

Heat Flow Leaves Two-Hidden-Unit Restricted Boltzmann Models

Alper Ferudun*

20 October 2026

Abstract

Independent bit-flip heat flow preserves binary restricted Boltzmann machines with one hidden unit, by their classical representation as mixtures of two product distributions. We give explicit evolving parameters. With two hidden units and at least four visible units, the behavior is different: for every fixed positive vector of flip rates, a generic positive parameter leaves even the algebraic closure of the model immediately. A sixteen-by-sixteen integer Jacobian determinant proves transversality in the four-visible case. A coordinate-permutation argument covers arbitrary positive rates, and conditioning extends the obstruction to every larger visible dimension. We distinguish a rank witness from a pointwise smoothness certificate, so the generic-exit argument does not assume that a parametrized tangent is the tangent of the full model at the numerical witness. For the logarithmic-mean graph transport metric, the result gives a generically nonzero residual in projected entropy flow. It concerns a finite-state heat equation, not a general continuum-PDE approximation theorem.

1 Models, heat flow, and the result

For $x \in \{0, 1\}^n$, positive parameters $\theta = (a_i, c_j, w_{ji})$ define

$$f_\theta(x) = \prod_{i=1}^n a_i^{x_i} \prod_{j=1}^m \left(1 + c_j \prod_{i=1}^n w_{ji}^{x_i} \right), \quad p_\theta(x) = \frac{f_\theta(x)}{Z_\theta}, \quad Z_\theta = \sum_x f_\theta(x). \quad (1)$$

These are the exponentiated biases and weights of a binary restricted Boltzmann machine (RBM). Let $\mathcal{R}_{n,m}^+$ be the finite-parameter model, $\overline{\mathcal{R}}_{n,m}$ its Euclidean closure in the probability simplex, and $V_{n,m}$ its complex projective Zariski closure. All initial distributions below have full support.

For $\kappa_i \geq 0$, define the independent bit-flip generator

$$(L_\kappa f)(x) = \sum_{i=1}^n \kappa_i (f(x \oplus e_i) - f(x)), \quad T_t = \exp(tL_\kappa). \quad (2)$$

It acts on mass functions by the same formula, since its matrix is symmetric.

Theorem 1. *Fix $n \geq 4$ and $\kappa_i > 0$. There is a nonzero real polynomial Q_κ in the $3n + 2$ exponentiated parameters of $\mathcal{R}_{n,2}^+$ such that, for every positive θ with $Q_\kappa(\theta) \neq 0$, there is $\varepsilon_\theta > 0$ satisfying*

$$[T_t p_\theta] \notin V_{n,2} \quad (0 < t < \varepsilon_\theta).$$

*Mercury Software GmbH. Email: alper@mercurycodelab.com. GitHub: [AlperTheKing](https://github.com/AlperTheKing). Unrefereed preprint. CC BY 4.0.

In particular $T_t p_\theta \notin \overline{\mathcal{R}}_{n,2}$ at those times. For $n = 4$, the assertion holds for any nonnegative rate vector other than zero.

The exceptional parameter set has empty interior and Lebesgue measure zero. The polynomial and the exit times are existential; neither is asserted to be uniform. This is a noninvariance theorem for a fixed-width statistical model, not a failure of the heat equation itself.

The algebraic description of RBMs as Hadamard products of secant varieties is classical [1, 2]. In particular $V_{4,2}$ is a degree-110 hypersurface [1, Section 2]. Our argument does not require its implicit equation, or its degree. It uses an exact Jacobian certificate and includes the passage from a parametrized tangent to generic exit from the whole model.

2 One hidden unit: an explicit exact evolution

Proposition 2. *For every n and nonnegative rate vector, $T_t \mathcal{R}_{n,1}^+ \subseteq \mathcal{R}_{n,1}^+$ for all $t \geq 0$. More generally, independent two-state stochastic channels preserve this family whenever the transformed conditional component probabilities remain in $(0, 1)$.*

Proof. Conditioning on the hidden bit gives the classical two-product mixture

$$p(x) = (1 - \lambda) \prod_i r_{i0}^{x_i} (1 - r_{i0})^{1-x_i} + \lambda \prod_i r_{i1}^{x_i} (1 - r_{i1})^{1-x_i},$$

where

$$0 < \lambda < 1, \quad r_{i0} = \frac{a_i}{1 + a_i}, \quad r_{i1} = \frac{a_i w_i}{1 + a_i w_i}.$$

Independent channels preserve each product component and its mixture weight. For (2),

$$r_{ih}(t) = \frac{1}{2} + (r_{ih}(0) - \frac{1}{2})e^{-2\kappa_i t}. \quad (3)$$

The following positive parameters represent the resulting mixture:

$$\begin{aligned} a_i(t) &= \frac{r_{i0}(t)}{1 - r_{i0}(t)}, & w_i(t) &= \frac{r_{i1}(t)}{1 - r_{i1}(t)} a_i(t)^{-1}, \\ c(t) &= \frac{\lambda}{1 - \lambda} \prod_i \frac{1 + a_i(t)}{1 + a_i(t)w_i(t)}. \end{aligned} \quad (4)$$

Indeed the two unnormalized components have total masses $\prod_i (1 + a_i)$ and $c \prod_i (1 + a_i w_i)$, with ratio $\lambda/(1 - \lambda)$. Equation (3) stays in $(0, 1)$ at every finite time. The same calculation proves the channel statement. $\square$

This proposition is a direct consequence of the known mixture representation [1, 2]; it is included to contrast exact one-hidden evolution with the obstruction at two hidden units.

3 An integer transversality certificate

For $n = 4, m = 2$, put $z_j(x) = c_j \prod_i w_{ji}^{x_i}$ and $v(x) = \prod_i a_i^{x_i}$. Define columns indexed by the sixteen binary vertices:

$$f = v(1 + z_1)(1 + z_2), \quad g_1 = v z_1(1 + z_2), \quad g_2 = v z_2(1 + z_1).$$

They are unnormalized masses and hidden-bias logarithmic derivatives. The visible-bias derivatives are $x_i f$, and the weight derivatives are $x_i g_j$, with entrywise multiplication. Use the fifteen-column matrix

$$B_\theta = [f, x_1 f, x_2 f, x_3 f, x_4 f, g_1, g_2, x_1 g_1, x_2 g_1, x_3 g_1, x_4 g_1, x_1 g_2, x_2 g_2, x_3 g_2, x_4 g_2]. \quad (5)$$

The first column supplies scaling; the others are all fourteen parameter derivatives.

Lemma 3. *Order vertices lexicographically, with x_4 varying fastest. At*

$$a = (1, 1, 1, 1), \quad c = (1, 1), \quad (w_{1i}) = (3, 1, 2, 2), \quad (w_{2i}) = (3, 3, 3, 2), \quad (6)$$

one has

$$\det[B_\theta, (L_1 + L_2 + L_3 + L_4)f_\theta] = 11241871441920 \neq 0,$$

where $L_i f(x) = f(x \oplus e_i) - f(x)$.

Proof. All entries in (5) are integers. Evaluating its formulas gives

$$f = (4, 9, 12, 35, 8, 21, 30, 95, 16, 49, 70, 247, 40, 133, 196, 715), \quad Z = 1680.$$

For completeness the coordinate determinants are

i	$\det[B_\theta, L_i f_\theta]$
1	14555740372992
2	0
3	551809843200
4	-3865678774272

and their sum is the asserted integer. A compact cofactor certificate is

$$\begin{aligned} \nu = & (-22932, 66885, 15540, -37905, 11466, -28665, -6216, 13965, \\ & -49875, 6327, 25480, -2695, 19950, -2331, -9100, 931). \end{aligned} \quad (7)$$

The signed cofactors of the last column in the sixteen-by-sixteen matrix are 334430208ν , and direct multiplication gives

$$\nu^T B_\theta = 0, \quad \nu^T (L_1 + L_2 + L_3 + L_4)f_\theta = 33615.$$

Expanding in the last column proves the determinant. The supplementary checker reconstructs all entries from (1) and also by summing over the four hidden states. It verifies the cofactors, the determinant by two exact elimination methods, and four modular reductions. $\square$

In particular $\text{rank } B_\theta = 15$. It follows that the normalized parameter derivative has rank fourteen and the heat vector is outside its range. This alone does *not* establish that the particular mass vector in the lemma is a smooth point of the full algebraic model. We use the certificate to prove a generic statement, not an unverified pointwise exit claim at (6).

4 From rank to exit from the full model

Lemma 4. *If $D_\kappa(\theta) = \det[B_\theta, L_\kappa f_\theta]$ is not the zero polynomial, Theorem 1 holds for four visible units and those rates.*

Proof. The rational normalization map $\theta \mapsto p_\theta$ has irreducible complex Zariski closure X in the affine hyperplane $A = \{\sum_x p_x = 1\}$. Its dimension is at most fourteen. At a parameter where $D_\kappa \neq 0$, quotienting the fifteen independent columns of B_θ by the scale column gives $\text{rank } Dp_\theta = 14$. Thus X is an irreducible hypersurface of the fifteen-dimensional space A .

The smooth locus of X is nonempty and Zariski open in characteristic zero. Its preimage under the dominant rational parameter map is nonempty and Zariski open. Intersecting with $D_\kappa \neq 0$ remains nonempty. This intersection meets the positive real parameter orthant: the excluded sets are proper algebraic sets and the positive orthant is Zariski dense. At such a parameter the normalized derivative has range $T_{p_\theta}X$.

Moreover $L_\kappa p_\theta$ is outside that range. Otherwise multiplying a tangent representation by Z_θ would put $L_\kappa f_\theta$ in the span of the scale and parameter derivative columns, contradicting $D_\kappa \neq 0$.

Let F be a real irreducible defining polynomial for X in A . It may be chosen real since X is the Zariski closure of a real parametrization. At the smooth transverse parameter just found,

$$dF_{p_\theta}(L_\kappa p_\theta) \neq 0.$$

This expression is rational in the positive exponentiated parameters. After clearing powers of Z_θ , its numerator Q_κ is a nonzero real polynomial. For every positive θ with $Q_\kappa(\theta) \neq 0$, analyticity in time gives

$$F(T_t p_\theta) = t dF_{p_\theta}(L_\kappa p_\theta) + O(t^2) \neq 0 \quad (0 < t < \varepsilon_\theta) \quad (8)$$

for a sufficiently small $\varepsilon_\theta > 0$. Since X is the normalized affine part of $V_{4,2}$, this is the required exit. The Euclidean model closure is contained in X . $\square$

Proof of Theorem 1. First let $n = 4$. At (6), write $d_i = \det[B_\theta, L_i f_\theta]$. Their sum is strictly positive. Permuting visible coordinates produces, up to a common row/column sign for each permutation, the rate combinations $\sum_i \kappa_i d_{\sigma(i)}$, $\sigma \in S_4$. Their sum is

$$6 \left(\sum_i \kappa_i \right) \left(\sum_i d_i \right) > 0$$

for any nonnegative rate vector other than zero. Hence at least one permuted certificate has $D_\kappa \neq 0$. Lemma 4 proves the four-visible assertion.

Now let $n > 4$. Choose four coordinates with their positive rates and a four-visible parameter at which (8) has nonzero linear coefficient. Set $a_i = w_{1i} = w_{2i} = 1$ on the remaining coordinates. The initial law is $p_4 \otimes u_{n-4}$, with u_{n-4} uniform, and its heat evolution is $T_t^{(4)} p_4 \otimes u_{n-4}$.

Let $\tilde{F}$ be the homogeneous defining polynomial of $V_{4,2}$ corresponding to F . Conditioning any two-hidden-unit RBM on the last $n - 4$ visible bits yields a two-hidden-unit RBM on the first four. In particular the polynomial

$$G(p) = \tilde{F}((p_{x,0}, \dots, 0)_{x \in \{0,1\}^4})$$

vanishes on the entire n -visible model and on $V_{n,2}$. At the embedded product law,

$$G(T_t(p_4 \otimes u_{n-4})) = 2^{-(n-4) \deg \tilde{F}} \tilde{F}(T_t^{(4)} p_4),$$

whose derivative at zero is nonzero. The rational expression $dG_{p_\theta}(L_\kappa p_\theta)$ therefore has a nonzero polynomial numerator in all $3n + 2$ parameters after normalization denominators are cleared. Its nonvanishing yields the same Taylor argument. $\square$

Remark 5. The proof does not identify Q_κ , prove smoothness at the displayed integer witness, or supply an effective common exit-time bound. These are not assumptions in the generic statement.

5 Projected graph transport flow

The invariance question has a direct implication for finite-state transport-based model reduction. Assume $\kappa_i > 0$, put $\pi(x) = 2^{-n}$ and $\rho = p/\pi$, and use $\langle u, v \rangle_\pi = \sum_x \pi(x)u(x)v(x)$. Define the logarithmic mean by

$$\Lambda(a, b) = \frac{a - b}{\log a - \log b} \quad (a \neq b), \quad \Lambda(a, a) = a.$$

For a positive density set

$$(K_\rho \psi)(x) = \sum_i \kappa_i \Lambda(\rho(x), \rho(x \oplus e_i)) (\psi(x) - \psi(x \oplus e_i)).$$

Connectedness makes K_ρ positive definite on zero-mean vectors. The logarithmic-mean graph metric and relative entropy are

$$g_\rho(s, t) = \langle s, K_\rho^\dagger t \rangle_\pi, \quad \mathcal{H}(\rho) = \langle \rho, \log \rho \rangle_\pi.$$

The identity $-K_\rho \log \rho = L_\kappa \rho$ is the finite-state entropy-gradient identity of Maas [3, Theorem 1.2]. The uniform rate normalization used for a Markov kernel can equivalently be absorbed by time rescaling.

On an identifiable smooth model chart with derivative $J = D\rho_\theta$, the pullback gradient equation is

$$(J^* K_\rho^\dagger J) \dot{\theta} = -J^* \log \rho. \quad (9)$$

Here the adjoint uses the Euclidean parameter inner product and the π -weighted density inner product. This is the usual statistical-manifold pullback construction; see [4] for graph transport natural gradients. We specify the logarithmic mean explicitly, rather than identifying all graph mobility conventions.

Let $u = L_\kappa \rho$ and $y = J\dot{\theta}$. Pairing (9) with any parameter direction shows that $u - y$ is g_ρ -orthogonal to the image tangent. Thus

$$\frac{d}{dt} \mathcal{H}(\rho) = -\|y\|_\rho^2, \quad \|u\|_\rho^2 = \|y\|_\rho^2 + \|u - y\|_\rho^2.$$

These are projection identities, not additional approximation theorems. The explicit lift in Proposition 2 gives $u = y$ on regular one-hidden image charts. Theorem 1 yields $u \neq y$ generically for two hidden units and $n \geq 4$: a tangent vector annihilates every polynomial vanishing on the model, whereas the polynomial used in (8) has nonzero directional derivative. Redundant raw RBM parameters must be replaced by an image chart or have their derivative kernel removed; their pullback matrix need not be invertible.

This is the dynamical logarithmic-mean transport geometry, not ordinary static discrete W_2 . For example, on two states at distance one, $W_2^2(\text{Ber}(r), \text{Ber}(s)) = |r - s|$, so a differentiable curve with nonzero derivative has infinite speed in that static metric. It cannot be substituted into (9).

6 Scope and reproducibility

The AIM Boltzmann-machines workshop discussed Wasserstein natural gradients and their computational implementation [5, p. 3]. The present result addresses exactness and an obstruction for one specific finite-state heat equation. It does not resolve a general question about using RBMs

to approximate PDEs, prove a continuum limit, or provide hidden-width-dependent approximation error estimates. No assertion is made for more than two hidden units, or for generic noninvariance with fewer than four visible units.

The contribution is the explicit heat-transversality certificate and its generic-exit, rate-permutation, and conditioning consequences. The mixture representation, algebraic RBM framework, graph metric, and pullback construction are credited above. A bounded literature search did not locate this precise result; this is not an absolute priority claim.

The accompanying Python files use only the standard library. They check the Jacobian by direct hidden-state enumeration, compare integer Bareiss elimination with rational Gaussian elimination, verify the cofactor normal and four modular determinants, and test the product-channel formula and parameter reconstruction on 120 exact rational cases. A further 100 integer parameter cases check derivative reconstruction. Ordinary and optimized interpreter runs agree. These checks support the displayed finite certificate; they are not a substitute for the generic algebraic argument or a formal proof-assistant certification. AI assistance was used in derivation, literature search, exact-code verification, and manuscript preparation. The note is unrefereed.

References

- [1] M. A. Cueto, J. Morton, and B. Sturmfels, *Geometry of the restricted Boltzmann machine*, in Algebraic Methods in Statistics and Probability II, Contemporary Mathematics **516** (2010), 135–153. doi:10.1090/conm/516/10172. Preprint [arXiv:0908.4425](https://arxiv.org/abs/0908.4425).
- [2] G. Montúfar, *Restricted Boltzmann machines: introduction and review*, in Information Geometry and Its Applications IV (2018), 75–115. [arXiv:1806.07066](https://arxiv.org/abs/1806.07066).
- [3] J. Maas, *Gradient flows of the entropy for finite Markov chains*, Journal of Functional Analysis **261** (2011), 2250–2292. doi:10.1016/j.jfa.2011.06.009.
- [4] W. Li and G. Montúfar, *Natural gradient via optimal transport*, Information Geometry **1** (2018), 181–214. doi:10.1007/s41884-018-0015-3. [arXiv:1803.07033](https://arxiv.org/abs/1803.07033).
- [5] American Institute of Mathematics, *Boltzmann machines*, workshop report, 2018. <https://aimath.org/pastworkshops/boltzmannrep.pdf>.