Universal Spatial Distribution Approximation in Equilibrium Lattice Models with Local Pairwise Interactions

Cameron Chalk
California Institute of Technology, Pasadena, CA, USA

Erik Winfree
California Institute of Technology, Pasadena, CA, USA

Abstract

Universality is one of the defining hallmarks of a useful neural network model: it means that the model is expressive enough, in principle, to learn arbitrary target behavior rather than being limited to a narrow class. Inspired by rigorous formal analogies between stochastic neural networks and lattice models in which molecules occupy discrete sites, contribute species-dependent interaction energies with neighbors, and rearrange according to the Boltzmann distribution, we ask whether interactions can be programmed to yield arbitrary spatial distributions. At the level of fine-grained spatial arrangements of species, universality fails: arbitrary spatial microstate distributions cannot be represented, as shown by a counting argument and explicit construction of unrepresentable distributions. Yet, for a natural observable – spatial patterns of molecular labels – universality can be achieved. We show that a grand-canonical lattice model of anisotropic molecules is universal for arbitrary label distributions, and a similar model with isotropic molecules is also universal given a simple symmetry-breaking boundary condition. Our results suggest that equilibrium systems of multivalent molecules undergoing stochastic rearrangement governed by local pairwise interactions can act similarly to probabilistic neural networks: with sufficient microstate detail in hidden dimensions, they can be universal in a meaningful macrostate observable.

2012 ACM Subject Classification Hardware → Neural systems; Applied computing → Physics; Applied computing → Systems biology

Keywords and phrases universality, molecular programming, neural networks, self-assembly

Digital Object Identifier 10.4230/LIPIcs.DNA.32.11

Funding *Cameron Chalk*: NSF CCF/FET 2212546
Erik Winfree: NSF CCF/FET 2212546

Acknowledgements The authors would like to thank David Doty and Damien Woods for tangential yet inspirational discussions.

1 Introduction

Life is not a homogeneous solution of chemicals. Living cells are cities, with roads [37], boundaries [38], factories [31], warehouses [3], scaffolds [11], and gradients [15], all features that change throughout the lifetime of the cell. Control over *where* molecules are helps transform ordinary chemistry into spatially-organized molecular systems capable of replication, evolution, adaptability, computation, and homeostasis. Concretely, control of molecular spatial distribution allows the cell to regulate reaction rates by colocalizing reactants [25], compartmentalize strictly regulated biomolecule populations into membraned or membraneless organelles [4], direct transport of cargo along tracks [21], change its geometry via the cytoskeleton [8], and break symmetries for morphogenetic and developmental purposes [23]. It follows that constructing prescribed molecular spatial distributions is essential for building synthetic systems with lifelike behavior.

Engineering accomplishments suggest that the combination of programmable attractive forces and random Brownian motion enables molecularly precise assembly. Thermodynamic self-assembly methods such as DNA origami [33] program molecular binding interactions towards a free energy minimum with (ideally) a high yield of the single target structure. Multivalent DNA nanostars with programmable interactions condense into aggregates with short-timescale reconfiguration [6]. However, spatial arrangement in the cell undergoes dynamic reorganization in response to, e.g., environmental changes [30]. Although engineered self-assembly can be dynamically responsive (see, e.g., reconfigurable DNA origami [40] and dynamic nanostar condensates [2]), approaches are currently specialized for particular tasks. What general programming framework might allow the design of molecules capable of complex, adaptable, and reconfigurable spatial arrangements – putting the right pieces in the right places under the right contexts?

Prior work [28, 29, 7] has situated programming of molecular interactions (chemical reactions, binding, diffusion) in an equilibrium setting as analogous to the training and usage of Boltzmann machines [1], a stochastic, statistical mechanics-inspired neural network model that stands as a precursor to modern artificial neural network approaches. When formalized in terms of local pairwise interactions of molecules in a spatial lattice [7], the model and learning framework was called *Boltzmann liquids*, and is summarized in Figure 1. A key element learned from classical Boltzmann machines is that stochastically generated high-dimensional probability distributions can be probed simply by clamping known variables, which results in probabilistic conditioning of the remaining variables. For example, a Boltzmann machine that is trained on the MNIST handwritten digit dataset can both classify input digits and generate digits given a class, and both of these tasks are achieved by conditioning the distribution. Figure 1(a) shows conditioning of a toy probability distribution to give intuition about how conditioning affects a probability distribution. In the Boltzmann liquids framework, a system of pairwise interacting molecules, typically at low molecular counts (e.g., tens to thousands of molecules as is the case in some biological condensates [18]), at equilibrium takes on a probability distribution of fine-grained molecular arrangements, as shown in Figure 1(b), left. If an environmental stimulus is introduced to the system, like a surface with a fixed arrangement (Figure 1(b), middle), or the presence of a polymer with a certain sequence of domains (Figure 1(b), right), the system assumes a conditional distribution. In other words, the molecules can store information via their interaction energies, and thus their distribution, that can be "accessed" or "queried" by stimuli that condition the distribution. Additionally, work on Boltzmann machines inspired a contrastive learning rule for Boltzmann liquids, as shown in Figure 1(c), which finds parameters for the system to match a target distribution. This framework enabled the design and simulation of systems of pairwise interacting molecules which performed Hopfield network-like associative recall and classification of MNIST handwritten digit patterns expressed as volume fractions of the constituent species types. In short, programming a stochastic system to take on a complex, high-dimensional distribution is an effective method of programming, which is a perspective supported by the success of stochastic neural network models like the Boltzmann machine.

In the study of neural networks, machine learning, and statistical modeling, the expressivity of the model used is a key question; single-layer perceptrons, for example, can only classify linearly-separable data [32], and thus multi-layer perceptrons were suggested for more useful classification. The multi-layer perceptron is known to be *universal* [9]. Universality means that a model can effectively represent or approximate any distribution or input-output relation, usually with some number of caveats. A proof of universality of a model provides assurance that the model does not have any fundamental limits inhibiting a classification or

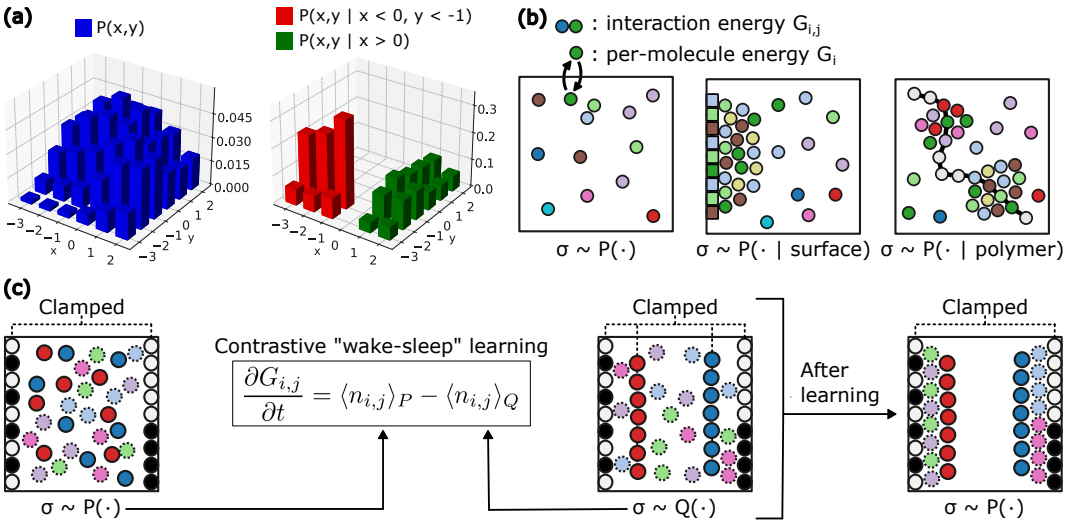

Figure 1 A cartoon summary of Boltzmann liquids. **(a)** On the left, a two-dimensional probability mass function (PMF). On the right, two different conditional distributions generated from the same PMF. **(b)** A cartoon representation of an equilibrium pairwise interaction model. Near-enough molecules of species i and j contribute interaction energy $G_{i,j}$, and exchange with a reservoir of molecules occurs dependent on a per-molecule energy (chemical potential) G_i, where more negative $G_{i,j}$ and G_i are more favorable. On the left, a sample configuration σ drawn from the unconditional Boltzmann distribution P, indicating that in P the molecules are dilute with high probability. In the middle, a surface on the left wall is fixed in space, inducing a conditional distribution $P(\cdot \mid \text{surface})$ in which condensation occurs. On the right, the presence of a polymer in the system induces a conditional distribution $P(\cdot \mid \text{polymer})$ in which condensation of different compositions occurs depending on the polymer's sequence. **(c)** A learning algorithm finds the parameter gradient that makes a distribution more like a target. In this cartoon, there is no molecular exchange with the reservoir, i.e., there is a fixed count of each species. In this example, *clamped* means the molecules are fixed in space; more general notions of clamping to an observable have been studied [7]. On the left, the colored molecules are left unclamped. In the middle, the distribution Q to be learned is sampled – Q is a distribution defined by the current parameters and clamping of a target marginal distribution Q' that specifies where the red and blue molecules should be but does not specify where the "hidden" molecules shown by dotted circles should be. Average statistics $\langle n_{i,j} \rangle$ are collected from each distribution, where $n_{i,j}$ counts the number of $i - j$ interactions. The learning algorithm yields parameters (if possible) which make the system behave more like Q as shown on the right, where the hidden species learn to help organize the red and blue species according to Q'.

generative task, and thus leaves only the remaining questions of required model size, time spent learning (i.e., how much of the data the model sees), and correctness of its learning algorithm. Well-mixed stochastic chemical reaction networks have been shown to be universal for distributions of counts [5], establishing a baseline for viewing molecular systems in terms of probabilistic machine learning theory.

If equilibrium systems are indeed universal for spatial distributions, a consequence is that molecular systems may not require energy to maintain complex spatial patterning. For example, the diversity of acceptable spatial arrangements of an engineered organism could be encoded entirely in its equilibrium distribution, and its response to environmental stimuli could be achieved by conditioning of that distribution (see Figure 2 for an imaginative example). Another consequence of organization as an equilibrium phenomenon is self-healing: the system naturally recovers from perturbations by relaxing to equilibrium. A further consequence is that these aspects of molecular behavior can be viewed as a form

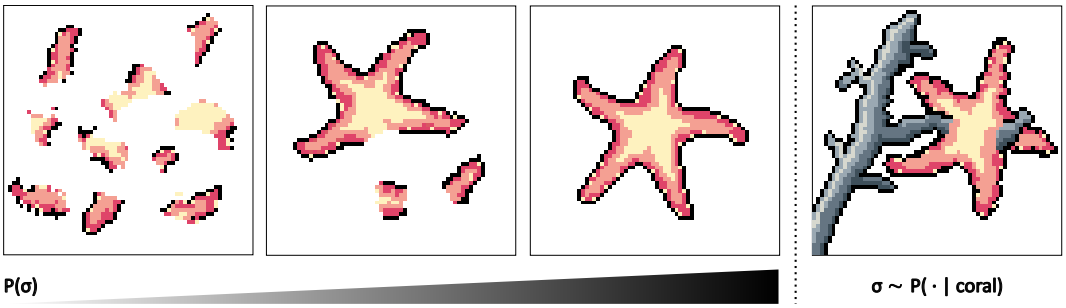

Figure 2 A cartoon of potential outcomes of encoding complex spatial patterning into equilibrium distributions. The Boltzmann distribution encodes the starfish's structure as shown in the first three panels, with probability increasing from left to right. Connected, five-armed configurations of its constituent molecules are favored (right). A natural consequence is self-healing, where equilibrium relaxation alone is sufficient to restore structure after perturbations like damage to its arms (middle), or even complete disassembly (left). The presence of and interactions with a branch of coral condition the starfish's Boltzmann distribution, resulting in a high probability of the starfish everting its stomach to feed on the coral.

of "thought" insofar as their governing mathematics are akin to those of neural networks. Of course, energy consumption could and must still be utilized to condition the spatial distribution into particular states, to drive the system through its various configurations, and to metabolically synthesize the relevant molecules in the first place. That said, maintenance of complex structure at equilibrium opposes the idea that equilibrium generically yields disorder. Surprisingly, the disorder intuition is most often taught in terms of lattice gas models close to those studied here. In our case, the key differences which introduce complex spatial order at equilibrium are heterogeneous interaction energies between molecular species, and large species diversity: the inclusion of many species, where an increase in the number of species increases spatial distribution complexity.

In this paper, we introduce several standard lattice models of pairwise molecular interactions at equilibrium, including the model initially studied in the Boltzmann liquids [7] work, and study their universality or lack thereof for representing spatial distributions. Our models are standard generalizations of basic lattice gas models, which include many species types with heterogeneous pairwise interaction energies. Stochastic modeling captures the noise of Brownian motion, and equilibrium assumptions imply the probability of each state of the system is given analytically by the Boltzmann distribution. The question then is whether interaction energies can be tuned so that a target distribution over spatial molecular arrangements can be achieved. Our work first presents two types of impossibility results, arguing that pairwise interactions alone are not sufficient to encode arbitrary distributions over fine-grained molecular arrangements. We then present a notion of "labeling" molecules – one can think of each of many species types being labeled with one of two fluorescent molecules – where we show that any distribution over spatial label configurations can be approximated by programming interaction energies. This universality result for label distributions stands as initial evidence that Boltzmann liquids – molecules rearranging according to pairwise interactions at equilibrium – are not limited to a simple set of distributions, and as such may be programmed effectively for complex computation. As a caveat that will be discussed later, such distributions may use strong energies and thus not behave as a "liquid" (i.e., may not sample many configurations) on reasonable timescales.

2 Lattice Models of Pairwise Interactions

2.1 Model assumptions

Our model assumptions are as follows. First, we consider space to consist of discrete positions. Second, we consider the molecules to be *monodisperse*, meaning all molecules are the same size, the solvent species is treated as any other species, and each molecule occupies one position. It is worth mentioning that the monodispersity assumption can be relaxed by an extension to the model not discussed in this work, and further that in the limit of finer lattice spacing relative to molecule size, the model approximates continuous space. There are two bifurcations to the model assumptions that we consider: the first is isotropy vs. anisotropy. In the isotropic case, we view molecules as rotationally symmetric, meaning species i and j in neighboring positions will always contribute a symmetric interaction energy $G_{i,j}$. We will also consider the anisotropic case, which will explicitly model molecules as rotatable cubes in a cubic lattice, in which case species i and j in neighboring positions contribute an interaction energy that depends on their orientations. The second bifurcation is whether the system is in thermal equilibrium with a heat bath but is otherwise isolated (the canonical ensemble in statistical thermodynamics) vs. in thermal *and chemical* equilibrium with an infinite reservoir wherein molecules exchange with the reservoir (the grand canonical ensemble in statistical thermodynamics). Thus, the reader can keep in mind four models: canonical isotropic, canonical anisotropic, grand-canonical isotropic, and grand-canonical anisotropic. Lastly, we assume that the energy of the system is captured solely by interaction energies $G_{i,j}$ and chemical potentials G_i, and because the system is in equilibrium, the probability mass is given by the Boltzmann distribution, to be formalized shortly.

2.2 Model formalization

We represent space as a set of positions V. If positions p_1 and p_2 are such that species in those positions induce a pairwise interaction, we say $p_1 \sim p_2$. Equivalently, the space is represented by an undirected graph $G = (V, E)$. Commonly, V is taken as a lattice $V = \mathbb{N}^H \times \mathbb{N}^W \times \mathbb{N}^L$, and edges are restricted to be local, e.g., the *von Neumann* neighbors. Formally, $p_1, p_2 \in \mathbb{Z}^3$ are von Neumann neighbors if $||p_1 - p_2||_1 = 1$.

Given a set of chemical *species* M, we consider spatial *configurations* of species in positions, such that a configuration is formally $\sigma : V \to M$. The solvent is modeled explicitly as one of the species in M. The assumption that the system is in equilibrium implies that the probability mass over configurations is given by the Boltzmann distribution $P(\sigma) \propto \exp(-G(\sigma)/kT)$ for an energy function $G(\sigma)$, where T is the temperature in Kelvin and k is Boltzmann's constant. This energy function is expressed differently in the isotropic vs. anisotropic cases as follows. Figure 3 provides visual examples that may aid the reader in understanding the models and distinguishing the isotropic vs. anisotropic cases.

2.2.1 Isotropic model

Here we use a standard generalization of a lattice gas Ising model, equivalently a Potts model with heterogeneous weights. This model has been used in prior work in the context of liquid-liquid phase separation of more than two species (multicomponent liquids) [12], as well as in the Boltzmann liquids work [7]. Here, the energy of a configuration is given by:

$$G(\sigma) = \frac{1}{2} \sum_{p_1 \in V} \sum_{\substack{p_2 \in V \\ \text{s.t. } p_1 \sim p_2}} G_{\sigma(p_1),\sigma(p_2)} + \sum_{p \in V} G_{\sigma(p)},$$

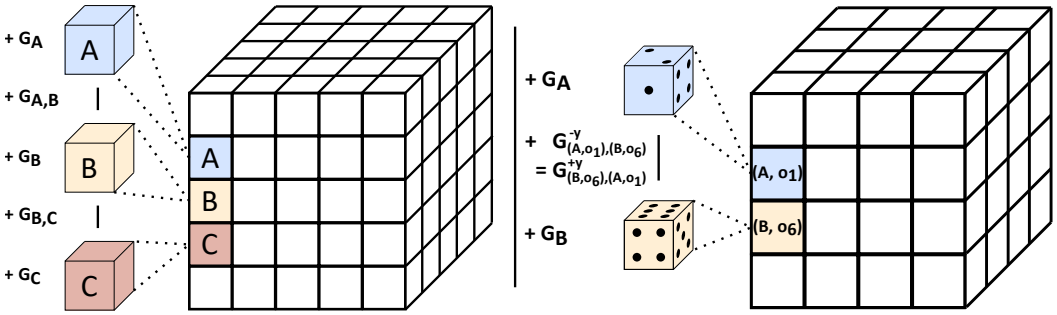

Figure 3 Two example partial configurations. The configurations are partial in the sense that our models do not allow empty positions: solvents, or equivalently "empty" species are explicit species of the set of species M, and each position must be assigned a species in M. Positions are left empty here for visual clarity. On the left, a configuration σ of the isotropic model and the contributions to the configuration energy $G(\sigma)$ provided by the visualized species and their interactions. On the right, a configuration σ of the anisotropic model. Species can be thought of as rotatable cubes (i.e., as dice), with six distinct faces. A position is assigned a species type (here, A, B) and an orientation (here, o_1 and o_6), where orientations correspond to the 24 discrete rotations of the cube, $SO(3,\mathbb{Z})$. Here, $+y = (0,1,0)$ and $-y = (0,-1,0)$ denote the relative direction in the lattice of the two species which are providing interaction energy. In the anisotropic models, constraints on interaction energies such as $G^{-y}_{(A,o_1),(B,o_6)} = G^{+y}_{(B,o_6),(A,o_1)}$ give the model equivalence to a model which defines $G_{A5,B6}$ as the energy provided when species A and B have their die faces 5 and 6 abutting, as pictured, noting that die face 5 is occluded but is opposite face 2.

where $G_{i,j} = G_{j,i}$ is the energy contribution of a molecule of species $i \in M$ interacting with a molecule of species $j \in M$, and where G_i is the internal energy (equivalently, chemical potential) of species $i \in M$. The convention is more negative $G_{i,j}$ and G_i are more favorable.

In our analysis, it will be useful to write this energy in an alternate form which clarifies that key statistics ($n_{i,j}$ and n_i, defined shortly) of a configuration are sufficient to determine its energy. The Kronecker delta function will be useful, which returns one if its subscripts are equal, otherwise zero: $\delta_{i,j} = 1$ if $i = j$, else $\delta_{i,j} = 0$. A configuration σ induces a value $n_{i,j}$ for all $i, j \in M$, which is the number of neighbor interactions between molecules i and j. We define:

$$\hat{n}_{i,j} = \sum_{p_1 \in V} \sum_{\substack{p_2 \in V \\ \text{s.t. } p_1 \sim p_2}} \delta_{\sigma(p_1),i} \delta_{\sigma(p_2),j} \ .$$

As written, $\hat{n}_{i,j}$ double counts the case $i = j$, which we correct as:

$$n_{i,j} = \begin{cases} \frac{1}{2}\hat{n}_{i,j} & i = j \\ \hat{n}_{i,j} & i \neq j \ . \end{cases}$$

Additionally, a configuration σ induces a count of each species in the lattice,

$$n_i = \sum_{p \in V} \delta_{\sigma(p),i} \ .$$

Finally, we can see that lattice configuration σ has the corresponding energy:

$$G(\sigma) = \sum_{i \leq j \in M} n_{i,j} G_{i,j} + \sum_{i \in M} n_i G_i \tag{1}$$

where each independent energy parameter appears exactly once. This formulation makes clear the idea that any two configurations with equal $n_{i,j}$ and n_i must have equal energy and thus equal probability.

If we assume the system's dynamics obey detailed balance with respect to the energy we have defined, the Boltzmann distribution provides an analytical formula for the probability of a given configuration:

$$P(\sigma) = \frac{1}{Z} e^{-G(\sigma)/kT}, \quad \text{where } Z \text{ is the partition function,} \quad Z = \sum_{\sigma \in \Omega} e^{-G(\sigma)/kT}.$$

Here, Ω is the set of valid configurations, which reflect assumptions about physical constraints. If the system is considered to be in thermal equilibrium with a constant-temperature heat bath (*canonical*), then species do not diffuse from or into the heat bath, so the molecular counts are fixed. If we let $\vec{c} \in \mathbb{N}^M$ be the vector representing the count of each species, the *canonical configuration space* is the set $\Omega = \{\sigma : V \to M \mid \forall i \in M, n_i(\sigma) = \vec{c}(i)\}$. If instead the system is in thermal and chemical equilibrium with a reservoir (*grand canonical*), species diffuse to and from the reservoir, so the counts of each species type can change, and thus the set of configurations is $\Omega = \{\sigma \mid \sigma : V \to M\}$.

2.2.2 Anisotropic model

For simplicity, in the anisotropic model we restrict the set of positions to the cubic lattice $V = \mathbb{N}^H \times \mathbb{N}^W \times \mathbb{N}^L$ with von Neumann neighbors. We then effectively consider molecules to be rotatable cubes with asymmetric internal structure (see Figure 3, right). Formally, each species is augmented with an *orientation*: the set of species $M = M_{\text{type}} \times M_{\text{orientation}}$ consists of elements formed by the Cartesian product of a set of *species types* M_{type} and a set of *orientations* $M_{\text{orientation}}$. Since we consider molecules as rotatable cubes, we take $M_{\text{orientation}} = SO(3, \mathbb{Z})$, where $SO(3, \mathbb{Z})$ is the group of orientation-preserving cube rotations. With $i = (a, o_1), j = (b, o_2) \in M$, the interaction energies must be specified with a direction $\hat{u} \in \Delta$, where Δ is a set of *displacements* that must be symmetric in the sense that $-\hat{u} \in \Delta$ if and only if $\hat{u} \in \Delta$. For our von Neumann neighborhood, $\Delta = \{\hat{u} \in \mathbb{Z}^3 \mid ||\hat{u}||_1 = 1\}$. Then $G^{\hat{u}}_{i,j}$ describes the interaction energy between species a and b when they are in orientations o_1 and o_2, respectively, and are displaced by $\hat{u}$ in the lattice. Note that the rotation group $SO(3, \mathbb{Z})$ also acts on displacements; then we require that the interaction energies are invariant to rotations, formally:

$$\forall R \in SO(3, \mathbb{Z}), \quad G^{R\hat{u}}_{(a,Ro_1),(b,Ro_2)} = G^{\hat{u}}_{(a,o_1),(b,o_2)}.$$

In addition, we require interaction symmetry as in the isotropic case: $G^{\hat{u}}_{(a,o_1),(b,o_2)} = G^{-\hat{u}}_{(b,o_2),(a,o_1)}$. Finally, we require that the internal energy is independent of orientation, formally $G_{(a,o)} = G_a$ for all $o \in SO(3, \mathbb{Z})$. In the canonical anisotropic model, the counts of each species type in M_{type} are fixed, but orientations can change.

The energy of a configuration is

$$G(\sigma) = \frac{1}{2} \sum_{p \in V} \sum_{\substack{\hat{u} \in \Delta \\ p + \hat{u} \in V}} G^{\hat{u}}_{\sigma(p),\sigma(p+\hat{u})} + \sum_{p \in V} G_{\sigma(p)}.$$

As in the isotropic case, it is useful to define sufficient statistics for the energy. For each $i, j \in M$ and $\hat{u} \in \Delta$, we define:

$$n^{\hat{u}}_{i,j}(\sigma) = \sum_{\substack{p \in V \\ p + \hat{u} \in V}} \delta_{\sigma(p),i} \delta_{\sigma(p+\hat{u}),j},$$

which counts the number of ordered neighboring pairs of type $i \to j$ in the displacement direction $\hat{u}$. The number of molecules of species type i in the lattice is defined as in the isotropic case: $n_i(\sigma) = \sum_{o \in SO(3,\mathbb{Z})} \sum_{p \in V} \delta_{\sigma(p),(i,o)}$. Noting that $n^{\hat{u}}_{i,j} = n^{-\hat{u}}_{j,i}$, we can write the energy as:

$$G(\sigma) = \frac{1}{2} \sum_{\hat{u} \in \Delta} \sum_{i,j \in M} n^{\hat{u}}_{i,j} G^{\hat{u}}_{i,j} + \sum_{i \in M_{\text{type}}} n_i G_i.$$

2.2.3 A unified expression for energy

To discuss all models simultaneously, it is helpful to write the energy of all models in a standard form. For each model, we define the *parameters* θ and the *sufficient statistics* t. The parameters θ are a vector and are defined as below for each model. Here, $[x_a]\, \forall a \in A$ denotes the vector obtained by fixing an ordering of the set A and listing one entry x_a for each $a \in A$. In the isotropic canonical model,

$$\theta = [G_{i,j}/kT] \quad \forall i,j \in M \text{ with } i \leq j,$$

in the isotropic grand-canonical model,

$$\theta = [G_{i,j}/kT, G_m/kT] \quad \forall i,j \in M, \forall m \in M \text{ with } i \leq j,$$

in the anisotropic canonical model,

$$\theta = [G^{\hat{u}}_{i,j}/kT] \quad \forall i,j \in M \text{ with } i \leq j, \forall \hat{u} \in \Delta,$$

and in the anisotropic grand-canonical model,

$$\theta = [G^{\hat{u}}_{i,j}/kT, G_m/kT] \quad \forall i,j \in M \text{ with } i \leq j, \forall m \in M_{\text{type}}, \forall \hat{u} \in \Delta.$$

The sufficient statistic vector t is a function $t : \Omega \to \mathbb{Z}^{|\theta|}_{\geq 0}$ that takes as input a configuration and outputs its statistics $n_{i,j}$, n_i, and/or $n^{\hat{u}}_{i,j}$. We define t so that it aligns with θ (e.g., if $G_{i,j}$ is element k of θ, then $n_{i,j}$ is element k of $t(\sigma)$) and thus write for any model:

$$G(\sigma)/kT = \theta \cdot t(\sigma), \qquad P(\sigma) = \frac{1}{Z} \exp(-\theta \cdot t(\sigma)).$$

This formulation makes clear the idea that, for fixed parameters θ, the Boltzmann distribution P depends only on the sufficient statistics t, hence their name; this dependence narrows possible target distributions, as discussed in the next section.

3 Impossibility of Microstate Universality

Here, we formalize and prove the impossibility of the models to represent all microstate distributions, even those that satisfy intuitive symmetries. The most straightforward formulation of universality for distributions is given in terms of the relative entropy, or Kullback-Leibler divergence between target distribution Q and approximate distribution P, where

$$D_{\text{KL}}(Q||P) = \sum_{\sigma} Q(\sigma) \log \frac{Q(\sigma)}{P(\sigma)} \ .$$

Then:

▶ **Definition 1.** *A model is* universal for microstate distributions *if for all* $\varepsilon > 0$*,* V*,* M*, and* Q *over* Ω*, there exists* θ *such that* $D_{KL}(Q||P) < \varepsilon$*.*

Of course, this formulation is beyond reach: as written, Q may require different probabilities for configurations that agree on sufficient statistics. For example, in the isotropic models, there is a rotational symmetry in which rotating a configuration yields another configuration with the same sufficient statistics. Similarly, if periodic boundary conditions are used in any model, then there is also a translational symmetry. We thus define the *uniform-within-fibers* constraint on distributions as follows: two configurations with equal sufficient statistics, $t(\sigma_1) = t(\sigma_2)$, must have equal probability. This constraint captures all translational and rotational symmetries of the models, in addition to configurations which share t but are not translations/rotations (see Appendix A.1 for an example). The set of configurations with sufficient statistics equal to t is called a *fiber* over t, formally $F_t = \{\sigma \in \Omega : t(\sigma) = t\}$, so we say a distribution Q over Ω is *fiber uniform* if:

$$\forall t, \forall \sigma_1, \sigma_2 \in F_t, Q(\sigma_1) = Q(\sigma_2).$$

Thus we define:

▶ **Definition 2.** *A model is* universal for fiber-uniform microstate distributions *if for all* $\varepsilon > 0$*,* V*,* M*, and fiber-uniform* Q *over* Ω*, there exists* θ *such that* $D_{KL}(Q||P) < \varepsilon$*.*

First, we provide a simple counting argument that, for nontrivial lattice sizes $|V| = \Theta(|M|^2)$, there are not enough parameters to encode every fiber-uniform target distribution. This counting argument proves that not every fiber-uniform target distribution can be the result of a parameter setting if parameters have bounded precision, and provides intuition that the number of parameters must be allowed to scale independently of the number of target distributions, motivating our problem reformulation in Section 4. Later in this section, we provide concrete, simple fiber-uniform distributions which suffice to prove that approximate microstate universality as defined above is not possible. These counterexamples are constructed by identifying a constraint on distributions that must be satisfied in order for them to be achieved, and thus provide value beyond the counting argument by illustrating what *kind* of microstate distributions are impossible, in addition to removing the assumption of bounded-precision parameters and illustrating impossibility even of approximation.

3.1 Bounded-precision exact microstate universality requires an exponential number of parameters

In the study of *in-silico* probabilistic models, it is standard to assume that parameters have bounded precision (i.e., that each parameter uses a constant number of bits b). If we make the same assumption for our models, with the physical interpretation that interaction energies are not tunable on a true continuum but instead have some underlying discrete characterization, we can apply a standard counting argument to show that the number of parameters required for exact representation of every fiber-uniform distribution should require a number of parameters that scales exponentially in the support of the distribution (the potential configurations).

For any model, define the set of feasible sufficient statistics vectors $\mathcal{T}_\Omega := \{t(\sigma) \mid \sigma \in \Omega\}$. For each probability vector q over $\mathcal{T}_\Omega$, define the following fiber-uniform microstate distribution: $Q_q(\sigma) := \frac{q(t(\sigma))}{|F_{t(\sigma)}|}$. The set of Q_q defined this way is the set of possible fiber-uniform target distributions. In order to argue that there are more targets than our parameter choices allow, it suffices to consider a subset of the total set of targets whose size is easier to count, so we define:

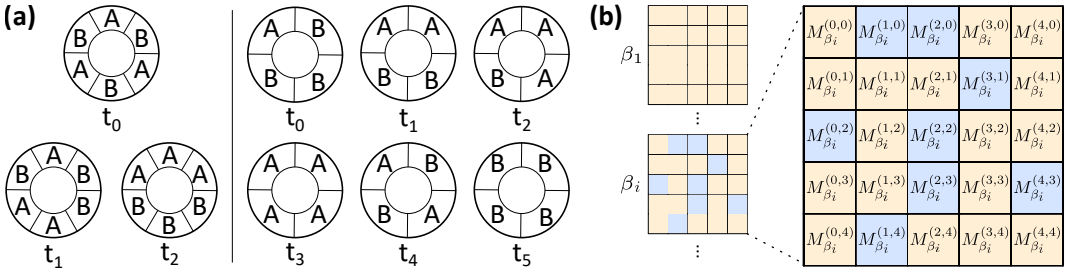

Figure 4 **(a)** On the left, one representative configuration from each of the fibers t_0, t_1, and t_2 in a canonical isotropic model with V a ring of six positions as shown. On the right, one representative configuration from each of the six fibers in a grand canonical isotropic model with V a ring of four positions as shown. **(b)** An example of Theorem 5's construction, in two dimensions for visual clarity. There are two labels, yellow and blue. $\beta_1 \dots \beta_i \dots$ denote the elected label configurations for each Euclidean equivalence class of label configurations. The intended configuration for β_i is shown in the pop-out. Because each species $M_{\beta_i}^{(x,y)}$ is adjacent to its proper neighbors, no energy contribution Λ is made, and the energy of the microstate configuration is given by the energy $-kT \ln Q_K(\beta_i)$ of the leader species $M_{\beta_i}^{(0,0)}$.

$$\mathcal{D}_{\mathcal{T}} = \left\{ q : \mathcal{T}_\Omega \to [0,1] \;\middle|\; q(t) = \frac{z_t}{|\mathcal{T}_\Omega|}, z_t \in \mathbb{Z}_{\geq 0}, \sum_{t \in \mathcal{T}_\Omega} z_t = |\mathcal{T}_\Omega| \right\}.$$

Using standard combinatorics (see Appendix A.2), we can count the size of $\mathcal{D}_{\mathcal{T}}$:

$$|\mathcal{D}_{\mathcal{T}}| = \binom{2|\mathcal{T}_\Omega| - 1}{|\mathcal{T}_\Omega| - 1} \geq 2^{|\mathcal{T}_\Omega| - 1}.$$

Now, let $n = |\theta|$, the number of parameters of the model. If each parameter is allowed b bits, the set of all possible parameter choices of the model has size 2^{bn}. For every fiber-uniform distribution to be achieved exactly by some parameter choice would require $2^{bn} \geq 2^{|\mathcal{T}_\Omega| - 1}$, therefore $n \geq \frac{|\mathcal{T}_\Omega| - 1}{b}$. In Appendix A.3, we prove in each model that with $|V| = \Theta(|M|^2)$ (the lattice size scaling quadratically in the number of species) that $|\mathcal{T}_\Omega| \geq 2^{\Omega(n)}$, where Ω on the right-hand side of the inequality refers to asymptotic scaling. Therefore, for fixed b and scaling n, the inequality $n \geq \frac{|\mathcal{T}_\Omega| - 1}{b}$ cannot be satisfied.

Next, we provide explicit counterexample distributions that allow proofs against even approximate microstate universality without the assumption of bounded precision.

3.2 Unbounded-precision approximate microstate universality is prohibited by explicit construction of unachievable distributions

Here we show a constraint that distributions must satisfy beyond being fiber uniform, which allows construction of distributions which cannot be approximated, even with unbounded-precision parameters. The *affine log-probabilities* constraint is that the log-probabilities of the Boltzmann distribution are affine ($f(x) = Ax + b$) in form:

$$\log P(\sigma) = -\theta \cdot t(\sigma) - \log Z,$$

since $\log Z$ is constant with respect to σ. Next, we use this constraint to construct microstate distributions in each model variant that are uniform within fibers but are not approximable, inhibiting universality.

Consider in the canonical isotropic model a ring of six positions with left-right edges but no top-bottom edges with three copies of species A and three copies of species B (Figure 4(a), left). There are three *feasible* fibers, meaning values of $t = (n_{A,A}, n_{A,B}, n_{B,B})$ for which $|F_t| > 0$, i.e., sufficient statistics that can arise from valid microstate configurations. These three feasible fibers are illustrated in Figure 4(a), left. Consider the following target distribution Q, described in terms of an assignment of probability to each fiber's microstates:

If $t(\sigma) = t_0 = (0, 6, 0)$, set $Q(\sigma) = 1/32$.
If $t(\sigma) = t_1 = (1, 4, 1)$, set $Q(\sigma) = 1/16$.
If $t(\sigma) = t_2 = (2, 2, 2)$, set $Q(\sigma) = 1/32$.

By this assignment method, this distribution satisfies the uniform-within-fibers constraint. Additionally, the probabilities sum to one: since $|F_{t_0}| = 2, |F_{t_1}| = 12$, and $|F_{t_2}| = 6$, we have $2/32 + 12/16 + 6/32 = 1$. But does the distribution satisfy the affine log-probabilities constraint? We show that the answer is no, by showing that it violates the *midpoint relation* of affine functions:

$$f\left(\frac{x_1 + x_2}{2}\right) = \frac{f(x_1) + f(x_2)}{2}.$$

To show this, first note that we can uniquely define the *fiber microstate probability* $u(t) := Q(\sigma)$ such that $\sigma \in F_t$, since Q is uniform within fibers. Note that $t_1 = \frac{t_0+t_2}{2}$. Towards contradiction, assume Q satisfies the affine log-probabilities constraint, so that it must satisfy the affine midpoint relation:

$$\log u(t_1) = \log u\left(\frac{t_0 + t_2}{2}\right) = \frac{\log u(t_0) + \log u(t_2)}{2},$$

and by exponentiating, we have

$$u(t_1)^2 = u(t_0)u(t_2),$$

but $(1/16)^2 \neq (1/32)^2$, thus a contradiction. Therefore Q, which satisfies the fiber-uniform constraint but does not satisfy the affine log-probabilities constraint, cannot have parameters $\theta = (G_{A,A}, G_{A,B}, G_{B,B})$ that produce it. In Appendix A.4, we provide similar examples in each model variant. Although these counterexamples use lattices with restricted edges for simplicity of presentation, we expect that counterexamples can be constructed on nontrivial lattices in general, e.g., on the typically-considered cubic lattices with von Neumann neighbors.

Next, we utilize the counterexamples to formally argue against universality as per Definition 2:

▶ **Theorem 3.** *No model variant is universal for fiber-uniform microstate distributions.*

Proof. Towards contradiction, assume at least one model variant is universal for fiber-uniform microstate distributions. Thus, for any fiber-uniform Q, there is an infinite sequence of parameter choices $\theta_1, \theta_2, \ldots$ such that the KL divergence between Q and P_{θ_n} gets arbitrarily close to zero as n goes to infinity: $\lim_{n\to\infty} D_{\mathrm{KL}}(Q||P_{\theta_n}) = 0$. Since P_{θ_n} and Q are discrete distributions, we have $\lim_{n\to\infty} P_{\theta_n}(\sigma) = Q(\sigma)$, i.e., the probability assigned to σ by P_{θ_n} gets arbitrarily close to the probability assigned by Q (this can be proven using Pinsker's inequality [36]: $\frac{1}{2}\sum_\sigma |Q(\sigma) - P_{\theta_n}(\sigma)| \leq \sqrt{\frac{1}{2} D_{\mathrm{KL}}(Q||P_{\theta_n})}$). Let the lattice V, species M, and Q be defined as in the counterexample for the assumed-universal model variant. We write $u_n(t) := P_{\theta_n}(\sigma)$ (and as above, $u(t) := Q(\sigma)$) such that $\sigma \in F_t$. As with $u(t)$,

$u_n(t)$ is well defined, since each P_{θ_n} must be fiber uniform. Since $\lim_{n\to\infty} P_{\theta_n}(\sigma) = Q(\sigma)$, we have $\lim_{n\to\infty} u_n(t) = u(t)$. Since P_{θ_n} is the Boltzmann distribution of one of our models, it must satisfy the log-affine constraint, and as above we have $\lim_{n\to\infty} u_n(t_1)^2 = \lim_{n\to\infty} u_n(t_0)u_n(t_2)$, so $u(t_1)^2 = u(t_0)u(t_2)$ which, as above, is a contradiction by the counterexample's construction. Therefore no model variant is universal for fiber-uniform microstate distributions. ◀

Aside from the parameter scaling issue raised by the counting argument, these counterexamples violating the affine log-probabilities constraint provide an initial investigation into the structure of exactly which microstate distributions can and cannot be represented. In the following section, we provide a macrostate picture in which we achieve universal distribution approximation by circumventing both flavors of impossibility results.

4 A Natural Observable Enables Parameter Scaling and Macrostate Universality

As shown by the counting argument in the previous section, universal distribution approximation requires a way of scaling parameters independently of the domain of the target distribution. Here we describe *macrostates*, which are partitions of the microstates into collections which have the same *observable outcomes*, introducing the possibility for the number of molecular species and thus parameters to be scaled independently of the number of possible observable macrostates.

Formally, we follow [7] in defining macrostates in general. Given a set of configurations Ω, an *observable* is *any* function $a : \Omega \to A$. We call $\alpha \in A$ an observable outcome. For example, an observable could be $a(\sigma) = n_i(\sigma)$, the count of species i in configuration σ (experimentally, species i could be fluorescently labeled, and thus our observable would be proportional to fluorescence measurement). Given an observable a and an outcome $\alpha \in A$, we define the set of microstates *consistent* with the outcome α as $M_a(\alpha) = \{\sigma \in \Omega \mid a(\sigma) = \alpha\}$. Given parameters θ and thus a Boltzmann distribution P over configurations, the probability of seeing a particular observable outcome α is the marginal probability $P_a(\alpha) = \sum_{\sigma \in M_a(\alpha)} P(\sigma)$, i.e., the sum of corresponding microstates.

Our choice of observable is as follows. We *label* each species with a label from a set of labels K, formally by a labeling function $\mathcal{K} : M \to K$ that assigns each species $m \in M$ to a label $k \in K$. In the anisotropic case, labels are meant to ignore orientation, so we enforce $\mathcal{K}((a, o)) = \mathcal{K}((a, o'))$ for all $a \in M_{\text{type}}$ and $o, o' \in M_{\text{orientation}}$. This labeling observable captures phenomena such as fluorescently-labeled molecules or molecules with functional domains independent of their interacting domains. A *label configuration* is a function $\beta : V \to K$ that maps lattice positions to labels. Given a set of species M, a labeling function $\mathcal{K}$, and a configuration σ, we define the label configuration of σ as $\beta_\sigma = \mathcal{K} \circ \sigma$. Given parameters θ, we can write the probability of a label configuration β as:

$$P_{\mathcal{K}}(\beta) = \sum_{\sigma \in \Omega \mid \mathcal{K} \circ \sigma = \beta} P(\sigma),$$

the sum of microstate configurations that map to β by the labeling function $\mathcal{K}$.

In this section, we focus on cubic lattices $V = \{0, \ldots, L-1\}^3$ with periodic von Neumann interaction neighborhoods. The purpose of equal edge lengths of the lattice is to simplify the explanation of the requirement that target label distributions must be rotationally invariant, as will be described. The purpose of considering periodicity is to highlight the construction's

requirement that target label distributions are translation invariant, as will be described; we expect the result holds with non-periodic boundary conditions, as well as on other graphs (positions/neighbors).

We define a distribution Q over label configurations as *translationally invariant* as follows. Let a lattice translation $\tau_d : V \to V$ be defined as $\tau_d(p) = p + d \mod L$ for $d \in \mathbb{Z}^3$. This induces a label translation $\tau_d\beta : V \to K$ defined as $\tau_d\beta(p) = \beta(\tau_d^{-1}(p))$. A distribution Q over label configurations is translationally invariant if $Q(\beta) = Q(\tau_d\beta)$ for all $d \in \mathbb{Z}^3$. Any distribution induced by a setting of parameters for our models on a periodic lattice must satisfy translational invariance because translations do not change energy nor probability: $G(\sigma) = G(\tau_d\sigma)$ and thus $P(\sigma) = P(\tau_d\sigma)$, where $\tau_d\sigma(p) = \sigma(\tau_d^{-1}(p))$.

Further, we define a distribution Q over label configurations as *rotationally invariant* as follows. Consider a lattice rotation $R \in SO(3, \mathbb{Z})$. We define the rotation of a label configuration as $R\beta(p) = \beta(R^{-1}p \mod L)$. Then a distribution Q over label configurations is rotationally invariant if for all $R \in SO(3, \mathbb{Z})$, $Q(\beta) = Q(R\beta)$. This constraint must be satisfied by distributions induced by our model because rotations do not change energy nor probability: in the isotropic model, $G(\sigma) = G(R\sigma)$ and thus $P(\sigma) = P(R\sigma)$, where $R\sigma(p) = \sigma(R^{-1}p \mod L)$. In the anisotropic model, a lattice rotation also rotates the internal orientation of each molecule: given an anisotropic species (i, o), define $\hat{R}(i, o) = (i, Ro)$ and allow it to act elementwise on configurations σ, so that $R\sigma(p) = \hat{R}(\sigma(R^{-1}p \mod L))$, and equivalently if $\sigma(R^{-1}p \mod L) = (i, o)$ then $R\sigma(p) = (i, Ro)$. Since we constrain the anisotropic interaction energies to be rotationally invariant, rotations do not change energy: $G(\sigma) = G(R\sigma)$ and $P(\sigma) = P(R\sigma)$.

We call a label distribution *Euclidean invariant* if it is both translationally and rotationally invariant. Thus we define universality for label distributions as follows:

▶ **Definition 4.** *A model is* universal for Euclidean-invariant label distributions *if for all $\varepsilon > 0$, $V = \{0, \ldots, L-1\}^3$, von Neumann neighbors E, a set of labels K, and Euclidean-invariant label distributions Q_K, there exists a set of species M, parameters θ, and labeling function $\mathcal{K}$ such that $D_{KL}(Q_K || P_{\mathcal{K}}) < \varepsilon$, where $P_{\mathcal{K}}$ is the label distribution induced by M, θ, V, E, and $\mathcal{K}$.*

Note that by this definition, the set of species M can scale independently of V and K, so that the number of parameters $|\theta|$ can scale independently of the domain of the target distribution Q_K, circumventing the impossibility via counting argument of Section 3.1. Next, we show by construction that both the anisotropic and isotropic grand-canonical models are universal for label distributions. Our particular definitions prohibit canonical models from universality: since a target distribution fixes the lattice size $|V|$ and canonical configuration spaces consider a set of at most $|V|$ distinct species, scaling the number of labels K increases the number of label configurations beyond the fixed number of parameters. In the discussion, we illustrate an alternative observable that may allow universality in the canonical models.

4.1 Universality in the anisotropic grand-canonical model

Here we provide a construction that shows that the anisotropic grand-canonical model is universal for Euclidean-invariant label distributions. The construction is a "brute-force", "hard-coded" construction, meaning that each label configuration's probability is effectively encoded into a single parameter. As a consequence, the number of parameters (and thus unique species) scales with the number of label configurations with nonzero probability in the target distribution. This may be considered analogous to Boltzmann machine universality proofs where the probability of each visible node configuration is effectively encoded in a

single hidden neuron [16]. The role of "lookup-table" constructions such as this one and the potential for next questions related to structured distributions that do not require so many parameters is elaborated on in the Discussion.

▶ **Theorem 5.** *The anisotropic grand-canonical model is universal for Euclidean-invariant label distributions.*

Proof. Figure 4(b) illustrates the construction in two dimensions, in case it is helpful to the reader. Given a target label distribution Q_K over the set of label configurations $\{\beta_1, \dots\}$, our goal is to construct a set of species M and a labeling function $\mathcal{K}$, and to set the parameters θ such that $P_{\mathcal{K}}$ approximates Q_K. We define the *Euclidean equivalence class* on label configurations as follows. We say $\beta_i \hat{=} \beta_j$ if there exist a lattice rotation $R \in SO(3,\mathbb{Z})$ and lattice translation τ_d for $d \in \mathbb{Z}^3$ such that $\tau_d(R\beta_i) = \beta_j$. Then the Euclidean equivalence class of β_i, denoted $\hat{\beta}_i$, is the equivalence class given the $\hat{=}$ relation, formally $\hat{\beta}_i = \{\beta_j \mid \beta_i \hat{=} \beta_j\}$. This partitions the set of label configurations into a set of equivalence classes. For each equivalence class, we choose one *elected configuration* and denote it $\hat{\beta}_i^*$.

We now construct the set of species M. We first define the *species for a label configuration* β_i. For each elected configuration $\hat{\beta}_i^*$ such that $Q_K(\hat{\beta}_i^*) > 0$, for each position in the lattice $(x,y,z) \in V$, we include a species type $M_{\beta_i}^{(x,y,z)}$. We define the *intended position* $\mathrm{pos}\left(M_{\beta_i}^{(x,y,z)}\right) = (x,y,z)$ and the *intended label configuration* $\mathrm{cfg}\left(M_{\beta_i}^{(x,y,z)}\right) = \hat{\beta}_i^*$. The labeling function $\mathcal{K}$ is defined so that a species type has a label matching its intended position in its intended label configuration, formally $\mathcal{K}\left(M_{\beta_i}^{(x,y,z)}\right) = \hat{\beta}_i^*((x,y,z))$.

We set the interaction energies $G_{i,j}^{\hat{u}}$ as follows. First, we pick a large positive energy Λ, to be quantified later with respect to the target ε. We will consider species of some *default orientation* $R_0 \in SO(3,\mathbb{Z})$. Our interactions will be local von Neumann interactions, so that $G_{i,j}^{\hat{u}} = 0$ for all $\hat{u} \notin \{\hat{u} \mid ||\hat{u}||_1 = 1\}$. For all $i = (M_{\beta_a}^{p_2}, R_0)$ and $j = (M_{\beta_b}^{p_1}, R_0)$, we set $G_{i,j}^{\hat{u}} = \Lambda$, unless the species types correspond to the same label configuration, formally $a = b$, and their intended positions are adjacent and $\hat{u}$ apart, formally $p_1 - p_2 = \hat{u}$, in which case we set $G_{i,j}^{\hat{u}} = 0$. We call a configuration σ *intended* if all species are next to their six intended von Neumann neighbors and all agree on the intended label configuration, formally $\forall u, v \in V$, $\mathrm{cfg}(\sigma(u)) = \mathrm{cfg}(\sigma(v))$, and there exists a lattice translation and rotation $\tau_d R$ such that $\forall u \in V$, $\mathrm{pos}(\sigma(u)) = \tau_d R(u)$; otherwise we call configurations *unintended*. By construction, the contribution of the parameters $G_{i,j}^{\hat{u}}$ to the energy of an intended configuration is zero and to unintended configurations is at least Λ.

We set the per-molecule energies G_i as follows. For all species types m, if $\mathrm{pos}(m) = (0,0,0)$, we set $G_m = -kT \ln Q_K(\mathrm{cfg}(m))$, i.e., we designate the species type $M_{\beta_i}^{(0,0,0)}$ as the *leader of label configuration* β_i, and set its per-molecule energy to $-kT \ln Q_K(\hat{\beta}_i^*)$, with the intention that $M_{\beta_i}^{(0,0,0)}$ is the only energy contribution to its intended configuration. For all other species types m such that $\mathrm{pos}(m) \neq (0,0,0)$, we set $G_m = 0$.

With the set of molecule types M and parameters thus defined, we now aim to argue correctness. First note that for each elected configuration $\hat{\beta}_i^*$ with $Q_K(\hat{\beta}_i^*) > 0$, there is an intended configuration σ_i consisting of the species types $M_{\beta_i}^{(x,y,z)}$ each in their intended positions. Because it is an intended configuration, this configuration has zero energy contribution from $G_{i,j}^{\hat{u}}$ parameters, and has one G_i contribution from $G_{M_{\beta_i}^{(0,0,0)}}$, so $G(\sigma_i) = -kT \ln Q_K(\hat{\beta}_i^*)$, and so $P(\sigma_i) = \frac{1}{Z}\exp(-G(\sigma_i)/kT) = \frac{1}{Z}Q_K(\hat{\beta}_i^*)$. For any lattice translation τ_d and lattice rotation R, the label configuration $\beta_j = (\tau_d \circ R)\hat{\beta}_i^*$ matches labels with the microstate configuration $\sigma_j = (\tau_d \circ R)\sigma_i$, and since rotations/translations do not affect energy, $G(\sigma_j) = G(\sigma_i)$, and so $P(\sigma_j) = P(\sigma_i)$.

Finally, let $\mathcal{I}$ be the set of intended configurations and $\mathcal{U}$ be the set of unintended configurations. By construction, for every label configuration β there is exactly one intended configuration σ_β with matching labels, namely the translation/rotation of the elected configuration for the Euclidean equivalence class of β. Further, $\exp(-G(\sigma_\beta)/kT) = Q_K(\beta)$. The total Boltzmann weight of all intended configurations is $\sum_{\sigma\in\mathcal{I}} \exp(-G(\sigma)/kT) = \sum_\beta Q_K(\beta) = 1$. By construction, every unintended configuration has at least one $G^{\hat{u}}_{i,j}$ contributing Λ to its energy. Note that each G_m is nonnegative since $0 < Q_K(\beta) \leq 1$ so $\ln Q_K(\beta) \leq 0$. Thus each unintended configuration $\sigma \in \mathcal{U}$ has $G(\sigma) \geq \Lambda$. The goal is to set Λ so that unintended configurations contribute negligibly to the partition function Z. Since the set of all configurations Ω is finite and $|\Omega| = |M|^{|V|}$, we have $\sum_{\sigma\in\mathcal{U}} \exp(-G(\sigma)/kT) \leq |\Omega| \exp(-\Lambda/kT)$. Therefore we can bound the partition function:

$$1 \leq Z = \sum_{\sigma\in\mathcal{I}} \exp(-G(\sigma)/kT) + \sum_{\sigma\in\mathcal{U}} \exp(-G(\sigma)/kT) \leq 1 + |\Omega| \exp(-\Lambda/kT).$$

We choose Λ such that $|\Omega| \exp(-\Lambda/kT) < \exp(\varepsilon) - 1$ so that $Z < \exp(\varepsilon)$. For any label configuration β, since there is one intended configuration σ_β with $\mathcal{K} \circ \sigma_\beta = \beta$, we have

$$P_{\mathcal{K}}(\beta) = \sum_{\sigma | \mathcal{K}\circ\sigma=\beta} P(\sigma) \geq P(\sigma_\beta) = \frac{Q_K(\beta)}{Z},$$

thus,

$$D_{\mathrm{KL}}(Q_K || P_{\mathcal{K}}) = \sum_\beta Q_K(\beta) \log \frac{Q_K(\beta)}{P_{\mathcal{K}}(\beta)} \leq \sum_\beta Q_K(\beta) \log Z = \log Z < \varepsilon. \quad \blacktriangleleft$$

4.2 Universality in the isotropic grand-canonical model with symmetry-breaking boundary conditions

With isotropic molecules, a construction conceptually similar to the anisotropic construction also yields universality for label distributions. However, the isotropic model needs a symmetry-breaking boundary condition to effectively orient the molecules, since anisotropy cannot be used. The global orientation is achieved by propagating orientation from the boundary inward. This boundary condition removes the requirement that the target distribution must be Euclidean invariant. It is worth noting that, as in the picture of environmental clamping in the Boltzmann liquids framework discussed in the Introduction and Figure 1, the boundary condition inducing orientation can be seen as a type of conditioning of the probability distribution. We first formalize the boundary condition. Given a cubic lattice $V^{+2} = \{-1, \ldots, L\}^3$, we define the *boundary* $\partial V^{+2} = \{(x, y, z) \in V^{+2} \mid x \in \{-1, L\}$ or $y \in \{-1, L\}$ or $z \in \{-1, L\}\}$. The *boundary condition* is that for each position in the boundary $p_\partial \in \partial V^{+2}$, we include a species M^{p_∂} that is fixed in that position.

▶ **Theorem 6.** *The isotropic grand-canonical model with a boundary condition is universal for label distributions.*

The proof can be found in Appendix A.5. The construction and proof are similar to those of Theorem 5, with an additional inductive argument that the boundary condition enforces a global orientation. We expect that the boundary condition can be relaxed to, e.g., a constant-sized corner of the cubic lattice. The boundary condition here is a special case of clamping as discussed in the Boltzmann liquids framework [7]. This result suggests that symmetry-breaking by a fixed boundary condition, which could be introduced by the

environment, can play a role in the complexity of spatial distributions of isotropic molecules internal to the boundary. We note that this is counterintuitive: small examples such as Figure 1(a) give the impression that conditioning a distribution reduces complexity by reducing the state space. However, this result suggests that conditioning reduces the state space in a way that is useful or expressive, in this case by breaking symmetries and likely increasing the class of representable distributions.

5 Discussion

The label-distribution formalism allows the system's number of parameters to scale independently of the support of the target distribution by including more underlying species per label. However, our universality construction uses a number of species exponential in the lattice volume and number of labels: $|M| = \Theta(|K|^{|V|})$. This is not unexpected: counting arguments like the one in Section 3.1 suggest that the exponential number of parameters is required. When considering all possible distributions, most look algorithmically random [17], in the sense that most have high Kolmogorov complexity: the shortest programs on a universal Turing machine which output the distributions must effectively hard-code them. In other words, most distributions look like uncorrelated noise. Future work, especially work including the use of learning algorithms such as the contrastive learning used in the Boltzmann liquids framework [7], will help identify the types of low-Kolmogorov complexity distributions that can be achieved with a reasonable number of parameters. In life, coding theory arguments suggest that genotype-phenotype maps may naturally favor low-Kolmogorov complexity phenotypes [13], aligning with our expectation that lifelike dynamic spatial distributions may be learnable. Exponential-parameter universality results such as ours are not prescriptions for suggested constructions, but instead suggest that there are no formal barriers to representing arbitrary macrostate distributions and thus invite the study of new universal macrostate pictures as well as the investigation into which spatial distributions are practical for equilibrium systems of programmable pairwise local interactions.

Our universality result for label distributions is conceptually a "hard-coded lookup table": a disjoint subset of the parameters directly encode each target configuration of the distribution. While lookup-table results are effective for proving the absence of formal restrictions against arbitrary distributions, neural-network literature includes stronger approximation results for structured classes which argue that approximation error can reduce at a controlled rate as more parameters are utilized (e.g., see Section 6 of [27] for a survey of smooth approximation results for multi-layer perceptrons). What quantitative structural properties of spatial distributions might define a "distribution complexity" such that, given a target approximation error, the required number of parameters scales smoothly with that complexity? How might physical constraints such as locality guide the identification of such properties?

The hard-coding of our universality construction calls to mind uniquely-addressed self-assembly, where a two-dimensional square lattice or three-dimensional cubic lattice consisting of unique species for each position is constructed, and shapes are assembled by leaving out subsets of the components [41, 14, 24]. Beyond unique addressing, although kinetic in nature, algorithmic self-assembly [10] has articulated how the number of unique species types required relates to the complexity of a target shape, with theoretical constructions achieving optimal results [34]. Another studied task beyond assembly of shapes is to design a set of self-assembling tiles which assembles a pattern of labels, where species types are labeled as in our labeling observable. The Pattern self-assembly Tile-set Synthesis (PATS) problem [19] aims to synthesize the smallest set of species that assembles a target label pattern. Our

universality construction can be loosely viewed as an extension of this problem to the self-assembly of *distributions* of label patterns, as opposed to deterministic self-assembly of a single label pattern. Future work could establish this connection rigorously and provide synthesis methods for the distribution variant of the PATS problem, or inversely, algorithmic self-assembly theory may inform how spatial distribution complexity relates to the number of parameters required in our lattice models.

The utilization of large positive (repulsive) interactions Λ in our universality results suggests that with any reasonably-defined dynamics, the system will have large energy barriers between its intended configurations. As such, this construction is much more like crystalline self-assembly than liquid-like molecular rearrangement, where we define liquid-like as able to explore its configuration space on short timescales. We posit, along with [7], that restricting to parameters that enable liquid-like behavior is an important desideratum to implement effective neural networks of this kind, as short-timescale reconfiguration corresponds to how quickly the molecules will react to an environmental stimulus, and further, systems with slow rearrangement dynamics do not explore their whole Boltzmann distribution on reasonable timescales and so will be sensitive to initial condition and require another kind of analysis. Thus, a motivated task is to identify a variant of universality (appropriate choice of observable, or limitation on target distributions) that can be achieved while maintaining liquid-like interaction energies (i.e., energies $\sim kT$).

Prior work [22, 35, 26] has identified *equilibrium dimerization networks* as a fruitful direction for programming well-mixed, non-spatial molecular information processing. These networks consist of dimerization reactions like $X_i + X_j \leftrightharpoons D_{i,j}$, where the equilibrium is determined by energies of each X_i and $D_{i,j}$ species, effectively tuned by the pairwise interaction energy between univalent X_i and X_j molecules. Such networks can be expressed in our anisotropic model, since univalence can be expressed by setting $G^{\hat{u}}_{i,j}$ appropriately. Further, our model can also be extended to include reactions: when molecules of species A and B are neighbors, they may transition to species C and D. Reactions of this form can induce complex reachability constraints in canonical models, potentially enhancing their distribution representation capability [28, 29]. That our anisotropic lattice model captures well-mixed reaction networks situates it as a prototype or potential candidate for a unified model of molecular information processing.

Although our labeling observable and universality definitions allow positive constructions exclusively for the grand-canonical models, we expect that the canonical models can be universal in distributions defined by other observables. For example, an enticing observable is to consider a subset of lattice positions, e.g. a two-dimensional surface on a three-dimensional lattice, as *visible* observed positions P, and define the rest of the positions as *hidden* positions H. Then, scaling $|H|$ while keeping $|P|$ fixed allows an increased number of species and thus parameters. We anticipate such a formulation may yield universality over microstate distributions of the visible positions in canonical models. In addition, this formulation allows for conditioning the distribution by physically clamping regions of the visible surface P, in contrast to our label observable that lacks a clear clamping mechanism.

Our explicit constructions of impossible distributions shown in Section 3.2 provide an initial exploration into the class of microstate distributions that can be represented by our models. These counterexamples can be seen as "the XOR of Boltzmann liquids", in analogy to the classical example of XOR as a non-linearly-separable dataset that illuminates the limits of single-layer perceptrons [20]. Future work may extend these counterexamples into a complete characterization of achievable microstate distributions.

Our universality result encourages pushing the limits of complexity in programmed equilibrium spatial distributions. Future work aligned with the Boltzmann liquids' notion of clamping [7] will illustrate how these complex molecular spatial arrangements may respond dynamically to environmental cues as the cues condition their Boltzmann distribution. Ultimately, this work provides theoretical evidence that programming equilibrium distributions of pairwise-interacting molecules may be sufficient to yield a complexity of spatial distributions matching or surpassing those demonstrated by life.

References

1 David H. Ackley, Geoffrey E. Hinton, and Terrence J. Sejnowski. A learning algorithm for Boltzmann machines. *Cognitive Science*, 9(1):147–169, 1985. doi:10.1207/S15516709COG0901_7.

2 Siddharth Agarwal, Dino Osmanovic, Mahdi Dizani, Melissa A Klocke, and Elisa Franco. Dynamic control of DNA condensation. *Nature Communications*, 15(1):1915, 2024.

3 Paul C. Bethke and Russell L. Jones. Vacuoles and prevacuolar compartments. *Current Opinion in Plant Biology*, 3(6):469–475, 2000.

4 Clifford P. Brangwynne, Christian R. Eckmann, David S. Courson, Agnieszka Rybarska, Cathrin Hoege, Javad Gharakhani, Frank Jülicher, and Anthony A. Hyman. Germline P granules are liquid droplets that localize by controlled dissolution/condensation. *Science*, 324(5935):1729–1732, 2009.

5 Daniele Cappelletti, Andrés Ortiz-Muñoz, David F Anderson, and Erik Winfree. Stochastic chemical reaction networks for robustly approximating arbitrary probability distributions. *Theoretical Computer Science*, 801:64–95, 2020. doi:10.1016/J.TCS.2019.08.013.

6 Aria S Chaderjian, Sam Wilken, and Omar A Saleh. Diverse, distinct, and densely packed DNA nanostar droplets. *Proceedings of the National Academy of Sciences*, 123(7):e2523462123, 2026.

7 Cameron Chalk, Salvador Buse, Krishna Shrinivas, Arvind Murugan, and Erik Winfree. Learning and inference in a lattice model of multicomponent condensates. In *30th International Conference on DNA Computing and Molecular Programming (DNA 30)*, pages 5–1. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.DNA.30.5.

8 Priyamvada Chugh, Andrew G. Clark, Matthew B. Smith, Davide A. D. Cassani, Kai Dierkes, Anan Ragab, Philippe P. Roux, Guillaume Charras, Guillaume Salbreux, and Ewa K. Paluch. Actin cortex architecture regulates cell surface tension. *Nature Cell Biology*, 19(6):689–697, 2017.

9 George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of Control, Signals and Systems*, 2(4):303–314, 1989. doi:10.1007/BF02551274.

10 David Doty. Theory of algorithmic self-assembly. *Communications of the ACM*, 55(12):78–88, 2012. doi:10.1145/2380656.2380675.

11 Daniel A Fletcher and R Dyche Mullins. Cell mechanics and the cytoskeleton. *Nature*, 463(7280):485–492, 2010.

12 William M Jacobs and Daan Frenkel. Phase transitions in biological systems with many components. *Biophysical Journal*, 112(4):683–691, 2017.

13 Iain G Johnston, Kamaludin Dingle, Sam F Greenbury, Chico Q Camargo, Jonathan PK Doye, Sebastian E Ahnert, and Ard A Louis. Symmetry and simplicity spontaneously emerge from the algorithmic nature of evolution. *Proceedings of the National Academy of Sciences*, 119(11):e2113883119, 2022.

14 Yonggang Ke, Luvena L Ong, William M Shih, and Peng Yin. Three-dimensional structures self-assembled from DNA bricks. *Science*, 338(6111):1177–1183, 2012.

15 Anna Kicheva and James Briscoe. Control of tissue development by morphogens. *Annual Review of Cell and Developmental Biology*, 39(1):91–121, 2023.

16 Nicolas Le Roux and Yoshua Bengio. Representational power of restricted Boltzmann machines and deep belief networks. *Neural Computation*, 20(6):1631–1649, 2008. doi:10.1162/NECO.2008.04-07-510.

17 Ming Li and Paul Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Texts in Computer Science. Springer, Cham, 4 edition, 2019.

18 Andrew S. Lyon, William B. Peeples, and Michael K. Rosen. A framework for understanding the functions of biomolecular condensates across scales. *Nature Reviews Molecular Cell Biology*, 22(3):215–235, 2021.

19 Xiaojun Ma and Fabrizio Lombardi. Synthesis of tile sets for DNA self-assembly. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 27(5):963–967, 2008. doi:10.1109/TCAD.2008.917973.

20 Marvin Minsky and Seymour A. Papert. *Perceptrons: An Introduction to Computational Geometry*. The MIT Press, 2017.

21 Sahil Nagpal, Karthikeyan Swaminathan, Daniel Beaudet, Maud Verdier, Samuel Wang, Christopher L. Berger, Florian Berger, and Adam G. Hendricks. Optogenetic control of kinesin-1, -2, -3 and dynein reveals their specific roles in vesicular transport. *Cell Reports*, 43(8):114649, 2024.

22 Maxim P Nikitin. Non-complementary strand commutation as a fundamental alternative for information processing by DNA and gene regulation. *Nature Chemistry*, 15(1):70–82, 2023.

23 Shigenori Nonaka, Hidetaka Shiratori, Yukio Saijoh, and Hiroshi Hamada. Determination of left-right patterning of the mouse embryo by artificial nodal flow. *Nature*, 418(6893):96–99, 2002.

24 Luvena L Ong, Nikita Hanikel, Omar K Yaghi, Casey Grun, Maximilian T Strauss, Patrick Bron, Josephine Lai-Kee-Him, Florian Schueder, Bei Wang, Pengfei Wang, et al. Programmable self-assembly of three-dimensional nanostructures from 10,000 unique components. *Nature*, 552(7683):72–77, 2017.

25 Vidhi Pareek, Hua Tian, Nicholas Winograd, and Stephen J. Benkovic. Metabolomics and mass spectrometry imaging reveal channeled de novo purine synthesis in cells. *Science*, 368(6488):283–290, 2020.

26 Jacob Parres-Gold, Matthew Levine, Benjamin Emert, Andrew Stuart, and Michael B Elowitz. Contextual computation by competitive protein dimerization networks. *Cell*, 188(7):1984–2002, 2025.

27 Allan Pinkus. Approximation theory of the MLP model in neural networks. *Acta Numerica*, 8:143–195, 1999.

28 William Poole, Andrés Ortiz-Muñoz, Abhishek Behera, Nick S. Jones, Thomas E. Ouldridge, Erik Winfree, and Manoj Gopalkrishnan. Chemical Boltzmann machines. In Robert Brijder and Lulu Qian, editors, *DNA Computing and Molecular Programming*, pages 210–231, Cham, 2017. Springer International Publishing. doi:10.1007/978-3-319-66799-7_14.

29 William Poole, Thomas Ouldridge, Manoj Gopalkrishnan, and Erik Winfree. Detailed balanced chemical reaction networks as generalized Boltzmann machines. *arXiv preprint arXiv:2205.06313*, 2022. doi:10.48550/arXiv.2205.06313.

30 Catherine Rabouille and Simon Alberti. Cell adaptation upon stress: the emerging role of membrane-less compartments. *Current Opinion in Cell Biology*, 47:34–42, 2017.

31 Venki Ramakrishnan. Ribosome structure and the mechanism of translation. *Cell*, 108(4):557–572, 2002.

32 Frank Rosenblatt. The perceptron: a probabilistic model for information storage and organization in the brain. *Psychological Review*, 65(6):386, 1958.

33 Paul WK Rothemund. Folding DNA to create nanoscale shapes and patterns. *Nature*, 440(7082):297–302, 2006.

34 David Soloveichik and Erik Winfree. Complexity of self-assembled shapes. *SIAM Journal on Computing*, 36(6):1544–1569, 2007. doi:10.1137/S0097539704446712.

35 Alexei V Tkachenko, Bortolo Matteo Mognetti, and Sergei Maslov. Evolutionary chemical learning in dimerization networks. *arXiv preprint arXiv:2506.14006*, 2025. doi:10.48550/arXiv.2506.14006.

36 Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Series in Statistics. Springer, New York, 2009. doi:10.1007/B13794.

37 Ronald D Vale. The molecular motor toolbox for intracellular transport. *Cell*, 112(4):467–480, 2003.

38 Gerrit Van Meer, Dennis R Voelker, and Gerald W Feigenson. Membrane lipids: where they are and how they behave. *Nature Reviews Molecular Cell Biology*, 9(2):112–124, 2008.

39 Carl G Wagner. *A first course in enumerative combinatorics*, volume 49. American Mathematical Society, 2020.

40 Jianbang Wang, Zhenzhen Li, and Itamar Willner. Dynamic reconfigurable DNA nanostructures, networks and materials. *Angewandte Chemie International Edition*, 62(18):e202215332, 2023.

41 Bryan Wei, Mingjie Dai, and Peng Yin. Complex shapes self-assembled from single-stranded DNA tiles. *Nature*, 485(7400):623–626, 2012.

A Appendix

A.1 Configurations in the same fiber which are not translations/rotations

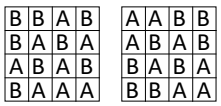

Figure 5 Two configurations with the same sufficient statistics which are not translations/rotations of each other. The neighborhood is von Neumann with periodic boundaries.

Here we provide an example of two configurations, shown in Figure 5, with the same sufficient statistics vector $t = (n_A, n_B, n_{A,A}, n_{A,B}, n_{B,B}) = (8, 8, 4, 12, 4)$, which are not rotations or translations. A short proof that they are not rotations or translations is to note that in the configuration in Figure 5 right, no B has four B neighbors, whereas the configuration in Figure 5 left, has a B with four B neighbors.

A.2 Counting the size of probability vector set $\mathcal{D}_{\mathcal{T}}$

Recall

$$\mathcal{D}_{\mathcal{T}} = \left\{ q : \mathcal{T}_\Omega \to [0,1] \;\middle|\; q(t) = \frac{z_t}{|\mathcal{T}_\Omega|}, z_t \in \mathbb{Z}_{\geq 0}, \sum_{t \in \mathcal{T}_\Omega} z_t = |\mathcal{T}_\Omega| \right\}.$$

Let $N = |\mathcal{T}_\Omega|$ and enumerate $\mathcal{T}_\Omega = \{t_1, \ldots, t_N\}$. Consider the function:

$$\Phi : q \to (z_1, \ldots, z_N), \qquad z_i = Nq(t_i).$$

Φ is a bijection between $\mathcal{D}_{\mathcal{T}}$ and the set of *weak compositions of N into N parts*:

$$\left\{ (z_1, \ldots, z_N) \in \mathbb{Z}_{\geq 0}^N \mid \sum_{i=1}^{N} z_i = N \right\}.$$

The set of weak compositions of N into N parts can be shown to have size $\binom{2N-1}{N-1}$ via stars-and-bars counting (see [39]).

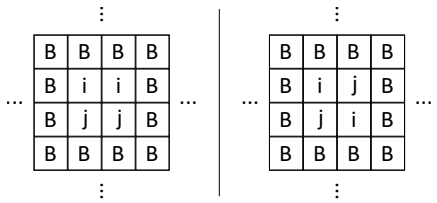

Figure 6 Two subconfigurations of a subset of positions of a lattice, with $i, j \in M$. Assuming von Neumann neighbors, the internal 2×2 positions in the left subconfiguration contribute $(n_{i,i}, n_{i,j}, n_{j,j}) = (1, 2, 1)$, while the internal 2×2 positions on the right contribute $(n_{i,i}, n_{i,j}, n_{j,j}) = (0, 4, 0)$.

A.3 The set of feasible sufficient statistics vectors $\mathcal{T}_\Omega$ scales exponentially in the number of parameters in the regime $|V| = \Theta(|M|^2)$

Here, for the sake of the counting argument of Section 3.1, we argue that $\mathcal{T}_\Omega := \{t(\sigma) \mid \sigma \in \Omega\}$, the set of sufficient statistics feasible for a given configuration space Ω, scales exponentially with the number of parameters $n = |\theta|$ in a construction where $|V| = \Theta(|M|^2)$. To do so, we consider a subset of configurations, where each configuration has a fixed count of species (and thus is a subset of both canonical and grand-canonical configuration spaces). Each configuration consists of 2×2 blocks of species, separated by a "buffer" species B that keeps the blocks separated (see Figure 6).

Assume the lattice is large enough to contain $\binom{|M|}{2}$ disjoint 2×2 blocks, one for each unordered pair of species $i, j \in M$, separated by buffer B. As shown in Figure 6, each i, j block can have one of at least two distinct arrangements, which give distinct contributions to $t(\sigma)$: $(n_{i,i}, n_{i,j}, n_{j,j}) = (1, 2, 1)$ or $(n_{i,i}, n_{i,j}, n_{j,j}) = (0, 4, 0)$. The gadgets are disjoint, so each choice can be made independently. Therefore $|\mathcal{T}_\Omega| \geq 2^{\binom{|M|}{2}}$.

In the isotropic canonical model, the number of parameters is $\frac{|M|(|M|+1)}{2}$, and in the isotropic grand-canonical model, the number of parameters is $|M| + \frac{|M|(|M|+1)}{2}$. In the anisotropic models with fixed, finite neighborhood and orientations sets, there is a constant multiplicative factor to the number of parameters. Thus, in each model variant, $n = \Theta(|M|^2)$. Since $\binom{|M|}{2} = \Theta(|M|^2)$, we have $|\mathcal{T}_\Omega| \geq 2^{\Omega(n)}$, i.e., the number of feasible sufficient statistic vectors t scales exponentially in the number of parameters, if $|V|$ is allowed to scale as $|M|^2$.

A.4 Microstate universality counterexamples

A.4.1 Isotropic grand-canonical model

Consider a ring of four positions with left-right edges and no top-bottom edges (see Figure 4(a), right). The species are A and B, so that $\theta = (G_A, G_B, G_{A,A}, G_{A,B}, G_{B,B})$ and $t(\sigma) = (n_A, n_B, n_{A,A}, n_{A,B}, n_{B,B})$. There are six feasible t values as shown in Figure 4(a), right, and the counterexample distribution is defined as follows:

If $t(\sigma) = t_0 = (1, 3, 0, 2, 2)$, set $Q(\sigma) = \frac{1}{64}$.

If $t(\sigma) = t_1 = (2, 2, 1, 2, 1)$, set $Q(\sigma) = \frac{1}{32}$.

If $t(\sigma) = t_2 = (3, 1, 2, 2, 0)$, set $Q(\sigma) = \frac{1}{64}$.

If $t(\sigma) = t_3 = (4, 0, 4, 0, 0)$, set $Q(\sigma) = \frac{1}{8}$.

If $t(\sigma) = t_4 = (2, 2, 0, 4, 0)$, set $Q(\sigma) = \frac{1}{4}$.

If $t(\sigma) = t_5 = (0, 4, 0, 0, 4)$, set $Q(\sigma) = \frac{1}{8}$.

Note that $t_1 = \frac{t_0+t_2}{2}$ and that the probabilities sum to 1. Towards contradiction, assume Q satisfies the affine log-probabilities constraint. As in the isotropic canonical model example in the main body, this means Q satisfies $u(t_1)^2 = u(t_0)u(t_2)$, however, $\left(\frac{1}{32}\right)^2 \neq \left(\frac{1}{64}\right)^2$, thus a contradiction. Therefore there does not exist θ for the isotropic grand canonical model such that the Boltzmann distribution is Q.

A.4.2 Anisotropic canonical model

The approach is to embed the isotropic canonical counterexample (Section 3.2) into the anisotropic case. Let V be a ring of 6 positions as in the isotropic canonical counterexample (Figure 4(a), left). The set of species types is $M_{\text{type}} = \{A, B\}$. We denote a standard orientation for each type: $A_0 = (A, o_1), B_0 = (B, o_1)$. We fix the canonical counts to have three of A and three of B. With direction $+\hat{x} = (1, 0, 0)$, we consider the subset of statistics $t^{+\hat{x}} = (n^{+\hat{x}}_{A_0,A_0}, n^{+\hat{x}}_{A_0,B_0}, n^{+\hat{x}}_{B_0,A_0}, n^{+\hat{x}}_{B_0,B_0})$. Note that the statistics $n^{-\hat{x}}_{X,Y} = n^{+\hat{x}}_{Y,X}$. Since we consider V as a one-dimensional ring and the following specifications of $t^{+\hat{x}}$ are for configurations in which all species are in their standard orientations A_0, B_0, each $t^{+\hat{x}}$ specifies a full fiber since the $-\hat{y}, +\hat{y}, -\hat{z}, +\hat{z}$ directions are unavailable for their respective statistics (e.g., $n^{-\hat{y}}_{A_0,B_0} = 0$) and non-standard oriented species have zero for their statistics (e.g., $n^{+\hat{x}}_{(A,o_2),(B,o_2)} = 0$). The $A_0B_0A_0B_0A_0B_0$ configurations have

$$t_0^{+\hat{x}} = (0, 3, 3, 0),$$

the $A_0A_0B_0B_0A_0B_0$ configurations have

$$t_1^{+\hat{x}} = (1, 2, 2, 1),$$

and the $A_0A_0A_0B_0B_0B_0$ configurations have

$$t_2^{+\hat{x}} = (2, 1, 1, 2).$$

These satisfy $t_1^{+\hat{x}} = \frac{t_0^{+\hat{x}}+t_2^{+\hat{x}}}{2}$.

We now define our fiber-uniform target distribution Q by assigning a microstate probability $u(t)$ to each fiber F_t. We set $u(t_0^{+\hat{x}}) = u(t_2^{+\hat{x}}) = p$, and $u(t_1^{+\hat{x}}) = 2p$ such that $p > 0$. Probabilities for each other fiber are set to arbitrary positive real numbers such that the sum of probabilities is one. As in Section 3.2, assume Q satisfies the affine log-probabilities constraint, then we have $u(t_1^{+\hat{x}})^2 = u(t_0^{+\hat{x}})u(t_2^{+\hat{x}})$, therefore $(2p)^2 = p^2$, which gives a contradiction since $p > 0$.

A.4.3 Anisotropic grand-canonical model

Here, we embed the isotropic grand-canonical counterexample distribution (Appendix A.4.1) in the anisotropic case. Let V be a ring of four positions, as in the isotropic grand-canonical counterexample (Figure 4(a), right). The set of species types is $M_{\text{type}} = \{A, B\}$. We denote a standard orientation for each type: $A_0 = (A, o_1), B_0 = (B, o_1)$. With direction $+\hat{x} = (1, 0, 0)$, we consider the subset of statistics $t^{+\hat{x}} = (n_A, n_B, n^{+\hat{x}}_{A_0,A_0}, n^{+\hat{x}}_{A_0,B_0}, n^{+\hat{x}}_{B_0,A_0}, n^{+\hat{x}}_{B_0,B_0})$. Note that the statistics $n^{-\hat{x}}_{X,Y} = n^{+\hat{x}}_{Y,X}$. Since we consider V as a one-dimensional ring and the

following specifications of $t^{+\hat{x}}$ are for configurations in which all species are in their standard orientations A_0, B_0, each $t^{+\hat{x}}$ specifies a full fiber since the $-\hat{y}, +\hat{y}, -\hat{z}, +\hat{z}$ directions are unavailable for their respective statistics (e.g., $n^{-\hat{y}}_{A_0,B_0} = 0$) and non-standard oriented species have zero for their statistics (e.g., $n^{+\hat{x}}_{(A,o_2),(B,o_2)} = 0$). The $A_0B_0B_0B_0$ configurations have

$$t_0^{+\hat{x}} = (1, 3, 0, 1, 1, 2),$$

the $A_0A_0B_0B_0$ configurations have

$$t_1^{+\hat{x}} = (2, 2, 1, 1, 1, 1),$$

and the $A_0A_0A_0B_0$ configurations have

$$t_2^{+\hat{x}} = (3, 1, 2, 1, 1, 0).$$

These satisfy $t_1^{+\hat{x}} = \frac{t_0^{+\hat{x}} + t_2^{+\hat{x}}}{2}$.

We now define our fiber-uniform target distribution Q by assigning a microstate probability $u(t)$ to each fiber F_t. We set $u(t_0^{+\hat{x}}) = u(t_2^{+\hat{x}}) = p$, and $u(t_1^{+\hat{x}}) = 2p$ such that $p > 0$. Probabilities for each other fiber are set to arbitrary positive real numbers such that the sum of probabilities is one. As in Section 3.2, assume Q satisfies the affine log-probabilities constraint, then we have $u(t_1^{+\hat{x}})^2 = u(t_0^{+\hat{x}})u(t_2^{+\hat{x}})$, therefore $(2p)^2 = p^2$, which gives a contradiction since $p > 0$.

A.5 Proof of Theorem 6

Proof. The construction is similar to the proof of Theorem 5, except with the addition of interactions with the boundary species. Note the boundary condition as defined just before this theorem's statement. For each label configuration β_i such that $Q_K(\beta_i) > 0$, we include species $M_{\beta_i}^{(x,y,z)}$ for $(x, y, z) \in V$ and call $\text{pos}\left(M_{\beta_i}^{(x,y,z)}\right) = (x, y, z)$ the *intended position* and $\text{cfg}\left(M_{\beta_i}^{(x,y,z)}\right) = \beta_i$ the *intended label configuration*. We set $G_{i,j} = 0$ if $\text{pos}(i)$ and $\text{pos}(j)$ are von Neumann neighbors and $\text{cfg}(i) = \text{cfg}(j)$, and otherwise set $G_{i,j} = \Lambda$. Additionally, for all β_i, we set $G_{i,j} = 0$ for $i = M^{p_{\hat{u}}}$ and $j = M^{p}_{\beta_i}$ for $p_{\hat{u}}$ on the boundary and p such that p and $p_{\hat{u}}$ are von Neumann neighbors. We call a (microstate) configuration σ *intended* if all species are next to their six intended von Neumann neighbors and all agree on intended label configuration, formally for all $u \in V^{+2}$, $\text{pos}(\sigma(u)) = u$ and for all $u, v \in V$, $\text{cfg}(\sigma(u)) = \text{cfg}(\sigma(v))$, otherwise we call configurations *unintended*. For each β_i, for species $M_{\beta_i}^{(0,0,0)}$, we set $G_{M_{\beta_i}^{(0,0,0)}} = -kT \ln Q_K(\beta_i)$. For all other species i, we set $G_i = 0$.

The key is to show that any unintended configuration σ must have $G(\sigma) \geq \Lambda$. Let $S_k = \{(x, y, z) \mid x \in \{k-1, L-k\} \vee y \in \{k-1, L-k\} \vee z \in \{k-1, L-k\}\}$, i.e., S_k contains the kth hollow cube of points moving inwards from the boundary S_0. We define $H(k)$ as the proposition: if there is no Λ contribution from interactions between species in $S_0 \bigcup \cdots \bigcup S_k$, then all these species agree on label configuration and are in their intended positions. $H(0)$ is always true since the boundary is fixed and provides no Λ contribution. We induct on $H(k)$. Assuming $H(k-1)$, if there is no Λ contribution in $S' = S_0 \bigcup \cdots \bigcup S_{k-1}$, we know all species in S_{k-1} agree on label configuration and are in their intended positions. Consider the contrapositive of the inductive hypothesis $H(k)$: if some species in S_k disagrees on label configuration or is not in its intended position, then there is a Λ contribution. The contrapositive is true since disagreement on label configuration between a species in S_k and one of its von Neumann neighbors in S_{k-1} contributes a Λ term, and similarly a species in S_k

which is not in its intended position contributes a Λ term with its von Neumann neighbors in S_{k-1} which are in their intended positions. Therefore $H(k)$ is true for all k, and so the contrapositive "if some species do not agree on label configuration or are not in their intended positions, then there is a Λ contribution from interactions between species in $S_0 \bigcup \cdots \bigcup S_k$" is true. This implies that every unintended configuration σ has $G(\sigma) \geq \Lambda$.

Finally, we set Λ in order to achieve KL divergence less than ε. Let $\mathcal{I}$ be the set of intended configurations and $\mathcal{U}$ be the set of unintended configurations. For each label configuration β, there is exactly one intended configuration σ_β with $\mathcal{K} \circ \sigma_\beta = \beta$, namely the configuration in which each interior position $p \in V$ contains the species M_β^p, while the boundary positions contain their fixed boundary species. By construction, every interaction term in this intended configuration contributes zero, except for the single leader species $M_\beta^{(0,0,0)}$, whose per-molecule energy is $G_{M_\beta^{(0,0,0)}} = -kT \ln Q_K(\beta)$. Therefore $G(\sigma_\beta) = -kT \ln Q_K(\beta)$ so $\exp(-G(\sigma_\beta)/kT) = Q_K(\beta)$. Summing over intended configurations gives

$$\sum_{\sigma \in \mathcal{I}} \exp(-G(\sigma)/kT) = \sum_{\beta} Q_K(\beta) = 1.$$

By the inductive proof above, every unintended configuration has at least one Λ penalty, and all other per-molecule energy contributions are nonnegative, since $Q_K(\beta) \leq 1$ implies $-kT \ln Q_K(\beta) \geq 0$. Therefore every $\sigma \in \mathcal{U}$ satisfies $G(\sigma) \geq \Lambda$. Since the configuration space is finite, we have $\sum_{\sigma \in \mathcal{U}} \exp(-G(\sigma)/kT) \leq |\Omega| \exp(-\Lambda/kT)$. We choose Λ so that $|\Omega| \exp(-\Lambda/kT) < \exp(\varepsilon) - 1$. Then $Z < \exp(\varepsilon)$. For any label configuration β, $P_\mathcal{K}(\beta) = \sum_{\sigma | \mathcal{K} \circ \sigma = \beta} P(\sigma) \geq P(\sigma_\beta) = \frac{Q_K(\beta)}{Z}$. Finally,

$$D_{\mathrm{KL}}(Q_K || P_\mathcal{K}) = \sum_{\beta} Q_K(\beta) \log \frac{Q_K(\beta)}{P_\mathcal{K}(\beta)} \leq \sum_{\beta} Q_K(\beta) \log Z = \log Z < \varepsilon. \qquad \blacktriangleleft$$