one of the ergodic distributions decomposes into conditionally independent components. In de Finetti’s theorem, these conditionally independent components are the elements X_i of the sequence.

As in the sequence case, the integral (2.8) in the general case represents a two-stage sampling scheme:

$$\begin{aligned} \Theta &\sim \nu \\ X_\infty \mid \Theta &\sim \mathbf{p}_\Theta. \end{aligned} \quad (2.10)$$

For Bayesian modeling, this means:

A Bayesian model for an exchangeable random structure X_∞ with representation (2.8) is characterized by a prior distribution on \mathbf{T} .

Suppose the prior ν concentrates on a subset $\mathcal{T} \subset \mathbf{T}$, that is, \mathcal{T} is the smallest subset to which the prior assigns probability 1. Then \mathcal{T} defines a subset

$$\mathcal{P} := \{\mathbf{p}_\theta : \theta \in \mathcal{T}\} \quad (2.11)$$

of ergodic measures. We thus have defined a Bayesian model on \mathbf{X}_∞ , with prior ν and observation model \mathcal{P} . In summary:

- \mathbf{T} is the natural parameter space for Bayesian models of X_∞ , and the prior distribution is a distribution on \mathbf{T} .
- The observation model \mathcal{P} is a subset of ergodic measures. An exchangeability theorem characterizing the ergodic measures therefore also characterizes the observation model, because each distribution \mathbf{p}_θ is completely determined by θ .

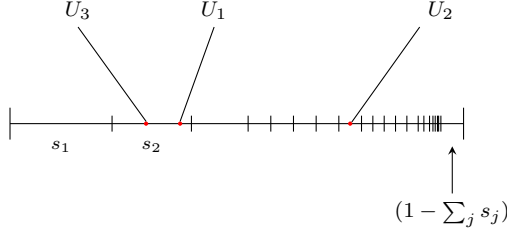


Fig 1: Sampling from a paint-box distribution with parameter $\mathbf{s} = (s_1, s_2, \dots)$. Two numbers i, j are assigned to the same block of the partition if the uniform variables U_i and U_j are contained in the same interval.

- If the set \mathcal{T} on which the prior concentrates its mass is a finite-dimensional subspace of \mathbf{T} , the resulting Bayesian model is **parameteric**. If \mathcal{T} has infinite dimension, the model is **nonparametric**.
- The representation (2.8) is typically complemented by a convergence results: A specific function of the samples converges to Θ almost surely as $n \rightarrow \infty$, generalizing Theorem 2.2. In particular, the parameter space \mathbf{T} can be interpreted as the set of all possible limit objects.

2.3. Exchangeable partitions. An illustrative example of an exchangeable random structure is an exchangeable partition. Bayesian nonparametric clustering models are based on such exchangeable random partitions. We again define the exchangeable structure as an infinite object: Suppose X_1, X_2, \dots is a sequence of observations. To encode a clustering solution, we have to specify which observations X_i belong to which cluster. To do so, it suffices to record which index i belongs to which cluster, and a clustering solution can hence be expressed as a partition $\pi = (b_1, b_2, \dots)$ of the index set \mathbb{N} . Each of the sets b_i , called **blocks**, is a finite or infinite subset of \mathbb{N} ; every element of \mathbb{N} is contained in exactly one block. An **exchangeable partition** is a random partition X_∞ of \mathbb{N} which is invariant under permutations of \mathbb{N} . Intuitively, this means the probability of a partition depends only on the relative sizes of its blocks, but not on which elements are in which block.

Kingman [39] showed that exchangeable random partitions can again be represented in the form of Eq. (2.8). The parameter space \mathbf{T} consists of all sequence $\theta := (s_1, s_2, \dots)$ of scalars $s_i \in [0, 1]$ which satisfy

$$s_1 \geq s_2 \geq \dots \quad \text{and} \quad \sum_i s_i \leq 1. \quad (2.12)$$

Let $\bar{s}_n := \sum_{i=1}^n s_i$. Then θ defines a partition of $[0, 1]$ into intervals

$$I_j := [\bar{s}_{j-1}, \bar{s}_j) \quad \text{and} \quad \bar{I} := (1 - \bar{s}_\infty, 1], \quad (2.13)$$

as shown in Fig. 1. Each ergodic distribution \mathbf{p}_θ is defined as the distribution of the following random partition of \mathbb{N} :

1. Generate $U_1, U_2, \dots \sim_{\text{iid}} \text{Uniform}[0, 1]$.

2. Assign $n \in \mathbb{N}$ to block b_j if $U_n \in I_j$. Assign every remaining element (those n with $U_n \in \bar{I}$) to its own block of size one.

Kingman called this distribution a **paint-box distribution**.

THEOREM 2.3 (Kingman). *Let X_∞ be random partition of \mathbb{N} .*

1. X_∞ is exchangeable if and only if

$$\mathbb{P}(X_\infty \in \cdot) = \int_{\mathbf{T}} \mathbf{p}_\theta(\cdot) \nu(d\theta), \quad (2.14)$$

for some distribution ν on \mathbf{T} , where \mathbf{p}_θ is the paint-box distribution with parameter $\theta \in \mathbf{T}$.

2. If X_∞ is exchangeable, the scalars s_i can be recovered asymptotically as limiting relative block sizes

$$s_i = \lim_{n \rightarrow \infty} \frac{|b_i \cap \{1, \dots, n\}|}{n}. \quad (2.15)$$

□

Part 1) is of course the counterpart to de Finetti's theorem, and part 2) corresponds to Theorem 2.2. In (2.15), we compute averages *within* a single random structure, having observed only a substructure of size n . Nonetheless, we can recover the parameter θ asymptotically from data. This is a direct consequence of exchangeability, and would not generally be true for an arbitrary random partition.

EXAMPLE 2.4 (Chinese restaurant process). A well-known example of a random partition is the Chinese restaurant process (CRP; see e.g. [30, 53] for details). The CRP is a one-parameter discrete-time stochastic process that induces a partition of \mathbb{N} . The parameter $\alpha > 0$ is called the *concentration*; different values of α correspond to different distributions $\mathbb{P}(X_\infty \in \cdot)$ in Eq. (2.14). If X_∞ is generated by a CRP, the paint-box parameter Θ is essentially the sequence of weights generated by the “stick-breaking” construction of the Dirichlet process [30]—with the difference that the elements of Θ are ordered by size, whereas stick-breaking weights are not. In other words, if X_∞ in (2.14) is a CRP, we can sample from ν by (1) sampling from a stick-breaking representation and (2) ordering the sticks by length. The lengths of the ordered sticks are precisely the scalars s_i in the theorem. ◁

2.4. “Non-exchangeable” data. Exchangeability seems at odds with many types of data; for example, a sequence of stock prices over time will be poorly modeled by an exchangeable sequence. Nonetheless, a Bayesian model of a time series will almost certainly imply an exchangeability assumption—the crucial question is which components of the overall model are assumed to be exchangeable.

EXAMPLE 2.5 (Discrete times series and random walks). Another important type of exchangeability property [19, 68] is defined for sequences X_1, X_2, \dots taking values in a

countable space \mathbf{X} . Such a sequence is called **Markov exchangeable** if the probability of observing an initial trajectory x_1, \dots, x_n depends only on the initial state x_1 and, for every pair $y, y' \in \mathbf{X}$, on the number of transitions $t_{y,y'} = \#\{j < n : x_j = y, x_{j+1} = y'\}$. In particular, the probability does not depend on when each transition occurs. Diaconis and Freedman [19] showed the following:

If a (recurrent) process is Markov exchangeable, it is a mixture of Markov chains.

(Recurrence means that each visited state is visited infinitely often if the process is run for an infinite number of steps.) Thus, each ergodic distribution \mathbf{p}_θ is the distribution of a Markov chain, and a parameter value θ consists of a distribution on \mathbf{X} (the distribution of the initial state) and a transition matrix. If a Markov exchangeable process is substituted for the Markov chain in a hidden Markov model, i.e., if the Markov exchangeable variables are latent variables of the model, the resulting model can express much more general dependencies than Markov exchangeability. The infinite hidden Markov model [9] is an example; see [24]. Recent work by Bacallado, Favaro, and Trippa [8] constructs prior distributions on random walks that are Markov exchangeable and almost surely reversible. \triangleleft

A very general approach to modeling is to assume that an exchangeability assumption holds marginally at each value of a covariate variable z , e.g., a time or a location in space: Suppose \mathbf{X}^∞ is a set of structures as described above, and \mathbf{Z} is a space of covariate values. A **marginally exchangeable random structure** is a random measurable mapping

$$\xi : \mathbf{Z} \rightarrow \mathbf{X}_\infty \quad (2.16)$$

such that, for each $z \in \mathbf{Z}$, the random variable $\xi(z)$ is an exchangeable random structure in \mathbf{X}_∞ .

EXAMPLE 2.6 (Dependent Dirichlet process). A popular example of a marginally exchangeable model is the dependent Dirichlet process (DDP) of MacEachern [48]. In this case, for each $z \in \mathbf{Z}$, the random variable $\xi(z)$ is a random probability measure whose distribution is a Dirichlet process. More formally, \mathbf{Y} is some sample space, $\mathbf{X}_\infty = \mathbf{M}(\mathbf{Y})$, and the DDP is a distribution on mappings $\mathbf{Z} \rightarrow \mathbf{M}(\mathbf{Y})$; thus, the DDP is a random conditional probability. Since $\xi(z)$ is a Dirichlet process if z is fixed, samples from $\xi(z)$ are exchangeable. \triangleleft

Eq. (2.16) is, of course, just another way of saying that ξ is a \mathbf{X}_∞ -valued stochastic process indexed by \mathbf{Z} , although we have made no specific requirements on the paths of ξ . The interpretation as a path is more apparent in the next example.

EXAMPLE 2.7 (Coagulation/fragmentation models). If ξ is a coagulation or fragmentation process, \mathbf{X}_∞ is the set of partitions of \mathbb{N} (as in Kingman's theorem), and $\mathbf{Z} = \mathbb{R}_+$. For each $z \in \mathbb{R}_+$, the random variable $\xi(z)$ is an

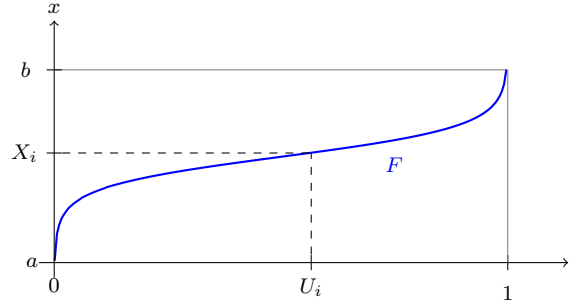


Fig 2: de Finetti's theorem expressed in terms of random functions: If F is the inverse CDF of the random measure Θ in the de Finetti representation, X_i can be generated as $X_i := F(U_i)$, where $U_i \sim \text{Uniform}[0, 1]$.

exchangeable partition—hence, Kingman's theorem is applicable marginally in time. Over time, the random partitions become consecutively finer (fragmentation processes) or coarser (coagulation processes): At random times, a randomly selected block is split, or two randomly selected blocks merge. We refer to [11] for more details and to [61] for applications to Bayesian nonparametrics. \triangleleft

2.5. Random functions vs random measures. De Finetti's theorem can be equivalently formulated in terms of a random function, rather than a random measure, and this formulation provides some useful intuition for Section 3. Roughly speaking, this random function is the inverse CDF of the random measure Θ in de Finetti's theorem; see Fig. 2.

More precisely, suppose that $\mathbf{X} = [a, b]$. A measure μ on $[a, b]$ can be represented by its cumulative distribution function (CDF), defined as $\psi(x) := \mu([a, x])$. Hence, sampling the random measure Θ in de Finetti's theorem is equivalent to sampling a random CDF Ψ . A CDF is not necessarily an invertible function, but it always admits a so-called right-continuous inverse $\overline{\psi^{-1}}$, given by

$$\overline{\psi^{-1}}(u) = \inf \{x \in [a, b] \mid \psi(x) \geq u\}. \quad (2.17)$$

This function inverts ψ in the sense that $\psi \circ \overline{\psi^{-1}}(u) = u$ for all $u \in [0, 1]$. It is well-known that any scalar random variable X_i with CDF ψ can be generated as

$$X_i \stackrel{d}{=} \overline{\psi^{-1}}(U_i) \quad \text{where } U_i \sim \text{Uniform}[0, 1]. \quad (2.18)$$

In the special case $\mathbf{X} = [a, b]$, de Finetti's theorem therefore translates as follows: If X_1, X_2, \dots is an exchangeable sequence, then there is a random function $F := \Psi^{-1}$ such that

$$(X_1, X_2, \dots) \stackrel{d}{=} (F(U_1), F(U_2), \dots), \quad (2.19)$$

where U_1, U_2, \dots are i.i.d. uniform variables.

It is much less obvious that the same should hold on an arbitrary sample space, but that is indeed the case:

COROLLARY 2.8. *Let X_1, X_2, \dots be an infinite, exchangeable sequence of random variables with values in a*

space \mathbf{X} . Then there exists a random function F from $[0, 1]$ to \mathbf{X} such that, if U_1, U_2, \dots is an i.i.d. sequence of uniform random variables,

$$(X_1, X_2, \dots) \stackrel{d}{=} (F(U_1), F(U_2), \dots). \quad (2.20)$$

□

As we will see in the next section, this random function representation generalizes to the more complicated case of array data, whereas the random measure representation in Eq. (2.3) does not.

3. Exchangeable Graphs, Matrices, and Arrays.

Random arrays are a very general type of random structure, which include important special cases, such as random graphs and random matrices. The representation theorem for exchangeable arrays is the Aldous-Hoover theorem. In this section, we focus on 2-arrays, matrices and graphs. The general case for d -arrays is conceptually similar, but considerably more technical, and we postpone it until Section 6.

A **d -array** is a collection of elements $x_{i_1, \dots, i_d} \in \mathbf{X}$ indexed by d indices $i_1, \dots, i_d \in \mathbb{N}$. A sequence is a 1-array. In this section, we assume the random structure X_∞ to be a random 2-array

$$X_\infty = (X_{ij}) = \begin{pmatrix} X_{11} & X_{12} & \dots \\ X_{21} & X_{22} & \dots \\ \vdots & \vdots & \ddots \end{pmatrix}. \quad (3.1)$$

A **random matrix** is a random 2-array, although the term matrix usually implies that \mathbf{X} has the algebraic structure of a field. A **random graph** is a random matrix with $\mathbf{X} = \{0, 1\}$. As in the sequence case, we assume X_∞ is infinite in size, and the statistical interpretation is that an observed, finite array is a sub-array of X_∞ . In network analysis problems, for example, an observed graph with n vertices would be interpreted as a random induced sub-graph of a huge underlying graph X_∞ , which is so large that it is modeled as infinite.

In this section, we are interested in the characterization of random arrays whose distributions are invariant to permutations reordering the rows and columns. For a 2-array, there are two natural ways to define exchangeability: we can ask that the distribution of the array be invariant only to the *joint* (simultaneous) permutation of the rows and columns or also to *separate* permutations of rows and columns.

DEFINITION 3.1. A random 2-array (X_{ij}) is called **jointly exchangeable** if

$$(X_{ij}) \stackrel{d}{=} (X_{\pi(i)\pi(j)}) \quad (3.2)$$

for every permutation π of \mathbb{N} , and **separately exchangeable** if

$$(X_{ij}) \stackrel{d}{=} (X_{\pi(i)\pi'(j)}) \quad (3.3)$$

for every pair of permutations π, π' of \mathbb{N} . ◁

Invariance to all separate permutations of the rows and the columns is an appropriate assumption if rows and columns correspond with two *distinct* sets of entities, such as in a collaborative filtering problem, where rows may correspond to users and columns to movies. On the other hand, if (X_{ij}) is the adjacency matrix of a random graph on the vertex set \mathbb{N} , we would require joint exchangeability, because there is only a single set of entities—the vertices of the graph—each of which corresponds both to a row and a column of the matrix.

Note that if the distribution of an array is invariant to an arbitrary permutation of its elements, then, by de Finetti's theorem, the elements are conditionally i.i.d. If the row and column structure is presumed to be important, invariance only to row and column permutations is a more appropriate modeling assumption.

3.1. The Aldous-Hoover theorem. The analogue of de Finetti's theorem for exchangeable arrays is the **Aldous-Hoover theorem** [2, 34]. It has two versions, for jointly and for separately exchangeable arrays.

THEOREM 3.2 (Aldous-Hoover). A random array (X_{ij}) is jointly exchangeable if and only if it can be represented as follows: There is a random function $F : [0, 1]^3 \rightarrow \mathbf{X}$ such that

$$(X_{ij}) \stackrel{d}{=} (F(U_i, U_j, U_{\{i,j\}})) , \quad (3.4)$$

where $(U_i)_{i \in \mathbb{N}}$ and $(U_{\{i,j\}})_{i,j \in \mathbb{N}}$ are, respectively, a sequence and an array of i.i.d. Uniform $[0, 1]$ random variables. □

Because the variables $U_{\{i,j\}}$ are indexed by a set, the indices are unordered, and we can think of the array $(U_{\{i,j\}})$ as an upper-triangular matrix with i.i.d. uniform entries. If the function F is symmetric in its first two arguments, then X_∞ is symmetric—that is, if $F(x, y, \cdot) = F(y, x, \cdot)$ for all x and y , then $X_{ij} = X_{ji}$ for all i and j . In general, however, a jointly exchangeability matrix or 2-array need not be symmetric.

Separately exchangeable arrays can also be given a precise characterization using Theorem 3.2:

COROLLARY 3.3 (Aldous). A random array (X_{ij}) is separately exchangeable if and only if it can be represented as follows: There is a random function $F : [0, 1]^3 \rightarrow \mathbf{X}$ such that

$$(X_{ij}) \stackrel{d}{=} (F(U_i^{\text{row}}, U_j^{\text{col}}, U_{ij})) , \quad (3.5)$$

where $(U_i^{\text{row}})_{i \in \mathbb{N}}$, $(U_j^{\text{col}})_{j \in \mathbb{N}}$ and $(U_{ij})_{i,j \in \mathbb{N}}$ are, respectively, two sequences and an array of i.i.d. Uniform $[0, 1]$ random variables. □

It is not hard to see that this follows directly from the jointly exchangeable case: If we choose two disjoint, infinite subsets C and R of \mathbb{N} , there exist bijections $r : \mathbb{N} \rightarrow R$ and $c : \mathbb{N} \rightarrow C$. Separate exchangeability of (X_{ij}) then implies $(X_{ij}) \stackrel{d}{=} (X_{r_i c_j})$. Because a separately exchangeable array is jointly exchangeable, it can be represented as in Eq. (3.4), and substituting in $(X_{r_i c_j})$, R and C yields (3.5).

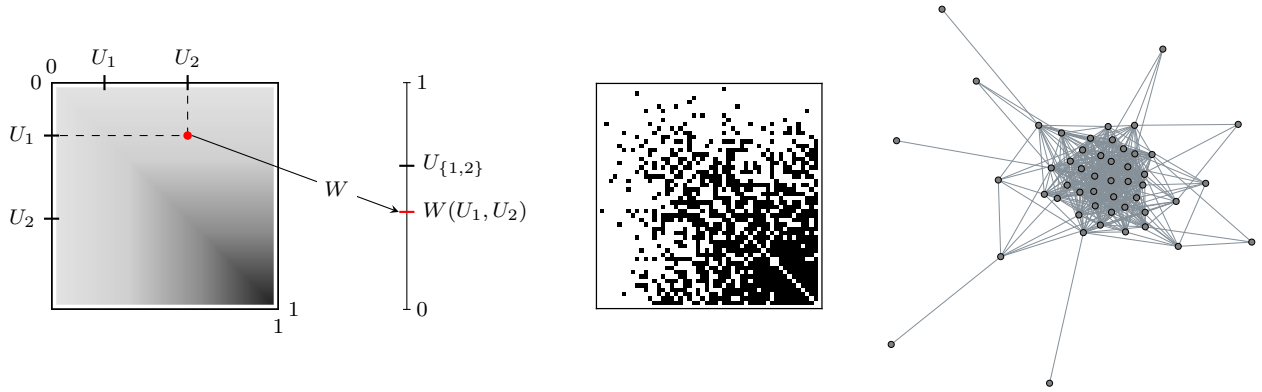


Fig 3: Sampling an exchangeable random graph according to Eq. (3.8). *Left:* A heat-map visualization of the random function W on $[0, 1]^2$, given here by $W(x, y) = \min\{x, y\}$. In the case depicted here, the edge $\{1, 2\}$ is not present in the graph, because $U_{\{1,2\}} > W(U_1, U_2)$. *Middle:* The adjacency matrix of a 50-vertex random graph, sampled from the function on the left. Rows (and columns) in the matrix have been ordered by their underlying U_i value, resulting in a matrix resembling W . *Right:* A plot of the random graph sample. The highly connected vertices plotted in the center correspond to values lower right region in $[0, 1]^2$. The minimum function example, due to Lovász [45], is chosen as a particularly simple symmetric function which is not piece-wise constant. See Section 4 for examples with more structure.

Because separate exchangeability treats rows and columns independently, the single sequence (U_i) of random variables in Eq. (3.4) is replaced by two distinct sequences $(U_i^{\text{row}})_{i \in \mathbb{N}}$ and $(U_j^{\text{col}})_{j \in \mathbb{N}}$, respectively. Additionally, for each pair of distinct indices i and j , the single variable $U_{\{i,j\}}$ in the joint case is now replaced by a pair of variables U_{ij} and U_{ji} . The index structure of the uniform random variables is the only difference between the jointly and separately exchangeable case.

EXAMPLE 3.4 (Collaborative filtering). In the prototypical version of a collaborative filtering problem, users assign scores to movies. Scores may be binary (“like/don’t like”, $X_{ij} \in \{0, 1\}$), have a finite range (“one to five stars”, $X_{ij} \in \{1, \dots, 5\}$), etc. Separate exchangeability then simply means that the probability of seeing any particular realization of the matrix does not depend on the way in which either the users or the movies are ordered. \triangleleft

Like de Finetti’s and Kingman’s theorem, the representation results are complemented by a convergence result, due to Kallenberg [35, Theorem 3]. The general result involves some technicalities and is not stated here; the special case for random graphs is discussed in Section 5.

REMARK 3.5 (Non-uniform sampling schemes). Is it important that the random variables (U_i) and (U_{ij}) are uniformly-distributed in $[0, 1]$? Can different distribution be chosen? If (V_i) is an i.i.d.- μ sequence and (V_{ij}) an independent i.i.d.- τ array, then the array $(G(V_i, V_j, V_{ij}))$ will be separately exchangeable if G is independent of (V_i) and (V_{ij}) , irrespective of the distribution μ of each V_i or distribution τ of each V_{ij} . However, Corollary 3.3 implies there is also a representation in terms of uniform random variables. On the other hand, for an arbitrary separately exchangeable array (X_{ij}) , there may not be a random function F such that $(X_{ij}) \stackrel{d}{=} (F(V_i, V_j, V_{ij}))$. While it is not necessary that μ and τ be the uniform distribution on $[0, 1]$,

it is sufficient. Technically speaking, it suffices for μ and τ to be atomless probability measures on a standard Borel space. The choice of $[0, 1]$ and the uniform distribution is a canonical one in probability. One could also choose \mathbb{R} and any Gaussian distribution. The resemblance between functions on $[0, 1]^2$ and empirical graph distributions (see Fig. 3) makes the unit square convenient for purposes of exposition. \triangleleft

3.2. Exchangeable Graphs. A particular important type of random structures are random graphs defined on a nonrandom vertex set. For a graph on a countably infinite vertex set, we can consider the vertex set to be \mathbb{N} itself, without any loss of generality. A random graph G is then given by a random edge set, which is a random subset of $\mathbb{N} \times \mathbb{N}$. A natural symmetry property of a random graph is the invariance of its distribution to a relabeling/permutation of its vertex set. In this case, G is said to be an **exchangeable graph**. Informally, G can be thought of as a random graph up to isomorphism and so its distribution is determined by the frequency of edges, triangles, five-stars, etc., rather than by where these finite subgraphs appear.

It is straightforward to check that G is an exchangeable graph if and only if its adjacency matrix is jointly exchangeable. More carefully, let (X_{ij}) be an array of binary random variables and put $X_{ij} = 1$ if and only if there is an edge between vertices $i, j \in \mathbb{N}$ in G . Then a simultaneous permutation of the rows and columns of X is precisely a relabeling of the vertex set of G .

An important special case is when G is simple—i.e., undirected and without self-loops. In this case, X is symmetric with a zero diagonal, and its representation (3.4) can be simplified: Let F be a random function satisfying (3.4). Without loss of generality, we may assume that F is symmetric in its first two arguments. Consider the (random) function W from $[0, 1]^2$ to $[0, 1]$ given by $W(x, x) := 0$

on the diagonal and otherwise by

$$W(x, y) := \mathbb{P}[F(x, y, U) = 1 | F] \quad (3.6)$$

$$= \int_0^1 F(x, y, u) du,$$

where $U \sim \text{Uniform}[0, 1]$ is independent of F . Then W is random symmetric function from $[0, 1]^2$ to $[0, 1]$, and, by construction,

$$(F(U_i, U_j, U_{\{i,j\}})) \stackrel{d}{=} (\mathbb{1}\{U_{\{i,j\}} < W(U_i, U_j)\}) . \quad (3.7)$$

We thus obtain the following specialization of the Aldous-Hoover theorem for exchangeable simple graphs:

COROLLARY 3.6. *Let G be a random simple graph with vertex set \mathbb{N} and let X be its adjacency matrix. Then G is an exchangeable graph if and only if there is a random function W from $[0, 1]^2$ to $[0, 1]$ such that*

$$(X_{ij}) \stackrel{d}{=} (\mathbb{1}\{U_{\{i,j\}} < W(U_i, U_j)\}), \quad (3.8)$$

where U_i and $U_{\{i,j\}}$ are independent i.i.d. uniform variables as in (3.4). \square

The representation (3.8) yields the following generative process:

1. Sample a random function $W \sim \nu$.
2. For every vertex $i \in \mathbb{N}$, sample an independent uniform random variable U_i , independent also from W .
3. For every pair of vertices $i < j \in \mathbb{N}$, sample

$$X_{ij} \mid W, U_i, U_j \sim \text{Bernoulli}(W(U_i, U_j)) , \quad (3.9)$$

where $X_{ij} = 1$ indicates the edge connecting i and j is present; if $X_{ij} = 0$, it is absent.

Fig. 3 illustrates the generation of random simple graph.

We call a (measurable) function from $[0, 1]^2$ to $[0, 1]$ a **graphon**. Thus, every exchangeable graph is represented by a random graphon W . In the language of integral decompositions, the ergodic distributions of exchangeable simple graphs are parametrized by graphons. In (2.8), we could take ν to be a distribution on the space of graphons. We will see in Section 5 that θ has an interpretation as a limit of the empirical adjacency matrices for larger and larger subgraphs.

3.3. Application to Bayesian Models. The representation results above have fundamental implications for Bayesian modeling—in fact, they provide a general characterization of Bayesian models of array-valued data:

Statistical models of exchangeable arrays can be parametrized by functions from $[0, 1]^3 \rightarrow [0, 1]$. Every Bayesian model of an exchangeable array is characterized by a prior distribution on the space of functions from $[0, 1]^3$ to $[0, 1]$.

For the special case of simple graphs, we can rephrase this idea in terms of graphons:

Statistical models of exchangeable simple graphs are parametrized by graphons. Every Bayesian model of an exchangeable simple graph is characterized by a prior distribution on the space of graphons.

In a Bayesian model of an exchangeable 2-array, the random function F plays the role of the random parameter Θ in (2.10), and the parameter space \mathbf{T} is the set of measurable function $[0, 1]^3 \rightarrow [0, 1]$. Every possible value f of F defines an ergodic distribution \mathbf{p}_f : In the jointly exchangeable case, for example, Theorem 3.2 shows that X_∞ can be sampled from \mathbf{p}_f by sampling

$$\forall i \in \mathbb{N} : \quad U_i \sim_{\text{iid}} \text{Uniform}[0, 1] \quad (3.10)$$

$$\forall i, j \in \mathbb{N} : \quad U_{\{i,j\}} \sim_{\text{iid}} \text{Uniform}[0, 1] \quad (3.11)$$

and computing X as

$$\forall i, j \in \mathbb{N} : \quad X_{ij} := f(U_i, U_j, U_{\{i,j\}}) . \quad (3.12)$$

Similarly, the ergodic distributions for separately exchangeable 2-arrays are given by (3.5). In the special case of exchangeable simple graphs, the parameter space \mathbf{T} can be reduced to the set of graphons, and the ergodic distribution \mathbf{p}_w defined by a graphon w is given by (3.8).

To the best of our knowledge, Hoff [31] was the first to invoke the Aldous-Hoover theorem for statistical modeling. The problem of estimating the distribution of an exchangeable graph can be formulated as a regression problem on the unknown function w . This perspective was proposed in [43], where the regression problem is formulated as a Bayesian nonparametric model with a Gaussian process prior. The regression model need not be Bayesian, however, and recent work formulates the estimation of w under suitable modeling conditions as a maximum likelihood problem [65].

REMARK 3.7 (Beyond exchangeability). Various types of array-valued data depend on time or some other covariate. In this case, joint or separate exchangeability might be assumed to hold marginally, as described in Section 2.4. E.g., in the case of a graph evolving over time, one could posit the existence of a graphon $W(\cdot, \cdot, t)$ depending also on the time t . More generally, the discussion in 2.4 applies to joint and separate exchangeability just as it does to exchangeable sequences. On the other hand, sometimes exchangeability will not be an appropriate assumption, even marginally. In Section 7, we highlight some reasons why exchangeable graphs may be poor models of very large sparse graphs. \triangleleft

3.4. Uniqueness of representations. In the representation Eq. (3.8), random graph distributions are parametrized by functions $w : [0, 1]^2 \rightarrow [0, 1]$. This representation is not unique, as illustrated in Fig. 4. The technical problems raised by the lack of uniqueness play an important role in the theory of graph limits, which we briefly survey in Section 5. In graph limit theory, two graphons are called **weakly isomorphic** if they parametrize the same random graph. From a statistical perspective, the graphon is not identifiable when regarded as a model parameter, although it is possible to treat the estimation problem up to equivalence of functions [35, Theorem 4].

To see that the representation by w is not unique, simply note that the graphon $w'(x, y) = w(1 - x, 1 - y)$ is weakly

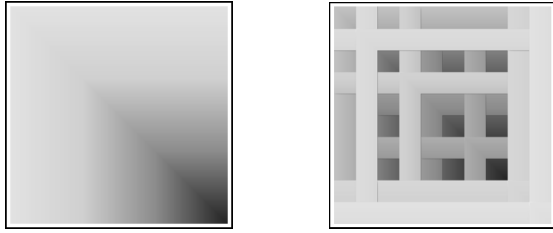


Fig 4: Non-uniqueness of representations: The function on the left parametrizes a random graph as in Fig. 3. On the right, this function has been modified by dividing the unit square into 10×10 blocks and applying the same permutation of the set $\{1, \dots, 10\}$ simultaneously to rows and columns. Since the random variables U_i in Eq. (3.8) are i.i.d., sampling from either function defines one and the same distribution on random graphs.

isomorphic to w because $(U_i) \stackrel{d}{=} (V_i)$ when $V_i = 1 - U_i$ for $i \in \mathbb{N}$. More generally, let $\phi : [0, 1] \rightarrow [0, 1]$ be a **measure-preserving transformation** (MPT), i.e., a map such that $\phi(U)$ is uniformly distributed when U is. By the same argument as above, the graphon w^ϕ given by $w^\phi(x, y) = w(\phi(x), \phi(y))$ is weakly isomorphic to w . Fig. 4 shows an example of a function w and its image under a MPT.

Although any graphon w^ϕ obtained from w by a MPT ϕ is weakly isomorphic to w , the converse is *not* true: For two weakly isomorphic graphons, there need not be a MPT that transforms one into the other [see 45, Example 7.11].

REMARK 3.8 (Monotonization is not applicable). A question that often arises in this context is whether a unique representation can be defined through “monotonization”: On the interval, every bounded real-valued function can be transformed into a monotone left-continuous functions by a measure-preserving transformation, and this left-continuous representation is unique [e.g. 45, Proposition A.19]. It is well known in combinatorics that the same does *not* hold on $[0, 1]^2$ [15, 45]. More precisely, one might attempt to monotonize w on $[0, 1]^2$ by first considering its projection $v(x) := \int_0^1 w(x, y) dy$. The

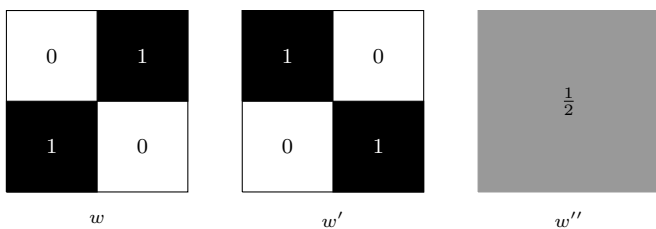


Fig 5: The functions w and w' are distinct but parametrize the same random graph (an almost surely bipartite graph). Both remain invariant and hence distinct under monotonization, which illustrates that monotonization does not yield a canonical representation (see Remark 3.8 for details). Additionally, function w'' shows that the projections do not distinguish different random graphs: w'' projects to the same constant functions as w and w' , but parametrizes a different distribution (an Erdős-Rényi graph with edge probability $1/2$).

function v can be transformed into a monotone representation by a unique MPT ϕ , which we can then apply to both arguments of w to obtain w^ϕ . It is tempting to think that the graphon w^ϕ resulting from this process depends only of the weak-isomorphism class of w , but this approach does not yield a canonical representation: Fig. 5 shows two distinct graphons w and w' that have identical projections $v = v'$ (the constant function $1/2$) and thus identical MPTs $\phi = \phi'$. Therefore, $w^\phi \neq w'^{\phi'}$, illustrating that the MPT given by monotonizing the projection does not yield a canonical representative. \triangleleft

4. Models in the Machine Learning Literature.

The representation theorems show that any Bayesian model of an exchangeable array can be specified by a prior on functions. Models can therefore be classified according to the type of random function they employ. This section surveys several common categories of such random functions, including random piece-wise constant (p.w.c.) functions, which account for the structure of models built using Chinese restaurant processes, Indian buffet processes and other combinatorial stochastic processes; and random continuous functions with, e.g., Gaussian process priors. Special cases of the latter include a range of matrix factorization and dimension reduction models proposed in the machine learning literature. Table 2 summarizes the classes in terms of restrictions on the random function and the values it takes.

4.1. Cluster-based models. Cluster-based models assume that the rows and columns of the random array $X := (X_{ij})$ can be partitioned into (disjoint) classes, such that the probabilistic structure between every row- and column-class is homogeneous. Within social science, this idea is captured by assumptions underlying **stochastic block models** [33, 64].

The collaborative filtering problem described in Example 3.4 is a prototypical application: here, a cluster-based model would assume that the users can be partitioned into classes/groups/types/kinds (of users), and likewise, the movies can also be partitioned into classes/groups/types/kinds (of movies). Having identified the underlying partition of users and movies, each class of user would be assumed to have a prototypical preference for each class of movie.

Because a cluster-based model is described by two partitions, this approach to modeling exchangeable arrays is closely related to clustering, and many well-known non-parametric Bayesian stochastic processes—e.g. the Dirichlet process and Pitman-Yor process, or their combinatorial counterpart, the Chinese restaurant process—are common components of cluster-based models. Indeed, we will begin by describing the Infinite Relational Model [38, 66], the canonical nonparametric, cluster-based, Bayesian model for arrays.

To our knowledge, the Infinite Relational Model, or simply IRM, was the first nonparametric Bayesian model of an exchangeable array. The IRM was introduced in 2006

Model class	Random function F	Distribution of values
Cluster-based (Section 4.1)	p.w.c. on random product partition	exchangeable
Feature-based (Section 4.2)	p.w.c. on random product partition	feature-exchangeable
Piece-wise constant (Section 4.3)	p.w.c. general random partition	arbitrary
Gaussian process-based (Section 4.4)	continuous	Gaussian

TABLE 2

Important classes of exchangeable array models. (Note that p.w.c. stands for piecewise constant.)

independently by Kemp, Tenenbaum, Griffiths, Yamada and Ueda [38], and then by Xu, Tresp, Yu and Kriegel [66]. (Xu et al. referred to their model as the Infinite Hidden Relational Model, but we will refer to both simply by IRM.) The IRM can be seen as a nonparametric generalization of parametric stochastic block models introduced by Holland, Laskey and Leinhardt [33] and Wasserman and Anderson [64]. In the following example, we describe the model for the special case of a $\{0, 1\}$ -valued array.

EXAMPLE 4.1 (Infinite Relational Model). Under the IRM, the generative process for a finite subarray of binary random variables X_{ij} , $i \leq n$, $j \leq m$, is as follows: To begin, we partition the rows (and then columns) into clusters according to a **Chinese restaurant process**, or simply CRP. (See Pitman’s excellent monograph [53] for an in-depth treatment of the CRP and related processes.) In particular, the first and second row are chosen to belong to the same cluster with probability proportional to 1 and to belong to different clusters with probability proportional to a parameter $c > 0$. Subsequently, each row is chosen to belong to an existing cluster with probability proportional to the current size of the cluster, and to a new cluster with probability proportional to c . Let $\Pi := \{\Pi_1, \dots, \Pi_\kappa\}$ be the random partition of $\{1, \dots, n\}$ induced by this process, where Π_1 is the cluster containing 1, and Π_2 is the cluster containing the first row not belonging to Π_1 , and so on. Note that the number of clusters, κ , is also a random variable. Let $\Pi' := \{\Pi'_1, \dots, \Pi'_{\kappa'}\}$ be the random partition of $\{1, \dots, m\}$ induced by this process on the *columns*, possibly with a different parameter $c' > 0$ determining the probability of creating new clusters. Next, for every pair (k, k') of cluster indices, $k \leq \kappa$, $k' \leq \kappa'$, we generate an independent beta random variable $\theta_{k,k'}$.¹ Finally, we generate each X_{ij} independently from a Bernoulli distribution with mean $\theta_{k,k'}$, where $i \in \Pi_k$ and $j \in \Pi'_{k'}$. As we can see, $\theta_{k,k'}$ represents the probability of links arising between elements in clusters k and k' .

The Chinese restaurant process (CRP) generating Π and Π' is known to be exchangeable in the sense that the distribution of Π is invariant to a permutation of the underlying set $\{1, \dots, n\}$. It is then straightforward to see that the distribution on the subarray is exchangeable. In addition, it is straightforward to verify that, were we to have generated an $(n+1) \times (m+1)$ array, the marginal distribution on the $n \times m$ subarray would have agreed with that of the above process. This implies that we have defined a

so-called projective family and so results from probability theory imply that there exists an infinite array and that the above process describes the distribution of every finite subarray. \triangleleft

The IRM model can be seen to be a special case of exchangeable arrays that we will call **cluster-based**. We will define this class formally, and then return to the IRM example, re-describing it in this new language where the exchangeability is manifest. To begin, we first introduce a subclass of cluster-based models, called **simple cluster-based** models:

DEFINITION 4.2. We say that a Bayesian model of an exchangeable array is *simple cluster-based* when, for some random function F representing X , there are random partitions B_1, B_2, \dots and C_1, C_2, \dots of the unit interval $[0, 1]$ such that:

1. On each block $A_{i,j} := B_i \times C_j \times [0, 1]$, F is constant. Let f_{ij} be the value F takes on block $A_{i,j}$.
2. The block values (f_{ij}) are themselves an exchangeable array, and independent from (B_i) and (C_j) .

We call an array simple cluster-based if its distribution is.² \triangleleft

Most examples of simple cluster-based models in the literature—including, e.g., the IRM—take the block values f_{ij} to be conditionally i.i.d. (and so the array (f_{ij}) is then trivially exchangeable). As an example of a more flexible model for (f_{ij}) , which is merely exchangeable, consider the following:

EXAMPLE 4.3 (exchangeable link probabilities). For every block i in the row partition, let u_i be an independent and identically distributed Gaussian random variable. Similarly, let (v_j) be an i.i.d. sequence of Gaussian random variables for the column partitions. Then, for every row and column block i, j , put $f_{ij} := \text{sig}(u_i + v_j)$, where $\text{sig}: \mathbb{R} \rightarrow [0, 1]$ is a sigmoid function. The array (f_{ij}) is obviously exchangeable. \triangleleft

Like with cluster-based models of exchangeable sequences, if the number of classes in each partition is bounded, then a simple cluster-based model of an exchangeable array is a mixture of a finite-dimensional family

¹For simplicity, assume that we fix the hyperparameters of the beta distribution, although this assumption can be relaxed if one is careful not to break exchangeability or projectivity.

²Those familiar with the theory of exchangeable partitions might note that our model does not allow for singleton blocks (aka *dust*). This is a straightforward generalization, but complicates the presentation.

of ergodic distributions. Therefore, mixtures of an infinite-dimensional family must place positive mass on partitions with arbitrarily many classes.

In order to define the more general class of cluster-based models, we relax the piecewise constant nature of the random function. In particular, we will construct an exchangeable array (X_{ij}) from a corresponding array (θ_{ij}) of parameters, which will have a simple cluster-based model. The parameter θ_{ij} could, e.g., determine the probability of an interaction $X_{ij} \in \{0, 1\}$. More generally, the parameters index a family of distributions on \mathbf{X} .

To precisely define such models, we adapt the notion of a *randomization* from probability theory [36]. Intuitively, given a family $\{P_t : t \in \mathbf{T}\}$ of distributions on \mathbf{X} and a collection of random variables θ_i in \mathbf{T} , we can generate a conditionally independent collection of random variable Y_i , each with distribution P_{θ_i} . The following definition makes this idea more precise.

DEFINITION 4.4 (randomization). Let \mathbf{T} be a parameter space, let P be a family $\{P_t : t \in \mathbf{T}\}$ of distributions on \mathbf{X} , and let $\theta := (\theta_i : i \in I)$ be a collection of random variables taking values in \mathbf{T} , indexed by elements of a set I . (E.g., $I = \mathbb{N}^2$) We say that a collection $Y := (Y_i : i \in I)$ of random variables, indexed by the same set I , is a ***P*-randomization of θ** when the elements Y_i are conditionally independent given θ , and $Y_i \mid \theta \sim P_{\theta_i}$ for all $i \in I$. \triangleleft

Thus, a generative model for the collection Y is to first generate θ , and then, for each $i \in I$, to sample Y_i independently from the distribution P_{θ_i} . It is straightforward to prove that, if θ is an exchangeable array and Y is a randomization of θ , then Y is exchangeable. We may therefore define:

DEFINITION 4.5 (cluster-based models). We say that a Bayesian model for an exchangeable array $X := (X_{ij})$ in \mathbf{X} is **cluster-based** when X is a *P*-randomization of a simple cluster-based exchangeable array $\theta := (\theta_{ij})$ taking values in a space \mathbf{T} , for some family $\{P_t : t \in \mathbf{T}\}$ of distributions on \mathbf{X} . We say an array is cluster-based when its distribution is. \triangleleft

The intuition is that the cluster memberships of every pair i, j of individuals determines a parameter θ_{ij} , which in turn determines a distribution $P_{\theta_{ij}}$. The actual observed relationship X_{ij} is then a sample from $P_{\theta_{ij}}$.

Let X , θ , \mathbf{T} and P be defined as above. We may characterize the random function F for X as follows: Let $\phi : \mathbf{T} \times [0, 1] \rightarrow \mathbf{X}$ be such that $\phi(t, U)$ is P_t -distributed for every $t \in \mathbf{T}$, when U is uniformly distributed in $[0, 1]$. (Such a function ϕ is sometimes called a **sampling function**.) Then, if G is the random function representing the exchangeable array (θ_{ij}) then

$$F(x, y, z) = \phi(G(x, y, z), z) \quad (4.1)$$

is a function representing X . (Recall that $G(x, y, z) = G(x, y, z')$ for almost all x, y, z, z' by part 1 of Definition 4.2.)

The next example describes a model that generates the random partitions using a Dirichlet process.

EXAMPLE 4.6 (Infinite Relational Model continued). We may alternatively describe the IRM distribution on exchangeable arrays as follows: Let P be a family $\{P_t : t \in \mathbf{T}\}$ of distributions on \mathbf{X} (e.g., a family of Bernoulli distributions indexed by their means in $[0, 1]$) and let H be a prior distribution on the parameter space $[0, 1]$ (e.g., a Beta distribution, so as to achieve conjugacy). The IRM model of an array $X := (X_{ij})$ is cluster-based, and in particular, is a *P*-randomization of a simple, cluster-based exchangeable array $\theta := (\theta_{ij})$ of parameters in \mathbf{T} .

In order to describe the structure of θ , it suffices to describe the distribution of the partitions (B_k) and (C_k) as well as that of the block values. For the latter, the IRM simply chooses the block values to be i.i.d. draws from the distribution H . (While the block values can be taken to be merely exchangeable, we have not seen this generalization in the literature.) For the partitions, the IRM utilizes the stick-breaking construction of a Dirichlet process [58].

In particular, let W_1, W_2, \dots be an i.i.d. sequence of $\text{Beta}(1, \alpha)$ random variables, for some concentration parameter $\alpha > 0$. For every $k \in \mathbb{N}$, we then define

$$V_k := (1 - W_1) \cdots (1 - W_{k-1}) W_k. \quad (4.2)$$

With probability one, we have $V_k \geq 0$ for every $k \in \mathbb{N}$ and $\sum_{k=1}^{\infty} V_k = 1$ almost surely, and so the sequence (V_k) characterizes a (random) probability distribution on \mathbb{N} . We then let (B_k) be a sequence of contiguous intervals that partition of $[0, 1]$, where B_k is the half-open interval of length V_k . In the jointly exchangeable case, the random partition (C_k) is usually chosen either as a copy of (B_k) , or as partition sampled independently from the same distribution as (B_k) .

The underlying discrete partitioning of G induces a partition on the rows and columns of the array under the IRM model. In the IRM papers themselves, the clustering of rows and columns is described directly in terms of a Chinese restaurant process (CRP) like we did in our first description of the IRM, rather than as above, where we described the model in terms of an explicit list of probabilities. To connect the random probabilities (V_k) for the rows with the CRP, note that V_k is the limiting fraction of rows in the k th cluster Π_k as the number of rows tends to infinity. \triangleleft

4.2. Feature-based models. Feature-based models of exchangeable arrays have similar structure to cluster-based models. Like cluster-based models, feature-based models partition the rows and columns into clusters, but unlike cluster-based models, feature-based models allow the rows and columns to belong to multiple clusters simultaneously. The set of clusters that a row belongs to are then called its **features**. The interaction between row i and column j is then determined by the features that the row and column possess.

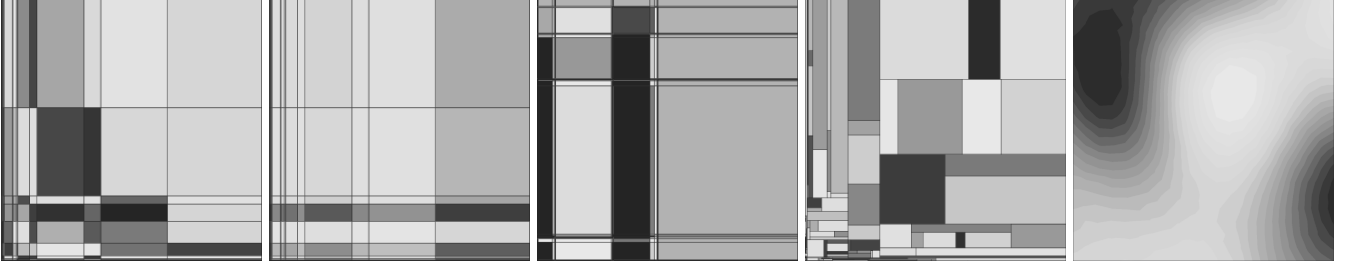


Fig 6: Typical directing random functions underlying, from left to right, 1) an IRM (where partitions correspond with a Chinese restaurant process) with conditionally i.i.d. link probabilities; 2) a more flexible variant of the IRM with merely *exchangeable* link probabilities as in Example 4.3; 3) a LFRM (where partitions correspond with an Indian buffet process) with feature-exchangeable link probabilities as in Example 4.10; 4) a Mondrian-process-based model with a single latent dimension; 5) a Gaussian-processed-based model with a single latent dimension. (Note that, in practice, one would use more than one latent dimension in the last two examples, although this complicates visualization. In the first four figures, we have truncated each of the “stick-breaking” constructions at a finite depth, although, at the resolution of the figures, it is very difficult to notice the effect.)

The stochastic process at the heart of most existing feature-based models of exchangeable arrays is the Indian buffet process, introduced by Griffiths and Ghahramani [28]. The Indian buffet process (IBP) produces an allocation of features in a sequential fashion, much like the Chinese restaurant process produces a partition in a sequential fashion. In the follow example, we will describe the Latent Feature Relational Model (LFRM) of Miller et al. [49], one of the first nonparametric, feature-based models of exchangeable arrays. For simplicity, we will describe the special case of a $\{0, 1\}$ -valued, separately-exchangeable array.

EXAMPLE 4.7 (Latent Feature Relational Model). Under the LFRM, the generative process for a finite subarray of binary random variables X_{ij} , $i \leq n$, $j \leq m$, is as follows: To begin, we allocate features to the rows (and then columns) according to an IBP. In particular, the first row is allocated a Poisson number of features, with mean $\gamma > 0$. Each subsequent row will, in general, share some features with earlier rows, and possess some features not possessed by any earlier row. Specifically, the second row is also allocated a Poisson number of altogether new features, but with mean $\gamma/2$, and, for every feature possessed by the first row, the second row is allocated that feature, independently, with probability $1/2$. In general, the k th row: is allocated a Poisson number of altogether new features, with mean γ/k ; and, for every subset $K \subseteq \{1, \dots, k-1\}$ of the previous rows, and every feature possessed by exactly those rows in K , is allocated that feature, independently, with probability $|K|/n$. (We use the same process to allocate a distinct set of features to the m columns, though potentially with a different constant $\gamma' > 0$ governing the overall number of features.)

We now describe how the features possessed by the rows and columns come to generate the observed subarray. First, we number the row- and column- features arbitrarily, and for every row i and column j , we let $N_i, M_j \subseteq \mathbb{N}$ be the set of features they possess, respectively. For every pair (k, k') of a row- and column- feature, we generate an independent and identically distributed Gaussian random variable $w_{k,k'}$. Finally, we generate each $X_{i,j}$

independently from a Bernoulli distribution with mean $\text{sig}(\sum_{k \in N_i} \sum_{k' \in M_j} w_{k,k'})$. Thus a row and column that possess feature k and k' , respectively, have an increased probability of a connection as $w_{k,k'}$ becomes large and positive, and a decreased probability as $w_{k,k'}$ becomes large and negative.

The exchangeability of the subarray follows from the exchangeability of the IBP itself. In particular, define the family of counts Π_N , $N \subseteq \{1, \dots, n\}$, where Π_N is the number of features possessed by exactly those row in N . We say that $\Pi := (\Pi_N)$ is a **random feature allocation** for $\{1, \dots, n\}$. (Let Π' be the random feature allocation for the columns induced by the IBP.) The IBP is exchangeable in the sense that

$$(\Pi_N) \stackrel{d}{=} (\Pi_{\sigma(N)}) \quad (4.3)$$

for every permutation π of $\{1, \dots, n\}$, where $\sigma(N) := \{\sigma(n) : n \in N\}$. Moreover, the conditional distribution of the subarray given the feature assignments (N_i, M_j) is the same as the conditional distribution given the feature allocations (Π_N, Π'_M) . It is then straightforward to verify that the subarray is itself exchangeable. Like with the IRM example, the family of distributions on subarrays of different sizes is projective, and so there exists an infinite array and the above process describes the distribution of every subarray. \triangleleft

We will cast the LFRM model as a special case of a class of models that we will call **feature-based**. From the perspective of simple cluster-based models, simple feature-based models also have a block structured representing function, but relax the assumption that values of each block form an exchangeable array. To state the definition of this class more formally, we begin by generalizing the notion of a partition of $[0, 1]$. (See [16] for recent work characterizing exchangeable feature allocations.)

DEFINITION 4.8 (feature allocation). Let U be a uniformly-distributed random variable and $E := (E_1, E_2, \dots)$ a sequence of subsets of $[0, 1]$. Given E , we

say that U has feature n when $U \in E_n$. We call the sequence E a **feature allocation** if

$$\mathbb{P}\left\{U \notin \bigcup_{k \geq n} E_k\right\} \rightarrow 1 \quad \text{as } n \rightarrow \infty. \quad (4.4)$$

<

The definition probably warrants some further explanation: A partition is a special case of a feature allocation, in which the sets E_n are disjoint and represent blocks of a partition. The relation $U \in E_k$ then indicates that an object represented by the random variable U is in block k of the partition. In a feature allocation, the sets E_k may overlap. The relation $U \in E_n$ now indicates that the object has feature n . Because the sets may overlap, the object may possess multiple features. However, condition Eq. (4.4) ensures that the number of features per object remains finite (with probability 1).

A feature allocation induces a partition if we equate any two objects that possess exactly the same features. More carefully, for every subset $N \subset \mathbb{N}$ of features, define

$$E_{(N)} := \bigcap_{i \in N} E_i \cap \bigcap_{j \notin N} ([0, 1] \setminus E_j). \quad (4.5)$$

Then, two objects represented by random variables U and U' are equivalent iff $U, U' \in E_{(N)}$ for some finite set $N \subset \mathbb{N}$. As before, we could consider a simple, cluster-based representing function where the block values are given by an $(f_{N,M})$, indexed now by finite subsets $N, M \subseteq \mathbb{N}$. Then $f_{N,M}$ would determine how two objects relate when they possess features N and M , respectively.

However, if we want to capture the idea that the relationships between objects depend on the individual features the objects possess, we would not want to assume that the entries of $f_{N,M}$ formed an exchangeable array, as in the case of a simple, cluster-based model. E.g., we might choose to induce more dependence between $f_{N,M}$ and $f_{N',M}$ when $N \cap N' \neq \emptyset$ than otherwise. The following definition captures the appropriate relaxation of exchangeability:

DEFINITION 4.9 (feature-exchangeable array). Let $Y := (Y_{N,M})$ be an array of random variables indexed by pairs $N, M \subseteq \mathbb{N}$ of finite subsets. For a permutation π of \mathbb{N} and $N \subseteq \mathbb{N}$, write $\pi(N) := \{\pi(n) : n \in N\}$ for the image. Then, we say that Y is **feature-exchangeable** when

$$(Y_{N,M}) \stackrel{d}{=} (Y_{\pi(N), \pi(M)}), \quad (4.6)$$

for all permutations π of \mathbb{N} .

<

Informally, an array Y indexed by sets of features is feature-exchangeable if its distribution is invariant to permutations of the underlying feature labels (i.e., of \mathbb{N}). The following is an example of a feature-exchangeable array, which we will use when we re-describe the Latent Feature Relational Model in the language of feature-based models:

EXAMPLE 4.10 (feature-exchangeable link probabilities). Let $w := (w_{ij})$ be a conditionally i.i.d. array of random variables in \mathbb{R} , and define $\theta := (\theta_{N,M})$ by

$$\theta_{N,M} = \text{sig}(\sum_{i \in N} \sum_{j \in M} w_{ij}), \quad (4.7)$$

where $\text{sig}: \mathbb{R} \rightarrow [0, 1]$ maps real values to probabilities via, e.g., the sigmoid or probit functions. It is straightforward to verify that θ is feature-exchangeable.

<

We can now define simple feature-based models:

DEFINITION 4.11. We say that a Bayesian model of an exchangeable array X is **simple feature-based** when, for some random function F representing X , there are random feature allocations B and C of the unit interval $[0, 1]$ such that, for every pair $N, M \subseteq \mathbb{N}$ of finite subsets, F takes the constant value $f_{N,M}$ on the block

$$A_{N,M} := B_{(N)} \times C_{(M)} \times [0, 1], \quad (4.8)$$

and the values $f := (f_{N,M})$ themselves form a feature-exchangeable array, independent of B and C . We say an array is simple feature-based if its distribution is.

<

We can relate this definition back to cluster-based models by pointing out that simple feature-based arrays are simple cluster-based arrays when either i) the feature allocations are partitions or ii) the array f is exchangeable. The latter case highlights the fact that feature-based arrays relax the exchangeability assumption of the underlying block values.

As in the case of simple cluster-based models, nonparametric simple feature-based models will place positive mass on feature allocations with an arbitrary number of distinct sets. As we did with general cluster-based models, we will define general feature-based models as randomizations of simple models:

DEFINITION 4.12 (feature-based models). We say that a Bayesian model for an exchangeable array $X := (X_{ij})$ in \mathbf{X} is **feature-based** when X is a P -randomization of a simple, feature-based, exchangeable array $\theta := (\theta_{ij})$ taking values in a space T , for some family $\{P_t : t \in \mathbf{T}\}$ of distributions on \mathbf{X} . We say an array is feature-based when its distribution is.

<

Comparing Definitions 4.5 and 4.12, we see that the relationship between random functions representing θ and X are the same as with cluster-based models. We now return to the LFRM model, and describe it in the language of feature-based models:

EXAMPLE 4.13 (Latent Feature Relational Model continued). Recall that a feature-based model is determined by the randomization family P , the distribution of the underlying feature-exchangeable array f of link probabilities, and the distribution of the random feature allocation. In the case of the LFRM, P is the family Bernoulli(p) distributions, for $p \in [0, 1]$ (although this is easily generalized,

and does not represent an important aspect of the model). The underlying feature-exchangeable array f is that described in Example 4.10.

The random feature allocations underlying the LFRM can be described in terms of so-called “stick-breaking” constructions of the Indian buffet process. One of the simplest stick-breaking constructions, and the one we will use here, is due to Teh, Görür, and Ghahramani [60]. (See also [62], [51] and [52].)

Let W_1, W_2, \dots be an i.i.d. sequence of $\text{Beta}(\alpha, 1)$ random variables for some concentration parameter $\alpha > 0$. For every k , we define $V_k := \prod_{j=1}^k W_j$. (The relationship between this construction and Eq. (4.2) highlights one of several relationships between the IBP and CRP.) It follows that we have $1 \geq V_1 \geq V_2 \geq \dots \geq 0$. The allocation of features then proceeds as follows: for every $k \in \mathbb{N}$, we assign the feature with probability V_k , independently of all other features. It can be shown that $\sum_k V_k$ is finite with probability one, and so we have a valid feature allocation; every object has a finite number of features with probability one.

We can describe a feature allocation (B_n) corresponding with this stick-breaking construction of the IBP as follows: Put $B_1 = [0, V_1)$, and then inductively, for every $n \in \mathbb{N}$, put

$$B_{n+1} := \bigcup_{j=1}^{2^n-1} [b_j, (b_{j+1} - b_j) \cdot V_{n+1}) \quad (4.9)$$

where $B_n = [b_1, b_2) \cup [b_3, b_4) \cup \dots \cup [b_{2^n-1}, b_{2^n})$. (As one can see, this representation obscures the conditional independence inherent in the feature allocation induced by the IBP.) \triangleleft

4.3. Piece-wise constant models. Simple partition- and feature-based models have piecewise-constant structure, which arises because both types of models posit prototypical relationships on the basis of a *discrete* set of classes or features assignments, respectively. More concretely, a partition of $[0, 1]^3$ is induced by partitions of $[0, 1]$.

An alternative approach is to consider partitions of $[0, 1]^3$ directly, or partitions of $[0, 1]^3$ induced by partitions of $[0, 1]^2$. Rather than attempting a definition capturing a large, natural class of such models, we present an illustrative example:

EXAMPLE 4.14 (Mondrian-process-based models [56]). A Mondrian process is a partition-valued stochastic process introduced by Roy and Teh [56]. (See also Roy [55, Chp. V] for a formal treatment.) More specifically, a **homogeneous Mondrian process on $[0, 1]^2$** is a continuous-time Markov chain $(M_t : t \geq 0)$, where, for every time $t \geq 0$, M_t is a floorplan-partition of $[0, 1]^2$ —i.e., a partition of $[0, 1]^2$ comprised of axis-aligned rectangles of the form $A = B \times C$, for intervals $B, C \subseteq [0, 1]$. It is assumed that M_0 is the trivial partition containing a single rectangle.

Every continuous-time Markov chain is characterized by the mean waiting times between jumps and the discrete-time Markov process of jumps (i.e., the *jump chain*) em-

bedded in the continuous-time chain. In the case of a Mondrian process, the mean waiting time from a partition composed of a finite set of rectangles $\{B_1 \times C_1, \dots, B_k \times C_k\}$ is $\sum_{j=1}^k (|B_j| + |C_j|)^{-1}$. The jump chain of the Mondrian process is entirely characterized by its transition probability kernel, which is defined as follows: From a partition $\{B_1 \times C_1, \dots, B_k \times C_k\}$ of $[0, 1]^2$, we choose to “cut” exactly one rectangle, say $B_j \times C_j$, with probability proportional to $|B_j| + |C_j|$; Choosing j , we then cut the rectangle vertically with probability proportional to $|C_j|$ and horizontally with probability proportional to $|B_j|$; Assuming the cut is horizontal, we partition B_j into two intervals $B_{j,1}$ and $B_{j,2}$, uniformly at random; The jump chain then transitions to the partition where $B_j \times C_j$ is replaced by $B_{j,1} \times C_j$ and $B_{j,2} \times C_j$; The analogous transformation occurs in the vertical case.

As is plain to see, each partition is produced by a sequence of cuts that hierarchically partition the space. The types of floorplan partitions of this form are called **guillotine partitions**. Guillotine partitions are precisely the partitions represented by kd -trees, the classical data structure used to represent hierarchical, axis-aligned partitions.

The Mondrian process possesses several invariances that allow one to define a Mondrian process M_t^* on all of \mathbb{R}^2 . The resulting process is no longer a continuous-time Markov chain. In particular, for all $t > 0$, M_t^* has a countably infinite number of rectangles with probability one. Roy and Teh [56] use this extended process to produce a nonparametric prior on random functions as follows:

Let (ψ_n) be an exchangeable sequence of random variables in \mathbf{X} , let M be a Mondrian process on \mathbb{R}^2 , independent of (ψ_n) , and let (A_n) be the countable set of rectangles comprising the partition of \mathbb{R}^2 given by M_c for some constant $c > 0$. Roy and Teh propose the random function F from $[0, 1]^3$ to $[0, 1]$ given by $F(x, y, z) = \psi_n$, where n is such that $A_n \ni (-\log x, -\log y)$. An interesting property of F is that the partition structure along any axis-aligned slice of the random function agrees with the stick-breaking construction of the Dirichlet process, presented in the IRM model example. Roy and Teh present results in the case where the ψ_n are Beta random variables, and the data are modeled as a Bernoulli randomization of an array generated from F . (See [55] and [56] for more details.) \triangleleft

4.4. Gaussian-process-based models. Up until now, we have discussed classes of models for exchangeable arrays whose random functions have piece-wise constant structure. In this section we briefly discuss a large and important class of models that relax this restriction by modeling the random function as a Gaussian process.

We begin by recalling the definition of a Gaussian process [e.g. 54]. Let $G := (G_i : i \in I)$ be an indexed collection of \mathbb{R} -valued random variables. We say that G is a **Gaussian process on I** when, for all finite sequences of indices $i_1, \dots, i_k \in I$, the vector $(G(i_1), \dots, G(i_k))$ is Gaussian, where we have written $G(i) := G_i$ for notational convenience. A Gaussian process is completely specified by two function-valued parameters: a **mean function** $\mu : I \rightarrow \mathbb{R}$,

satisfying

$$\mu(i) = \mathbb{E}(G(i)), \quad i \in I, \quad (4.10)$$

and a positive semidefinite **covariance function** $\kappa: I \times I \rightarrow \mathbb{R}_+$, satisfying

$$\kappa(i, j) = \text{cov}(G(i), G(j)). \quad (4.11)$$

DEFINITION 4.15 (Gaussian-process-based exchangeable arrays). We say that a Bayesian model for an exchangeable array $X := (X_{ij})$ in \mathbf{X} is **Gaussian-process-based** when, for some random function F representing X , the process $F = (F_{x,y,z}; x, y, z \in [0, 1])$ is Gaussian on $[0, 1]^3$. We will say that an array X is Gaussian-process-based when its distribution is. \triangleleft

In the language of Eq. (3.12), a Gaussian-process-based model is one where a Gaussian process prior is placed on the function f . The definition is stated in terms of the space $[0, 1]^3$ as domain of the uniform random variables U to match our statement of the Aldous-Hoover theorem and of previous models. In the case of Gaussian processes, however, it is arguably more natural to use the real line instead of $[0, 1]$, and we note that this is indeed possible: Given an embedding $\phi: [0, 1]^3 \rightarrow J$ and a Gaussian process G on J , the process G' on $[0, 1]^3$ given by $G'_{x,y,z} = G_{\phi(x,y,z)}$ is Gaussian. More specifically, if the former has a mean function μ and covariance function κ , then the latter has mean $\mu \circ \phi$ and covariance $\kappa \circ (\phi \otimes \phi)$. We can therefore talk about Gaussian processes on spaces J that can be put into correspondence with the unit interval.

The above definition also implies that the array X is conditionally Gaussian, ruling out, e.g., the possibility of $\{0, 1\}$ -valued arrays. This restriction is overcome by considering randomizations of Gaussian-process-based arrays. Indeed, in the $\{0, 1\}$ -valued case, the most common type of randomization can be described as follows:

DEFINITION 4.16 (noisy sigmoidal/probit likelihood). Let ξ be a Gaussian random variable with mean $m \in \mathbb{R}$ and variance $v > 0$, and let $\sigma: \mathbb{R} \rightarrow [0, 1]$ be a sigmoidal function. We define the family $\{L_r: r \in \mathbb{R}\}$ of distributions on $\{0, 1\}$ by $L_r\{1\} = \mathbb{E}(\sigma(r + \xi))$. \triangleleft

Many of the most popular parametric models for exchangeable arrays of random variables can be constructed as (randomizations of) Gaussian-process-based arrays. For a catalog of such models and several nonparametric variants, as well as their covariance functions, see [43]. Here we will focus on the parametric **eigenmodel**, introduced by Hoff [31, 32], and its nonparametric cousin, introduced by Xu, Yan and Qi [67]. To simplify the presentation, we will consider the case of a $\{0, 1\}$ -valued array.

EXAMPLE 4.17 (Eigenmodel [31, 32]). In the case of a $\{0, 1\}$ -valued array, both the eigenmodel and its

nonparametric extension can be interpreted as an L -randomizations of a Gaussian-process-based array $\theta := (\theta_{ij})$, where L is given as in Definition 4.16 for some mean, variance and sigmoid. To complete the description, we define the Gaussian processes underlying θ .

The eigenmodel is best understood in terms of a zero-mean Gaussian process G on $\mathbb{R}^d \times \mathbb{R}^d$. (The corresponding embedding $\phi: [0, 1]^3 \rightarrow \mathbb{R}^d \times \mathbb{R}^d$ is $\phi(x, y, z) = \Phi^{-1}(x)\Phi^{-1}(y)$, where Φ^{-1} is defined so that $\Phi^{-1}(U) \in \mathbb{R}^d$ is a vector of independent doubly-exponential (aka Laplacian) random variables when U is uniformly distributed in $[0, 1]$.) The covariance function $\kappa: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}_+$ of the Gaussian process G underlying the eigenmodel is simply

$$\kappa(u, v; x, y) = \langle u, x \rangle \langle v, y \rangle, \quad u, v, x, y \in \mathbb{R}^d, \quad (4.12)$$

where $\langle \cdot, \cdot \rangle: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ denotes the dot product, i.e., Euclidean inner product. This corresponds with a more direct description of G : in particular,

$$G(x, y) = \langle x, y \rangle_\Lambda \quad (4.13)$$

where $\Lambda \in \mathbb{R}^{d \times d}$ is a $d \times d$ array of independent standard Gaussian random variables and $\langle x, y \rangle_\Lambda = \sum_{n,m} x_n y_m \Lambda_{n,m}$ is an inner product. \triangleleft

A nonparametric counterpart to the eigenmodel was introduced by Xu *et al.* [67]:

EXAMPLE 4.18. The Infinite Tucker Decomposition model [67] defines the covariance function on $\mathbb{R}^d \times \mathbb{R}^d$ to be

$$\kappa(u, v; x, y) = \kappa'(u, x) \kappa'(v, y), \quad u, v, x, y \in \mathbb{R}^d, \quad (4.14)$$

where $\kappa': \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is some positive semi-definite covariance function on \mathbb{R}^d . This change can be understood as generalizing the inner product in Eq. (4.12) from \mathbb{R}^d to a (potentially, infinite-dimensional) reproducing kernel Hilbert space (RKHS). In particular, for every such κ' , there is an RKHS \mathcal{H} such that

$$\kappa'(x, y) = \langle \phi(x), \phi(y) \rangle_{\mathcal{H}}, \quad x, y \in \mathbb{R}^d. \quad (4.15)$$

\triangleleft

A related nonparametric model for exchangeable arrays, which places fewer restrictions on the covariance structure and is derived directly from the Aldous-Hoover representation, is described in [43].

5. Convergence, Concentration and Graph Limits. We have already noted that the parametrization of random arrays by functions in the Aldous-Hoover theorem is not unique. Our statement of the theorem also lacks an asymptotic convergence result such as the convergence of the empirical measure in de Finetti's theorem. The tools to fill these gaps have only recently become available in a new branch of combinatorics which studies objects known as graph limits. This section summarizes a few elementary

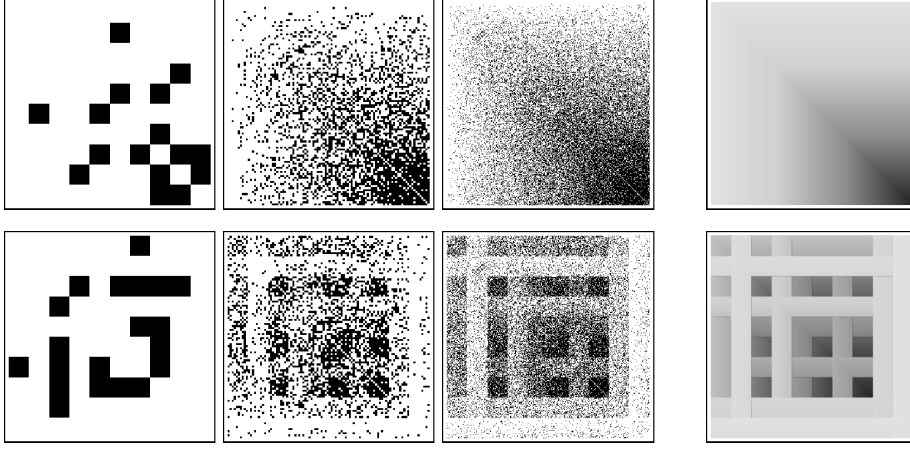


Fig 7: For graph-valued data, the directing random function F in the Aldous-Hoover representation can be regarded as a limit of adjacency matrices: The adjacency matrix of a graph of size n can be represented as a function on $[0, 1]^2$ by dividing the square into $n \times n$ patches of equal size. On each patch, the representing function is constant, with value equal to the corresponding entry of the adjacency matrix. (In the figure, a black patch indicates a value of one and hence the presence of an edge.) As the size of the graph increases, the subdivision becomes finer, and converges to the function depicted on the right for $n \rightarrow \infty$. Convergence is illustrated here for the two functions from Fig. 4. Since the functions are equivalent, the two random graphs within each column are equal in distribution.

notions of this rapidly emerging field and shows how they apply to the Aldous-Hoover theorem for graphs.

Convergence results in statistics study the behavior of the empirical distribution (2.5). The corresponding object for exchangeable graphs is an empirical estimate of the graphon, which is a “checkerboard function”: Given a finite graph g_n with n vertices, we subdivide $[0, 1]^2$ into $n \times n$ square patches, resembling the $n \times n$ adjacency matrix. We then define a function w_{g_n} with constant value 0 or 1 on each patch, equal to the corresponding entry of the adjacency matrix. We call w_{g_n} the **empirical graphon** of g_n . Examples are plotted in Fig. 7. Since w_{g_n} is a valid graphon, it parameterizes an infinite random graph, even though n is finite. Aldous-Hoover theory provides a graph counterpart to the law of large numbers:

THEOREM 5.1 (Kallenberg [35]). *Let w be a graphon. Suppose we sample a random graph from w one vertex at a time, and G_n is the graph given by the first n vertices. Then the distributions defined by w_{G_n} converge weakly to the distribution defined by w with probability 1.* \square

A recent development in combinatorics and graph theory is the theory of graph limits [46, 47]. This theory defines a distance measure δ_\square on graphons (more details below). The metric can be applied to finite graphs, since the graphs can be represented by the empirical graphon. It is then possible to study conditions under which sequences of graphs converge to a limit. It turns out that limits of graphs can be represented by graphons, and the convergence of graphs corresponds precisely to the weak convergence of the distributions defined by the empirical graphons. This theory refines the Aldous-Hoover theory with a large toolbox of powerful results. We describe a few aspects in the following. The authoritative (and very

well-written) reference is [45].

5.1. Metric definition of convergence. The most convenient way to define convergence is by defining a metric: If d is a distance measure, we can define w as the limit of w_{g_n} if $d(w, w_{g_n}) \rightarrow 0$ as $n \rightarrow \infty$. The metric on functions which has emerged as the “right” choice for graph convergence is called the **cut metric**, and is defined as follows: We first define a norm as

$$\|w\|_\square := \sup_{S, T \subset [0, 1]} \int_{S \times T} w(x, y) dx dy. \quad (5.1)$$

(The integral is with respect to Lebesgue measure dx because the variables U_i are uniformly distributed.) Intuitively—if we assume for the moment that w can indeed be thought of as a limiting adjacency matrix— S and T are subsets of nodes. The integral (5.1) measures the total number of edges between S and T in the “graph” w . Since a partition of the vertices of a graph into two sets is called a **cut**, $\|\cdot\|_\square$ is called the **cut norm**. The distance measure defined by $d_\square(w, w') := \|w - w'\|_\square$ is called the **cut distance**.

Suppose w and w' are two distinct functions which parametrize the same distribution on graphs. The distance d_\square in general perceives such functions as different: The functions in Fig. 7, for instance, define the same graph, but have non-zero distance under d_\square . Hence, if we were to use d_\square to define convergence, the two sequences of graphs in the figure would converge to two different limits. We therefore modify d_\square as follows: For any given w , let $[w]$ be the set of all functions w' which define the same random graph.

$$\delta_\square(w_1, w_2) := \inf_{w' \in [w_2]} d_\square(w_1, w'). \quad (5.2)$$

Informally, we can think of the functions in $[w_2]$ as functions obtained from w_2 by a “rearrangement” like the one

illustrated in Fig. 4. The definition above says that, before we measure the distance between w_1 and w_2 using d_\square , we rearrange w_2 in the way that aligns it most closely with w_1 . In Fig. 4, this closest rearrangement would simply reverse the permutation of blocks, so that the two functions would look identical.

The function δ_\square is called the **cut pseudometric**: It is not an actual metric, since it can take value 0 for two distinct functions. It does, however, have all other properties of a metric. By definition, $\delta_\square(w, w') = 0$ holds if and only if w and w' parametrize the same random graph.

DEFINITION 5.2. We say that a sequence $(g_n)_{n \in \mathbb{N}}$ of graphs converges if $\delta_\square(w_{g_n}, w) \rightarrow 0$ for some measurable function $w : [0, 1]^2 \rightarrow [0, 1]$. The function w is called the limit of (g_n) , and often referred to as a **graph limit**. \triangleleft

Clearly, the graph limit is a graphon, and the two terms are used interchangeably. This definition indeed provides a metric counterpart to convergence of exchangeable graph distributions:

THEOREM 5.3. A function w is the graph limit of a sequence $(g_n)_{n \in \mathbb{N}}$ of graphs if and only if the random graph distributions defined by the empirical graphons w_{g_n} converge weakly to the distribution defined by w . \square

5.2. Unique parametrization in the Aldous-Hoover theorem. Recall that two graphons can be equivalent, in the sense that they are distinct functions but define the same random graph (they are weakly isomorphic in the language of Section 3.4). The equivalence classes $[w]$ form a partition of the space \mathbf{W} of all graphons, which motivates the definition of a “quotient space”: We can define a new space $\widehat{\mathbf{W}}$ by collapsing each equivalence class to a single point. Each element $\widehat{w} \in \widehat{\mathbf{W}}$ corresponds to all functions in one equivalence class, and hence to one specific random graph distribution. Since the pseudometric δ_\square only assigns distance 0 to two distinct functions if they are equivalent, it turns into a metric on $\widehat{\mathbf{W}}$, and $(\widehat{\mathbf{W}}, \delta_\square)$ is a metric space. Although the elements of this space are abstract objects, not actual functions, the space has remarkable analytical properties, and is one of the central objects of graph limit theory.

Since $\widehat{\mathbf{W}}$ contains precisely one element for each ergodic distribution on exchangeable graphs, we can obtain a *unique* parametrization of exchangeable graph models by using $\mathbf{T} := \widehat{\mathbf{W}}$ as a parameter space: If $w \in \mathbf{W}$ is a graphon and \widehat{w} the corresponding element of $\widehat{\mathbf{W}}$ —the element to which w was collapsed in the definition of $\widehat{\mathbf{W}}$ —we define a family $\{\mathbf{p}_{\widehat{w}} : \widehat{w} \in \widehat{\mathbf{W}}\}$ of distributions on exchangeable arrays by taking $\mathbf{p}_{\widehat{w}}$ to be the distribution induced by the uniform sampling scheme described by Eq. (3.8) when $W = w$.

Although the existence of such a probability kernel is not a trivial fact, it follows from a technical result of Orbanz and Szegedy [50]. In particular, the Aldous-Hoover

theorem for an exchangeable random graph G can then be written as a unique integral decomposition

$$\mathbb{P}(G \in \cdot) = \int_{\widehat{\mathbf{W}}} \mathbf{p}_{\widehat{w}}(\cdot) \nu(d\widehat{w}), \quad (5.3)$$

in analogy to the de Finetti representation.

5.3. Regularity and Concentration. All convergence results we have seen so far provide only asymptotic guarantees of convergence, but no convergence rates. We give two examples of concentration results from graph limit theory, which address similar questions as those asked in mathematical statistics and empirical process theory: How large a graph do we have to observe to obtain reliable estimates?

Underlying these ideas is one of the deepest results of modern graph theory, Szemerédi’s regularity lemma: For every very large graph g , there is a small, weighted graph \hat{g} that summarizes all essential structure in g . The only condition is that g is sufficiently large. In principle, this means that \hat{g} can be used as an approximation or summary of g , but unfortunately, the result is only valid for graphs which are much larger than possible in most conceivable applications. There are, however, weaker forms of this result which hold for much smaller graphs.

To define \hat{g} for a given graph g , we proceed as follows: Suppose $\Pi := \{V_1, \dots, V_k\}$ is a partition of $\mathbf{V}(g)$ into k sets. For any two sets V_i and V_j , we define p_{ij} as the probability that two vertices $v \in V_i$ and $v' \in V_j$, each chosen uniformly at random from its set, are connected by an edge. That is,

$$p_{ij} := \frac{\# \text{ edges between } V_i, V_j}{|V_i| \cdot |V_j|}. \quad (5.4)$$

The graph \hat{g}_Π is now defined as the weighted graph with vertex set $\{1, \dots, k\}$ and edge weights p_{ij} for edge (i, j) . To compare this graph to g , it can be helpful to blow it up to a graph g_Π of the same size as g , constructed as follows:

- Each node i is replaced by a clique of size $|V_i|$ (with all edges weighted by 1).
- For each pair V_i and V_j , all possible edges between the sets are inserted and weighted by p_{ij} .

If we measure how much two graphs differ in terms of the cut distance, g can be approximated by g_Π as follows:

THEOREM 5.4 (Weak regularity lemma [27]). Let $k \in \mathbb{N}$ and let g be any graph. There is a partition Π of $\mathbf{V}(g)$ into k sets such that $d_\square(g, g_\Pi) \leq 2(\sqrt{\log(k)})^{-1}$. \square

This form of the result is called “weak” since it uses a less restrictive definition of what it means for g and g_Π to be close than Szemerédi’s original result. The weaker hypothesis makes the theorem applicable to graphs that are, by the standards of combinatorics, of modest size.

A prototypical concentration result based on Theorem 5.4 is the following:

THEOREM 5.5 ([44, Theorem 8.2]). *Let f be a real-valued function on graphs, which is smooth in the sense that $|f(g) - f(g')| \leq d_{\square}(g, g')$ for any two graphs g and g' defined on the same vertex set. Let $G(k, g)$ be a random graph of size k sampled uniformly from g . Then the distribution of $f(G(k, g))$ concentrates around some value $f_0 \in \mathbb{R}$, in the sense that*

$$\mathbb{P}\left\{|f(G(k, g)) - f_0| > \frac{20}{\sqrt{k}}\right\} < 2^{-k}. \quad (5.5)$$

□

The relevance of such results to statistics becomes evident if we think of f as a statistic of a graph or network (such as the edge density) which we try to estimate from an observed subgraph of size k . Results of this type, for graphs and other random structures, are collectively known under the term *property testing*, and are covered by a sizeable literature in combinatorics and theoretical computer science [6, 45].

6. Exchangeability and higher-dimensional arrays*. *The theory of exchangeable arrays extends beyond 2-dimensional arrays, and, indeed, some of the more exciting implications and applications of the theory rely on the general results. In this section we begin by defining the natural extension of (joint) exchangeability to higher dimensions, and then give higher-dimensional analogues of the theorems of Aldous and Hoover due to Kallenberg. These theorems introduce exponentially-many additional random variables as the dimension increases, but a theorem of Kallenberg's shows that only a linear number are necessary to produce an arbitrarily good approximation. The presentation owes much to Kallenberg [35].*

DEFINITION 6.1 (jointly exchangeable d -arrays). Let (X_{k_1, \dots, k_d}) be a d -dimensional array (or simply d -array) of random variables in \mathbf{X} . We say that X is **jointly exchangeable** when

$$(X_{k_1, \dots, k_d}) \stackrel{d}{=} (X_{\pi(k_1), \dots, \pi(k_d)}) \quad (6.1)$$

for every permutation π of \mathbb{N} . ◁

As in the 2-dimensional representation result, a key ingredient in the characterization of higher-dimensional jointly exchangeable d -arrays will be an indexed collection U of i.i.d. latent random variables. In order to define the index set for U , let $\tilde{\mathbb{N}}^d$ be the space of multisets $J \subseteq \mathbb{N}$ of cardinality $|J| \leq d$. E.g., $\{\{1, 1, 3\}\} \in \tilde{\mathbb{N}}^3 \subseteq \tilde{\mathbb{N}}^4$. Rather than two collections—a sequence (U_i) indexed by \mathbb{N} , and a triangular array $(U_{\{i, j\}})$ indexed by multisets of cardinality 2—we will use a single i.i.d. collection U indexed by elements of $\tilde{\mathbb{N}}^d$. For every $I \subseteq [d] := \{1, \dots, d\}$, we will write \tilde{k}_I for the multiset

$$\{\{k_i : i \in I\}\} \quad (6.2)$$

and write

$$(U_{\tilde{k}_I}; I \in 2^{[d]} \setminus \emptyset) \quad (6.3)$$

for the element of the function space $[0, 1]^{2^{[d]} \setminus \emptyset}$ that maps each nonempty subset $I \subseteq [d]$ to the real $U_{\tilde{k}_I}$, i.e., the element in the collection U indexed by the multiset $\tilde{k}_I \in \tilde{\mathbb{N}}^{|I|} \subseteq \tilde{\mathbb{N}}^d$.

THEOREM 6.2 (Aldous, Hoover). *Let U be an i.i.d. collection of uniform random variables indexed by multisets $\tilde{\mathbb{N}}^d$. A random d -array $X := (X_k; k \in \mathbb{N}^d)$ is jointly exchangeable if and only if there is random measurable function $F : [0, 1]^{2^{[d]} \setminus \emptyset} \rightarrow \mathbf{X}$ such that*

$$(X_k; k \in \mathbb{N}^d) \stackrel{d}{=} (F(U_{\tilde{k}_I}; I \in 2^{[d]} \setminus \emptyset); k \in \mathbb{N}^d). \quad (6.4)$$

□

When $d = 2$, we recover Theorem 3.2 characterizing two-dimensional exchangeable arrays. Indeed, if we write $U_i := U_{\{i\}}$ and $U_{ij} := U_{\{i, j\}}$ for notational convenience, then the right hand side of Eq. (6.4) reduces to

$$(F(U_i, U_j, U_{ij}); i, j \in \mathbb{N}) \quad (6.5)$$

for some random $F : [0, 1]^3 \rightarrow \mathbf{X}$. When $d = 3$, we instead have

$$(F(U_i, U_j, U_k, U_{ij}, U_{ik}, U_{jk}, U_{ijk}); i, j, k \in \mathbb{N}) \quad (6.6)$$

for some random $F : [0, 1]^7 \rightarrow \mathbf{X}$, where we have additionally taken $U_{ijk} := U_{\{i, j, k\}}$ for notational convenience.

6.1. Separately exchangeable d -arrays. As in the two-dimensional case, arrays with certain additional symmetries can be treated as special cases. In this section, we consider separate exchangeability in the setting of d -arrays, and in the next section we consider further generalizations. We begin by defining:

DEFINITION 6.3 (separately exchangeable d -arrays). We say that d -array X is **separately exchangeable** when

$$(X_{k_1, \dots, k_d}) \stackrel{d}{=} (X_{\pi_1(k_1), \dots, \pi_d(k_d)}) \quad (6.7)$$

for every collection π_1, \dots, π_d of permutations of \mathbb{N} . ◁

For every $J \subseteq [d]$, let 1_J denote its indicator function (i.e., $1_J(x) = 1$ when $x \in J$ and 0 otherwise), and let the vector $k_J \in \mathbb{Z}_+^d := \{0, 1, 2, \dots\}^d$ be given by

$$k_J := (k_1 1_J(1), \dots, k_d 1_J(d)). \quad (6.8)$$

In order to represent separately exchangeable d -arrays, we will use a collection U of i.i.d. uniform random variables indexed by vectors \mathbb{Z}_+^d . Similarly to above, we will write

$$(U_{k_I}; I \in 2^{[d]} \setminus \emptyset) \quad (6.9)$$

for the element of the function space $[0, 1]^{2^{[d]} \setminus \emptyset}$ that maps each nonempty subset $I \subseteq [d]$ to the real U_{k_I} , i.e., the element in the collection U indexed by the vector k_I . Then we have:

COROLLARY 6.4. *Let U be an i.i.d. collection of uniform random variables indexed by vectors \mathbb{Z}_+^d . A random d -array $X := (X_k; k \in \mathbb{N}^d)$ is separately exchangeable if and only if there is random measurable function $F : [0, 1]^{2^{[d]} \setminus \emptyset} \rightarrow \mathbf{X}$ such that*

$$(X_k; k \in \mathbb{N}^d) \stackrel{d}{=} (F(U_{k_I}; I \in 2^{[d]} \setminus \emptyset); k \in \mathbb{N}^d). \quad (6.10)$$

□

We can consider the special cases of $d = 2$ and $d = 3$ arrays. Then we have, respectively,

$$(F(U_{i0}, U_{0j}, U_{ij}); i, j \in \mathbb{N}) \quad (6.11)$$

for some random $F : [0, 1]^3 \rightarrow \mathbf{X}$; and

$$(F(U_{i00}, U_{0j0}, U_{00k}, U_{ij0}, U_{i0k}, U_{0jk}, U_{ijk}); i, j, k \in \mathbb{N}) \quad (6.12)$$

for some random $F : [0, 1]^7 \rightarrow \mathbf{X}$. As we can see, jointly exchangeable arrays, which are required to satisfy fewer symmetries than their separately exchangeable counterparts, may take $U_{ij0} = U_{0ij} = U_{i0j} = U_{jio} = \dots$. Indeed, one can show that these additional assumptions make jointly exchangeable arrays a strict superset of separately exchangeable arrays, for $d \geq 2$.

6.2. *Further generalizations.* In applications, it is common for the distribution of an array to be invariant to permutations that act simultaneously on *some but not all* of the dimensions. E.g., if the first two dimensions of an array index into the same collection of users, and the users are *a priori* exchangeable, then a sensible notion of exchangeability for the array would be one for which these first two dimensions could be permuted jointly together, but separately from the remaining dimensions.

More generally, we consider arrays that, given a partition of the dimensions of an array into classes, are invariant to permutations that act jointly within each class and separately across classes. More carefully:

DEFINITION 6.5 (π -exchangeable d -arrays). Let $\pi = \{I_1, \dots, I_m\}$ be a partition of $[d]$ into disjoint classes, and let $p = (p^I; I \in \pi)$ be a collection of permutations of \mathbb{N} , indexed by the classes in π . We say that a d -array X is π -exchangeable when

$$(X_{k_1, \dots, k_d}; k \in \mathbb{N}^d) \stackrel{d}{=} (X_{p^{\pi_1}(k_1), \dots, p^{\pi_d}(k_d)}; k \in \mathbb{N}^d), \quad (6.13)$$

for every collection p of permutations, where π_i denotes the subset $I \in \pi$ containing i . ◁

We may now cast both jointly and separately exchangeable arrays as π -exchangeable arrays for particular choices of partitions π . In particular, when $\pi = \{[d]\}$ we recover joint exchangeability, and when $\pi = \{\{1\}, \dots, \{d\}\}$, we recover separate exchangeability. Just as we characterized

jointly and separately exchangeable arrays, we can characterize π -exchangeable arrays.

Let π be a partition of $[d]$. In order to describe the representation of π -exchangeable d -arrays, we will again need a collection U of i.i.d. uniform random variables, although the index set is more complicated than before: Let $\mathcal{V}(\pi) := \mathbf{X}_{I \in \pi} \tilde{\mathbb{N}}^{|I|}$ denote the space of functions taking classes $I \in \pi$ to multisets $J \subseteq \mathbb{N}$ of cardinality $J \leq |I|$. We will then take U to be a collection of i.i.d. uniform random variables indexed by elements in $\mathcal{V}(\pi)$.

When $\pi = \{[d]\}$, $\mathcal{V}(\pi)$ is equivalent to the space $\tilde{\mathbb{N}}^d$ of multisets of cardinality no more than d , in agreement with the index set in the jointly exchangeable case. The separately exchangeable case is also instructive: there $\pi = \{\{1\}, \dots, \{d\}\}$ and so $\mathcal{V}(\pi)$ is equivalent to the space of functions from $[d]$ to $\tilde{\mathbb{N}}^1$, which may again be seen to be equivalent to the space \mathbb{Z}_+^d of vectors, where 0 encodes the empty set $\emptyset \in \tilde{\mathbb{N}}^1 \cap \tilde{\mathbb{N}}^0$. For a general partition π of $[d]$, an element in $\mathcal{V}(\pi)$ is a type of generalized vector, where, for each class $I \in \pi$ of dimensions that are jointly exchangeable, we are given a multiset of indices.

For every $I \subseteq [d]$, let $\tilde{k}_{\pi I} \in \mathcal{V}(\pi)$ be given by

$$\tilde{k}_{\pi I}(J) = \tilde{k}_{I \cap J}, \quad J \in \pi, \quad (6.14)$$

where \tilde{k}_J is defined as above for jointly exchangeable arrays. We will write

$$(U_{\tilde{k}_{\pi I}}; I \in 2^{[d]} \setminus \emptyset) \quad (6.15)$$

for the element of the function space $[0, 1]^{2^{[d]} \setminus \emptyset}$ that maps each nonempty subset $I \subseteq [d]$ to the real $U_{\tilde{k}_{\pi I}}$, i.e., the element in the collection U indexed by the generalized vector $\tilde{k}_{\pi I}$. Then we have:

COROLLARY 6.6 (Kallenberg [35]). *Let π be a partition of $[d]$, and let U be an i.i.d. collection of uniform random variables indexed by generalized vectors $\mathcal{V}(\pi)$. A random d -array $X := (X_k; k \in \mathbb{N}^d)$ is π -exchangeable if and only if there is random measurable function $F : [0, 1]^{2^{[d]} \setminus \emptyset} \rightarrow \mathbf{X}$ such that*

$$(X_k; k \in \mathbb{N}^d) \stackrel{d}{=} (F(U_{\tilde{k}_{\pi I}}; I \in 2^{[d]} \setminus \emptyset); k \in \mathbb{N}^d). \quad (6.16)$$

□

6.3. *Approximations by simple arrays.* These representational results require a number of latent random variables exponential in the dimension of the array, i.e., roughly twice as many latent variables are needed as the entries generated in some subarray. Even if a d -array is sparsely observed, each observation requires the introduction of potentially 2^d variables. (In a densely observed array, there will be overlap, and most latent variables will be reused.)

Regardless of whether this blowup poses a problem for a particular application, it is interesting to note that exchangeable d -arrays can be approximated by arrays with much simpler structure, known as **simple arrays**.

DEFINITION 6.7 (simple d -arrays). Let $U = (U_k^I; I \in \pi, k \in \mathbb{N})$ be an i.i.d. collection of uniform random variables. We say that a π -exchangeable d -array X is **simple** when there is a random function $F: [0, 1]^{[d]} \rightarrow \mathbf{X}$ such that

$$(X_k; k \in \mathbb{N}^d) \stackrel{d}{=} (F(U_{k_1}^{\pi_1}, \dots, U_{k_d}^{\pi_d}); k \in \mathbb{N}^d), \quad (6.17)$$

where π_j is defined as above. \triangleleft

Again, it is instructive to study special cases: in the jointly exchangeable case, taking $U_j := U_j^{\{[d]\}}$, we get

$$(F(U_{k_1}, \dots, U_{k_d}); k \in \mathbb{N}^d) \quad (6.18)$$

and, in the separately exchangeable case, we get

$$(F(U_{k_1}^1, \dots, U_{k_d}^d); k \in \mathbb{N}^d), \quad (6.19)$$

taking $U_j^i := U_j^{\{i\}}$. We may now state the relationship between general arrays and simple arrays:

THEOREM 6.8 (simple approximations, Kallenberg [35, Thm. 2]). *Let X be a π -exchangeable d -array. Then there exists a sequence of simple π -exchangeable arrays X^1, X^2, \dots such that, for all finite sub-arrays $X_J := (X_k; k \in J)$, $J \subseteq \mathbb{N}^d$, the distributions of X_J and X_J^n are mutually absolutely continuous, and the associated densities tend uniformly to 1 as $n \rightarrow \infty$ for fixed J .* \square

7. Sparse random structures and networks. *Exchangeable random structures are not “sparse”. In an exchangeable infinite graph, for example, the expected number of edges attached to each node is either infinite or zero. In contrast, graphs representing network data typically have a finite number of edges per vertex, and exhibit properties like power-laws and “small-world phenomena”, which can only occur in sparse graphs. Hence, even though exchangeable graph models are widely used in network analysis, they are inherently misspecified. Since the lack of sparseness is a direct mathematical consequence of exchangeability, networks and sparse random structures pose a problem that seems to require genuinely non-exchangeable models. The development of a coherent theory is, despite intense efforts in mathematics, a largely unsolved problem. In this section, we make the problem more precise and describe how, at least in principle, exchangeability might be substituted by other symmetry properties. The topic raises a host of challenging questions to which, in most cases, we have no answers.*

7.1. Dense vs Sparse Random Structures. In an exchangeable structure, events either never occur, or they occur infinitely often with a fixed (though unknown) probability. The simplest example is an exchangeable binary sequence: By de Finetti’s theorem, the binary variables are conditionally i.i.d. with a Bernoulli distribution. If we sample infinitely often, conditionally on the random Bernoulli parameter taking value $p \in [0, 1]$, the fraction of

ones in the sequence will be precisely p . Therefore, we either observe a constant proportion of ones (if $p > 0$) or no ones at all (if $p = 0$).

In an exchangeable graph, rather than ones and zeros, we have to consider the possible subgraphs (single edges, triangles, five-stars, etc). Each possible subgraph occurs either never, or infinitely often. Since an infinite graph may have infinitely many edges even if it is sparsely connected, the number of edges is best quantified in terms of a rate:

DEFINITION 7.1. Let (g_n) be a sequence of graphs $g_n = (v_n, e_n)$, where g_n has n vertices. We say that the sequence is **sparse** if, as n increases, $|e_n|$ is of size $O(n)$ (is upper-bounded by $c \cdot n$ for some constant c). It is called **dense** if $|e_n| = \Theta(n^2)$ (lower-bounded by $c \cdot n^2$ for some constant c). \triangleleft

If a random graph is sampled step-wise one vertex at a time, the partial graphs at each step also form a sequence, and we can refer to the random graph as dense or sparse. (This definition has to be used with caution, since changing the order in which vertices are generated may affect the rate.) A typical example of dense random graphs are infinite random graphs in which each vertex has infinite degree. Random graphs with bounded degrees are sparse. Many important types of graph and array data are inherently sparse: In a social network with billions of users, individual users do not, on average, have billions of friends.

FACT 7.2. Exchangeable graphs are not sparse. If a random graph is exchangeable, it is either dense or empty. \triangleleft

The argument is simple: Let G_n be an n -vertex random undirected graph sampled according to Eq. (3.8). The expected proportion of edges in present in G_n , out of all $\binom{n}{2} = \frac{n(n-1)}{2}$ possible edges, is independent of n and given by $\varepsilon := \frac{1}{2} \int_{[0,1]^2} W(x, y) dx dy$. (The factor $\frac{1}{2}$ occurs since W is symmetric.) If $\varepsilon = 0$, it follows that G_n is empty with probability one and therefore trivially sparse. On the other hand, if $\varepsilon > 0$, we have $\varepsilon \cdot \binom{n}{2} = \Theta(n^2)$ edges in expectation and so, by the law of large numbers, G_n is dense with probability one.

REMARK 7.3 (Graph limits are dense). The theory of graph limits described in Section 5 is intimately related to exchangeability, and is inherently a theory of dense graphs: If we construct a sequence of graphs with sparsely growing edge sets, convergence in cut metric is still well-defined, but the limit object is always the empty graphon, i.e., a function on $[0, 1]^2$ which vanishes almost everywhere. \triangleleft

One possible way to generate sparse graphs is of course to modify the sampling scheme for exchangeable graphs to generate fewer edges.

EXAMPLE 7.4 (The BJR model [13]). There is a very simple way to modify the Aldous-Hoover approach into one that generates sparse random graphs: Suppose we sample a finite graph with a fixed number n of vertices. We simply multiply the probability in our usual sampling scheme by $1/n$:

$$X_{ij} \sim \text{Bernoulli}\left(\frac{1}{n}w(U_i, U_j)\right) \quad \text{for} \quad i, j \leq n.$$

Comparison with our argument why exchangeable graphs are dense immediately shows that a graph sampled this way is sparse. More generally, we can multiply w by some other rate function ρ_n (instead of specifically $\rho_n = 1/n$), and ask how this model behaves for $n \rightarrow \infty$. Statistical properties of this model are studied by Bickel, Chen, and Levina [12], who consider the behavior of moment estimators for the edge density, triangle density and other subgraph densities. \triangleleft

An obvious limitation of the BJR model is that it does not actually attempt to model network structure: It can equivalently be sampled by sampling from a graphon as in (3.8) and then deleting each edge independently at random, with probability $(1 - \rho_n)$. (In the parlance of random graph theory, this is exchangeable sampling followed by i.i.d. bond percolation.) In other words, the BJR model modifies an exchangeable graph to fit a first-order statistic of a network (the number of edges), but it cannot generate typical network structures, such as power laws.

7.2. *Beyond exchangeability: Symmetry and ergodic theory.* The example of networks and sparse structures shows that there are important random structures which are not exchangeable. This raises the question whether integral decompositions and statistical models, which we have throughout derived from exchangeability, can be obtained in a similar manner for structures that are not exchangeable. In principle, that is possible: Exchangeability is a special case of a probabilistic symmetry. It turns out that integral decompositions can be derived from much more general symmetries than exchangeability.

A **probabilistic symmetry** is defined by choosing a group \mathbb{G} of transformations $g : \mathbf{X}_\infty \rightarrow \mathbf{X}_\infty$. A random structure X_∞ is called **invariant to \mathbb{G}** or **\mathbb{G} -symmetric** $g(X) \stackrel{d}{=} X$ for all $g \in \mathbb{G}$. If so, we also say that the distribution of X_∞ is \mathbb{G} -invariant. For example, a sequence of random variables is exchangeable if and only if the distribution of the sequence is invariant under the group of permutations of \mathbb{N} acting on the indices of the sequence. Exchangeability of arrays (as in the Aldous-Hoover theorem) corresponds with a subgroup generated by row and column permutations. Invariant measures play a key role in several fields of mathematics, especially ergodic theory.

A very general result, the *ergodic decomposition theorem*, shows that integral decompositions of the form (2.8) are a general consequence of probabilistic symmetries, rather than specifically of exchangeability. The general theme is that there is some correspondence of the form

$$\text{invariance property} \longleftrightarrow \text{integral decomposition}.$$

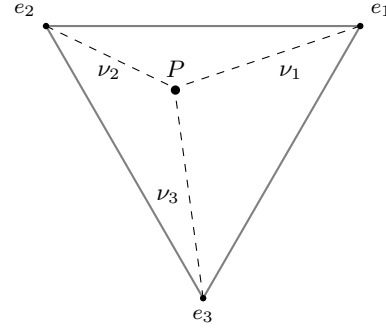


Fig 8: If \mathcal{E} is finite, the de Finetti mixture representation Eq. (2.3) and the more general representation Eq. (7.1) reduce to a finite convex combination. The points inside the set—i.e., the distributions P with the symmetry property defined by the group \mathbb{G} —can be represented as convex combinations $P = \sum_{e_i \in \mathcal{E}} \nu_i e_i$, with coefficients $\nu_i \geq 0$ satisfying $\sum_i \nu_i = 1$. When \mathcal{E} is infinite, an integral is substituted for the sum.

In principle, Bayesian models can be constructed based on any type of symmetry, as long as this symmetry defines a useful set of ergodic distributions. The following statement of the ergodic decomposition theorem glosses over various technical details; for a precise statement, see [37, Theorem A1.4].

THEOREM 7.5 (Varadarajan [63]). *Let \mathbb{G} be a “nice” group acting on a space \mathbf{X}^∞ of infinite structures. Then there exists a family $\mathcal{E} := \{\mathbf{p}_\theta : \theta \in \mathbf{T}\}$ of distributions on \mathbf{X}^∞ such that, if the distribution of a random structure X_∞ is \mathbb{G} -invariant, it has a representation of the form*

$$\mathbb{P}(X_\infty \in \cdot) = \int_{\mathbf{T}} \mathbf{p}_\theta(\cdot) \nu(d\theta) \quad (7.1)$$

for a unique distribution ν on \mathbf{T} . The distributions $\mathbf{p}_\theta \in \mathcal{E}$ are the so-called ergodic distributions associated with \mathbb{G} . \square

We have already encountered the components of (7.1) in Section 2: In Bayesian terms, \mathbf{p}_θ again corresponds to the observation distribution and ν to the prior. Geometrically, the integral representation Eq. (7.1) can be regarded as convex combination. Fig. 8 illustrates this idea for a toy example with three ergodic measures. A special case of the ergodic decomposition theorem is well-known in Bayesian theory as a result of Freedman [25, 26]:

EXAMPLE 7.6 (Freedman’s theorem). Consider a sequence X_1, X_2, \dots as in de Finetti’s theorem. Now replace invariance under permutations by a stronger condition: Let $O(n)$ be the orthogonal group of rotations and reflections on \mathbb{R}^n , i.e., the set of $n \times n$ orthogonal matrices. We now demand that, if we regard any initial sequence of n variables as a random vector in \mathbb{R}^n , then rotating and/or reflecting this vector does not change the distribution of the sequence: That is, for every $n \in \mathbb{N}$ and $M \in O(n)$,

$$(X_1, X_2, \dots) \stackrel{d}{=} (M(X_1, \dots, X_n), X_{n+1}, \dots). \quad (7.2)$$

In the language of Theorem 7.5, the group G is the set of all rotations and reflections acting on all finite prefixes of a sequence. For every $\sigma > 0$, let \mathcal{N}_σ be the distribution of zero-mean Gaussian random variable with standard deviation σ . Freedman showed that, if X^∞ satisfies Eq. (7.2), then its distribution is a scale mixture of Gaussians:

$$\mathbb{P}(X^\infty \in \cdot) = \int_{\mathbb{R}_+} \mathcal{N}_\sigma^\infty(\cdot) \nu_{\mathbb{R}_+}(d\sigma). \quad (7.3)$$

Thus, \mathcal{E} contains all factorial distributions of zero-mean Gaussian distributions on \mathbb{R} , \mathbf{T} is the set $\mathbb{R}_{>0}$ of variances, and ν a distribution on $\mathbb{R}_{>0}$. \triangleleft

Compared to de Finetti’s theorem, where \mathbb{G} is the group of permutations, Freedman’s theorem increases the size of \mathbb{G} : Any permutation can be represented as an orthogonal matrix, but here rotations have been added as well. In other words, we are strengthening the hypothesis by imposing more constraints on the distribution of X^∞ . As a result, the set \mathcal{E} of ergodic measures shrinks from all factorial measures to the set of factorials of zero-mean Gaussians. This is again an example of a general theme:

larger group \longleftrightarrow more specific representation

In contrast, the Aldous-Hoover theorem *weakens* the hypothesis of de Finetti’s theorem—in the matrix case, for instance, the set of all permutations of the index set \mathbb{N}^2 is restricted to those which preserve rows and columns—and hence yields a more general representation.

REMARK 7.7 (Symmetry and sufficiency). An alternative way to define symmetry in statistical models is through sufficient statistics: Intuitively, a symmetry property identifies information which is not relevant to the statistical problem; so does a sufficient statistic. For example, the empirical distribution retains all information about a sample except for the order in which observations were recorded. A model for random sequences is hence exchangeable if and only if the empirical distribution is a sufficient statistic. In an exchangeable graph model, the empirical graphon (the checkerboard function in Fig. 7) is a sufficient statistic. If the sufficient statistic is finite-dimensional and computes an average $\frac{1}{n} \sum_i S_0(x_i)$ over observations for some function S_0 , the ergodic distributions are exponential family models [41]. A readable introduction to this topic is given by Diaconis [18]. The definitive reference is the monograph of Lauritzen [42], who refers to the set \mathcal{E} of ergodic distributions as an *extremal family*. \triangleleft

Not every probabilistic symmetry is applicable in statistics in the same way as exchangeability is. To be useful to statistics, the symmetry must satisfy two conditions:

1. The set \mathcal{E} of ergodic measures should be a “small” subset of the set of symmetric measures.
2. The measures \mathbf{p}_θ should have a tractable representation, such as Kingman’s paint-box or the Aldous-Hoover sampling scheme.

Theorem 7.5 guarantees neither. If (1) is not satisfied, the representation is useless for statistical purposes: The integral representation Eq. (7.1) means that the information in X_∞ is split into two parts, the information contained in the parameter value θ (which a statistical procedure tries to extract) and the randomness represented by \mathbf{p}_θ (which the statistical procedure discards). If the set \mathcal{E} is too large, Θ contains almost all the information in X_∞ , and the decomposition becomes meaningless. We will encounter an appealing notion of symmetry for sparse networks in the next section—which, however, seems to satisfy neither condition (1) or (2). It is not clear at present whether there are useful types of symmetries which do not imply some form of invariance to a group of permutations. Although the question is abstract, the incompatibility of sparseness and exchangeability means it is directly relevant to Bayesian statistics.

7.3. Stationary networks and involution invariance. Is there a form of invariance that yields statistical models for network data? There is indeed a very natural notion of invariance in networks, called *involution invariance*, which we describe in more detail below. This property has interesting mathematical properties and admits an ergodic decomposition as in Theorem 7.5, but it seems to be too weak for applications in statistics.

A crucial difference between network structures and exchangeable graphs is that, in most networks, location in the graph matters. If conditioning on location is informative, exchangeability is broken. Probabilistically, location is modeled by marking a distinguished vertex in the graph. A **rooted graph** (g, v) is simply a graph g in which a particular vertex v has been marked. A very natural notion of invariance for networks is called *involution invariance* [1] or *unimodularity* [10], and can be thought of as a form of stationarity:

DEFINITION 7.8. Let P be the distribution of a random rooted graph, and define a distribution \tilde{P} as follows: A sample $(G, w) \sim \tilde{P}$ is generated by sampling $(G, v) \sim P$, and then sampling w uniformly from the neighbors of v in G . The distribution P is called **involution invariant** if $P = \tilde{P}$. \triangleleft

The definition says that, if an observer randomly walks along the graph G by moving to a uniformly selected neighbor in each step, the *distribution* of the network around the observer remains unchanged (although the actual neighborhoods in a sampled graph may vary).

Involution invariance is a symmetry property which admits an ergodic decomposition, and Aldous and Lyons [1] have characterized the ergodic measures. This characterization is abstract, however, and there is no known “nice” representation resembling, for example, the sampling scheme for exchangeable graphs. Thus, of the two desiderata described in Section 7.2, property (2) does not seem to hold. We believe that property (1) does not hold either: Although we have no proof at present, it seems

that involution invariance is too weak a constraint to yield interesting statistical models (i.e., the set of ergodic distributions is a “large” subset of the involution invariant distributions).

Since exchangeability and involution invariance are the only well-studied probabilistic symmetries for random graphs, the question how statistical models of networks can be characterized is an open problem:

Is there a notion of probabilistic symmetry whose ergodic measures in (7.1) describe useful statistical models for sparse graphs with network properties?

There is a sizeable literature on sparse random graph models which can model power laws and other network properties; see, for example [21]. These are probability models and can be simulated, but estimation from data is often intractable, due to stochastic dependencies between the edges in the random graph. On the other hand, *some* dependence between edges is necessary to obtain a power law and similar properties. Hence, a suitable notion of symmetry would have to restrict dependencies between edges sufficiently to permit statistical inference, but not to the full conditional independence characteristic of the exchangeable case.

8. Further References. Excellent non-technical references on the general theory of exchangeable arrays and other exchangeable random structures are two recent surveys by Aldous [4, 5]. His well-known lecture notes [3] also cover exchangeable arrays. The most comprehensive available reference on the general theory is the monograph by Kallenberg [37] (which presupposes in-depth knowledge of measure-theoretic probability). Kingman’s original article [39] provides a concise reference on exchangeable random partitions. A thorough, more technical treatment of exchangeable partitions can be found in [11].

Schervish [57] gives an insightful discussion of the application of exchangeability to Bayesian statistics. There is a close connection between probabilistic symmetries (such as exchangeability) and sufficient statistics, which is covered by a substantial literature. See Diaconis [18] for an introduction and further references. For applications of exchangeability results to machine learning models, see [24], who discuss applications of the partial exchangeability result of Diaconis and Freedman [19] to the infinite hidden Markov model [9].

The theory of graph limits in its current form was initiated by Lovász and Szegedy [46, 47] and Borgs *et al.* [14]. It builds on work of Frieze and Kannan [27], who introduced both the weak regularity lemma (Theorem 5.4) and the cut norm d_\square . In the framework of this theory, the Aldous-Hoover representation of exchangeable graphs can be derived by purely analytic means [46, Theorem 2.7]. The connection between graph limits and Aldous-Hoover theory was established, independently of each other, by Diaconis and Janson [20] and by Austin [7]. An accessible introduction to the analytic perspective is the survey [44], which assumes basic familiarity with measure-

theoretic probability and functional analysis, but is largely non-technical. The monograph [45] gives a comprehensive account.

Historically, the Aldous-Hoover representation was established in independent works of Aldous and Hoover in the late 1970s. Aldous’ proof used probability-theoretic methods, whereas Hoover, a logician, leveraged techniques from model theory. In 1979, Kingman [40] writes

...a general solution has now been supplied by Dr David Aldous of Cambridge. [...] The proof is at present very complicated, but there is reason to hope that the techniques developed can be applied to more general experimental designs.

Aldous’ paper [2], published in 1981, attributes the idea of the published version of the proof to Kingman. The results were later generalized considerably by Kallenberg [35].

References.

- [1] Aldous, D. and Lyons, R. (2007). Processes on unimodular random networks. *Electron. J. Probab.*, **12**, no. 54, 1454–1508.
- [2] Aldous, D. J. (1981). Representations for partially exchangeable arrays of random variables. *J. Multivariate Anal.*, **11**(4), 581–598.
- [3] Aldous, D. J. (1985). Exchangeability and related topics. In P. L. Hennequin, editor, *École d’Été de Probabilités de Saint-Flour XIII - 1983*, number 1117 in Lecture Notes in Mathematics, pages 1–198. Springer.
- [4] Aldous, D. J. (2009). More uses of exchangeability: Representations of complex random structures.
- [5] Aldous, D. J. (2010). Exchangeability and continuum limits of discrete random structures. In *Proceedings of the International Congress of Mathematicians*.
- [6] Alon, N. and Spencer, J. H. (2008). *The Probabilistic Method*. J. Wiley & Sons.
- [7] Austin, T. (2008). On exchangeable random variables and the statistics of large graphs and hypergraphs. *Probab. Surv.*, **5**, 80–145.
- [8] Bacallado, S. A., Favaro, S., and Trippa, L. (2013). Bayesian nonparametric analysis of reversible Markov chains. *Ann. Statist.* To appear.
- [9] Beal, M. J., Ghahramani, Z., and Rasmussen, C. (2002). The infinite hidden Markov model. In T. G. Dietterich, S. Becker, and Z. Ghahramani, editors, *Advances in Neural Information Processing Systems*, volume 14, pages 577–584.
- [10] Benjamini, I. and Schramm, O. (2001). Recurrence of distributional limits of finite planar graphs. *Electron. J. Probab.*, **6**, no. 23, 13 pp. (electronic).
- [11] Bertoin, J. (2006). *Random Fragmentation and Coagulation Processes*. Cambridge University Press.
- [12] Bickel, P. J., Chen, A., and Levina, E. (2011). The method of moments and degree distributions for network models. *Ann. Statist.*, **39**(5), 2280–2301.
- [13] Bollobás, B., Janson, S., and Riordan, O. (2007). The phase transition in inhomogeneous random graphs. *Random Structures Algorithms*, **31**(1), 3–122.
- [14] Borgs, C., Chayes, J., Lovász, L., Sós, V. T., Szegedy, B., and Vesztegombi, K. (2005). Graph limits and parameter testing. In *Topics in discrete mathematics*, volume 25 of *Algorithms Combin.* Springer.
- [15] Borgs, C., Chayes, J., and Lovász, L. (2010). Moments of two-variable functions and the uniqueness of graph limits. *Geometric And Functional Analysis*, **19**(6), 1597–1619.
- [16] Broderick, T., Jordan, M. I., and Pitman, J. (2013). Feature allocations, probability functions, and paintboxes. To appear in *Bayesian Anal.*, arXiv:1301.6647.
- [17] Bühlmann, H. (1960). *Austauschbare stochastische Variablen und ihre Grenzwertsätze*. Ph.D. thesis. University of California Press, 1960.
- [18] Diaconis, P. (1992). Sufficiency as statistical symmetry. In F. Browder, editor, *Proc. 100th Anniversary American Mathematical Society*, pages 15–26. American Mathematical Society.

- [19] Diaconis, P. and Freedman, D. (1980). De Finetti's theorem for Markov chains. *The Annals of Probability*, **8**(1), pp. 115–130.
- [20] Diaconis, P. and Janson, S. (2008). Graph limits and exchangeable random graphs. *Rendiconti di Matematica, Serie VII*, **28**, 33–61.
- [21] Durrett, R. (2006). *Random Graph Dynamics*. Cambridge University Press.
- [22] de Finetti, B. (1930). Funzione caratteristica di un fenomeno aleatorio. *Atti Reale Accademia Nazionale dei Lincei*, **VI**(4), 86–133.
- [23] de Finetti, B. (1937). La prévision: ses lois logiques, ses sources subjectives. *Ann. Inst. H. Poincaré*, **7**(1), 1–68.
- [24] Fortini, S. and Petrone, S. (2012). Predictive construction of priors in Bayesian nonparametrics. *Braz. J. Probab. Stat.*, **26**(4), 423–449.
- [25] Freedman, D. A. (1962). Invariants under mixing which generalize de Finetti's theorem. *Ann. Math. Statist.*, **33**, 916–923.
- [26] Freedman, D. A. (1963). Invariants under mixing which generalize de Finetti's theorem. *Ann. Math. Statist.*, **34**(1194–1216).
- [27] Frieze, A. and Kannan, R. (1999). Quick approximation to matrices and applications. *Combinatorica*, **19**(2), 175–220.
- [28] Griffiths, T. L. and Ghahramani, Z. (2006). Infinite latent feature models and the Indian buffet process. In *Adv. in Neural Inform. Processing Syst. 18*, pages 475–482. MIT Press, Cambridge, MA.
- [29] Hewitt, E. and Savage, L. J. (1955). Symmetric measures on Cartesian products. *Transactions of the American Mathematical Society*, **80**(2), 470–501.
- [30] Hjort, N., Holmes, C., Müller, P., and Walker, S., editors (2010). *Bayesian Nonparametrics*. Cambridge University Press.
- [31] Hoff, P. (2008). Modeling homophily and stochastic equivalence in symmetric relational data. In *Adv. Neural Inf. Process. Syst. 2007*.
- [32] Hoff, P. D. (2011). Separable covariance arrays via the Tucker product, with applications to multivariate relational data. *Bayesian Analysis*, **6**(2), 179–196.
- [33] Holland, P. W., Laskey, K. B., and Leinhardt, S. (1983). Stochastic blockmodels: First steps. *Social Networks*, **5**(2), 109–137.
- [34] Hoover, D. N. (1979). Relations on probability spaces and arrays of random variables. Technical report, Institute of Advanced Study, Princeton.
- [35] Kallenberg, O. (1999). Multivariate sampling and the estimation problem for exchangeable arrays. *J. Theoret. Probab.*, **12**(3), 859–883.
- [36] Kallenberg, O. (2001). *Foundations of Modern Probability*. Springer, 2nd edition.
- [37] Kallenberg, O. (2005). *Probabilistic Symmetries and Invariance Principles*. Springer.
- [38] Kemp, C., Tenenbaum, J., Griffiths, T., Yamada, T., and Ueda, N. (2006). Learning systems of concepts with an infinite relational model. In *Proc. of the Nat. Conf. on Artificial Intelligence*, volume 21, page 381.
- [39] Kingman, J. F. C. (1978). The representation of partition structures. *J. London Math. Soc.*, **2**(18), 374–380.
- [40] Kingman, J. F. C. (1979). Discussion of: "on the reconciliation of probability assessments" by D. V. Lindley, A. Tversky and R. V. Brown. *Journal of the Royal Statistical Society. Series A*, **142**(2), 171.
- [41] Küchler, U. and Lauritzen, S. L. (1989). Exponential families, extreme point models and minimal space-time invariant functions for stochastic processes with stationary and independent increments. *Scand. J. Stat.*, **16**, 237–261.
- [42] Lauritzen, S. L. (1988). *Extremal Families and Systems of Sufficient Statistics*. Lecture Notes in Statistics. Springer.
- [43] Lloyd, J. R., Orbanz, P., Roy, D. M., and Ghahramani, Z. (2012). Random function priors for exchangeable arrays. In *Adv. in Neural Inform. Processing Syst. 25*.
- [44] Lovász, L. (2009). Very large graphs. In D. Jerison, B. Mazur, T. Mrowka, W. Schmid, R. Stanley, and S. T. Yau, editors, *Current Developments in Mathematics*, pages 67–128. International Press.
- [45] Lovász, L. (2013). *Large Networks and Graph Limits*. American Mathematical Society.
- [46] Lovász, L. and Szegedy, B. (2006). Limits of dense graph sequences. *J. Combin. Theory Ser. B*, **96**, 933–957.
- [47] Lovász, L. and Szegedy, B. (2007). Szemerédi's lemma for the analyst. *Geometric And Functional Analysis*, **17**(1), 252–270.
- [48] MacEachern, S. N. (2000). Dependent Dirichlet processes. Technical report, Ohio State University.
- [49] Miller, K. T., Griffiths, T. L., and Jordan, M. I. (2009). Nonparametric latent feature models for link prediction. *Advances in Neural Information Processing Systems (NIPS)*, pages 1276–1284.
- [50] Orbanz, P. and Szegedy, B. (2012). Borel liftings of graph limits. Preprint.
- [51] Paisley, J., Zaas, A., Woods, C., Ginsburg, G., and Carin, L. (2010). A stick-breaking construction of the beta process. In *Proc. Int. Conf. on Machine Learning*.
- [52] Paisley, J., Blei, D., and Jordan, M. (2012). Stick-breaking beta processes and the poisson process. In *Proc. Int. Conf. on A.I. and Stat.*.
- [53] Pitman, J. (2006). *Combinatorial stochastic processes*, volume 1875 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin. Lectures from the 32nd Summer School on Probability Theory held in Saint-Flour, July 7–24, 2002, With a foreword by Jean Picard.
- [54] Rasmussen, C. E. and Williams, C. K. I. (2006). *Gaussian Processes for Machine Learning*. The MIT Press.
- [55] Roy, D. M. (2011). *Computability, inference and modeling in probabilistic programming*. Ph.D. thesis, Massachusetts Institute of Technology.
- [56] Roy, D. M. and Teh, Y. W. (2009). The Mondrian process. In *Advances in Neural Information Processing Systems*, volume 21, page 27.
- [57] Schervish, M. J. (1995). *Theory of Statistics*. Springer.
- [58] Sethuraman, J. (1994). A constructive definition of Dirichlet priors. *Statist. Sinica*, **4**(2), 639–650.
- [59] Teh, Y. W. and Jordan, M. I. (2010). Hierarchical Bayesian nonparametric models with applications. In N. L. Hjort, C. Holmes, P. Müller, and S. G. Walker, editors, *Bayesian Nonparametrics*. Cambridge University Press.
- [60] Teh, Y. W., Görür, D., and Ghahramani, Z. (2007). Stick-breaking construction for the Indian buffet process. In *Proc. of the 11th Conf. on A.I. and Stat.*.
- [61] Teh, Y. W., Blundell, C., and Elliott, L. (2011). Modelling genetic variations using fragmentation-coagulation processes. In *Adv. Neural Inf. Process. Syst.*, pages 819–827.
- [62] Thibaux, R. and Jordan, M. I. (2007). Hierarchical beta processes and the Indian buffet process. In *Proc. of the 11th Conf. on A.I. and Stat.*.
- [63] Varadarajan, V. S. (1963). Groups of automorphisms of Borel spaces. *Transactions of the American Mathematical Society*, **109**(2), pp. 191–220.
- [64] Wasserman, S. and Anderson, C. (1987). Stochastic a posteriori blockmodels: Construction and assessment. *Social Networks*, **9**(1), 1–36.
- [65] Wolfe, P. J. and Olhede, S. C. (2013). Nonparametric graphon estimation.
- [66] Xu, Z., Tresp, V., Yu, K., and Kriegel, H.-P. (2006). Infinite hidden relational models. In *Proc. of the 22nd Int. Conf. on Uncertainty in Artificial Intelligence (UAI)*.
- [67] Xu, Z., Yan, F., and Qi, Y. (2012). Infinite Tucker Decomposition: Nonparametric Bayesian Models for Multiway Data Analysis. In *Proceedings of the 29th International Conference on Machine Learning*.
- [68] Zabell, S. L. (1995). Characterizing Markov exchangeable sequences. *J. Theoret. Probab.*, **8**(1), 175–178.