ON CONTAMINATION OF SYMBOLIC DATASETS*

ANTONY PEARSON[†] AND MANUEL E. LLADSER[‡]

Abstract. Data taking values on discrete sample spaces are the embodiment of modern biological research. “Omics” experiments produce millions of symbolic outcomes in the form of reads (i.e., DNA sequences of a few dozens to a few hundred nucleotides). Unfortunately, these intrinsically non-numerical datasets are often highly contaminated, and the possible sources of contamination are usually poorly characterized. This contrasts with numerical datasets where Gaussian-type noise is often well-justified. To overcome this hurdle, we introduce the notion of latent weight, which measures the largest expected fraction of samples from a contaminated probabilistic source that conform to a model in a well-structured class of desired models. We examine various properties of latent weights, which we specialize to the class of exchangeable probability distributions. As proof of concept, we analyze DNA methylation data from the 22 human autosome pairs. Contrary to what it is usually assumed, we provide strong evidence that highly specific methylation patterns are overrepresented at some genomic locations when contamination is taken into account.

Key words. Categorical data, Contamination, DNA methylation, Exchangeability, Hypothesis testing, p-value, Symbolic data, Truthiness

AMS subject classifications. 62-07, 62F03, 62F10, 62G10, 62P10, 92D20

1. Introduction. Symbolic data is the epitome modern biological datasets due to the advent of high-throughput sequencing assays. These assays generate millions of comparatively short DNA sequences of a few dozen to a few hundred nucleotides and allow scientists to assess various microscopic processes such as investigating the relative abundance of unculturable organisms in an environment [26, 14], or pinpointing the location of the enzymes that are actively transcribing DNA into RNA along a genome [24]—among many other possibilities [35, 6, 28, 32]. Unfortunately, these datasets are often very noisy, and it is often unclear how to describe the noise because the possible sources of data corruption can be so intricate that there is little motivation for any specific and let alone universal representation of it. In contrast, Gaussian errors are often well-justified with continuous numerical data.

There is a rich history of using mixtures to describe deviations from idealized continuous models [27, 33, 16, 17]. Such mixtures are usually of the form: $P = (1 - \epsilon) \cdot Q + \epsilon \cdot R$, where P is the probabilistic source producing the data, Q is a Gaussian distribution, and R is some contaminating probability distribution from which an observation is drawn with some small probability ϵ . When the mixture is unspecified, and P ought to be estimated from data, a highly specific structure for R is usually needed, e.g. a Gaussian with known mean, to make ϵ and Q identifiable [29].

In this manuscript, we address the problem of assessing contamination (or its opposite, “purity,”) in symbolic datasets—which are intrinsically discrete. Like prior work on continuous data, we model contamination as a mixture. Unlike previous lines of work, we treat contamination as incidental; in particular, we do not commit to any prespecified form for it. To do so, we introduce the notion of latent weight with respect to a given structured class of probabilistic models (e.g., exchangeable probability distributions, which are the main focus of this manuscript). Broadly

*Submitted to the editors DATE.

Funding: This work was partially supported by the NSF Graduate Research Fellowship Program grant No. 2016198773 (Pearson), and the NSF IGERT grant No. 1144807.

[†]Department of Applied Mathematics, University of Colorado, Boulder, CO.

[‡]Department of Applied Mathematics, University of Colorado, Boulder, CO (email: manuel.lladser@colorado.edu).

speaking, this is the largest weight a model in the structured class can have as a component of the source producing the data. In particular, it describes the largest expected fraction of samples from a contaminated random sample which conform to a probabilistic model in the structured class.

We argue that latent weights are always identifiable, and allow one to represent unstructured probabilistic models as mixtures with a well-structured component; in particular, when this component carries a substantial latent weight, most samples from the mixture can be attributed to it. Latent weights offer therefore a measure of the truthiness of a hypothesis, which may not be strictly true.

To fix ideas, consider binary random variables X and Y with joint probability distribution given by the matrix

$$(1.1) \quad P := \begin{pmatrix} \overset{Y=0}{1/10} & \overset{Y=1}{3/10} \\ \underset{X=0}{1/10} & \underset{X=1}{1/2} \end{pmatrix}.$$

Since $X \sim \text{Bernoulli}(3/5)$ and $Y \sim \text{Bernoulli}(4/5)$, it is straightforward to check that X and Y are not independent. Nevertheless, we can ask whether or not P can be represented as a mixture with a component with independent marginals. This is related to determining if the latent weight of P with respect to the class $\mathcal{Q}$ of product measures of the form $(\mu \otimes \nu)$, with μ and ν probability models supported on $\{0, 1\}$, is positive or not (see Definition 2.1). It turns out the largest weight one can give to the model $\text{Bernoulli}(3/5) \otimes \text{Bernoulli}(4/5)$ formed using the marginals of (X, Y) is $5/6 \approx 83\%$. Indeed:

$$P = \frac{5}{6} \cdot \text{Bernoulli}(3/5) \otimes \text{Bernoulli}(4/5) + \frac{1}{6} \cdot \begin{pmatrix} 1/5 & 1/5 \\ 0 & 3/5 \end{pmatrix}.$$

The latent weight of P w.r.t. $\mathcal{Q}$ must be therefore at least $5/6$. In fact, a simple calculation reveals that

$$P = \frac{24}{25} \cdot \text{Bernoulli}(5/8) \otimes \text{Bernoulli}(5/6) + \frac{1}{25} \cdot \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix},$$

i.e., the latent weight of P with respect to $\mathcal{Q}$ is at least $24/25 = 96\%$. Further analysis would show that it is precisely this; in particular, up to a hidden event with 96% probability, X and Y behave independently. This finding is noteworthy for many reasons.

On one hand, perhaps unexpectedly, the marginal distributions of X and Y are not associated with the latent weight of P with respect to $\mathcal{Q}$. On the other hand, if we derived our beliefs about the independence of X and Y from a large sample from P , e.g. using a chi-squared test of independence, at moderate significance levels we would typically reject the hypothesis that X and Y are independent, missing that most of the time this is not the case. In this regard, latent weights could be used as a proxy for the truthiness of a hypothesis, without being tied to the absolutes of truth or falsity in classical hypothesis testing approaches. Finally, if P were estimated from a large but corrupted sample, as is usually the case with modern biological datasets, the high weight of $\text{Bernoulli}(5/8) \otimes \text{Bernoulli}(5/6)$ as a component of P would suggest modeling X and Y as independent—as opposed to a more complicated model with a spurious correlation induced by a seemingly 4% contamination

Paper organization. Since our primary interest is on symbolic data, we restrict ourselves to the setting of finite sample spaces. From a technical point of view, this helps to develop the theory but without trivializing it. Section 2 introduces and analyzes the notion of latent weight associated with a general class of probability models. Sections 3-5 are exclusively devoted to the class of exchangeable probability models. Section 3.1 introduces some upper-bounds that may be useful for assessing an exchangeable latent weight when the sample space is too large relative to the number of observations from the model. Section 4 develops the statistical machinery necessary for inference of exchangeable latent weights when the probabilistic source (producing the data) is observed only indirectly through a random sample. Finally, Section 5 demonstrates the use of latent weights to assess the exchangeability of DNA methylation, a near-universal assumption in epigenomic analyses.

2. Latent weights. In what follows, $\mathcal{P}$ denotes the set of all probability measures over certain finite non-empty sample space Ω . As such, $\mathcal{P}$ is compact under any norm induced metric; in particular, the total variation norm. Recall that for $P_1, P_2 \in \mathcal{P}$, this norm is defined as [22]:

$$\|P_1 - P_2\| := \max_{A \subset \Omega} |P_1(A) - P_2(A)| = \frac{1}{2} \sum_{x \in \Omega} |P_1(x) - P_2(x)|.$$

In what remains of this manuscript, $\mathcal{Q} \subset \mathcal{P}$ denotes a closed non-empty subset of probability measures. In particular, $\mathcal{Q}$ is a compact subset of $\mathcal{P}$. Instances like this include, for example, singletons, as well as models with independent marginals (when Ω is a product space), among various other possibilities.

DEFINITION 2.1. Let $P \in \mathcal{P}$. We define the (latent) weight of $\mathcal{Q}$ in P as the coefficient $\lambda_{\mathcal{Q}}(P) := \sup\{\lambda \text{ such that } P \geq \lambda \cdot Q \text{ for some } Q \in \mathcal{Q}\}$, where $P \geq \lambda \cdot Q$ means that $P(\omega) \geq \lambda \cdot Q(\omega)$, for all $\omega \in \Omega$.

This definition resembles that given in [5, 25] but for the very different purpose of representing a long-lasting Markov chain by shorter-lived independent chains.

Clearly, $0 \leq \lambda_{\mathcal{Q}}(P) \leq 1$. Latent weights have various other properties which we now state.

THEOREM 2.2. $\lambda_{\mathcal{Q}}(P) = 1$ if and only if $P \in \mathcal{Q}$.

Proof. If $\lambda_{\mathcal{Q}}(P) = 1$ then there exists a sequence of real numbers $(\lambda_n)_{n \geq 1}$ such that $1 - 1/n < \lambda_n \leq 1$, and $P \geq \lambda_n \cdot Q_n$ for some $Q_n \in \mathcal{Q}$. Without loss of generality assume that $\lambda_n < 1$. In particular, if we define $R_n := (P - \lambda_n \cdot Q_n)/(1 - \lambda_n)$ then $R_n \in \mathcal{P}$ and $P = \lambda_n \cdot Q_n + (1 - \lambda_n) \cdot R_n$. As a result: $(P - Q_n) = (1 - \lambda_n) \cdot (R_n - Q_n)$, which implies that $\|P - Q_n\| = (1 - \lambda_n) \cdot \|R_n - Q_n\| < 1/n$. Hence $P \in \mathcal{Q}$ because $\mathcal{Q}$ is closed. The converse is immediate because $\lambda_{\mathcal{Q}}(Q) = 1$ for all $Q \in \mathcal{Q}$. $\square$

THEOREM 2.3.

$$(2.1) \quad \lambda_{\mathcal{Q}}(P) = \sup_{Q \in \mathcal{Q}} \min_{\omega \in \Omega} \frac{P(\omega)}{Q(\omega)},$$

where any division by zero (including zero-over-zero) is to be interpreted as $+\infty$, and this supremum is achieved; in particular, there is $Q \in \mathcal{Q}$ such that $P \geq \lambda_{\mathcal{Q}}(P) \cdot Q$ and there is $R \in \mathcal{P}$ such that

$$(2.2) \quad P = \lambda_{\mathcal{Q}}(P) \cdot Q + (1 - \lambda_{\mathcal{Q}}(P)) \cdot R.$$

Further, if $\mathcal{Q}$ is convex and $\lambda_{\mathcal{Q}}(P) < 1$ then $\lambda_{\mathcal{Q}}(R) = 0$; in particular, $R \notin \mathcal{Q}$.

Proof. Fix $P \in \mathcal{P}$, and define $\lambda^* := \sup_{Q \in \mathcal{Q}} \min_{\omega \in \Omega} P(\omega)/Q(\omega)$. For each $Q \in \mathcal{Q}$, note that $P \geq \lambda \cdot Q$ if and only if $\min_{\omega \in \Omega} P(\omega)/Q(\omega) \geq \lambda$; in particular, $\lambda^* \geq \lambda_{\mathcal{Q}}(P)$. Define $\epsilon := \lambda^* - \lambda_{\mathcal{Q}}(P)$. If $\epsilon > 0$ then, from the definition of λ^* , there would be $Q \in \mathcal{Q}$ such that $\min_{\omega \in \Omega} P(\omega)/Q(\omega) \geq \lambda^* - \epsilon/2$. In particular, for all $\omega \in \Omega$, $P(\omega) \geq (\lambda^* - \epsilon/2)Q(\omega)$, hence $\lambda_{\mathcal{Q}}(P) \geq \lambda^* - \epsilon/2 > \lambda_{\mathcal{Q}}(P)$, a contradiction. Therefore $\epsilon = 0$ i.e. $\lambda^* = \lambda_{\mathcal{Q}}(P)$.

Next, note that for each $Q \in \mathcal{Q}$:

$$(2.3) \quad \min_{\omega \in \Omega} \frac{P(\omega)}{Q(\omega)} = \min_{\omega: Q(\omega) > 0} \frac{P(\omega)}{Q(\omega)}.$$

Further, if $\lim_{n \rightarrow \infty} Q_n = Q$ in total variation distance then $\lim_{n \rightarrow \infty} Q_n(\omega) = Q(\omega)$ uniformly for all $\omega \in \Omega$. Hence, the transformation $Q \in \mathcal{Q} \rightarrow \min_{\omega \in \Omega} P(\omega)/Q(\omega)$ from $\mathcal{Q}$ to $\mathbb{R}$ is continuous; in particular, because $\mathcal{Q}$ is compact, the supremum in equation (2.1) is achieved, and there is $Q \in \mathcal{Q}$ such that $P \geq \lambda_{\mathcal{Q}}(P) \cdot Q$. For brevity, define $\lambda := \lambda_{\mathcal{Q}}(P)$. Following an argument similar to the proof of Theorem 2.2, it follows that there is $R \in \mathcal{P}$ such that $P = \lambda \cdot Q + (1 - \lambda) \cdot R$. Likewise, if $\alpha := \lambda_{\mathcal{Q}}(R)$ then there is $Q' \in \mathcal{Q}$ and $R' \in \mathcal{P}$ such that $R = \alpha \cdot Q' + (1 - \alpha) \cdot R'$. As a result, if $\mathcal{Q}$ is convex then $P \geq \lambda \cdot Q + (1 - \lambda)\alpha \cdot Q' = (\lambda + (1 - \lambda)\alpha) \cdot Q''$ for some $Q'' \in \mathcal{Q}$. But then $\lambda \geq \lambda + (1 - \lambda)\alpha$, which implies that $(1 - \lambda)\alpha = 0$. Thus, if $\lambda < 1$ then $\alpha = 0$ as claimed. $\square$

We emphasize that the probability measure Q in Theorem 2.3 is not necessarily unique. For instance, if $\mathcal{Q}$ is the space of probability measures over $\Omega = \{0, 1\}^2$ with i.i.d. marginals, and P is the uniform probability measure over the set $\{(0, 0), (1, 1)\}$ then $\lambda_{\mathcal{Q}}(P) = 1/2$, and $\delta_{(0,0)}$ and $\delta_{(1,1)}$, the point masses at $(0, 0)$ and $(1, 1)$, respectively, both achieve the supremum in equation (2.1). Likewise, since $P = \delta_{(0,0)}/2 + \delta_{(1,1)}/2$, this counterexample shows the importance of convexity to guarantee that the probability measure R in equation (2.2) does not belong to $\mathcal{Q}$.

The identity in equation (2.2) implies that any probability model P over Ω admits a mixture representation with a component in $\mathcal{Q}$ of weight $\lambda_{\mathcal{Q}}(P)$. (This motivates the terminology of “latent weight.”) The weight of Q in P may be interpreted therefore as the largest expected fraction of observations from P that can be attributed to a single model in $\mathcal{Q}$.

We also note that a large latent weight is indicative of closeness in total variation distance to $\mathcal{Q}$. In fact, if Q and R are as in equation (2.2) then (see proof of Theorem 2.2): $\|P - Q\| = (1 - \lambda_{\mathcal{Q}}(P)) \cdot \|R - Q\|$. So, if $\lambda_{\mathcal{Q}}(P)$ is close to 1 then P is close to $\mathcal{Q}$ in total variation distance. The converse is not necessarily true, however. For instance, consider $\Omega = \{0, 1\}^d$, with finite $d > 1$, and let 0^d and 1^d denote the sequences of d zeros and ones, respectively. Define $P(x) := (2^d - 2)^{-1}$ for $x \in \Omega \setminus \{0^d, 1^d\}$. In particular, $P(0^d) = P(1^d) = 0$, and $\|P - \text{Uniform}(\{0, 1\}^d)\| = 2^{2-d}$, i.e. P is very close in total variation distance to $\mathcal{Q}$. Nevertheless, the latent weight of P with respect to the class $\mathcal{Q}$ of i.i.d. distributions over $\{0, 1\}^d$ is 0. Indeed, for all $\lambda > 0$ and $Q := \otimes_{k=1}^d \text{Bernoulli}(q)$, with $0 \leq q \leq 1$, it is not possible to have $P \geq \lambda \cdot Q$. Otherwise, because $Q(0^d) = (1 - q)^d$ and $Q(1^d) = q^d$, it would follow that $q = 1$ and $q = 0$ simultaneously.

3. Exchangeable weights. The remaining of the manuscript focuses on the class of exchangeable probabilistic models. Recall that a finite sequence of random variables $X_1, \dots, X_d$ is called exchangeable if for any permutation σ of $(1, \dots, d)$ the random vector $(X_{\sigma(1)}, \dots, X_{\sigma(d)})$ has the same distribution as $(X_1, \dots, X_d)$.

Exchangeability is a common a priori assumption in Bayesian statistics [8], permutation hypothesis testing [12], and coalescent theory [20]. The class of exchangeable models contains all finite sequences of independent and identically distributed (i.i.d.) random variables, but is generally much larger. In fact, sampling colored balls from an urn without replacement produces a random sequence of colors which is exchangeable but not necessarily independent.

In what follows, $(X_1, \dots, X_d)$, with $d > 1$ finite, is a random vector with probability distribution P . Each X_i is assumed to take values in a certain finite set $\mathcal{X}$ of cardinality $k > 1$. In addition, $\mathcal{P}$ denotes the set of all probability models over $\mathcal{X}^d$, and $\mathcal{E} \subset \mathcal{P}$ the subset of exchangeable models, which is clearly closed. We refer to the latent weight of P with respect to $\mathcal{E}$, $\lambda_{\mathcal{E}}(P)$, as the exchangeable weight of P .

DEFINITION 3.1. *For each $x \in \mathcal{X}^d$, let $[x] \subset \mathcal{X}^d$ denote the set of all vectors of the form $(x_{\sigma(1)}, \dots, x_{\sigma(d)})$, with σ a permutation of $(1, \dots, d)$.*

The set $\mathcal{X}^d$ has $\binom{k+d-1}{d}$ permutation-equivalence classes. (The number of permutation equivalence classes equals the number of ways to place d unlabelled balls in k labelled urns.) In fact, $\mathcal{E}$ is a simplex, in particular, also a convex set, with $\binom{k+d-1}{d}$ extreme points, each of which is a probability model having uniform mass over a single permutation-equivalence class [9].

In what follows, $[\mathcal{X}^d]$ denotes the set of permutation equivalence classes of $\mathcal{X}^d$. Next we show various properties of exchangeable weights, starting with the following explicit formula.

THEOREM 3.2. *For all $P \in \mathcal{P}$:*

$$(3.1) \quad \lambda_{\mathcal{E}}(P) = \sum_{x \in \mathcal{X}^d} \min_{y \in [x]} P(y) = \sum_{z \in [\mathcal{X}^d]} |z| \cdot \min_{x \in z} P(x).$$

If $\lambda_{\mathcal{E}}(P) > 0$ then exactly one probability measure $Q \in \mathcal{E}$ is associated with the exchangeable weight of P : for each $x \in \mathcal{X}^d$, $Q(x) = \min_{y \in [x]} P(y) / \lambda_{\mathcal{E}}(P)$.

Proof. Define $\lambda^* := \sum_{x \in \mathcal{X}^d} \min_{y \in [x]} P(y)$. Suppose $\beta \geq 0$ and $Q' \in \mathcal{E}$ are such that $P \geq \beta \cdot Q'$. Since exchangeability implies $Q'(y) = Q'(x)$ for all $y \in [x]$, it follows that

$$\beta \cdot Q'(x) \leq \min_{y \in [x]} P(y).$$

Moreover, because Q' is a probability measure, we have that

$$\beta = \sum_{x \in \mathcal{X}^d} \beta \cdot Q'(x) \leq \sum_{x \in \mathcal{X}^d} \min_{y \in [x]} P(y) = \lambda^*,$$

which implies that $\lambda_{\mathcal{E}}(P) \leq \lambda^*$. But observe that $P \geq \lambda^* \cdot Q$, where $Q(x) := \min_{y \in [x]} P(y) / \lambda^*$. Since Q is an exchangeable probability measure over $\mathcal{X}^d$, it follows that $\lambda_{\mathcal{E}}(P) \geq \lambda^*$, hence $\lambda_{\mathcal{E}}(P) = \lambda^*$. The second identity in equation (3.1) is now direct from this equality.

Finally, assume that $\lambda^* > 0$ and suppose that $S \in \mathcal{E}$ is such that $P \geq \lambda^* \cdot S$. We show using an argument by contradiction that $S \leq Q$. Indeed, if there were $x \in \mathcal{X}^d$ such that $S(x) > Q(x)$ then, because Q and S are exchangeable, the definition of Q would imply that there is $y \in [x]$ such that $S(y) > Q(y)$ and $\lambda^* \cdot Q(y) = P(y)$. In particular, $P(y) < \lambda^* \cdot S(y)$, which is not possible. Thus $S \leq Q$, which implies that $S = Q$ as claimed. $\square$

Due to Theorem 2.3, whenever P is not itself exchangeable i.e. $\lambda_{\mathcal{E}}(P) < 1$, P also has a unique component $R := (P - \lambda \cdot Q)/(1 - \lambda)$, which we call the unexchangeable component, which admits no exchangeable component of its own. In this sense P can be distilled entirely into its exchangeable and unexchangeable parts. This property allows one to combine an exchangeable probability model with one that is totally unexchangeable so that the resultant source has a desired exchangeable weight:

COROLLARY 3.3. *If $P = \beta \cdot Q' + (1 - \beta) \cdot R'$, where $0 \leq \beta \leq 1$, $Q' \in \mathcal{E}$, and $R' \in \mathcal{P} \setminus \mathcal{E}$ is such that $\lambda_{\mathcal{E}}(R') = 0$, then $\lambda_{\mathcal{E}}(P) = \beta$.*

Proof. Since $\lambda_{\mathcal{E}}(R') = 0$, the identity in equation (3.1) implies that for each $z \in [\mathcal{X}^d]$ there is $y \in z$ such that $R'(y) = 0$. In particular, because Q' is exchangeable, reusing equation (3.1) we obtain that:

$$\lambda_{\mathcal{E}}(P) = \sum_{x \in \mathcal{X}^d} \left(\beta \cdot Q'(x) + (1 - \beta) \cdot \min_{y \in [x]} R'(y) \right) = \beta \cdot \sum_{x \in \mathcal{X}^d} Q'(x) = \beta. \quad \square$$

3.1. Bounds on exchangeable weights. In what follows, for each non-empty $I \subset \{1, \dots, d\}$, P_I denotes the marginal distribution of $(X_i)_{i \in I}$. Clearly, P_I is exchangeable when P is exchangeable.

The following result may be useful to estimate an upper-bound on the exchangeable weight of a probabilistic source on $\mathcal{X}^d$. Indeed, estimating $\lambda_{\mathcal{E}}(P)$ when P is known only indirectly through data requires estimating $(k^d - 1)$ free parameters, which may be infeasible in practice. However, marginalizing X to a random sub-vector of dimension $s < d$, one may significantly reduce the dimension of the estimation problem.

THEOREM 3.4. *The exchangeable weight of a full d -dimensional joint distribution is a lower bound on the exchangeable weight of any marginal, i.e. $\lambda_{\mathcal{E}}(P) \leq \lambda_{\mathcal{E}}(P_I)$, for all $P \in \mathcal{P}$ and non-empty $I \subset \{1, \dots, d\}$.*

In particular, if $\lambda_{\mathcal{E}}(P_I)$ is small for some I then so is $\lambda_{\mathcal{E}}(P)$; in which case, only very few of the data produced by P could be attributed to an exchangeable source.

Proof. For each $x \in \mathcal{X}^{|I|}$ and $\alpha \in \mathcal{X}^{d-|I|}$, let $x\alpha$ be the vector $y \in \mathcal{X}^d$ such that $(y_i)_{i \in I} = x$ and $(y_i)_{i \notin I} = \alpha$. If Q is the exchangeable probability measure given in equation (3.1) then

$$P_I(x) = \sum_{\alpha \in \mathcal{X}^{d-|I|}} P(x\alpha) \geq \sum_{\alpha \in \mathcal{X}^{d-|I|}} \lambda_{\mathcal{E}}(P) \cdot Q(x\alpha) = \lambda_{\mathcal{E}}(P) \cdot Q_I(x),$$

for each $x \in \mathcal{X}^{|I|}$. Hence, since Q_I is exchangeable, $\lambda_{\mathcal{E}}(P) \leq \lambda_{\mathcal{E}}(P_I)$. $\square$

The tightness of the inequality in the theorem is related to the notion of extendibility [10]. Let $1 \leq s \leq d$. An (exchangeable) probability model μ on $\mathcal{X}^s$ is called d -extendible if there is an exchangeable probability measure ν on $\mathcal{X}^d$ and $I \subset \{1, \dots, d\}$ of cardinality s such that $\mu = \nu_I$. Necessary and sufficient conditions for extendibility can be found in [11].

COROLLARY 3.5. *For any $P \in \mathcal{P}$ with $\lambda_{\mathcal{E}}(P) > 0$, and $I \subset \{1, \dots, d\}$, if $\lambda_{\mathcal{E}}(P) = \lambda_{\mathcal{E}}(P_I)$ then the exchangeable component of P_I is d -extendible.*

Proof. Suppose that $\lambda_{\mathcal{E}}(P) = \lambda_{\mathcal{E}}(P_I) > 0$, with $|I| = s$. Let Q and $\tilde{Q}$ denote the exchangeable components of P and P_I , respectively. From the proof of Theorem 3.4, $P_I \geq \lambda_{\mathcal{E}}(P) \cdot Q_I = \lambda_{\mathcal{E}}(P_I) \cdot Q_I$; in particular, due to Theorem 3.2, $\tilde{Q} = Q_I$. $\square$

The converse in the corollary is not necessarily true, however. For a counterexample, consider $P \in \mathcal{P}(\{0, 1\}^3)$ such that $P(1, 0, 1) = P(1, 1, 0) = P(1, 1, 1) = 1/3$; in particular, $\lambda(P) = 1/3$. If $I = \{1, 2\}$ then $P_I(1, 0) = 1/3$ and $P_I(1, 1) = 2/3$, hence $\lambda(P_I) = 2/3$, and $\tilde{Q} = \delta_{(1,1)}$, which is 3-extendible to $\delta_{(1,1,1)}$ yet $\lambda(P) \neq \lambda(P_I)$.

Finally, another way to bound the exchangeable weight of a probability measure on $\mathcal{X}^d$ is to lump states in $\mathcal{X}$ as follows.

THEOREM 3.6. *Let $\mathcal{Y}$ be a finite set and $r : \mathcal{X} \rightarrow \mathcal{Y}$ a function, and define $\Phi : \mathcal{X}^d \rightarrow \mathcal{Y}^d$ as $\Phi(x) := (r(x_1), \dots, r(x_d))$. Then, for each $P \in \mathcal{P}(\mathcal{X}^d)$, $\lambda_{\mathcal{E}}(P) \leq \lambda_{\mathcal{E}}(P \circ \Phi^{-1})$, where $P \circ \Phi^{-1}$ is the forward measure of P by Φ .*

Proof. If Q denotes the exchangeable component of P then

$$\begin{aligned} \lambda_{\mathcal{E}}(P \circ \Phi^{-1}) &= \sum_{[y] \subset \mathcal{Y}^d} |[y]| \cdot \min_{z \in [y]} P(\Phi^{-1}(z)) \\ &\geq \sum_{[y] \subset \mathcal{Y}^d} |[y]| \cdot \min_{z \in [y]} \lambda_{\mathcal{E}}(P) \cdot Q(\Phi^{-1}(z)) \\ &= \lambda_{\mathcal{E}}(P) \cdot \sum_{[y] \subset \mathcal{Y}^d} |[y]| \cdot \min_{z \in [y]} Q(\Phi^{-1}(z)) = \lambda_{\mathcal{E}}(P), \end{aligned}$$

where for the very last identity we have used that Q is exchangeable; in particular, $\min_{z \in [y]} Q(\Phi^{-1}(z)) = Q(\Phi^{-1}(y))$, for each $y \in \mathcal{Y}^d$. $\square$

4. Estimation of exchangeable weights. In this section, $X_1, \dots, X_n$ denote d -dimensional i.i.d. samples from a probability measure P defined over $\mathcal{X}^d$. Define $\lambda := \lambda_{\mathcal{E}}(P)$.

Let $\hat{P}_n := \sum_{i=1}^n \delta_{X_i}/n$ denote the empirical measure associated with the sample. A natural estimator of λ is $\hat{\lambda}_n := \lambda_{\mathcal{E}}(\hat{P}_n)$. Since $\hat{P}_n$ is a maximum likelihood estimator (MLE) of P , $\hat{\lambda}_n$ is a MLE of λ . Moreover, as noted in the proof of Theorem 2.3, $\lambda_{\mathcal{E}}(\cdot)$ is continuous. In particular, since $\hat{P}_n \rightarrow P$ almost surely, $\hat{\lambda}_n \rightarrow \lambda$ also almost surely. Since $0 \leq \hat{\lambda}_n \leq 1$ for all $n \geq 1$, it follows that $\hat{\lambda}_n$ is an asymptotically unbiased estimator of λ . Nevertheless, because for each $z \in [\mathcal{X}^d]$ the transformation $P \rightarrow \min_{y \in z} P(y)$ is concave down, equation (3.1) implies that $\lambda_{\mathcal{E}}(\cdot)$ is concave down. Thus, by Jensen's inequality, $E(\hat{\lambda}_n) \leq \lambda$, i.e. $\hat{\lambda}_n$ is a negatively biased estimator of λ .

For each $x \in \mathcal{X}^d$, consider the quantities

$$m_x := \min_{y \in [x]} P(y); \quad \sigma_x^2 := m_x(1 - m_x).$$

Since these quantities remain constant within each permutation equivalence class, we sometimes abuse the notation and write for $z \in [\mathcal{X}^d]$: m_z and σ_z^2 to mean m_x and σ_x^2 , with $x \in z$, respectively.

Our next result characterizes implicitly the asymptotic distribution of $\hat{\lambda}_n$.

THEOREM 4.1. *If for each $z \in [\mathcal{X}^d]$, $C_z := \{x \in z \text{ such that } P(x) = m_x\}$ and $c_z := |C_z|$ then*

$$(4.1) \quad \lim_{n \rightarrow \infty} \sqrt{n}(\hat{\lambda}_n - \lambda) \stackrel{d}{=} \sum_{z \in [\mathcal{X}^d]} |z| \cdot \min_{x \in C_z} Z_z^{(x)},$$

where $(Z_z^{(x)})_{z \in [\mathcal{X}^d], x \in C_z}$ is a normal random vector such that, for each $z \in [\mathcal{X}^d]$, $(Z_z^{(x)})_{x \in C_z}$ is a c_z -dimensional zero-mean exchangeable normal random vector with

variance-covariance matrix Σ_z such that $\Sigma_z(x, x) = \sigma_z^2$ and $\Sigma_z(x, y) = -m_z^2$, for all $x, y \in C_z$ with $x \neq y$, and for $z_1, z_2 \in [\mathcal{X}^d]$ with $z_1 \neq z_2$, $\text{cov}(Z_{z_1}^{(x)}, Z_{z_2}^{(y)}) = -m_{z_1} m_{z_2}$, for all $x \in z_1$ and $y \in z_2$.

Proof. Our arguments follow closely those in [13].

For the sake of notation, we remove the sub-index n from quantities defined in terms of $\hat{P}_n$; in particular, we write $\hat{P}$ instead of $\hat{P}_n$ and $\hat{\lambda}$ instead of $\hat{\lambda}_n$. For $x \in \mathcal{X}^d$, define $\hat{M}_x := \min_{y \in C_x} \hat{P}(y)$ and $\hat{m}_x := \min_{y \in [x]} \hat{P}(y)$. According to the Law of Large Numbers, $\hat{P} \rightarrow P$ almost surely; in particular:

$$(4.2) \quad \mathbb{P}(\text{there is } x \in \mathcal{X}^d \text{ such that } \hat{M}_x \neq \hat{m}_x) = o(1).$$

Define $\delta := \sqrt{n}(\hat{\lambda} - \lambda)$ and $\Delta := \sqrt{n}(\sum_{[x] \subset \mathcal{X}^d} |[x]| \cdot \min_{y \in C_x} \hat{P}(y) - \lambda)$. That is, Δ is proportional to the difference between $\hat{\lambda}$ and an estimator of it computed from only those $\hat{P}(x)$ corresponding to outcomes $x \in \mathcal{X}^d$ for which $P(x) = m_x$.

Fix $t \in \mathbb{R}$. By equations (3.1) and (4.2),

$$(4.3) \quad |\mathbb{P}(\delta \leq t) - \mathbb{P}(\Delta \leq t)| = o(1).$$

Furthermore, due to the well-known Central Limit Theorem for the multinomial distribution:

$$(4.4) \quad \begin{aligned} \mathbb{P}(\Delta \leq t) &= \mathbb{P}\left(\sum_{[x] \subset \mathcal{X}^d} |[x]| \cdot \min_{y \in C_x} \sqrt{n} \cdot (\hat{P}(y) - m_x) \leq t\right) \\ &\rightarrow \mathbb{P}\left(\sum_{[x] \subset \mathcal{X}^d} |[x]| \cdot \min_{y \in C_x} Z_{[x]}^{(y)} \leq t\right), \end{aligned}$$

where $(Z_z^{(x)})_{z \in [\mathcal{X}^d], x \in C_z}$ is a zero-mean normal random vector with variance-covariance matrix as described above. The theorem is now a direct consequence of equations (4.3) and (4.4). $\square$

Let $\mathcal{P}_U \subset \mathcal{P}$ denote the set of $P \in \mathcal{P}$ such that $c_x = 1$ for each $x \in \mathcal{X}^d$, i.e. for each $z \in [\mathcal{X}^d]$ there is a unique $y \in z$ such that $P(y) = m_z$. The following result is an almost direct consequence of the previous theorem. This result also follows, albeit less directly, from the multivariate delta method [34].

COROLLARY 4.2. *If $P \in \mathcal{P}_U$ then*

$$\lim_{n \rightarrow \infty} \sqrt{n}(\hat{\lambda} - \lambda) \stackrel{d}{=} Z,$$

where Z is a zero-mean normal random variable with variance

$$V(Z) = \sum_{[x] \subset \mathcal{X}^d} |[x]|^2 \sigma_x^2 - \sum_{[x] \neq [y]} |[x]| |[y]| m_x m_y.$$

Proof. Due to the hypothesis on P ,

$$\sqrt{n}(\hat{\lambda} - \lambda) \xrightarrow{d} Z = \sum_{[x] \subset \mathcal{X}^d} |[x]| \cdot Z_{[x]},$$

where $(Z_{[x]})_{[x] \subset \mathcal{X}^d}$ is an $\binom{k+d-1}{d}$ -dimensional exchangeable normal random vector such that $E(Z_{[x]}) = 0$, $V(Z_{[x]}) = m_x(1 - m_x)$, and for $[x] \neq [y]$, $\text{cov}(Z_{[x]}, Z_{[y]}) = -m_x m_y$. In particular, Z has a normal distribution, from which the corollary follows. $\square$

Let $\mathbb{X} := (X_1, \dots, X_n)$ denote an i.i.d. sample from $P \in \mathcal{P}$, $\mathbb{X}^* := (X_1^*, \dots, X_n^*)$ denote a single resample with replacement from $\mathbb{X}$, and $\hat{\lambda}^*$ denote the exchangeable weight associated with the empirical measure $\hat{P}^* := \sum_{i=1}^n \delta_{X_i^*}/n$. In the next theorem and corollary we characterize the asymptotic distribution of the bootstrap distribution estimator of $\sqrt{n}(\hat{\lambda}^* - \hat{\lambda})$. In what follows, $Z = (Z_z^{(x)})_{z \in [\mathcal{X}^d], x \in C_z}$ denotes the normal random vector described in Theorem 4.1, which has dimension $\sum_{z \in [\mathcal{X}^d]} c_z$.

THEOREM 4.3. *Fix $t \in \mathbb{R}$, and associate with each vector $v := (v_z^{(x)})_{z \in [\mathcal{X}^d], x \in C_z}$, the function:*

$$\psi(v) := \mathbb{P} \left(\sum_{z \in [\mathcal{X}^d]} |z| \cdot \left[\min_{x \in C_z} (Z_z^{(x)} + v_z^{(x)}) - \min_{x \in C_z} v_z^{(x)} \right] \leq t \right).$$

Note that $\psi(0) = \lim_{n \rightarrow \infty} \mathbb{P}(\sqrt{n}(\hat{\lambda} - \lambda) \leq t)$, the C.D.F. of the limiting random variable described in Theorem 4.1. The bootstrap estimator of $\psi(0)$ is

$$\hat{\psi} := \mathbb{P}(\sqrt{n}(\hat{\lambda}^* - \hat{\lambda}) \leq t \mid \mathbb{X}),$$

and

$$(4.5) \quad \hat{\psi} \xrightarrow{d} \psi(Y),$$

where Y is an independent copy of Z .

The proof of this theorem resembles closely the arguments given in [13]. We first require the following lemma.

LEMMA 4.4. *Define $\hat{M}_x^* := \min_{y \in C_x} \hat{P}^*(y)$ and $\hat{m}_x^* := \min_{y \in [x]} \hat{P}^*(y)$, i.e. $\hat{m}_x^*$ is the minimum probability estimated from a bootstrap resample in each equivalence class, and $\hat{M}_x^*$ is the same, estimated only from C_x . Then*

$$\mathbb{P}(\text{there is } x \in \mathcal{X}^d \text{ such that } \hat{M}_x^* \neq \hat{m}_x^* \mid \mathbb{X}) = o_p(1).$$

Proof. If $P \in \mathcal{E}$, $\hat{M}_x = \hat{m}_x$, and $\hat{M}_x^* = \hat{m}_x^*$ with probability 1 for each x , and the claim follows.

Instead, for $P \in \mathcal{P} \setminus \mathcal{E}$, define

$$\xi := \min_{[x] \subset \mathcal{X}^d \text{ s.t. } |[x]| < c_x} \left\{ \min_{y \in [x] \setminus C_x} p_y - m_x \right\},$$

that is, the smallest difference between p_x , for $x \notin C_x$, and m_x . Note that ξ is positive. Similarly, define $\hat{\xi} := \min_{[x] \subset \mathcal{X}^d \text{ s.t. } |[x]| < c_x} \{\min_{y \in [x] \setminus C_x} \hat{p}_y - \hat{M}_x\}$. (If $\hat{\xi} < 0$, then $\hat{M}_x > \hat{m}_x$ for some $[x]$.) The Law of Large Numbers guarantees that $\mathbb{P}(\hat{\xi} \geq \xi/2) = (1 - o(1))$.

Define $Y_n := \mathbb{P}(\text{there is } x \in \mathcal{X}^d \text{ such that } \hat{M}_x^* \neq \hat{m}_x^* \mid \mathbb{X})$, a random variable taking values in $[0, 1]$. Fix $\epsilon > 0$. Then

$$\mathbb{P}(Y_n \geq \epsilon) = \mathbb{P}(Y_n \geq \epsilon \mid \hat{\xi} \geq \xi/2) \cdot (1 - o(1)) + \mathbb{P}(Y_n \geq \epsilon \mid \hat{\xi} < \xi/2) \cdot o(1). \quad \square$$

A second application of the Law of Large Numbers now yields $\mathbb{P}(Y_n \geq \epsilon \mid \hat{\xi} \geq \xi/2) \rightarrow 0$, proving the lemma.

Proof of Theorem 4.3. Define

$$\begin{aligned}\delta^* &:= \sqrt{n}(\hat{\lambda}^* - \hat{\lambda}) \\ \Delta^* &:= \sqrt{n} \left(\sum_{z \in [\mathcal{X}^d]} |z| \cdot (\hat{M}_z^* - \hat{M}_z) \right).\end{aligned}$$

Due to Lemma 4.4:

$$(4.6) \quad \mathbb{P}(\delta^* \leq t \mid \mathbb{X}) = \mathbb{P}(\Delta^* \leq t \mid \mathbb{X}) + o_p(1).$$

Then it follows:

$$\begin{aligned}(4.7) \quad \mathbb{P}(\Delta^* \leq t \mid \mathbb{X}) &= \mathbb{P} \left(\sqrt{n} \sum_{z \in [\mathcal{X}^d]} |z| \cdot (\hat{M}_z^* - \hat{M}_z) \leq t \mid \mathbb{X} \right) \\ &= \mathbb{P} \left(\sum_{z \in [\mathcal{X}^d]} |z| \cdot \left[\min_{x \in C_z} \{ \sqrt{n}(\hat{p}_x^* - \hat{p}_x) + \sqrt{n}(\hat{p}_x - m_z) \} \right. \right. \\ &\quad \left. \left. - \min_{x \in C_z} \sqrt{n}(\hat{p}_x - m_z) \right] \leq t \mid \mathbb{X} \right) \\ (4.8) \quad &= \psi \left((\sqrt{n}[\hat{p}_x - m_z])_{z \in [\mathcal{X}^d], x \in C_z} \right) + o_p(1).\end{aligned}$$

The last equality is due to the fact that, for almost every sample sequence $X_1, X_2, \dots$, $\sqrt{n}(\hat{P}^* - \hat{P})$ has the same conditional limiting distribution given $\mathbb{X}$ that $\sqrt{n}(\hat{P} - P)$ does unconditionally, because $\hat{P}$ has finite second moments and mean P [4, Theorem 2.2]. That is, for any $s \in \mathbb{R}^{k^d}$, $\mathbb{P}(\sqrt{n}(\hat{P}^* - \hat{P}) \leq s \mid \mathbb{X}) \rightarrow \lim_{n \rightarrow \infty} \mathbb{P}(\sqrt{n}(\hat{P} - P) \leq s)$ almost surely (and therefore in probability). For this reason, we can replace $\sqrt{n}(\hat{P}^* - \hat{P}) \mid \mathbb{X}$ in line (4.7) by its limiting random vector, which introduces some random perturbation to the conditional C.D.F. which converges in probability to 0.

Finally, because $\psi(v)$ is a continuous function of v , we take equation (4.8) with equation (4.6) to find that

$$\mathbb{P}(\delta^* \leq t \mid \mathbb{X}) = \psi \left((\sqrt{n}[\hat{p}_x - m_z])_{z \in [\mathcal{X}^d], x \in C_z} \right) + o_p(1) \xrightarrow{d} \psi((Y_z^{(x)})_{z \in [\mathcal{X}^d], x \in C_z}),$$

as claimed. $\square$

The above convergence is weak, however, by requiring $P \in \mathcal{P}_U$, we can ensure convergence in probability of the bootstrap distribution estimator. The following corollary follows simply from the fact that when $c_z = 1$ for each $z \in [\mathcal{X}^d]$, $\min_{x \in C_z} (Z_z^{(x)} - Y_z^{(x)}) - \min_{x \in C_z} Y_z^{(x)} = Z_z^{(x)}$.

COROLLARY 4.5. *When $P \in \mathcal{P}_U$,*

$$\mathbb{P}(\sqrt{n}(\hat{\lambda}^* - \hat{\lambda}) \leq t \mid \mathbb{X}) \xrightarrow{p} \mathbb{P}(\sqrt{n}(\hat{\lambda} - \lambda) \leq t),$$

for each $t \in \mathbb{R}$.

Let $\mathcal{P}_N := \mathcal{P} \setminus \mathcal{P}_U$ denote the set of sources which have at least one permutation class $z \in [\mathcal{X}^d]$ where $c_z > 1$, i.e., sources for which at least one m_z is not uniquely achieved by $y \in z$. When $P \in \mathcal{P}_N$, a more explicit characterization of the distribution

of Z seems very elusive, and in particular, is not Gaussian. To see why, observe that [2, Corollary 5] implies that $\min_{x \in C_z} Z_z^{(x)}$ has probability density function (p.d.f.):

$$f_z(t) = \frac{c_z}{\sigma_z} \cdot \varphi\left(\frac{t}{\sigma_z}\right) \cdot \Phi_{c_z-1}\left(\frac{t\sqrt{m_z}}{\sigma_z^2}, \dots, \frac{t\sqrt{m_z}}{\sigma_z^2}; \rho_x \mathbb{I}_{c_z-1} + \frac{\Sigma_x}{\sigma_z^2}\right),$$

where $\varphi(\cdot)$ is the p.d.f. of a standard Normal random variable, $\Phi_k(\cdot; \Sigma)$ is the cumulative distribution function (c.d.f.) of a zero-mean k -dimensional multivariate normal distribution with variance-covariance matrix Σ , and $\mathbb{I}_k$ is the k -dimensional identity matrix.

Unfortunately, the above probability densities are not enough to describe the distribution of Z due to the correlation between the minima in equation (4.1). Furthermore, Theorem 4.3 does not hold when $P \in \mathcal{P}_N$, because whenever $c_x > 1$, $\min_{x \in C_z} (Z_z^{(x)} + Y_z^{(x)}) - \min_{x \in C_z} Y_z^{(x)} \neq \min_{x \in C_z} Z_z^{(x)}$ with probability 1. As a result, the weak convergence in the last line of equation (4.8) is to a version of the distribution of Z , but with some Gaussian perturbation.

Luckily, however, the bootstrap estimator of Z when $P \in \mathcal{P}_U$ can be made consistent by choosing a resample size n_0 of order $o(n)$. In this case, we would redefine $\delta^* := \sqrt{n_0}(\hat{\lambda}_n^* - \hat{\lambda})$ and $\Delta^* := \sqrt{n_0}(\sum_{z \in [\mathcal{X}^d]} |z| \cdot (\hat{M}_z^* - \hat{M}_z))$, so that

$$\begin{aligned} \hat{\psi} &= \mathbb{P}(\delta^* \leq t \mid \mathbb{X}) = \mathbb{P}(\Delta^* \leq t \mid \mathbb{X}) + o_p(1) \\ &= \psi\left(\left(\sqrt{n_0}(\hat{p}_x - p_x)_{z \in [\mathcal{X}^d], x \in C_z}\right)\right) + o_p(1) \\ &\xrightarrow{P} \psi(0), \end{aligned}$$

because $\sqrt{n_0}(\hat{p}_x - p_x) \xrightarrow{P} 0$.

Based on simulations, we have found that it is usually more accurate to use a full size- n resample for the purposes of correcting the bias of $\hat{\lambda}_n$, even at moderate sample sizes and when $P \in \mathcal{P}_N$. Additionally, Monte Carlo bootstrap estimates of $V(\hat{\lambda})$ tend to be more accurate than the asymptotic formula given in derived from Corollary 4.2. With very large samples, users may wish to try using a size $n_0 := 2\sqrt{n}$ resample. Users might also explore ad hoc methods for combined estimators of $\min_{y \in [x]} p_y$ when there is strong reason to believe that $P \in \mathcal{P}_N$. In what follows in this manuscript, bias and variance of $\hat{\lambda}$ are approximated from a Monte Carlo estimate of the bootstrap distribution $\sqrt{n}(\hat{\lambda}_n^* - \hat{\lambda}) \mid \mathbb{X} \stackrel{d}{\approx} \sqrt{n}(\hat{\lambda} - \lambda)$.

An explicit Berry-Esseen type bound on the error using the limiting normal distribution to approximate the sampling distribution of $\hat{\lambda}_n$ remains elusive, therefore we recommend selecting sample size by simulation. The largest difficulty in estimating $\hat{\lambda}$ is controlling negative bias, especially when $P \in \mathcal{P}_N$. Therefore we recommend to simulate data from several test sources $T \in \mathcal{E}$ to approximate a worst-case sampling distribution for a given sample size. In particular, it is suitable to choose T with $T(\{a\}^d) = 0$ for each $a \in \mathcal{X}$, and uniform mass elsewhere, because each of the outcomes $\{a\}^d$ belongs to a singleton equivalence class. Any observation of these outcomes can only increase the estimate of $\lambda(T)$. In our experience on a variety of sample spaces $\mathcal{X}^d$, estimation of $\lambda(T)$ has the largest bias and standard deviation of any source in $\mathcal{P}(\mathcal{X}^d)$.

We suggest the following heuristic: first, select several candidate sample sizes $(n_i)_{i=1,2,\dots}$. For each candidate sample size, repeatedly simulate n_i outcomes from the test source T described above to get an empirical estimate of standard error and

Fig. 1: Diagram of a typical WGBS experiment. The blue rectangle represents a segment of ssDNA, with the location of CpGs on that strand marked by black vertical bars. Black horizontal line segments represent reads mapped to a reference strand, open and closed circles represent the partially-observed joint methylation status several CpGs.

bias of $\lambda(\hat{T})$ (the test source above satisfies $\lambda(T) = 1$). After selecting a sample size for which standard deviation and bias of $\hat{\lambda}_{n_i}(T)$ appear acceptably small, collect samples of this size for every source in a coarse grid over $\mathcal{P}(\mathcal{X}^d)$. This is to ensure that even on the most pathological sources we can obtain acceptable estimates of λ .

5. DNA methylation analysis. When a DNA sequence contains a cytosine residue (C) followed by a guanine residue (G) in the 5'-to-3' sense, this dimer is referred to as a CpG. The cytosine in a CpG may or may not have methyl group bonded to it in the 5' position of its pyrimidine ring. This methylation is regulated by reversible enzymatic processes and is known to modulate gene expression; increased methylation in gene promoters is associated with transcriptional silencing [19], and specific DNA methylation patterns have been linked to human disease [30]. In particular, certain aberrant methylation patterns are a hallmark of some cancers [18], and as such considerable effort has been expended to determine regions of DNA that have differential methylation under different cellular conditions.

One popular modern assay to assess DNA methylation is Whole-Genome Bisulfite Sequencing (WGBS) [23], a procedure in which unmethylated cytosines are chemically transformed into thymine (T) through treatment with a bisulfite catalyst. When bisulfite-treated DNA is then sequenced by high-throughput shotgun technology, methylated CpGs can be distinguished from unmethylated ones by the observation of a “CG” dimer versus a “TG” dimer, as depicted in Figure 1.

When attempting to describe DNA methylation, it is routine to use a sliding window approach [15, 1] wherein all observations of methylated and unmethylated CpGs are counted in the window, typically 1 Kb in length, to summarize local methylation. It is common to compare methylation between two different biological samples using, e.g., Fisher’s exact test [1], concluding that a window is differentially methylated if the null hypothesis of equal distribution can be rejected. This approach assumes that in a single window the methylation status of each CpG contributes identically in its biological effect. For example, if we denote an unmethylated CpG as a ‘0’ and a methylated CpG as a ‘1’ and consider ten consecutive CpGs in different tissues, with

the first tissue always producing the configuration ‘1111100000’ and the second always producing the configuration ‘0000011111’, the above approach would be unable to identify this locus as differentially methylated. So current approaches for differential methylation implicitly assume that binary sequences representing methylation inside each window have an exchangeable distribution.

To evaluate this assumption, we examined WGBS data from 121 experimental replicates representing 77 unique biological samples, publicly available from ENCODE [3, 7]. These replicates include clinical tissue samples, cell lines, and primary cells. We selected replicates using single-end reads which were not flagged by ENCODE as having low coverage or insufficient read length. The list of the sample identifiers (`ENCODE.IDs.xlsx`) and processed datasets can be found on GitHub (<https://github.com/antonypearson/OnContaminationofSymbolicDatasets>).

Each replicate is associated with a BAM file generated by mapping reads to GRCh38 using Bismark [21]. We used the MethPipe methylation software suite [31] to convert BAM files into MethPipe format and generate epiread files, an efficient format reporting the genomic index and methylation status of each CpG contained in a read.

To investigate the exchangeability of local DNA methylation we focused on sets of 3 consecutive CpGs, which we call “triplets.” Due to Theorem 3.4, a genomic region containing a highly unexchangeable triplet must have highly unexchangeable methylation overall.

For each replicate, we used the epireads file generated to extract data from “well-covered” triplets—i.e. those where all three CpGs are jointly covered by at least 100 reads, discarding reads which report a CpG with ambiguous methylation status. In all datasets, we observed 637,612 well-covered triplets, representing 72,815 unique loci. For each autosome in each sample we estimated the exchangeable weight of each well-covered triplet and corrected for estimator bias using a sample mean of $N = 1000$ full bootstrap resamples. Although each estimate of a triplet’s exchangeable weight $\hat{\lambda}$ lies in $[0, 1]$, the bias-adjusted estimate $(\hat{\lambda} - \hat{\lambda}^*)$ may be larger than 1 or smaller than 0. Therefore we truncate these estimates to $[0, 1]$. Available online are Numpy files containing processed triplets corresponding to each BAM file ID. Each row corresponds to a well-covered triplet, with columns corresponding to 1) chromosome number, 2) index of the triplet on the chromosome, 3) an estimate of the total variation distance to the class of exchangeable distributions, 4) an estimate of the exchangeable weight of the triplet (bias-corrected), 5) a bootstrap estimate of the standard deviation of $\hat{\lambda}$, 6-13) the counts of each of the 8 possible triplet configurations (ordered lexicographically, i.e. ‘000’, ‘001’, etc.), and 14-21) an estimate of the largest exchangeable component.

Estimates of triplet exchangeable weight are depicted in Figure 2. As seen in the figure, in some chromosomal regions, particularly e.g. on chromosomes 6 and 13, there are triplets whose exchangeable weights are very small. In fact, some appear completely unexchangeable.

As seen in Table 1, triplet exchangeability does not appear strongly correlated with the genomic distance to the nearest promoter. Further, as seen in Figure 3, within each chromosome, and within each dataset, the correlation between triplet exchangeability and distance from a promoter is usually small. There is a noticeable trend, however, that triplet exchangeable weight is more likely to be negatively correlated with distance from a promoter. Indeed, both a two-sided Wald test and a Spearman rank-order test of the null hypothesis that TSS proximity and estimated

Fig. 2: Estimated exchangeable weights of well-covered triplets by chromosome. Dashed red lines denote the mean, and green plots the histograms associated with these weights.

triplet exchangeable weight are uncorrelated give very small p-values ($p \ll 10^{-10}$). That is, despite the small magnitude of the effect, we can detect that triplets close to promoters tend to have more-exchangeable methylation configurations.

Table 1: Correlation per chromosome between estimated exchangeable weight of well-covered triplets, and distance between their center and the nearest transcription start site (TSS).

Chromosome	Correlation	Chromosome	Correlation
1	-1.200×10^{-1}	12	$+2.676 \times 10^{-2}$
2	$+1.236 \times 10^{-1}$	13	-2.406×10^{-1}
3	$+4.738 \times 10^{-2}$	14	-2.977×10^{-1}
4	-1.742×10^{-1}	15	-1.999×10^{-1}
5	-1.584×10^{-1}	16	$+4.356 \times 10^{-2}$
6	$+7.900 \times 10^{-2}$	17	$+9.088 \times 10^{-2}$
7	$+7.423 \times 10^{-2}$	18	$+8.553 \times 10^{-2}$
8	$+2.869 \times 10^{-2}$	19	-6.973×10^{-2}
9	$+2.586 \times 10^{-2}$	20	-8.210×10^{-2}
10	-3.123×10^{-2}	21	$+6.563 \times 10^{-2}$
11	$+5.987 \times 10^{-2}$	22	-2.552×10^{-1}

Fig. 3: Top, plot of estimated exchangeable weight for each well-covered triplet versus its distance from the nearest TSS. Bottom, correlation between distance to the nearest TSS and exchangeable weight of each triplet, per chromosome (left) and per dataset (right).

To confirm that these highly unexchangeable loci are not due to uncertainty in estimation of the exchangeable weight, we simulated data from the uniform distribution over binary triplets except ‘000’ or ‘111’. Based on empirical study, this source is the worst case for estimating the exchangeable weight in terms of bias and standard error. Nevertheless, as seen in Figure 4, the empirical distribution of estimated exchangeable weights of all triplets gives much greater probability mass near 0 than the corresponding sampling distribution of the synthetic data. That is, uncertainty from statistical estimation does not account for the apparent phenomenon of highly unexchangeable loci.

In general, it is impossible to disentangle contamination which is caused by, e.g., sequencing errors or incomplete enzymatic conversion of unmethylated cytosines, from biological processes discriminating specific configurations of methylation. Under the assumption that contamination of the former kind is small, i.e. that we have a truly accurate picture of how methylation is configured in cells, we would expect triplets to have exchangeable weights close to one if overall methylation levels govern biological function. This might mean that in some cell types methylation far away from

Fig. 4: Histogram of estimated exchangeable weights of all triplets from aggregated WGBS data (green) and synthetic samples of size $n = 100$ from the uniform distribution on $\{(0, 0, 1), (0, 1, 0), (0, 1, 1), (1, 0, 0), (1, 0, 1), (1, 1, 0)\}$ (red). The averages of estimated exchangeable weights from real and synthetic data are given by the dashed green and red lines, respectively.

promoters (and likely far from CpG islands) is “locked in,” and specific patterns of methylation rather than overall methylation levels modulate biological function.

Identifying the biological reason for highly unexchangeable loci remains an open question, which may not have a universal answer. We conclude that there are some loci which are far from exchangeable—that is, some configurations of methylation are discriminated at these triplets. The identification of these loci opens opportunities for more high-resolution understanding of methylation patterns. In particular, these loci represent regions where very specific configurations of methylation may regulate function.

REFERENCES

- [1] A. AKALIN, M. KORMAKSSON, S. LI, F. E. GARRETT-BAKELMAN, M. E. FIGUEROA, A. MELNICK, AND C. E. MASON, *methylKit: a comprehensive R package for the analysis of genome-wide DNA methylation profiles*, Genome Biology, 13 (2012), p. R87.
- [2] R. B. ARELLANO-VALLE AND M. G. GENTON, *On the exact distribution of the maximum of absolutely continuous dependent random variables*, Statistics & Probability Letters, 78 (2008), pp. 27–35.
- [3] B. BERNSTEIN, E. BIRNEY, I. DUNHAM, E. GREEN, C. GUNTER, M. SNYDER, ENCODE PROJECT CONSORTIUM, AND T. HUBBARD, *An integrated encyclopedia of DNA elements in the human genome*, NATURE, 489 (2012), pp. 57–74.
- [4] P. J. BICKEL AND D. A. FREEDMAN, *Some asymptotic theory for the bootstrap*, Ann. Statist., 9 (1981), pp. 1196–1217.
- [5] S. CHESTNUT AND M. E. LLADSER, *Occupancy distributions via Doeblin’s ergodicity coefficient*, in Discrete Mathematics and Theoretical Computer Science Proceedings, vol. AM, 2010, pp. 79–92.
- [6] L. J. CORE, J. J. WATERFALL, AND J. T. LIS, *Nascent RNA sequencing reveals widespread*

- pausing and divergent initiation at human promoters*, Science, 322 (2008), pp. 1845–1848.
- [7] C. A. DAVIS, B. C. HITZ, C. A. SLOAN, E. T. CHAN, J. M. DAVIDSON, I. GABDANK, J. A. HILTON, K. JAIN, U. K. BAYMURADOV, A. K. NARAYANAN, K. C. ONATE, K. GRAHAM, S. R. MIYASATO, T. R. DRESZER, J. S. STRATTAN, O. JOLANKI, F. Y. TANAKA, AND J. M. CHERRY, *The encyclopedia of DNA elements (ENCODE): data portal update*, Nucleic Acids Research, 46 (2017), pp. D794–D801.
- [8] B. DE FINETTI, *La prévision: ses lois logiques, ses sources subjectives*, Annales de l’institut Henri Poincaré, 7 (1937), pp. 1–68.
- [9] P. DIACONIS, *Finite forms of de Finetti’s theorem on exchangeability*, Synthese, 36 (1977), pp. 271–281.
- [10] P. DIACONIS AND D. FREEDMAN, *Finite exchangeable sequences*, Ann. Probab., 8 (1980), pp. 745–764.
- [11] A. V. GNEDIN, *A class of exchangeable sequences*, Statistics & Probability Letters, 28 (1996), pp. 159 – 164.
- [12] P. I. GOOD, *Extensions of the concept of exchangeability and their applications*, Journal of Modern Applied Statistical Methods, 1 (2002), p. 34.
- [13] P. HALL, W. HÄRDLE, AND L. SIMAR, *On the inconsistency of bootstrap distribution estimators*, CORE Discussion Papers RP 1062, Université catholique de Louvain, Center for Operations Research and Econometrics (CORE), 1993, <https://EconPapers.repec.org/RePEc:cor:louvnp:1062>.
- [14] J. HAMPTON AND M. E. LLADSER, *Estimation of Distribution Overlap of Urn Models*, PLoS ONE, 7 (2012), p. e42368.
- [15] K. D. HANSEN, B. LANGMEAD, AND R. A. IRIZARRY, *BSmooth: from whole genome bisulfite sequencing reads to differentially methylated regions*, Genome Biology, 13 (2012), p. R83.
- [16] P. J. HUBER, *Robust estimation of a location parameter*, Ann. Math. Statist., 35 (1964), pp. 73–101.
- [17] P. J. HUBER, *A robust version of the probability ratio test*, Ann. Math. Statist., 36 (1965), pp. 1753–1758.
- [18] P. A. JONES, *Functions of DNA methylation: islands, start sites, gene bodies and beyond*, Nature Reviews Genetics, 13 (2012), pp. 484–492.
- [19] P. L. JONES, G. C. J. VEENSTRA, P. A. WADE, D. VERMAAK, S. U. KASS, N. LANDSBERGER, J. STROUBOULIS, AND A. P. WOLFFE, *Methylated dna and mecp2 recruit histone deacetylase to repress transcription*, Nature genetics, 19 (1998), p. 187.
- [20] J. F. C. KINGMAN, *On the genealogy of large populations*, Journal of Applied Probability, 19 (1982), p. 27–43.
- [21] F. KRUEGER AND S. R. ANDREWS, *Bismark: a flexible aligner and methylation caller for Bisulfite-Seq applications*, Bioinformatics, 27 (2011), pp. 1571–1572.
- [22] T. LINDVALL, *Lectures on the Coupling Method*, Wiley Series in Probability and Statistics - Applied Probability and Statistics Section, Wiley, 1992.
- [23] R. LISTER, R. C. O’MALLEY, J. TONTI-FILIPPINI, B. D. GREGORY, C. C. BERRY, A. H. MILLAR, AND J. R. ECKER, *Highly integrated single-base resolution maps of the epigenome in arabidopsis*, Cell, 133 (2008), pp. 523–536.
- [24] M. E. LLADSER, J. G. AZOFEIFA, M. A. ALLEN, AND R. D. DOWELL, *RNA Pol II transcription model and interpretation of GRO-seq data*, Journal of Mathematical Biology, 74 (2017), pp. 77–97.
- [25] M. E. LLADSER AND S. CHESTNUT, *Approximation of sojourn-times via maximal couplings: Motif frequency distributions*, Journal of mathematical biology, 69 (2013).
- [26] M. E. LLADSER, R. GOEUE, AND J. REEDER, *Extrapolation of Urn Models via Poissonization: Accurate Measurements of the Microbial Unknown*, PLoS One, 6 (2011).
- [27] S. NEWCOMB, *A generalized theory of the combination of observations so as to obtain the best result*, American Journal of Mathematics, 8 (1886), pp. 343–366.
- [28] P. J. PARK, *Chip-seq: advantages and challenges of a maturing technology*, Nature reviews genetics, 10 (2009), p. 669.
- [29] A. PUNZO AND P. D. MCNICHOLAS, *Parsimonious mixtures of multivariate contaminated normal distributions*, Biometrical Journal, 58 (2016), pp. 1506–1537.
- [30] K. D. ROBERTSON, *DNA methylation and human disease*, Nature Reviews Genetics, 6 (2005), p. 597.
- [31] Q. SONG, B. DECATO, E. E. HONG, M. ZHOU, F. FANG, J. QU, T. GARVIN, M. KESSLER, J. ZHOU, AND A. D. SMITH, *A reference methylome database and analysis pipeline to facilitate integrative and comparative epigenomics*, PLoS ONE, 8 (2013), p. e81148.
- [32] M. M. SUZUKI AND A. BIRD, *DNA methylation landscapes: provocative insights from epigenomics*, Nature Reviews Genetics, 9 (2008), p. 465.

- [33] J. W. TUKEY, *A Survey of Sampling From Contaminated Distributions*, Contributions to Probability and Statistics (in: I. Olkin et al., eds), (1960), pp. 448–485.
- [34] A. W. VAN DER VAART, *Asymptotic statistics*, Cambridge Series in Statistical and Probabilistic Mathematics, Cambridge University Press, 1998.
- [35] Z. WANG, M. GERSTEIN, AND S. M., *RNA-Seq: a revolutionary tool for transcriptomics*, Nat Rev Genet, 10 (2009), pp. 57–63.